Consumer Search Costs and Market Performance

Douglas D. Davis and Charles A. Holt*

December 1993
(revised subsequently for the *Economic Journal*)

Abstract

Laboratory markets are conducted to evaluate the effects of consumer search costs on market performance. The primary research goal is to assess the behavioral relevance of Diamond's (1971) paradoxical conclusion that the injection of a small consumer search cost alters the equilibrium price prediction from competitive to monopoly levels. Although monopoly prices are not consistently observed, we find that search costs do tend to raise prices. Additional experimentation indicates that below-monopoly prices are not explained by sellers' price-posting reputations, but that prices increase as search costs are raised. We conjecture that the Diamond prediction fails because sellers neither immediately appreciate, nor readily learn the necessary recursive "price-plus-search-cost" reasoning.

1. Introduction

One of the seminal developments leading to the information economics revolution in the 1970's and 1980's was Diamond's (1971) theoretical result that, in the absence of publicly posted price information, the existence even a small search cost could lead to monopoly pricing. The intuition is straightforward: No buyer with one price quote would want to search for a second, unless the new quote is expected to be lower by at least the amount of the search cost. Thus each seller has an incentive to price slightly *above* any common price, and the noncooperative equilibrium in a single-stage game yields a monopoly price. This result is viewed as a paradox, since a "small" search cost produces high prices, but a zero search cost would produce the usual Bertrand incentives that drive prices to competitive levels, in the absence of capacity constraints

* Virginia Commonwealth University and University of Virginia, respectively. This research was supported by grants from the University of Virginia Bankard fund and the National Science Foundation (grants SES 9012694 and SES 9012691).

and other imperfections.[1]

Recent theoretical work has focused on finding a "resolution" to the paradox. Stahl (1989), for example, generates a smooth transition between competitive and monopoly outcomes as the fraction of consumers with zero search costs increases from 0 to 1. Interestingly, Stahl finds that, holding the number of informed buyers constant, increasing the number of sellers makes pricing more monopolistic. Bagwell and Ramey (1992) propose an alternative resolution that applies to an infinitely repeated market game where sellers who raise prices can develop a reputation that can affect sales in future periods. Bagwell and Ramey show that buyers can obtain lower equilibrium prices by following a "loyalty-boycott" search rule, where low-pricing sellers are rewarded with repeated purchases, and high-pricing sellers are punished with switching. The price predictions vary continuously between monopoly and competitive levels, depending on the size of the consumer search cost.

The policy relevance of the Diamond paradox depends on whether there are realistic market environments in which the provision of public price information will actually lower prices and improve efficiency.[2] Due to the difficulty of controlling and measuring information flows in natural markets, the laboratory represents an ideal place to evaluate the Diamond paradox and proposed resolutions. The only experimental analysis of the Diamond paradox of which we are aware is reported by Grether, Schwartz and Wilde (1988), who observed monopoly pricing in 3 of 4 predicted cases. Their results are suggestive, but not definitive, for reasons discussed below. In particular, all four cases involved the same group of subjects "to the extent possible" (p. 328).

This paper reports an experiment consisting of twelve market sessions designed to assess the effects of public price information. The markets are conducted as normal posted-offer

---

[1] Much subsequent theoretical research focused on the sensitivity of this paradox to alterations in the information-transmission technology. The resulting models identify conditions under which monopoly, competitive and heterogenous prices are predicted, depending on how information is disseminated. See, e.g., Butters (1977), Salop and Stiglitz (1977), and Wilde and Schwartz (1979).

[2] Both intuition and the Diamond result would suggest that if goods are sold on a posted-price basis, then public information about those prices should improve performance. But it should not be assumed that the effect of public information is independent of other institutional aspects of the market. For example, Hong and Plott (1982) show that market performance is impaired, not improved, by the transition from unstructured bilateral negotiation (with no public price information) to posted prices (with public price information). In contrast, the Diamond paradox pertains to a comparison between costly and public price information, within a common institutional framework in which prices are posted on a take-it-or-leave-it basis.

markets, except that buyers must pay a small cost each time they approach a different seller. In baseline "Posted-Offer" treatments, prices are publicly displayed to buyers, while in "Search" treatments, prices are not publicly displayed. In brief, we found that prices approach competitive levels under the Posted-Offer treatment, and prices are significantly higher under the Search treatment. Nevertheless, the Diamond prediction of monopoly pricing under Search is not observed, although increases in search costs do raise prices. These results indicate that economists should take a more careful look at attempts to resolve the Diamond paradox by introducing other factors that may impede monopoly pricing in markets with consumer search.

The paper is organized as follows. Our two primary treatments are outlined in section 2, and results are presented in section 3. The failure to observe the Diamond monopoly price prediction motivated additional research on the effects of reputations and search costs, which is presented in section 4. The final section contains a conclusion.

2. Experiment Design and Procedures

The first six sessions to be discussed involved the Posted-Offer and Search treatments mentioned above. Each session consisted of two sequences of 20 market periods, with prices being posted publicly to buyers in one sequence and not in the other. Sellers could not see one another's posted prices in either treatment. The treatment order was alternated in every other session in this "within-group" design. In all sessions the buyers' shopping cost was set at 15 cents per seller approach.

The Search treatment implemented a standard sequential search setup: a buyer had to specify a seller and pay the 15 cent cost in order to see the seller's price. The buyer could then either purchase at the posted price by pressing "p" on the keyboard, or shop elsewhere by pressing "s." In the Posted-Offer treatment, a buyer could see the all sellers' prices at no cost, and the displays indicated whether a seller was out of stock. The buyer would still have to pay the 15 cent shopping cost to approach a seller. This cost is like a travel cost in the sense that it is independent of the number of units purchased.[3]

[3] Thus we compare performance across price-publicity conditions, holding search costs constant. This differs from the usual theoretical discussion, which focuses on the effects of injecting costly search into an environment where prices are not public. We chose not to evaluate performance in a no-search cost/private-information baseline environment largely

It is worth emphasizing the way in which our Search treatment differs from the Grether, Schwartz, and Wilde (1988) design. In their Search ("Monopoly") treatment, buyers were shown the complete list of posted prices at the beginning of each period, without seller identifications. A buyer could avoid a search cost by making a purchase (if profitable) from one seller randomly selected by the experimenter. Alternatively, the buyer could pay a cost (10 cents) to obtain a sample of 2 or more randomly selected sellers, so that a purchase could be made at the lower price. We decided not to reveal any prices to buyers unless they paid a search cost in a sequential search setup. We were motivated in part by Stahl's (1989, p. 700) argument against "... the dubious assumption that consumers can 'see' deviations by firms before they actually search."[4] In any event, our design is unquestionably as legitimate as the Grether, Schwartz and Wilde design, and is perhaps more appealing from the perspective of realism.

Each session consisted of 3 buyers (B1-B3) and 3 sellers (S1-S3), with values and costs indicated by the trader identification for each unit on the market demand and supply curves shown on the left side of figure 1. The buyers are symmetric, each with units valued 70 and 30 cents above the competitive prediction, P_e, which is normalized to zero. Similarly, sellers faced identical costs, with a single unit at a marginal cost 30 cents below P_e, and three additional units with costs at P_e. The large excess supply of 6 units at supracompetitive prices makes the market very competitive; one seller will be left out of the market if sellers offer at least 2 units.[5]
In particular, the design violates the standard assumption that sellers produce at constant marginal cost, and that buyers have a constant and uniform reservation value. The high-value and low-cost

---

for procedural reasons: It takes each buyer several minutes to elicit costless price quotes from sellers each period. Not only would this treatment make it difficult examine a pair of treatments in a 2-hour session, but any observed differences in performance from that observed in a standard posted-offer market could be attributable to (uncontrolled) time-costs of shopping.

[4] There are numerous other differences between our search treatment, and that reported by Grether, Schwartz and Wilde. In addition to using the same participants for all sessions (as mentioned in the introduction), and posting the price distribution to buyers, these authors used thicker markets (5-8 sellers and 25 buyers); buyers each had constant marginal valuation for a single unit; sellers had U-shaped cost curves (induced via a fixed cost and a capacity constraint), and sellers could not stock-out in a period. Rather, sellers were obligated to satisfy all realized demand even if they had to incur the fixed production costs again to start a new production run in the same period.

[5] When buyers see the posted prices, it is a weak Nash equilibrium for each seller to offer 2 units at the competitive price.

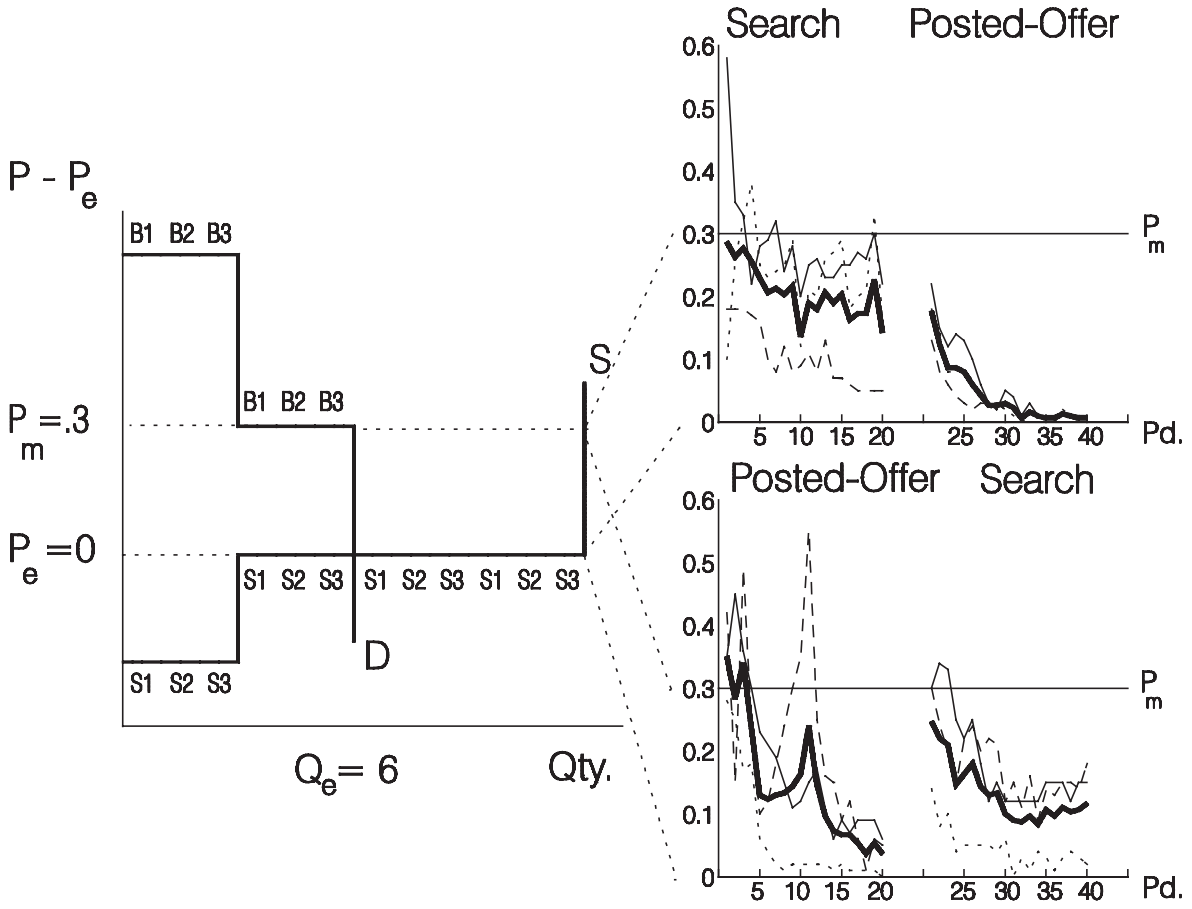

Figure 1. The Search design, and initial results. (Results key: —— indicates experienced sessions; ---, ··· indicate inexperienced sessions; ■ indicates the overall mean.)

This market design does not exactly implement any specific search model in the literature. steps were added to give the competitive and monopoly predictions a realistic chance of being observed, since stable outcomes in which one side of the market earns nothing are rarely observed. The design is anchored on a minimum earnings of 25 cents per trader per period at each of these outcomes. For this reason, the cost step for the first unit is 25 cents below P_e. Similarly, the placement of buyer's first unit value at 40 cents above the monopoly price, P_m, guarantees each buyer a minimum earning of 25 cents at the monopoly outcome, after subtracting out the 15 cent shopping cost.

Subjects were University of Virginia students who were recruited from economics classes. Buyer and seller roles were determined by random draws. Then subjects were seated at visually

isolated personal computers. Instructions were presented on the displays as an experimenter read aloud from a monitor display.[6] After the initial 20-period treatment, supplemental instructions were read for the final 20-period sequence. The final period was not announced for either sequence.[7] Subjects were paid $6.00 for showing up, in addition to earnings from trading. Earnings averaged $24.10 per subject, and ranged from $12.50 to $35.25. Payments were made in private immediately after the session. Two of the six sessions were conducted with subjects who had previously participated in a laboratory posted-offer market session (but in a different design), while participants in the remaining four sessions had no previous experience in an economics laboratory experiment.

The configuration of treatments by session is summarized by the 3-part identifiers in column (1) of table 1. Each identifier consists of a two-letter prefix ('SP' or 'PS') to indicate whether the search or posted-offer sequence came first, followed by a number indicating the order in sequence, and an 'x' if the session used experienced participants. Thus, for example session PS3x in the third row to table 1 refers to the third session in the posted-offer/search sequence with experienced participants.

3. Results

Mean price paths are shown on the right side of figure 1, with results for the Search/Posted-Offer treatment sequence shown above those for the Posted-Offer/Search results. The bolded line in each panel represents the mean price path averaged across the three sessions for each treatment sequence. Mean price paths for the individual sessions are indicated by the light lines. In each panel, the light, solid line represents a session with experienced subjects.

Several conclusions are immediately apparent. First, consider mean prices in the Search and Posted-Offer sequences. As seen in figure 1, mean price deviations are higher in the search

---

[6] A printed version of instructions for the baseline posted-offer treatment are in Davis and Holt (1993, Appendix A4.2). The additional instructions that were read aloud to participants to create the search treatment are available on request.

[7] In retrospect, we might have announced the final period, since we are evaluating a static prediction. Uncertainty about the final period would at least implicitly induce discounting, and therefore might prompt price increases via the use of trigger strategies. The generally low prices in the Posted-Offer treatment suggest that tacit collusion was not a problem.

Table 1. Mean Price Deviations for the Final 5 Periods: Search and Posted Offer Treatments

	Search Sequence	Posted-Offer Sequence
(1)	(2)	(3)
Session	$P-P_e$	$P-P_e$
SP1	5	1
SP2	22	1
SP3x	26	1
AVG.	18	1
PS1	15	6
PS2	3	1
PS3x	14	8
AVG.	11	5

sequence of each session than in the corresponding posted-offer sequence. This result is also clear from a comparison of mean price deviations for the last 5 periods each session, shown in table 1.[8] Mean price deviations in the search sequence of each session (shown in column 2) are larger than the mean price deviations for the corresponding posted-offer sequence (shown in column 3). This outcome allows rejection of the null hypothesis of no treatment (public price-

[8] We confine our attention to the last 5 periods of each treatment sequence as a crude means of controlling for learning.

posting) effect at a 95% confidence level, using the nonparametric Wilcoxon signed-ranks test.[9] This observation supports our first conclusion:

Conclusion 1: *Search matters: nonpublic postings raises prices.*

Consider next the relationship of prices in the Search treatment to the monopoly prediction, P_m. When Search precedes Posted Offer (the left side of the upper panel), the bolded overall price path is roughly 2/3 of the distance from the competitive to the monopoly price. In contrast, overall average prices are only about 1/3 of that distance in the Search treatment when it follows Posted Offer (the right side of the lower panel). Moreover, even "near monopoly" prices are the exception rather than the rule. Consider, for example, the mean price deviations $P-P_e$ for these 6 sessions, listed in column (1) of table 1. Even deleting the most obvious outlier in each treatment sequence leaves mean prices more than a nickel below the monopoly prediction in the Search/Posted Offer sessions, and slightly more than 15 cents below from the monopoly prediction in the Posted-Offer/Search treatment. To summarize:

Conclusion 2: *"Diamond-Result" monopoly prices are not consistently observed in this environment.*

The heterogeneity of outcomes both within and across treatments indicates that factors other than the presence of public price information affect pricing behavior. Some of these factors may be largely procedural. An order-of-treatment effect is suggested by the mean price information in table 1. Average deviations are higher in sequences where the search sequence

[9] This test assumes independence of the treatment pairs. The test statistic is created as follows (see, e.g., Conover, 1980, p. 280). First, rank the sessions by the absolute values of the differences between the search and posted-offer treatments. Then sum the ranks for the sessions in which the treatment difference is positive. The null hypothesis of no treatment effect is rejected if the test statistic exceeds a critical value determined by the number of sessions. In this case, the null hypothesis can be rejected, as the test statistic of 21 exceeds the 95% critical value of 19. The intuition behind the reported confidence level is easy to motivate for the special case observed in the table, where prices are higher under the search treatment in all six sessions. This outcome is analogous to getting six heads in a row from a series of coin tosses. If the coin were fair (e.g., if there was no treatment effect) this event would occur with a probability of one half raised to the sixth power, or about .03.

occurred first (18 cents vs. 11 cents), but the difference is not significant.[10] Similarly, there may be an experience effect, since the lowest two price deviations were generated in sessions where participants were inexperienced.

The variability of performance within treatments suggests that differences in participant behavior may be important. Sellers, for example, may vary substantially in their appreciation of the recursive, "price-plus-shopping-cost" reasoning that generates Diamond's monopoly pricing prediction. Buyers, on the other hand, may be able to force prices down by "punishing" sellers. As suggested by Bagwell and Ramey, buyers may either refuse to purchase if a posted price is unacceptably high, or may purchase, but shop elsewhere in other periods.

There is at least anecdotal evidence that buyers employed punishment behavior of this sort. Consider for example, the sequence of contracts for the search sequence of session PS3x in figure 2. Data for the 20 trading periods are separated by vertical bars. Within periods, price postings for sellers S1, S2 and S3 are represented, in respective order, by crosses (+). Contracts for single units are denoted by small dots (▪); multiple units sold at the same price show as an overlap of dots to the right of the price postings. Thus, for example, period 1 of session PS3x is illustrated between the left-most pair of vertical bars in figure 2. In this period, sellers S1 and S2 posted prices 5 cents below P_m, and S3 posted a price 20 cents above P_m. Sellers S1 and S2 subsequently sold 2 units each, while S3 sold a single unit.

In sequence PS3x, seller S2 posted prices well above P_m in periods 2 to 5, and posted prices at P_m in periods 7 and 10. Although this seller occasionally made sales at high prices (e.g,. in periods 2, 3 and 7), the buyers became wary of these high postings, and failed to approach this seller in periods 10 through 17, this despite successive price decreases in periods 11, 14 and 17 (where price was lowered to the cost of the low cost unit). Although punishment was not immediate in this session, seller S2 clearly paid for high price postings in the early periods of the session.

Other instances of punishment were observed in the other sessions, albeit less dramatic. It is difficult to develop a reasonable statistic for switching away in response to "high" prices,

---

[10] Using the Mann-Whitney test, the null hypothesis of no order-of-sequence effect can be rejected at only an 80% confidence level. Intuitively, there are $\binom{6}{3}=20$ possible ways that average price deviations for the 3 SP and the 3 PS sessions could be ranked, and there are 4 outcomes which generate a smaller sum of ranks than that observed.

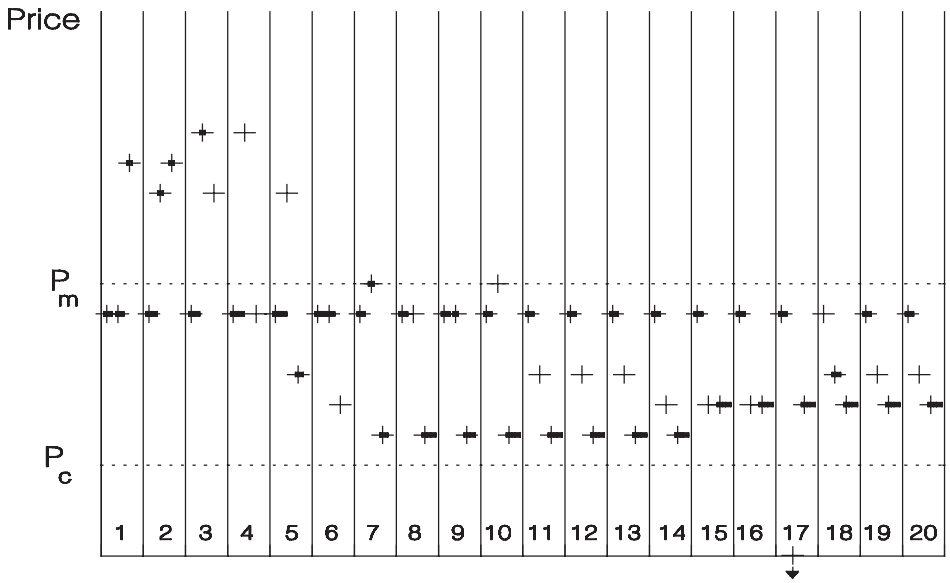

Figure 2. Prices and Contracts for the Search Treatment Periods of Session PS3x. (Key: + indicates price postings, ▪ indicates contracts.)

since, given shopping costs, a "high" price in a given period is a subjective assessment on the part of the buyer. In figure 2, for example one buyer (B1) consistently purchased units from S1, despite the fact that this seller posted the highest price in periods 11 to 20.

Nevertheless the effects of reputation may be evaluated by examining behavior in additional laboratory sessions where reputations are controlled. A primary goal of the research described in the next section, is to isolate the effects of sellers' reputations.

4. Search Costs and Reputations

Although the absence of price information tends to raise prices in an environment where shopping is costly, prices are not raised to monopoly levels as implied by the Diamond paradox. This raises the question of whether there is a reasonable baseline condition in which the absence of public price information clearly generates monopoly prices. Toward identifying such a baseline, we conducted two additional 3-session treatments. In each case, we attempted to control for reputations by disguising seller identities. To examine the comparative-statics effects of a search cost increase, the second treatment differs from the first in that shopping costs are doubled. These sessions are summarized by the identifiers listed in columns (1) and (3) of table 2. These identifiers roughly follow the labeling convention in table 1: The two-letter "SN"

prefix indicates that price information was not publicly displayed (e.g. 'search'), and that sellers reputations were disguised ('no reputations'). This prefix is followed by an L (search cost = 15 cents) or H (search cost = 30 cents), a number in sequence, and an 'x' if experienced participants were used.

Table 2. Mean Price Deviations for the Final 5 Periods: Search with No Reputations

Search, No Reputations, c=15		Search, No Reputations, c=30	
(1)	(2)	(3)	(4)
Session	$P-P_e$	Session	$P-P_e$
SNL1	18	SNH1x	26
SNL2	6	SNH2x	31
SNL3x	15	SNH3x	21
AVG.	13	AVG.	26

Controlling for Reputations: The most important difference between the Search treatment described in the previous section and the SNL treatment listed in columns (1) and (2) of table 2 is that seller identities were disguised so that buyers would be unable to determine which sellers posted which prices. Higher prices in this treatment than in the preceding treatments would indicate that reputation effects in fact tend to lower prices.[11]

To disguise seller identities, the following procedures were used. At the beginning of the session, sellers were visually isolated from the buyers and were given a colored marble "identifier." Prior to each period a monitor drew the marbles in sequence from the urn and

[11] Although this treatment may be expected to shed some light on the importance of seller reputations on pricing outcomes, it should not be interpreted as a test of the Bagwell-Ramey model. An exact implementation of the Bagwell-Ramey model would require too many design alterations to allow comparison of results with our existing sessions (in particular, sessions must be indefinitely repeated).

assigned the role of seller S1 to the first marble drawn, seller S2 to the second marble, and S3 to the third marble. The sellers then took their seats and posted prices. Further, to prevent sellers from divulging their identity via very rapid or very slow price postings, the monitor made the final price confirmation for each seller, once all sellers had finished posting prices. In this way, the posting sequence was terminated in a pre-announced, anonymity-preserving fashion.

In an effort to give monopoly price outcomes a reasonable chance, we made two additional procedural changes: First, given the likelihood of sequencing effects, we did not precede any of the sessions with a posted-offer sequence.[12] Second, we decided to let sellers see each other's price postings. While the relevant theory is silent on the matter of the amount of price information available to sellers, we decided to give this information to sellers, in an effort to facilitate learning of the recursive "price-plus-shopping-cost" reasoning that goes into generating Diamond results. In all other respects procedures were identical to our initial search treatments: Shopping costs were 15 cents per approach, each sequence consisted of 20 periods, and the experience profile consisted of two sessions with inexperienced participants, and a single experienced session.

As in the preceding section, results are obvious, and follow almost without comment from the mean contract price sequences shown in the left panel of figure 3 (formatted as figure 1). As indicated by the bolded overall mean contract price path lying roughly halfway between P_m and P_e, mean prices are no higher when reputations are disguised. In fact, comparing price data in sessions SNL1-3x and SP1-3x, one can see that mean prices were actually 5 cents lower on average in the "no-reputation" sessions than in the reputation sessions. Although results are too mixed to allow us to make a reasonable statistical claim that prices are actually higher when seller identities are not hidden, the results provide absolutely no support for the hypothesis that disguising information tends to raise prices. This motivates our third observation:

Conclusion 3: *Seller reputations do not lower prices. Prices are no closer to the monopoly*

[12] Actually, we decided to focus on the effects of reputation in the initial session after observing almost perfectly competitive prices in the "search" sequence of a "no-reputation" pilot session where the search sequence followed a posted-offer sequence. The reported sequences were followed with a variety of pilot treatments, none of which merit comment here.

prediction when seller identities are disguised than when prices are publicly displayed.

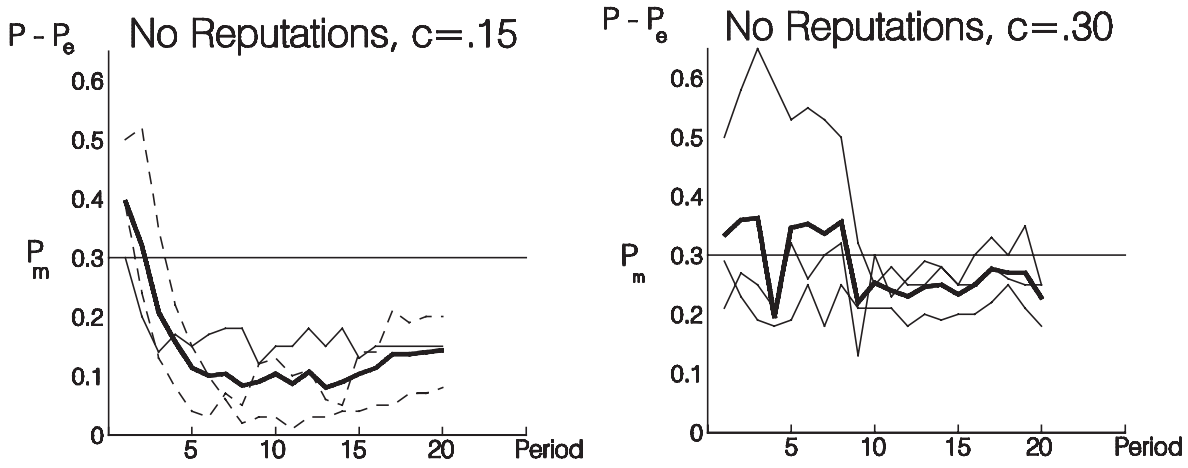

Figure 3. Mean Contract Price Sequences for the No-reputation Sessions (Key: —— indicates experienced participants; --- indicates inexperienced participants; ■ indicates the overall mean.)

The failure to observe monopoly price outcomes, even when seller identities are hidden and when sellers can see the prices posted by others, suggests that sellers fail to appreciate the recursive "price-plus-search-cost" reasoning underlying the Diamond prediction. Nevertheless, search costs provide sellers with some pricing discretion, much in the manner of switching costs or other entry impediments. To the extent this is true,: The magnitude of the impediment determines the extent of sellers' price discretion. As a final treatment, we examine this prediction by increasing search costs.

Search Costs, Comparative Statics Effects: Our final treatment consists of 3 sessions conducted as in the no-reputation treatment, except that switching costs were increased to 30 cents. With minor exceptions, other procedures were exactly as described in the reputation control sessions.[13]

Results of these "SNH" sessions are summarized in the right panel of figure 3, and in

[13] There were two exceptions. First, our experience profile consisted entirely of experienced participants. Second, buyer earnings were supplemented with a private, one-time $5 payment in the middle of each session. We were prompted to provide this supplement while observing very low buyer earnings in session SNH1x, and which threatened to diminish buyer interest.

column (4) of table 2. A comparison of the left and right panels of figure 3 indicates that increases in search costs raise prices. In the high-search-cost sessions, the mean price deviation is 26 cents above P_e, twice the 13 cent deviation from P_e observed in the low search cost sessions. Further, the mean price in each SNH session is higher than the highest price in the SNL treatment. The uniformity of results allows rejection of the null hypothesis that increasing search costs does not affect prices at a 95% confidence level using the nonparametric Mann-Whitney test.[14] This leads to our fourth, and final conclusion.

Conclusion 4: *Other things constant, increases in search costs tend to increase prices.*

5. Discussion

This paper has examined the results of 12 market sessions designed to evaluate the behavioral robustness of the Diamond prediction. The two primary lessons of this research are that (1) the absence of public price information raises prices when shopping is costly, but that (2) the monopoly prices implied by the "Diamond paradox" are not generally observed. Efforts to find a baseline treatment where monopoly price outcomes yielded two additional conclusions: (3) Controlling for seller identities is not sufficient to generate monopoly prices, and (4) increases in the search costs will raise prices.

In an important sense, the results of our initial sessions are consistent with the intuition motivating Bagwell and Ramey's efforts to resolve the Diamond paradox. The implied thrust of their work is that the "Diamond Paradox" is too extreme, and that in richer contexts non-monopoly prices will be observed. But equally important, results of our latter session are also inconsistent with the reasoning underlying Bagwell and Ramey's theoretical modifications. We observe non-monopoly prices, even when reputations cannot form. What then causes below-monopoly prices? We conjecture that, as a behavioral matter, the Diamond result breaks down because sellers neither immediately appreciate nor do they readily learn the recursive "price-plus-shopping-cost" reasoning necessary to generate the monopoly outcome. Rather, it appears that

[14] Intuitively, of the $\binom{6}{3}=20$ possible ways that average price deviations for the 3 SP and the 3 PS sessions could be ranked, the most extreme was observed (all SNH session above all SNL session). Under the null hypothesis of no treatment effect, this would occur 1 time out of 20.

search costs impede market performance much in the same way as other barriers to entry: Search costs make it difficult to buyers to leave a seller. Thus, the higher the search cost, the more market power sellers possess.

Our future research in this area will proceed in two directions. The first involves examining just where and how the Diamond prediction breaks down. It is possible that there are some contexts in which sellers do come to recognize the recursive reasoning necessary to generate monopoly outcomes. Obvious steps in this direction are to evaluate the differences between our treatments and the Grether, Schwartz and Wilde design (where monopoly prices were observed). Possible alterations include increasing the number of buyers and sellers, decreasing the capacity of buyers, and perhaps prohibiting seller stock-outs in a period.

The second direction for future research involves evaluation of predictions in the switching cost literature. Costs incurred when a buyer changes sellers may generate price outcomes consistent with those observed here, even when postings are public.[15]

[15] Klemperer (1992) provides a nice summary of issues in the switching cost literature.

References

Bagwell, Kyle and Gary Ramey (1992) "The Diamond Paradox: A Dynamic Resolution," Discussion Paper 1013, Northwestern University.

Butters, Gerald (1977) "Equilibrium Distributions, Sales and Advertising Prices," *Review of Economic Studies, 64*, 465-491.

Conover, W. J., *Practical Nonparametric Statistics*, John Wiley and Sons, New York, 1980.

Davis, Douglas D., and Charles A. Holt, (1993a) "Equilibrium Cooperation in Three-Person, Choice-of-Partner Games," *Games and Economic Behavior*, forthcoming.

Davis, Douglas D., and Charles A. Holt (1993b) *Experimental Economics*, Princeton: Princeton University Press.

Diamond, Peter A. (1971) "A Model of Price Adjustment," *Journal of Economic Theory, 3*, 158-168.

Grether, David M., Alan Schwartz, and Louis L. Wilde (1988) "Uncertainty and Shopping Behavior: An Experimental Analysis," *Review of Economic Studies, 55*, 323-342.

Hong, James T., and Charles R. Plott (1982) "Rate Filing Policies for Inland Water Transportation: An Experimental Approach," *The Bell Journal of Economics, 13*, 1-19.

Klemperer, Paul (1992) "Competition when Consumers Have Switching Costs: An Overview," working paper, Oxford University.

Salop, Stephen and Joseph Stiglitz (1977) "Bargains and Ripoffs: A Model of Monopo-listically Competitive Price Dispersion," *Review of Economic Studies, 44*, 493-510.

Stahl, Dale O., (1989) "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review, 79*, 700-712.

Wilde, Louis and Alan Schwartz (1979) "Equilibrium Comparison Shopping," *Review of Economic Studies, 46*, 543-53.