

LEARNING AND NOISY EQUILIBRIUM BEHAVIOR
IN AN EXPERIMENTAL STUDY OF IMPERFECT PRICE COMPETITION

C. Monica Capra, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt
January 2000

Department of Economics
114 Rouss Hall
University of Virginia
Charlottesville VA 22903-3288

ABSTRACT

This paper considers a duopoly price-choice game in which the unique Nash equilibrium is the Bertrand outcome. Price competition, however, is imperfect in the sense that the market share of the high-price firm is not zero. Economic intuition suggests that price levels should be positively related to the market share of the high-price firm. Although this relationship is not predicted by standard game theory, it is implied by a generalization of the Nash equilibrium that results when players make noisy (logit) best responses to expected payoff differences. This logit equilibrium model was used to design a laboratory experiment with treatments that correspond to changing the market share of the high-price firm. The model predicts the final-period price averages for both treatments with remarkable accuracy. Moreover computer simulations of a naive learning model were used, *ex ante*, to predict the observed differences in the time paths of average prices.

JEL Classification: C72, C92.

Keywords: laboratory experiments, simulation, decision error, learning, logit equilibrium.

LEARNING AND NOISY EQUILIBRIUM BEHAVIOR
IN AN EXPERIMENTAL STUDY OF IMPERFECT PRICE COMPETITION

C. Monica Capra, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt*

I. Introduction

Many economic situations have the property that payoffs are determined by the minimum of all agents' decisions. For example, in a market for a homogeneous product, firms' sales only depend on the lowest posted price. In the absence of capacity constraints and imperfect information, the low-price firm will obtain all sales, and the familiar Bertrand outcome results.¹ This paper considers a model of price competition in which the firm with the lower price has a larger market share but, unlike the usual Bertrand setup, the market share for the high-price firm is not zero. This situation may occur when some buyers will seek the low-price firm but others have contracts that prevent switching if the seller is willing to match the lower price. For example, buyers with "meet-or-release" contracts cannot abandon a seller who meets a competitor's price cut. In this case, the firm that posts the higher price must sell at a price that matches the other's low price. With constant marginal cost, the Nash equilibrium for this game is a price that equals marginal cost, i.e. the Bertrand/Nash outcome, regardless of how close the market share of the high-price firm is to one half. In contrast, it is intuitively plausible that prices should increase on average as the market share of the high-price firm approaches one half.

In the literature, there have been two approaches to explaining the Bertrand paradox, i.e. changing the assumptions about industry structure or changing the assumptions about pricing behavior. Examples of the first approach include product differentiation (e.g. Anderson, de Palma, and Thisse, 1992), consumer search costs (e.g. Diamond, 1971), capacity constraints (e.g. Kreps and Scheinkman, 1983), and trigger-price strategies in repeated games (e.g. Friedman, 1971; Porter, 1983). These changes in structural assumptions add realism and richness to price competition models, but they cannot be the whole story, since pricing above marginal cost persists in experiments that implement the standard Bertrand model in finite horizon settings with

* This project was funded in part by the National Science Foundation (SBR-9617784, SBR-9818683). We wish to thank Jim Cox, John Kagel, and Dale Stahl for useful suggestions, and John Turner for research assistance.

¹ Alternatively, in a production process with perfect complementarities the group production is determined by the minimum of individual workers' efforts.

small numbers of sellers (Duwfenberg and Gneezy, 2000).

The second approach to the Bertrand paradox generally involves some relaxation of the assumption of perfect rationality, e.g., by introducing a set of "near-maximizing agents" (Akerlof and Yellen, 1985), ϵ -equilibria (Radner, 1980), or probabilistic choice models of boundedly rational equilibrium behavior (Rosenthal, 1989; McKelvey and Palfrey, 1995). The intuition behind the Akerlof and Yellen model is that some firms may not respond optimally to changes in an exogenous parameter, but this inertia may have only a second-order effect if profit functions are continuous and prices are near optimal. Radner's ϵ -equilibrium allows strategy combinations with the property that unilateral deviations cannot yield payoff increases that exceed (some small amount) ϵ . The idea behind this approach is similar to Simon's (1955) notion of "satisficing behavior," i.e. it may not be worthwhile to search for opportunities that result in only small gains. Behavior in an ϵ -equilibrium is "discontinuous" in the sense that deviations do not occur unless the gain is greater than ϵ , in which case they occur with probability one. In contrast, the probabilistic choice approach is based on the idea that choice probabilities are smooth, increasing functions of expected payoffs. McKelvey and Palfrey's quantal response equilibrium (QRE) incorporates probabilistic choice into an equilibrium framework. The analysis of this paper is an application of QRE to a pricing game with a continuum of feasible choices.

In particular, we use a logit probabilistic choice rule to inject some noise into best-response behavior. The effect of noise is determined by an error parameter. As the error parameter is reduced to zero, the equilibrium behavior in this "logit equilibrium" model converges to the Bertrand/Nash equilibrium. When the error parameter is positive, there will be a probability distribution of price decisions for any given beliefs about other's prices. In equilibrium, the distributions of price decisions that come out of the logit probabilistic choice rule match the distributions that represent players' beliefs.² Since the Nash equilibrium is a limiting case of the logit equilibrium, econometric estimates of the error parameter can be used to evaluate the extent to which behavior in laboratory experiments is "close" to Nash predictions.

² Rosenthal (1989) first proposed the incorporation of probabilistic choice into the equilibrium framework of game theory. The properties of such equilibria are derived in the classic paper by McKelvey and Palfrey (1995), for a general class of "quantal response equilibria" that includes the logit equilibrium as a special case. Their analysis pertains to matrix games; applications to games with continuous strategies can be found in Anderson, Goeree, and Holt (1997, 1999).

In addition to the equilibrium analysis, we use computer simulations of a dynamic learning model to predict trends in adjustment to equilibrium. These predictions are derived by using an error rate and a learning parameter estimated from a prior laboratory experiment (Capra *et al.*, 1999). The predictions are then evaluated with a new experiment in which the parameter representing the market share of the high price firm is changed. Finally, the data from the experiment are used to estimate the logit error parameter in two different ways: with the equilibrium model applied to data for the final periods and with the learning model applied to data for all periods. The difference between these two approaches is that beliefs in the equilibrium model are determined by the requirement that they are consistent with decision distributions, whereas beliefs in the learning model evolve over time in response to observed decisions of other subjects. We use simulation techniques to analyze the long-run outcomes of the learning model and show that the steady-state choice distributions are similar to the logit equilibrium, although systematic differences exist.

The order of topics is: description of the game and the logit equilibrium and derivation of comparative statics (section II), description of a naive learning model (section III), report on the experimental results (section IV), econometric and simulation analysis (section V), and conclusion (section VI).

II. A Model of Noisy Equilibrium Behavior

Consider a market game in which firms 1 and 2 simultaneously choose prices p_1 and p_2 in the range $[p_L, p_H]$. Demand is assumed to be perfectly inelastic and the sales quantity of the firm with the low price, p_{\min} , is normalized to be one, so the low-price firm earns an amount equal to its price. The high-price firm only earns αp_{\min} , where $\alpha < 1$ is inversely related to the degree of buyer responsiveness to low price, i.e. a high value of α corresponds to the case of low responsiveness. In the event of a tie, the $1+\alpha$ sales units are shared equally, so each seller earns $\frac{1}{2}(1+\alpha)p_{\min}$. This type of payoff structure can arise in a market in which some buyers are protected by "meet-or-release" contracts. These contracts require a seller to meet a rival's lower

price or release the buyer from the contract.³ With this admittedly simplified setup, a seller is, of course, willing to meet the competitor's price, and earn αp_{\min} , instead of releasing the buyers and earn nothing. This model is not intended to address subtle issues in contract theory, but rather, it is intended to provide a parametric example of a situation in which market shares are not perfectly sensitive to relative price.⁴

As long as the high-price firm obtains less than half of the market sales ($\alpha < 1$), the Nash equilibrium is for both firms to set the lowest possible price p_L . To see this, note that at any common price, firms have an incentive to undercut the other by a small amount to increase market share. Moreover, a unilateral price increase will not raise the price that buyers pay, since they will demand to be released from their contracts. Therefore, the usual Bertrand logic applies, and the unique Nash equilibrium involves both firms charging the lowest possible price. The harsh competitive nature of the Nash prediction seems to go against simple economic intuition that the degree of buyer inertia will affect the behavior of firms. When $\alpha = 0.8$ say, the loss from having the higher price is relatively small, and firms should be more likely to set prices above p_L when there is a small chance that rivals will do the same. Indeed, in the extreme case when $\alpha = 1$ it becomes a weakly dominant strategy for both firms to choose the highest possible price p_H . While a standard Nash analysis predicts no change as long as $\alpha < 1$ (and then an

³ For example, suppose that a proportion β of buyers are tied to each seller by meet-or-release contracts, where $\beta < 1/2$. The remaining $1 - 2\beta$ fraction of buyers are free to switch. The firm with the higher price would rather meet the low price than end up with zero sales, so the low-price seller will only sell to a fraction, $1 - \beta$, of the buyers. For notational simplicity, we normalize the sales quantity of the low-price seller to be 1, and the sales quantity of the high-price seller is therefore $\beta/(1-\beta) = \alpha$.

⁴ Meet-or-release contracts are sometimes used in producer goods markets, where buyers wish to ensure delivery in periods of temporary shortage without running the risk that other buyers (their downstream competitors) obtain inputs at a lower price. Industrial organization economists have conjectured that meet-or-release contracts may have anti-competitive effects because they protect a seller from losing existing sales when a rival cuts price. For example, Holt and Scheffman (1987) point out that meeting a rival's price cut allows a firm to maintain its sales quantity as its rival cuts price. This quantity-holding ability transforms the price competition game into one that is more analogous to a Cournot game. To see this, note that if one firm cuts price and conjectures that the other will hold its own quantity by matching the cut, then the only new sales will come from the entry of new buyers or additional units purchased by existing buyers. The possible anti-competitive effects of these contracts, must of course, be weighed against any efficiencies provided by having stable buyer-seller relationships. A more general analysis would also have to explain why some buyers sign the contracts and others do not. For example, supply-side uncertainty could provide some buyers with stable, high valuations to seek meet-or-release contracts in order to ensure supply in periods of unanticipated excess demand. Buyers with uncertain needs would be less likely to want to sign such contracts. This insurance role of meet-or-release contracts could produce efficiencies that offset some or all anti-competitive effects.

abrupt change when $\alpha \geq 1$), it seems plausible that prices will gradually rise with α . The model presented in this section captures this comparative static feature.⁵

The crucial assumption underlying the usual Bertrand result is that firms respond optimally to any potential gain in profit, *no matter how small*. Suppose instead that decision-making is less perfect: firms choose better options more often than worse ones, but do not necessarily choose the best one with probability one. Just as important, firms are assumed to realize that others' decisions are not perfectly predictable when the costs of "errors" are small. To formalize the interactive effect of these assumptions, we need a rule (other than perfect-maximization) that relates expected payoffs to decision probabilities. Let $\pi_i^e(p)$ denote the expected payoff from choosing a price p , which depends on the distribution of the rival's price, denoted by $F_j(p)$, with density $f_j(p)$. The expected payoff consists of two terms, depending on whether or not the firm has the lowest price:

$$(1) \quad \pi_i^e(p) = \alpha \int_{p_L}^p y f_j(y) dy + p(1 - F_j(p)), \quad i, j = 1, 2, \quad i \neq j.$$

The first term on the right corresponds to the case where the firm has the smaller market, and the second term corresponds to the case where it has the larger share. The (imperfect) decision rule that we will use is the familiar logit rule:

$$(2) \quad f_i(p) = \frac{\exp(\pi_i^e(p)/\mu)}{\int_{p_L}^{p_H} \exp(\pi_i^e(y)/\mu) dy}, \quad i = 1, 2,$$

where μ is an "error parameter" that determines how sensitive firms are with respect to differences in expected profits. When μ is very large, payoff differences get washed out and non-optimal decisions become more likely, i.e. behavior becomes more random. At the other extreme, as μ tends to zero, the decision rule in (2) limits to the perfect-maximization rule; the best option is chosen with probability one. Note that (2) is not an explicit solution since the

⁵ It can be shown that there is no mixed-strategy Nash equilibrium when the maximum price is specified to be finite. There would be a mixed-strategy equilibrium if demand were inelastic at *any* price no matter how high, but it can be shown that the properties of this equilibrium are unintuitive, with a high value of α resulting in lower prices.

densities $f_i(p)$ on the left side also appear on the right side (through the expected payoff function that appears in the exponential terms).⁶ By differentiating both sides of (2) with respect to p , one obtains the "logit differential equation:"

$$(3) \quad \mu f_i'(p) = \pi_i^{e'}(p) f_i(p), \quad i = 1, 2.$$

Thus, in equilibrium, the density of decisions is increasing in p when the expected payoff function is increasing, and vice versa, so their relative maxima would coincide. Taking the derivative of the expected payoff in (1), and substituting the result in (3) provides a differential equation for the equilibrium choice density:

$$(4) \quad \mu f_i'(p) = [1 - F_j(p) - (1 - \alpha)p f_j(p)] f_i(p), \quad i, j = 1, 2, \quad i \neq j.$$

Existence of a solution to (4) is ensured by Theorem 1 of Anderson, Goeree, and Holt (1999). As is the case with the Nash equilibrium, the logit equilibrium is unique and symmetric:

Proposition 1. The logit equilibrium is unique and symmetric across players.

Proof. We first prove symmetry. Let F_1 and F_2 denote the distributions of players 1 and 2 respectively. Suppose, in contradiction, that the distribution functions are not everywhere the same. Without loss of generality, assume that $F_1 > F_2$ for some prices, as shown in Figure 1. Any region of divergence between the distribution functions will have a maximum *vertical* difference, as indicated by the vertical dashed line at price p^* . The first-order condition for the distance to be maximized at p^* is that the slopes of the distribution functions be identical at p^* , i.e. $f_1(p^*) = f_2(p^*)$. The second-order condition is that the slope of F_2 increases no slower than F_1 , i.e. $f_2'(p^*) \geq f_1'(p^*)$. However, since $F_1(p^*) > F_2(p^*)$ and $f_1(p^*) = f_2(p^*)$, equation (4) implies

⁶ Although it is reasonable to expect choice probabilities to be increasing functions of expected payoffs, the use of exponential functions seems somewhat arbitrary, despite the econometric convenience of the logit formulation. One motivation for using exponential functions in (2) is that additive changes in the payoffs for all decisions have no effect on choice probabilities, which is one of the requirements imposed by Luce (1959) in his axiomatic derivation of the logit rule. Alternatively, Anderson, Goeree, and Holt (1997) show that the logit rule results as a steady-state of a dynamic process in which players alter their decisions in the direction of better responses but make some error in doing so. In that model, gradient-based evolution with normal noise yields the logit choice rule with exponential functions.




Figure 1. A Configuration with $F_1 > F_2$.

that $f_1'(p^*) > f_2'(p^*)$, which yields the desired contradiction. Next we prove that there is at most one symmetric equilibrium. Suppose in contradiction that there are two symmetric equilibria, distinguished by "I" and "II" subscripts. Dropping the player-specific subscripts from (4) yields the following differential equations for the two candidate solutions:

$$(5) \quad \begin{aligned} \mu f_I' &= f_I(1 - F_I - (1 - \alpha) p f_I), \\ \mu f_{II}' &= f_{II}(1 - F_{II} - (1 - \alpha) p f_{II}). \end{aligned}$$

Without loss of generality, assume $F_I(x)$ is lower on some interval. Any region of divergence between the distribution functions will have a maximum *horizontal* difference, see Figure 2. The conditions for the horizontal distance to be maximized at height F^* are: $f_I'(p_I) = f_{II}'(p_{II})$, and that $f_I'(p_I) \geq f_{II}'(p_{II})$. However, since $p_I > p_{II}$, the logit differential equations in (5) imply that $f_I'(p_I) < f_{II}'(p_{II})$, a contradiction. Q.E.D.

Unlike the Nash equilibrium, the logit equilibrium will be sensitive to changes in the buyer inertial parameter, α , as can be seen from (4). The following proposition shows that the logit equilibrium has the intuitive property that prices rise when buyers are less responsive.

Proposition 2. In the logit equilibrium, an increase in the buyer inertia parameter α results in higher equilibrium prices in the sense of first-degree stochastic dominance.

Proof. Suppose that $\alpha_I < \alpha_{II}$, and let the corresponding symmetric equilibrium distributions be denoted by $F_I(p)$ and $F_{II}(p)$. The proof requires showing that $F_I(p)$ produces stochastically lower prices, i.e. that $F_I(p) > F_{II}(p)$. Suppose, in contradiction, that $F_I(p)$ is lower on some interval, as shown in Figure 2. Any region of divergence between the distribution functions will have a maximum *horizontal* difference, as indicated by the horizontal dashed line at the height of F^* , i.e. $F^* = F_I(p_I) = F_{II}(p_{II})$ for $p_I > p_{II}$. The first and second order conditions for the distance to be maximized at a height of F^* are that the slopes of the distribution functions be identical at F^* , i.e. $f_I(p_I) = f_{II}(p_{II})$, and that $f_I'(p_I) \geq f_{II}'(p_{II})$. In order to obtain a contradiction, recall that the distribution functions are determined by the logit differential equation in (4), evaluated at the appropriate level of α :

$$(6) \quad \begin{aligned} \mu f_I' &= f_I(1 - F_I - (1 - \alpha_I) p f_I), \\ \mu f_{II}' &= f_{II}(1 - F_{II} - (1 - \alpha_{II}) p f_{II}). \end{aligned}$$

Since $F_I(p_I) = F_{II}(p_{II})$ and $f_I(p_I) = f_{II}(p_{II})$, everything except for α_I , α_{II} , p_I , and p_{II} on the right sides of the equations in (6) are identical, when these equations are evaluated at p_I and p_{II} respectively. By assumption, $\alpha_I < \alpha_{II}$ and $p_I > p_{II}$, and it follows that $(1 - \alpha_I)p_I > (1 - \alpha_{II})p_{II}$. Therefore, the equations in (6) imply that $f_I'(p_I) < f_{II}'(p_{II})$, which contradicts the second-order condition for the maximum horizontal difference. Q.E.D.

This comparative statics result forms the basis for the experiment, which involved treatments with both low and high buyer inertia parameters: $\alpha = .2$ and $\alpha = .8$, with prices constrained to the interval [60, 160]. To obtain quantitative predictions about the level of average prices for these two treatments, we used a previously estimated error rate parameter,




Figure 2. A Configuration with $F_{II} > F_I$.

$\mu = 8.3$, from a different experiment (Capra *et al.*, 1999).⁷ For these parameter values, the differential equation in (4) can be solved numerically to obtain the logit equilibrium densities, see Figure 3. From these densities it is straightforward to calculate the average price and its standard deviation: 78 (± 7) for the $\alpha = 0.2$ treatment and 128 (± 6) for the $\alpha = 0.8$ treatment.⁸ In fact, the values for the treatment parameter α were chosen to ensure that price predictions for the final periods in the two treatments would be on opposite sides of the range of feasible prices.

It is natural to ask how the logit equilibrium compares to Radner's (1980) ε -equilibrium in this context. Recall that the two approaches are similar in the sense that both relax the assumption of perfect maximization, i.e. the assumption that behavior is completely determined by signs (not magnitudes) of payoff differences. However, while the continuous probabilistic choice approach produces a unique logit equilibrium, the discontinuous nature of the responses implied by an ε -equilibrium may result in multiple equilibria. In fact, for $\varepsilon < (1-\alpha)p_L$ there is

⁷ The game used in Capra *et al.* (1999) is of a similar degree of complexity as the game considered in this paper, and the experimental procedures (random matching, subject pools, etc.) were also comparable.

⁸ The standard deviation of the average is based on a sample size of 10, which is the number of subjects in each session.




Figure 3. Logit Equilibrium Densities for Both Treatments ($\mu = 8.3$)

no pure-strategy ε -equilibrium other than the Bertrand/Nash equilibrium, while for $\varepsilon \geq (1-\alpha)p_L$ there is a continuum of pure-strategy ε -equilibria in which the difference between the two prices does not exceed ε . To see this, note that the high price seller earns α times the low price, p_{\min} , so a deviation to undercut that price and sell a quantity of 1 would not occur if $p_{\min} - \alpha p_{\min} \leq \varepsilon$, or if $p_{\min} \leq \varepsilon/(1-\alpha)$. Similarly, the low-price will not bother to adjust upward if the two prices are within ε of each other. With a lower bound on prices, p_L , pure-strategy ε -equilibria of this type can exist if and only if $\varepsilon \geq (1-\alpha)p_L$ (to ensure that $p_{\min} \geq p_L$). Since in the experiment, $p_L = 60$, this lower bound on ε is unrealistically high (48 cents) for the high- α treatment, and it is 12 cents for the low- α treatment. In addition, there is a continuum of mixed-strategy ε -equilibria.⁹ For instance, the logit equilibrium described above for $\mu = 8.3$, can be interpreted as such a mixed-strategy ε -equilibrium for an ε as low as 3 cents for the high- α treatment (the expected payoff in the logit equilibrium is 110 while the expected payoff resulting from a best

⁹ See Baye and Morgan (1999) for a derivation of a particular family of mixed-strategy ε -equilibria for the standard Bertrand model.

response to the logit equilibrium is 113). To conclude, while the intuitive motivation for ε -equilibria is similar to that for the logit equilibrium, we prefer to use the latter since its uniqueness facilitates estimation and simulation of learning behavior. In addition, the logit equilibrium is the unique steady-state of a fictitious-play learning model, which is the topic of the next section.

III. A Geometric Fictitious Play Learning Model

The complicated appearance of the logit differential equation (4) raises the question of how boundedly rational players will conform to such a model. Our approach is to consider a naive learning model in which players use observations of rivals' past prices to update their beliefs about others' future actions. As before, the expected payoffs based on these beliefs determine players' choice probabilities via a logit rule. This model is used to simulate behavior in an experiment, in order to derive predictions about the nature of convergence to the logit equilibrium described above. We also use this dynamic model for econometric estimation of learning and error parameters, see section V below.

To obtain a tractable model, the feasible price range [60, 160] is divided into 101 one-cent categories. Players assign weights to each category and use observations of their rival's choices to update these weights as follows: each period *all* weights are "discounted" by a factor ρ and the discounted weight of the observed category is increased by 1. In other words, the weight, w , of an observed category is updated as $w \rightarrow \rho w + 1$, whereas the other weights are simply discounted by ρ : $w \rightarrow \rho w$. The belief probabilities in each period are obtained by dividing the weight of each category by the sum of all weights. Hence, the model is one of "geometric fictitious play" in which the learning parameter, ρ , determines the importance of new observations relative to previous information.¹⁰ Generally ρ will be between 0 and 1. When $\rho = 0$, the observations prior to the most recent one are ignored, and the model is one of Cournot best response. At the other extreme, when $\rho = 1$, the model reduces to "fictitious play" in which each observation is given equal weight, regardless of the number of periods that have elapsed

¹⁰ See e.g. Cheung and Friedman (1997) for a test of the geometric fictitious play learning model in 2×2 games. A generalization of the geometric fictitious play model can be found in Camerer and Ho (1999); see also Chen and Tang (1998).

since that observation. For intermediate values of ρ , the weight given to past observations declines geometrically over time.

The expected payoffs based on the updated beliefs in any period determine each player's decision probabilities via the logit rule in (2) with the integral replaced by a sum:

$$(7) \quad P_i(j|\rho) = \frac{\exp(\pi_i^e(j|\rho)/\mu)}{\sum_{k=1}^{101} \exp(\pi_i^e(k|\rho)/\mu)}, \quad i = 1, 2, \quad j = 1, \dots, 101,$$

where the ρ notation indicates the dependence of probabilities and expected payoffs on the learning parameter. In this dynamic model, beliefs and hence probabilistic price choices depend on the history of what has been observed up to that point. Since individual histories are realizations of a stochastic process, the predictions of this model will be stochastic and can be analyzed with simulation techniques.¹¹

The structure of the computer simulation program matches that of the experiment to be reported below: for each session or "run" there are 10 simulated subjects who are randomly matched in a sequence of 10 periods. We specify initial prior beliefs for each subject to be uniform on the integers in the set [60, 160]. These priors determine expected payoffs for each price, which in turn, determine the choice probabilities via the discrete logit rule in (7). The simulation begins by determining each simulated player's actual price choice for period 1 as a draw from the logit probabilistic response to the payoffs for priors that are uniform on [60, 160]. The simulated players are randomly divided into five pairs, and each one "sees" the other's actual price choice. These price observations are used to update players' beliefs using the naive learning rule explained above, with the estimated value of the learning parameter $\rho = 0.75$, which was also taken from Capra *et al.* (1999).¹² The updated beliefs, which become the priors for

¹¹ Previous computer simulations of adaptive learning in signaling games has been quite successful in explaining patterns of laboratory data, even where the observed human data converge to the equilibria that are ruled out by standard refinements of the Nash equilibrium. See Brandts and Holt (1996) who use adaptive learning with logistic decision error. Cooper, Garvin, and Kagel (1994) use adaptive learning with an assumption that some players are able to recognize and avoid dominated strategies.

¹² These simulations are used here to make *ex ante* predictions about dynamic adjustment paths. The data reported in the present paper will be used in section V to re-estimate these learning and error parameters.




Figure 4. Simulated Average Prices Obtained From a 1,000 Simulations (dark lines) Plus or Minus Two Standard Deviations (dotted lines) and a Typical Run (lines connecting squares)

period 2, will not all be the same if the simulated subjects encountered different price choices in period 1. Next, the process is repeated, with the period 2 priors determining expected payoffs, which in turn determine the logit choice probabilities, and hence the observed price realizations for that period. The whole process is repeated for 10 periods. Figure 4 shows the sequences of average prices (dark lines) obtained from 1,000 simulations together with plus or minus two standard deviations of the average (dotted lines). In addition, a typical "run" or simulation for each treatment (lines connecting squares) is shown. These simulation results predict that average prices decline in the $\alpha = 0.2$ treatment and stay the same in the $\alpha = 0.8$ treatment. The relation between the long-run steady-state of the simulated learning process and the logit equilibrium will be discussed in section V.

IV. The Experiment

The data were collected from six different groups of 10 subjects, who were recruited from

undergraduate economics classes at the University of Virginia. Participants were told that they would be paid a \$6 participation fee plus all additional money earned. Upon arrival, students were seated in isolated booths. After the reading of the instructions (reproduced in Appendix A for $\alpha = 0.8$), each group participated in a series of ten identical games during a session that lasted about one hour. As shown in Table I, there were two different treatments: $\alpha = 0.8$ for three groups of ten subjects, and $\alpha = 0.2$ for three other groups.

Table I. Experimental Design

	Treatment	Range of Feasible Prices
Sessions 1, 4, 5	$\alpha = 0.80$	\$0.60 to \$1.60
Sessions 2, 3, 6	$\alpha = 0.20$	\$0.60 to \$1.60

The price decisions were referred to as "numbers," and the earnings calculations were explained in the instructions without reference to any market context. At the start of each period, subjects were asked to record their choices for that period on their decision sheets. Numbers were required to be any amount between and including 60 and 160 cents, with decimals being used to indicate fractions of cents. Subjects were told that individual earnings would consist of two parts: a fixed payment of 25 cents per period and a percentage of the minimum number chosen by the corresponding pair of participants. The person choosing the lower number would receive the fixed payment plus the number chosen, while the subject choosing the higher number would receive the fixed payment, plus a fraction α of the lower number. If the numbers were equal, each player would receive the fixed payment of $\frac{1}{2}(1+\alpha)p_{\min}$, which is 90% of the minimum price for the high- α treatment and 60% of the low-alpha treatment.

Participants were told that the session would consist of 10 periods and that they would be randomly paired with another person in each period. After decisions for a period were made and the sheets were collected, draws of pairs of numbered ping pong-balls (without replacement) were used to determine the random matchings. The "other person's" decision was recorded on each person's decision sheet, and earnings were calculated before we returned the sheets at the end of the period. Subjects only saw the decision made by the person with whom they were matched in a given period, not the other's identity. Total earnings ranged from \$2.48 to \$12.40

per subject, to which we added a \$6.00 initial payment for coming on time.¹³




Figure 5. Average Prices by Session (dashed lines) and Treatment (dark line)

Figure 5 shows the average prices per period for sessions with $\alpha = 0.8$ and $\alpha = 0.2$. The averages for the three sessions in each treatment are plotted as dashed lines and the average prices for all sessions in a treatment are plotted as a solid line. The Nash equilibrium is 60 cents regardless of treatment, as indicated by the horizontal line at 60 cents near the bottom of the figure. A simple non-parametric test for the effect of the buyer inertia parameter α on average prices is based on the average prices in the final five periods. It is obvious from Figure 5 that the mean prices in the later periods of the three high- α sessions are all above the mean prices in the three low- α sessions. There are "six-take-three" = 20 ways that two groups of three objects can be ranked. Each of these 20 rankings would be equally likely under the null hypothesis of no treatment effect. Of these 20 rankings, we observed the one that is most

¹³ This experiment was followed by a series of one-shot matrix games that were unrelated to the game reported here. The earnings for these one-shot games generally boosted subjects' total earnings for the two-hour session into the \$20-\$30 range.

extreme in the direction of a treatment effect, so the probability of this is $1/20 = 5\%$. Thus the null hypothesis can be rejected at the 5% level, which is the same significance level that would be obtained by application of a Wilcoxon test. To summarize: *an increase in the buyer inertia parameter, α , results in higher prices, a result that is not predicted by Nash but is consistent with a logit equilibrium analysis.*

In contrast to the observed prices in low- α treatment, average prices in the high- α treatment almost always stay above 100 cents. The averages in the final five periods for both treatments are given in Table II. The actual averages are remarkably close to the logit equilibrium predictions of 78 (± 7) and 128 (± 6) that were based on equation (5) with an error parameter estimated from a different experiment.¹⁴

Table II. Average Prices in Periods 6-10 (Standard Deviations)

	session 1	session 2	session 3	Pooled	Logit Predictions
low- α treatment	63 (14)	72 (20)	73 (32)	69 (13)	78 (7)
high- α treatment	102 (14)	126 (31)	134 (17)	121 (13)	128 (6)

When described in words, these results are not surprising; that is, prices should be lower when buyers are more responsive to prices. What is somewhat surprising is the accuracy of the logit point predictions that were based on out-of-sample data. The dynamic predictions derived from the naive learning model are also quite accurate, as is apparent from the actual average price sequences for each treatment in Figure 5. There is a tendency for average prices in the low- α treatment to decrease and for those in the high- α treatment to stay roughly constant, as predicted by the computer simulations in Figure 4. Even the price averages in the very first period of the experiment, (75,90,100) and (105,123,127) for the low and high inertia treatments, are quite close to the period 1 predictions of 93 (± 6) and 126 (± 7) (numbers between parentheses again denote standard deviations of averages based on a sample size of $n = 10$) that are based

¹⁴ Notice that for both the low and high- α treatment, 2 of the 3 session averages are within one standard deviation of the logit prediction, and that only one session average is more than two standard deviations off.

on a logit best response to an admittedly arbitrary uniform belief density on $[60, 160]$.¹⁵ To summarize: *computer simulations of the learning model with a diffuse initial prior explain both the final-period price levels and the most salient feature of the adjustment pattern, i.e. the flat trajectory of prices in the high- α treatment and the clear price declines in the low- α treatment.*

V. Econometric Estimates of Learning and Error Parameters

As noted above, it is straightforward to obtain numerical solutions for the equilibrium densities for specific values of the error parameter μ , given the buyer responsiveness parameter α . This solution then can be used to calculate the densities associated with the particular prices selected by all subjects under each treatment. These densities, in turn, are used to calculate the value of the likelihood function, i.e. the likelihood of drawing the prices actually selected, given the assumed value of μ .¹⁶ Then a grid search yields a maximum likelihood estimate of $\mu = 6.7$ with a standard error of 0.3 (see Table III). This is of the same magnitude as the value of the error parameter (8.3) used from Capra *et al.* (1999) to calculate the logit equilibrium predictions in section II above, which explains the accuracy of those predictions. It should be noted that the equilibrium model was estimated using only the data for the final five periods, since it is apparent from Figure 5 that the data do not stabilize until after period 5.

The equilibrium approach determines beliefs via equilibrium (consistency-of-action-and-beliefs) conditions. As noted above, an alternative is to model the learning process that generates beliefs, using data from all periods. To estimate the magnitude of error and learning parameters, we estimated the dynamic model from Section III, using data for *all* periods, since much of the learning occurs in early periods.¹⁷ The probability $P_i(j|\rho)$ that player i chooses a price in the

¹⁵ With a uniform belief distribution the expected payoff in (1) becomes $\pi_i^e(p) = -(2-\alpha)/200 (p-160/(2-\alpha))^2$, and the logit density is therefore a (truncated) normal on $[60, 160]$ with mean $160/(2-\alpha)$.

¹⁶ The equilibrium model stipulates that the price decisions are independent draws from a stationary distribution. Thus the likelihood is the product of the densities evaluated at each of the prices selected by each subject in the last 5 periods.

¹⁷ The data are reproduced in Appendix B. In each of the six tables, a column corresponds to one of the ten subjects, S1,...,S10, and a row to one of the ten periods. The tables list both the actual decisions of subjects and the choice made by the subject they were matched with (given by the number in parentheses). The latter information is used by players to update their beliefs in each period.

j th category is given by (7), and the likelihood function is simply a product of such probabilities over all periods and all players. The log-likelihood function is therefore given by the sum:

$$(8) \quad \text{Log}(L) = \sum_{i=1}^{10} \sum_{t=1}^{10} \log(P_i(p_{i,t}|\rho)),$$

where $p_{i,t}$ denotes player i 's price in period t . Table III gives the maximum-likelihood estimation results. The overall estimate of the error parameter is: $\mu = 8.4$, which is roughly in line with previous estimates we have obtained for other data sets.¹⁸ One obvious conclusion to be drawn from the table is that the null hypothesis of no error ($\mu = 0$) can be rejected at very low significance levels. This is a rejection of the perfect-rationality assumption that would lead to the Bertrand-Nash outcome.

Table III. Maximum-Likelihood Estimates

	error parameter, μ	learning parameter, ρ
logit equilibrium model	6.7 (.3)	NA
learning model	8.4 (.4)	.72 (.03)

The dynamic model applied to all of the data provides an estimate of the same magnitude as the error parameter (6.7) that was obtained from the equilibrium model (Table III). Recall that the equilibrium model was estimated for the last five periods with a completely different approach, i.e. solving for the equilibrium beliefs instead of modeling them as evolving over time in response to experience. The dynamic model takes into account the "history dependence" that

¹⁸ Both of these estimates are approximately the same as the error parameter estimates obtained for data from a "traveler's dilemma" experiment reported in Capra, Goeree, Gomez, and Holt (1997): $\mu = 8.3 (\pm 0.3)$ for the equilibrium model and $\mu = 10.9 (\pm 0.5)$ for the dynamic model, where the standard errors are in parentheses. In general, the amount of noise in the data should depend on the subject pool, the complexity of the game and the importance of un-modeled factors in the decision-making process. The traveler's dilemma experiments were also conducted at the University of Virginia with a 10-period length and random matching. In comparing μ estimates across experiments, it is important to adjust for the way that payoffs are measured. We typically express payoffs in pennies, but if payoffs are in dollars, the μ estimates would be 100 times smaller than the ones obtained from entering the payoffs in pennies, i.e. 0.085 instead of 8.5. Another caveat is that McKelvey and Palfrey (1995) and others report the reciprocal of μ , which is therefore a precision parameter for which high values indicate low error.

may cause individual decisions to fail to be independent in a statistical sense, and hence is useful for econometric purposes.

Given that the error structure is similar for both the equilibrium and learning models, one might wonder what the learning model implies about the long-run steady-state distribution of price decisions. In particular, will geometric learning generate a price distribution that corresponds to the logit equilibrium distribution? To investigate these issues, it is useful to consider the extreme cases of the geometric learning model, i.e. standard fictitious play ($\rho = 1$) and Cournot best response ($\rho = 0$). Since there is no "forgetfulness" in fictitious play, any steady state distribution of decisions will eventually be fully learned by all players, i.e. the empirical frequencies of price draws from the distribution will converge to that distribution. Since, in this case, each player is making a logit probabilistic best response to the empirical distribution, a steady-state of the fictitious play learning model is necessarily a logit equilibrium. Our estimate of the learning parameter, $\rho = .72$, however, implies that only the most recent five or six observations receive much weight, so the above argument does not guarantee that the logit equilibrium is also a steady state of the geometric learning process. We simulated the learning process with a cohort of 10 randomly matched players, using $\alpha = .8$ and the estimated $\rho = .72$. After 1,000 periods, the empirical choice frequencies for all ten players were quite close to the logit equilibrium distribution, but there was a slight tendency for the empirical distributions to be too flat at the mode. This can be seen from Figure 6, where the logit prediction is the thin line and the simulated frequency for a typical simulated player are plotted as the thick line.

The intuition behind the extra flatness in the simulated data with $\rho < 1$ can be understood by considering the other extreme of $\rho = 0$, i.e. Cournot type beliefs that only place weight on the most recent observation. To determine whether a logit equilibrium would be a steady state for this case, recall that the noise-free Cournot best response to any price observation is just that observation (minus ϵ), so the distribution of noise-free best responses to draws from a logit equilibrium distribution will be that distribution. But with $\mu > 0$, there is some noise in the responses, which would cause a logit equilibrium to be mapped into a more dispersed distribution. This suggests that the steady state of the Cournot learning model would be flatter than the logit prediction for a given value of $\mu > 0$. This intuition is consistent with the simulations that we have run, as can be seen in Figure 7, where the flatter solid line represents




Figure 6. The Logit Equilibrium Distribution (Thin Line) and Simulated Distribution of Prices Under Geometric Fictitious Play with $\rho = .72$ (Thick Line)

the empirical choice frequencies of a representative player and the other solid line is the logit equilibrium. This disparity raises the issue of exactly what is the long-run steady state of the learning process, i.e. the "stochastic learning equilibrium." With Cournot learning, beliefs are determined by the most recent draw, so the steady state is a distribution that is a noisy best response to point beliefs determined by draws from the same distribution. Seen this way, the stochastic learning equilibrium can be defined as the fixed point of an integral equation, which can be used to calculate the theoretical long-run steady-state distribution. This distribution is shown as the dotted line in Figure 7, which is almost indistinguishable from the solid line that shows the simulated price frequencies for $\rho = 0$.¹⁹

¹⁹ The relevant integral equation that determines the stochastic learning equilibrium is:

$$f(p) = \int_{p_L}^{p_H} LBR(p, p') f(p') dp'$$




Figure 7. The Logit Equilibrium Distribution (Solid Line)
 Simulated Distribution of Prices Under Cournot Type Beliefs (Solid Line)
 and the Steady-State Learning Equilibrium for $\rho = 0$ (Dotted Line)

VI. Conclusion

Many strategic situations of interest to economists have the property that the minimum decision drives all players' payoffs. This paper considers a game in which the firm with the lowest price captures a larger market share, but a fixed fraction of buyers have price-protection contracts that allows their current seller to match a competitor's price cut. Simple intuition suggests that the resulting buyer inertia may be anti-competitive, since increases in the fraction of buyers who do not switch to the low-price seller will reduce the profitability of unilateral price cuts. In contrast, the unique Nash equilibrium for this game produces the Bertrand outcome, irrespective of non-critical changes in buyer inertia. This paper uses theory, experiments, and

where the logit best response $LBR(p, p')$ to point-mass beliefs at p' is given by

$$LBR(p, p') = \frac{\exp(\pi(p, p')/\mu)}{\int_{p_i} \exp(\pi(q, p')/\mu) dq}$$

simulation methods to evaluate the tension between economic intuition and the sharp predictions made by a standard game-theoretic analysis.

Our previous work on related games indicates that subjects in laboratory experiments are not perfectly rational, in the sense that estimated error parameters are significantly greater than zero. We used these parameter estimates to run *ex ante* computer simulations of experiments with ten periods of random matchings. These simulated price distributions were then used to design the experiments with human subjects that are reported here. The human data contain some surprises, but the big picture is that the average price decision sequences for our student subjects conform to the predictions based on the simulations: prices fall to low levels when a high fraction of the demand goes to the low-priced seller, whereas prices stay approximately level in the upper half of the range when a low fraction of demand goes to the low-price seller. The simulations use learning and error parameters estimated from the data of a previous experiment (Capra *et al.*, 1999) to explain why behavior in one treatment falls towards the unique Nash prediction and why behavior in the other treatment stays well away from the Nash prediction. The fact that Nash works well for one parametrization and not for another suggests that this equilibrium concept should be generalized.

The overall price averages for each treatment with human subjects are well explained by a stochastic generalization of the Nash equilibrium, the logit equilibrium proposed by McKelvey and Palfrey (1995). We use maximum-likelihood methods to estimate a logit error parameter, using the equilibrium model that imposes consistency of action and belief distributions, and using a dynamic model in which beliefs evolve over time via a naive Bayesian learning process. The parameter estimates from these two very different procedures are quite close, and are of the same magnitude as estimates we have obtained in experiments for different games with similar (ten-period, random matching) procedures.

Appendix A: Instructions

You are going to take part in an experimental study of decision making. The funding for this study has been provided by several foundations. The instructions are simple, and by following them carefully, you may earn a considerable amount of money. At this time, you will be given \$6 for coming on time. All the money that you earn subsequently will be yours to keep, and your earnings will be paid to you in cash today at the end of this experiment. We will start by reading the instructions, and then you will have the opportunity to ask questions about the procedures described.

Earnings

The experiment consists of a sequence of periods. In each period, you will be randomly matched with another participant in the room. The decisions that you and the other participant make will determine the amount earned by each of you. At the beginning of each period, you will be asked to choose a number between 60 and 160 cents and write down it on a decision sheet that is attached to these instructions. The number you can choose may be any amount between and including 60 and 160 cents. That is, we allow fractions of cents. The person who you are matched with will also choose a number between and including 60 and 160 cents. Each of you receives 25 cents plus an amount that depends on the number chosen. This amount equals a percentage of the minimum of your number and the other's number. If the numbers are equal, then the percentage for you and the other person each equals the 90% of the number chosen. If you are the person choosing the lower number, the percentage you receive equals the 100% of the number you chose. If you are the person choosing the higher number, the percentage you receive equals the 80% of the other person's number.

Example: Suppose that your number is X and the other's number is Y .

If $X = Y$, you get $0.9 \cdot X$, and the other gets $0.9 \cdot Y$.

If $X > Y$, you get $0.8 \cdot Y$, and the other gets Y .

If $X < Y$, you get X , and the other gets $0.8 \cdot X$.

(In each case, 25 cents will be added to your earnings.)

Record of Results

Now, each of you should examine the record sheet. This sheet is the last one attached to these instructions. Your identification number is written in the top-right part of this sheet. Please look at the columns of your record sheet. Going from left to right, you will see columns for the "period," "your number," "other's number," "minimum number," "your earnings," and "plus 25 cents." You begin by writing down your own number in the appropriate column. As mentioned above, this number must be greater than or equal to 60 and less than or equal to 160 cents, and the number may be any amount in this range, (i.e. fractions of cents are allowed). Use decimals to separate fractions of cents. For example, $wx.z$ cents indicates wx cents, and a fraction $z/10$ of a cent. Similarly, $x.z$ cents indicates x cents and a fraction $z/10$ of a cent.

After you make and record your decision for period one, we will collect all decision sheets. Then we will draw numbered ping pong balls to match each of you with another person. Here we have a container with ping pong balls, each ball has one of your identification numbers on it. We will draw the ping pong balls to determine who is matched with whom. After we have matched someone with you, we will write the other's number, the minimum number, and your earnings in the relevant columns

of your decision sheet and return it to you. Then, you make and record your decision for period two, we collect all decision sheets, draw ping pong balls to randomly match you with another person, write the other's number, minimum number, and earnings in your decision sheet and return it to you. This same process is repeated a total of ten times.

Summary

To begin, each participant chooses and records a number by writing it in the appropriate column of the decision sheet. Then the decision sheets are collected and participants are randomly matched using draws of numbered ping pong balls. Once the matching is done, the other's decision, the minimum number, and the earnings are written on each person's decision sheet. After all decision sheets are returned, participants choose and record their numbers for the next period. The decisions determine each person's earnings as described above. You will receive an amount that equals 25 cents plus a percentage of the minimum of your number and the other's number. If you chose the lower number, the percentage will be 100% of the minimum number. If you chose the higher number, the percentage will be 80% of the minimum number. And if you and the other person chose the same number, the percentage will be 90% of that number. Note that a new random matching is done in each period.

Final Remarks

This experiment will be followed by another, quite different experiment in which you will have additional opportunity to make decisions that can increase your earnings.

At the end of today's session, we will pay to you, privately in cash, the amount that you have earned. You have already received the \$6 participation payment. Therefore, if you earn an amount X during the exercise that follows, you will receive a total amount of $\$6.00 + X$. Your earnings are your own business, and you do not have to discuss them with anyone.

During the experiment, you are not permitted to speak or communicate with the other participants. If you have a question while the experiment is going on, please raise your hand and one of us will come to your desk to answer it. At this time, do you have any questions about the instructions or procedures? If you have a question, please raise your hands and one of us will come to your seat to answer it.

Identification Number: _____

Please choose a number that is greater than or equal to 60 and less than or equal to 160, using decimals to indicate fractions; e.g. wx.z or x.z.

Period	Your number	Other's number	Minimum number	Your earnings	Plus 25 cents
1					
2					

Appendix B: Individual Decisions

Session 1: $\alpha = 0.80$

period	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
1	140 (158)	75 (150)	130 (120)	160 (110.6)	150 (75)	110.6 (160)	150 (80)	80 (150)	158 (140)	120 (130)
2	140 (115)	150 (120)	130 (120)	160 (95)	120 (130)	110.6 (150)	150 (110.6)	95 (160)	120 (150)	115 (140)
3	159.9 (155)	150 (109)	110 (105)	155 (159.9)	109 (150)	110.6 (130)	150.5 (122)	105 (110)	130 (110.6)	122 (150.5)
4	140 (110)	160 (100)	110 (150.5)	155 (109.9)	109.9 (155)	150.5 (110)	151.7 (145)	100 (160)	110 (140)	145 (151.7)
5	140 (115)	150 (157)	120 (125)	157 (150)	109.9 (149)	149.9 (152.8)	152.8 (149.9)	115 (140)	125 (120)	149 (109.9)
6	110 (120)	160 (109.8)	120 (110)	155 (120)	109.8 (160)	149.9 (153.8)	153.8 (149.9)	125 (133)	120 (155)	133 (125)
7	120 (151.7)	160 (120)	120 (160)	149 (125)	109.8 (135)	151.7 (120)	154.8 (130)	130 (154.8)	135 (109.8)	125 (149)
8	130 (151.7)	160 (130)	120 (109.6)	158 (128)	109.6 (120)	151.7 (130)	145 (135)	135 (145)	130 (160)	128 (158)
9	120 (132)	160 (130)	115 (155)	155 (115)	109.5 (130)	135.9 (140)	130 (160)	140 (135.9)	130 (109.5)	132 (120)
10	125 (157)	160 (138)	110 (122)	157 (125)	109.5 (149.9)	149.9 (109.5)	135 (109)	138 (160)	109 (135)	122 (110)

Session 2: $\alpha = 0.20$

period	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
1	61.13 (60)	60 (95)	95 (60)	92 (75)	60 (61.13)	90 (61)	60 (99.9)	61 (90)	99.9 (60)	75 (92)
2	60 (60)	60 (60)	70 (99.9)	80 (65)	60 (75)	80 (60)	60 (80)	65 (80)	99.9 (70)	75 (60)
3	60 (65)	60 (64)	78 (69.9)	60 (75)	65 (60)	60 (60)	60 (60)	64 (60)	69.9 (78)	75 (60)
4	60.25 (60)	60 (65)	68 (65)	65 (68)	60 (60)	60 (60.25)	60 (69.8)	60 (60)	69.8 (60)	65 (60)
5	60 (72.99)	60 (60)	72.99 (60)	70 (60)	60 (60)	60 (60)	60 (70)	60 (69)	69 (60)	60 (60)
6	63 (60)	60 (63)	63 (60)	65 (65)	60 (65)	70 (60)	65 (65)	65 (60)	60 (70)	60 (63)
7	65.1 (60)	60 (60)	60 (60)	64.9 (60)	60 (60)	60 (64.9)	60 (60)	60 (65.1)	60 (60)	60 (60)
8	61 (60)	60 (60)	60 (61)	60 (60)	60 (63)	64.8 (60)	60 (64.8)	63 (60)	60 (60)	60 (60)
9	60 (60)	60 (60)	60 (60)	60 (60)	60 (60)	60 (60)	60 (60)	60 (60)	60 (60)	60 (60)
10	60 (60)	60 (60)	60 (60)	159 (60)	60 (60)	60 (60)	60 (60)	60 (60)	60 (159)	60 (60)

Session 3: $\alpha = 0.20$

period	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
1	120 (125)	69 (65)	85 (95.8)	65 (69)	95.8 (85)	68 (100)	120 (155.5)	155.5 (120)	100 (68)	125 (120)
2	100 (60.5)	81 (95.1)	95.1 (81)	94 (120)	82.6 (94)	94 (82.6)	120 (94)	60.5 (100)	60 (110)	110 (60)
3	60 (90.3)	62 (75.5)	90.3 (60)	90 (87.1)	87.1 (90)	137 (60)	110 (80)	75.5 (62)	60 (137)	80 (110)
4	80 (75.4)	69 (60)	75.4 (80)	85 (89.9)	89.9 (85)	142 (60.5)	100 (80)	70.5 (142)	60 (69)	80 (100)
5	65 (98)	64 (80)	82.3 (60)	88.5 (70)	84.9 (75)	98 (65)	75 (84.9)	70 (88.5)	60 (82.3)	80 (64)
6	70 (75.5)	65 (70)	72.4 (62.3)	160 (75)	77.7 (80)	62.3 (72.4)	75 (160)	75.5 (70)	70 (65)	80 (77.7)
7	75 (60)	67 (61.8)	60.5 (60)	72.5 (160)	79.6 (75)	61.8 (67)	75 (79.6)	160 (72.5)	60 (60.5)	60 (75)
8	75 (60)	60 (95.9)	60 (75)	95.9 (60)	70.2 (69.9)	61.7 (60)	75 (60)	69.9 (70.2)	60 (61.7)	60 (75)
9	75 (73.2)	60.64 (60)	73.2 (75)	65 (60)	62.9 (60)	60 (62.9)	70 (70)	70 (70)	60 (60.64)	60 (65)
10	75 (60)	60.64 (66.7)	74.5 (60)	60 (62)	62 (60)	60 (75)	66.7 (60.64)	69.9 (60)	60 (74.5)	60 (69.9)

Session 4: $\alpha = 0.80$

period	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
1	150 (90)	145 (100)	125 (120)	108.9 (150)	160 (81)	81 (160)	100 (145)	150 (108.9)	120 (125)	90 (150)
2	150 (125)	95 (100)	125 (150)	128.9 (80)	100 (95)	84 (120)	61 (130)	80 (128.9)	130 (61)	120 (84)
3	145 (100)	99 (95)	120 (100)	118.9 (119.9)	100 (120)	95 (99)	100 (120)	120 (100)	119.9 (118.9)	100 (145)
4	125 (118.9)	95 (109.9)	119.9 (90)	118.9 (125)	100 (97)	97 (100)	109 (125)	90 (119.9)	109.9 (95)	125 (109)
5	125 (118.9)	100 (150)	99.98 (100)	118.9 (125)	100 (119)	98.5 (100)	119 (100)	100 (98.5)	100 (99.98)	150 (100)
6	125 (99.9)	125 (95)	94.98 (99.5)	114.9 (100)	100 (114.9)	99.5 (94.98)	99.9 (100)	95 (125)	99.9 (125)	100 (99.9)
7	125 (100)	100 (90.1)	89.49 (96)	114.9 (100)	100 (114.9)	90.1 (100)	99.5 (99.9)	96 (89.49)	99.9 (99.5)	100 (125)
8	120 (120)	95 (125)	89.48 (99.4)	118.9 (95)	120 (120)	93.5 (98.8)	98.8 (93.5)	95 (118.9)	99.4 (89.48)	125 (95)
9	120 (100)	95 (95)	89.48 (101)	89.99 (89.8)	101 (89.48)	94.9 (89.9)	89.8 (89.99)	95 (95)	89.9 (94.9)	100 (120)
10	100 (93.5)	94.9 (89.9)	89.47 (78.99)	78.99 (89.47)	100 (94)	93.5 (100)	84.49 (160)	94 (100)	89.9 (94.9)	160 (84.49)

Session 5: $\alpha = 0.80$

period	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
1	150 (160)	75.5 (100)	97 (80)	160 (60)	160 (150)	67 (100)	60 (160)	80 (97)	100 (75.5)	100 (67)
2	155 (128)	75.5 (160)	128 (155)	159.9 (98)	160 (75.5)	98 (159.9)	160 (65)	65 (160)	99 (88)	88 (99)
3	156 (135)	95.5 (160)	135 (159)	159.1 (136)	150 (99.9)	136 (159.1)	160 (95.5)	135 (156)	99.9 (150)	159 (135)
4	157.9 (160)	160 (157.9)	145 (159)	140 (160)	160 (140)	67 (70)	159 (145)	70 (67)	99.99 (120)	120 (99.99)
5	159 (158)	155 (140)	150 (66)	159.5 (99.7)	140 (155)	66 (150)	158 (159)	75 (98)	99.7 (159.5)	98 (75)
6	160 (72)	149.5 (62)	145 (140)	155 (157)	140 (145)	62 (149.5)	157 (155)	85 (99.5)	99.5 (85)	72 (160)
7	160 (91)	139 (140)	140 (139)	150 (156)	140 (152)	152 (140)	156 (150)	86 (99.4)	99.4 (86)	91 (160)
8	159 (140)	139 (140)	140 (139)	150 (93)	140 (159)	160 (155)	155 (160)	77 (118)	118 (77)	93 (150)
9	156 (90)	139.5 (92)	138 (154)	130 (117)	150 (60)	60 (150)	154 (138)	90 (156)	117 (130)	92 (139.5)
10	155 (138.5)	138.5 (155)	145 (94)	125 (140)	140 (125)	60 (100)	150 (116)	100 (60)	116 (150)	94 (145)

Session 6: $\alpha = 0.20$

period	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
1	100 (160)	60 (80.5)	82 (60)	160 (100)	75 (72.3)	109.9 (100)	72.3 (75)	60 (82)	80.5 (60)	100 (109.9)
2	90 (60)	60 (69.9)	75 (60)	60 (99.9)	60 (75)	99.9 (60)	69.9 (60)	100 (60)	60 (100)	60 (90)
3	70 (160)	60 (99.8)	60 (60)	70 (60)	74.9 (94.9)	99.8 (60)	160 (70)	60 (60)	94.9 (74.9)	60 (70)
4	80 (160)	60 (60)	60 (60)	60 (99.8)	84.9 (68)	99.8 (60)	68 (84.9)	160 (80)	60 (85)	85 (60)
5	100 (60)	60 (70.2)	60 (70)	60 (100)	65 (60)	60 (65)	70.2 (60)	60 (60)	60 (60)	70 (60)
6	79 (60)	60 (79)	60 (60)	60 (60)	65 (60)	60 (60)	70 (60)	60 (60)	60 (60)	60 (70)
7	60 (60)	60 (60)	60 (60)	60 (60)	160 (60)	60 (60)	60 (60)	60 (160)	60 (60)	60 (60)
8	60 (60)	60 (60)	60 (60)	60 (160)	160 (60)	60 (60)	60 (160)	160 (60)	60 (60)	60 (60)
9	60 (60)	60 (60)	60 (60)	60 (60)	160 (60)	60 (160)	72 (60)	60 (60)	60 (60)	60 (72)
10	60 (160)	60 (60)	60 (60)	60 (60)	160 (60)	60 (60)	60 (60)	160 (60)	60 (160)	60 (60)

References

- Akerlof and Yellin (1985) "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?," *American Economic Review*, 75, 708-720.
- Anderson, Simon P., André de Palma, and Jacques-François Thisse (1992) *Discrete Choice Theory of Product Differentiation*, Cambridge: MIT Press.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1997) "Stochastic Game Theory: Adjustment and Equilibrium with Bounded Rationality," working paper, University of Virginia.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1999) "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies," Working Paper, University of Virginia.
- Baye, Micheal R. and John Morgan (1999) "Bounded Rationality in Homogeneous Product Pricing Games," Discussion Paper, Indiana University.
- Brandts, Jordi and Charles A. Holt (1996) "Naive Bayesian Learning and Adjustment to Equilibrium in Signalling Games," working paper, University of Virginia.
- Camerer, Colin and Teck-Hua Ho (1999) "Experience Weighted Attraction Learning in Normal-Form Games," *Econometrica*, 67, 827-874.
- Capra, C. Monica, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt (1999) "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, 89(3), June, 678-690.
- Cooper, David J., Susan Garvin, and John H. Kagel (1994) "Adaptive Learning vs. Equilibrium Refinements in an Entry Limit Pricing Game," *RAND Journal of Economics*, 106(3), 662-683.
- Chen, Yan and Fang-Fang Tang (1998) "Learning and Incentive Compatible Mechanisms for Public Goods Provision: An Experimental Study," *Journal of Political Economy*, 106, 633-662.
- Cheung, Yin-Wong and Daniel Friedman (1997) "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, 19, April, 46-76.
- Diamond, Peter A. (1971) "A Model of Price Adjustment," *Journal of Economic Theory*, 3, 158-168.
- Erev, Ido and Alvin E. Roth (1997) "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88(4), 848-881.

- Friedman, James (1971) "A Non-cooperative Equilibrium for Supergames," *Review of Economic Studies*, 78, 1-12.
- Holt, Charles A. and David Scheffman (1987) "Facilitating Practices: The Effects of Advance Notice and Best-Price Policies," *RAND Journal of Economics*, 18(2), 187-197.
- Kreps, David and Jose Scheinkman (1983) "Quantity Precommitment and Bertrand Competition and Yield Cournot Outcomes," *Bell Journal of Economics*, 14, Autumn 1983, 326-337.
- Luce, R. Duncan (1959) *Individual Choice Behavior*, New York: John Wiley & Sons.
- McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- Porter, Robert H. (1983) "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886," *Bell Journal of Economics*, 14, 301-314.
- Radner, Roy (1980) "Collusive Behavior in Oligopolies with Long But Finite Lives," *Journal of Economic Theory*, 22, 136-156.
- Rosenthal, Robert W. (1989) "A Bounded Rationality Approach to the Study of Noncooperative Games," *International Journal of Game Theory*, 18, 273-292.
- Roth, Alvin and Ido Erev (1995) "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164-212.
- Selten, Reinhard and Joachim Buchta (1998) "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions," in *Games and Human Behavior, Essays in Honor of Amnon Rapoport*, D. Budescu, I. Erev, and R. Zwick, eds., New Jersey: Lawrence Erlbaum Associates, Inc., 105-142.
- Siegel, Sidney and N. John Castellan, Jr. (1988) *Nonparametric Statistics for the Behavioral Sciences*, New York: McGraw-Hill.
- Simon, Herbert A. (1957) *Models of Man*, New York: John Wiley & Sons.