The Strategies Behind their Actions: A Method to Infer Repeated-Game Strategies and an Application to Buyer Behavior

Jim Engle-Warnick
McGill University

Bradley J. Ruffle
Ben-Gurion University

December 2005

Abstract: We introduce a Bayesian method to infer repeated-game strategies in the form of if-then statements that best describe individuals' observed actions. We apply this method to buyer behavior in posted-offer market experiments. While the strategies of one-quarter of the buyers in our experiments correspond to the game-theoretic prediction of passive price-taking, for three-quarters of the buyers we infer repeated-game strategies that condition on time, price, and combinations of time and price. Our analysis fills a gap in a literature that studies the convergence of pricing behavior in posted-offer markets but has not addressed the market as a repeated game. We propose that strategy inference should at least complement existing methods of statistical inference on observed strategic behavior.

keywords: strategy inference, repeated-game strategies, Bayesian methods, strategic buyer behavior, binary decision tree, experimental posted-offer market.

JEL classification nos. C91, D42.

Contact Information: Engle-Warnick: Department of Economics, 855 Sherbrooke St. W., McGill University, Montreal, QC, H3A 2T7, Canada, e-mail: `jim.engle-warnick@mcgill.ca`, tel.: (514) 398-1559; fax: (514) 398-4938; Ruffle: Department of Economics, Ben-Gurion University, Beer Sheva, Israel, 84105, e-mail: `bradley@bgu.ac.il`, tel.: 972-8-6472308, fax: (917) 591-8151.

Acknowledgements: We thank David Balan, John Duffy, Erik Eyster, Frank Gerhard, Margaret Meyer, Avi Weiss and seminar participants at the Economic Science Association meetings in Tucson and Amsterdam, and seminar participants in economics departments at Koc University, Lehigh University, McGill University, University of East Anglia, University College London, University of Oxford, Purdue University, Royal Holloway College, University of London, and University of Nottingham for valuable comments. We also thank Moty Fania for programming the posted-offer market software used to conduct the experiments. Engle-Warnick was supported in part by the University of Pittsburgh, Department of Economics through an Andrew W. Mellon fellowship. Ruffle thanks grants from the Russell Sage Foundation and Ben-Gurion University for funding.

1 Introduction

The flourishing of game theory in the late 1970s and 1980s provided experimental economists with a wealth of opportunities to test its predictions under controlled laboratory conditions. Indeed, many economics experiments consist of duplicating the assumptions of game-theoretic models and then establishing that subjects' play does not converge to equilibrium predictions, despite repeated play and monetary incentives (see, e.g., Camerer, 2001; Davis and Holt, 1993; Kagel and Roth, 1995). Through nonparametric statistical tests and regression analysis, researchers conclude that observed behavior differs significantly from the equilibrium outcomes of the model.

When departures from equilibrium play are observed, two approaches are typical. Sometimes, subjects' responses from questionnaires are analyzed in an attempt to glean information about their motivations and intentions. More often, follow-up experiments (e.g., variations in the original game in which one variable at a time is altered) are conducted to separate out competing explanations for the observed deviations from equilibrium. Data from numerous follow-up experiments accumulates and new theories are advanced to explain behavior and unify the body of evidence.

In this paper, to understand deviations from predicted outcomes in our market experiments, we make use of subjects' observed actions to infer repeated-game strategies that best describe their play. Our goal is threefold: first, to develop an easily implementable technique to infer repeated-game strategies from observed actions; second, to illustrate the usefulness of the technique for categorizing subject behavior in a way that is intuitive and interpretable as a strategy thereby providing new insights into the behavior of subjects in this environment; third, to evaluate the success of different strategies.

The game we examine is a posted-offer market in which a monopolist faces a small number of buyers, either two or four; monopolists are either informed or uninformed as to the number of buyers in the market.[1] According to the rules of the posted-offer market, the monopolist posts a price and a quantity of a good to make available at that posted price. Observing the posted price (but not the quantity), the randomly ordered buyers then proceed one at a time to make

[1] The experiments test the role of buyer concentration on pricing in markets. This topic, and buyer countervailing power more generally, have recently caught the attention of economists and antitrust policymakers due, in large part, to the advent of mega-retailers. Ruffle (forthcoming) and Snyder (forthcoming) survey this literature.

privately the number of purchases that each desires. The take-it-or-leave-it nature of this market institution limits buyer strategic behavior to the rejection of profitable purchases, that is, the rejection of purchases at a price below the buyer's valuation, referred to as *demand withholding*. In a finitely repeated game, the unique subgame-perfect equilibrium strategy requires a rational buyer to make all profitable purchases in every period; demand withholding therefore should not be observed. However, withholding behavior has been observed, even intensely, in a number of studies (see Ruffle, 2000, and the references therein); buyers withhold demand in the hope of bringing prices down in subsequent periods.

We attempt to understand withholding behavior in this environment by inferring unobserved, repeated-game withholding strategies from the observed withholding actions of buyers in 30-round experiments. These buyer withholding strategies take the form of (possibly nested) if-then statements. An example of a simple strategy would be to withhold two units of demand if the posted price exceeds a threshold level, but to make all profitable purchases otherwise. To identify strategies that best describe the observed behavior of each individual buyer, we combine a Bayesian procedure with a technique from the binary classification tree literature (Devroye, Gyorfi and Lugosi, 1996). For the first time we are able to formulate repeated-game strategies that people may actually be using in this environment. We report distributional information such as the probability that a buyer's strategy is of different complexities, the probability that the strategy contains certain conditioning variables, the most likely strategy that generated the data and how well it fits the observed decisions.

For about one-quarter of our buyers, we find that the equilibrium prediction of no withholding is accurate. For the remaining three-quarters of the buyers we fit on an individual basis repeated-game strategies that most accurately describe the variables upon which the buyer conditions her withholding decisions.

We find evidence that subjects' actions are consistent with strategies that condition on time, price, and combinations of the two variables. Furthermore, the more complex the strategy (where complexity is measured by the number of binary tests that comprise the withholding strategy), the lower are market prices and monopolist profits. Among simpler strategies, we find that unconditional and intense withholding early in the game is more effective against the monopolist than strategies that trigger withholding above price thresholds or that intensify withholding over time. Since buyer profits are not significantly correlated with these strategy characteristics, and since the observed prices in these sessions were significantly below the monopoly price, the early withholding strategy appears to be an effective counteracting response to monopoly power.

Our procedure builds on the work on identifying subject heterogeneity in experiments by El-Gamal and Grether (1995), who used a Bayesian procedure to estimate decision rules in a population of subjects faced with a static, individual choice task. Their goal was to discover whether people are Bayesian, and they did this convincingly by assuming a functional form for the decision rule and using their Bayesian technique to draw inference regarding the number and types of rules that generated their data. By contrast, our application is to repeated-game strategies, which requires both a different strategy model and computational technique for inference.[2]

Houser, Keane, and McCabe (2004) provide a method to draw inference regarding the number and types of decision rules in a population of subjects playing a dynamic game. They estimate a polynomial approximation of the value function (that is, the expected value of future payoffs) in state variables similar to Geweke and Keane (1999, 2001), allowing individual decision rules to differ by the parameters in their value functions. They illustrate their technique with a game that subjects play against nature based on a model of school choice, and find evidence for interesting behavioral types. This approach is very flexible because the researcher does not have to specify *a priori* the functional form of the decision rules, rather, one simulates the rules to interpret the behavior. Our application gains some efficiency (at the cost of flexibility) by specifying a strategy model. Our strategy model is useful for a rather wide class of multi-player games, and it covers both strategies predicted by theory and simpler rules of thumb. And for games in which if-then statements can be useful characterizations of decision-making, it is rather simple to implement.[3]

The direct inference methodology in this paper is complementary to several existing approaches.

[2] Other probabilistic choice models include Stahl and Wilson (1995) who study heterogeneity in levels of reasoning in games solvable through iterated dominance specifying both the form and number of decision rules, and McKelvey and Palfrey (1995) who introduce Quantal Response Equilibrium, which makes it possible to study subject behavior in deviation from optimality (though not subject heterogeneity).

[3] Houser, Keane, and McCabe note that the two types of procedures can be complementary: one could use their procedure to identify subject types and then an El-Gamal and Grether-type procedure to describe the decision rules.

In the strategy method of Selten et al. (1997), strategy choices are made observable through elicitation; a second approach is to identify strategies by tracking the manner in which subjects collect and process information (Costa-Gomes, Crawford and Broseta, 2001; Johnson, Camerer, Sen and Rymon, 2001); in a third approach a probabilistic choice model is estimated from the data (e.g., El-Gamal and Grether, 1995; Engle-Warnick, 2003; Selten and Stoecker, 1986; Stahl and Wilson, 1995); a fourth approach is to report estimates of probabilistic choice models through a well-specified econometric model (Manski, 2002; McKelvey and Palfrey, 1995); and of course there are the classic approaches of experimental manipulation and protocol responses.

We view the posted-offer market as a laboratory environment well-suited for the display of our method's comparative advantages (and disadvantages). More familiar techniques like the strategy method and the analysis of protocol responses are insightful in simple games. However, the posted-offer institution is sufficiently complex and subjects' behavior is often dynamic responding to feedback during the play of the game that these familiar methods are inadequate.

We begin with a description of the experimental design and a summary of the main qualitative results in sections 2 and 3. Sections 4 and 5 detail the strategy model followed by the inference method. We present the results of strategy inference, beginning with estimates of the distributions of the inferred strategies and ending with examples of specific best-fitting strategies in section 6. Section 7 concludes.

2 Experimental Design and Procedure

In our experiments, subjects were randomly assigned to the role of buyer or seller, and randomly grouped to form markets. Each market consisted of a single seller (i.e., a monopolist) and either two or four buyers. We ask whether differences in buyer concentration can lead to differences in pricing policies in markets, and our simple measure of buyer concentration is the number of buyers that exist in the market.

The monopolists were given units of production to sell in the market, each with an associated cost. The buyers were given units of demand to purchase, each with an associated valuation. To induce subjects to trade, they were told that the sellers would earn, in cash, the difference between the selling price and the cost of each unit sold, and the buyers would earn, in cash, the difference between the valuation and selling price on each unit purchased (see Smith, 1976, for the induced value methodology). The cost of unsold units was not deducted from the monopolists' profits.

Figures 1a and 1b display the monopolist's induced marginal cost curve and the buyers' induced aggregate demand curve for the two-buyer and four-buyer treatments, respectively. The midpoint of the competitive price range has been normalized to 0 and, for convenience, will be referred to as the competitive price. All costs, valuations and prices will be expressed as deviations from this price. In the actual experiments, a constant was added to all costs and valuations so that the subjects traded in positive currency units. The supply and demand curves were induced by providing the buyers with four units of demand, which were valued at +0.35, +0.20, +0.20, and +0.05 normalized units of currency. The sellers in the two-buyer treatment were given eight units of production at a cost of −0.85, −0.85, −0.85, −0.65, −0.65, −0.45, −0.45, and −0.05 normalized units of currency. In the four-buyer treatment, the demand curves were identical, while the monopolist's marginal cost curve was altered to maintain the same monopoly price and competitive price equilibrium.[4]

[insert Figures 1a and 1b here]

The subjects played 30 rounds of a two-stage game. In the first stage, the monopolist is asked to choose a price and a quantity of units to make available for sale at that price.[5] In the second stage, the randomly ordered buyers proceeded sequentially, each deciding how many units to purchase at the posted price. The position of each buyer in the shopping queue was randomly determined in each round.

At the end of the second stage, the monopolist learned the total number of units purchased by the buyers and the profit earned from these sales. Buyers did not learn the number of units other buyers purchased. In each round, subjects remained grouped in the same cohort, played the same role in the game, and were given exactly the same individual marginal cost and demand curves.

[4] See Engle-Warnick and Ruffle (2005) for further discussion and a more complete justification of the experimental design.

[5] The choice of quantity is in fact trivial: the monopolist should make available all units whose cost is less than (or equal to) the selected price. With very rare exception, all sellers in all rounds chose such a quantity and therefore we focus on the monopolists' pricing behavior in all subsequent analysis.

Note well the private nature of buyers' purchasing and withholding decisions. Only the individual buyer knows how many units she purchased and how many profitable units she passed up. This institutional detail makes it impossible for buyers to coordinate, even implicitly, their responses to the monopolist. More importantly for the purposes of our paper, the independent nature of buyers' actions allows us to infer individual buyer strategies.

The unique Nash equilibrium of the stage game involves the monopolist posting the monopoly price (+0.20) in each round and the buyers accepting all profitable purchases at this price. According to our design, each buyer (again independent of treatment) is able to make three profitable purchases at the monopoly price. The buyer earns a total of +0.15 (+0.15 on the first purchase and 0 on the second and third purchases). The monopolist, on the other hand, sells six units in the two-buyer treatment and earns +5.50, compared to 12 sales in the four-buyer treatment and +6.40 in earnings.

When the game is repeated it is no longer a dominant strategy to accept passively all profitable purchases because with future periods left in the game a buyer can punish the monopolist by withholding profitable units of demand. This type of strategy is more costly to implement as fewer periods remain in the game because the potential gain from lower prices diminishes. Thus, while there is no cooperation in equilibrium in the complete-information game, the incomplete information framework in Kreps et al. (1982) for finitely repeated games may be relevant. Our contribution is to search for empirical evidence of this type of strategic behavior.

The curves of the relevant treatment as well as the individual buyers' demand curves were made common knowledge by providing each subject with a table of costs and values of all subjects and by reading aloud the contents of the table. The market structure in these two treatments was also common knowledge so that the monopolist (and the buyers) knew precisely how many (other) buyers were in the market.[6]

All experimental sessions were computerized and conducted at Ben-Gurion University. At the end of the experiment, subjects were paid a 15 NIS (New Israeli Shekel) showup payment in addi-

[6] The results of these "informed" two-buyer and four-buyer treatments led us to design two additional treatments in which the monopolist only was uninformed of the number of buyers in the market. These treatments, in which the monopolist was not told the number of buyers in the market, will be referred to as "uninformed" treatments.

tion to their experimental earnings.[7] Average seller earnings (including the showup payment) were 121 NIS compared to 67 NIS for the buyers. Sessions lasted on average one hour and thirty minutes. Seven two-buyer (informed) sessions were conducted along with eight four-buyer (informed) sessions. All subjects were economics or business majors and had taken at least an introductory course in microeconomics. Participation was restricted to one session only per subject.

3 Experimental Results

We conducted first the pair of two-buyer and four-buyer treatments in which the monopolist in each session was informed of the precise number of buyers she faced. Prices in both of these informed treatments were typically well below the monopoly price. Moreover, as Figure 2 shows, prices in the two-buyer, informed treatment are significantly lower than prices in the four-buyer informed treatment. What is more, buyers in the two-buyer, informed sessions achieved these lower prices without withholding more than those in the four-buyer, informed sessions: the average per buyer, per period number of units withheld is identical in the two treatments.[8]

[insert Figure 2 here]

Given the controls built into the experimental design, there are two possible explanations why prices in the two-buyer, informed treatment are lower, despite identical levels of withholding in the two treatments. First, the monopolist may simply price more cautiously when confronted with two buyers than when faced with four. Since the monopolist earns zero on sales lost to demand withholding, she may choose to post lower prices in the two-buyer sessions for fear of provoking their withholding. To explore this hypothesis, we conducted a second pair of "uninformed" treatments with identical marginal cost and demand parameters to those employed in the "informed" treatments. The sole difference between the "uninformed" and the previous pair of "informed" treatments was that in the former, the monopolist was not told how many buyers she faced; instead, she was told in both the two-buyer and four-buyer uninformed treatments that she faced "a

[7] At the time these experiments were conducted 4 NIS was equivalent to approximately $1 U.S.

[8] Since our focus is on the development and application of the strategy inference technique, we present just enough of the experimental results by treatment to appreciate the application of the technique. Engle-Warnick and Ruffle (2005) provide a detailed analysis and discussion of the experimental results by treatment.

small number of buyers, but more than one". We conducted eight two-buyer, uninformed sessions and seven four-buyer, uninformed sessions.

The most striking result in the uninformed treatments is that the disappearance in the initial and middle rounds of the price gap between the two-buyer and four-buyer informed treatments, as seen in Figure 3. This suggests that the observed difference in initial pricing in the informed treatments is, at least in part, due to the monopolist pricing more cautiously when confronted with only two buyers.

[insert Figure 3 here]

A second possible explanation for lower prices in the two-buyer, informed treatment is a difference in the *quality* of withholding between the treatments. For instance, perhaps the buyers in the two-buyer sessions condition their decisions to withhold on different variables than those in the four-buyer sessions, and these strategies are more effective in bringing prices down. The strategy inference technique developed in the next two sections will allow us to address this hypothesis.

4 The Strategy Model

We model buyer strategies as binary classification trees (see Breiman, Friedman, Olshen and Stone, 1984, for a variety of applications), and introduce the strategy model by way of example with an actual inferred buyer strategy, shown in Figure 4. This strategy contains three *relational nodes*, which are represented by filled circles. Relational nodes are binary tests that consist of a variable, a relation, and a coefficient. The relational node at the top of the tree (called the root node) tests whether the price at time t, $P(t)$, is less than or equal to 0.01. The strategy also contains four *action nodes* at the bottom of the tree, each marked by an empty circle. The left column of values below each action node lists each possible buyer decision; "0" represents the decision to withhold zero units of demand, "1" to withhold one unit, and "+" to withhold more than one unit.[9] The

[9] There is a theoretical reason for limiting the choices to 0, 1 and more than 1 unit of demand withheld. We endowed each buyer with "market power" by making it possible to lower profitably the market-clearing price by unilaterally withholding two units of demand. A buyer who withholds 0 units of demand is thus a passive price-taker, a buyer who withholds 1 unit is not passive but does not exercise market power, and a buyer who withholds 2 or more units exercises market power. Holt (1989) proposes an analogous definition of seller market power, which we adapt to buyers.

corresponding right column indicates the number of rounds (out of a total of 30) in which the buyer made each of the corresponding decisions: at the right-most action node this buyer withheld no units zero times, one unit zero times, and more than one unit five times.

[insert Figure 4 here]

The tree in Figure 4 thus defines both the functional form of an actual strategy and classifies the observed actions of an actual buyer. Evaluation of the expression begins with the root node. If the expression $P(t) \leq 0.01$ is true, then evaluation proceeds down to the root node's left-hand descendant node, which is the tree's left-most action node. If the expression $P(t) \leq 0.01$ is false, then evaluation proceeds down to the root node's right descendant node, which is a nested relational node that specifies the test $P(t) \leq 0.04$. Taken together these two relational tests represent the compound expression IF $P(t) > 0.01$ AND $P(t) \leq 0.04$. If this compound expression is true, then evaluation proceeds to the tree's left-center action node, and if $P(t) > 0.01$ AND $P(t) > 0.04$ a third relational node labelled $t \leq 10$ is reached. Evaluation proceeds at this point as before.

The grammar implicit in constructing the strategies is quite general: since the boolean operators AND and NOT ($>$ is the same as NOT $\leq$) are both included in this grammar, without loss of generality, any boolean expression may be the result of combinations of relational and action nodes. It is the case that every decision in the data will always fall to exactly one action node, hence the strategy is a plan of action for all observed contingencies in the repeated game.[10]

The binary tree representation of the strategy lends itself well to behavioral interpretation. The buyer whose actions are represented in Figure 4 never withheld demand whenever the posted price was low enough (below 0.01), increased withholding intensity to one unit of demand for intermediate prices (prices within the range 0.01 to 0.04), and after round 10 withheld demand intensely for high prices (above 0.04).

The class of strategies we consider consists of no more than four relational nodes. We restrict the number of relational nodes to four due to the impracticality of further classifying thirty decisions with more relational nodes since the amount of data at each terminal node drops exponentially.

[10] We abstract from semantics and refer to these decision rules as strategies, noting well the fact that in many cases, including our application, we observe only a single history of each repeated game and thus cannot reconstruct the rule that a subject would have used for any possible contingency.

In practice this restriction was not binding as the probability of four-node strategies in the data turned out to be very low and 89/90 buyers' inferred strategies consist of *strictly* less than four relational nodes.

We selected the round number, t; the round t price, $P(t)$; the difference in price from one round to the next, $\Delta P(t)$; and the withholding decision in the previous round, $W(t-1)$. We allowed any coefficient actually realized during the experimental session. Since $\Delta P(t)$ never occurred in any inferred strategy, and in the few instances that $W(t-1)$ did appear, it did not contribute to a useful interpretation of what the subjects may have been doing, we focus on t and $P(t)$ in the subsequent analysis.

The current round price is perhaps the most obvious variable upon which to base one's withholding decisions. We included the round number to allow the strategies to vary with time. The appearance of time in an inferred strategy may suggest that the subject adopted one strategy at a certain point in the game and later discarded it for another strategy. That subjects may intentionally vary their demand withholding over time may be seen in the numerous buyers who, independent of prevailing prices, withheld early to signal toughness, and then later in the session, as the expected future gains from withholding dwindled, ceased to withhold.[11]

Due to the vast number of logically possible repeated-game strategies, the solution to the inference problem always requires a reduction of the dimensionality of the problem. We performed this reduction in at least three ways: we limited the types of demand withholding strategies through the experimental design itself (buyers do not see the choices of other buyers and the number of conditioning variables available to the buyers are few); we modeled the strategy choices as withholding 0, 1 or more than 1 unit of demand (recall the theoretical reason from footnote 9 for doing so); and we limited the maximum number of relational nodes to four.[12]

[11] We do not include multi-period punishment strategies of the variety that withhold demand for a fixed length of time independent of what happens during the punishment phase because the posted-offer institution renders such forms of punishment ineffective: the monopolist is unaware that it is the same buyer who withheld, say, two units in each of the last four periods. In principle, however, it is possible to include multi-period punishments by adding a length of implementation variable at the bottom of the tree (this length is implicitly set to one round in our model).

[12] Other successful studies have followed dimension reduction strategies as well. El-Gamal and Grether (1995) approach the problem through the *a priori* specification of functional forms of decision rules. Houser, Keane and McCabe (2004) do so through experimental design. Their subjects gain experience playing a game against a known stationary stochastic process, thus the subjects' strategies are assumed to be stationary, and the stochastic process

5 The Bayesian Inference Method

The inference task is to find an unobserved strategy that best describes the observed actions of a subject. A difficulty is that adding complexity to a strategy will always weakly improve its ability to describe the data, but we lack a theory to guide us in selecting the appropriate level of complexity for a strategy; thus we need a method that tells us when additional complexity does not improve the fit of the strategy enough to be worthwhile. We approach this problem with a Bayesian estimation of repeated-game strategies analogous to the procedure in El-Gamal and Grether (1995), and adopt their notation.[13] We assign priors to each of the assignment problems inherent in constructing a strategy, and update by observing a sequence of decisions by an individual in a repeated game.

5.1 Defining the Assignment Problem

Let us illustrate the assignment problem by constructing a single strategy. To begin, we select the universe of n conditioning variables (an operation we call assignment zero). To construct a complete strategy requires three further assignments: (1) the *number* k of relational nodes, (2) the conditioning *variable* at each relational node, and (3) the *configuration* of the relational nodes in the tree representation of the strategy.

The first assignment, the number of relational nodes, can in theory be any nonnegative integer, k. In Figure 5, we choose $k = 2$ for illustrative purposes.

[insert Figure 5 here]

The second assignment involves choosing k variables for use in the k distinct relational nodes. Anticipating the fact that no two relational nodes that contain the same variable will have the same

provides power for strategy inference. In neither of these cases did simplification inhibit what was learned about decision making in economic markets.

[13] In their application to an individual choice problem, a subset of k decision rules is selected from a set of n candidate decision rules, denoted C, to best-fit the actions of a population of subjects in an experiment. There are three assignments to be made in this problem: (1) the *number* of decision rules k that should be used to best-fit the data, (2) the *specific* decision rules to take from the candidate set C to best-fit the data, and (3) the best *assignment* of these k specific decision rules to each individual subject. El-Gamal and Grether (1995) approach the problem by forming priors over each of these assignments, and then finding the posterior mode estimate of their joint assignment. Our approach extends their technique by inferring *repeated-game* strategies on a *subject-by-subject* basis. We also report the estimate of the posterior distribution for each individual's strategy.

coefficient, we construct a candidate set C of $k \cdot n$ elements that contains k replications of each of the n explanatory variables. We draw without replacement from this set of variables to select the relational nodes. The example in Figure 5, in which $k = 2$ and $n = 2$, has $C = \{P_1(t), P_2(t), t_1, t_2\}$. Ignoring order until the third assignment, it is easy to see that there are six possible ways to choose two elements from this set of relational nodes, as shown in Figure 5.

The third assignment is to take the k variables from the second assignment and configure them into a strategy. The lower part of Figure 5 shows the four ways that this can be done. Notice that both the configuration of the tree and the location of the variables within the configuration vary.

5.2 Forming Priors Over the Three Assignments

We first define the prior for k, the number of relational nodes in the strategy. In theory k can be any nonnegative integer (where $k = 0$ represents no binary tree), however we impose the restriction $k \in \{0, 1, 2, 3, 4\}$, and specify a Poisson distribution for the prior $P(k)$ truncated at $k = 4$. The Poisson distribution allows us to weight strategies with fewer relational nodes more heavily while varying only one free parameter, which is intuitively appealing if one believes that more complex strategies are less likely (see Denison, Mallick, and Smith, 1998 for priors in binary regression trees and Camerer, Ho, and Chong, 2001 for heterogeneity in bounded reasoning):

$$P(k) = \frac{\frac{\lambda^k}{e^{\lambda} x!}}{\sum_{j=0}^{4} \frac{\lambda^j}{e^{\lambda} j!}},$$

where λ is the mean of the distribution. Thus the specification of this prior amounts to the specification of the mean λ of the number of relational nodes. We will report the sensitivity of the inference results to a range of λ.

We assign an uninformative (i.e., uniform) prior over the selection of k decision nodes from the set C of $k \cdot n$ nodes for use as strategy components. For this we need to calculate the number of such possible choices that exist. The calculation is straightforward: the number of combinations of $k \cdot n$ objects taken k at a time is

$$S_n^k = \begin{pmatrix} k \cdot n \\ k \end{pmatrix} = \frac{k \cdot n!}{k!(k \cdot n - k)!}.$$

In the previous example there are 6 ways to select $k = 2$ relational nodes from the set of $n = 2$ candidate nodes.

We also assign an uninformative prior over the possible configurations of the selected relational nodes. In general, for the chosen k relational nodes, there are k ways to select the first node of a decision tree (by choosing one of the k nodes). Now whenever a relational node is added to a tree, two action nodes are also added, but since the new relational node uses one of the existing action nodes, the net addition to the number of action nodes is one. Thus after selecting the first node, $k-1$ variables remain for assignment at one of two action nodes. Repeating at the third level, $k-2$ variables remain for assignment at one of three action nodes. This process continues until only one variable remains for assignment at k action nodes. The number of tree configurations that can be formed from k decision nodes is thus given by,

$$T^k = \prod_{j=1}^{k} j \cdot (k - (j-1)).$$

In summary, to construct a strategy we select the number of nodes, k, to be used in the formation of a strategy; we hypothesize a set of n candidate relational node variables, and choose k specific nodes from them; and we then construct all possible trees from k nodes using the chosen variables. Thus the total number of strategies with k nodes that can be formed from n candidate variables is the product of the two preceding results:

$$U_n^k = S_n^k \cdot T^k = \frac{(k \cdot n)!}{k!(k \cdot n - k)!} \cdot \prod_{j=1}^{k} j \cdot (k - (j-1)).$$

In our previous example, there are $6 \cdot 4 = 24$ possible ways to construct a strategy with $k = 2$ decision nodes from a set of $n = 2$ relational nodes.

Thus the prior probability distributions are as follows: the prior probability of k relational nodes $P(k)$ is given by the Poisson distribution truncated at $k = 4$, the prior probability of selecting any k relational nodes from a set of n variables is $1/S_n^k$, and the prior probability of any particular arrangement of k relational nodes into a strategy is $1/T^k$.[14]

[14] Our assumption regarding the error rate in strategy implementation will be that each possible rate is equally likely; this makes the procedure simpler to implement. Chipman, George and McCulloch (2001) provide a Markov Chain Monte Carlo method assuming a multinomial distribution for the errors; we tested our data using their software and found similar strategies.

5.3 Building the Likelihood Function

For the probability model, let $x_t = (x_t^{(1)}, \ldots, x_t^{(d)}) = (P(t), t)$ be the set of explanatory variables observed at time t. Let $y_t = W(t) \in \{0, 1, +\}$ be the action taken by the subject at time t. The set of actions corresponds to the decisions to withhold 0, 1, or more than one unit of demand in a given round.

We consider a class of strategies called binary decision trees, which consist of a set of relational nodes and a set of action nodes. Each relational node is a test of the form $x_t^{(i)} \leq \alpha_j^{(i)}$, where $\alpha_j^{(i)}$ is a coefficient. The set of coefficients that was actually realized in the experiment for each variable constitutes the set of admitted coefficients. If the evaluation of the test is true (false) then a left (right) descendant node is reached. Descendant nodes of relational nodes may be either relational nodes or action nodes. Action nodes specify an action to take and have no descendant nodes. Hence a relational test may be followed by a subsequent (i.e., nested) relational test, or it may be followed by the specification of an action to take.

A strategy $g \in G$ can be thought of as a possibly nested if-then statement that always specifies an action to take conditional on realizations of variables determined to be important to decision-making. Every data point y_t, x_t, when dropped through the decision tree, always reaches exactly one action node. Thus we can compare the actions specified by the strategy with every action observed in the data. It follows that we interpret the strategy as a plan of action for every possible observed contingency in the game.

Let $D_g|x_t$ denote the action specified by the strategy given the observed data at time t. Define the variable

$$x_{g,t} = \begin{cases} 1 & \text{if } y_t = D_g|x_t \\ 0 & \text{otherwise.} \end{cases}$$

That is, $x_{g,t} = 1$ if the action specified by the strategy agrees with the action taken by the subject. Summing the number of actions that agree with the strategy yields $X_g = \sum_{t=1}^{T} x_{g,t}$.

A subject takes the action specified by a strategy with probability $1 - \epsilon$, and randomizes with equal probability among each possible action with total probability ϵ whenever she deviates from the strategy. As in El-Gamal and Grether (1995), a strategy $g \in G$ and error rate ϵ define a probability function $f^{g,\epsilon} : X \to [0,1]$, where the index k represents the number of relational nodes contained by the function. Noting again the private nature of both individual buyers' withholding decisions and the monopolist's quantity sold effectively rules out dependence among different buyers' withholding decisions, we specify the likelihood function for each individual buyer:

$$\hat{g}, \hat{\epsilon} = \arg\max_{g,\epsilon} f^{g,\epsilon}(x_1, \ldots, x_T) = (1 - \frac{2\epsilon}{3})^{X_g} \times (\frac{2\epsilon}{3})^{T - X_g}.$$

The estimate for $\hat{\epsilon}$ is found by forcing the rule to make the decision that occurs most often at each action node, summing the number of decisions in the data that do not agree with the strategy specification, and then dividing by the total number of decisions made.

5.4 Computing the Posterior Mode Estimate

The posterior mode estimate of the joint assignment of the number of nodes, the specific nodes, the arrangement of the nodes in the tree, and the error is:

$$\hat{k}, \hat{s_n^k}, \hat{t^k} = \arg\max_{k,s,t,\hat{g}} \{X_{\hat{g}} \log(1 - \frac{2\hat{\epsilon}}{3}) + (T - X_{\hat{g}}) \log(\frac{2\hat{\epsilon}}{3}) - \log(P(k)) - \log(S_n^k) - \log(T^k)\}.$$

The first two terms specify the fitness of the strategy, while each of the last three terms corresponds to a prior probability for one of the three assignment problems. The priors have the effect of penalizing the likelihood function for the model complexity, since each of the three terms is a decreasing function of the number of relational nodes in the strategy. We therefore force a tradeoff between strategy fitness and complexity, where strategy complexity is crudely approximated by the number of relational nodes. A possible outcome of this tradeoff is that the degenerate strategy with no conditioning nodes at all is inferred from the decisions of the subject. Our conservative approach to strategy inference as implied by the triple penalty for each additional relational node implies that those strategies that are inferred perform well at categorizing the subject's withholding behavior.

5.5 Computational Strategy

We compute the posterior likelihood for each possible strategy formulation for $k \in \{0, 1, 2, 3, 4\}$ for each buyer.[15] We report the posterior probability for each of the possible values of k, as well as the probability that each explanatory variable along with its impact on withholding occurs in the strategy. For robustness we report these results as a function of the mean, λ, of the Poisson prior for k, where λ varies from 1 to 3 in increments of 0.1. We report summary statistics for the modal strategies inferred at a selected value of the mean, and we report the modal strategy for all subjects in selected experimental sessions.

6 Strategy Inference Results

We present the results in three subsections. First, we present inference results aggregated across experimental treatments. Second, we present posterior distributions for each buyer in the form of strategy characteristic statistics such as, for example, the probability that the buyer's strategy contains k nodes, $k = 0, 1, 2, 3, 4$, and that it contains specific strategy characteristics or components. We also explore correlations between specific strategy characteristics and market outcomes. Third, we display inferred strategies that correspond to the posterior mode estimate for each subject in selected sessions.[16] For the purposes of introducing this strategy inference method and for brevity we focus primarily on buyer strategies.

6.1 Aggregate Results on Strategy Complexity and Strategy Composition

6.1.1 Strategy Complexity

On average, the complexity of the inferred strategies is greater than the degenerate case of zero relational nodes, and is relatively insensitive to its prior. Figure 6 presents the number of relational

[15] An exhaustive search is possible due to our limitation of $k \leq 4$; the number of possible strategies was typically on the order of 100,000. Our approach is not however limited to cases in which an exhaustive search is possible. The software for estimation was programmed in Ox 3 Professional (Doornik, 2001) and is available upon request from the authors.

[16] In fact, we computed the modal strategy for all subjects in all sessions. Due to space restrictions, we present the modal strategies for all buyers in two specific sessions in section 6.4.

nodes in the strategy that corresponds to the posterior mode estimate, averaged across all subjects in all treatments. This measure attests to the relative stability of the number of nodes in the inferred modal strategies to changes in the prior mean number of nodes: it varies by only 0.6 nodes, from 0.9 to 1.5, as the prior mean number of relational nodes varies from one to three. (We varied the prior from one to three because we limited non-zero probabilities between zero and four nodes.) The expected number of relational nodes (computed as the weighted average of strategy size over all possible strategies) is similarly stable: Figure 6 also shows that the expected number of relational nodes varies by only 0.6 nodes, from 1.6 to 2.2, as the prior varies from one to three. The fact that the average strategy size for both of these measures consists of more than one relational node despite the triple cost of complexity suggests the inferred repeated-game strategies fit well buyers' decisions.

[insert Figure 6 here]

Figures 7a and 7b display the average number of relational nodes in the modal strategies according to the experimental treatment variable. Figure 7a reveals that modal strategies are slightly more complex in the informed experimental treatments. This result was unexpected as it is the information given to the monopolist, not the buyers, that is being manipulated. One conjecture is that since the monopolist does not know how many buyers she faces in the uninformed treatments, buyers feel they must adopt simpler, more transparent strategies and stick with them to signal clearly to the monopolist what is unacceptable; whereas buyers in the informed treatments have the luxury to be able to fine tune (i.e., make more complex) their withholding behavior. Strategy complexity as a function of the number of buyers does not reveal a similarly evident relationship, as shown in Figure 7b.

[insert Figures 7a and 7b here]

6.1.2 Strategy Composition

For a look at the composition of the buyer strategies, Figures 8a and 8b display the probability that the buyer strategies contain a relational node that increases and decreases withholding with time respectively. We denote these node types "Time +" and "Time -", respectively, and similarly define the node types "Price +" and "Price -". For example, the $t \leq 10$ relational node in Figure 4 is a Time + node because the buyer increases her withholding after round 10. Each of the graphs in Figures 8a and 8b corresponds to one of the four treatments. Figures 9a and 9b present the identical statistics for relational nodes that condition on price.

[insert Figures 8a and 8b here]

These figures reveal important differences between the experimental treatments. Most notably, Figure 8a shows that nodes decreasing withholding with time are more likely to occur in strategies in the two-buyer, informed treatment than in the four-buyer, informed treatment. Figure 8b shows that the opposite is true for increasing withholding with time, namely, that buyers in the four-buyer, informed treatment are more likely to increase their withholding with time than buyers in the two-buyer, informed treatment. This marked difference in the strategies employed by buyers in the two-buyer and four-buyer informed treatments points to a possible explanation for the lower observed prices in the two-buyer, informed sessions: decreasing withholding over time appears to be a more effective strategy against the monopolist than increasing withholding over time. We will have more to say about this hypothesis in the next two subsections.

In addition, Figure 9a shows that decreasing withholding with price is equally highly unlikely in all treatments; this intuitive result is reassuring since we do not expect subjects to buy more units (i.e., decrease their withholding) as the price increases. Figure 9b shows that increasing withholding with price is more likely to occur when monopolists are informed.

[insert Figures 9a and 9b here]

6.2 Strategy Characteristics for Individual Buyers

In this subsection we examine more closely the inferred strategy characteristics at the subject level by reporting the estimates of the posterior distributions of inferred strategies for individual buyers. Since it is impractical to present this analysis over the entire range of prior means of relational nodes (i.e., 1 - 3), we choose the midpoint of the range, two, for the analysis. The good news is that all of the strategy complexity and strategy composition results from the previous subsection were shown to be insensitive to the prior mean number of relational nodes.

Table 1 displays the results for the two-buyer, informed treatment, and Tables 2, 3, and 4 contain identical information for the remaining three treatments. Each row reports the strategy characteristics and observed game variables for a different buyer. From left to right the tables reveal the session and subject identification numbers, the median session price from the last five rounds, the mean seller and buyer per round profits, the breakdown of each buyer's observed withholding decisions (i.e., the number of rounds in which the buyer withhold 0, 1 and multiple units of demand), the number of relational nodes inferred in the modal strategy (labeled "size"), the error rate of the modal strategy (to be discussed below), and the posterior probability that the buyer's strategy contains 0, 1, 2, 3, and 4 relational nodes (which sum to 1). The last set of columns report the posterior probability that the strategy contains the conditioning nodes Time - , Time +, Price - and Price +. These may sum from zero (when the posterior probability of a zero node strategy is one) to four (when only strategies that contain at least one of each of the four nodes have non-zero posterior probabilities).

Table 1 reveals the richness of information regarding buyer heterogeneity that can be gleaned from a strategy inference procedure. Overall, for three-quarters of the buyers in the experiments the modal strategy is a non-degenerate strategy with at least one relational node. There is much heterogeneity among these subjects with respect to the size of the inferred modal strategy (between 1 and 4 relational nodes) and the posterior probabilities of the strategy characteristics. That the inferred modal strategies differ across buyers both in terms of the strategy characteristics and in size despite the common prior for all buyers suggests that the strategies are not being overwhelmed by the prior and that they describe the buyers' decisions well. Further evidence that the strategies fit the decisions well can be seen by the error rate, which appears in the "modal strategy characteristics" column of Tables 1–4. Across all treatments, the average error rate is approximately 0.16, meaning that strategies on average classify correctly 84% or 25/30 of the buyer's observed withholding decisions. By comparison, if we fit each buyer with a single relational node, best-fit strategy (either withhold 0, withhold 1 or withhold more than 1 unit), the average error rate is 0.33, implying that the extra complexity our inference method permits cuts the error rate by 50%.

[insert Tables 1–4 here]

In the two-buyer, informed treatment, the Time - variable appears far more frequently than any other conditioning variable in the inferred modal strategies of buyers: for 9/14 subjects, the probability that their strategy contains a Time - node exceeds 0.8. By contrast, there is not a single buyer who conditions on Time + with probability greater than 0.8. Two subjects increase their withholding with an increasing price (Price +) with probability greater 0.8. For non-degenerate strategies, error rates vary from 0.033 for a two-relational node strategy to 0.367 for a one-relational node strategy. The actions of four buyers varied so little that we could only construct the degenerate zero-node strategy. Notice that although Buyer 2 in Session 1 shows considerable variance in her withholding actions — she withheld one or more units in 13 rounds — her apparently somewhat random withholding pattern did not admit the inference of a non-degenerate modal strategy. As a result, the error rate for this buyer is 0.433. As a matter of fact, Session 1 was the lone session in this treatment for which we failed to infer a non-degenerate strategy for both buyers.

In the four-buyer, informed treatment (Table 3), there is a noticeable shift from Time - nodes to Time + nodes in comparison with the two-buyer, informed treatment. Ten out of 32 buyers condition on Time - with probability greater than 0.8. This same fraction conditions on Time + (compared to no buyers in the two-buyer, informed treatment). Figures 8a and 8b confirm that this treatment is an outlier both in terms of the high posterior probability that an average buyer's strategy contains a Time + node and the low posterior probability that it contains Time -. In addition, six buyers employ Price + with probability greater than 0.8. We inferred non-degenerate modal strategies for two or more buyers in all eight sessions.

In the two-buyer, uninformed treatment (Table 2), strategies again vary by complexity and composition. Indeed the patterns of strategy complexity and composition are broadly very similar to those in the two-buyer, informed treatment. Again there is a single session (Session 29) in which we failed to infer a strategy for both buyers.

In the four-buyer, uninformed treatment (Table 4), there are far fewer Price + nodes with probability close to 0.8 or greater (only two). High posterior probabilities are more spread out among Time -, Time + and Price - nodes: for instance, with probability greater than or equal to 0.8, 17/28 buyers condition on Time - and four buyers condition on Time +.

6.3 Strategy Effectiveness against the Monopolist

We investigated correlations between strategy characteristics and session price, and seller and buyer profits. The results aggregated across all treatments are presented in Table 5. Cells marked in boldface represent statistical significance at the 10% level according to the Spearman Rank Coefficient Test (two-tailed test). The table reveals that both the modal and expected strategy size are negatively and significantly correlated with the median session price: the more complex the strategies used by buyers, the lower the session prices. Similarly, strategy complexity is negatively and significantly correlated with seller profit, but correlation between complexity and buyer profit is not significant. It seems that the lower prices that buyers achieve through withholding roughly compensate them for the foregone profit from the withholding. A further indication of the effectiveness of withholding strategies is given by the positive (0.380) and significant correlation between inferring a degenerate strategy (P(Zero Nodes) in Table 5) and session price.

[insert Table 5 here]

Correlating the different conditioning nodes with session price suggests that the most effective strategy component against seller pricing is Time - . In fact, Time - is the only strategy component significantly (and negatively) correlated with price: when strategies decrease withholding over time, session price (and seller profits) tend to be significantly lower. The only strategy component that is significantly correlated with buyers' profits is Price - . The negative correlation between the two highlights the obvious fact that if a buyer increases her purchases (i.e., decreases her withholding) when the price increases, her profits will be relatively low.

To investigate further relationships between strategy components and market outcomes we ran regressions with the median session price from the last five rounds of the game, the average seller profit, and average buyer profit as dependent variables. Possible independent variables were the probability that strategies contained Time -, Time +, Price +, and Price - nodes, the size of the modal strategies, and dummies for the two-buyer and informed sessions. All variables were averaged across the individual values for each buyer found in Tables 1–4. We used the software package PcGets (Hendry and Krolzig, 2001) to test down to the final model presented below, and for simplicity report results from one-stage estimation.[17]

The results are displayed in Table 6. The table presents the estimated coefficients with standard errors in parentheses. All coefficients are significant at least at the 5% level.

[insert Table 6 here]

The table reveals that two buyers, the probability of a Time - node, and the probability of a Price + node all negatively influence the market price, as measured by the median price from the last five rounds. All three effects are of similar magnitude. Not surprisingly, these same three variables negatively affect seller profits. Concerning buyer profits, only the number of buyers has a significant effect: median buyer profits are 0.3 NIS per round higher in the two-buyer than the four-buyer treatments. This follows from lower prices in the two-buyer sessions, while the quantity of demand withholding does not differ between treatments.

The results provide insight into the effectiveness (and lack thereof) of various withholding strategies. Increasing complexity and decreasing withholding over time (i.e., early withholding) are negatively correlated with seller prices and seller profits. These characteristics are not significantly correlated with buyer profit, even though buyers forego profitable purchases by employing these strategies. Combining this evidence with the existence of remarkably low session prices paints a picture of strategies that influence the monopolist pricing decision at no discernible cost to the buyers.

One plausible explanation for the effectiveness of early withholding is that it signals the buyers' unwillingness to accept existing prices. The monopolist thus responds by *lowering* her price to

[17] PcGets automatically selects a final (i.e., specific) model that is congruent with data evidence by starting with a congruent general model, eliminating statistically insignificant variables, and checking the validity of the reductions with diagnostic tests. We used PcGets to remove subjectivity in our choice of the final model to present.

increase sales and profits. If buyers do not withhold early on, feeling encouraged, the monopolist may raise her price and continue to do so until she is met with resistance.

6.4 Strategies in Selected Sessions

6.4.1 Buyer Strategies

The actual modal strategies inferred in the individual sessions reflected the heterogeneity described by the posterior distributions reported in Tables 1-4. We selected two sessions with highly interpretable results to complement visually the distributional and statistical results reported in the previous subsection. We present the modal strategies inferred in both a two-buyer, informed session and a four-buyer informed session. We chose informed sessions because they constitute the most likely conditions in which the buyers can convey a clear message to the sellers through strategic withholding. The contrast between a two-buyer and a four-buyer session illustrates the relative effectiveness of withholding strategies employed by buyers in the two-buyer sessions.

The strategies inferred in two-buyer session 2B7 are presented in Figure 10, and the price series is shown in Figure 11. The modal strategy contains three relational nodes for Buyer 1 and one relational node for Buyer 2. Buyer 1 did not withhold through period 5, but became price sensitive thereafter: whenever the price exceeded -0.30, he withheld multiple units of demand nine out of eleven times (Price +). Whenever the price was less than -0.30, he withheld a single unit of demand up to period 24 and zero units after period 24 (Time -). Buyer 2 withheld multiple units of demand 16 times in the first 21 periods and purchased all profitable units thereafter eight out of nine times (Time -). Notice from Figure 11 that prices in Session 7 vary substantially up to period 21; thus, Buyer 2's unconditional withholding strategy up to period 21 is not an artifact of a lack of variation in the explanatory price variable. The intense early withholding in the game combined with one buyer's sensitivity to an already low price (-0.30) appears to have driven the monopolist's price down to a remarkably low level in this session.

[insert Figures 10 and 11 here]

The strategies inferred in four-buyer session 4B10 are presented in Figure 12, with the price series shown in Figure 13. Buyer 3 did not withhold demand on even a single instance. The three remaining buyers exhibited behavior consistent with sensitivity to a price threshold. The inferred price thresholds of 0.04, 0.03 0.00, -0.02 and -0.03 all fall within the competitive price tunnel, well below the monopoly price of 0.20. Two of these three buyers' strategies (Buyers 1 and 4) also contain a time relational node that reflects unconditional demand withholding early in the game. Again from Figure 13 one can see that this is not for want of price variance. Like Session 2B7, the willingness to withhold demand early combined with price sensitivity appears to have kept pricing far below the monopoly price, and in the region of the competitive range.

[insert Figures 12 and 13 here]

6.4.2 Seller Strategies

A natural question to ask is what strategies might the monopolist be employing. To answer this, we ran the inference procedure on the sellers, allowing the seller strategies to condition on the number of sales lost to withholding at time $t-1$ and the round number. We report the best-fitting strategies from Sessions 7 and 10 in Figure 14. The strategies are interpreted in the same way as the buyer strategies; seller decisions are represented by a "-" indicating the decision to lower the price from time $t-1$ to time t, a "0" indicating no change in the price, and a "+" indicating an increase in the price.

The left-hand portion of Figure 14 reveals that the seller in two-buyer Session 7 tends to lower the price in response to demand withholding in excess of one unit. The right-hand portion of Figure 14 presents the inferred strategy of the relatively active seller in Session 10: if more than three units of demand are withheld the monopolist lowers the price, otherwise the modal decision is to increase the price. And these are not isolated examples: the best-fitting strategy is non-trivial for the majority of the sellers in our experiments.

[insert Figure 14 here]

7 Conclusions

In this paper, we apply a Bayesian inference method that estimates repeated-game strategies for individuals based on their observed play. The method is quite general: the researcher need only have access to the set of possible conditioning variables and their realized levels observed by the decision makers and the actual decisions made. With this in hand, the researcher must formulate a logical way in which to construct strategies for the purpose of forming priors as well as a probability model to compute a likelihood and the computation of Bayes' rule. The method is easy to use and supplements classical methods of inference in informationally deficient environments.

The potential output from the application of this inference method is also quite general: the researcher can report many different outputs, depending on the question of interest, including the posterior mode estimate of the strategy, the probability distribution over all of the different possible strategies, the probability of specific strategy components and even, if desired, the probability of combinations of different components occurring in the strategy.

We apply this method to 30-round, posted-offer monopoly experiments to infer unobserved, repeated-game strategies from the observed actions of buyers and sellers. As a first application of this technique, the posted-offer market is well suited: buyers independently make their purchase decisions with their only available, non-trivial action being the rejection of a profitable purchase, referred to as demand withholding. On the basis of buyers' independently made and observable withholding decisions and sellers' pricing responses, we are able to infer repeated-game strategies for individual buyers and sellers.

From our application, we report active, repeated-game, demand withholding strategies for three-quarters of the buyers in the experiments, while inferring the passive price-taking strategy for the remaining one-quarter of the buyers. The inferred withholding strategies are diverse in their degrees of complexity and in the variables upon which withholding is conditioned. In comparing these diverse strategies, certain regularities emerge. More complex strategies seem to lead to lower seller prices and lower seller profits, apparently at no overall cost the buyers. Equally interesting, we find evidence for the relative effectiveness of certain strategy characteristics compared to others. For instance, withholding that decreases with time (early withholding) is more successful in bringing down prices than withholding that increases over time or withholding triggered by a price threshold. Higher buyer concentration induces the more successful strategy.

The results from our procedure are robust, conservative and plausible. They are robust because they hold for a wide range of priors. They are conservative because our procedure imposed a triple penalty for strategy complexity. They are plausible because they are not dominated by the priors; in fact, we observed much heterogeneity across subjects' inferred strategies. This work serves as a non-invasive and complementary approach to better understanding strategic responses to market conditions. Nothing in our technique interferes with the subjects' decision-making processes and the results are both intuitive and interpretable.

References

[1] Breiman, L., H. Friedman, A. Olshen, and C.J. Stone (1984) *Classification and Regression Trees*, Wadsworth International, Belmont, CA.

[2] Camerer, C. (2003) *Behavioral Game Theory: Experiments on Strategic Interaction*, Princeton University Press, Princeton.

[3] Camerer, C., T.-H. Ho, and J.K. Chong (2001) "Behavioral Game Theory: Thinking, Learning, and Teaching", forthcoming in S. Huck editor, *Essays in Honour of Werner Guth.*

[4] Chipman, H. and E. George and R. McCulloch (1998) "Bayesian CART Model Search", *Journal of the American Statistical Association*, 93:443, 935-960.

[5] Costa-Gomes, M., V. Crawford, and B. Broseta (2001) "Cognition and Behavior in Normal-form Games: An Experimental Study", *Econometrica*, 69:5, 1193-1235.

[6] Davis, D. and C. Holt (1993) *Experimental Economics*, Princeton University Press, Princeton, NJ.

[7] Denison, D. and B. Mallick and A. Smith (1998) "A Bayesian CART Algorithm", *Biometrika*, 85:2, 363-377.

[8] Devroye, L., L. Gyorfi, and G. Lugosi (1996) *A Probabilistic Theory of Pattern Recognition*, Springer-Verlag, NY.

[9] Doornik, J. (2001) *Ox: An Object-Oriented Matrix Language*, 4th edition, Timberlake Consultants Press, London.

[10] Duffy, J. and J. Engle-Warnick (2002) “Using Symbolic Regression to Infer Strategies in Experimental Data”, in *Evolutionary Computation in Economics and Finance*, S.H. Chen editor, Springer-Verlag, NY.

[11] El-Gamal, M. and D. Grether (1995) “Are People Bayesian? Uncovering Behavioral Strategies”, *Journal of the American Statistical Association*, 90:432, 1137-1145.

[12] Engle-Warnick, J. (2003) “Inferring Strategies from Observed Actions: A Nonparametric, Binary Tree Classification Approach”, *Journal of Economic Dynamics and Control*, 27, 2151-2170.

[13] Engle-Warnick, J. and B. Ruffle (2005) “Buyer Concentration as a Source of Countervailing Power: Evidence from Experimental Posted-Offer Markets”, unpublished manuscript.

[14] Geweke, J. and M. Keane (1999) “Bayesian Inference for Dynamic Discrete Choice Models Without the Need for Dynamic Programming”, in Mariano, Schuermann and Weeks, (eds), *Simulation Based Inference and Econometrics: Methods and Applications*, Cambridge University Press, Cambridge.

[15] Geweke, J. and M. Keane (2001) “Computationally Intensive Methods for Integration in Econometrics”, in J. Heckman and E. Leamer, (eds), *Handbook of Econometrics*, Vol. 5, North Holland, 3463-3568.

[16] Hendry, D. and H.-M. Krolzig (2001) *Automatic Econometric Model Selection Using PcGets*, Timberlake Consultants Ltd., London, UK.

[17] Houser, D., M. Keane and K. McCabe (2004) “Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm”, *Econometrica*, 72:3, 781-822.

[18] Holt, C. (1989) “The Exercise of Market Power in Laboratory Experiments”, *Journal of Law and Economics*, 32, S107-S131.

[19] Inderst, Roman and Christian Wey (2003) “Bargaining, Mergers, and Technology Choice in Bilaterally Oligopolistic Industries,” *RAND Journal of Economics*, 34, 1-19.

[20] Kagel, J. and A. Roth (1995) *The Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ.

[21] Kreps, D., P. Milgrom, J. Roberts, and R. Wilson (1982) “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma”, *Journal of Economic Theory*, 27, 245-252.

[22] Johnson, E., C. Camerer, S. Sen, and T. Rymon (2001) “Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining”, *Journal of Economic Theory*, 104:1, 16-47.

[23] Manksi, C. (2002) “Identification of Decision Rules in Experiments on Simple Games of Proposal and Response”, *European Economic Review*, 46, 880-891.

[24] McKelvey, R. and T. Palfrey (1995) “Quantal Response Equilibrium for Normal Form Games’, *Games and Economic Behavior*, 10, 6-38.

[25] Ruffle, B. (forthcoming) "Buyer Countervailing Power: A Survey of the Theory and Experimental Evidence," in *Experiments and Antitrust Policy*, J. Hinloopen and H.-T. Normann (eds.), Cambridge, U.K.: Cambridge University Press.

[26] Ruffle, B. (2000) "Some Factors Affecting Demand Withholding in Posted-Offer Markets", *Economic Theory*, 16:3, 529-544.

[27] Selten, R., M. Mitzkewitz and G. R. Ulrich (1997) "Duopoly Strategies Programmed by Experienced Traders", *Econometrica*, 65, 517-555.

[28] Selten, R. and R. Stoecker (1986) "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames", *Journal of Economic Behavior and Organization*, 38, 217-236.

[29] Smith, V. L. (1982) "Experimental Economics: Induced Value Theory", *American Economic Review Papers and Proceedings*, 66:2, 274-279.

[30] Snyder, C. (forthcoming) "Countervailing Power," in S. Durlauf and L. Blume (eds.), *The New Palgrave Dictionary of Economics*, 2nd edition, London: Palgrave Macmillan.

[31] Stahl, D. and P. Wilson (1995) "On Players' Models of Other Players: Theory and Experimental Evidence", *Games and Economic Behavior*, 10, 218-254.

Table 1: Strategy Characteristics and Game Variable Summary for Each Buyer in the Two-Buyer Informed Treatment

Session Information			Mean Profit Per Round		Number of Decisions			Modal Strategy Characteristics		Probability of Strategy Complexity by Number of Relational Nodes					Probability of Strategy Containing Specific Relational Nodes			
Session	**Buyer**	**Price**	**Seller**	**Buyer**	**0**	**1**	**+**	**Size**	**Error**	**0**	**1**	**2**	**3**	**4**	**Time -**	**Time +**	**Price -**	**Price +**
1	1	-0.01	4.381	0.735	30	0	0	0	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	2			0.707	17	4	9	0	0.433	0.009	0.131	0.548	0.218	0.094	0.946	0.158	0.340	0.004
2	1	0.05	2.285	1.185	29	1	0	0	0.033	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	2			0.033	3	5	22	2	0.100	0.000	0.025	0.319	0.306	0.351	0.719	0.238	0.920	0.341
3	1	-0.25	2.443	0.990	14	6	10	1	0.367	0.003	0.091	0.526	0.248	0.132	0.962	0.053	0.145	0.351
	2			1.287	14	12	4	2	0.233	0.001	0.063	0.429	0.303	0.204	0.990	0.050	0.101	0.269
4	1	-0.02	4.088	0.467	22	1	7	2	0.167	0.009	0.113	0.493	0.292	0.094	0.899	0.220	0.270	0.189
	2			0.578	30	0	0	0	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
5	1	0.03	4.093	0.689	29	1	0	0	0.033	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	2			0.516	15	4	11	1	0.333	0.003	0.100	0.598	0.194	0.104	0.891	0.045	0.215	0.598
6	1	-0.22	2.956	1.545	26	4	0	2	0.033	0.001	0.013	0.207	0.386	0.394	0.990	0.000	0.214	0.922
	2			0.858	11	5	14	1	0.367	0.003	0.059	0.333	0.309	0.297	0.935	0.450	0.115	0.474
7	1	-0.3	2.236	1.235	10	9	11	3	0.161	0.002	0.019	0.195	0.292	0.493	0.968	0.784	0.027	0.874
	2			0.937	11	2	17	1	0.200	0.000	0.035	0.439	0.291	0.234	0.995	0.358	0.256	0.476

Each row reports the observed game variables (median price from the last five rounds, "Price", mean seller and buyer per round profit) and the strategy characteristics for a different buyer according to session. The "Number of Decisions" column indicates the breakdown of a buyer's observed withholding decisions. "Size" indicates the number of relational nodes in the inferred modal strategy, while Error indicates the fraction of decisions this strategy does not categorize correctly. The next set of columns reports the posterior probability that the buyer's strategy contains 0, 1, 2, 3, and 4 relational nodes. These sum to one. The last set of columns report the posterior probability that the strategy contains the different conditioning variables. These may sum from zero (when the posterior probablity of a zero node is one) to four (when only rules that contain at least one of each of the four nodes have non-zero posterior probability).

Table 2: Strategy Characteristics and Game Variable Summary for Each Buyer in the Two-Buyer Uninformed Treatment

Session Information			Mean Profit Per Round		Number of Decisions			Modal Strategy Characteristics		Probability of Strategy Complexity by Number of Relational Nodes					Probability of Strategy Containing Specific Relational Nodes			
Session	**Buyer**	**Price**	**Seller**	**Buyer**	**0**	**1**	**+**	**Size**	**Error**	**0**	**1**	**2**	**3**	**4**	**Time -**	**Time +**	**Price -**	**Price +**
23	1	-0.22	3.370	0.673	14	4	12	2	0.167	0.000	0.002	0.267	0.342	0.389	0.593	0.953	0.175	0.724
	2			0.914	24	2	4	2	0.100	0.003	0.043	0.357	0.360	0.237	0.839	0.373	0.585	0.259
24	1	-0.03	4.296	0.892	25	5	0	2	0.100	0.009	0.107	0.433	0.274	0.178	0.981	0.173	0.207	0.005
	2			0.753	19	1	10	1	0.267	0.002	0.080	0.445	0.293	0.179	0.974	0.181	0.035	0.265
25	1	-0.2	3.172	0.823	7	13	10	2	0.333	0.008	0.122	0.545	0.245	0.080	0.861	0.686	0.029	0.185
	2			1.402	28	1	1	0	0.067	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
26	1	-0.15	2.730	0.620	17	8	5	1	0.367	0.005	0.096	0.510	0.261	0.128	0.700	0.766	0.329	0.176
	2			0.508	10	9	11	4	0.167	0.005	0.071	0.398	0.263	0.263	0.874	0.425	0.031	0.577
27	1	-0.05	2.832	0.475	12	5	13	2	0.233	0.000	0.018	0.352	0.302	0.328	0.991	0.556	0.388	0.276
	2			0.633	16	6	8	3	0.100	0.000	0.001	0.026	0.439	0.534	1.000	0.662	0.105	0.181
28	1	0.02	4.598	0.736	17	12	1	1	0.233	0.000	0.115	0.558	0.248	0.080	0.969	0.085	0.024	0.223
	2			0.746	30	0	0	0	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
29	1	-0.1	3.970	1.100	30	0	0	0	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	2			1.035	26	1	3	0	0.133	0.016	0.183	0.504	0.223	0.074	0.984	0.000	0.000	0.000
30	1	-0.07	1.775	0.106	5	1	24	0	0.200	0.008	0.129	0.509	0.255	0.099	0.895	0.134	0.487	0.022
	2			0.426	13	2	15	2	0.133	0.000	0.013	0.330	0.411	0.245	0.626	0.458	0.480	0.915

Each row reports the observed game variables and the strategy characteristics (see the caption of Table 1 for an explanation) for a different buyer in the two-buyer uninformed treatment.

Table 3: Strategy Characteristics and Game Variable Summary for Each Buyer in the Four-Buyer Informed Treatment

Session Information			Mean Profit Per Round		Number of Decisions			Modal Strategy Characteristics		Probability of Strategy Complexity by Number of Relational Nodes					Probability of Strategy Containing Specific Relational Nodes			
Session	**Buyer**	**Price**	**Seller**	**Buyer**	**0**	**1**	**+**	**Size**	**Error**	**0**	**1**	**2**	**3**	**4**	**Time -**	**Time +**	**Price -**	**Price +**
8	1	0.01	4.094	0.679	15	5	0	2	0.050	0.002	0.039	0.451	0.384	0.125	0.350	0.951	0.896	0.046
	2			0.438	5	8	7	1	0.400	0.023	0.253	0.630	0.087	0.007	0.568	0.457	0.383	0.223
	3			0.560	9	8	3	0	0.550	0.020	0.224	0.650	0.099	0.007	0.674	0.238	0.425	0.270
	4			0.635	19	0	1	0	0.050	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
9	1	0.15	5.434	0.291	5	3	12	2	0.050	0.000	0.008	0.729	0.255	0.008	0.276	0.990	0.008	0.976
	2			-0.483	18	2	0	0	0.100	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	3			0.318	12	3	5	2	0.150	0.001	0.018	0.401	0.258	0.322	0.422	0.928	0.609	0.536
	4			0.327	20	0	0	0	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
10	1	-0.02	3.677	0.489	6	6	18	2	0.167	0.000	0.025	0.397	0.375	0.203	0.992	0.150	0.288	0.396
	2			0.725	15	11	4	3	0.167	0.001	0.015	0.263	0.302	0.420	0.361	0.571	0.449	0.923
	3			0.781	30	0	0	0	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	4			0.694	14	8	8	2	0.233	0.000	0.024	0.297	0.154	0.525	0.942	0.000	0.245	0.871
11	1	0	3.157	0.562	14	6	10	1	0.433	0.009	0.152	0.567	0.205	0.067	0.979	0.106	0.105	0.092
	2			0.760	9	19	2	2	0.200	0.002	0.071	0.562	0.271	0.095	0.925	0.698	0.267	0.012
	3			0.738	12	14	4	2	0.300	0.005	0.086	0.548	0.271	0.091	0.913	0.831	0.104	0.039
	4			0.580	11	6	13	1	0.367	0.006	0.105	0.568	0.216	0.105	0.947	0.220	0.077	0.453
12	1	0.04	4.567	0.615	13	4	3	2	0.150	0.003	0.053	0.675	0.236	0.033	0.982	0.000	0.498	0.011
	2			0.574	6	10	4	1	0.400	0.030	0.270	0.605	0.091	0.004	0.552	0.770	0.139	0.121
	3			0.736	18	2	0	0	0.100	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	4			0.738	19	1	0	0	0.050	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
13	1	0.04	4.213	0.328	10	3	7	3	0.050	0.001	0.009	0.153	0.829	0.008	0.042	0.991	0.002	0.950
	2			0.321	4	5	11	2	0.150	0.003	0.102	0.680	0.201	0.015	0.825	0.706	0.016	0.315
	3			0.316	15	2	3	1	0.150	0.012	0.261	0.598	0.124	0.004	0.725	0.834	0.150	0.000
	4			0.334	9	4	7	2	0.100	0.001	0.022	0.813	0.161	0.003	0.095	0.995	0.013	0.873
14	1	0.15	4.544	0.224	15	0	15	2	0.033	0.000	0.000	0.307	0.404	0.289	0.829	1.000	0.099	0.412
	2			0.238	30	0	0	0	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	3			0.207	11	5	14	2	0.133	0.000	0.037	0.418	0.346	0.199	0.443	1.000	0.080	0.625
	4			0.234	24	4	2	1	0.133	0.002	0.086	0.513	0.296	0.104	0.903	0.707	0.306	0.002
15	1	0.01	4.008	0.528	15	10	5	3	0.033	0.000	0.000	0.001	0.821	0.178	0.132	0.999	0.999	0.988
	2			0.518	20	6	4	1	0.233	0.002	0.079	0.504	0.281	0.134	0.708	0.923	0.312	0.032
	3			0.424	10	10	10	3	0.200	0.003	0.033	0.303	0.374	0.287	0.699	0.980	0.270	0.075
	4			0.495	8	12	10	2	0.300	0.006	0.100	0.517	0.243	0.135	0.769	0.344	0.022	0.667

Each row reports the observed game variables and the strategy characteristics (see the caption of Table 1 for an explanation) for a different buyer in the four-buyer informed treatment.

Table 4: Strategy Characteristics and Game Variable Summary for Each Buyer in the Four-Buyer Uninformed Treatment

Session Information			Mean Profit Per Round		Number of Decisions			Modal Strategy Characteristics		Probability of Strategy Complexity by Number of Relational Nodes					Probability of Strategy Containing Specific Relational Nodes			
Session	**Buyer**	**Price**	**Seller**	**Buyer**	**0**	**1**	**+**	**Size**	**Error**	**0**	**1**	**2**	**3**	**4**	**Time -**	**Time +**	**Price -**	**Price +**
16	1	-0.01	3.194	0.499	19	8	3	2	0.200	0.001	0.019	0.418	0.326	0.235	0.718	0.379	0.548	0.501
	2			0.414	19	5	6	2	0.200	0.001	0.018	0.329	0.344	0.307	0.964	0.296	0.342	0.178
	3			0.554	28	2	0	0	0.067	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	4			0.481	19	7	4	2	0.267	0.002	0.040	0.356	0.275	0.326	0.857	0.317	0.482	0.317
17	1	0.15	4.247	0.240	27	2	1	0	0.100	0.013	0.156	0.471	0.259	0.101	0.975	0.000	0.173	0.000
	2			0.235	28	2	0	0	0.067	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	3			0.195	21	2	7	2	0.167	0.002	0.029	0.425	0.375	0.170	0.995	0.498	0.150	0.000
	4			0.207	18	7	5	0	0.400	0.008	0.098	0.482	0.289	0.124	0.984	0.325	0.146	0.000
18	1	-0.01	4.203	0.767	21	7	2	2	0.167	0.004	0.053	0.395	0.330	0.219	0.988	0.409	0.203	0.006
	2			0.593	3	14	13	1	0.200	0.000	0.080	0.524	0.265	0.131	1.000	0.000	0.195	0.092
	3			0.800	23	6	1	1	0.100	0.000	0.068	0.429	0.339	0.163	0.821	0.997	0.274	0.000
	4			0.819	29	1	0	0	0.033	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
19	1	0.02	3.299	0.498	3	19	8	1	0.300	0.006	0.103	0.520	0.205	0.166	0.902	0.288	0.448	0.126
	2			0.582	18	3	9	1	0.333	0.007	0.123	0.551	0.222	0.096	0.951	0.057	0.274	0.174
	3			0.422	9	4	17	2	0.233	0.003	0.073	0.534	0.260	0.129	0.882	0.157	0.064	0.594
	4			0.450	3	10	17	3	0.161	0.001	0.041	0.375	0.302	0.280	0.955	0.709	0.345	0.238
20	1	-0.03	4.468	0.542	11	6	13	1	0.400	0.003	0.083	0.480	0.294	0.140	0.578	0.935	0.263	0.136
	2			0.732	29	1	0	0	0.033	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	3			0.725	25	5	0	2	0.067	0.001	0.018	0.296	0.384	0.302	0.852	0.554	0.595	0.299
	4			0.701	10	19	1	1	0.133	0.000	0.081	0.557	0.276	0.086	0.797	0.076	0.154	0.994
21	1	0	2.667	0.400	21	0	9	1	0.033	0.000	0.046	0.306	0.375	0.274	1.000	0.000	0.348	0.001
	2			0.284	15	3	12	1	0.400	0.006	0.100	0.542	0.240	0.113	0.925	0.179	0.273	0.224
	3			0.294	6	6	18	3	0.200	0.002	0.032	0.256	0.382	0.328	0.983	0.454	0.113	0.502
	4			0.030	0	1	29	0	0.033	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
22	1	-0.03	4.139	0.440	14	6	10	1	0.200	0.000	0.031	0.246	0.321	0.401	0.579	0.998	0.592	0.000
	2			0.644	13	14	3	3	0.100	0.000	0.001	0.078	0.491	0.430	0.450	0.946	0.561	0.788
	3			0.703	28	2	0	0	0.067	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	4			0.702	28	2	0	0	0.067	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Each row reports the observed game variables and the strategy characteristics (see the caption of Table 1 for an explanation) for a different buyer in the four-buyer uninformed treatment.

Table 5: Strategy Characteristics and Market Outcome Correlations

	All Sessions		
	Price	Seller Profit	Buyer Profit
Modal Size	**-0.333**	**-0.420**	-0.027
	0.071	**0.009**	0.315
Expected Size	**-0.437**	**-0.550**	-0.043
	0.013	**0.002**	0.458
Error	-0.327	-0.356	-0.041
	0.198	0.109	0.597
P(Zero Nodes)	**0.380**	**0.488**	0.080
	0.023	**0.005**	0.204
P(Time-)	**-0.484**	**-0.579**	0.044
	0.004	**0.003**	0.746
P(Time+)	-0.006	0.023	-0.190
	0.960	0.966	0.021
P(Price-)	0.038	**-0.356**	**-0.141**
	0.693	**0.067**	**0.078**
P(Price+)	-0.415	**-0.352**	0.038
	0.157	**0.063**	0.479

Spearman rank correlations between the posterior probability estimates of various strategy characteristics (rows) and market outcomes aggregated across all sessions (columns). Each cellreports the correlation coefficient (top entry) and the p-value (bottom entry). Significant results appear in bold.

Table 6: Regressions of Market Outcomes on Strategy Characteristics

	Independent Variables			
Dependent Variable	Constant	2-Buyer	Time -	Price +
Price	0.149	-0.113	-0.130	-0.170
	(0.043)	(0.03)	(0.066)	(0.077)
Seller Profit	5.630	-0.610	-1.810	-2.740
	(0.363)	(0.228)	(0.483)	(0.854)
Buyer Profit	0.480	0.306		
	(0.063)	(0.089)		

OLS regressions of median session price from the last five rounds of the game ("Price"), average seller profit, and average buyer profit on a dummy variable for two-buyer sessions ("2-Buyer"), and the posterior probability estimates of different strategy characteristics. We used the software package PcGets to select a final regression specification, thereby eliminating subjectivity in the choice of a model.

Figure 1a: Two-Buyer Treatment Parameters

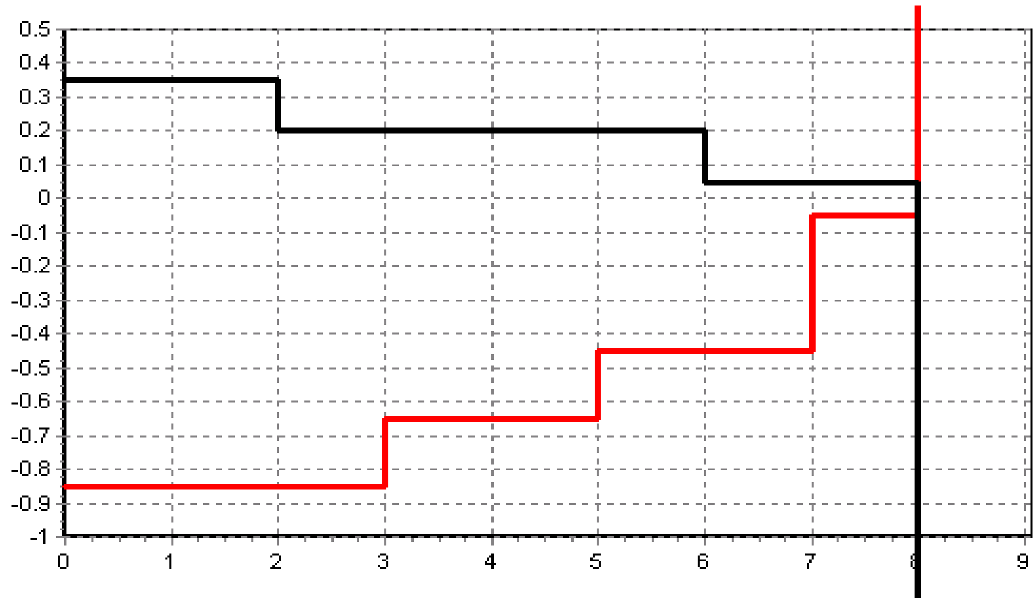

The monopolist's marginal cost and the buyers' demand curve in the two-buyer sessions. All costs and valuations are expressed as deviations from the competitive price range's midpoint, which is normalized to 0. The competitive price range consists of the interval –0.05 and 0.05. Each of the two symmetric buyers possesses four units of demand, the first unit of which is valued at +0.35, the second and third units have values of +0.20 each, and the fourth unit has a value of +0.05.

Figure 1b: Four-Buyer Treatment Parameters

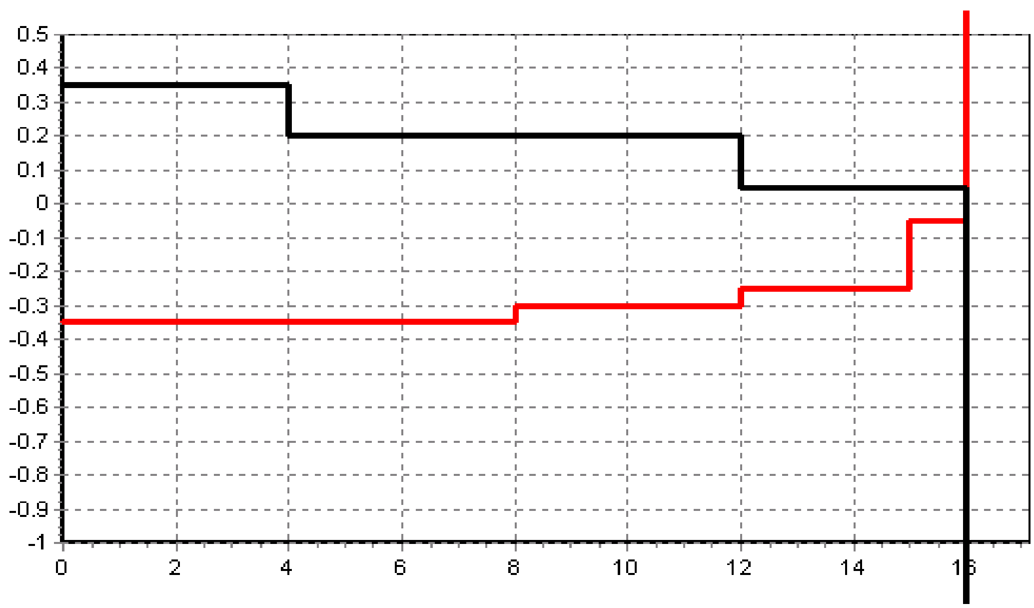

The monopolist's marginal cost and the buyers' demand curve in the four-buyer sessions. The competitive price range and individual buyers' demand curves are identical to the two-buyer treatment. All costs and valuations are expressed as deviations from the midpoint of the competitive price range.

Figure 2: Time Series Plot of Median Price in Informed Sessions

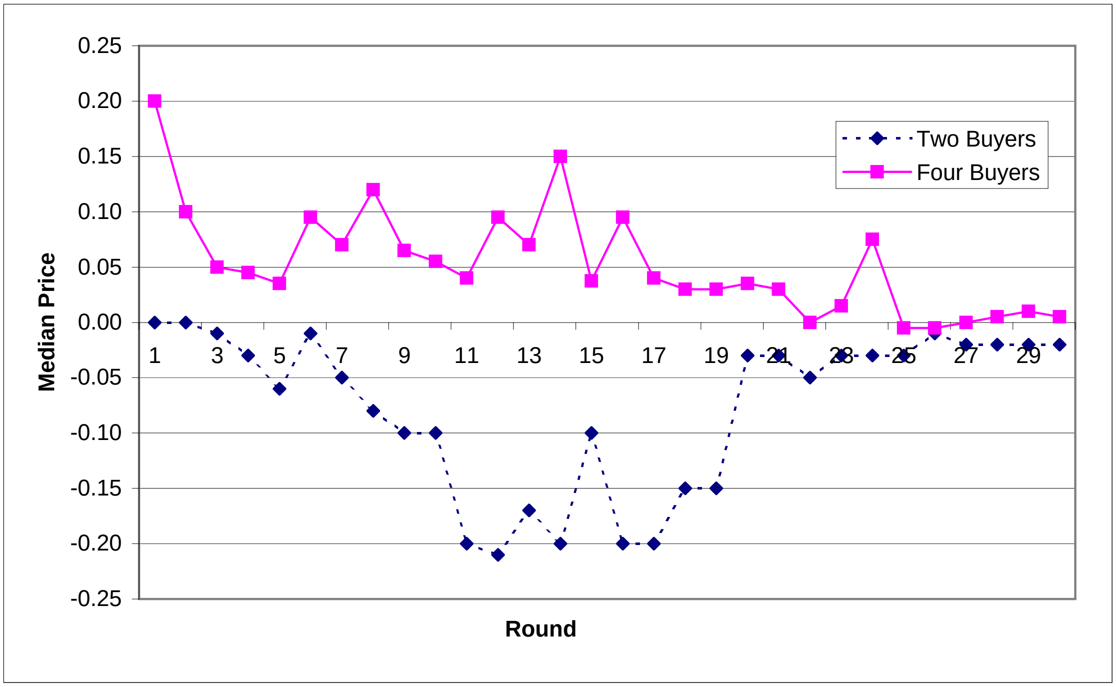

Time series plot of the median posted price for the two-buyer informed and the four-buyer informed treatments. The median price is computed per round from the pooled prices of all of the sessions in each treatment.

Figure 3: Time Series Plot of Median Price in Uninformed Sessions

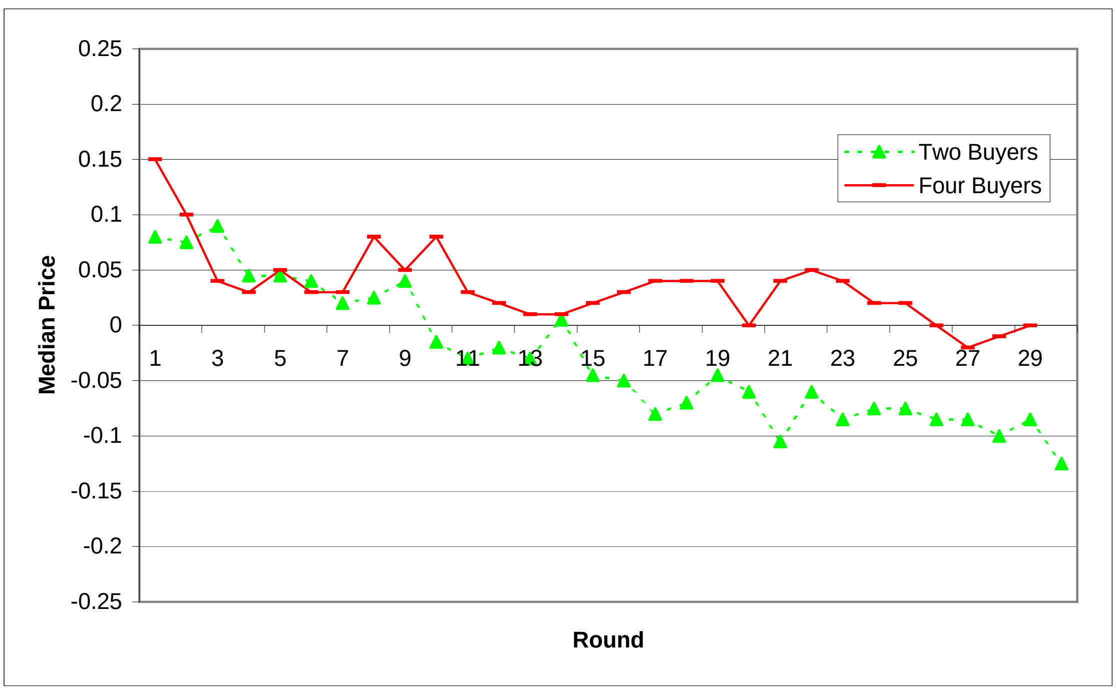

Time series plot of the median posted price for the two-buyer uninformed and the four-buyer uninformed treatments. The median price is computed per round from the pooled prices of all of the sessions in each treatment.

Figure 4: An Example Strategy

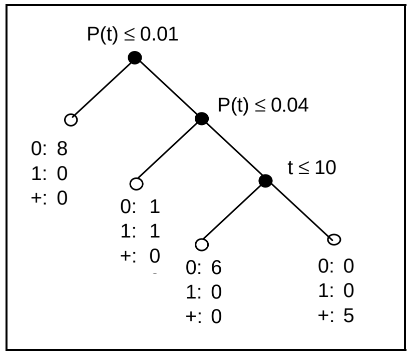

An actual inferred buyer withholding strategy containing three relational nodes (two prices thresholds and a time node). The behavioral interpretation of this strategy suggests the buyer never withheld demand when the price was low enough (below 0.01), increased withholding to one unit for intermediate prices (between 0.01 and 0.04), and withheld demand intensely after round 10 for high prices (above 0.04).

Figure 5: Constructing a Strategy

Assignment 0: the set of possible conditioning variables: {P(t), t}

Assignment 1: k, the number of relational nodes
Choices: k ∈ {0,1,2,3,4}

Example choice: k = 2

Assignment 2: Select k = 2 variables from Variable Set C = {P1(t), P2(t), t1, t2}
Choices: {P1(t), P2(t)}, {P1(t), t1}, P1(t), t2} {P2(t), t1}, {P2(t), t2}, {t1, t2}

Example choice: {P1(t), t1}

Assignment 3: Node Configuration Choices:

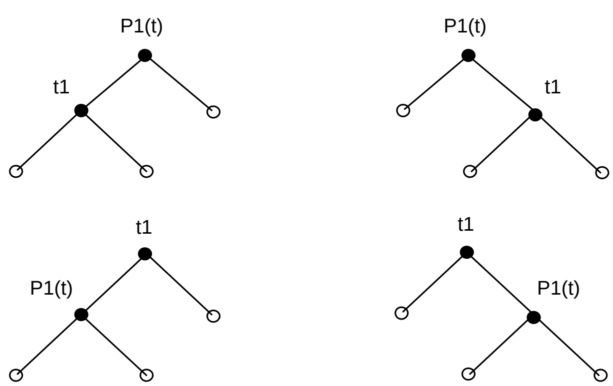

An illustration of the assignment problems involved in constructing a single strategy. We form priors over each possible assignment. We construct each possible strategy up to the number of relational nodes k=4. We compute the likelihood of each strategy, and then estimate the posterior probability of each strategy. From this inbformation, we report estimates of the probability that a strategy is of a certain complexity, or contains specific elements.

Figure 6: Average Complexity of the Modal and Expected Strategies

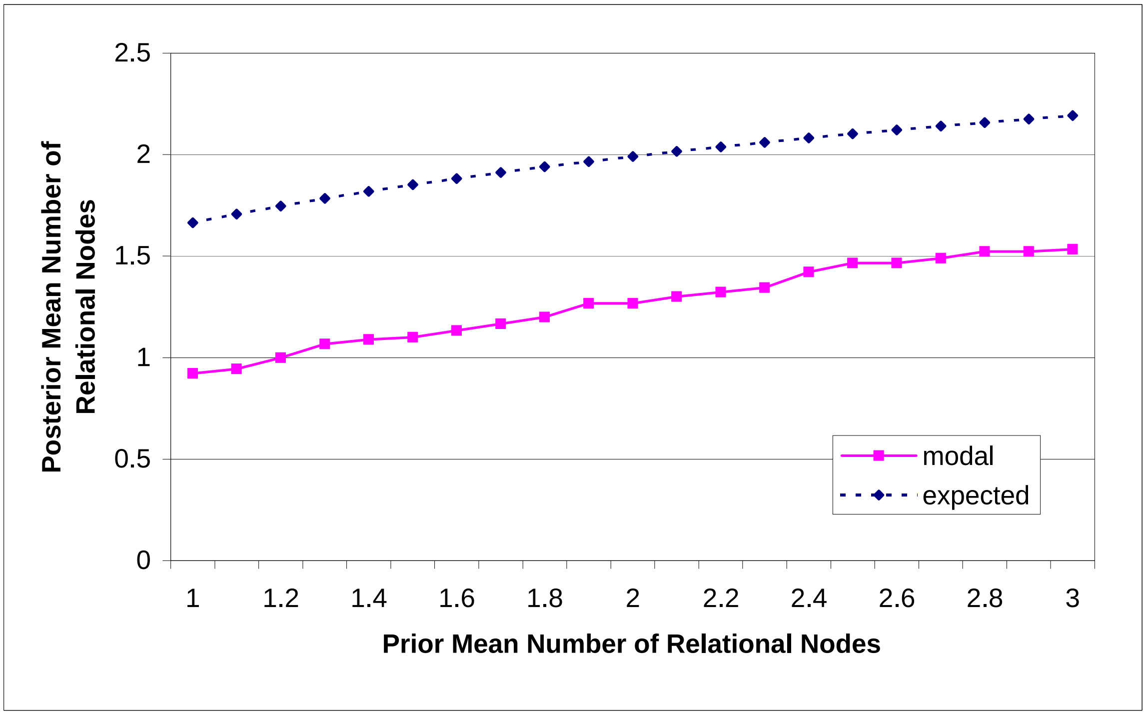

The average number of relational nodes in the posterior modal and posterior expected strategies (pooled across all subjects in all treatments) when the prior mean number of nodes varies between 1 and 3.

Figure 7a: Average Strategy Complexity by Informational Treatment

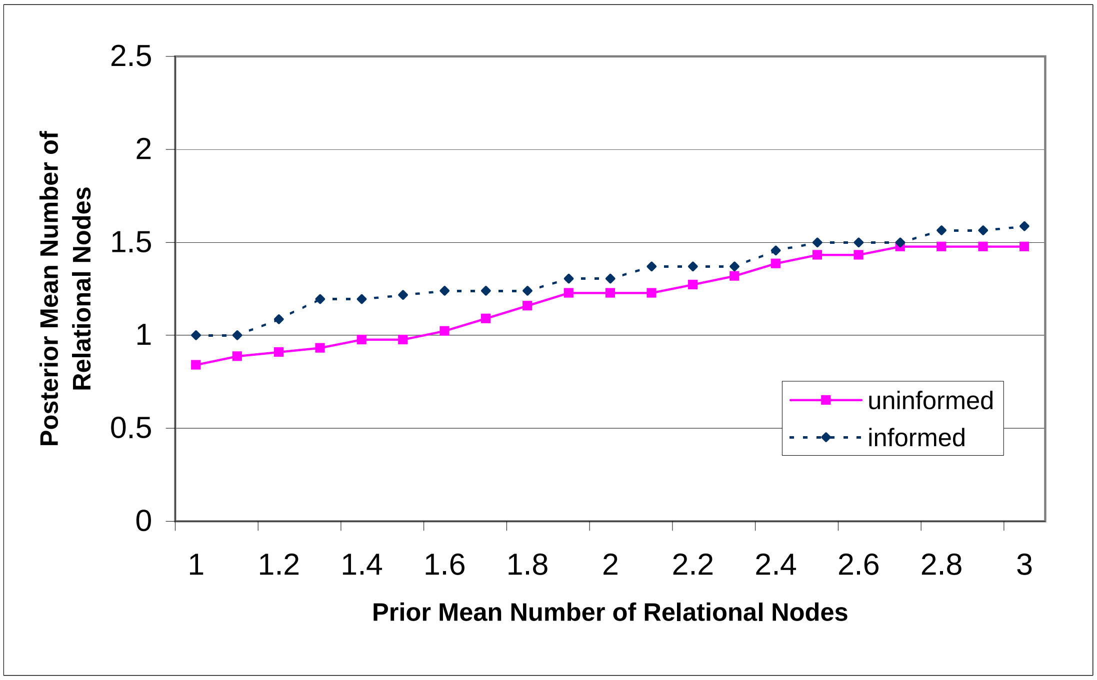

The average number of relational nodes in the posterior modal strategies for the informed (two-buyer and four-buyer) and uninformed (two-buyer and four-buyer) treatments, when the prior mean number of nodes varies from 1 to 3.

Figure 7b: Average Strategy Complexity by Number of Buyers

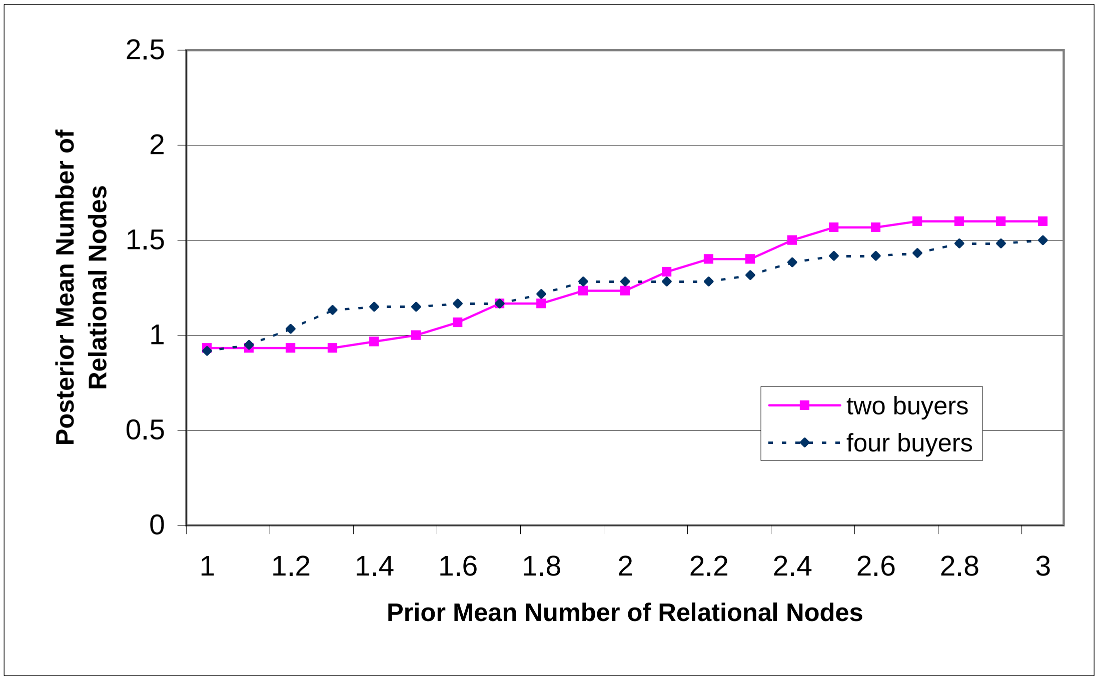

The average number of relational nodes in the posterior modal strategies for the two-buyer (informed and uninformed) and four-buyer (informed and uninformed) treatments, when the prior mean number of nodes varies from 1 to 3.

Figure 8a: Average Probability That Strategies Contain Time- Relational Node

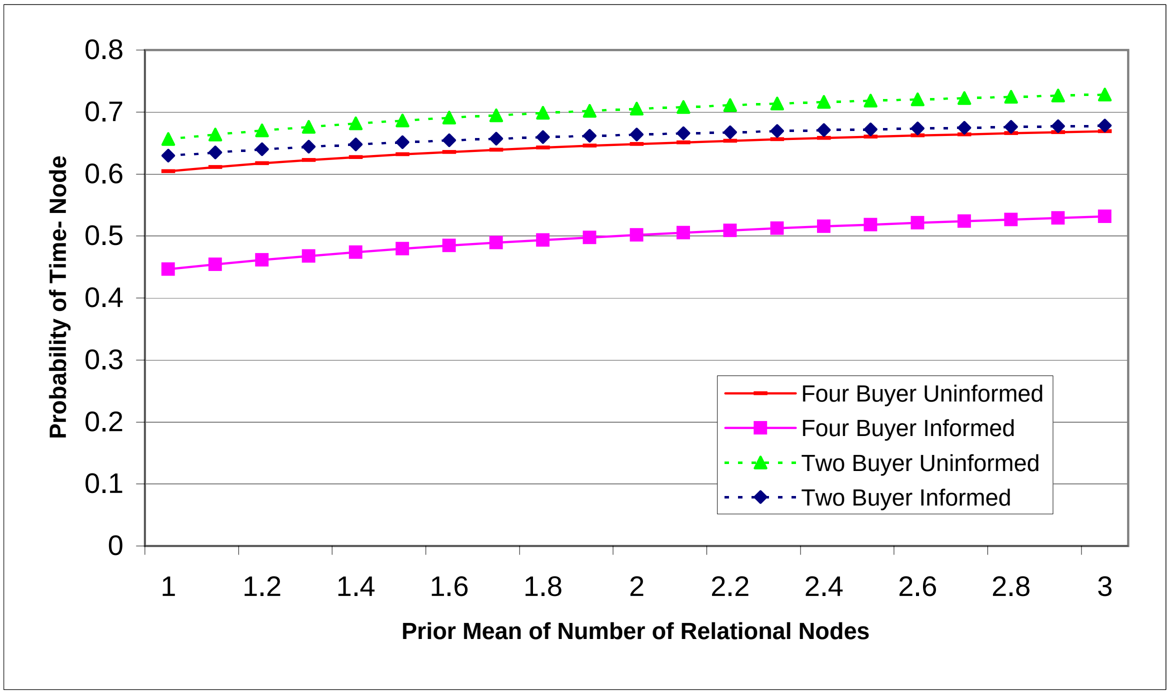

The posterior probability that the buyers' strategies (pooled by treatment) contain the relational node that decreases withholding with time (Time -), when the prior mean number of relational nodes is allowed to vary between 1 and 3.

Figure 8b: Average Probability That Strategies Contain Time+ Relational Node

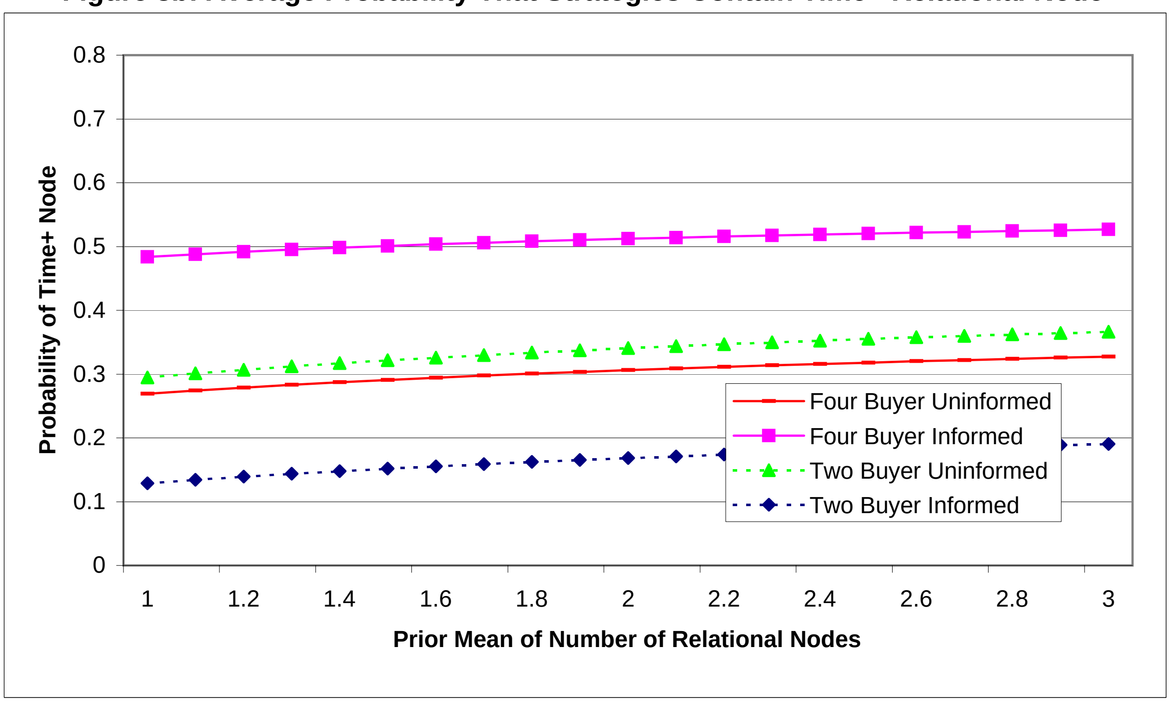

The posterior probability that the buyers' strategies (pooled by treatment) contain the relational node that increases withholding with time (Time +), when the prior mean number of relational nodes is allowed to vary between 1 and 3.

Figure 9a:Average Probability that Strategies Contain Price- Relational Node

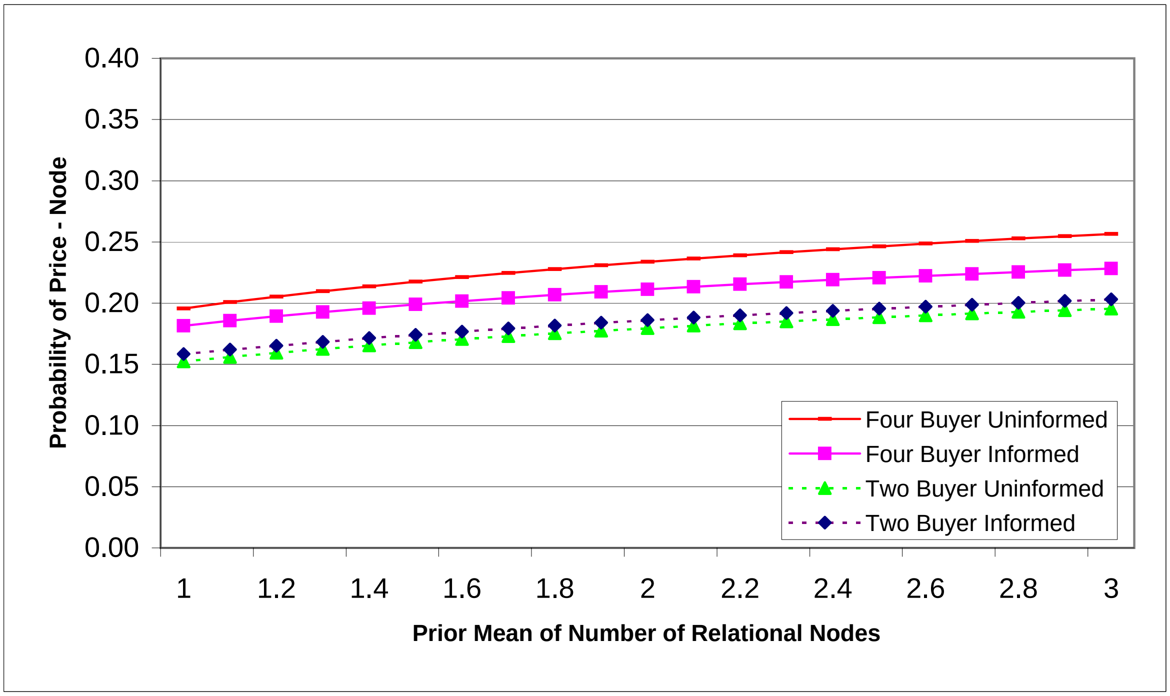

The posterior probability that the buyers' strategies (pooled by treatment) contain the relational node that decreases withholding with price (Price -), when the prior mean number of relational nodes is allowed to vary between 1 and 3.

Figure 9b: Average Probability that Strategies Contain Price+ Relational Node

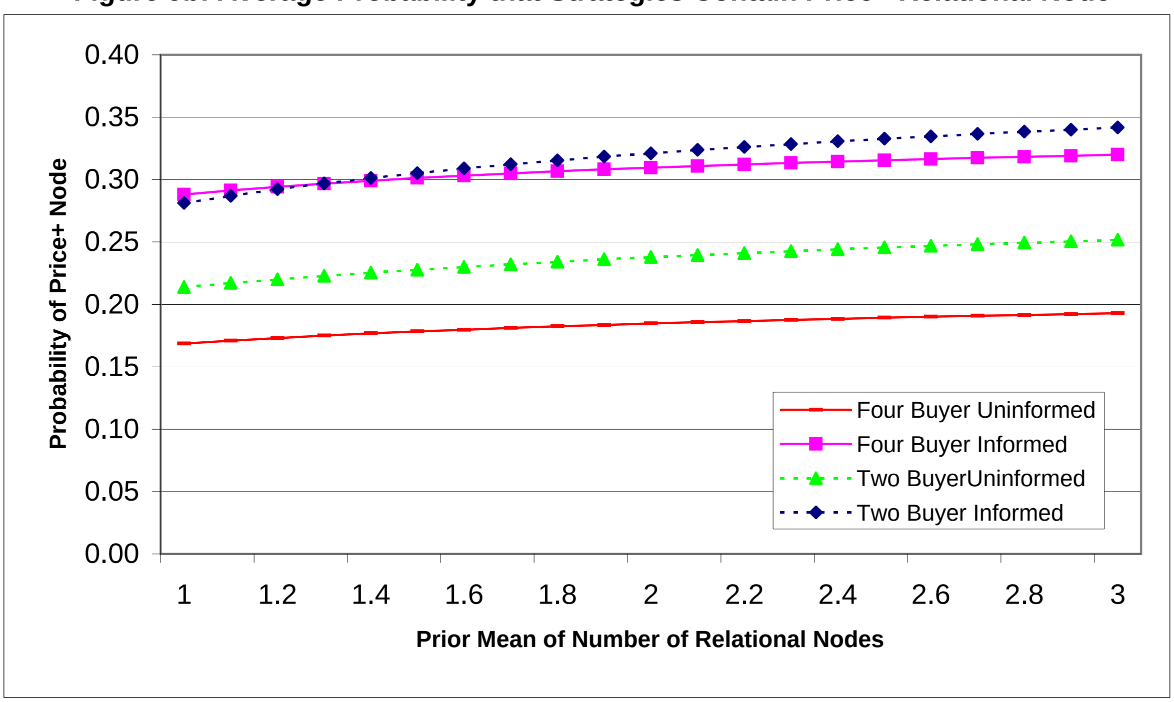

The posterior probability that the buyers' strategies (pooled by treatment) contain the relational node that increases withholding with price (Price +), when the prior mean number of relational nodes is allowed to vary between 1 and 3.

Figure 10: Buyer Strategies in Two-Buyer Informed Session 7

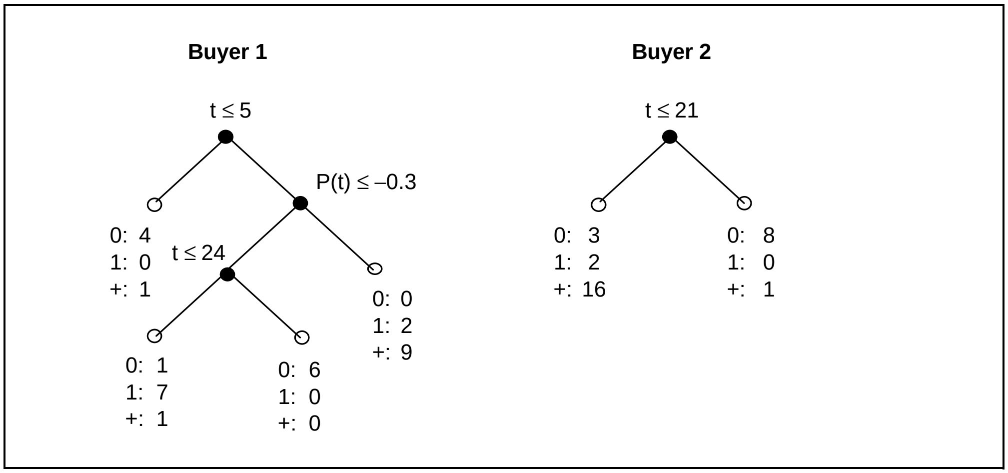

The modal strategies inferred for each buyer in this two-buyer session. Buyer 1's strategy reflects price sensitivity after period 5. Buyer 2 is an intense early withholder.

Figure 11: Time Series of Prices in Two-Buyer Informed Session 7

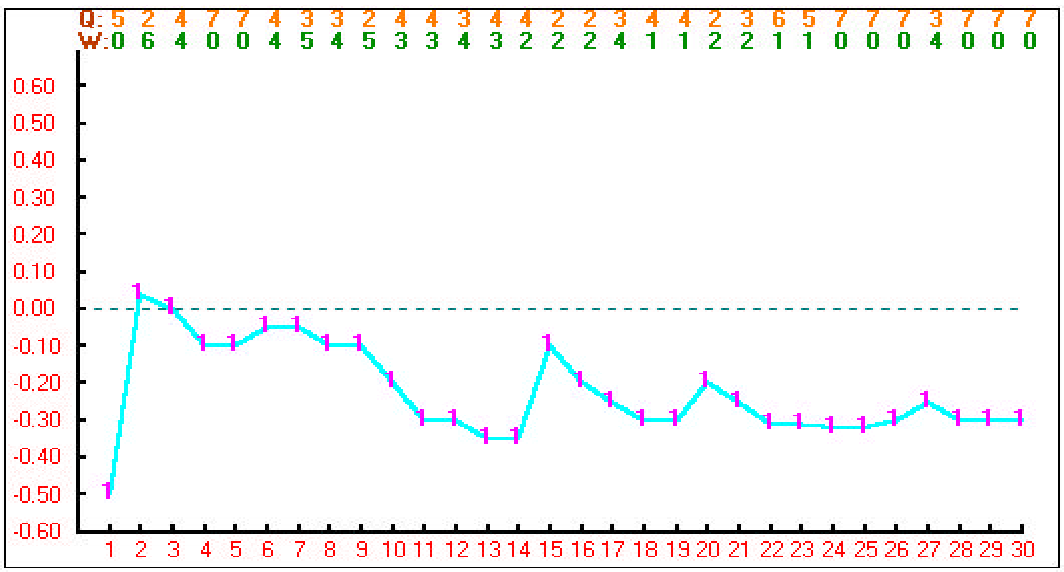

The monopolist's posted prices in each of the 30 rounds. The series "Q and "W" indicate the number of units sold and the number of sales lost to demand withholding by period, respectively.

Figure 12: Buyer Strategies in Four-Buyer Informed Session 10

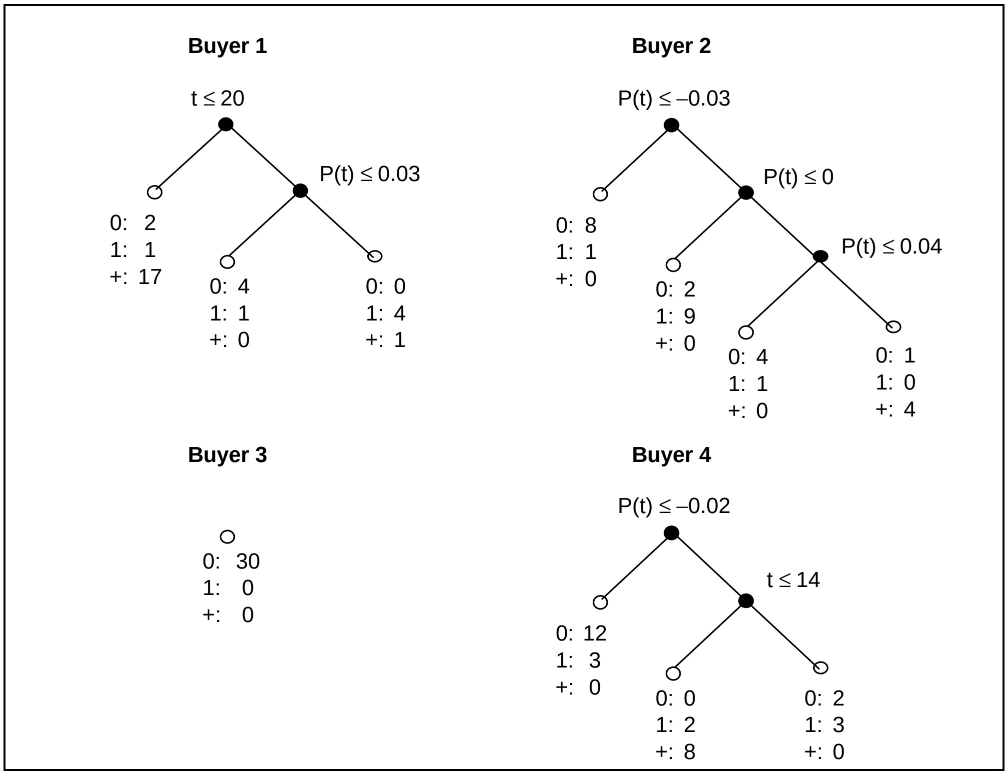

The modal strategies inferred for each buyer in this four-buyer session. For Buyer 3, the degenerate strategy of making all profitable purchases in every period is the modal strategy. Price thresholds appear in the other strategies. Buyers 1 and 4 display early, unconditional demand withholding.

Figure 13: Time Series of Prices in Four-Buyer Informed Session 10

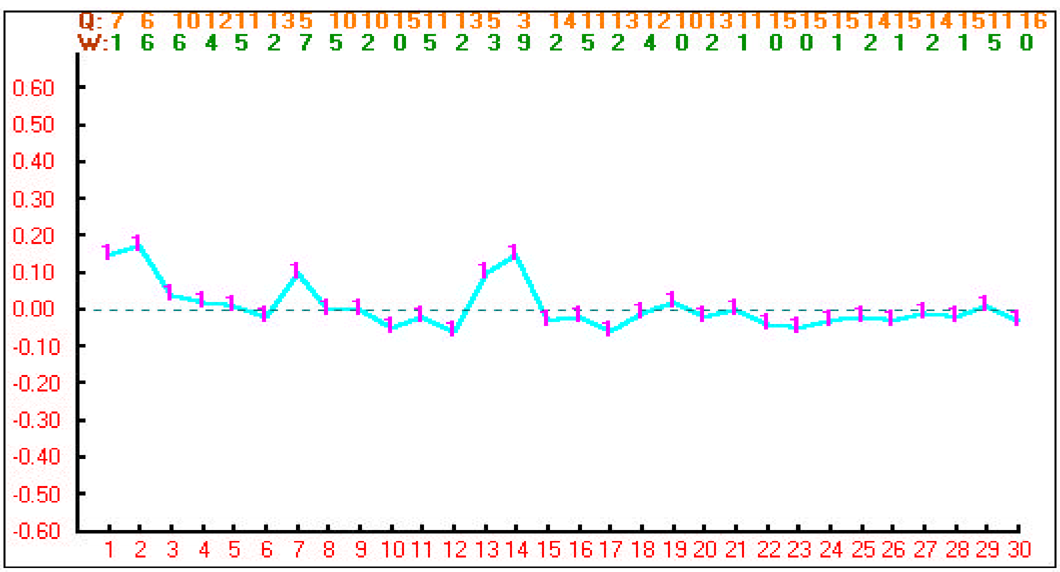

The monopolist's posted prices in each of the 30 rounds. The series "Q and "W" indicate the number of units sold and the number of sales lost to demand withholding by period, respectively.

Figure 14: Seller Strategies

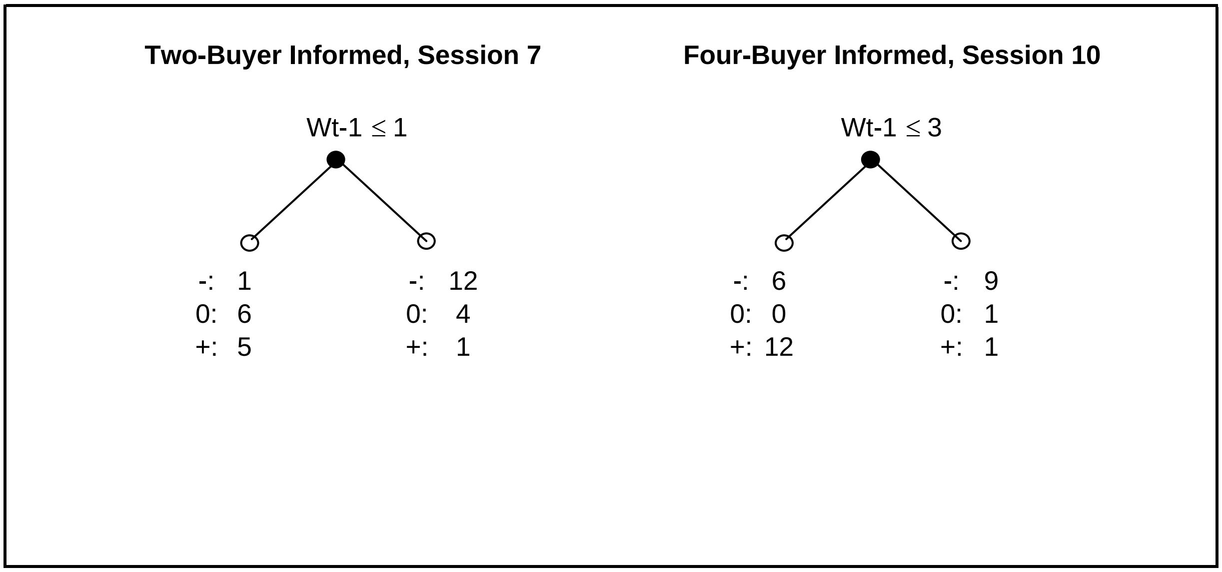

The modal strategies inferred for the sellers in Sessions 7 and 10. Session 7 seller strategy reflects a price reduction whenever more than one unit of demand is withheld. Session 10 seller increases the price when 3 or less units are withheld, and lowers the price otherwise.