

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

Jim Engle-Warnick
Department of Economics
Nuffield College, Oxford University
New Road, Oxford OX1 1NF, UK
jim.ingle-warnick@nuf.ox.ac.uk

Bradley J. Ruffle*
Department of Economics
Ben-Gurion University, P.O.B. 653
Beer Sheva, 84105, Israel
bradley@bgumail.bgu.ac.il

February 2002
Comments Welcome

Abstract

We infer unobserved, repeated-game strategies from the observed actions of buyers in posted-offer market experiments to evaluate their effectiveness against a monopolist. 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 different non-trivial, repeated-game strategies that condition on time, price, and combinations of the two variables. The strategies are generally effective: the use of strategies and their complexity correlate negatively with market prices and monopolist profits. One strategy is particularly effective: the unconditional and intense forgoing of profitable purchases early in the game is more effective as a counteracting response to monopoly power than punishments that trigger when the market price exceeds a threshold. We propose that strategy inference should at least complement existing methods of statistical inference on observed strategic behavior.

keywords: strategy inference, strategic buyer behavior, binary tree classification, experimental economics, posted-offer market

JEL classification nos. C63, C91, D42.

*We thank Erik Eyster, Frank Gerhard, Meg Meyer and seminar participants at the Economic Science Association 2000 meetings in Tucson and Amsterdam, University College London and Oxford University. We also thank Moty Fania for programming the 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 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 predictions 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, we take a very different, though complementary tack. 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 observed play. Our goal is twofold: first, to categorize subject behavior in a way that is intuitive and interpretable as a strategy thereby providing insights into the strategies employed by subjects in this environment; second, to evaluate the success 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, the randomly ordered buyers then proceed one at a time to make the number of purchases that each desires. The take-it-or-leave-it nature of this market institution limits buyer strategic behavior to

¹ The experiments test the role of buyer concentration on pricing in markets, a topic on which there is a growing literature in the form of empirical (e.g. Schumacher, 1991, and Chipty and Snyder, 1999), theoretical (e.g. Snyder, 1996, and Bjornerstedt and Stennek, 2001), and experimental studies (Ruffle, 2000).

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, backward induction would lead a rational buyer to make all profitable purchases; 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 infer unobserved 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. The institutional limitation of strategic buyer behavior to rejecting profitable purchases and the observed variation in game parameters help us to identify strategic behavior. To do so, we apply a Bayesian inference technique from the binary classification tree literature (Breiman, Friedman, Olshen, and Stone, 1999; Chipman, George, and McCulloch, 1998; Denison, Mallick, and Smith, 1998). 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 conditional variables, and the most likely strategy that generated the data. Our experiments thus combine traditional laboratory methods with a new method of data analysis.

One advantage of our approach is that we do not discard the game-theoretic prediction of passive price taking in every round, but rather characterize the precise nature of the deviation from equilibrium play on a subject-by-subject basis. For some buyers (about one-quarter of them), 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 his withholding decisions. Characterizing play with strategies offers insights unavailable in traditional econometric techniques: it places inference into the language of game theory, thus bridging a gap between theory and observed behavior.

We find evidence that subjects use, either implicitly or explicitly, 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. 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.

The direct inference methodology in this paper is complementary to several existing approaches. In the strategy method of Selten et al. (1997), strategy choices are made observable through elicitation; a second approach is to validate the inferred strategies by tracking the manner in which subjects collect and process information (i.e., to collect attentional data as in Costa-Gomez, 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, 2002; 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, 2001; McKelvey and Palfrey, 1992); and of course there are the classic approaches of experimental manipulation and protocol responses. Taken together, all of these complementary approaches provide different methods to advance the understanding of how people play games.

What distinguishes our strategy inference approach from others is that to infer repeated-game strategies from actions subject by subject, we neither elicit strategies from players nor place a layer of decision-making in front of their choice of actions in the game. Instead we rely solely upon the observed actions of subjects. We therefore minimize – indeed eliminate – the interaction between the inference methodology and subject decision-making. This suggests that our approach to studying repeated-game strategies may be useful in field studies in which the researcher can neither intervene nor collect additional information to help decipher the subject’s strategy.

A limitation of our approach is that by bringing no information besides observed actions to bear on the problem, we encounter difficulties in identifying strategies from the data. We discuss

the steps we took to alleviate this difficulty throughout the paper, and present the evidence for success in the conclusion.

We begin with a brief description of the experimental design and 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

In each round of our 30-round, posted-offer experiments, the monopolist posts a price and offers a quantity for sale at that price. The buyers are then randomly ordered. Each one in turn determines how many units to purchase at the available price. Buyers' purchasing and withholding decisions are made privately. Only the individual buyer knows how many units he purchased and how many profitable units he passed up. This institutional detail makes it impossible for buyers to coordinate 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.

Figures 1a and 1b display the static monopolist's marginal cost and the buyers' aggregate demand curves for the two-buyer and four-buyer treatments, respectively. These curves as well as the buyers' individual demand curves (identical for all buyers in all treatments) 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.²

[insert Figures 1a and 1b here]

Every effort has been made to hold constant across treatments all variables believed to be

²These "informed" treatments are to be distinguished from the "uninformed" treatments (to be discussed in the next section) in which the monopolist only did not know the number of buyers in the market.

important to demand withholding and seller pricing. In this regard, notice that both treatments share the same ten-unit, competitive price range. The midpoint of the competitive range has been normalized to zero, with all other prices, costs and valuations henceforth expressed as deviations from this competitive price. Moreover, both treatments share the identical monopoly price of 20 units above the competitive price (+0.20).

In addition, individual buyer demand curves are identical for all buyers regardless of treatment. Specifically, every buyer 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.

The surplus division between buyers and the monopolist is also held fixed across treatments at 6:1 in favor of the monopolist: at the competitive price, each buyer earns +0.80 compared with +4.80 for the monopolist. Ruffle (2000) shows that the surplus inequality at the competitive equilibrium is an important determinant of the degree of buyer withholding and sellers' prices. Since individual buyers possess the same demand curve, to hold constant the surplus division across treatments requires changing the monopolist's marginal cost curve.

Finally, in both treatments, the buyers face a monopolist. This eliminates possible seller concerns and uncertainty about the simultaneous price choice of additional sellers, thereby allowing us to concentrate on the impact of the buyers' decisions on monopolist pricing without the complication of competition between multiple sellers.

All experimental sessions were computerized and conducted at Ben-Gurion University. The software randomly assigned individual subjects to the role of buyer or seller. Buyers earned the difference between their valuations and the price they paid on each unit they purchased. The monopolist earned the difference the posted price and his cost on each unit sold. The costs of unsold units were not subtracted from the monopolists profits. At the end of the experiment, subjects were paid a 15 NIS (New Israeli Shekel) showup payment in addition to their experimental earnings.³ Average seller earnings (including the showup payment) were 121 NIS compared to 67 NIS for the buyers. Sessions lasted between one hour and one hour and thirty minutes. Seven two-buyer (informed) sessions were conducted along with eight four-buyer (informed) sessions.

³At the time these experiments were conducted 4 NIS was equivalent to approximately \$1 U.S.

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 Brief Summary of the 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 he 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.⁴ 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.

[insert Figure 2 here]

Given the controls built into the experimental design, there remain 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, he 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 is that in the former, the monopolist was not told how many buyer he faced; instead, he was told in both the two-buyer and four-buyer uninformed treatments that he faced “a 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 price gap in the initial and middle rounds between the two-buyer and four-buyer informed treatments has disappeared, as seen

⁴ Engle-Warnick and Ruffle (2002) provides a detailed analysis of the experimental results by treatment.

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 variables are more effective in bringing prices down. The strategy inference technique employed in this paper will allow us to address this hypothesis. We now turn to the repeated-game withholding strategies employed by buyers in these experiments to gain a more complete understanding of the dynamics that led to the observed price levels in the different treatments.

4 The Strategy Model

We 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 always involve a variable, a relation, and a coefficient. Each relational node is a test; 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 followed by a colon; “0:” represents the decision to withhold zero units of demand, “1:” to withhold one unit, and “+:” to withhold more than one unit. The corresponding right column indicates the number of times 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 the 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.⁵

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 intensively 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 because estimation suffers from the “curse of dimensionality”; with each additional relational node the amount of data at the terminal nodes reduces exponentially. With at most 30 observations for each individual, fitting more complex strategies is of questionable value. An indication that this constraint is not too restrictive is that 89/90 buyer strategies inferred consist of strictly less than four relational nodes; a four-relational node strategy was inferred for

⁵ We abstract from semantics and refer to these decision rules as strategies, noting well the fact that we only observe a single history of each repeated game and thus cannot reconstruct the rule that a subject would have used for any possible contingency.

only a single buyer.⁶

We selected the time t price, $P(t)$, and the round number, t , for candidate relational node variables. We allowed any coefficient actually realized during the experimental session.⁷ Perhaps the most obvious variable upon which to base one’s withholding decisions is the current period price. We included the round number to allow the strategies to vary with time. The appearance of time in an inferred strategy suggests that the subject adopted one strategy at a certain point in the game and later discarded it for another strategy. Put differently, subjects may intentionally vary their demand withholding over time; for instance, perhaps a buyer withholds early, independent of prevailing prices, in order to signal toughness, and then later in the session, as the expected future gains from withholding dwindle, ceases to withhold.

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 appropriately complex strategies; 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).⁸ Our Bayesian approach assigns priors to each of the assignment problems inherent in formulating a strategy, and then finds the posterior mode estimate of the joint assignment.⁹

⁶ Another way to look at the problem of increasing complexity approximated by the number of relational nodes is that the number of strategies in the class increases roughly exponentially with the number of relational nodes, making it increasingly difficult to identify strategies in the data. For one of our subjects, there were 34 possible trees with one relational node, 999 with two nodes, 29,740 with three nodes, and 299,386 with four nodes.

⁷ We ran the procedure including also $\Delta P(t)$, the difference in price from one period to the next, and $W(t - 1)$, the withholding decision in the previous period. The former variable never occurred in an inferred strategy, and the latter did not contribute to a useful interpretation of what the subjects may have been doing.

⁸ In their application, a subset of k decision rules is selected from a set of n candidate decision rules, denoted C , to best–fit the actions 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 find the posterior mode estimate of the joint assignment.

⁹ See also Denison, Mallick, and Smith (1998) and Chipman, George, and McCulloch (1998) for Bayesian stochastic search algorithms. A related approach is “bagging” (Breiman, 1994), in which multiple versions of a predictor are used

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 *variable* at each relational node, and (3) the *configuration* of the relational nodes in the tree representation of the strategy. The idea is that *a priori* we do not know the functional form of a subject’s strategy, so we form priors over each of the assignments involved in constructing the functional form and update with the likelihood.

The first assignment can in theory be any nonnegative integer, k . As noted above, in practice we impose the restriction $k \leq 4$. In Figure 5 we choose $k = 2$ for illustrative purposes.

[insert Figure 5 here]

The second assignment involves choosing k variables for use in k unique relational nodes. Anticipating the fact that no two relational nodes that contain the same variable will contain the same 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 $C = \{P_1(t), P_2(t), t_1, t_2\}$.¹⁰ 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 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.

This illustrates the construction of a single strategy. In order to estimate the distribution of

to form an aggregate predictor. However, the interpretability of the output as a strategy is lost with this approach. Also related is “bumping” (Tibshirani and Knight, 1997), which is a bootstrap-based procedure. Our application lacks the number of observations necessary to perform bumping, but see Engle–Warnick (2002) for a similar approach applied to experimental data.

¹⁰ We abstract from the number of coefficients for the purpose of forming priors because we wish to avoid penalizing a rule for complexity just because there happen to be a large number of possible coefficients. For example, when we include time as the conditioning variable, we would have to divide prior probabilities over 30 different round numbers. The effect would be a rather steep penalty for a phenomenon that does not appear to us to be related to the actual complexity of the strategy.

the strategy that best characterizes the subject’s observed decisions, and in order to locate the best-fitting strategy, we conduct an exhaustive search over all possible strategies.

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 as noted above we imposed the restriction $k \in \{0, 1, 2, 3, 4\}$. We follow the literature (see, e.g., Denison, Mallick, and Smith, 1998) and specify the Poisson distribution for the prior $p(k)$, truncated at $k = 4$. Thus the specification of this prior amounts to the specification of the mean λ of the number of relational nodes. We report the sensitivity of the inference results to a range of means for k .

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 possible such choices that exist. The calculation is straightforward: the number of combinations of $k \cdot n$ objects taken k at a time is

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

In the example above 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, and again calculate the number of such choices that exist. 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). Whenever a relational node is added to a tree, two action nodes are also added. However, 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 (1+1) 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; we then construct all possible trees from k nodes using the chosen variables. Given the analysis above, 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 simple 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 probability of k relational nodes is

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

where λ is the mean of the distribution. The probability of selecting any k relational nodes from a set of n variables is $1/S_n^k$, and the probability of any particular arrangement of k relational nodes into a strategy is $1/T^k$.

5.3 Building the Likelihood Function

For the probability model, let $x_t = (x_t^{(1)}, \dots, x_t^{(d)}) = (P(t), t)$ 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. The division of the action space into these three categories is behaviorally meaningful. Withholding zero units is the passive price-taking behavior predicted by backward-induction reasoning. Withholding a single unit while qualitatively different from passive price-taking requires little sacrifice since the buyer's last unit is his least profitable one, yielding, for instance, only 0.05 units of profit at the competitive price. Two or more units of demand withheld,

by contrast, signals a committed attempt to force down the monopolist's price.¹¹ From a theoretical viewpoint, it is also natural to group together the decision to withhold two or more units since such a decision indicates an exercise of market power.

We consider a class of strategies, as noted above, 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. We consider strategies with no more than four relational nodes.

A strategy $g \in G$ can be thought of as a possibly nested if-then statement which 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

¹¹ Furthermore, we do not separate withholding two units from withholding three or four units since this distinction is sometimes artificial rather than a reflection of buyer intent. For example, if the monopolist posts a price above the competitive range, then it can be shown that the buyer last in the queue (independent of treatment) can withhold at most two units. Similarly, for prices below the competitive range, the last buyer in the queue can withhold no more than three units.

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

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

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. It can be shown that this procedure maximizes the log-likelihood for strategy g .

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:

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

Each of the three additional 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 no strategy 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 suggests 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.¹² We report the posterior probability for each of the possible values of k , as well as the probability that each explanatory variable 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 rules inferred at a selected value of the mean, and we report the modal rule for all subjects in selected experimental sessions.

6 Experimental Results

We present the results in three sub-sections. First, we present inference results aggregated across experimental treatments. Second, we present posterior distributional results for each buyer in the form of strategy characteristic statistics such as, for example, the probability that the buyer’s strategy contains k nodes, $k=0,1,2,3,4$. 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.

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 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 modal number of nodes in the modal strategies inferred 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

¹² 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 limited to cases in which an exhaustive search is possible. See the references listed in footnote 9 for stochastic search procedures.

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 cost of complexity discussed in section 5, 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. These figures reveal 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 he 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 liberty to be able to fine tune their withholding behavior.

Strategy complexity as a function of the number of buyers do not reveal a similarly evident relationship: Figure 7b shows that whether strategies are more complex with different numbers of buyers depends on the prior.

[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. 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 subsection.

In addition, Figure 9a shows that decreasing withholding with price is equally highly unlikely in all treatments; this intuitive result serves to validate the inference procedure, as 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 in strategies when monopolists are informed.

[insert Figures 9a and 9b here]

6.2 Strategy Characteristics for Individual Buyers

In the previous subsection we explored the strategy complexity and strategy composition on an aggregate level when we varied the prior mean number of relational nodes from one to three. In the next two subsections we examine more closely the inferred strategy characteristics at the subject level by reporting the posterior estimates of the distributions of inferred strategies for individual buyers. To report the distributions of strategy characteristics for each individual buyer over the entire range of prior mean numbers of relational nodes from one to three would be far too unwieldy. For the sake of parsimony, we must therefore select a specific prior mean number of relational nodes. The good news is that all of the strategy complexity and strategy composition observations

have been shown to be insensitive to the prior mean number of relational nodes. To avoid taking a strong stand on the choice of a prior, we choose two, the midpoint of the range of the priors explored.

Table 1 displays the results for the two-buyer, informed treatment, and Tables 2, 3, and 4 contain identical information for the remaining treatments. From left to right the table reports the session and subject identification numbers, the breakdown of each buyer’s observed withholding decisions, the median session price from the last five rounds, the mean seller and buyer per round profits, the number of relational nodes and the error rate of the modal strategy (to be discussed below), the posterior probability that the buyer’s strategy contains 0, 1, 2, 3, and 4 relational nodes, and the posterior probability that the strategy contains the conditioning nodes Time - , Time +, Price - and Price +.

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. On the whole, the strategies fit the decisions well as seen by the error rate which appears in the “modal rule characteristics” column of Tables 1-4. The error rate is calculated by taking the sum of the non-modal decisions at each of the action nodes and dividing by the total number of decisions made; in other words, the error rate is the proportion of decisions that the strategy misclassifies. Across all treatments we averaged the error rate and found it to be approximately 0.16, meaning that strategies on average classify correctly 25/30 of the buyer’s observed withholding decisions.

[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. (In Session 7, Buyer 1’s modal strategy contains Time + with probability 0.784.) Two subjects increase their withholding

with an increasing price (Price +) with probability greater 0.8. Error rates vary for non-degenerate strategies 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 his withholding actions – he withheld one or more units in 13 rounds – his 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. Indeed 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 +. (Recall that only one buyer in the two-buyer, informed treatment conditioned Time + with a probability close to 0.8.) 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, uniformed 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, 17/28 buyers condition on Time - with probability greater than or equal to 0.8 and four buyers condition on Time + with the same probability.

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 bold-face 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 buyers are able to lower the price through their withholding strategies, at no overall cost (or benefit) to themselves. Put differently, 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 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 reveals that the most effective strategy component against seller pricing is Time-: 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 his purchases (i.e. decreases his withholding) when the price increases, his profits will be relatively low.

To further investigate 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.¹³

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]

¹³ PcGets automatically selects a final (i.e., specific) model which 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.

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 affect: median buyer profits are 0.3 NIS per round higher in the two-buyer treatments than in the four-buyer treatments. This follows from the fact that prices in the two-buyer sessions are lower, 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 such prices. The monopolist thus responds by lowering his price to increase sales and profits. If buyers do not withhold early on, feeling encouraged, the monopolist may raise his price and continue to do so until he is met with resistance.

6.3 Strategies in Selected Sessions

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 with 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 12a. The modal strategy contains three relational nodes for Buyer 1 and one relational node for Buyer 2. Interpreting buyer 1's strategy, he became price sensitive after period five: whenever the price exceeded -0.3, he withheld multiple units of demand nine out of eleven times. Whenever the price was less -0.3, this buyer withheld a single unit of demand up to period 24 and zero units after period 24. 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. Notice from Figure 12a 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 implicitly coordinated intense early withholding in the game combined with one buyer's sensitivity to a resulting low price appears to have driven the monopolist's price down to a remarkably low level in this session.

[insert Figures 12a and 12b here]

The strategies inferred in four-buyer session 4B10 are presented in Figure 11, with the price series shown in Figure 12b. Buyers 1 and 4 exhibited behavior consistent with sensitivity to a price threshold. These thresholds of 0.03 and -0.02 both fall within the competitive price tunnel, well below the monopoly price of 0.20. The time relational nodes for both of these subjects showed a higher propensity for unconditional demand withholding early in the game. Again from Figure 12b one can see that this is not for want of price variance. The strategy for Buyer 2 shows a step-wise sensitivity to prices at the levels of -0.03 to 0.00 to 0.04. Buyer 3 did not withhold demand on even a single instance. While the withholding was slightly less coordinated in this session than in Session 2B7, again 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.

These examples visually demonstrate the strategic nature of interactions in these experiments. They paint a picture of buyers' reactions to the pricing experimentation of the monopolists. The threat of demand withholding is very real, and appears to be an effective response against the power of the monopolist. The strategies are even more remarkable when one considers the fact that

buyers were not informed of the decisions of other buyers in their market. The buyers may have believed that either their actions alone would affect the pricing decision or that other buyers in the market would respond in the same way thus increasing the effectiveness of the individual strategies. At any rate, the evidence suggests that buyers systematically deviated from passive price-taking. For the first time we are able to characterize their deviations with repeated-game strategies.

7 Conclusions

In this paper we apply a Bayesian strategy inference technique to infer unobserved strategies from the observed actions of buyers in posted-offer monopoly experiments. 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, we are able to infer repeated-game withholding strategies for individual buyers.

While for one-quarter of the buyers in these experiments, we infer the passive price-taking strategy, for three-quarters of the buyers, we find active, repeated-game, demand withholding strategies. 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 to the buyers. More interestingly, 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 strategies triggered by a price threshold.

This paper is part of a growing accumulation of evidence that models of repeated-game strategies can usefully describe observed behavior: Duffy and Engle-Warnick (2001) and Engle-Warnick (2002) fit a similar strategy model to ultimatum game data, and Engle-Warnick and Slonim (2001) fit finite automata to observed actions in a repeated trust game. This evidence suggests that we

can learn more about strategic behavior in repeated games by at least complementing existing analysis with strategy analysis such as the one in this paper. It also suggests the importance of future innovative research: reducing or solving the identification problem requires bringing more information to bear on the problem (Manski, 2001), including strict experimental control (such as the control we presented in our experiments) or direct elicitation (e.g., Nyarko and Schotter, 2000, who elicited beliefs from subjects). However, this work serves as a first 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.

As such we view this work as a first step toward applying strategy inference techniques of this type to field data. A better understanding of decisions to strike, provide credit, buy and sell stocks and bonds, or to implement monetary policy (Duffy and Engle-Warnick, 2001), for example, could result from methods such as this one which go a step further than more conventional methods of inference toward allowing the data to determine the form of the underlying strategy. The approach 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 constructs a probability model suitable for the computation of a likelihood from the data. Since the identification problem is severe, particularly in circumstances where relatively few observations are available, results based on experimental data serve to validate the procedure in an environment where a degree of control may be exercised over the decision-making environment.

References

- [1] Bjornerstedt, J. and J. Stennek (2001) "Bilateral Oligopoly", working paper No. 555, Research Institute of Industrial Economics, Stockholm, Sweden.
- [2] Breiman, L. (1994) "Bagging Predictors", Department of Statistics, University of California Berkeley, Technical Report No. 421.

- [3] Breiman, L., H. Friedman, A. Olshen, and C.J. Stone (1984) *Classification and Regression Trees*, Wadsworth International, Belmont, CA.
- [4] Camerer, C. (2001) *Behavioral Game Theory: Experiments on Strategic Interaction*, In press.
- [5] Chipman, H. and E. George and R. McCulloch (1998) “Bayesian CART Model Search”, *Journal of the American Statistical Association*, 93: 443, 935-960.
- [6] Chipty, T. and C. Snyder (1999) “The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry”, *Review of Economics and Statistics*, 81:2, 326-340.
- [7] Costa-Gomez, M., V. Crawford, and B. Broseta (2001) “Cognition and Behavior in Normal-form Games: An Experimental Study”, *Econometrica*, 69:5, 1193-1235.
- [8] Davis, D. and C. Holt (1993) *Experimental Economics*, Princeton University Press, Princeton, NJ.
- [9] Denison, D. and B. Mallick and A. Smith (1998) “A Bayesian CART Algorithm”, *Biometrika*, 85:2, 363-377.
- [10] Devroye, L., L. Györfi, and G. Lugosi (1996) *A Probabilistic Theory of Pattern Recognition*, Springer-Verlag, NY.
- [11] Duffy, J. and J. Engle-Warnick (2001) “Multiple Regimes in US Monetary Policy? A Non-parametric Approach”, Unpublished manuscript, Nuffield College, University of Oxford.
- [12] Duffy, J. and J. Engle-Warnick (2000) “Using Symbolic Regression to Infer Strategies in Experimental Data, forthcoming in *Evolutionary Computation in Economics and Finance*, S.H. Chen editor, Springer-Verlag, NY.
- [13] El-Gamal, M. and D. Grether (1995) “Are People Bayesian? Uncovering Behavioral Strategies”, *Journal of the American Statistical Association*, 90:432, 1137-1145.
- [14] Engle-Warnick, J. (2002) “Inferring Strategies from Observed Actions: A Nonparametric, Binary Tree Classification Approach”, forthcoming in *Journal of Economic Dynamics and Control*.
- [15] Engle-Warnick, J. and B. Ruffle (2002) “Buyer Countervailing Power versus Monopoly Power: Evidence from Experimental Posted-Offer Markets”, unpublished manuscript.
- [16] Engle-Warnick, J. and R. Slonim (2001) “Inferring Repeated Game Strategies from Actions: Evidence from Trust Game Experiments”, Nuffield College, Oxford University Discussion Paper No. 2001-W13.
- [17] Hendry, D. and H.-M. Krolzig (2001) *Automatic Econometric Model Selection Using PcGets*, Timberlake Consultants Ltd., London, UK.
- [18] Kagel, J. and A. Roth (1995) *The Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ.

- [19] Johnson, E., C. Camerer, S. Sen, and T. Rymon (2001) “Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining”, Unpublished manuscript, Columbia School of Business, Columbia University.
- [20] Manksi, C. (2001) “Identification of Decision Rules in Experiments on Simple Games of Proposal and Response”, Unpublished manuscript, Department of Economics and Institute for Policy Research, Northwestern University.
- [21] McKelvey, R. and T. Palfrey (1992) “An Experimental Study of the Centipede Game”, *Econometrica*, 112:803-836.
- [22] Nyarko, Y. and A. Schotter (2000) “An Experimental Study of Belief Learning Using Elicited Beliefs”, forthcoming *Econometrica*.
- [23] Ruffle, B. (2000) “Some Factors Affecting Demand Withholding in Posted-Offer Markets”, *Economic Theory*, 16:3, 529-544.
- [24] Schumacher, U. (1991) “Buyer Structure and Seller Performance in U.S. Manufacturing Industries”, *Review of Economics and Statistics*, 73:2, 277-284.
- [25] Selten, R., M. Mitzkewitz and G. R. Ulrich(1997) “Duopoly Strategies Programmed by Experienced Traders”, *Econometrica*, 65:517-555.
- [26] 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.
- [27] Snyder, C. (1996) “A Dynamic Theory of Countervailing Power”, *RAND Journal of Economics*, 27:4, 747-769.
- [28] Stahl, D. and P. Wilson (1995) “On Players’ Models of Other Players: Theory and Experimental Evidence”, *Games and Economic Behavior*, 10:218-254.
- [29] Tibshirani, R. and K. Knight (1997) “Model Search and Inference by Bootstrap “Bumping”, Unpublished manuscript, Department of Statistics and Department of Public Health Sciences, University of Toronto.

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 Rule 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

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 Rule 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

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 Rule 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

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

Top cell number: correlation; Bottom cell number: p-value, significant results in bold

Table 6: Regressions of Market Outcomes on Strategy Characteristics

Dependent Variable	Independent Variables			
	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)		

Top cell number: estimated coefficient, bottom cell number: std. error; all significant at 5%

Figure 1a: Two-Buyer Treatment Parameters

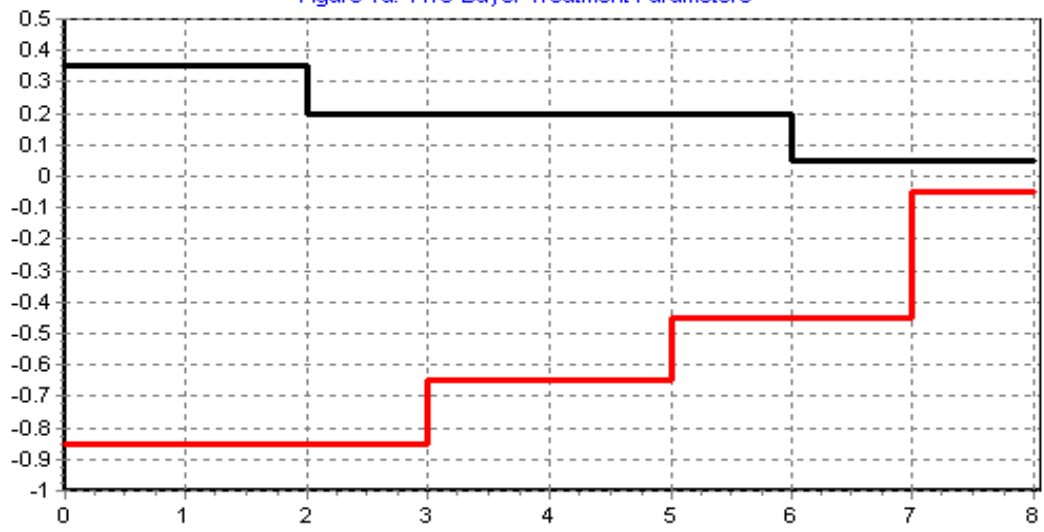
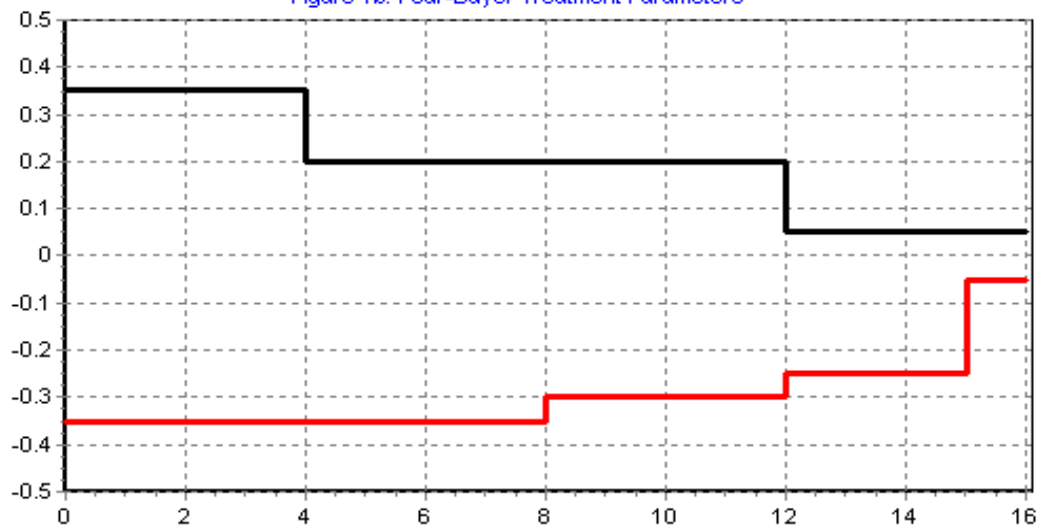


Figure 1b: Four-Buyer Treatment Parameters



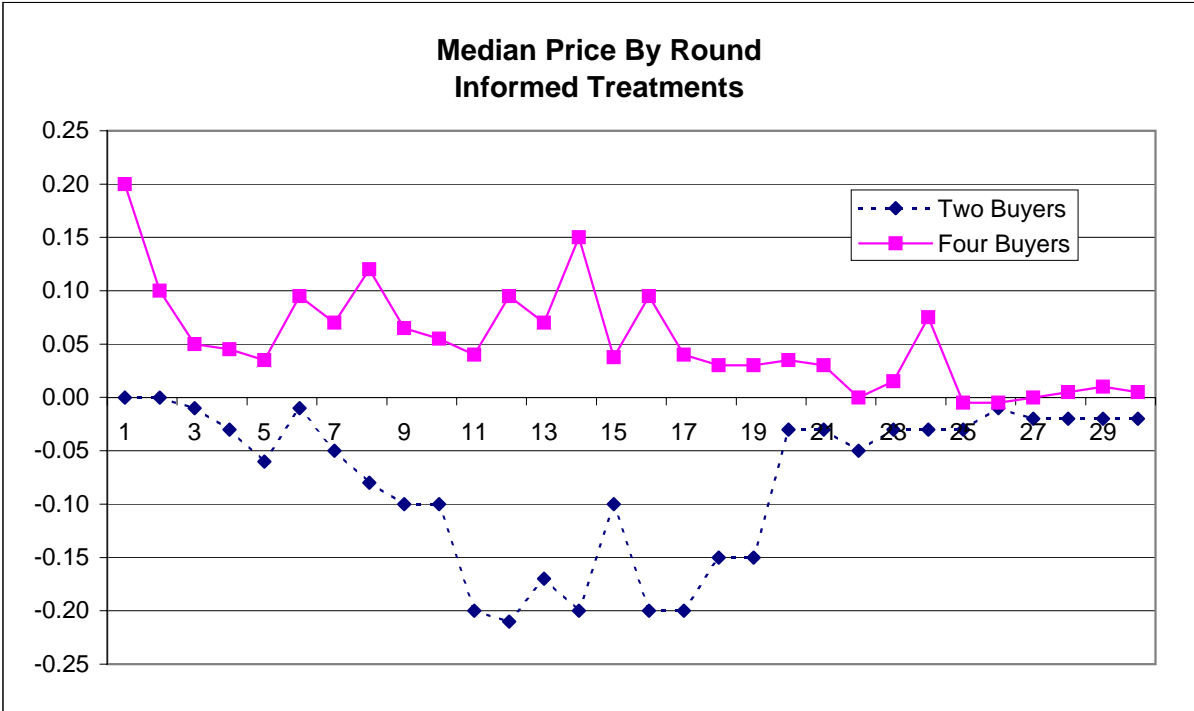


Figure 2: Median Price in Informed Sessions

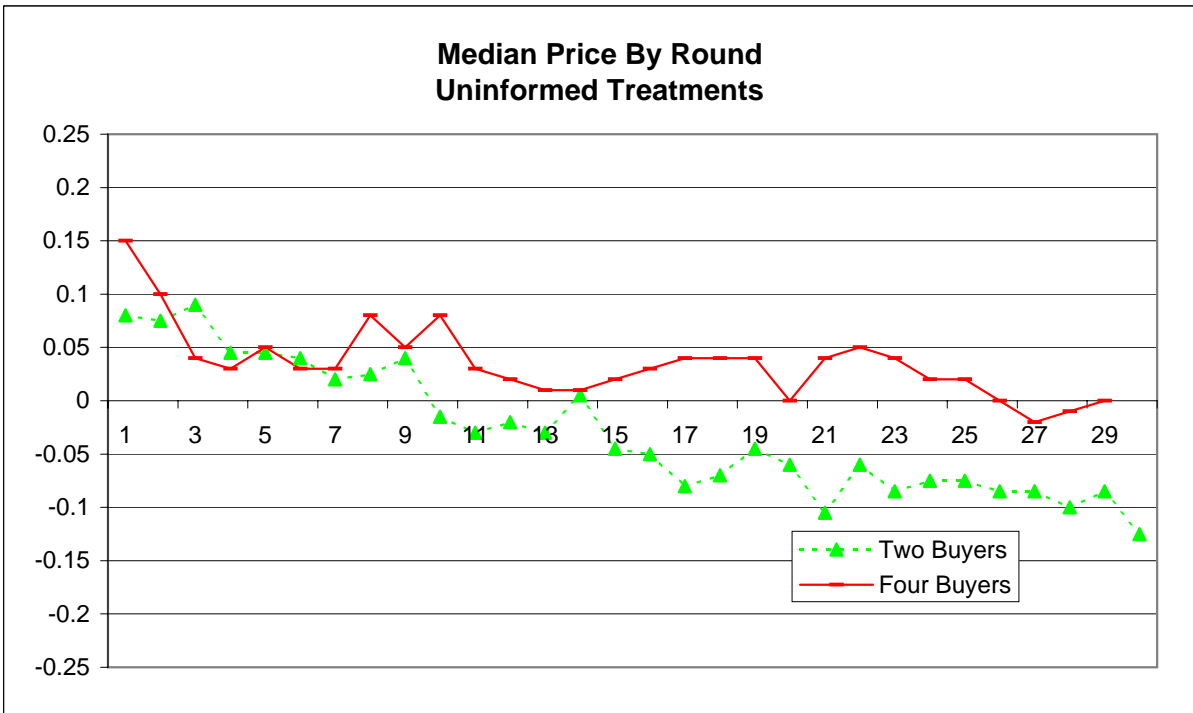


Figure 3: Median Price in Uninformed Sessions




Figure 4: An Example Strategy

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

Assignment 1: k , the number of relational nodes

Choices: $k \in \{0, 1, 2, 3, 4\}$

Example choice: $k = 2$

Assignment 2: Select $k = 2$ variables from Variable Set $C = \{P_1(t), P_2(t), t_1, t_2\}$

Choices: $\{P_1(t), P_2(t)\}, \{P_1(t), t_1\}, \{P_1(t), t_2\}, \{P_2(t), t_1\}, \{P_2(t), t_2\}, \{t_1, t_2\}$

Example choice: $\{P_1(t), t_1\}$

Assignment 3: Node Configuration Choices:

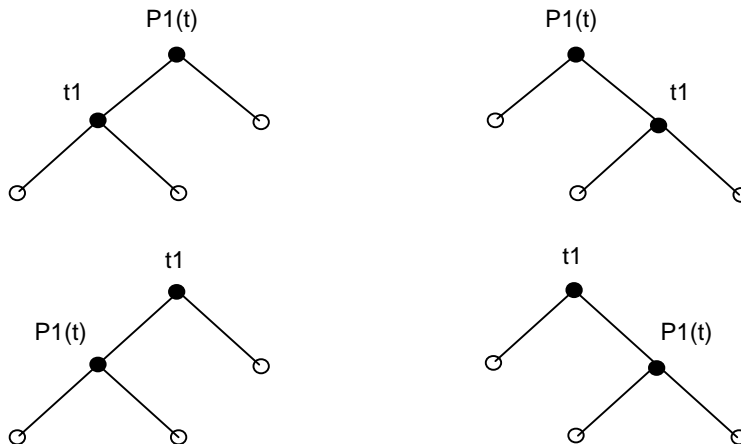


Figure 5: Constructing a Strategy

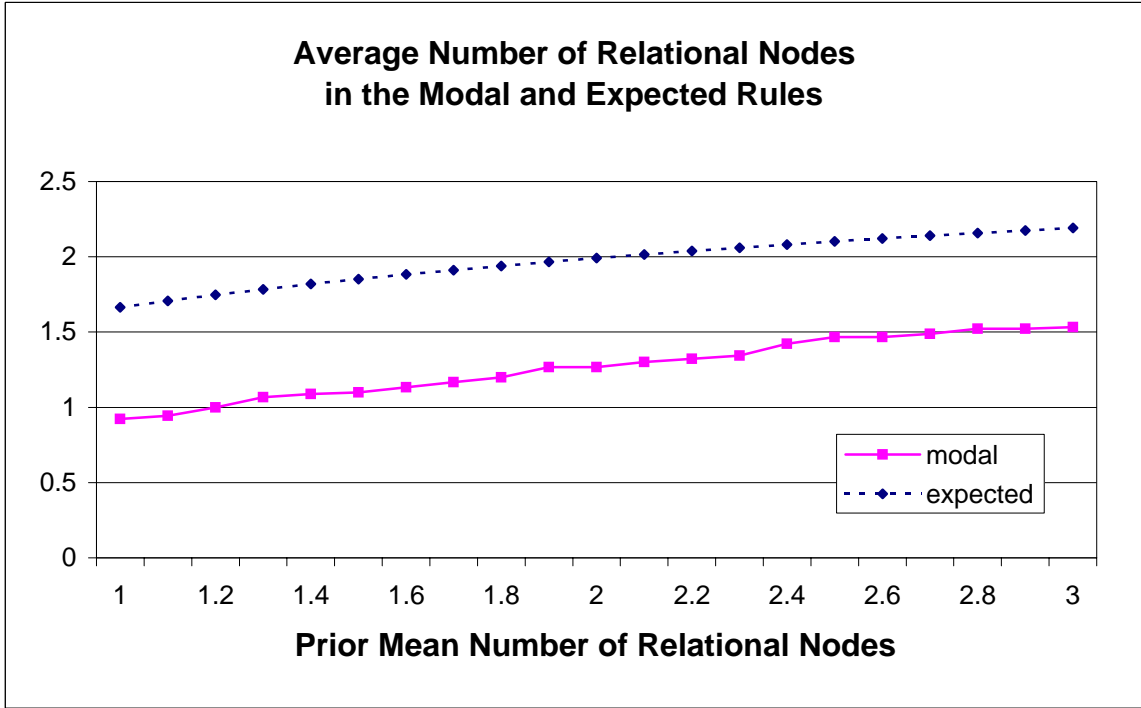


Figure 6: Average Complexity of the Strategies

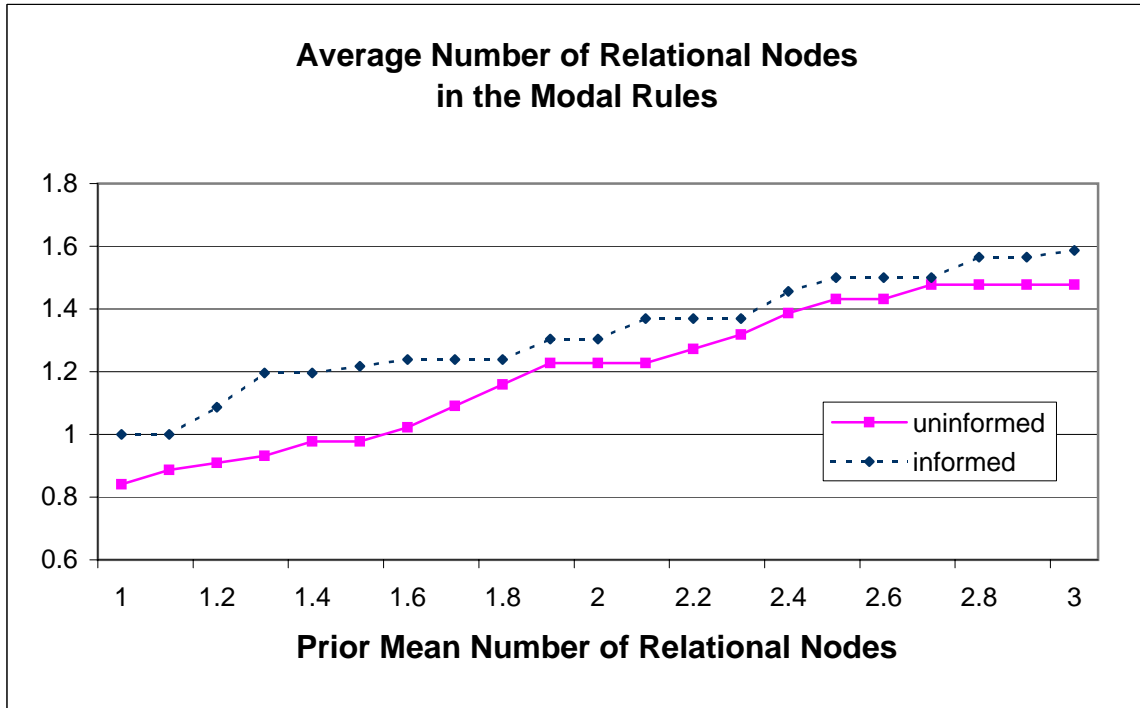


Figure 7a: Average Complexity by Informational Treatment




Figure 7b: Average Complexity by Number of Buyers

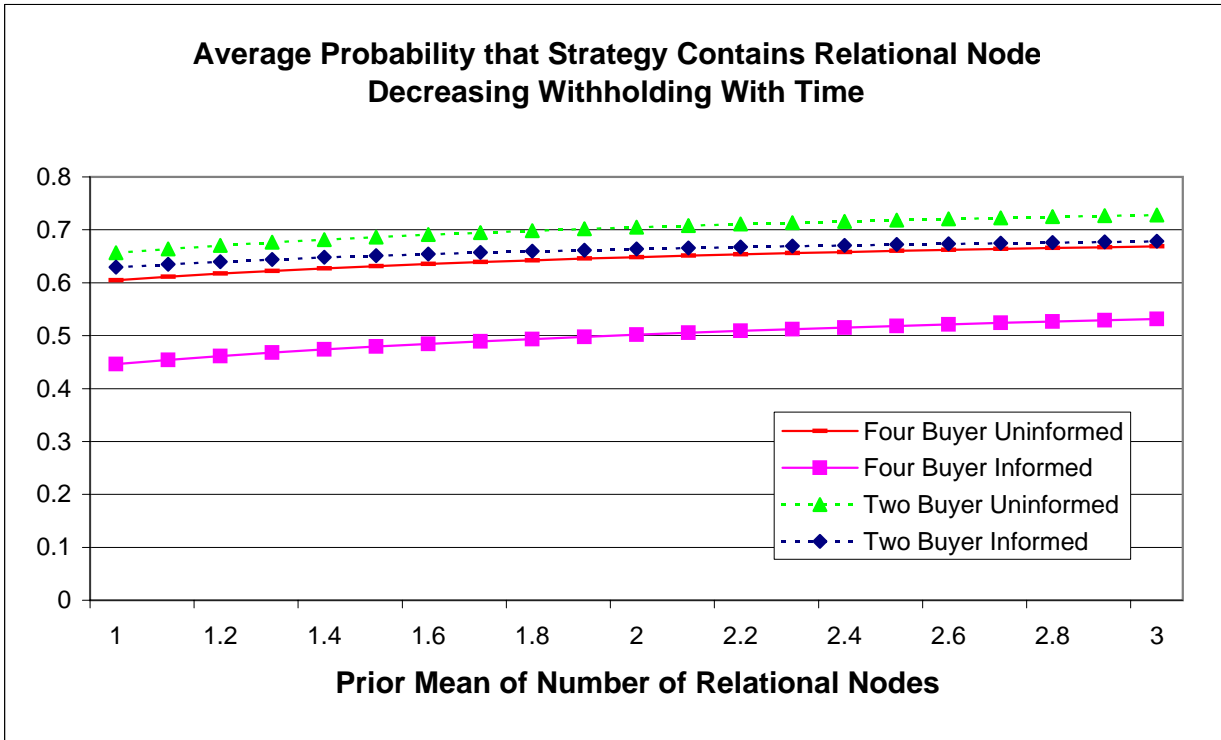


Figure 8a: Probability of Time- Relational Node

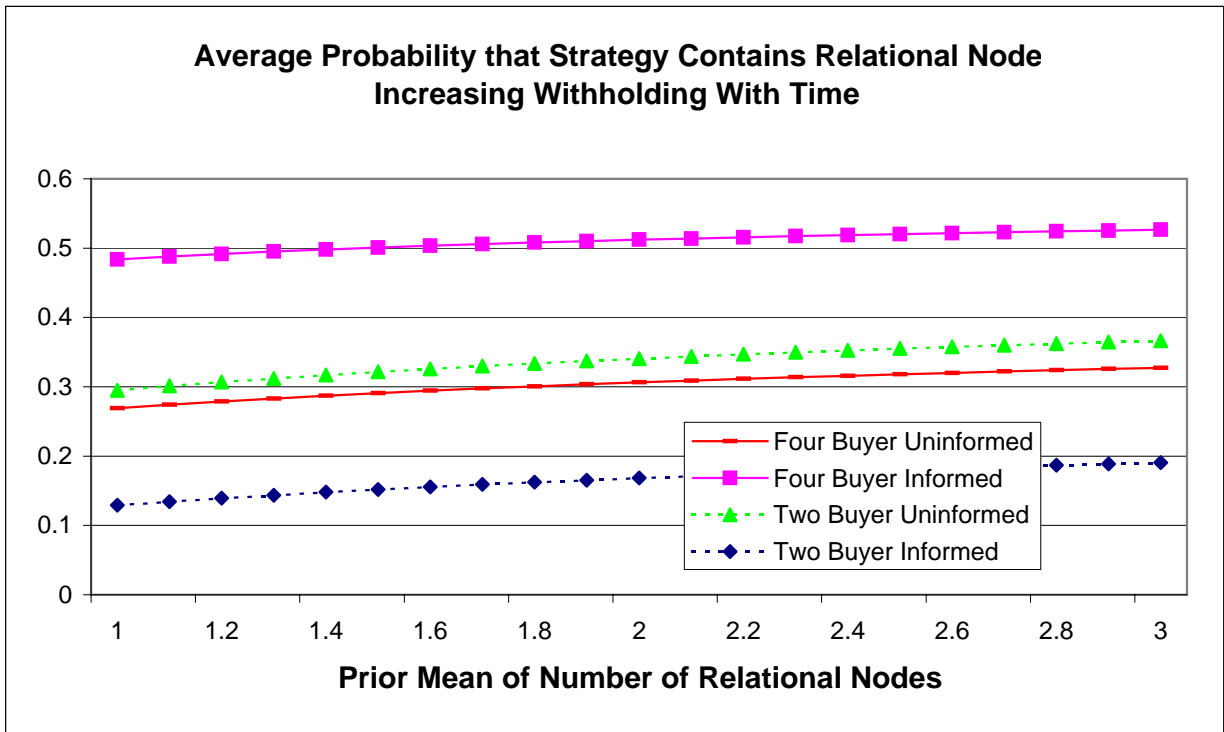


Figure 8b: Probability of Time+ Relational Node




Figure 9a: Probability of Price- Relational Node




Figure 9b: Probability of Price+ Relational Node

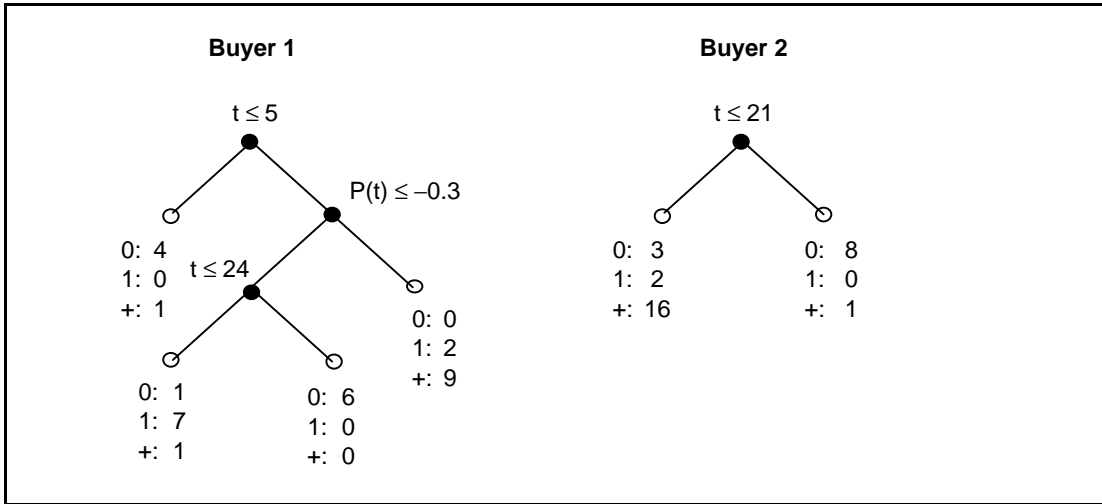


Figure 10: Strategies in Two Buyer Session 7




Figure 11: Strategies in Four Buyer Session 10

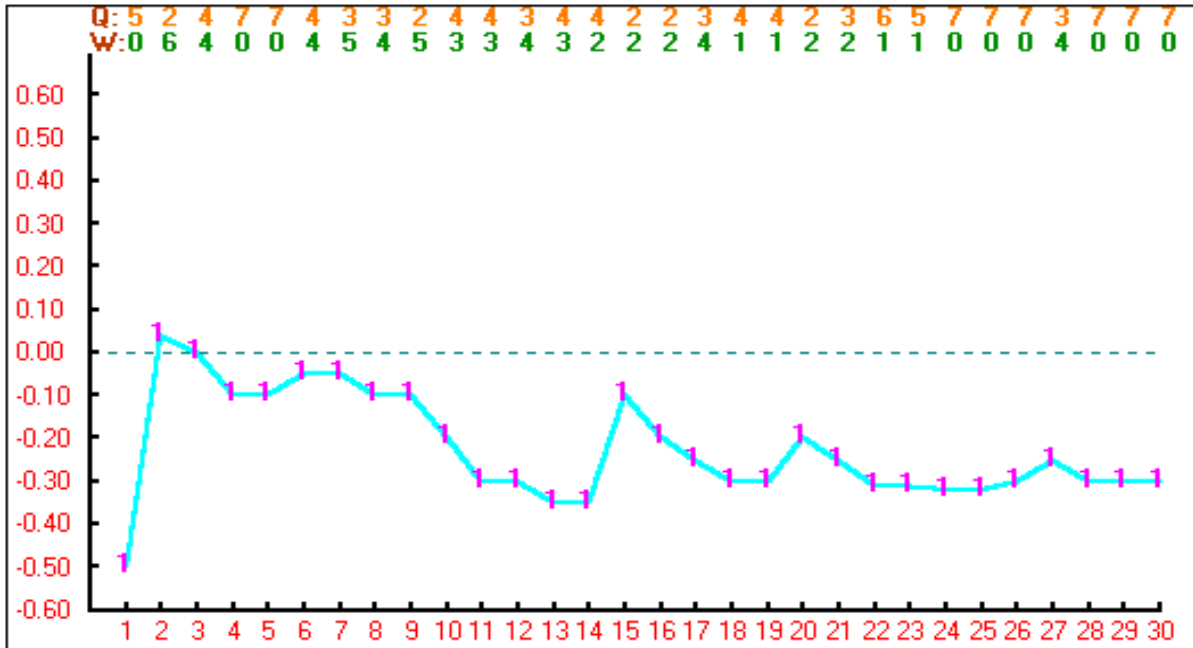


Figure 12a: The Monopolist's Posted Prices in Two-Buyer, Informed Session 7

Note: The sequence labelled Q at the top of the figure indicates the number of units sold by period, and W indicates the number of units lost to demand withholding by period.

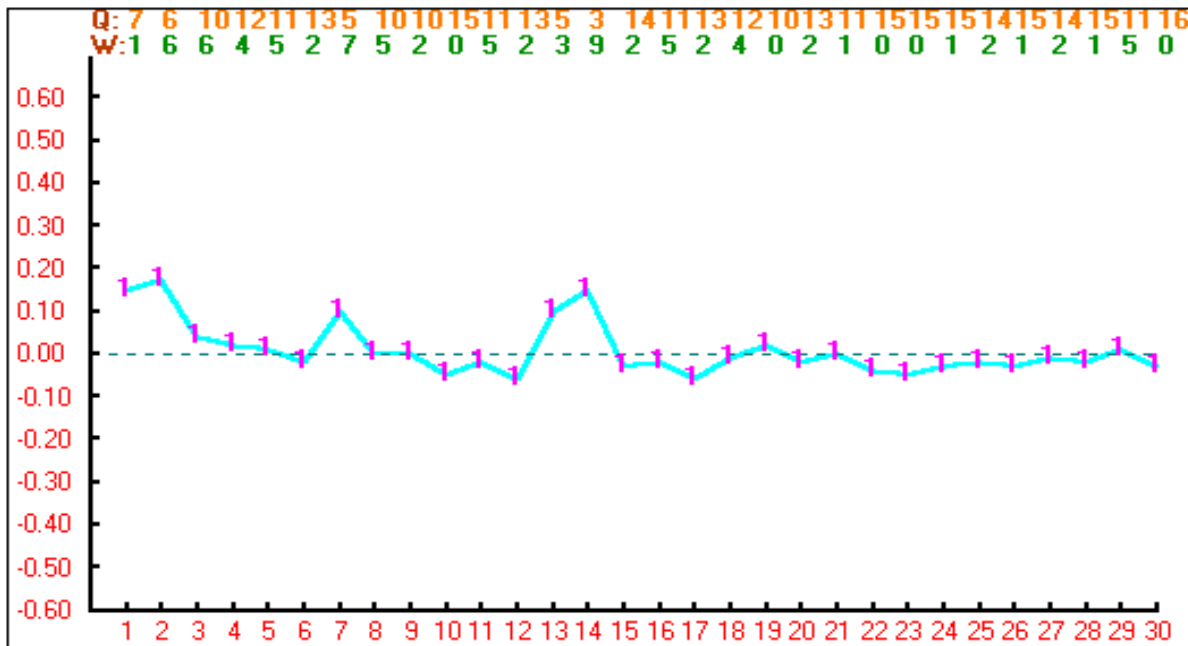


Figure 12b: The Monopolist's Posted Prices in Four-Buyer, Informed Session 10

Note: The sequence labelled Q at the top of the figure indicates the number of units sold by period, and W indicates the number of units lost to demand withholding by period.