

Durable-goods monopoly: laboratory market and bargaining experiments

Stanley S. Reynolds*

Results from single-period monopoly experiments (nondurable environment) are compared with results from multiperiod monopoly experiments that have features of a durable-goods environment. Average prices were below the static monopoly benchmark price in all settings. Observed initial prices were higher in multiperiod experiments than in single-period experiments, in contrast to equilibrium predictions. Prices in multiperiod experiments tended to fall over time; there was less price cutting in market experiments than in bargaining experiments. There was substantial demand withholding by buyers in multiperiod experiments. A version of bounded rationality is a promising candidate for explaining deviations from equilibrium predictions.

1. Introduction

■ In a seminal article, Ronald Coase (1972) pointed out how product durability may limit a seller's monopoly power. A prospective durable-good consumer who expects decreasing price offers over time would be willing to wait to purchase as long as the expected gain from the price decrease exceeds the opportunity cost of waiting to purchase. This may induce a monopoly seller to set relatively low prices. This line of argument led Coase to propose what has come to be known as the Coase conjecture. If the product is perfectly durable and if consumers are sufficiently patient, then monopoly prices will converge to marginal cost almost instantaneously.

The significance of product durability and monopoly power was recognized even before Coase's article. Major U.S. antitrust cases such as *U.S. v. Aluminum Co. of America* (1945) and *U.S. v. United Shoe Machinery Corp.* (1953) and criminal and civil antitrust cases involving IBM (see Fisher, McGowan, and Greenwood, 1983) were centered on issues of monopoly power and product durability. These issues continue to play a prominent role in public policy toward markets for durable products, such as markets for computer hardware and software and for consumer electronic goods, when monopoly power may be involved.

* University of Arizona; reynolds@bpa.arizona.edu.

I am grateful for suggestions and comments on various drafts from Mark Bagnoli, Bruno Broseta, Catherine Eckel, Charles Holt, Amnon Rapoport, Stephen Rassenti, three anonymous referees, and the Editor. Financial support was provided by the University of Arizona Small Grants Program and National Science Foundation grant no. SBR-9809110. Thanks to Timothy O. Dang, Viliam Druska, Zeynep Hansen, and Todd Swarthout for research assistance. I owe a special acknowledgment to Mike Miller, who provided valuable research assistance with the experiments and data analysis. Mike died unexpectedly in 1997.

Durable-goods monopoly theory has become a standard part of the set of tools of analysis for Industrial Organization economists.¹ However, predictions from this body of theory are sensitive to fine differences in model assumptions; such differences are likely to be difficult to observe in naturally occurring markets. For example, differences in assumptions about the period of time over which a seller can commit to a fixed price offer, the length of the time horizon for decision making, and the amount of information the seller has about buyers' willingness to pay can lead to radically different price predictions. At present there is relatively little empirical evidence about the predictive power of this theory—even about the most basic elements of the theory.²

Laboratory methods permit a researcher to control critical features of the environment (such as information about other agents' payoffs, time discounting, length of time horizon) that theoretical models identify as key determinants of outcomes. This article reports on laboratory experiments involving a monopolist who sells to one or more buyers who have values that are randomly assigned. Product "durability" is induced via the assignment of buyer values and the time structure of price offers and purchasing opportunities. The main issue is whether product durability serves to limit the monopoly power of the seller and induce prices below those set by a monopolist selling a nondurable good. A related issue is whether buyers wait to make purchases (this is referred to as *demand withholding* in the experimental literature) in a way that facilitates reduced monopoly power. This article reports results from laboratory environments with varying numbers of buyers and varying numbers of trading periods. The experiments are designed so that asymmetric information equilibrium theory can provide predictions of buyer and seller behavior.

The remainder of the article is organized as follows. Relevant theoretical background and evidence on durable-goods monopolies are discussed in Section 2. The experimental design and procedures are explained in Section 3. The predictions appear in Section 4. Section 5 provides a description and analysis of the results. Section 6 contains a discussion of the results and potential explanations for the results. Concluding comments are in Section 7.

2. Monopoly and durability

■ **Theoretical background.** Bulow (1982) compares outcomes for a monopoly seller and a monopoly lessor in a two-period setting. He shows that a seller sets lower prices (per period of usage) and earns lower profit than a lessor. This result is driven by buyers who understand that the seller has an incentive to offer a price reduction in the second period, after making sales to buyers with relatively high values in period one. Stokey (1981) and Gul, Sonnenschein, and Wilson (1986) show that when the time horizon is infinite, there is an equilibrium in which prices are arbitrarily close to seller marginal cost if the discount factor is sufficiently close to unity (equivalently, if the period of commitment is sufficiently short). This result seems to capture the spirit of the Coase conjecture.

Two situations have been identified for which durability does not necessarily limit a seller's monopoly power. Ausubel and Deneckere (1989) show that when the time horizon is infinite there are many equilibria, including an equilibrium in which the seller is able to maintain prices close to the static monopoly price. This result holds for a market model with a continuum of buyers and for a bargaining model with a

¹ All of the following articles published recently in this journal use durable-goods monopoly theory as a point of departure: Karp and Perloff (1996), Kühn and Padilla (1996), Hendel and Lizzeri (1999), and Denicolò and Garella (1999).

² Most empirical studies examine Alcoa's experience as a monopolist during the interwar period.

single buyer and a seller who is uninformed about the buyer's value. Bagnoli, Salant, and Swierzbinski (1989) show that when a seller has complete information about values for a finite number of prospective buyers, durability can provide an opportunity for the seller to price discriminate over time and extract *larger* profits than a "static monopolist" (e.g., a monopoly lessor) could obtain.

The experimental design focuses on settings in which existing theory indicates that durability should limit a seller's monopoly power. The analysis for the case of a single buyer and a finite time horizon appears in articles on bargaining under asymmetric information by Sobel and Takahashi (1983) and Fudenberg and Tirole (1983). The buyer is privately informed about his value for a single unit of the good; the seller knows the distribution from which the value is drawn. Rejected price offers cause a seller's belief about the buyer's value to be revised downward. Thus, buyer rejections lead to declining price offers over time in equilibrium. This asymmetric information theory is the basis for predictions for the bargaining experiments.

Many of the practical (e.g., antitrust) concerns regarding durable-goods monopoly involve market settings with multiple (perhaps many) buyers. Thus, it seems desirable to run laboratory experiments with multiple buyers. Reynolds (1997) shows how the asymmetric-information model of durable-goods monopoly may be extended to allow for multiple buyers.³ Under particular conditions, such as independent buyer values, equilibrium pricing and buyer behavior are essentially the same as in the single-buyer case. This extension is the basis for predictions for market experiments with multiple buyers.

□ **Experimental evidence.** Güth, Ockenfels, and Ritzberger (1995) report on experiments in which a single seller was matched with ten buyers. Each experiment consisted of five or six experimental games. In each game there was a different combination of maximum number of trading periods (two or three) and discount factors. In one experiment, subjects received prior training in analyzing durable-goods monopoly pricing games. The results for untrained subjects were grossly inconsistent with the theory. Prices failed to conform to comparative statics predictions, and prices tended to be much higher than predicted. The levels of prices with trained subjects were closer to theoretical predictions, but prices still failed to satisfy comparative statics predictions.

Rapoport, Erev, and Zwick (1995) report on bargaining experiments with time discounting, one-sided incomplete information, and an infinite (unlimited) time horizon. They report that (1) price offers tended to decline over time, as predicted by the sequential equilibrium (SE), (2) average initial prices were *higher* the higher the discount factor, contrary to the SE, and (3) for some discount factors the average initial price was above the static monopoly price, contrary to the SE. Rapoport, Erev, and Zwick suggest that the SE theory may be failing because some buyers are using suboptimal rules of thumb for making purchases and because of a social norm about "fair" divisions of surplus.

Smith (1981) reports results for monopoly experiments with an offer auction. The seller makes a sequence of price offers, one unit at a time, within each trading period; buyers either accept or remain silent in response to the offers. The offer auction yielded lower contract prices and higher efficiency than posted offer monopoly. It may be that

³ A continuum of buyers is assumed in most theoretical formulations. This assumption presents an obvious difficulty for experimental implementation. Dudey's (1995) model has a finite number of buyers, where each has a downward-sloping demand for a perfectly divisible durable good. Reynold's (1997) formulation, which involves a simple step demand for each buyer, is more straightforward to apply to an experimental setting with multiple buyers.

the process of making price offers in real time (possibly in rapid succession) leads to lower prices than would occur in a setting with a fixed period length and one offer per period.

Rejection of a price offer that would have provided a positive buyer payoff has been termed *demand withholding* in the experimental literature.⁴ Durable-goods monopoly analysis provides a theoretical rationale for demand withholding. Buyers reject some price offers that would provide positive payoffs because they expect prices to fall in future periods.

3. Experimental design and procedures

■ **Instructions.** The experiments were run on a computerized posted offer pricing program at the Economic Science Laboratory of the University of Arizona. Instructions were displayed on subjects' computer terminals. In addition, subjects were given two pages of written supplemental instructions that covered procedures and information conditions. At the beginning of each experiment, the experimenter read these supplemental instructions aloud in a room in which all of the subjects were present. (The complete instructions are available at <http://www.rje.org/main/sup-mat.html>.)

□ **Periods and sessions.** Each experiment is divided into "sessions" of trading activity. A session consists of one or more trading periods. The maximum number of trading periods was either one, two, or six. Sessions with a single trading period may be thought of as nondurable monopoly sessions. Sessions with multiple trading periods can be viewed as durable-goods monopoly settings. One reason for choosing experiments with a finite horizon is that there is a unique perfect Bayesian equilibrium (PBE) when the horizon is finite.⁵ Ausubel and Deneckere (1989) show that infinite horizon settings such as the experiments of Rapoport, Erev, and Zwick (1995) have multiple equilibria.

The maximum number of trading periods is public information for all subjects. There was one practice session and either nine or ten actual (i.e., for payments) sessions in each experiment. Subjects did not know the exact number of actual sessions in advance; they were told that the experiment would have between six and twelve actual sessions.

□ **Buyers.** Buyers value the items traded in the market because they can resell them to the experimenter for specified amounts. Product "durability" is introduced into the experiments through the assignment of redemption values and discount factors, coupled with restrictions on buyers' ability to purchase the good. The key feature of durability as it is put into operation here is that a buyer has the option of either making a purchase right away or waiting to purchase until later in the session when a more favorable price offer may be available.⁶

⁴ Brown-Kruse (1991) reports examples of demand withholding by buyers in contestable market experiments.

⁵ The PBE is unique for asymmetric-information games with one buyer. I am not aware of a proof of uniqueness for the asymmetric-information game with multiple buyers.

⁶ There are a variety of other features of durability that one might wish to examine. These features include the rate at which the product wears out, a seller's choice of product quality (which could influence the rate of depreciation), the care exercised by the buyer (which could influence the discounted present value of service flows), and the presence of a second-hand market for the good. The form of durability used in the experiments is identical to the way durability is introduced into most game-theoretic analyses.

Two types of experiments were run: bargaining experiments with a single buyer and market experiments with five buyers. Each buyer is assigned a redemption value (or reservation price), v , for a session.⁷ If the buyer makes a purchase in period t of the session at price p_t , then the buyer earns $\delta^{t-1}(v - p_t)$ for the session. A buyer can make at most one purchase during a session. If the buyer does not purchase during the session, then the buyer earns zero for the session. The term δ is a discount factor that is the same for all subjects during the experiments.

□ **Sellers.** A seller produces at zero cost in the experiments. If a seller sets a price p_t in period t of a session, then the seller earns $\delta^{t-1}p_t$ for each unit sold in t . A seller earns zero in a session if buyers do not elect to purchase during the session.

□ **Matching buyers and sellers.** The experiments involve repetitions of sessions in order to give subjects experience with the trading environment and also give them a chance to learn about what other subjects are doing. This kind of repetition would seem to be an essential part of giving an equilibrium theory its best chance of success. Results from experiments reported by Güth, Ockenfels, and Ritzberger (1995), which involved very few repetitions of the same game, were far from equilibrium predictions. However, repetition has the potential to generate reputation effects. For instance, a seller's pricing decisions in an early session could influence buyers' expectations about its prices in later sessions of the same experiment. This sort of seller reputation effect could cause prices to deviate from theoretical predictions of behavior within a single session.

To reduce the potential for reputation effects, a pool of subjects was recruited for each experiment. This permitted several groups to participate in a session simultaneously. Before each new session, subjects were matched together to form new groups. Subjects did not know the identity of the other subjects they were matched with during a session (the precise information provided to subjects about this matching is in the supplemental instructions). For market experiments, 18 subjects were recruited; three subjects were selected (at random) to be sellers and the remaining subjects were buyers for the entire experiment. This allowed me to run three groups through each session simultaneously, with one seller and five buyers per group. For bargaining experiments, ten subjects were recruited; one-half were selected (again, at random) to be sellers and the other one-half were buyers. This allowed me to run five groups simultaneously, with one seller and one buyer per group.

□ **Information.** A buyer's value for a session is a random draw from a uniform distribution over zero to \$12. The buyer observes the draw. The seller (and other buyers, if it is a market experiment) do not observe the draw; they know only the distribution from which the value is drawn. Buyers' values are independent draws in market experiments.

There are several other informational conditions that are important. The discount factor and the seller's marginal cost of zero are public information. The maximum number of trading periods per session is public information. As a session unfolds, each participant in the session observes the prices set by the seller and, at the end of each trading period within the session, the number of buyers who purchased during that period. At the end of each *session*, the prices and purchase periods of buyers for all groups were displayed on a large screen at the front of the laboratory for all subjects to see. This provided a large amount of common information for subjects. This type

⁷ This redemption value can be thought of as the gross, discounted present value of service flows that a consumer would derive from the durable good over the lifetime of the good.

of public information about the history of all groups was provided in the sequential equilibrium experiments of Camerer and Weigelt (1988) and Jung, Kagel, and Levin (1994). Camerer and Weigelt explain that providing this type of information may enable subjects to learn a sophisticated equilibrium concept.

□ **Design and parameters.** The experimental design and parameters are summarized in Table 1. The “periods” treatment permits a comparison of results from a nondurable setting (single period) with results from a durable-goods setting (multiple periods). Such a comparison was not possible in prior experimental studies. Discounting is imposed exogenously in multiperiod sessions. There is a proportional 20% reduction in payoffs after each period within a session.

□ **Other procedures.** Subjects were University of Arizona students from upper-level undergraduate business and economics classes. Subjects were paid a \$5.00 participation fee. Two experiments were run with experienced subjects. Experienced subjects played the same kind of role that they had played in the earlier experiment (i.e., buyer subjects continued as buyers, seller subjects continued as sellers).⁸ Subjects were paid in cash at the end of the experiment. All subjects except sellers in market experiments were paid an amount equal to their dollar earnings in the experiment; sellers in market experiments received one dollar in cash for every five dollars in experimental earnings. Subject payments ranged from \$18 to \$42 for a two-hour experiment.

4. Predictions

■ Predictions are straightforward for single-period sessions. The buyer is predicted to purchase if his value exceeds the price offer (i.e., no withholding). A risk-neutral seller maximizes expected payoff by selecting a price equal to one-half the upper bound of the support of the buyer’s value distribution; that is, a price of \$6. This is the static monopoly price.⁹

There are two principal types of predictions for multiperiod sessions. In a precommitment equilibrium (PE) (see Sobel and Takahashi, 1983) the seller commits in advance to prices for all periods in a session. A seller who can commit to prices can avoid the loss of monopoly power that Coase argued would be associated with product durability.¹⁰ In the experimental environment the PE involves setting the static monopoly price of \$6 in each period (as long as one or more buyers remains in the market). Any buyer who faces this seller strategy and who believes that the seller will maintain the \$6 price would purchase in period one if his value is above \$6 and would never purchase if his value is below \$6. There would be no withholding of demand by buyers in a PE. In the absence of some enforcement mechanism, a precommitment strategy would not be credible. A seller would have an incentive to alter her price following a rejected offer because the seller’s beliefs about the buyer’s value would change after a rejection.

⁸ There was one exception in the bargaining experiment with experienced subjects. Four of the sellers had been sellers in a prior bargaining experiment. The fifth seller had been a buyer in a prior bargaining experiment.

⁹ A risk-averse seller will set a lower price. If the seller has constant relative risk aversion with risk aversion index θ , then the predicted price is the fraction $(1 - \theta)/(2 - \theta)$ of the upper bound of the value support.

¹⁰ In naturally occurring markets, precommitment might be achieved through some contractual arrangement. For example, Butz (1990) shows how best-price provisions in contracts with buyers can induce the same result as precommitment. Contractual arrangements like this are not permitted in these experiments.

TABLE 1 Experimental Design and Parameters

Experiment Number	Maximum Number of Periods per Session	Total Number of Sessions	Number of Buyers	Subject Experience
1 & 2	1	100	1	Inexperienced
3 & 4	2	100	1	Inexperienced
5 & 6	6	110	1	Inexperienced
7	6	39	1	Experienced ^a
8 & 9	6	60	5	Inexperienced
10	6	30	5	Experienced ^b

Notes: Subjects earned one dollar for each dollar of experimental payoffs, except for sellers in experiments with five buyers; these sellers earned \$.20 for each dollar of experimental payoffs. The discount factor (δ) was .8 and seller marginal cost was zero in all experiments.

^a Subjects were recruited from experiments 5 and 6. Includes one inexperienced seller subject. Price data generated by this subject are included with price data from other inexperienced subjects.

^b Subjects were recruited from experiments 8 and 9.

The second type of prediction is a perfect Bayesian equilibrium. A PBE requires that players' strategies be mutual best responses and, in addition, that players update their beliefs about private information held by other agents in a way that is consistent with Bayes' rule and with the strategies of other players. The PE price of \$6 coupled with no buyer withholding is the unique PBE for single-period sessions. Perfect Bayesian equilibria for two-period and six-period sessions are described in Table 2 (the PBE is unique for single-buyer settings but may not be unique for five-buyer settings). The table lists the equilibrium price in each period and the seller's belief about the range of values for a buyer who is still in the market in that period. The table also indicates what a buyer would do in equilibrium. For example, in the two-period case a buyer facing an initial price of \$5.40 will purchase if his value is \$9 or more. Buyers with values between \$5.40 and \$9 are said to *withhold* demand because they do not purchase even though a purchase would yield positive surplus. A buyer does not withhold demand in the final period, since there are no future purchase possibilities. Note that the initial PBE price for two-period sessions is less than the static monopoly price of \$6 and that the initial PBE price for six-period sessions is even lower. Also, the PBE predicts that 55% of buyer types who have values above the price would reject the price offer for periods before the final period.

Reynolds (1997) extends the Sobel and Takahashi analysis to permit multiple buyers with independent private values. A PBE is shown to exist for which the path of prices and seller beliefs about the maximum of buyers' values is the same as in the single-buyer case. This yields two implications for the experiments. First, the sequence of price offers is predicted to be the same in bargaining and market experiments. Second, the predicted price in any period of a market experiment is independent of the number of active buyers.

5. Experimental results

■ **Pricing.** Average prices in each trading period are listed in Table 3 for each type of session. Average initial price offers (period one) are below the static monopoly price

TABLE 2 Prices, Seller Beliefs, and Buyer Behavior in Perfect Bayesian Equilibrium

Trading Period	$T = 2$				$T = 6$			
	Price	Belief	Withhold	Purchase	Price	Belief	Withhold	Purchase
1	5.40	$\nu \in [0, 12]$	$\nu \in [5.40, 9)$	$\nu \in [9, 12]$	4.19	$\nu \in [0, 12]$	$\nu \in [4.19, 8.53)$	$\nu \in [8.53, 12]$
2	4.50	$\nu \in [0, 9)$	None	$\nu \in [4.50, 9)$	3.11	$\nu \in [0, 8.53)$	$\nu \in [3.11, 6.12)$	$\nu \in [6.12, 8.53)$
3					2.36	$\nu \in [0, 6.12)$	$\nu \in [2.36, 4.44)$	$\nu \in [4.44, 6.12)$
4					1.83	$\nu \in [0, 4.44)$	$\nu \in [1.83, 3.27)$	$\nu \in [3.27, 4.44)$
5					1.47	$\nu \in [0, 3.27)$	$\nu \in [1.47, 2.46)$	$\nu \in [2.46, 3.27)$
6					1.23	$\nu \in [0, 2.46)$	None	$\nu \in [1.23, 2.46)$

of \$6 for all types of sessions. Average prices for two-period bargaining sessions are very close to PBE predicted prices. Moreover, 74 out of 100 sessions with $T = 2$ extend for two periods, which nearly matches the PBE prediction of $\frac{7}{12}$ (75%). However, average initial prices for six-period sessions are well above the PBE prediction of \$4.19, and average price for single-period sessions is *below* the static monopoly price (which is also the PBE prediction). In fact, holding the number of buyers and subject experience constant, average initial prices *rise* as the maximum number of trading periods goes from one to two to six. This is the opposite of PBE predictions, which have initial prices falling as the time horizon grows.

There is a tendency for average prices to fall over time in multiperiod sessions. However, average prices do not decline for all periods, and the overall amount of price reduction is small in some settings (e.g., market experiments with inexperienced subjects). Sessions with experienced subjects tend to yield lower prices than sessions with inexperienced subjects. But even with experienced subjects, prices are well above PBE predictions.

Price changes are reported in Table 4. The averages of actual price changes need not coincide with changes calculated from Table 3. The number of price observations

TABLE 3 Price Data

Trading Period	Bargaining			Bargaining	Market	Market
	$T = 1$	$T = 2$	$T = 6$	$T = 6$ Experienced	$T = 6$	$T = 6$ Experienced
1	\$4.67 (100)	\$5.44 (100)	\$5.78 (110)	\$5.25 (39)	\$5.37 (60)	\$5.16 (30)
2		\$4.75 (74)	\$5.00 (86)	\$4.54 (23)	\$4.94 (59)	\$5.29 (30)
3			\$4.29 (68)	\$3.79 (18)	\$4.66 (57)	\$4.62 (30)
4			\$4.33 (55)	\$3.18 (13)	\$4.72 (57)	\$4.40 (29)
5			\$3.83 (47)	\$2.45 (11)	\$4.88 (56)	\$4.33 (28)
6			\$4.64 (38)	\$3.31 (9)	\$4.79 (54)	\$4.21 (27)

Note: The dollar figure is the average price set by sellers in the trading period. The number in parentheses is the corresponding number of price observations.

TABLE 4 Price Changes

Trading Periods	Bargaining ($T = 2$)		Bargaining ($T = 6$)		Bargaining ($T = 6$) Experienced		Market ($T = 6$)		Market ($T = 6$) Experienced	
	Percent Change	Cut Fre- quency	Percent Change	Cut Fre- quency	Percent Change	Cut Fre- quency	Percent Change	Cut Fre- quency	Percent Change	Cut Fre- quency
1-2	-15%	73%	-19%	84%	-22%	100%	-7%	53%	2%	30%
2-3			-16%	80%	-18%	87%	-5%	49%	-14%	67%
3-4			-3%	63%	-13%	55%	3%	40%	-11%	62%
4-5			-13%	61%	-13%	56%	13%	38%	-8%	57%
5-6			31%	54%	45%	25%	2%	56%	-15%	59%

Note: The "percent change" figure is the average percentage change in prices for all pairs of successive prices. The "cut frequency" figure is the number of times the price is reduced from one period to the next divided by the total number of successive price pairs.

in successive periods may differ because a session may end before reaching the next trading period. Two types of statistics are reported in Table 4. The first is the percent change in price from one trading period to the next, averaged over all pairs of successive prices. The second is the percentage frequency that a pair of successive prices involves a price cut ("Cut Frequency"). Average price cuts between periods one and two of six-period bargaining experiments are substantially larger than the drop in average prices observed in Table 3. This is because relatively low initial offers were likely to be accepted, ending the session. Table 4 reveals a sharp increase in price in the final period of six-period bargaining experiments. For inexperienced subjects, this increase occurs even though more than one-half of sellers decrease their price. A minority of sellers chose a large price increase for period six; some sellers set a price of \$12 in the final period.

Figures 1 through 4 illustrate the price data and the predictions. These figures show PE prices as diamonds and PBE prices as triangles. The average price for a trading period is indicated by a black square. A 95% confidence interval for the data is indicated by a vertical bar extending above and below the square.¹¹ Several features of the results may be summarized based on Tables 3 and 4 and the figures.

Single-period sessions. The 95% confidence interval for the data lies below the static monopoly prediction.

Two-period sessions. Pricing is consistent with PBE and inconsistent with PE. The 95% confidence intervals for the data lie below PE prices for both periods.

Six-period sessions. Average prices and confidence intervals lie between the PE and PBE predictions in all periods. Average prices tend to fall after the first two trading periods. Average prices in market sessions are roughly constant during the last 3-4 periods. Average prices in bargaining sessions fall in period five and then rise in the final period.

Bargaining versus market sessions. Average prices in bargaining and market sessions are similar. The hypothesis that mean prices in bargaining and market sessions

¹¹ Data for experienced and inexperienced subjects are combined for the graphs. The confidence interval is constructed based on the variability of the observed price data. The interval is the mean price for the period plus or minus $1.96(s/(n)^{1/2})$, where s is the sample standard deviation of prices for the period and n is the number of price observations for the period.

FIGURE 1
PRICES IN BARGAINING EXPERIMENTS ($T = 1$)




FIGURE 2
PRICES IN BARGAINING EXPERIMENTS ($T = 2$)




FIGURE 3
PRICES IN BARGAINING EXPERIMENTS ($T = 6$)





FIGURE 4
PRICES IN MARKET EXPERIMENTS ($T = 6$)



are the same cannot be rejected at the 95% level for any period except period five. These results are largely consistent with the prediction in Reynolds (1997) that PBE prices are the same regardless of the number of buyers. The main difference is that prices in bargaining sessions started higher and declined more over time (except for the final period) than prices in market sessions.

Price paths for individual seller subjects are illustrated in Figures 5–10. The price data for each figure are from the seventh session of an experiment. Figures 5 and 6 show five price paths for two-period bargaining sessions (all inexperienced subjects). For Figure 5, two offers were accepted in period one (\$6 and \$4.50). For the three other paths in Figure 5, two have a drop in price in the second period; the third path is constant at \$5. Figures 7 and 8 show price paths for six-period bargaining sessions (Figure 7 has inexperienced subjects, Figure 8 has experienced subjects). In Figure 7, two paths begin at \$8.50. The other three have an initial price near \$6, with one of the three initial offers accepted. All of the paths are nonincreasing over time. In Figure 8 the initial offers range from \$5.25 to \$6.75. Two initial offers of \$6 were accepted. Only one path extends the full six periods. Figures 9 and 10 each show three price paths for market sessions (Figure 9 has inexperienced subjects, Figure 10 has experienced subjects). Initial prices are between \$4 and \$6 in both figures. Prices fall monotonically for only two out of six paths. The low price path in Figure 10 is fairly close to the PBE predicted path. Another path in this figure is almost a mirror image of this path, with prices *rising* throughout the session.

FIGURE 5
PRICE PATHS—BARGAINING ($T = 2$)
INEXPERIENCED SUBJECTS




FIGURE 6
PRICE PATHS—BARGAINING ($T = 2$)
INEXPERIENCED SUBJECTS




FIGURE 7
PRICE PATHS—BARGAINING ($T = 6$)
INEXPERIENCED SUBJECTS




FIGURE 8
PRICE PATHS—BARGAINING ($T = 6$)
EXPERIENCED SUBJECTS




FIGURE 9
PRICE PATHS—MARKET ($T = 6$)
INEXPERIENCED SUBJECTS





FIGURE 10
PRICE PATHS—MARKET ($T = 6$)
EXPERIENCED SUBJECTS



Figures 5–10 illustrate a high degree of heterogeneity in price-setting behavior across subjects. This heterogeneity tended to persist even though subjects were provided with complete price histories for all sellers in their experiment, and even for experienced subjects. The price paths for an individual subject did tend to converge as sessions were repeated in an experiment.

It may be that the failure of the PBE to predict prices is due to attempts by some sellers to establish a reputation for setting and maintaining high prices. Sellers could communicate to buyers only through their price choices, so any attempt to establish a reputation would involve some observable pattern of prices for a seller. Given the provision of price histories for all sellers in each experiment, a seller might attempt to acquire a reputation by repeatedly charging a notable price in the first period.

I investigated this issue by looking for cases in which a seller charged the same opening price in consecutive sessions of an experiment. In two-period bargaining experiments, two sellers had runs of three consecutive sessions with the same opening price. However, in each case this occurred in the last three sessions of the experiment. This suggests convergence to a strategy more than an attempt to establish a reputation. The longest runs in six-period bargaining experiments were four sessions and five

sessions. The four-session run began early, in session three. However, the seller set a low opening price of \$4 in these sessions. A five-session run by a different seller was for the last five periods and involved a price of \$6.50. Two factors suggest that this was not motivated by reputation concerns. First, it was late in the experiment, so that the seller had relatively few sessions to reap rewards from a reputation. Second, the price was not distinctive. There were three other identical opening prices during these five sessions. The longest run in market experiments was for four sessions in which a seller opened with a price of \$6, beginning in session six. Again, this was relatively late in the experiment and the price was not distinctive compared to other sellers' opening prices.

The pricing patterns observed in market experiments suggest another way that a seller might attempt to establish a reputation. One of the three experienced sellers established a pattern of opening around \$5 and increasing the price in each period (this is illustrated in Figure 10). This seems irrational and bizarre at first glance. However, it may have been motivated by an attempt to make it very clear that this seller was not offering any price cuts. The difficulty of distinguishing him- or herself was compounded in this experiment by the presence of a second seller who sometimes increased price in the second period (again, see Figure 10). The seller with increasing prices earned an average payoff of \$2.87 per session, which is closer to the PE payoff than to the PBE payoff.

□ **Seller profit.** Profits earned by subjects provide additional information about the extent to which sellers exercised monopoly power successfully. Predicted expected profits and actual average profit per session are listed in Table 5. Average seller profits were somewhat lower for multiperiod sessions than for single-period sessions. This is true even though initial prices were higher in multiperiod sessions. The multiperiod setting introduced substantial demand withholding and delays in purchasing that resulted in lower profits to sellers compared to seller profits with take-it-or-leave-it offers ($T = 1$). Buyers also earned less on average in multiperiod sessions than in single-period sessions. Since sellers set relatively low prices in single-period sessions, total surplus and buyers' share of total surplus were greater than would be predicted by static monopoly pricing. Product durability, as it is introduced in the experiments, leads to higher prices, lower buyer surplus, and *lower* seller profits because of delays in purchasing. So durability did reduce seller profits, but not in a way that ultimately benefited buyers. Market efficiency (as measured by total surplus) was decreasing in the maximum number of trading periods (T), in contrast to the PBE prediction.

□ **Buyer behavior and demand withholding.** There was a substantial amount of demand withholding observed in the experiments. Results are summarized in Table 6,

TABLE 5 Average Profit per Session

	PE Predicted	PBE Predicted ($T = 2$)	PBE Predicted ($T = 6$)	Bargaining ($T = 1$)	Bargaining ($T = 2$)	Bargaining ($T = 6$)	Market ($T = 6$)
Seller	\$3.00	\$2.70	\$2.10	\$2.89	\$2.56	\$2.37	\$2.66
Buyer	\$1.50	\$1.95	\$2.86	\$2.59	\$2.30	\$2.32	\$2.12

Note: For multiperiod sessions the profit is the expected discounted value of payoffs using the discount factor, $\delta = .8$.

TABLE 6 Demand Withholding

Trading Period	Bargaining $T = 1$	Bargaining $T = 2$	Bargaining $T = 6$	Bargaining $T = 6$ Experienced	Market $T = 6$	Market $T = 6$ Experienced
1	8% (71)	59% (64)	64% (58)	39% (31)	42% (182)	24% (94)
2		11% (45)	65% (46)	59% (17)	54% (87)	67% (18)
3			65% (34)	50% (12)	72% (53)	58% (24)
4			72% (25)	70% (10)	78% (45)	75% (20)
5			75% (24)	37% (8)	48% (33)	72% (18)
6			17% (18)	0% (4)	19% (33)	0% (20)
$t < T$		59% (64)	67% (187)	50% (78)	53% (400)	44% (174)

Note: The percent figure is the percentage of potential purchases withheld. The number in parentheses is the corresponding number of potential purchases.

where withholding is reported by time period for each type of session. The bottom row reports overall withholding for all periods before the final period. The share of units for which demand was withheld was 50% or more during many periods of multiperiod sessions. The extent of demand withholding is roughly comparable to the 55% withholding prediction of the PBE for periods before the final trading period (for sessions with $T = 2$ and $T = 6$).¹² The PBE predicts zero withholding in the final period of a session. This is precisely what happens in sessions with experienced subjects (although the total number of purchase opportunities in period six is small). There is a small amount of withholding in the final period for inexperienced subjects.

The nature of demand withholding is examined in more detail by formulating and estimating a simple, discrete-choice model of buyer purchase decisions. Estimation of this model reveals the responsiveness of purchase decisions to current buyer surplus. Financially motivated buyer subjects should be more likely to purchase the greater their current surplus, *ceteris paribus*. The model may also be used to consider the roles of other factors such as fairness perceptions and expectations of future price.

The discrete-choice model has the following utility specification:

Utility from purchase by buyer i in experimental session j in period t :

$$U_{ijt} = v_{ij} - p_{jt}$$

Utility from withholding for buyer i in experimental session j in period t :

$$W_{ijt} = \theta_t v_{ij} + \epsilon_{ijt}$$

The session index, j , is used to identify a match between a seller and buyer(s) in a

¹² The 55% figure is the predicted withholding rate along the equilibrium path for the PBE. An alternative comparison would be the withholding rate predicted for buyers in the PBE based on the prices observed in the experiments, i.e., the off-the-equilibrium predicted withholding rates.

particular session of an experiment. Since there is one seller per group, this index also identifies the seller that the buyer is matched with. The payoff associated with withholding is the sum of a term that depends on the buyer's value and an idiosyncratic term, ϵ_{ijt} .¹³ This idiosyncratic term could stem from mistakes in decision making, from experimentation on the part of subjects, from heterogeneity in expectations of future prices, and other sources.

A buyer's perception of the fairness of a price offer could be one source of withholding. For example, a buyer may elect to withhold if his surplus from a purchase would be less than some fraction γ of the seller's surplus (which is equal to the price) on the grounds that this would be an unfair division of total surplus. Then the median buyer (for whom $\epsilon_{ijt} = 0$) would prefer to withhold if $v_{ij} - p_{jt} < [\gamma/(1 + \gamma)]v_{ij}$. This kind of buyer behavior would yield an estimate of θ_t equal to $\gamma/(1 + \gamma)$. The perception that an offer is unfair is an explanation for withholding in the final period of a session.

In periods before the final period of a session there are additional motives for withholding. One motive may be that a buyer expects the price to fall (this is the only reason for withholding in the PBE analysis). Consider a reformulation of the utility associated with withholding:

$$\tilde{W}_{ijt} = \delta(v_{ij} - p_{j,t+1}) + \tilde{\epsilon}_{ijt}.$$

This reformulation specifies the utility as the discounted surplus the buyer expects next period plus an idiosyncratic term. Suppose in addition that $p_{j,t+1} = (1 - \alpha_t)p_{jt}$, where α_t is the fraction by which a buyer expects the price to be cut next period. Then the condition for choosing to purchase, $U_{ijt} \geq \tilde{W}_{ijt}$, is equivalent to $v_{ij} - p_{jt} \geq \theta_t v_{ij} + \epsilon_{ijt}$, where $\theta_t = \delta\alpha_t/(1 - \delta + \delta\alpha_t)$ and $\epsilon_{ijt} = \tilde{\epsilon}_{ijt}/(1 - \delta + \delta\alpha_t)$. An estimate of the price cut expected by buyers who intend to purchase next period may be inferred from the estimated coefficients of buyer surplus and buyer value. Another motive for withholding is that the buyer retains an option to purchase in a later period and that this option is valuable. A buyer faces strategic uncertainty about future prices during a session. From a buyer's point of view, future prices may appear to be random variables.¹⁴ If a price turns out to be relatively low, then the buyer can exercise his option to purchase. This value of the option to purchase would be increasing in the buyer's value. This option-value interpretation of withholding is consistent with specifying W_{ijt} as an increasing function of v_{ij} .

For estimation purposes, the ϵ_{ijt} terms are modelled as independent draws from a normal distribution with mean zero and variance σ^2 . This formulation yields the following probit model of purchasing decisions:

$$\begin{aligned} \Pr[\text{purchase by } i \text{ in period } t \text{ of session } j] &= \Pr[W_{ijt} \leq U_{ijt}] \\ &= \Pr[\epsilon_{ijt} \leq (v_{ij} - p_{jt}) - \theta_t v_{ij}] \\ &= \Phi\left[\frac{(v_{ij} - p_{jt}) - \theta_t v_{ij}}{\sigma}\right], \end{aligned}$$

where Φ is the cumulative distribution function for the standard normal distribution.

¹³ This term could be influenced by a variety of factors, including buyer characteristics and the history of play the buyer has observed to date (in the current session and any prior sessions the buyer has participated in).

¹⁴ There is no such strategic uncertainty in the PBE formulation. In equilibrium a buyer can predict future seller prices perfectly.

The arguments of Φ are current buyer surplus, $v_{ij} - p_{jt}$, and the buyer's value, v_{ij} . Note that the coefficient for the value is time dependent ($-\theta_t/\sigma$).

The model was estimated using data after the first two sessions for each type of experiment. I excluded the first two sessions because subjects are most likely to be learning about the environment in early sessions, so error terms governing forecasts in early sessions may be distributed differently from error terms in later sessions. Observations for which the buyer's value was less than the price offer are excluded from the sample. The experiment rules did not permit a subject to purchase if his value was less than the price. The variable labeled *SURPLUS* is equal to a buyer's value minus a price offered by the seller the buyer was matched with. The variable labeled *VALUE#* is equal to a buyer's value times a dummy variable equal to one if the trading period equals # and zero otherwise.

Estimation results are reported in Table 7. The *SURPLUS* coefficient is positive in each of the four estimated models and significant at the 1% level in three models. The *p*-value for the *SURPLUS* coefficient for the $T = 1$ case is .06. This suggests that buyers responded strongly to financial incentives for making purchases. The coefficients for the value terms for the final trading period (*VALUE1* for $T = 1$, *VALUE2* for $T = 2$, *VALUE6* for $T = 6$) are close to zero and statistically insignificant for each estimation. This suggests that fairness perceptions were not an important influence on withholding in these experiments. There was withholding in the final period (for inexperienced subjects), but its incidence did not rise as a buyer's surplus became small relative to the seller's surplus. The estimated coefficients for value terms before the final period are all negative and significant. The estimates imply positive estimates for the corresponding θ_t terms. The estimates for θ_t terms range from about .3 to .5. If withholding is based on an expectation of a purchase in the next period, then the anticipated price cuts that are implied by the θ_t estimates range from 35% to 50%. Average actual price

TABLE 7 Probit Estimates of Buyers' Purchase Decisions

Independent Variables	Bargaining ($T = 1$)	Bargaining ($T = 2$)	Bargaining ($T = 6$)	Market ($T = 6$)
<i>SURPLUS</i>	3.622 (1.957)	.518* (.136)	.571* (.083)	.599* (.058)
<i>VALUE1</i>	-.252 (.185)	-.254* (.058)	-.229* (.038)	-.175* (.025)
<i>VALUE2</i>		.006 (0.057)	-.231* (.040)	-.216* (.028)
<i>VALUE3</i>			-.263* (.044)	-.271* (.035)
<i>VALUE4</i>			-.296* (.049)	-.330* (.044)
<i>VALUE5</i>			-.289* (.053)	-.220* (.045)
<i>VALUE6</i>			-.087 (.065)	.132 (.070)
Log-likelihood	-4.81	-36.8	-124.08	-252.44
Number of observations	57	95	236	508

* Significant at the 1% level. Standard errors of coefficient estimates are in parentheses.

cuts reported in Table 4 were much smaller than this. This suggests that withholding before the final period was not based solely on the expectation of a price cut that would result in a purchase during the next period. The option-value explanation may account for the rest of observed withholding.

The assumption that the ϵ_{ijt} residual terms are independent is quite strong. I attempted to estimate a random-effects model to check this assumption. In the random-effects model the residual term is specified as, $\epsilon_{ijt} = \mu_i + \eta_{ijt}$, where each of the two components is a draw from an independent normal distribution. This specification permits correlation across residuals for the same buyer. The random-effects model could not be estimated for $T = 1$ because the number of observations with zero values for the dependent variable was too small. The random-effects model was estimated for the other three cases of Table 7 (the estimation results are not reported here). All of the estimated coefficients retained the same signs as the original estimates. The magnitudes of estimated coefficients tended to be slightly larger than the original estimates. The pattern of statistical significance of estimated coefficients was the same as in Table 7, except for six-period bargaining, for which none of the estimated coefficients was significant. The estimated correlation across a buyer's residuals was positive for the three models but was not statistically significant at the 5% level for any of the three models. I performed likelihood ratio tests to compare the overall performance of the random-effects model with the basic probit model. These tests permit me to reject the basic probit model in favor of the random-effects model (at the 5% level) for six-period bargaining and market data, but not for two-period bargaining data. The main effect of correlation of residuals is to increase the standard errors of coefficient estimates for six-period bargaining data, so that the estimates are no longer significant.

6. Discussion

■ While the PE and PBE predictions each help to explain some aspects of observed results, each one also fails in important respects. It may be that these theories fail because they neglect aspects of subjects' preferences, perceptions, and/or decision-making capabilities. This leads me to consider how risk preferences, fairness concerns, subjects' use of prior information, and bounded rationality might account for observed results.

□ **Risk preferences.** The equilibrium predictions assume risk neutrality. Subjects' risk preferences were not controlled by the experimenter, and subjects may have been either risk averse or risk loving. Sellers with risk-averse preferences could explain the observed pricing results for nondurable sessions ($T = 1$). However, risk aversion lowers PBE predictions of prices for multiperiod sessions, which would move the predictions for six-period sessions away from observed results. Thus, risk aversion by itself does not look like a promising candidate for reconciling game-theoretic predictions with experimental results.

□ **Fairness.** Preferences that involve concerns about fairness are a potential explanation for outcomes in bargaining environments. Roth et al. (1991) report that receivers often reject offers that yield positive payoffs in ultimatum game experiments and that equal splitting is a common outcome in these experiments. The potential impact of subjects' fairness concerns is more difficult to sort out in an asymmetric-information environment. A rejected offer does not necessarily mean that the buyer (or receiver) has left money on the table, as it does in an ultimatum game. For this reason, a seller would be less likely to interpret a rejected offer as an indication that the offer was

perceived to be unfair, compared to a proposer in an ultimatum game whose offer was rejected. So, even if buyers sometimes withhold because they view offers as unfair, it would be difficult for a seller to understand that this is happening and make corresponding adjustments to price decisions. It is true that (inexperienced) buyer subjects did sometimes leave money on the table by withholding in the final trading period. However, probit estimates of buyer behavior suggest that withholding was not influenced by asymmetry in the division of total surplus between buyer and seller.

There is an anomaly regarding seller pricing that may be related to sellers' perceptions of fairness. In the final period of six-period bargaining sessions, average prices were 31% above average prices in the previous period for inexperienced subjects, and 45% higher for experienced subjects. This is inconsistent with equilibrium theories of behavior. A price increase may be a (spiteful) response to a perception of unreasonable or unfair behavior on the part of the buyer the seller is matched with.¹⁵ The fact that a buyer is still active in period six of a bargaining session implies that all five of the seller's prior offers were rejected.

□ **Prior beliefs.** A subject's beliefs about another subject's private information may not be completely controlled by the experimenter, in spite of the experimenter's efforts to provide public information about the prior distribution. Camerer and Weigelt (1988) propose the idea of "homemade priors" for subjects as an adjustment to sequential equilibrium. They find that sequential equilibrium provides good predictions of their experimental results if subjects utilize a prior distribution that differs from the one the experimenters conveyed to subjects.

An adjustment that improves PBE predictions for six-period sessions is to shift the prior so that probability is moved away from very low values and distributed across higher values. Prices in six-period sessions may have been too high because sellers had overly optimistic prior beliefs about buyers' values. However, this type of homemade prior moves equilibrium predicted prices away from observed prices for one-period sessions. Or, to put it differently, a different type of homemade prior, one that shifts the distribution toward *lower* values, would be required to reconcile predictions with results for one-period sessions.¹⁶ There is no apparent reason why homemade priors should change as parameters governing the environment change. Jung, Kagel, and Levin (1994) find similar difficulties with homemade priors in their experimental study of a sequential equilibrium model of predatory pricing.

□ **Bounded rationality.** The experimental environment is complex, especially for six-period sessions. The combined difficulties of asymmetric information and a solution that involves backward recursion may prevent subjects from arriving at the PBE. Selten (1978) argued that one may observe behavior close to a perfect equilibrium in one- or two-period settings but fail to observe behavior close to a perfect equilibrium in settings with more periods due to complexity and subjects' *limited rationality*. These factors may certainly be playing a role in observed results for six-period sessions. However,

¹⁵ These price increases may be an example of punishing (or conveying displeasure toward) someone you perceive to be hurting you, as discussed in Rabin (1993). In contrast, prices did not rise (on average) in the final period of market sessions. This contrast in seller behavior is consistent with experimental results in Roth et al. (1991), who report that fairness concerns play a larger role in bargaining experiments than in market experiments.

¹⁶ Sobel and Takahashi (1983) show how to compute PBE predictions for bargaining games with a cumulative distribution of the form $F(v) = (v/12)^a$. When $a = 1$, the distribution is uniform, as in the experiments. A value of $a > 1$ shifts the distribution toward higher values; a value of $a < 1$ shifts the distribution toward lower values.

the results from short-horizon sessions do not conform to PBE predictions either. The important failure is that prices for the single-period sessions are well below the PBE prediction (which coincides with the PE prediction).¹⁷

McKelvey and Palfrey (1995, 1998) propose a games-with-errors approach for modelling bounded rationality that they term quantal response equilibrium (QRE). A quantal response is a smoothed-out best response to the strategies of other players; a player chooses actions that yield higher expected payoffs with higher probability but does not choose a best response with probability one. The equilibrium concept in extensive form games is termed an agent quantal response equilibrium (AQRE).

There are features of observed behavior and of equilibrium payoffs that suggest that the AQRE concept might be a useful way to view experimental results. Buyers are observed to withhold in the last period even though their payoff would have been higher had a purchase been made. Also, a smooth response such as a probit or logit distribution describes buyers' decision rules better than a best-response rule that implies purchase with probability one for all buyer surplus levels above a fixed cutoff. Figure 11 illustrates the estimated decision rule for buyers in the first period of $T = 2$ sessions. The smooth curve is the probit estimated probability of purchase as a function of buyer value, given that the seller sets a price of \$5.40. The PBE decision rule given a price of \$5.40 is a step function that goes from probability zero below \$9 to probability one above \$9. Finally, expected payoff for the seller as a function of price is highly asymmetric in the PBE. Consider, for example, the two-period case. The optimal price of \$5.40 yields expected seller payoff of \$2.70. The expected seller payoff drops smoothly to zero as the price declines to zero. As the price is raised above \$5.40 the expected payoff initially falls, but once the price reaches \$7.20, the payoff remains constant at \$2.40.¹⁸ In an AQRE formulation this type of payoff asymmetry would tend to induce the seller to set prices that are, on average, above the PBE price.

McKelvey and Palfrey (1998) provide some illustrative calculations of AQRE for relatively simple extensive-form games. For more complex games, such as the multi-period, asymmetric-information bargaining games studied in the present article, these calculations become enormously complicated. I have made some preliminary calculations of the AQRE for one-period and two-period games, using a discretized set of buyer values and seller prices. Two results are noteworthy. First, the initial expected seller price can be higher for the $T = 2$ game than the $T = 1$ game. In other words, the comparative statics prediction on the effect of the number of trading periods on initial prices for the PBE can be reversed in an AQRE. Second, the expected price (conditional on a rejected offer) in period two does not fall as much in the AQRE as it does in the PBE. This could account for the relatively slow descent of prices in two-period and six-period sessions.


7. Concluding remarks

■ I have examined monopoly pricing environments with varying numbers of buyers and trading periods. The “periods” treatment allowed me to compare results from static monopoly experiments to results from multiperiod experiments that capture essential features of durable-goods sales. The experimental results were consistent with some

¹⁷ A plausible adjustment to the theory that accounts for the one-period results is seller risk aversion. However, the level of risk aversion required to move predictions close to observed results for $T = 1$ would move predicted prices well below observed prices for $T = 2$.

¹⁸ This payoff asymmetry for the seller is accounted for by the fact that no buyer type would purchase in the continuation equilibrium if the initial price were above \$7.20. In this event the seller retains the valuable option to sell in the second period to a buyer whose value is uniformly distributed over [0, \$12].

FIGURE 11
ESTIMATED BUYER DECISION RULE (PERIOD ONE OF $T = 2$ SESSIONS)



qualitative implications of perfect Bayesian equilibrium. Buyers in multiperiod sessions were much more likely to withhold purchases than buyers in single-period sessions. Sellers in multiperiod sessions tended to offer prices that decreased over time. However, the results were inconsistent with the main comparative statics prediction of the effect of product durability. Initial seller prices were *higher* as the number of trading periods increased, controlling for the number of buyers and subject experience. Product durability, as implemented in the experiments, did not restrain seller pricing. Part of the reason for this is that prices in single-period sessions were well below the static monopoly price prediction. The combination of human buyer subjects and seller uncertainty about values appears to limit monopoly power, even for “one-shot” monopolists.

The failure of the prediction on the effect of changes in the number of trading periods is puzzling. A promising explanation for this failure is a version of bounded rationality, perhaps coupled with subject risk aversion. A games-with-errors approach can yield higher predicted average initial prices in two-period sessions than in one-period sessions. Such an approach can also provide a better fit than PBE to observed buyer choice behavior. An important direction for future research is to identify an experimental design for which calculations for a games-with-errors model are less complex and to run the experiments. A formulation that allows for decision errors may be a useful way to model buyer behavior and seller pricing in “naturally occurring” markets.

Some qualifications of these experimental results should be noted. Results may be sensitive to the subject-matching scheme that was used and/or to the informational conditions of the experiments. It is possible that the public information condition for session histories facilitated development of seller reputations. I provided evidence suggesting that reputations did not emerge in bargaining experiments, but seller reputations may have played a role in market experiments. The experience level of subjects may also be important for some aspects of results, such as the effect of the “periods” treatment.

This article does not address directly Coase’s conjecture on convergence to marginal cost pricing. It appears doubtful that any experimental design would generate results consistent with the Coase conjecture. My results for six-period games and those of Rapoport, Erev, and Zwick (1995) for unlimited-horizon games involve observed prices above equilibrium predictions. So even if an experimental design involved a very large number of periods and a high discount factor, prices may very well remain well above marginal cost. An experimental design that would give Coase’s conjecture its best shot would seem to require enough flexibility to allow many offers to be made in a short period. A “real-time” trading environment similar to Smith’s (1981) offer auction experiment might provide the right sort of flexibility.

References

- AUSUBEL, L.M. AND DENECKERE, R.J. "Reputation in Bargaining and Durable Goods Monopoly." *Econometrica*, Vol. 57 (1989), pp. 511–531.
- BAGNOLI, M., SALANT, S.W., AND SWIERZBINSKI, J.E. "Durable-Goods Monopoly with Discrete Demand." *Journal of Political Economy*, Vol. 97 (1989), pp. 1459–1478.
- BROWN-KRUSE, J.L. "Contestability in the Presence of an Alternative Market: An Experimental Examination." *RAND Journal of Economics*, Vol. 22 (1991), pp. 136–147.
- BULOW, J.I. "Durable-Goods Monopolists." *Journal of Political Economy*, Vol. 90 (1982), pp. 314–332.
- BUTZ, D.A. "Durable-Good Monopoly and Best-Price Provisions." *American Economic Review*, Vol. 80 (1990), pp. 1062–1076.
- CAMERER, C. AND WEIGELT, K. "Experimental Tests of a Sequential Equilibrium Reputation Model." *Econometrica*, Vol. 56 (1988), pp. 1–36.
- COASE, R.H. "Durability and Monopoly." *Journal of Law and Economics*, Vol. 15 (1972), pp. 143–149.
- DENICOLÒ, V. AND GARELLA, P.G. "Rationing in a Durable Goods Monopoly." *RAND Journal of Economics*, Vol. 30 (1999), pp. 44–55.
- DUDEY, M. "On the Foundations of Dynamic Monopoly Theory: Comment." *Journal of Political Economy*, Vol. 103 (1995), pp. 893–902.
- FISHER, F.M., MCGOWAN, J.J., AND GREENWOOD, J.E. *Folded, Spindled and Mutilated: Economic Analysis and U.S. v. IBM*. Cambridge, Mass.: MIT Press, 1983.
- FUDENBERG, D. AND TIROLE, J. "Sequential Bargaining with Incomplete Information." *Review of Economic Studies*, Vol. 50 (1983), pp. 221–247.
- GUL, F., SONNENSCHN, H., AND WILSON, R. "Foundations of Dynamic Monopoly and the Coase Conjecture." *Journal of Economic Theory*, Vol. 39 (1986), pp. 155–190.
- GÜTH, W., OCKENFELS, P., AND RITZBERGER, K. "On Durable Goods Monopolies: An Experimental Study of Intrapersonal Price Competition and Price Discrimination Over Time." *Journal of Economic Psychology*, Vol. 16 (1995), pp. 247–274.
- HENDEL, I. AND LIZZERI, A. "Interfering with Secondary Markets." *RAND Journal of Economics*, Vol. 30 (1999), pp. 1–21.
- JUNG, Y.-J., KAGEL, J.H., AND LEVIN, D. "On the Existence of Predatory Pricing: An Experimental Study of Reputation and Entry Deterrence in the Chain-Store Game." *RAND Journal of Economics*, Vol. 25 (1994), pp. 72–93.
- KARP, L.S. AND PERLOFF, J.M. "The Optimal Suppression of a Low-Cost Technology by a Durable-Good Monopoly." *RAND Journal of Economics*, Vol. 27 (1996), pp. 346–364.
- KÜHN, K.-U. AND PADILLA, A.J. "Product Line Decisions and the Coase Conjecture." *RAND Journal of Economics*, Vol. 27 (1996), pp. 391–414.
- MCKELVEY, R.D. AND PALFREY, T.R. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, Vol. 10 (1995), pp. 6–38.
- AND ———. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*, Vol. 1 (1998), pp. 9–42.
- RABIN, M. "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, Vol. 83 (1993), pp. 1281–1301.
- RAPOPORT, A., EREV, I., AND ZWICK, R. "An Experimental Study of Buyer-Seller Negotiation with One-Sided Incomplete Information and Time Discounting." *Management Science*, Vol. 41 (1995), pp. 377–394.
- REYNOLDS, S.S. "Durable Goods Monopoly: Multiple Buyers with Independent Values." Mimeo, Department of Economics, University of Arizona, 1997.
- ROTH, A.E., PRASNIKAR, V., OKUNO-FUJIWARA, M., AND ZAMIR, S. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *American Economic Review*, Vol. 81 (1991), pp. 1068–1095.
- SELTEN, R. "The Chain-Store Paradox." *Theory and Decision*, Vol. 9 (1978), pp. 127–159.
- SMITH, V.L. "An Empirical Study of Decentralized Institutions of Monopoly Restraint." In V.L. Smith, ed., *Papers in Experimental Economics*. New York: Cambridge University Press, 1981.
- SOBEL, J. AND TAKAHASHI, I. "A Multistage Model of Bargaining." *Review of Economic Studies*, Vol. 50 (1983), pp. 411–426.
- STOKEY, N.L. "Rational Expectations and Durable Goods Pricing." *Bell Journal of Economics*, Vol. 12 (1981), pp. 112–128.