ADVERTISING AND CONSUMER SEARCH

by

Ninghua Du

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF ECONOMICS

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2005

UMI Number: 3164761

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

THE UNIVERSITY OF ARIZONA ®
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation

prepared by Ninghua Du

entitled Advertising and Consumer Search

and recommend that it be accepted as fulfilling the dissertation requirement for the

Degree of Doctor of Philosophy

James C. Cox Date

Martin Dufwenberg Date

Ronald L. Oaxaca Date

Price V. Fishback Date

Date

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Dissertation Director: James C. Cox Date 4/27/2005

STATEMENT BY AUTHOR

SIGNED: Ninghua Du

ACKNOWLEDGEMENTS

I would like to thank Jim Cox, Martin Dufwenberg, Ron Oaxaca and Stan Reynolds for their mentoring and extensive support. I also thank Greg Crawford, Price Fishback, David Reiley, Mark Walker, and John Wooders for their constructive comments and encouragement. Todd Swarthout and Richard Kiser have provided valuable advice on software programming. Financial Support from the Economic Science Laboratory of the University of Arizona is gratefully acknowledged.

Finally, I wish to thank my family for their constant support and my friends at the University of Arizona who kept me company during the past several years. Especially, thanks to Yanwei for her care and understanding.

DEDICATION

To My Parents, Xiao Yu and Du Lian

TABLE OF CONTENTS

TABLE OF CONTENTS – *Continued*

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

This dissertation advances our understanding of interaction between advertising and consumer search. If advertising lowers consumer search costs, it can affect competition. Previous studies by Butters (1977) and Robert and Stahl (1993) show that giving sellers the option of price advertising can significantly lower equilibrium market prices. These models assume that sellers make two bundled decisions: sellers determine the proportion of buyers that receive advertisements (ads) and reveal the price that they intend to charge in such ads. However, the vast majority of advertising does not reveal product pricing. Chapter 1 argues that certain types of advertising may reduce consumer search costs without actually mentioning the price in the message. This leads me to propose a model in which firms first decide whether to advertise, and then set prices. In this model, the equilibrium price with advertising returns to the monopoly level.

Chapter 2 provides an experimental test of the theoretical model in the previous chapter. In the laboratory sessions, each period human sellers make two decisions: what price to set, and whether to advertise to eliminate consumer search costs for their product. Robot buyers then follow an optimal search rule (known to all sellers) to decide which price offer (if any) to accept. The two experimental conditions are (1) advertising the price, or (2) advertising before pricing. Data from ten sessions indicate that, as predicted, firms choose more often to advertise when advertising conveys price, and prices in the second treatment are significantly higher than prices in the first treatment.

Chapter 3 uses patterns of prices for books to investigate the empirical validity of the theoretical predictions that as the consumer search cost for a homogenous good goes

up, both the market price and the degree of price dispersion for that good will go up (Butters 1977, Robert and Stahl 1993). Search costs for books sold in retail bookstores are presumably higher than search costs for books sold online. By comparing prices on the internet and prices in retail bookstores for the same book, this study finds that the internet has a negative impact on average prices. This finding is consistent with the theoretical predictions. However, the empirical evidence shows that the internet does not lower the variance of prices. Several possible explanations for this observation are provided.

Chapter 1

ON ADVERTISING AND PRICE COMPETITION

1.1 Introduction

Consumer search costs and the availability of sellers' advertising choices are key factors that influence how a market works. The classical model of Bertrand competition where all sellers have the same constant marginal cost may be viewed as a simplified case where neither consumer search cost nor advertising plays a role. Under Bertrand competition, buyers have perfect information and take the lowest price in the market. Since the seller who offers the lowest price takes the whole market, sellers have an incentive to undercut their rivals and the market outcome is marginal cost pricing.

However, if one relaxes the assumption that consumers have perfect information about prices, price competition will be imperfect and the market share of the high-price firm may be nonzero. One extreme example is the price adjustment procedure studied by Peter Diamond (1971). Given positive consumer search costs (uniformly bounded below by an arbitrary small positive number), Diamond shows that the unique market equilibrium is monopoly pricing. The intuition is that when all buyers have positive search costs for prices, buyers will not switch to other sellers when they see a price that is slightly higher than other prices; when other sellers are charging a price below the monopoly price, choosing a slightly higher price will strictly improve one seller's profit.

One way to undermine the monopoly equilibrium derived by Diamond is to introduce price advertising. When advertising is available to sellers, sellers have an incentive to advertise their prices in order to increase their market shares. This way, buyers observe the advertised prices for free and buyers can switch to the advertised offers without paying a search cost. The studies by Butters (1977) and Robert and Stahl (1993) show that as the advertising costs go up, the degree of price dispersion (i.e., the variance of prices) in the market will go up. When advertising is free, the market outcome reverts to marginal cost pricing.

Both Butters and Robert & Stahl assume that sellers make two simultaneous decisions: sellers determine the proportion of buyers that receive advertisements (ads) and reveal the price that they intend to charge in such ads. Price advertising is modeled in a way that presumes that prices are contained in the ads. However, this is not necessarily the most reasonable way to model advertising. In fact, the vast majority of advertising does not reveal product pricing. Examples include TV advertising, newspaper advertising and ads in the Yellow Pages. Internet advertising is another example. Merchants use banners, pop-up windows and sponsored links to advertise their products. However, in most cases, the banners, pop-up windows and links themselves do not directly convey product prices to buyers. Even in the circumstance that ads contain prices, the nominal prices shown in the ads may not be informative. In many cases, prices are multi-dimensional. When people shop for cars and houses, they not only care about the nominal price, but also care about the warrantee, insurance and related fees. In order to obtain this important information, people need to visit car dealers or real estate agents instead of sitting at home reading newspaper ads column.

Without actually mentioning the price in the message, certain types of advertising still provide a way to reduce consumer search costs for advertisers' prices. An advertiser may leave its location and contact method in the ads. It is very common to see examples like the following in the Yellow Pages: "Insurance problems? Hassle free quotes! Call us first. We make house calls. Open seven days a week. Out of area call 1-800-xxx-xxxx." For internet advertising, by clicking the banners, pop-up windows and sponsored links, potential buyers are led to the advertiser's homepage that provide information about the product details, including prices. This way, consumers can easily find out product prices by several phone calls or several clicks. Ads substantially reduce consumers' time and effort during the search process.

For these important reasons, I propose and examine the impact of an alternative way to model advertising. Given that ads reduce consumer search costs without mentioning prices, it is plausible to think of sellers as making two decisions sequentially. First, sellers decide whether to advertise. Second, after observing rivals' advertising choices, sellers choose their prices. The first stage is that advertising a firm's existence eliminates the search cost for a consumer to shop at that firm and the firms choose prices after observing all other firms' advertising decisions. I show that, in contrast to the results by Butters and Robert & Stahl, the market outcome under the proposed Advertise-then-Price assumption reverts to monopoly pricing. [1]

1.2 The Backdrop: Two Previous Results of Search and Advertising

Diamond shows that in a market thatconsumers must search sequentially with a positive search cost, in finite time the prevailing price becomes that which maximizes firms' joint profits. However, one can avoid this surprising result by altering the assumptions. When consumers have other sources of free information, the monopoly equilibrium in the Diamond model may be undermined. Butters and Robert & Stahl study the interaction between advertising and consumer search and show that the market equilibrium goes back to marginal cost pricing when advertising is free.

In this section I use two one-shot games to recapture the main results by Diamond and by Butters and Robert & Stahl. The search cost game is an extension of a simple Bertrand pricing game, and the price advertising game is an extension of the search cost game. First, these two games serve as the backdrop for the advertise-then-price game in the next section, and allow me to compare the market outcome of the advertise-then-price game with the market outcomes from the previous studies under the same structure of modeling. Secondly, reinterpreting the results from the previous studies under the same game theoretic framework itself enriches the literature.

To make the games simple and illustrative, I assume there are two sellers, identified as 1 and 2, and each of them produces a homogeneous good with zero cost of production. The sellers' goal is to maximize profits. There is a population of identical buyers, who have reservation value v for one unit of the good. The buyers' goal is to maximize consumer surplus. I assume there is no resale among buyers. Therefore, the

[1] In my model, I assume that ads do not contain prices. However, one may wonder what will happen if

buyers' optimal strategies are identical and I can further assume there is only one buyer that represents the population. Sellers' costs and the buyer's reservation value are common knowledge. All of the players are risk neutral. These assumptions are kept throughout this section.

1.2.1 The Search Cost Game

When searching for new prices is costly for consumers, the market outcome moves from marginal cost pricing in Bertrand competition to monopoly pricing. Given positive search cost and the buyer's equilibrium belief that the unobserved price is not lower than the observed price, the buyer will not search for the unobserved price. In the search cost game, the buyer has search cost $c > 0$ and the value of this parameter is common knowledge. Assume $c < v$.

The game proceeds in two stages. First, sellers 1 and 2 simultaneously choose their prices P_1 and P_2. Then the buyer is randomly matched with one of the sellers and observes that seller's price for free. With probability of 0.5 the buyer observes seller 1's price for free, and the buyer observes seller 2's price for freewith the same probability. [2]

In the second stage, it is the buyer's turn to move. The buyer has three choices. She can take the observed offer, quit the market, or search for a new price. If the buyer takes the offer, then the buyer gets $v - P_i$ and seller i gets P_i. If the buyer quits, all players get zero. If the buyer decides to search, then the buyer pays search cost c and observes

sellers can *choose* whether to put prices in the ads. In section IV I show that giving sellers the option of choosing the format of advertising will not change the main conclusion of this chapter.

the other seller's price. Now the buyer can either take the offer P_1 or P_2, or quit the market.

It is natural to restrict attention to price offers that do not give the player losses if accepted, so I will assume that sellers make price offers $P_i \in [0, v]$, i=1, 2. It is further assumed that the buyer accepts the lower price if the buyer knows both prices. Then the buyer's strategy can be characterized by decision d_1 at the price offer observed for free. d_1 is a function that maps the offer observed for free to an action, d_1: [0, v]→{A, S, Q}, where A refers to accept, S refers to search and Q refers to quit. In order to guarantee the existence of Nash equilibrium, I will finally assume that the buyer always chooses A when the buyer is indifferent between A and S or the buyer is indifferent between A and Q. The following diagram illustrates the buyer's moves in the second stage given the buyer is matched with seller 1 and observes P_1 in the first stage.

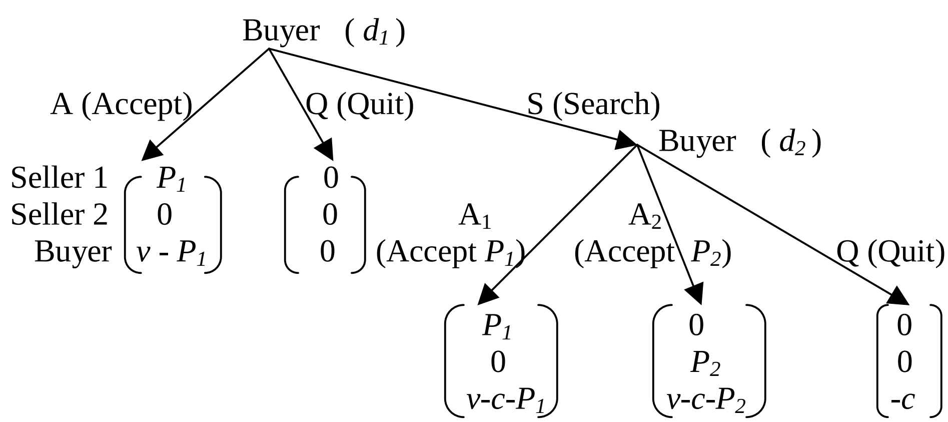

FIGURE 1.1. The Buyer's Actions and the Terminal Payoffs

---

[2] The price observed for free can be interpreted as the buyer's local store price. The buyer observes each seller's price for free with probability of 0.5 means that the population of buyers is evenly distributed between two sellers.

Since the buyer cannot observe the sellers' identities, at the decision node d_1 the buyer must have a unique cutoff price r for both seller 1 and 2's offers. At d_1, the buyer accepts the observed offer p when $p < r$ and the buyer searches for the unobserved offer when $p > r$. When $p = r$, the buyer is indifferent between A and S and the buyer chooses A according to the assumption stated previously. A Nash equilibrium in this game must be of the form "$P_1 = \bar{p}$, $P_2 = \bar{p}$, the buyer chooses cutoff price $\bar{p}$." It can be easily verified that in this game there are infinitely many Nash equilibria: for any $\bar{p} \in [0, v]$, the strategy profile "both sellers offer $\bar{p}$, the buyer chooses cutoff price $\bar{p}$" is a Nash equilibrium.

However, any cutoff price $\bar{p} < v$ is not a credible threat, since it is not consistent with the buyer's belief along the equilibrium path. The reasoning is as follows. Consider the Nash equilibrium "The buyer chooses cutoff price $\bar{p} < v$, both sellers offer $\bar{p}$." Suppose the buyer is matched with seller 1 and observes P_1 for free. On the equilibrium path, the buyer believes $P_2 = \bar{p}$. Then for any ex-post observation $\bar{p} < P_1 < \min(v, \bar{p} + c)$, the option 'search for P_2' is not optimal for the buyer, since the buyer's payoff by accepting P_1 is $v - P_1$, and the expected payoff by search is $v - \bar{p} - c$, which is less than $v - P_1$.

In order to cater to this anomaly, one needs to employ a refinement of Nash equilibrium. Burdett and Judd (1983) were the first to provide a game theoretic interpretation of Diamond's model and they defined a "search equilibrium," which has become the standard solution concept in the literature of consumer search and price dispersion. [3] Before introducing the concept of search equilibrium, it is helpful to first characterize the buyer's cutoff price that is consistent with her belief about the unobserved offer.

Observation 1.1 Suppose the buyer observes offer p for free and suppose the buyer's belief about the unobserved offer is a distribution of the form $F(q)$, $q \in [0, v]$. [4] Then the buyer's unique cutoff price r that is consistent with $F(q)$ is given by the equation $c = \int_0^r F(p)dp$.

Proof: Since the buyer takes the lower offer when she observes both 1 and 2's prices, the buyer's expected payment conditional on search is $E = \int_0^p qdF(q) + p\int_p^v dF(q)$. Therefore, the buyer accepts the observed offer when buyer's payoff by accepting offer p, which is v - p, is higher than buyer's expected payoff from search, which is v - c - E. This condition can be written as follows:

$$d_1 = \begin{cases} A & if \quad v - p > v - c - E \\ S & if \quad v - c - E > v - p \end{cases} \qquad \text{where } E = \int_0^p qdF(q) + p\int_p^v dF(q) \text{ or}$$

$$d_1 = \begin{cases} A & if \quad p - \int_0^p qdF(q) - p\int_p^v dF(q) < c \\ S & if \quad p - \int_0^p qdF(q) - p\int_p^v dF(q) > c \end{cases}$$

Notice that $p - \int_0^p qdF(q) - p\int_p^v dF(q) = \int_0^p (p-q)dF(q) = \int_0^p F(q)dq$.

---

[3] Examples include Stahl (1989), Cason and Friedman (2003) and Morgan, Orzen and Sefton (2001).

[4] In this section, the sellers do not mix their price choices. However, the notion of mixed strategy is useful since observation 1.1 will be applied to later sections, which involve mixed strategies.

The expression $\int_0^p F(q)dq$ is continuous in *p*, and strictly increasing on the support of *F*. The expression $\int_0^p F(q)dq$ is zero at $p = 0$, and is unbounded as $p \to \infty$. According to the intermediate value theorem, the equation $c = \int_0^r F(p)dp$ has unique solution *r*, since the search cost *c* is positive. Now we can rewrite the buyer's decision at the offer observed for free in the following way:

$d_1 = \begin{cases} A & if \quad p < r \\ S & if \quad p > r \end{cases}$ *r* is determined by $c = \int_0^r F(p)dp$.

Notice that $r = q + c$ when $F(p) = 0$ for $p < q$ and $F(p) = 1$ for $p \geq q$. Since $P_i \in [0, v]$ for $i=1, 2$, setting the cutoff price at $\min(v, r)$ represents the same buyer strategy as setting the cutoff price at *r*.

Q. E. D.

Given the buyer's cutoff price that is consistent with her belief on the unobserved offer, the search equilibrium that applies to the search cost game is stated as follows.

Definition The strategy profile (P_1^*, P_2^*, r^*) is a *search equilibrium* if:

(1) For sellers, $\pi_i (P_i^*, P_j^*, r^*) \geq \pi_i (P_i, P_j^*, r^*)$ for $\forall P_i$, $i = 1, 2$, $i \neq j$, where π_i is seller *i*'s payoff.

(2) For the buyer,

$$d_1 = \begin{cases} A & if \quad p < r^* \\ S & if \quad p > r^* \end{cases} \quad r^* = \min(v, r),$$

Where p is the price observed for free, r is determined by $c = \int_0^r F(q)dq$ and F is buyer's belief about the unobserved price.

(3) In equilibrium, the buyer's belief is consistent with the equilibrium prices.

If seller i's price is not observed by the buyer, then F(q) =0 for $q < P_i^*$ and F(q) =1 for $q \geq P_i^*$.

In the definition, conditions (1) and (2) state that a search equilibrium must be a Nash equilibrium. Condition (3) imposes a restriction on the buyer's belief on the unobserved price and this restriction is not required in Nash equilibrium. Therefore, the search equilibrium is a refinement.

Observation 1.2 The unique search equilibrium is ($P_1^* = P_2^* = v$, $r^* = v$). Both sellers get expected payoff 0.5v, and the buyer gets zero payoff.

Proof: 1) The Nash equilibrium "The buyer chooses cutoff price $\overline{p} < v$, both sellers offer $\overline{p}$" does not satisfy the definition of search equilibrium. According to the definition of search equilibrium, given both sellers offer $\overline{p}$ in equilibrium, the buyer's equilibrium cutoff price is min(v, $\overline{p}$ +c) > $\overline{p}$. This contradicts that the buyer chooses cutoff price $\overline{p}$ in equilibrium.

2) The strategy profile "The buyer chooses cutoff price $\overline{p} = v$, both sellers offer $\overline{p}$ " satisfies the definition of search equilibrium.

By 1) and 2), the unique search equilibrium is ($P_1^* = P_2^* = v$, $r^* = v$). *Q. E. D.*

Observation 1.2 indicates that the unique market outcome that satisfies the definition of search equilibrium is monopoly pricing. Two sellers split the market. The intuition is that choosing a slightly higher price strictly improves one seller's profit when its rival's price is less than v, since the buyer always accepts a slightly higher price offer given her positive search cost and market prices less than v.

1.2.2 The Price Advertising Game

Now let us introduce sellers' advertising options into the search cost game. When advertising is available to sellers, the monopoly outcome in the previous buyer search model will be undermined. Since the buyer observes the advertiser's price for free, the sellers have an incentive to advertise their prices in order to increase their market shares. In the price advertising game, both sellers can advertise their prices. The whole buyer population observes the advertisers' prices without paying a search cost. Advertising costs sellers nothing.

The timing of the game is as follows.

Stage 1: Each of the sellers makes two bundled decisions: 1) choose product price and 2) choose whether or not to reveal this price to the buyer in the ads. The sellers make their decisions simultaneously. Then the buyer population is evenly allocated between two sellers and the buyer observes one of the sellers' prices for free, as described in the search cost game.

Stage 2: The buyer then observes all of the advertised prices, and makes her decision. To illustrate the game clearly, let us discuss the three possible cases that can occur.

Case I: The buyer receives two ads. In this case the buyer gets full market price information: she observes both prices offered on the market, and she can either take any one of the offers without paying a search cost, or quit from the market. If the buyer takes seller i's offer, then the buyer gets $v - P_i$ and seller i gets P_i. If the buyer quits, all players get zero.

Case II: The buyer receives exactly one ad. Without loss of generality, suppose it is from seller 1. If the buyer is matched with seller 2 in stage 1, then the buyer observes both prices and it degenerates to the full information case. If the buyer is matched with 1 in stage 1, then the buyer cannot observe 2's price. The buyer can take 1's offer, quit from the market, or pay a search cost c to search for 2's price. In this case, the description of the buyer's moves is the same as in the search cost model.

Case III: The buyer receives no ad. The buyer's moves degenerates to that of the search cost game: no advertising, the buyer pays a search cost to observe a new price. For detailed descriptions see the search cost model, the timing of the game, Stage 2.

The sellers' strategies are characterized by (P_i, a_i), $P_i \in [0, v]$, $a_i \in \{\text{Ad, No Ad}\}$, where Ad means 'to advertise' and No Ad means 'not to advertise'. I assume that the buyer accepts the lower observed offer when she gets full price information, and the buyer accepts one of the two offers with probability of 0.5 when she observes a tie. Then the buyer's strategy is given by decision d_1 at the price offer observed for free. d_1: $((P_1, a_1), (P_2, a_2)) \rightarrow \{A, S, Q\}$. A refers to accepting the lower observed offer, S refers to search and Q refers to quit.

Given the specification on the player's strategies, we have the following observation.

Observation 1.3 In Nash equilibrium, both sellers must choose ($P_i = 0$, Ad), $i = 1, 2$.

Proof: Let us first show $P_i > 0$, $i = 1, 2$ cannot be part of a strategy in Nash equilibrium.

Suppose seller 1 chooses ($P_1 > 0$, Ad). Then choosing $P_2 < P_1$ cannot be part of seller 2's best response. By choosing $P_2 < P_1$, seller 2's payoff must be less than or equal to P_2. However, if seller 2 chooses ($\tilde{p}$, Ad), where $P_2 < \tilde{p} < P_1$, then seller 2 will get payoff $\tilde{p}$ which is strictly higher than P_2. Choosing $P_2 = P_1$ cannot be part of seller 2's best response. By choosing $P_2 = P_1$, seller 2's payoff must be less than or equal to $0.5P_2$. If seller 2 chooses ($0.75P_2$, Ad), then seller 2 will get payoff $0.75P_2$, which is strictly higher than $0.5P_2$. Finally, choosing $P_2 > P_1$ cannot be part of seller 2's best response since choosing $P_2 > P_1$ gives seller 2 zero payoff. In summary, ($P_1 > 0$, Ad) cannot be a strategy in Nash equilibrium. Similar arguments apply to ($P_1 > 0$, No Ad).

Secondly, ($P_i = 0$, No Ad), $i = 1, 2$ cannot be a strategy in Nash equilibrium.

Suppose seller 1 chooses (P_1 = 0, No Ad). Then choosing (0 < P_2 < c, Ad) gives seller 2 positive payoff, which is higher than zero payoff from P_2 = 0. Therefore P_2 > 0, and choosing P_1 = 0 cannot be part of seller 1's best response.

Finally, it can be easily verified that (P_i = 0, Ad), i =1, 2, and the buyer accepts one of the offers with probability 0.5 is a Nash equilibrium.

Q. E. D.

Observation 1.3 indicates that the market outcome becomes marginal cost pricing when ads effectively reach the whole buyer population and the ads cost nothing to advertisers. The intuition is that when other sellers are charging a price above the marginal cost, one seller has an incentive to advertise a slightly lower price in order to take the whole market. This result is consistent with the previous literature.

1.3 Advertise-then-Price

Finally, it is time to introduce the Advertise-then-Price game. Consider the following advertising format: ads do not directly contain prices; meanwhile, ads still remove the consumer search cost for prices. The ads are considered "contact method advertising" rather than direct price advertising. In pre-stage, sellers decide whether to advertise the search methods for their prices. In the second stage, the sellers' advertising strategies are identified to the public and the sellers choose their prices. Finally it is the buyer's turn to accept one of the observed prices, search for a new price or quit. The timing of the game is as follows.

Pre-stage: Sellers simultaneously decide whether to send ads to the buyer.

Stage 1: Both sellers observe their rival's pre-stage decision. Sellers simultaneously choose their prices P_1 and P_2. The buyer population is evenly distributed between two sellers and the buyer observes one seller's price for free, as described in the search cost model.

Stage 2: Now it is the buyer's turn to move. The description of the buyer's moves is exactly the same as that of the price advertising model.

The restrictions on players' strategy sets are the same as that of the price advertising model.

The game tree for *Advertise-then-Price* is as follows:

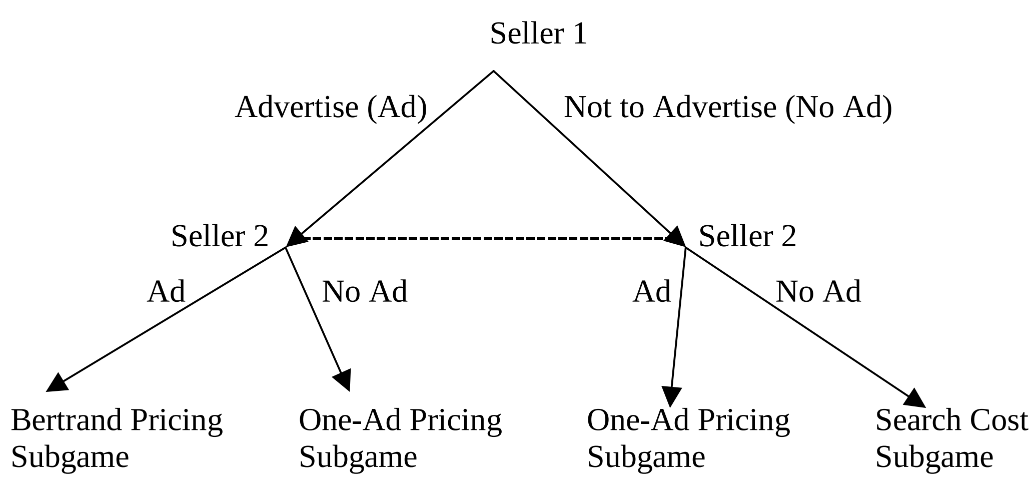

Note: the dottedine indicates seller 2 's information set.

FIGURE 1.2. Game Tree for *Advertise-then-Price*

As shown above, advertising choices are realized before sellers choose prices. Therefore we can break the original game into an "advertising game" and a "pricing game"—given every seller's advertising choice, there is a corresponding separable pricing subgame. What we need to do is to discuss three possible cases.

Case I: both of the sellers choose 'not to advertise.'

In this case the pricing subgame degenerates to the search cost model. As shown in Figure 2, this scenario is the "search cost subgame." We use the unique search equilibrium as our solution—that is, P_1*=P_2*= v, the buyer sets the cutoff price at v. Both sellers get the expected payoff 0.5v, and the buyer gets payoff zero.

Case II: both of the sellers choose 'to advertise.'

In this case, the pricing subgame degenerates to the Bertrand competition. As shown in Figure 2, this scenario is the "Bertrand pricing subgame." In equilibrium, P_1*=P_2*=0, and the buyer takes the offer. Both sellers get payoff zero, and the buyer gets payoff v.

Case III: One seller chooses 'to advertise,' but the other seller chooses 'not to advertise.'

As shown in Figure 2, this scenario is the "one-ad pricing subgame." Without loss of generality, let us assume that seller 1 chooses 'to advertise,' but seller 2 chooses 'not to advertise.' In this case, the buyer always observes P_1. However, the buyer observes 2's price with probability 0.5, since with chance of 50/50 the buyer is matched with 2. The buyer makes choices after P_1 and P_2 are chosen. Again, I restrict attention to price offers $P_i \in [0, v]$, i=1, 2. Given price P_1 and the buyer maximizes her payoff conditional on the price pair (P_1, P_2), seller 2's best response can be characterized as follows:

$$P_2 = \begin{cases} P_1 - \varepsilon & if \quad \frac{1}{2}(P_1 - \varepsilon) > (P_1 - c - \varepsilon) \\ P_1 - c - \varepsilon & if \quad \frac{1}{2}(P_1 - \varepsilon) < (P_1 - c - \varepsilon) \end{cases} \text{ or } P_2 = \begin{cases} P_1 - \varepsilon & if \quad P_1 < 2c \\ P_1 - c - \varepsilon & if \quad P_1 > 2c \end{cases}$$

Since the buyer always observes P_1, seller 2 always has an incentive to undercut. If seller 2 chooses a price which is higher than P_1, seller 2 will never make a sale. If seller 2 undercuts 1 by a sufficiently small amount ε, then the fully informed buyer will buy 2's product and seller 2 gets payoff 0.5(P_1 - ε). If seller 2 chooses to undercut by $c + \varepsilon$, then both the fully informed buyer and the buyer who does not observe P_2 will accept 2's offer and seller 2 gets payoff $P_1 - c - \varepsilon$. Therefore, when $0.5(P_1 - \varepsilon) > (P_1 - c - \varepsilon)$, or $P_1 < 2c$, seller 2 chooses to undercut by ε. When $0.5(P_1 - \varepsilon) < (P_1 - c - \varepsilon)$, or $P_1 > 2c$, seller 2 chooses to undercut by $c + \varepsilon$. When $P_1 = 2c$, seller 2 is indifferent.

Similarly, given price P_2 and the buyer maximizes her payoff conditional on the price pair (P_1, P_2), seller 1's best response can be characterized as follows:

$$P_1 = \begin{cases} P_2 - \varepsilon & if \quad (P_2 - \varepsilon) > \frac{1}{2}(P_2 + c - \varepsilon) \\ P_2 + c - \varepsilon & if \quad (P_2 - \varepsilon) < \frac{1}{2}(P_2 + c - \varepsilon) \end{cases} \text{ or } P_1 = \begin{cases} P_2 - \varepsilon & if \quad P_2 > c \\ P_2 + c - \varepsilon & if \quad P_2 < c \end{cases}$$

Seller 1 faces tradeoff between sales volume and markup. If seller 1 chooses to undercut 2 by a sufficiently small amount of ε, then seller 1 will take the whole market. However, to increase the price by $c - \varepsilon$ gives seller 1 higher profit margin without losing the market that the buyer does not observe P_2. Therefore, when $(P_2 - \varepsilon) > 0.5(P_2 + c - \varepsilon)$, or $P_2 > c$, seller 1 chooses to undercut by ε. When $(P_2 - \varepsilon) < 0.5(P_2 + c - \varepsilon)$, or $P_2 < c$, seller 1 chooses to raise the price by $c - \varepsilon$. When $P_2 = c$, seller 1 is indifferent.

Since seller 1's best response to any P_2 involves infinitely small quantity ε (same for seller 2's best response), there is no pure strategy Nash equilibrium in this pricing subgame. When P_1 and P_2 are higher than c, both seller 1 and seller 2 will undercut their rival. When the P_1 is at c, seller 2 will choose a price that is slightly lower than c. Given seller 2's choice, seller 1 will increase the price by the amount of $c - \varepsilon$. Because of seller 1 and 2's asymmetric positions, price dispersion exists in the one-ad pricing subgame.

I use the search equilibrium defined in section II as the solution concept for the one-ad pricing subgame. The equilibrium is described by the strategy profile $(F_1(P_1), F_2(P_2), r)$. $F_1(P_1)$ and $F_2(P_2)$ are seller 1 and seller 2's price distributions. r is the buyer's cutoff price when the buyer does not observe P_2. $P_i \in [P^i_L, P^i_H]$, $0 \leq P^i_L \leq P^i_H \leq v$, i=1, 2. $r = \min(v, r')$ where r' is determined by $c = \int_0^{r'} F_2(p)dp$. The buyer always accepts the lower price when she gets full price information. The unique search equilibrium in this one-ad pricing subgame is $F_1(p) = 1 - r/2p$, $p \in [0.5r, r)$, and $Prob(P_1 = r) = 0.5$ for seller 1; $F_2(p) = 2 - r/p$, $p \in [0.5r, r)$ for seller 2, and $r = \min(v, c/(1-ln2))$ for the buyer. Seller 1's expected payoff is $0.5r$ and seller 2's expected payoff is $0.25r$. The proof is in appendix A.

In the one-ad pricing subgame, both sellers get positive expected payoffs. The intuition is that when sellers choose different advertising strategies in the pre-stage, the advertiser keeps the market share from the buyer who does not observe the not advertised price and the seller who chooses not to advertise undercuts the advertiser in the fully informed market.

Now we have solved for the equilibrium in every single pricing subgame. Therefore, we can use the equilibrium market outcomes in each pricing subgame to substitute for the pricing subgames and simplify the Advertise-then-Price game as follows:

		Seller 2: Not to Advertise	Seller 2: To Advertise
Seller 1	Not to Advertise	$0.5v$, $0.5v$	$0.25\min(\frac{c}{1-\ln 2}, v)$, $0.5\min(\frac{c}{1-\ln 2}, v)$
	To Advertise	$0.5\min(\frac{c}{1-\ln 2}, v)$, $0.25\min(\frac{c}{1-\ln 2}, v)$	0, 0

TABLE 1.1. Payoffs for the Sellers, Advertise-then-Price Game

It is clear from the table that sellers' strategy 'not to advertise' weakly dominates the strategy 'to advertise.' When $c/(1-ln2) < v$, 'not to advertise' even strictly dominates the strategy 'to advertise.' The unique weakly dominant strategy equilibrium in this simplified game is (1 not to advertise, 2 not to advertise) and both sellers get payoff $0.5v$.

In contrast to the previous literature, the market outcome under Advertise-then-Price assumption reverts to monopoly pricing.

1.4 Variations

In this section, we briefly discuss three extensions of the Advertise-then-Price model to check the robustness of the result.

1.4.1 What will happen if the sellers can choose the format of advertising?

One may wonder what will happen if sellers can choose whether to put prices in the ads. In this subsection we show that giving sellers the option of choosing the format of advertising will not change the main conclusion of this paper. Suppose the procedure of the game is as follows.

Pre-stage: Sellers simultaneously decide whether to send ads to the buyer. If seller i decides to send ads, she can choose whether to put her price P_i in the ads.

Stage 1: Both sellers observe their rival's pre-stage decision. If seller i decides not to send ads or seller i decides to send ads without price in the pre-stage, then seller i will choose her price P_i in stage 1. The buyer population is evenly distributed between two sellers and the buyer observes one seller's price for free, as described in the search cost model in section II.

Stage 2: Now it is the buyer's turn to move. The description of the buyer's moves is exactly the same as that of the price advertising model in section II.

The restrictions and assumptions in the previous section apply to this game. Given every seller's pre-stage decision, there is a corresponding separable subgame in stage 1. The structure of this game is similar to that of the advertise-then-price game in the previous section. What we need to do is to discuss the new cases that at least one of the sellers decides to put her price in the ads. Without loss of generality, suppose in the pre-stage seller 1 decides to put her price P_1 in the ads.

Case I: Seller 2 decides not to advertise in the pre-stage.

In this case, seller 2 observes P_1 and undercuts seller 1. Following the discussion in the previous section, given positive P_1, seller 2's price is

$$P_2 = \begin{cases} P_1 - \varepsilon & if \quad P_1 < 2c \\ P_1 - c - \varepsilon & if \quad P_1 > 2c \end{cases}$$

In stage 2 the buyer observes P_1 for sure and observes P_2 with probability 0.5. When $P_1 > 2c$, seller 1 gets payoff zero and seller 2 gets payoff $P_1 - c - \varepsilon$. When $P_1 < 2c$, 1 gets $0.5P_1$ and 2 gets $0.5(P_1 - \varepsilon)$.

Case II: Seller 2 decides to send ads without price in the pre-stage.

In stage 1, seller 2 undercuts seller 1 by ε and takes the whole market. In stage 2 the buyer observes both P_1 and P_2. Seller 1 gets zero and seller 2 gets $P_1 - \varepsilon$.

Case III: In the pre-stage, seller 2 decides to put her price P_2 in the ads.

The prices are already chosen in the pre-stage and the buyer observes both P_1 and P_2. In this case no price will be chosen in stage 1. Play proceeds to stage 2 where the buyer will simply take the lower offer.

Combining with the analysis in the previous section, the market outcome in every subgame in stage 1 is unique. Now we can use this unique equilibrium market outcome in each subgame to substitute for the subgames in stage 1 and simplify the original game. In the simplified game, (1 not to advertise, 2 not to advertise) is a Nash equilibrium, which entails that both sellers subsequently choose the monopoly price in the search cost subgame as shown in section III. To see this, suppose seller 1 chooses 'not to advertise'. Then choosing 'not to advertise' gives seller 2 payoff 0.5v. Sending ads without price gives seller 2 0.5min(v, c/(1-ln2)). Seller 2's payoff can never exceed 0.5min(2c, v) when seller 2 decides to put P_2 in the ads, since $P_2 \leq v$. Therefore, in the simplified game 'not to advertise' is seller 2's best response to seller 1's choice 'not to advertise'. The same argument applies to seller 1. [5]

1.4.2 What will happen if advertising is not free?

In the previous literature, Butters (1977) shows the comparative statics result that as the advertising becomes more expensive, buyers pay for it in terms of higher prices. Robert and Stahl (1993) show that as the advertising cost goes up, the probability that

[5] One need to notice that in the simplified game, 'not to advertise' is not a dominant strategy for a seller. However, (seller 1 not to advertise, seller 2 not to advertise) is still a payoff dominant Equilibrium for sellers.

sellers choose the monopoly price will go up. Now let us introduce advertising cost into Advertise-then-Price**.** A seller must pay fixed cost *b* if it chooses to advertise its price. Then the structure of the simplified advertising game is as follows:

		Seller 2	
		Not to Advertise	To Advertise
Seller 1	Not to Advertise	$0.5v$, $0.5v$	$0.25\min(\frac{c}{1-\ln 2}, v)$, $0.5\min(\frac{c}{1-\ln 2}, v)$ - b
	To Advertise	$0.5\min(\frac{c}{1-\ln 2}, v)$ - b, $0.25\min(\frac{c}{1-\ln 2}, v)$	- b, - b

TABLE 1.2. Payoffs for the Sellers, Advertise-then-Price (advertising is not free)

From the table we can see that seller *i*'s strategy 'not to advertise' strictly dominates 'to advertise.' The unique strictly dominant strategy equilibrium in the simplified advertising game is (seller 1 not to advertise, seller 2 not to advertise). Not surprisingly, as advertising cost goes up, the monopoly price is the only market outcome. Therefore, analyzing the "boundary" case that advertising is free is enough to illustrate the characteristics of the Advertise-then-Price game.

1.4.3 What will happen if there are N sellers in the market?

Let us assume that there are N sellers making decisions and the buyer population is evenly distributed among N sellers, otherwise the game is exactly the same as the Advertise-then-Price game described in the previous section.

Given every seller's advertising choice in the pre-stage, there is a corresponding pricing subgame. We characterize the pricing subgames in three possible cases.

Case I: All of the sellers choose 'not to advertise.'

In this case the pricing subgame degenerates to the search cost game. The buyer has unique cutoff price r. When the lowest observed price is less than or equal to r, the buyer accepts the lowest observed price offer. Otherwise the buyer searches for a new price. The equilibrium cutoff price r^* depends on the equilibrium prices. '$P_1^* = \ldots = P_N^* = v$, the buyer sets the cutoff price at v' is the unique search equilibrium in this pricing subgame. Each of the sellers gets the expected payoff v/N, and the buyer gets payoff zero.

Case II: At least two of the sellers choose 'to advertise.'

The seller who advertises its price always has an incentive to undercut the other advertisers' prices in order to take the whole market. There is Bertrand competition among the advertisers. Finally, the advertisers set prices at zero and the sellers who choose not to advertise cannot make a sale. In Nash equilibrium, all sellers get payoff zero and the buyer gets payoff v.

Case III: One seller chooses 'to advertise,' but all other sellers choose 'not to advertise.'

The analysis of the one-ad pricing subgame is parallel to that of the two-seller case. Let the advertiser's price distribution be F and let the other seller's price distribution be G. The buyer's cutoff price is given by $r = \min(v, r')$, where r' is determined by $c = \int_0^{r'} G(p)dp$. The equilibrium payoff to the advertiser is r/N and the equilibrium payoff to the other sellers is r/N^2, where $r = \min(v, r')$ and r' is given by $r' = \frac{N-1}{N-1-\ln N}c$. The proof is collected in appendix A.

After the discussion of the equilibrium payoffs to sellers in all of the possible pricing subgames, we can easily verify the following observation.

Observation 1.4 In the N-seller Advertise-then-Price game, after substituting every pricing subgame by its solution, the unique weakly dominant strategy equilibrium in the simplified game is that every seller chooses not to advertise.

1.5 Concluding Remarks

The equilibrium prediction heavily depends on the way advertising is modeled. In this chapter, changing the way of modeling is much more than playing around with assumptions, since this new way captures important features in the real business practice.

There are some empirical studies that address the question whether price advertising reduces prices. Benham (1972) examines markets for eyeglasses in which advertising is prohibited and those in which advertising is allowed. Lower prices are found in markets that permit advertising, and there is no clear evidence to show that the quality of service is lowered in those markets, hence the null hypothesis that advertising restriction is a proxy of collusion cannot be rejected. Cady (1976) reports similar findings for retail prescription drug markets. Kwoka (1984) tests the claim that advertising lowers price and quality simultaneously, and hence others are forced to follow. The empirical evidence shows that the advertisers' prices and qualities are indeed lower, and while nonadvertisers' prices fall, their quality actually is greater. In Milyo and Waldfogel's "Rum and Vodka" paper (1999), they examine a single market with an exogenous regulation change. They find advertising stores substantially cut prices of their advertised products, and follow their rivals when their rivals cut prices, while nonadvertising stores do not do so.

The theoretical study in this chapter provides an alternative way to explain the empirical puzzle. Understanding the way to advertise is crucial to answer the question whether advertising reduces prices. If the format of advertising is to remove consumer search costs without directly revealing sellers' prices (for example, the "contact method advertising"), then it is possible for the sellers to choose not to advertise and to maintain a positive profit margin in the long run.

Chapter 2

EXPERIMENTAL EVIDENCE ON ADVERTISING AND PRICE COMPETITION

2.1 Introduction

This chapter provides an experimental test of the theoretical model in the previous chapter. 10 laboratory sessions were conducted at the Economic Science Laboratory (ESL), the University of Arizona. The two experimental conditions are (1) advertising the price, or (2) advertising before pricing. As stated previously, theory predicts that equilibrium prices will be lower when firms' advertising conveys the price than when it does not convey the price. Data from ten sessions indicate that, as predicted, firms choose more often to advertise when advertising conveys price, and prices in the advertise-then-price condition are significantly higher than prices in the advertise-with-price condition.

2.2 Experimental Design

In the laboratory sessions, all participants act as sellers. The buyers are computer-simulated and these automated buyers follow the equilibrium strategies described in the previous chapter.

There are several important reasons to use automated buyers instead of human buyers. First, my experiments incorporate a finite number of buyers. As argued by Coursey, Isaac and Smith (1984), incorporating a finite number of human buyers “could leave open the possibility that the competitive discipline of the markets is due not directly to contesting by sellers but rather to the actual (or merely anticipated) strategic withholding of demand by buyers." Using automated buyers allows me to control for strategic withholding of demand and gives the theory its best chance to survive. If one observes marginal cost pricing in laboratory sessions, this must be due directly to competition among sellers rather than the buyers' market power. Moreover, how the buyers' market power influences market prices is not the research question of this chapter. Secondly, when the buyers population is relatively large, the assumption that buyers reveal demand is quite realistic. It is hard to imagine that buyers strategically withhold demand when they shop in grocery stores and bookstores. Finally, using automated buyers substantially reduces the payments to subjects.

The main goal of this chapter is to answer the question of how the format of advertising matters. There are two games in the experiments, the advertise-then-price game and the advertise-with-price game.

In the advertise-then-price game, sellers first simultaneously decide whether to advertise their prices (i.e., reveal their prices). As shown in figure 2.1.A, they click button "Reveal" if they decide to advertise, and they click "Not to Reveal" if they decide not to advertise. Then the sellers will observe their rivals' advertising choices ("Reveal" or "Not to Reveal") and will enter their prices in the given text box, as shown in figure 2.1.B. In the advertise-with-price games, all sellers simultaneously make their advertising and pricing choices. As shown in figure 2.1.C, sellers click the option button to choose whether or not to advertise, and enter the price in the given text box.

In both the advertise-then-price and advertise-with-price games, advertising is free for sellers. The buyer search cost is fixed at 30 cents and the number of sellers in the market is fixed at three. All sellers have zero costs and have no capacity constraint. The buyers' reservation value is $2.00. According to the market games in the previous section, in the advertise-then-price games, 'all sellers choose not to advertise' is the weakly dominant strategy equilibrium in the simplified advertising game and the equilibrium price is $2.00. In the advertise-with-price games, 'all sellers choose to advertise, and sellers choose prices at zero' is the equilibrium prediction.

The number of automated buyers in the market is also three. Each automated buyer demands one unit of the good. The automated buyer's shopping rule, which is written in the instructions, is known to all sellers.

10 sessions were conducted at the Economic Science Laboratory (ESL), the University of Arizona, from November 2003 to January 2004. The experimental software was written in Visual Basic 6.0. 5 sessions are advertise-then-price games (identified as ATP1, ATP2,…,ATP5) and the other 5 sessions are advertise-with-price games (identified as AWP1, AWP2,…,AWP5). In each session, 6 human subjects are recruited as sellers. 20 trading periods are scheduled in each session. At the beginning of every trading period, 6 sellers are randomly divided into two markets, 3 sellers in each market. At the end of every trading period, sellers review their own profit or loss and observe the choices of the other sellers from the same market, as described in figure 2.1.D. Subjects were randomly chosen from the ESL experiment recruiter database. Those who registered in the database must have a valid the University of Arizona student ID card. Each subject

participated in only one session. The experimental treatments were implemented across-subjects; that is, different subjects participated in the two market games.

Each subject was paid a $5 show-up fee plus the earnings during the experiment. The earnings during the experiment were recorded in experimental dollars. The experimental dollars were convertible to USD at the rate of 0.5 USD per experimental dollar in the advertise-then-price games, and at the rate of 1 USD per experimental dollar in the advertise-with-price games. The average payment in the advertise-then-price games was $19.95, and the average payment in the advertise-with-price games was $11.76. The length of the advertise-with-price games was about 45 minutes. The length of the advertise-then-price games was about 1 hour and 20 minutes and the task was moderately more complicated.

Since there are only six subjects in each session, the common history of plays cannot be fully avoided. Each session of each market game is treated as one independent observation.

2.3 Results

Sellers need to make two decisions in the market games. They need to decide whether to advertise and they need to choose their prices. These two decisions are separately analyzed as follows.

2.3.1 Sellers' Advertising Choices

The frequencies of the advertising choices in advertise-with-price games are reported in figure 2.2. From figure 2.2 we can see that in each session, the frequency of "*all of the sellers decide to advertise*" ('all advertise' in figure 2.2) dominates the frequency of "*All of the sellers decide not to advertise*" ('no advertise' in figure 2.2) and dominates the frequency of "*some of the sellers decide to advertise, but the other sellers decide not to advertise.*" ('other' in figure 2.2). The overall frequencies are reported in figure 2.3. The proportions of the advertising choices in advertise-with-price games are reported in table 2.1. We can see that the proportions of 'no advertise' in all 5 sessions are zero. Table 2.2 shows that both means test and Kolmogorov-Smirnov test suggest that the proportion of 'all advertise' is greater than 0.5 in advertise-with-price games. The test statistics are statistically significant at the one percent level. These results are consistent with the theoretical prediction that in the advertise-with-price game, all sellers choose to advertise in equilibrium.

It is shown in figure 2.4 that in advertise-then-price games the frequency of 'no advertise' dominates the frequency of 'all advertise.' However, there are a lot of noises. In session ATP2 and session ATP4, the frequency of 'other' is higher than the frequency of 'no advertise.' Overall, the frequency of 'no advertise' is much higher than the frequency of choices of 'other' (figure 2.5). The proportions of the advertising choices in the advertise-then-price games are reported in table 2.3. Table 2.4 shows the tests on the proportion of 'no advertise' in advertise-then-price games. The mean test suggests that the proportion of 'no advertise' is greater than 0.5, though the test statistics is marginal significant (p-value is 0.06). The Kolmogorov-Smirnov test suggests that we cannot reject the null hypothesis that the proportion of 'no advertise' is equal to 0.5. From these results, we find weak evidence to support the theoretical prediction that in the advertise-then-price game, none of the sellers choose to advertise in equilibrium.

Table 2.5 reports the tests on the treatment effect on advertising. The test results show that the treatment effect is highly significant. Comparing the advertise-with-price games and the advertise-then-price games, the proportion of 'all advertise' decreases from 0.845 to 0.03, and the proportion of 'no advertise' increases from 0 to 0.755. Both means test statistics and Kolmogorov-Smirnov test statistics are statistically significant at the one percent level. These results are consistent with the hypothesis that sellers choose more often to advertise when advertising conveys price.

2.3.2 Sellers' Pricing Decisions

The mean transaction prices in the advertise-with-price games are shown in figure 2.6. The prices start high and decrease over periods. By the end of the sessions, the prices converge to zero. We find strong evidence to support marginal cost pricing in advertise-with-price games.

The mean transaction prices in the advertise-then-price games are shown in figure 2.7. In session ATP1, ATP3 and ATP5, subjects chose the weakly dominant strategy 'not to advertise' at the very beginning, and the prices are maintained close to the buyers' reservation price \$2.00. From session ATP2, the prices stay low in the first 10 periods. From period 11 to period 20, subjects start to figure out that 'not to advertise' gives them a better payoff and the prices start to climb. By the end of the session, the market prices are maintained at $2.00. There are many fluctuations in the session ATP4.

The comparison of the mean transaction prices between the advertise-with-price games and advertise-then-price games is shown in figure 2.8. We can see that the mean transaction prices in the five advertising-then-price sessions are all higher than mean transaction prices in the advertising-with-price sessions. The numerical figures are reported in table 2.6. The tests on mean transaction prices are reported in table 2.7. The treatment effect on the mean transaction prices is statistically significant. The mean transaction prices increases from $0.338 in the advertise-with-price games to $1.495 in the advertise-then-price games. All three test statistics (means test, Kolmogorov-Smirnov test and Mann-Whitney test) are statistically significant at one percent level. These results are consistent with the research hypothesis that the format of advertising has an impact on transaction prices.

2.4 Concluding Remarks

This chapter provides an experimental test of a recent theoretical model of advertising and price competition by Du (2004). In a controlled laboratory environment, a market for a homogeneous good is created in which three human sellers compete to sell to three robot buyers. Each seller makes two decisions: what price to set, and whether to advertise to eliminate consumer search costs for their product. Each robot buyer, who is constrained to buy at most one unit of the commodity, then accepts an observed price, drops out, or pays a search cost to find an unobserved price according to the optimal search rule known by all sellers. The two experimental conditions are (1) advertising the price, or (2) advertising before pricing. As the theory predicts, equilibrium prices will be lower when firms' advertising conveys the price than when it does not convey the price.

Data from ten laboratory sessions indicate that, as predicted, firms choose more often to advertise when advertising conveys price, and prices in the advertise-then-price treatment are significantly higher than prices in the advertise-with-price treatment. The empirical evidence from this chapter strongly supports the hypothesis that the market price is sensitive to the nature of the format of advertising. This gives a concrete example of "*institutions matter because the rules matter, and the rules matter because incentives matter*."

TABLE 2.1. Advertising Choices: Advertise-with-Price Games

	Proportion of 'No Advertise'	Proportion of 'All Advertise'	Proportion of Other Advertising Choices
AWP1	0	0.900	0.100
AWP2	0	0.950	0.050
AWP3	0	0.875	0.125
AWP4	0	0.650	0.350
AWP5	0	0.850	0.150
Average Proportion	0	0.845	0.155

TABLE 2.2. Tests on Average Proportion of 'All Advertise' (Advertise-with-Price Games)

Average Proportion of 'All Advertise' (P_A)	Alternative Hypothesis	Means Test (*t*)	Kolmogorov-Smirnov Test (*D*)
0.845	$P_A > 0.5$	6.70 (p=0.0013)	1 ($p < 0.01$)

TABLE 2.3. Advertising Choices: Advertise-then-Price Games

	Proportion of 'No Advertise'	Proportion of 'All Advertise'	Proportion of Other Advertising Choices
ATP1	1.000	0.000	0.000
ATP2	0.425	0.100	0.475
ATP3	0.900	0.000	0.100
ATP4	0.450	0.050	0.500
ATP5	1.000	0.000	0.000
Average Proportion	0.755	0.030	0.215

TABLE 2.4. Tests on Average Proportion of 'No Advertise' (Advertise-then-Price Games)

Average Proportion of 'No Advertise' (P_N)	Alternative Hypothesis	Means Test (t)	Kolmogorov-Smirnov Test (D)
0.755	$P_N > 0.5$	1.95 (p=0.06)	0.6 ($p > 0.05$)

TABLE 2.5. Tests on Average Proportions of Advertising Choices (Advertise-with-Price Games vs. Advertise-then-Price Games)

	Advertise-with-Price Games	Advertise-then-Price Games	Alternative Hypothesis	Means Test (*t*)	Kolmogorov-Smirnov Test (*D*)
Proportion of 'All Advertise' (P_A)	0.845	0.03	$P_A^{ATP} < P_A^{AWP}$	14.76 (p<0.0001)	1 ($p < 0.01$)
Proportion of 'No Advertise' (P_N)	0	0.755	$P_N^{AWP} < P_N^{ATP}$	5.77 (p=0.0022)	1 ($p < 0.01$)

TABLE 2.6. Mean Transaction Prices

	Mean Transaction Price		Mean Transaction Price
ATP1	1.944	AWP1	0.347
ATP2	0.877	AWP2	0.477
ATP3	1.696	AWP3	0.187
ATP4	0.967	AWP4	0.370
ATP5	1.993	AWP5	0.309
Advertise-then-Price Games	1.495	Advertise-with-Price Games	0.338

TABLE 2.7. Tests on Mean Transaction Prices (Advertise-with-Price Games vs. Advertise-then-Price Games)

Advertise-with-Price Games(P^{AWP})	Advertise-then-Price Games(P^{ATP})	Alternative Hypothesis	Means Test	Kolmogorov-Smirnov Test (D)	Mann-Whitney Test (U)
0.338	1.495	$P^{AWP} < P^{ATP}$	4.74 (p=0.005)	1 ($p < 0.01$)	0 ($p < 0.01$)

FIGURE 2.1.A. Seller Window for Advertising Choice (Advertise-then-Price)

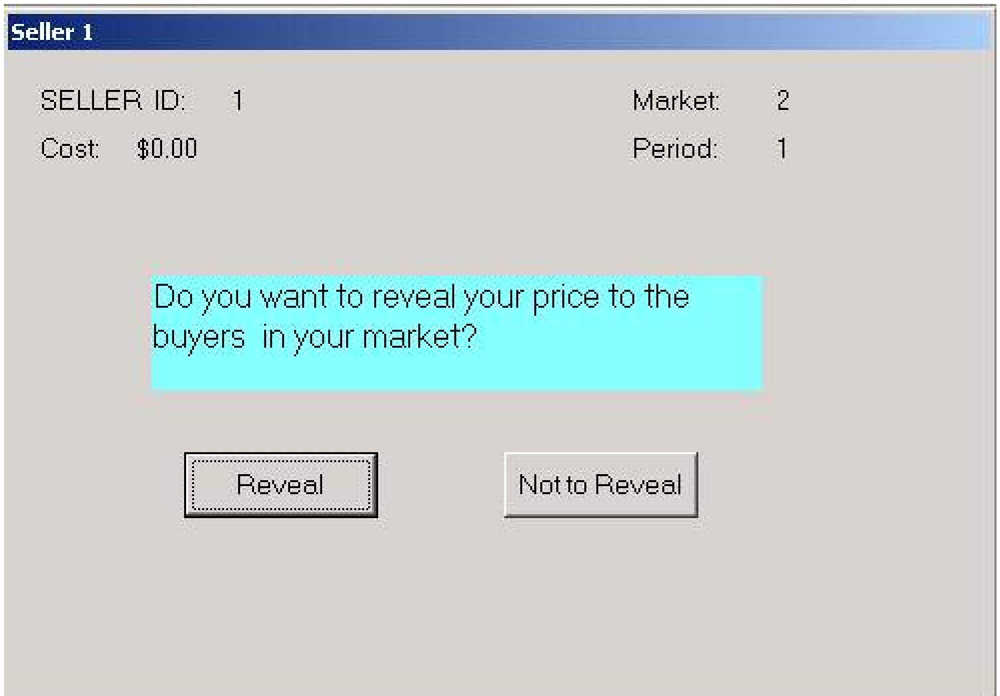

FIGURE 2.1.B. Seller Window for Choosing Price (Advertise-then-Price)

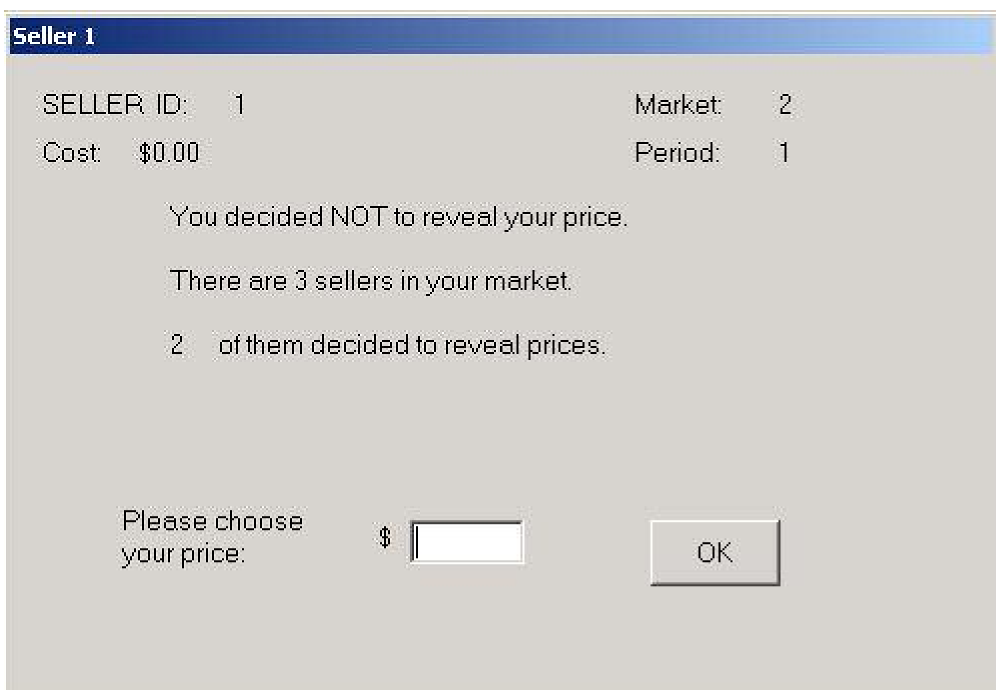

FIGURE 2.1.C. Seller Window (Advertise-with-Price)

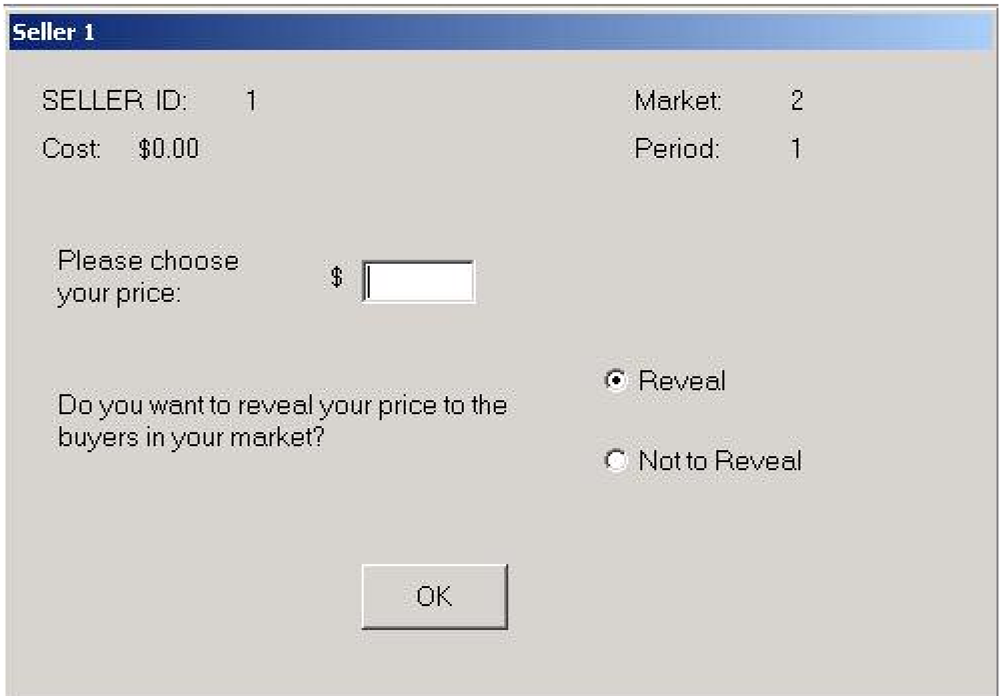

FIGURE 2.1.D. Profit Summary

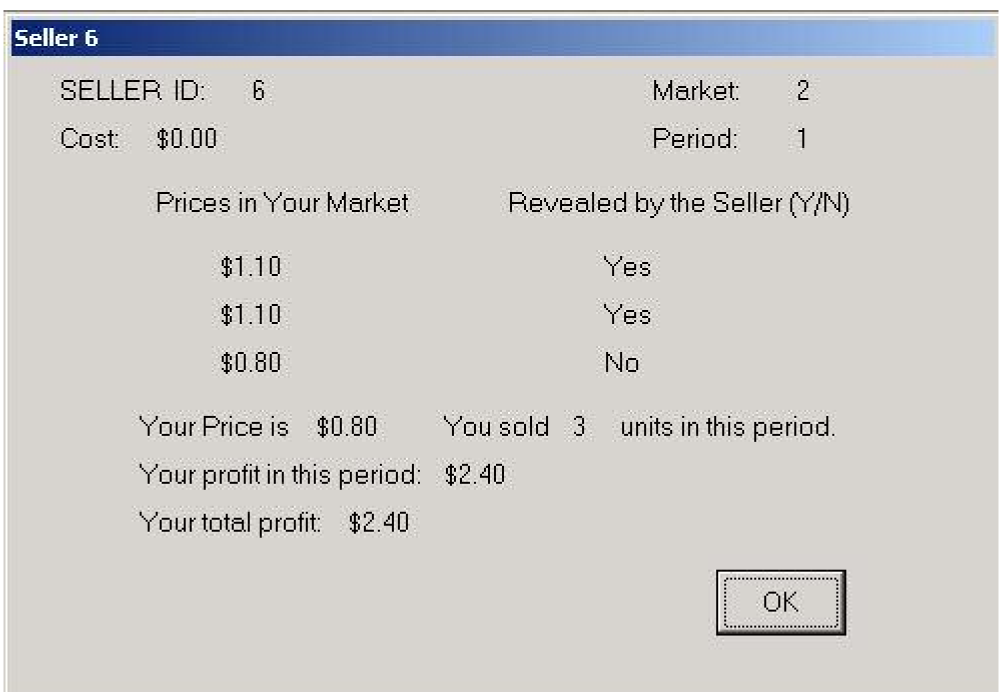

FIGURE 2.2.A. Advertising Choices: AWP1

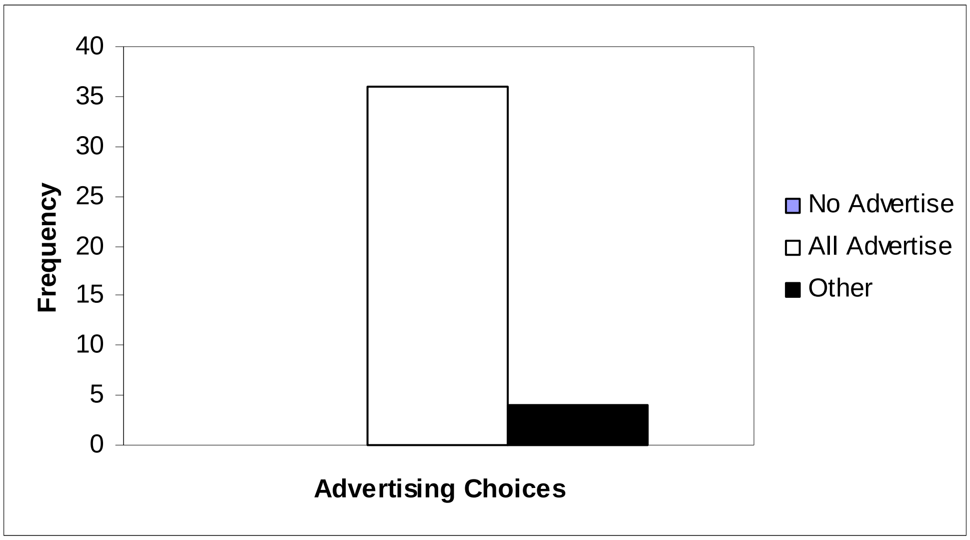

FIGURE 2.2.B. Advertising Choices: AWP2

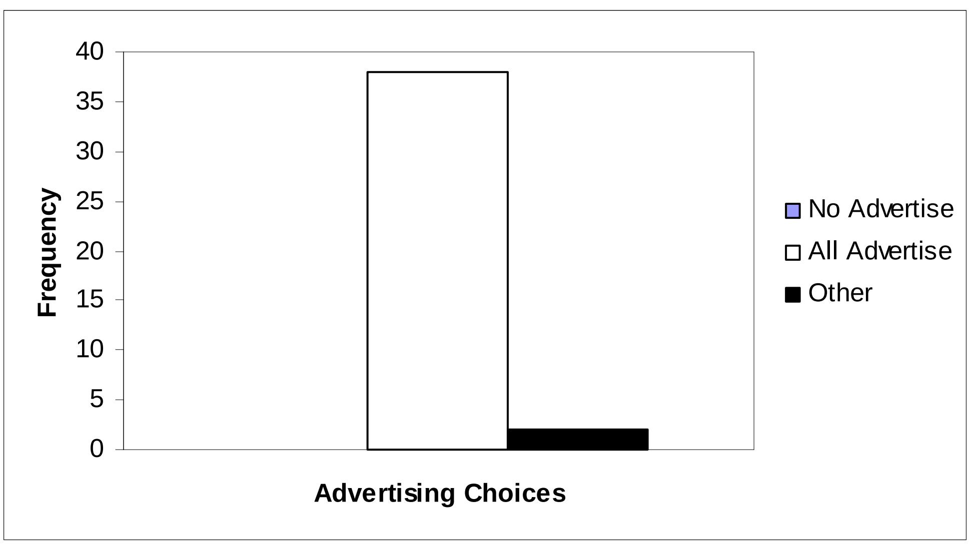

FIGURE 2.2.C. Advertising Choices: AWP3

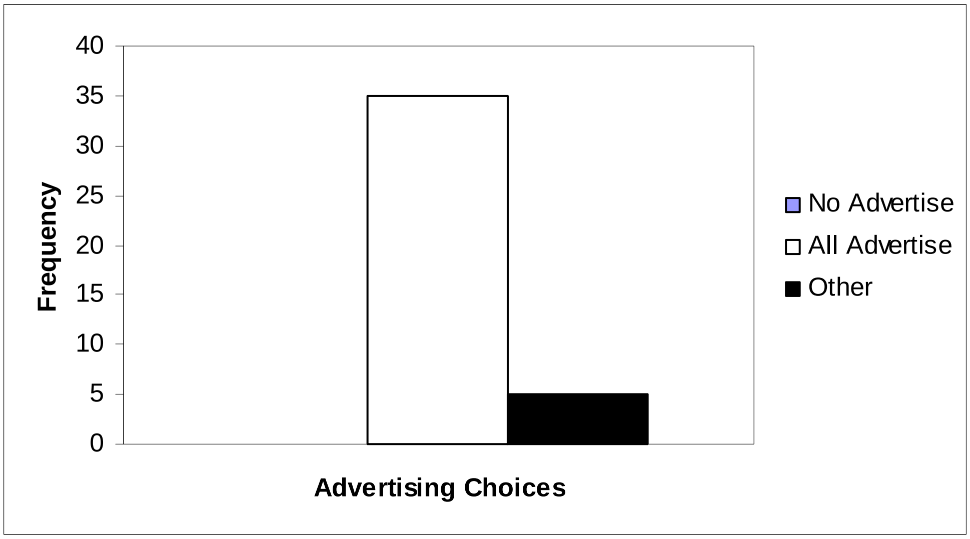

FIGURE 2.2.D. Advertising Choices: AWP4

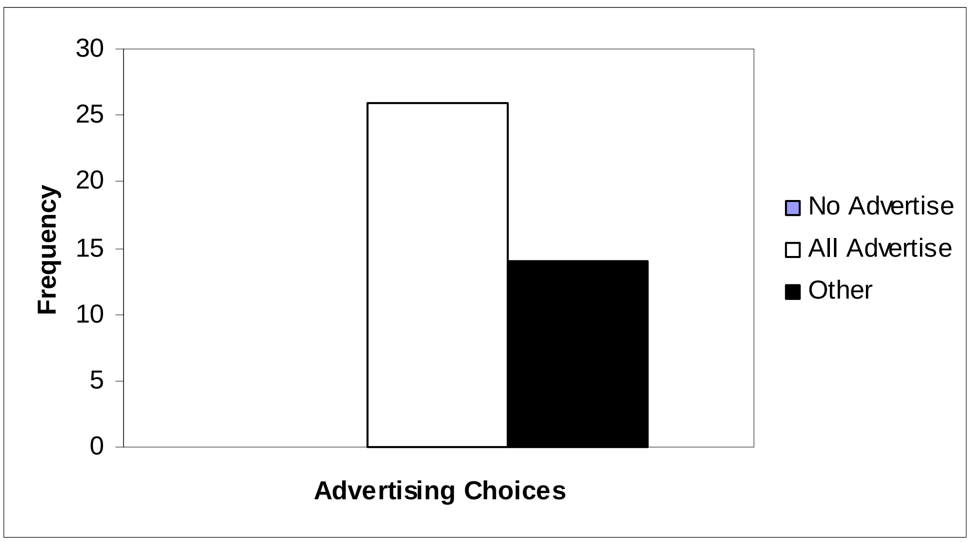

FIGURE 2.2.E. Advertising Choices: AWP5

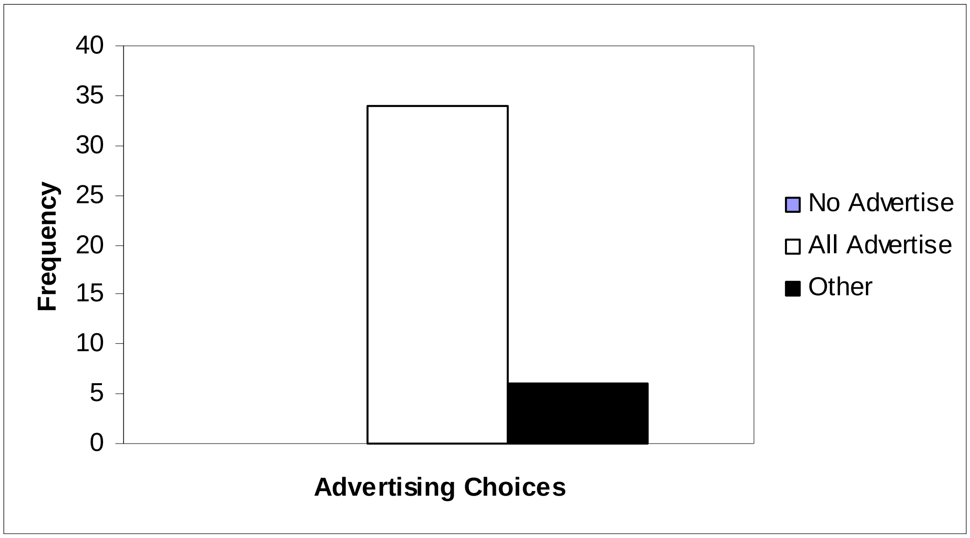

FIGURE 2.3. Advertising Choices: Advertise-with-Price Games

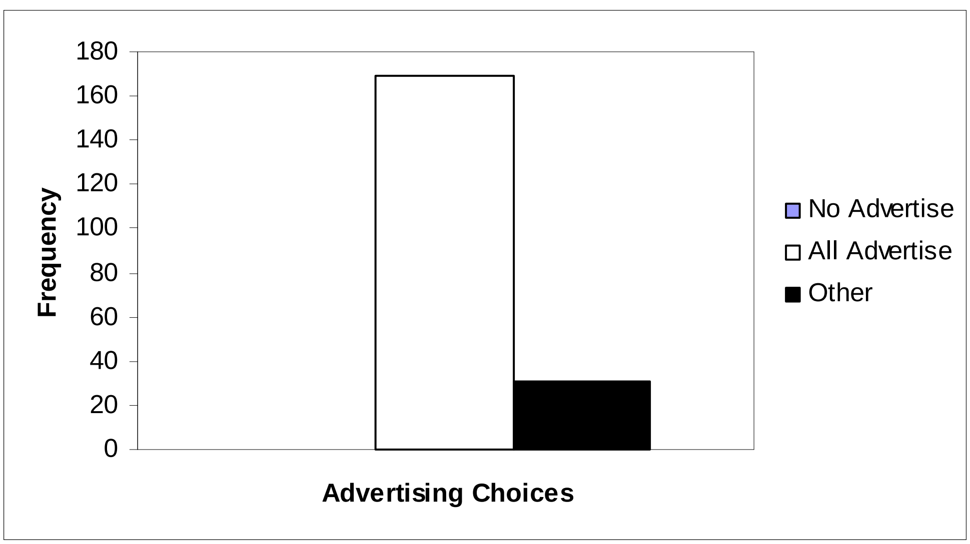

FIGURE 2.4.A. Advertising Choices: ATP1

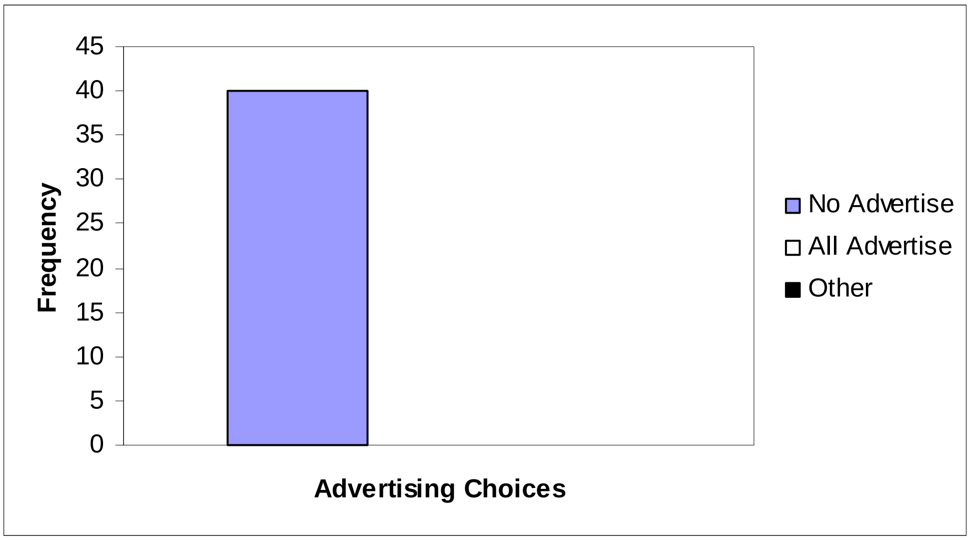

FIGURE 2.4.B. Advertising Choices: ATP2

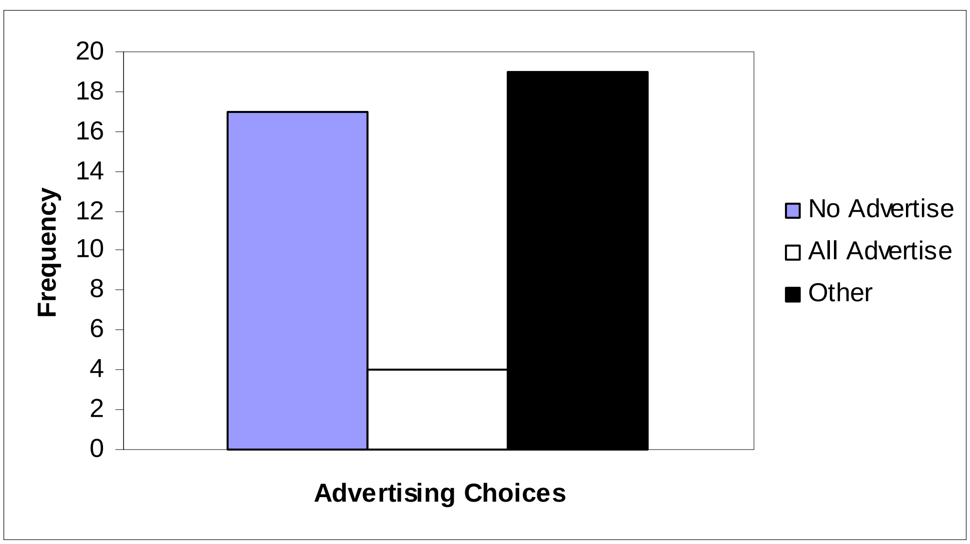

FIGURE 2.4.C. Advertising Choices: ATP3

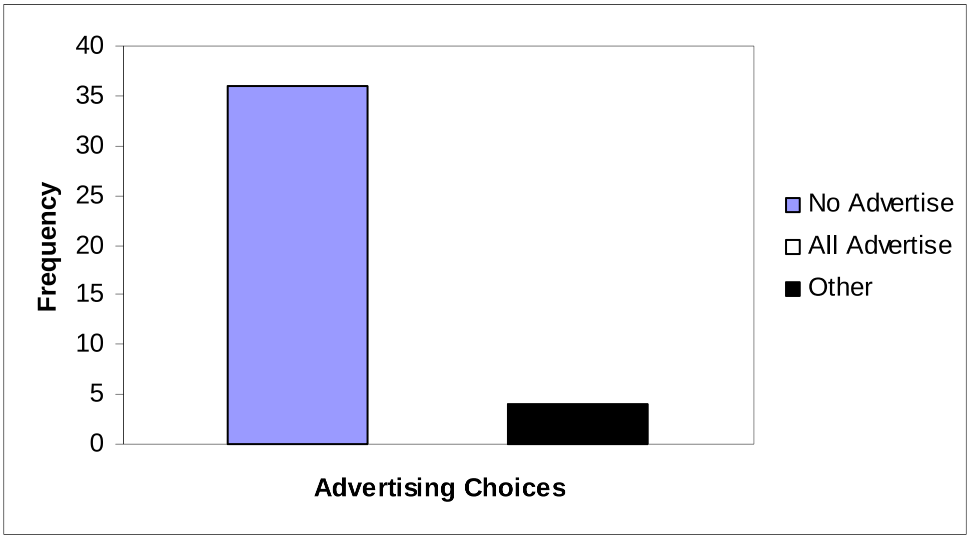

FIGURE 2.4.D. Advertising Choices: ATP4

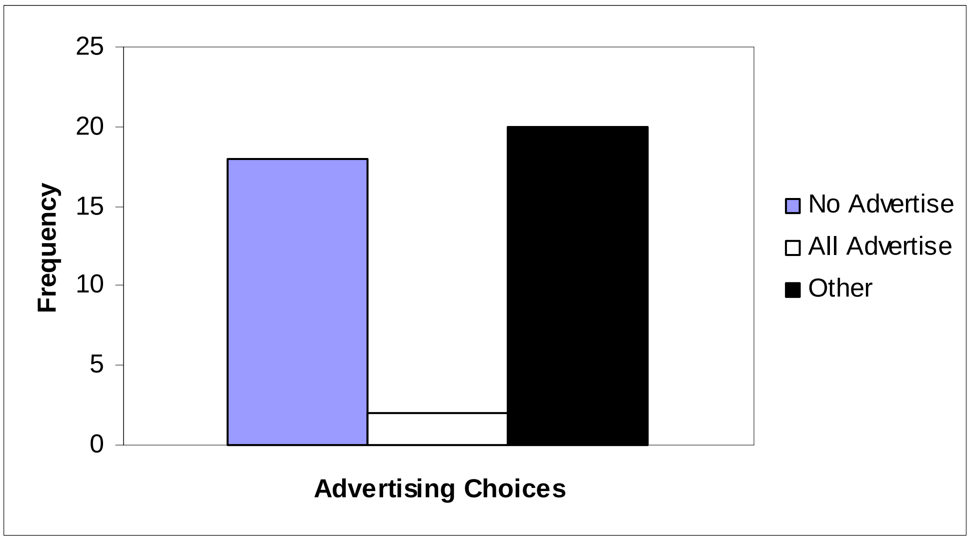

FIGURE 2.4.E. Advertising Choices: ATP5

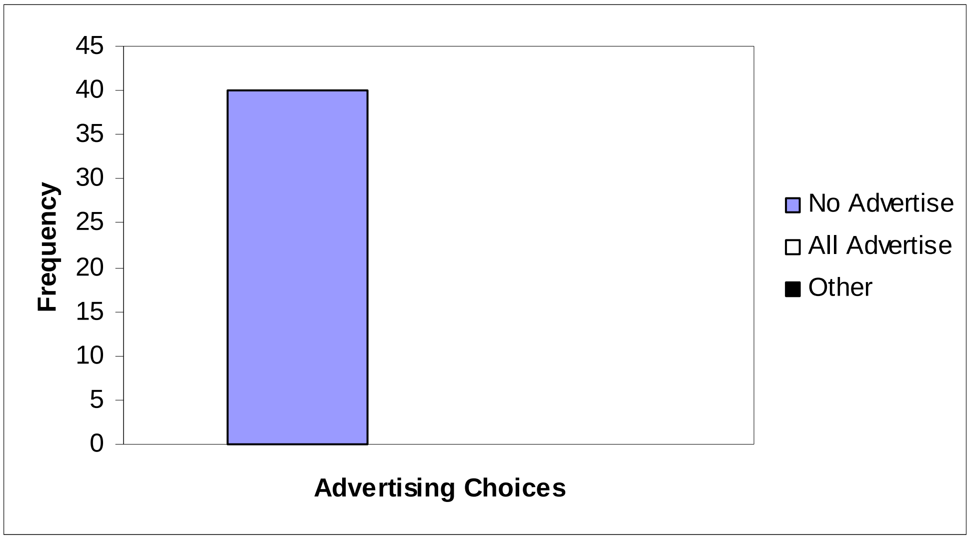

FIGURE 2.5. Advertising Choices: Advertise-then-Price Games

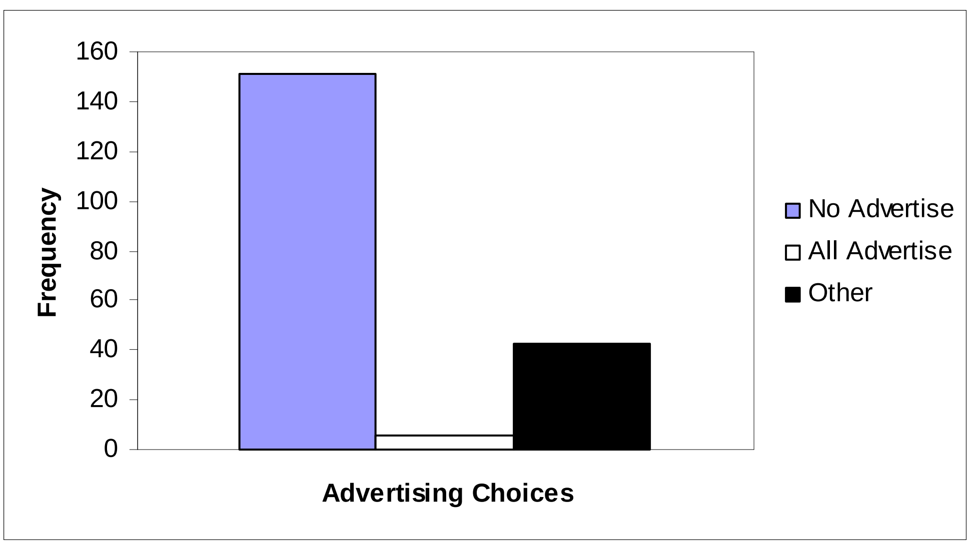

FIGURE 2.6.A. Mean Transaction Price: AWP1

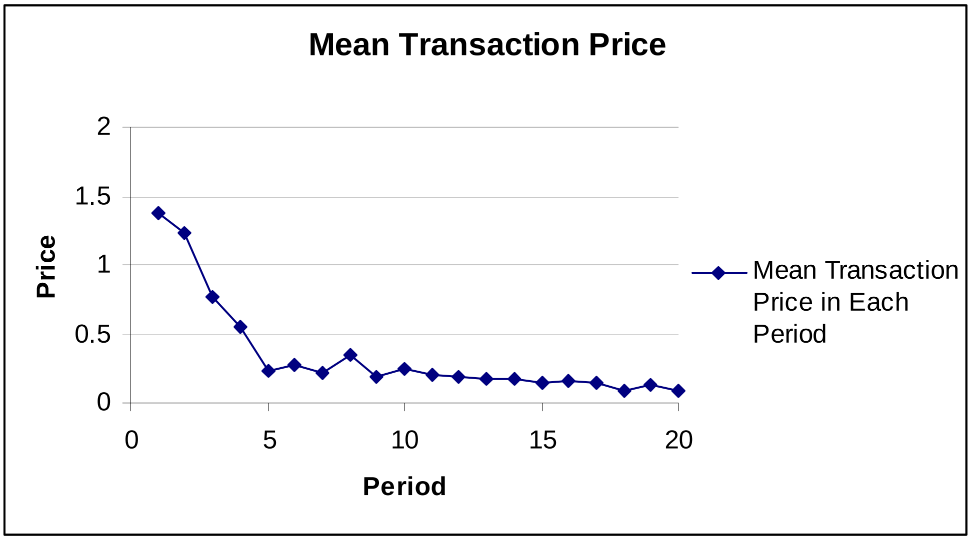

FIGURE 2.6.B. Mean Transaction Price: AWP2

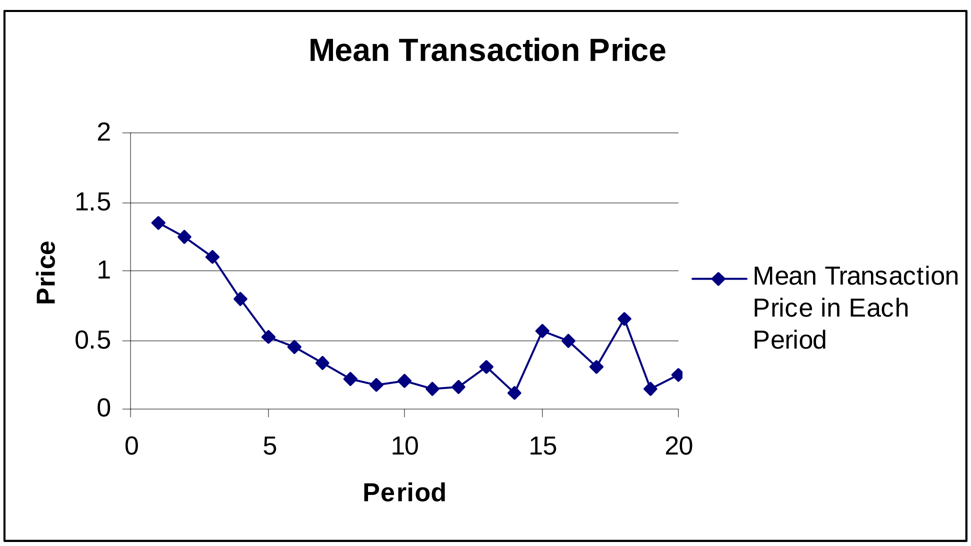

FIGURE 2.6.C. Mean Transaction Price: AWP3

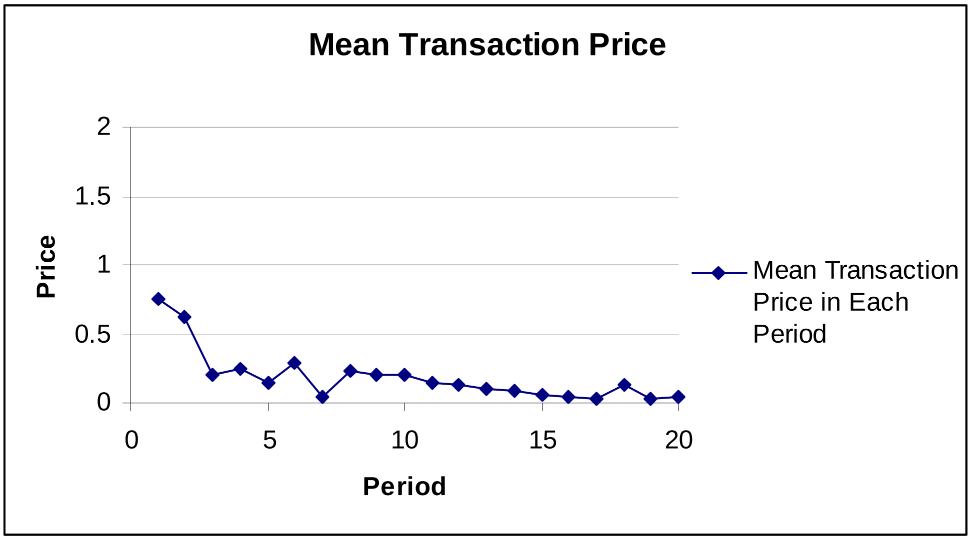

FIGURE 2.6.D. Mean Transaction Price: AWP4

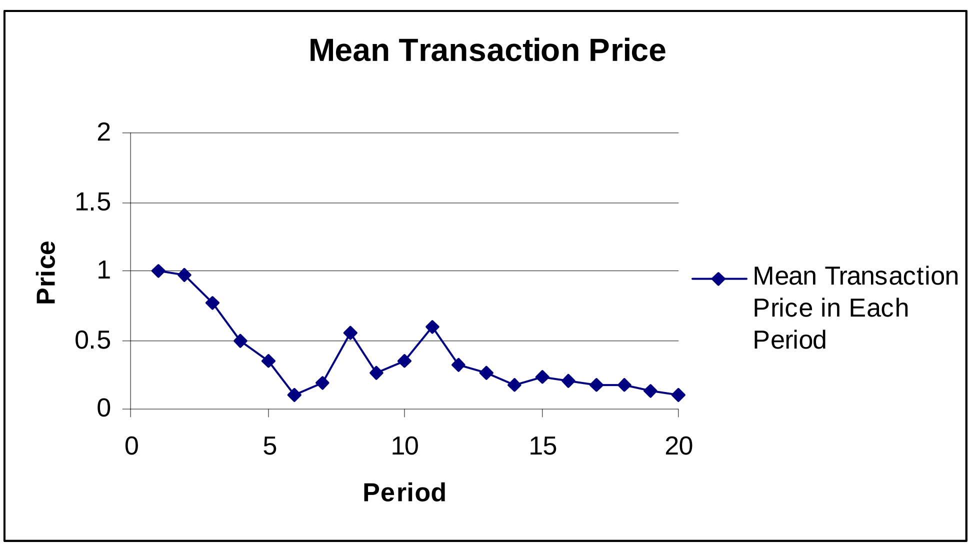

FIGURE 2.6.E. Mean Transaction Price: AWP5

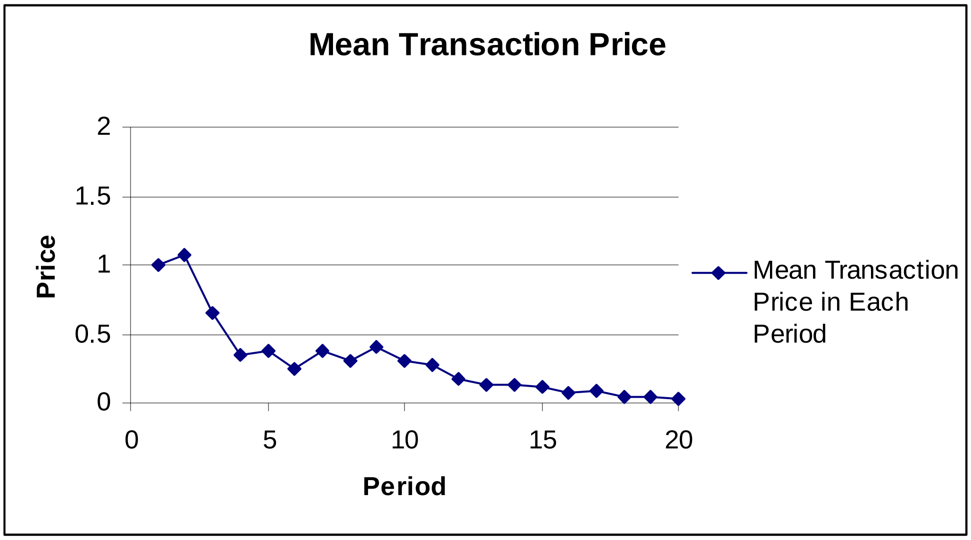

FIGURE 2.7.A. Mean Transaction Price: ATP1

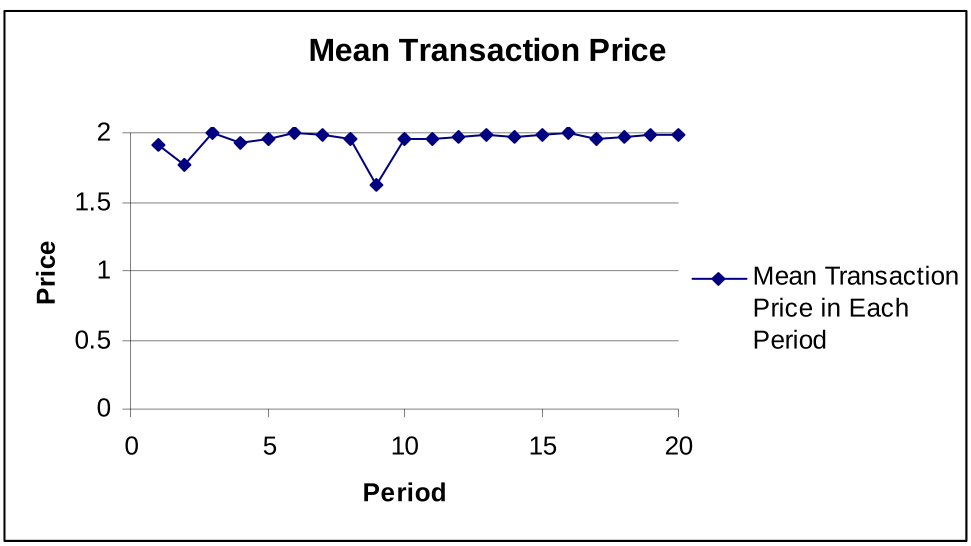

FIGURE 2.7.B. Mean Transaction Price: ATP2

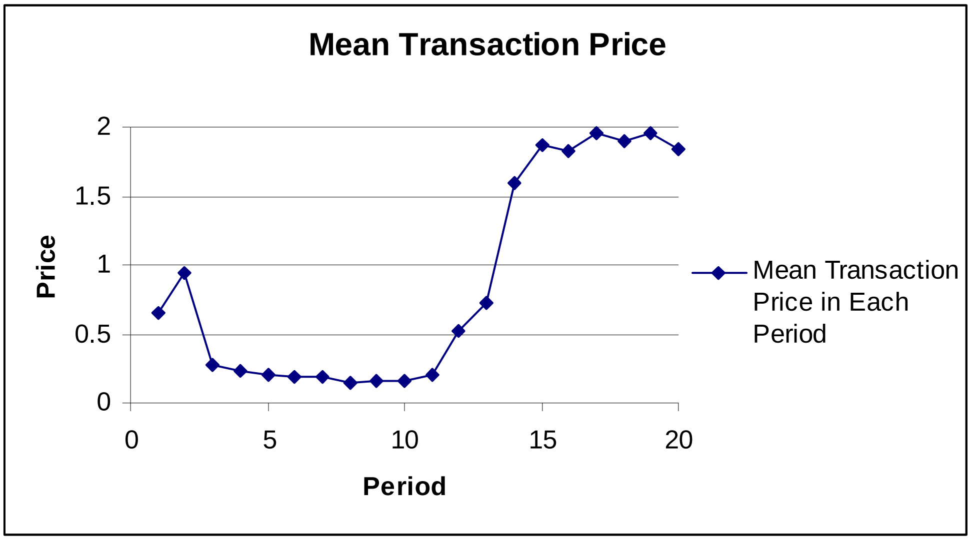

FIGURE 2.7.C. Mean Transaction Price: ATP3

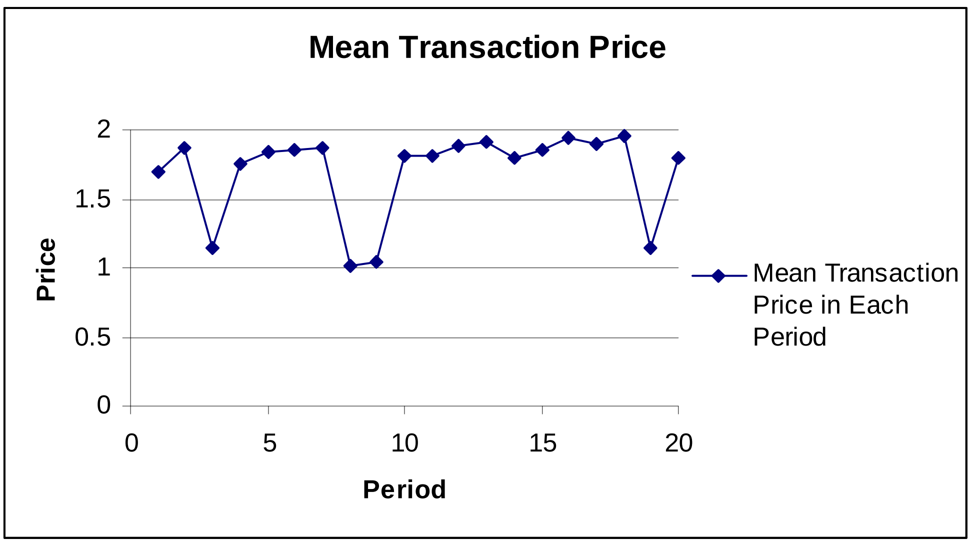

FIGURE 2.7.D. Mean Transaction Price: ATP4

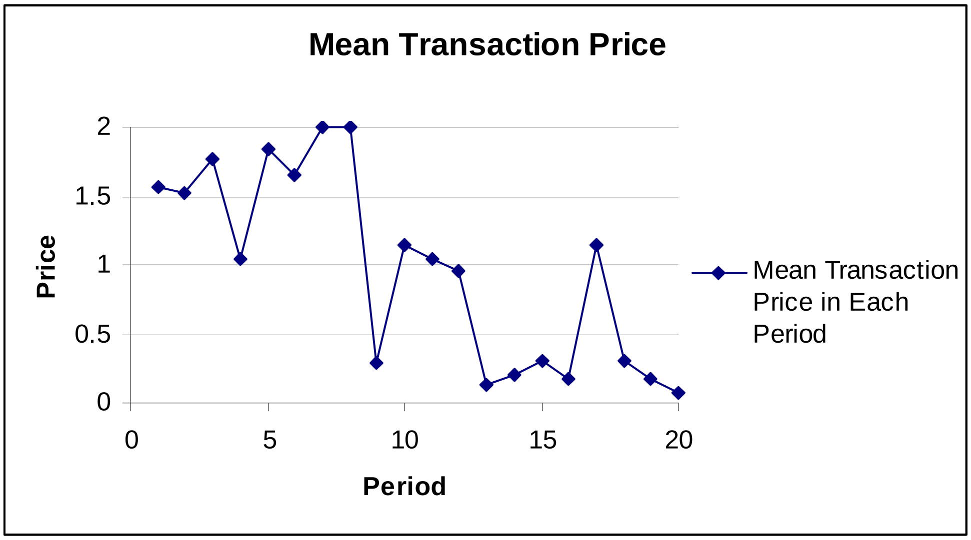

FIGURE 2.7.E. Mean Transaction Price: ATP5

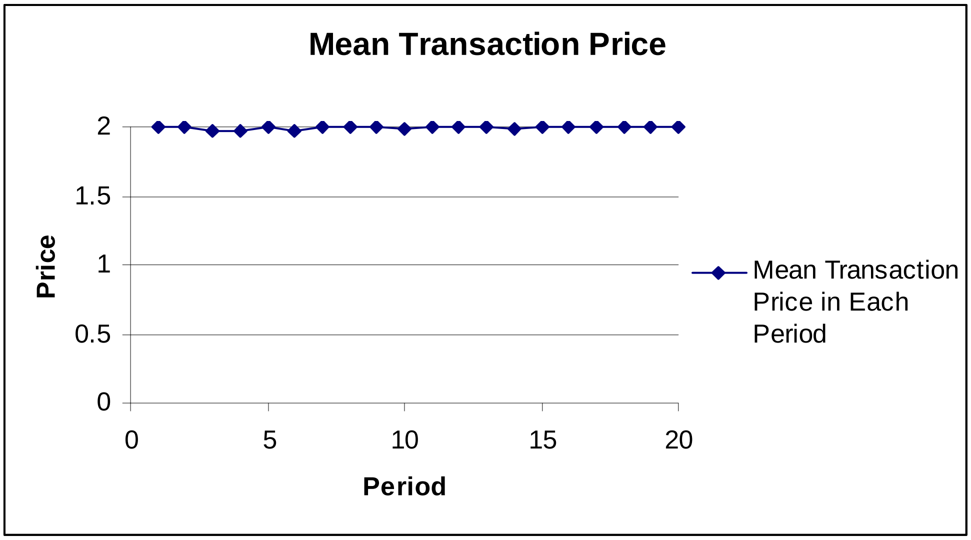

FIGURE 2.8. Mean Transaction Price: Advertise-with-Price Games vs. Advertise-then-Price Games

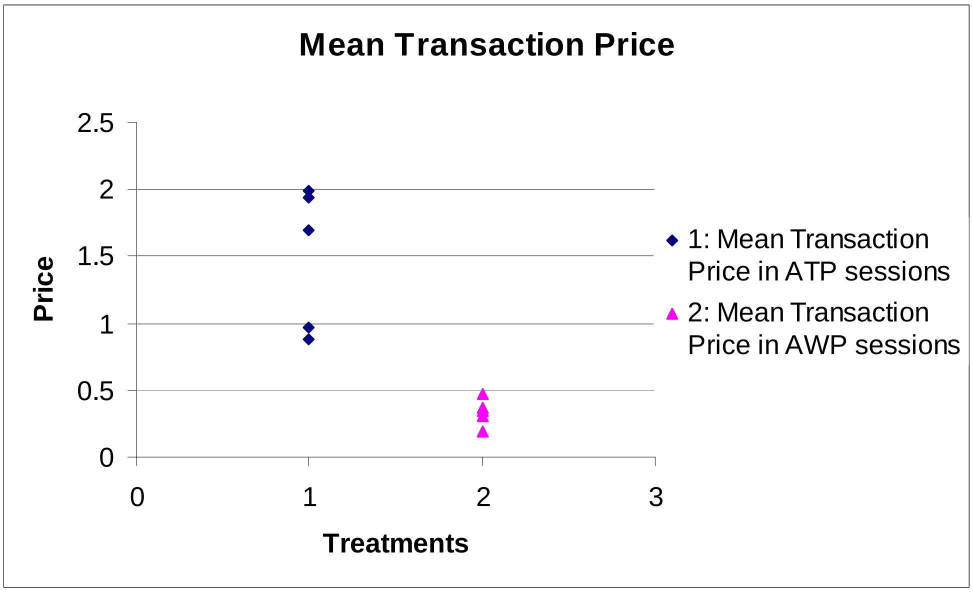

Chapter 3

DO HIGHER SEARCH COSTS LEAD TO HIGHER PRICES? *A PRICE COMPARISON BETWEEN ONLINE AND RETAIL BOOKSTORES*

3.1 Introduction

Electronic commerce plays an increasingly important role in our economy. Nowadays consumers are able to browse product prices at home and find what they are looking for without incurring transportation cost. The theory of information economics predicts that this source of free information reduces both price and price dispersion. However, the previous empirical work generates contradictive results. Some studies suggest that online prices are lower than prices at conventional brick-and-mortar stores, while other studies suggest that online prices are indeed higher (See Lee 1997, Bailey 1998, Brynjolfsson and Smith 2000 and Clay *et al.* 2002).

This chapter uses a new data set to test the research hypothesis that the internet reduces both the mean price and price dispersion. I collected the prices during the week of January 12, 2004 for 20 books identified by their ISBN number across six online bookstores and six physical bookstores in Tucson, Arizona. The results of the book price comparisons show that the internet has a negative impact on average prices, while the internet does not lower the variance of prices. The contribution of this chapter includes a new data set and the control for the costs of selected books, which have not been studied in the previous literature.

3.2 Theoretical Predictions

Stigler (1961) launched the literature of information economics and showed that advertising is a substitute for consumer search and results in a lower mean and variance of prices. When advertising is free for a homogeneous commodity, the market outcome reverts to marginal cost pricing. The studies by Butters (1977) and Robert and Stahl (1993) show that when the advertising cost is positive, there will be price dispersion in the market. The comparative statics results from Butters and Robert & Stahl show that consumers pay higher prices as search becomes more expensive and more costly search also implies greater price dispersion.

For completeness, I reproduce here the settings and main results by Butters (1977). Consider a homogeneous product market with M identical buyers and N identical sellers. Each buyer has reservation value v for one unit of the good, and sellers' costs are normalized to zero. It costs a seller $b > 0$ to advertise her price to one buyer. The search cost for a buyer to observe a not advertised price is $c > 0$. Let A(p) equal the number of ads sent out at prices less than or equal to p divided by the number of buyers, M. Let $p_{\min}$ be the minimum price at which ads are ever sent and let $p_{\max}$ be the maximum price at which ads are ever sent. Finally, let us use h to denote the proportion of buyers who receive at least one ad. Then 1 - h is the proportion of buyers who search. In Nash equilibrium, we have the following equations:

$$p_{\min} = bh \tag{3.2.1}$$

$$A(p) = \ln(\frac{p}{bh}) \tag{3.2.2}$$

$$c = b\left(\frac{h}{1-h} - \ln\left(\frac{1}{1-h}\right)\right) \tag{3.2.3}$$

$$\bar{p} = b\ln\left(\frac{1}{1-h}\right) \tag{3.2.4}$$

$$\mathrm{var}(p) = b^2\left(\frac{h^2}{1-h} - 2\ln\left(\frac{1}{1-h}\right)\right) \tag{3.2.5}$$

Equation (3.2.1) says that in equilibrium, a seller's minimum price must cover her advertising cost. Equation (3.2.2) indicates that the number of ads sent out by a seller is decreasing in her advertising cost. $\bar{p}$ is the mean of the sales price distribution and var(p) is the variance of the sales price distribution. From the equations (3.2.3), (3.2.4) and (3.2.5), we have the comparative statics results:

$$\frac{\partial h}{\partial c} = \frac{(1-h)^2}{bh} > 0 \tag{3.2.6}$$

$$\frac{\partial \bar{p}}{\partial c} = \frac{1-h}{h} > 0 \tag{3.2.7}$$

$$\frac{\partial\, \mathrm{var}(p)}{\partial c} = 2c\left(\frac{1-h}{h}\right) - bh\,;\ 0 < \frac{\partial\, \mathrm{var}(p)}{\partial c} < c(1-h) \tag{3.2.8}$$

The comparative statics results show that as the consumer search cost goes up, both the mean price and the variance of prices will go up. For mathematical details, see Butters (1977).

Both consumer search and advertising are important activities in electronic commerce. Although acquiring price information on the internet is much easier than acquiring price information from physical retail stores, the consumer search cost on the internet is still not trivial. It is true that consumers can use internet search engines to obtain prices from a variety of online stores. However, some consumers may not know about the engines. For consumers who use the engines, they may not take the time and effort to click through 10 page long search results (See Clay *et al.* 2002).

Most online stores use the mechanism of "sponsored links" to advertise their existence, instead of directly advertising prices. For example, Overture.com is the advertising department of Yahoo!, Inc. First, merchants select search terms and write titles and descriptions for their products. Second, merchants bid on search terms. The bid on each search term is also called Cost-Per-Click (CPC), which is the amount a merchant pays Overture when an actual shopper clicks on that merchant's product listing within the Overture.com shopping search results. If a consumer visits Overture.com and searches for a certain search term, then Overture.com will return a search result page. The search result page contains the links to the merchants who bid on that search term. The links are ordered by merchants' CPC. The higher the CPC, the better the placement is.

Given the fact that both the consumer search cost and the advertising cost are positive on the internet, theory predicts that there will be price dispersion online. However, since the search cost for prices on the internet is lower than the search cost for prices from conventional bricks-and-mortar stores, theory predicts that online prices will have a lower degree of price dispersion.

3.3 The Previous Empirical Results

The literature that examines the patterns of prices for homogeneous goods between internet retailers and traditional bricks-and-mortar retailers goes back to Lee

(1997). Lee studied an electronic marketplace AUCNET, which was introduced to reduce the search costs of buyers for used-car transactions in Japan. Lee found that the average contract price of secondhand cars sold through AUCNET was much higher than that of traditional, non-electronic markets. One explanation of this phenomenon is the coordination economy—higher contract prices attract many sellers to list their cars on AUCNET, and this in turn attracts more buyers since AUCNET offers a variety of purchase choices.

Bailey (1998) inaugurated the literature of the online book industry. Bailey analyzed prices for books, CDs, and software in Internet and conventional outlets from 1996 to 1997. He found that prices were actually higher online than in retail stores. Further, he found more price dispersion on the internet. Internet retailers can develop pricing strategies to differentiate themselves from their competitors and to price discriminate. A later study by Brynjolfsson and Smith (2000) uses a data set of over 8,500 price observations collected over a period of 15 months and across 41 Internet and conventional retail stores. They found that internet prices were 9-16% lower than prices in conventional stores. The most recent work by Clay *et al*. (2002) uses data collected in April 1999 on the prices of 107 books in thirteen online and two physical bookstores. They found similar average prices online and in physical stores. However, the online price dispersion was greater than in physical stores.

This chapter differs from the previous literature in two important ways. First, the data set is new. Bailey (1998) and Brynjolfsson and Smith (2000) compare internet prices with Boston local book market prices. Clay *et al*. (2002) choose two nation-wide chain stores, Borders and Barnes and Nobles in Pittsburgh, Pennsylvania as the physical retailers. I choose six bookstores in Tucson, AZ, including Borders and Barnes and Nobles and four local bookstores, as the physical retailers. My data set is collected during a more recent time period. Second, this chapter uses more advanced econometric techniques. There is no regression analysis involved in Baily (1998) and Brynjolfsson and Smith (2000). Clay *et al*. (2002) use a fixed effect model to analyze data. For explanatory variables, they focus on the firms' pricing strategies. (e.g., loyalty program, past-purchase based recommendations and survey-based recommendations, etc.) In this chapter, I use panel data analysis (both fixed effect and random effect model). Nevertheless, as will be introduced in section IV, cost variables such as pages, hardcover dummy variable and months in print, are included as explanatory variables.

3.4 Data

Since the theoretical models developed by Butters (1977) and Robert and Stahl (1993) were dealing with homogeneous products, I choose book prices for study. Since books can be uniquely identified by their ISBN number, each book is considered as a homogeneous good.

I select six online bookstores that sell books over the internet and six physical bookstores that sell books through conventional outlets. Within the online bookstore industry, the two dominant players are Amazon and Barnes & Noble (BN.com). According to Net Ratings, these two firms account for more than 85% of online book sales (See Chevalier and Goolsbee 2003). The online bookstores include Amazon.com and BN.com, and four other internet booksellers that are randomly selected from Google's "online bookstore" search results. These four other internet booksellers are Buy.com, Overstock.com, Booksamillion.com and A1books.com. Google is believed to be unbiased because retailers are listed at no charge. Six physical bookstores include nation-wide chain stores Borders and Barnes & Nobles, and four other local bookstores in Tucson, Arizona. These four other local bookstores are Antigone, University of Arizona bookstore, Reader's Oasis and Clues Unlimited. Two "hybrid" bookstores that maintain operations both on the internet and in conventional outlets are included. Beside Barnes & Noble, Borders is another "hybrid" bookstore since the e-commerce division of Borders is teamed with Amazon. Consumers who visit Borders.com will be directly referred to Amazon's web page. For these two "hybrid" bookstores, the prices posted on the retailer's internet web page are quite different from the prices offered in conventional outlets. For example, comparing to Barnes & Nobles in Tucson, BN.com offers higher discount rates. One thing to note is the online bookstore Booksamillion.com. This company has physical stores in 18 states and the District of Columbia. However, it does not have any stores in Arizona, and a consumer in Tucson would not have the ability to buy from their physical location. Therefore, Booksamillion.com is considered as an internet retailer. The list of bookstores is summarized in Table 3.1.

I select twenty book titles. Ten book titles were randomly selected from the New York Times best-seller list during the week of January 12, 2004. The other ten book titles were randomly selected from the shelf of Bookman's Used Books Music, the largest store selling used books in Tucson. I chose this design to balance the goal of randomness and availability. Bookman's Used Books Music is not included in my study. The list of book titles is summarized in Table 3.2.

Book prices were collected during the week of January 12, 2004. The final price from selected online bookstores includes the price posted on the retailer's internet web page and the shipping cost. The shipping costs of the online bookstores are summarized in Table 3.1. In this study, the shipping costs for online orders do not vary with distance. In Tucson, the tax rate for books is 7.6%. However, textbooks are free of tax. Textbooks are excluded from my study. Therefore, the final price from the selected physical bookstores is equal to the book price plus 7.6% tax. In this study, the "price" of a book will refer to the final price. All internet book prices were collected from the retailer's web page. The data for all of the physical bookstores was gathered directly from the brick-and-mortar stores in my personal visits.

There were three books not carried by Overstock.com and one book not carried by A1books.com. For those physical bookstores, not all books on my list were on the shelf. However, the employees of the bookstore told me that the book was in stock at the warehouse and could be delivered within one week; therefore, I still treat the book as available. My sample consists of 116 online store observations and 120 physical store observations. For selected online bookstores, prices vary across stores. Most selected physical stores sell books at the publisher's list price. However, some of the selected physical bookstores offer discounts for books on the best-seller list. Since it was right after New Year's Day, Clues Unlimited was offering 10% off for all products in its store.

From the statistical tests shown in Table 3.3 and Table 3.4, we can see that internet book prices are lower than physical store book prices, and the coefficient of variation ($\frac{\sigma_x}{\bar{x}}$) of internet book prices is higher than the coefficient of variation of physical store book prices.

Let variable x_o represent online book prices and let x_c represent physical store book prices. Notice that $\ln\frac{\sigma_{x_o}}{\bar{x}_o} - \ln\frac{\sigma_{x_c}}{\bar{x}_c} = \ln\frac{\sigma_{x_o}}{\sigma_{x_c}} + \ln\frac{\bar{x}_c}{\bar{x}_o}$.

From Table 3.3 and Table 3.4, we have

$$\ln\frac{\sigma_{x_o}}{\bar{x}_o} - \ln\frac{\sigma_{x_c}}{\bar{x}_c} = \ln(0.1608) - \ln(0.0764) = 0.744 \text{ and}$$

$$\ln\frac{\bar{x}_c}{\bar{x}_o} = \ln(\frac{17.37}{14.73}) = 0.165\,.$$

Therefore,

$$\ln\frac{\sigma_{x_o}}{\sigma_{x_c}} = 0.744 - 0.165 = 0.579 \text{ and}$$

$$\ln\frac{\sigma_{x_o}}{\sigma_{x_c}} / (\ln\frac{\sigma_{x_o}}{\bar{x}_o} - \ln\frac{\sigma_{x_c}}{\bar{x}_c}) = 0.579/0.744 = 0.778\,.$$

After decomposition, we can see that approximately 78% of the difference in the coefficient of variation comes from the difference in variance.

3.5 Results

There are two dimensions, bookstores and book titles, in my data set. I use fixed effect model to analyze the data. Beside store effects, the explanatory variables include three cost variables and one preference variable. The pages of a book (PAGES) and the format dummy variable (HARDCVR) serve as proxies of the variable costs of a book. The value of HARDCVR is one if the book format is hardcover, and the value is zero if the book format is paperback. Months in print (MOINPRIN) proxies for the average cost. There is a high fixed cost when the publisher first prints a book. The longer the book is in print, the lower the average cost. The weeks on New York Times best-seller list (WEEKBS) represents reader's preference for a book. The book is more favored by readers as it stays on the list longer.

3.5.1 The fixed Effect Model

The store fixed effects are estimated in the fixed effect model (Model 1). The description of the model is as follows:

$$P_{ij} = \alpha_i + \beta_1 PAGES_j + \beta_2 HARDCVR_j + \beta_3 MOINPRIN_j + \beta_4 WEEKBS_j + \varepsilon_{ij} \qquad (3.5.1)$$

where *i* is the index of the stores and *j* is the index of the book titles. α_i's are the bookstore fixed effects. The regression results are reported in Table 3.5. All of the cost variables are statistically significant. The variable WEEKBS is statistically significant at the ten percent level. The coefficients of PAGES, HARDCVR, and WEEKBS are positive. The coefficient of MOINPRIN is negative. These results are consistent with the presumption that the higher the cost, the higher the price.

The regression results show that the store effects are highly significant. The online store effects vary across online bookstores. However, the physical store effects are very close to each other. Notice that the store effect of Reader's Oasis, a Tucson local independent bookstore, and the store effect of Antigone, Tucson's only feminist bookstore, are identical. The reason is that both bookstores sell books at publisher's list price, and there is no discount offered.

3.5.2 Online Bookstore Price Dispersion

In order to test whether prices vary across online bookstores, I impose the restriction that all online bookstores have the same impact on prices. The description of the model is as follows:

$$P_{ij} = \alpha_i \sum_{k=1}^{6} D_{ik} + \gamma ONLINE_i + \beta_1 PAGES_j + \beta_2 HARDCVR_j + \beta_3 MOINPRIN_j + \beta_4 WEEKBS_j + \varepsilon_{ij} \tag{3.5.2}$$

α_i's are the physical bookstore fixed effects. $D_{ik}=1$ when $k=i$, otherwise $D_{ik}=0$. ONLINE is a dummy variable and the value of ONLINE is one if the bookstore is online. The regression results of the restricted model (Model 2) are reported in Table 3.6. An F-test is used to test the restricted model versus the fixed effect model. The F-test statistics is 5.389. Since the F-test statistic is greater than 3.11, the critical F value at the 1% significant level, we reject the restricted model at the 1% level of significance. The test result suggests that different online bookstores have different impacts on prices.

3.5.3 Physical Bookstore Price Dispersion

In order to test whether prices vary across physical bookstores, I impose the restriction that all physical bookstores have the same impact on prices. The description of the model is as follows:

$$P_{ij} = \alpha_i \sum_{k=7}^{12} D_{ik} + \gamma LOCAL_i + \beta_1 PAGES_j + \beta_2 HARDCVR_j + \beta_3 MOINPRIN_j + \beta_4 WEEKBS_j + \varepsilon_{ij} \tag{3.5.3}$$

α_i 's are the online bookstore fixed effects. $D_{ik} = 1$ when $k = i$, otherwise $D_{ik} = 0$. LOCAL is a dummy variable and the value of LOCAL is one if the bookstore is physical. The regression results of the restricted model (Model 3) are reported in Table 3.7. Again an F-test is used to test the restricted model versus the fixed effect model. The F-test statistic is 0.834. Since the F-test statistic is less than 2.26, the critical F value at the 5% significant level, we fail to reject the restricted model at 5% significant level. The test result suggests that different physical bookstores have the same impact on prices.

3.5.4 Impact of the Internet

In this sub section we study whether the data are consistent with the hypothesis that the lower the search cost, the lower the price. Though we do not expect search costs in internet channels to be trivially small, the new technologies provided by the internet reduce search costs. The description of the model is as follows:

$$P_{ij} = \alpha + \gamma ONLINE_i + \beta_1 PAGES_j + \beta_2 HARDCVR_j + \beta_3 MOINPRIN_j + \beta_4 WEEKBS_j + \varepsilon_{ij} \tag{3.5.4}$$

The regression results on the impact of the internet are reported in Table 3.8. The internet has a negative impact on prices, and the impact is statistically significant (at 1% level). This empirical finding is consistent with the theoretical predictions provided by Butters (1977) and Robert and Stahl (1993).

3.5.5 Fixed Effect Model vs. Random Effect Model

For completeness, the results of the random effect model are reported in Table 3.9. The Hausman test shows that the difference between the variance-covariance matrix of the fixed effect estimator and the variance-covariance matrix of the random effect estimator is nearly singular. From Table 3.9 we can see that the random effect estimates are almost the same as the fixed effect estimates. Therefore, the random effect model cannot be rejected.

However, in the random effect model, the store effects are not identified. Therefore, we cannot use the random effect model to investigate our research hypothesis that the internet has a negative impact on average prices. As Marc Nerlove argued in his book *Essays in Panel Data Econometrics*, one should always take the random effect model whenever the random effect model cannot be rejected, since the fixed effect model is less efficient and the fixed effect model eliminates the time invariant variables. This chapter provides a counter example.

3.6 Conclusions and Discussions

There are two important observations from the regression results.

Observation 3.1: *The internet has a negative impact on average prices.*

This observation is consistent with the theoretical predictions. The previous empirical studies yield conflicting results. My finding is consistent with the study by Brynjofsson and Smith (2000), while contradicts the finding by Clay *et al* (2002).

These disparate findings can be explained by differences in data collection. Both Brynjofsson & Smith and Clay *et al.* select bestseller book titles from New York Times bestseller list. I did the very same thing. However, Brynjofsson and Smith select nonbestseller book titles in a way to ensure that most physical and online stores carry them, while Clay *et al* randomly select nonbestseller book titles by taking the first English-language entry on every 100th page from the Author volumes of Books in Print. I randomly select nonbestseller book titles from the shelf of Bookman's Used Books Music, the largest store selling used books in Tucson. The purpose of doing this is to balance the goal of randomness and availability. For physical bookstores, Clay *et al* only focus on large chains (Barnes & Noble and Borders), whereas Brynjofsson & Smith and my study consider a wider range of retail bookstores.

The differences in treating data also influence the results. The price comparisons in this chapter and the comparisons by Clay *et al.* are unweighted. Brynjofsson and Smith present comparisons weighted by a proxy of volume of sales. When comparing total price, Brynjofsson and Smith include transportation cost to physical stores ($0.53) and assume that the shipping cost per book is one-third of the shipping cost per order since most online shoppers buy three books at a time. I and Clay *et al* assume that online shoppers buy one book at a time and they go to physical stores for free. The way that I deal with data in this chapter falls between Brynjofsson & Smith and Clay *et al*.

Observation 3.2: *The variance of online store prices is higher than the variance of physical store prices.*

All of the previous empirical studies show significant online price dispersion. In this chapter I find the same result. All these findings contradict the theoretical prediction that the lower the search cost, the lower the level of price dispersion. As a theorist, one needs to seriously take this fact into account and provide convincing explanations.

One possible explanation is product differentiation. The availability of low-cost information on the internet leads all online sellers to charge similar prices for mass produced physical goods. In order to avoid this unpleasant outcome, firms seek to differentiate themselves in order to avoid directly competing on price. Product differentiation can explain the existence of online price dispersion. However, it is hard to use this explanation to account for why online price dispersion is higher than the price dispersion in physical stores, because physical stores also have the same incentive to differentiate their products. Moreover, the empirical analysis by Clay *et al.* yields no clear results on product differentiation for online book sales.

Another possible explanation is that reputation pays. Shoppers may be willing to pay a premium to online sellers that they have dealt with previously and the online sellers with a good reputation of shipping on time. Once the consumers herd to the leading online book sellers as Amazon.com and BN.com, it is hard for other online book sellers to steal customers from the leading firms even if they cut their prices. The impact of reputation may be a plausible explanation. On the internet the traders cannot see each other face to face. Therefore, the reputation, namely the probability that sellers ship the commodity on time, is extremely important for online transactions. Houser and Wooders (2001) examine eBay auctions and find that reputation has a statistically significant effect on price. Different online bookstores charge different prices because of the differences in reputation. Different physical bookstores charge similar prices because the probability of shipping is not a key issue for hand-to-hand transactions.

Heterogeneity in consumer search costs is the third explanation. Some consumers buy from the stores offering the lowest price, whereas other consumers shop on the basis of their previous shopping experiences and have much higher search cost. The former are shoppers, while the latter are captive buyers. The captive buyers are the group of people with high earnings and high opportunity cost of time. Stores face a tradeoff between posting low prices to attract shoppers and posting higher prices to exercise market power in the captive buyer market. The theoretical study by Stahl (1989) shows that if in the market there are both buyers with zero search cost (shoppers) and buyers with positive search cost (captive buyers), then there is a unique symmetric Nash equilibrium in mixed strategies. The Nash equilibrium price distribution changes continuously from the monopoly price to the competitive price as the fraction of zero search cost buyers varies from 0 to 1. The laboratory experiments set by Morgan *et al*. (2001) provide strong support for the comparative static prediction given above.

If one can find data set on individual consumer characteristics, providing evidence showing that online consumers are more heterogeneous than physical store consumers, then the empirical puzzle that the variance of online store prices is higher than the variance of physical store prices will be completely solved. This is an avenue for further research.

TABLE 3.1. Bookstore List

Retailers	Category	Shipping Cost	Sales Tax
Buy.com	Internet	3.90	0
Overstock.com	Internet	1.40	0
Booksamillion.com	Internet	3.98	0
A1books.com	Internet	2.99	0
Amazon.com (Borders.com)	Internet—Hybrid *	3.99	0
BN.com	Internet—Hybrid	3.99	0
Borders	Physical—Hybrid*	0	7.6%
Barnes and Nobles	Physical—Hybrid	0	7.6%
Antigone	Physical	0	7.6%
University of Arizona bookstore	Physical	0	7.6%
Reader's Oasis	Physical	0	7.6%
Clues Unlimited	Physical	0	7.6%

* Both Amazon and Borders are considered as "hybrid" since the e-commerce division of Borders is teamed with Amazon.

TABLE 3.2. Book Titles

Title	Author	Publisher	ISBN
COLD MOUNTAIN	Frazier	Vintage	0375700757
LIFE OF PI	Martel	Harvest/Harcourt	0156027321
SEABISCUIT	Hillenbrand	Ballantine	0449005615
UNDER THE TUSCAN SUN	Mayes	Broadway	0767900383
JOURNALS	Cobain	Riverhead	157322359X
THE DA VINCI CODE	Brown	Doubleday	0385504209
THE BIG BAD WOLF	Patterson	Little, Brown	0316602906
SKIPPING CHRISTMAS	Grisham	Doubleday	0385508417
BENJAMIN FRANKLIN	Isaacson	Simon & Schuster	0684807610
HISTORY OF EVERYTHING	Bryson	Broadway	0767908171
Girls' Guide to Hunting & Fishing	Bank	Penguin	0140293248
Cosmic Serpent	Narby	Tarcher	0874779642
Madame Sadayakko	Downer	Gotham	1592400051
Golden Bough	Frazer	Touchstone	0684826305
Gitanjali	Tagore	Scribner	0684839342
Soul Mountain	Gao	Perennial	0060936231
Greatest Freethinkers	Frank	Random	0375425853
The Third Wave	Toffler	Bantam	0553246984
Wealth of Nations	Smith	Modern Library	0679783369
Art of War	Sun	Barnes & Noble	1593080166

Note: The first 10 titles are from NY Times Best-Seller list (January 7, 2004).

TABLE 3.3. t-test on Mean Book Prices

Product Market	Online Bookstore Price Mean (P_O)	Physical Bookstore Price Mean (P_{CON})	Alternative Hypothesis	*t*-test Significance
Books	14.73	17.37	$P_O < P_{CON}$	0.0002

TABLE 3.4. F-test on Coefficient of Variations of Book Prices

Product Market	Online Bookstore Price Coefficient of Variation (C_O)	Physical Bookstore Price Coefficient of Variation (C_{CON})	Alternative Hypothesis	*F*-test Significance
Books	0.1608	0.0764	$C_{CON} < C_O$	<0.0001

TABLE 3.5. Store Fixed Effects

Dependent Variable: Price (including tax and shipping cost)

Variable	Coefficient	Standard Error	P-value
PAGES	.0076	.0011	.0000
HARDCVR	6.7813	.6322	.0000
MOINPRIN	-.0142	.0056	.0120
WEEKSBS	.0151	.0077	.0524
Store Fixed Effects:			
Buy.com	9.2205	1.1422	.0000
Overstock.com	5.3957	1.2156	.0000
Booksamillion.com	10.6140	1.1422	.0000
A1books.com	8.1919	1.1535	.0000
Amazon.com	10.0615	1.1422	.0000
BN.com	11.8670	1.1422	.0000
Barnes and Noble	11.5900	1.1422	.0000
Borders	12.0355	1.1422	.0000
Antigone	12.8885	1.1422	.0000
UofA	11.0149	1.1422	.0000
Reader's Oasis	12.8885	1.1422	.0000
Clues Unlimited	11.0575	1.1422	.0000

TABLE 3.6. The Restricted Model—All Online Bookstores Have the Same Impact on Prices

Dependent Variable: Price (including tax and shipping cost)

Variable	Coefficient	Standard Error	P-value
PAGES	.0076	.0012	.0000
HARDCVR	6.7590	.6621	.0000
MOINPRIN	-.0148	.0059	.0121
WEEKSBS	.0145	.0081	.0737
Physical Stores Effects:			
Barnes and Noble	11.6187	1.1965	.0000
Borders	12.0642	1.1965	.0000
Antigone	12.9172	1.1965	.0000
UofA	11.0437	1.1965	.0000
Reader's Oasis	12.9172	1.1965	.0000
Clues Unlimited	11.0862	1.1965	.0000
Online Stores Effect:			
ONLINE	9.3626	.8101	.0000

TABLE 3.7. The Restricted Model—All Physical Bookstores Have the Same Impact on Prices

Dependent Variable: Price (including tax and shipping cost)

Variable	Coefficient	Standard Error	P-value
PAGES	.0076	.0011	.0000
HARDCVR	6.7813	.6311	.0000
MOINPRIN	-.0141	.0056	.0119
WEEKSBS	.0151	.0077	.0519
Online Stores Effects:			
Buy.com	9.2205	1.1401	.0000
Overstock.com	5.3957	1.2133	.0000
Booksamillion.com	10.6140	1.1401	.0000
A1books.com	8.1919	1.1514	.0000
Amazon.com	10.0615	1.1401	.0000
BN.com	11.8670	1.1401	.0000
Physical Stores Effect:			
LOCAL	11.9125	.7689	.0000

TABLE 3.8. The Impact of the Internet

Dependent Variable: Price (including tax and shipping cost)

Variable	Coefficient	Standard Error	P-value
Constant	11.9412	.8047	.0000
PAGES	.0076	.0012	.0000
HARDCVR	6.7590	.6604	.0000
MOINPRIN	-.0148	.0058	.0118
WEEKSBS	.0145	.0081	.0729
Impact of Internet:			
ONLINE	-2.5787	.5622	.0000

TABLE 3.9. The Random Effect Model

Dependent Variable: Price (including tax and shipping cost)

Variable	Coefficient	Standard Error	P-value
Constant	10.5940	.8882	.0000
PAGES	.0076	.0011	.0000
HARDCVR	6.7764	.6322	.0000
MOINPRIN	-.0144	.0056	.0102
WEEKSBS	.0149	.0077	.0536

APPENDIX A: MATHEMATICAL DETAILS

Claim 1: Consider the one-ad pricing subgame in two-seller advertise-then-price game. Suppose seller 1 chooses to advertise and seller 2 chooses not to advertise. Then the unique search equilibrium in this one-ad pricing subgame is $F_1(p)=1-r/2p$, $p\in[0.5r, r)$, and $Prob(P_1=r)=0.5$ for seller 1; $F_2(p) = 2-r/p$, $p\in[0.5r, r)$ for seller 2, and $r = \min(v, c/(1-ln2))$ for the buyer. Seller 1's expected payoff is $0.5r$ and seller 2's expected payoff is $0.25r$.

Proof:

Seller 1's expected payoff from charging price p is $0.5\psi + 0.5[1 - F_2(p)]p$. $\psi = p$ when $p\leq \min(v, r')$, otherwise $\psi = 0$. r' is given by $c = \int_0^{r'} F_2(p)dp$. The first component is from the buyer that does not observe P_2 and the second component is from the buyer with full price information.

First, it must be true that $P^2_H < P^1_H$, since seller 2 always undercuts seller 1. It is also true that $P^2_L \leq P^1_L$. Suppose $P^1_L < P^2_L$. Then charging a price p, $P^1_L < p < P^2_L$, gives seller 1 payoff p. Therefore, seller 1's payoff is increasing in p when $P^1_L < p < P^2_L$. This contradicts that p is in the support of $F_1(P_1)$.

Then we show that the upper limit of seller 1's price distribution $F_1(P_1)$ is the lesser of the buyer reservation value v and r', $P^1_H = \min(v, r')$. The reasoning is as follows. If $P^1_H < \min(v, r')$, then charging P^1_H gives seller 1 payoff $0.5P^1_H$. However, by charging a price p slightly higher than P^1_H, $P^1_H < p < \min(v, r')$, sellers 1 gets payoff $0.5p$ which is strictly higher than $0.5P^1_H$. This contradicts the fact that P^1_H is in the support of $F_1(P_1)$. If $r' < P^1_H$ and $r' < v$, then seller 1's payoff from P^1_H is 0 (notice that $1-F_2(P^1_H) = 0$, given $P_2 < P^1_H$). However, charging price r' gives seller 1 strictly positive payoff. Again this contradicts the fact that P^1_H is in the support of $F_1(P_1)$. Finally it is obvious that $v < P^1_H$ and $v \le r$ is impossible, since any price higher than the reservation value will be rejected by the buyer. Therefore, $P^1_H = \min(v, r')$. It follows that seller 1's equilibrium payoff is $0.5P^1_H$, which equals to $0.5\min(v, r')$.

Now we can solve for the equilibrium. For seller 1's price p in the support of $F_1(p)$, we must have $0.5p+0.5[1 - F_2(p)]p = 0.5\min(v, r')$. Then $F_2(p) = 2 - [\min(v, r')]/p$, $p\in[0.5\min(v, r'), \min(v, r'))$. For seller 2's price p in the support of $F_2(p)$, we have $0.5[1 - F_1(p)]p = \pi$, where π is seller 2's equilibrium payoff. Since $p=0.5\min(v, r')$ belongs to the support of $F_2(p)$ and $P^2_L \le P^1_L$, $\pi = 0.25\min(v, r')$. Therefore $F_1(p)=1-[\min(v,r)]/2p$, $p\in[0.5\min(v, r'), \min(v, r'))$, and $Prob(P_1=\min(v, r'))=0.5$.

Solving for r' from $c = \int_0^{r'} F_2(p)dp$, we have $r' = c/(1-ln2)$.

Q. E. D.

Claim 2: Consider the one-ad pricing subgame in N-seller advertise-then-price game. In search equilibrium, the advertiser's expected payoff is r/N and the expected payoff to the other sellers is r/N^2, where $r = \min(v, r')$, $r' = \dfrac{N-1}{N-1-\ln N}c$.

Proof:

The proof is parallel to that of the two-seller case.

Let the advertiser's price distribution be F and let the nonadvertiser's price distribution be G. The buyer's cutoff price is given by $r = \min(v, r')$, where r' is determined by $c = \int_0^{r'} G(p)dp$.

In equilibrium, the buyer never searches for a new price: first, the upper limit of the support of F is $\min(v, r')$; secondly, the upper limit of the support of G will not exceed $\min(v, r')$, since the other sellers always undercuts the advertiser. The advertiser's expected payoff from charging price p is $\frac{1}{N}p + \frac{N-1}{N}[1-G(p)]p$. The first component is from the buyer that only observes the advertised price and the second component is from the buyer that observes the advertised price and one not advertised price. The advertiser's expected payoff from charging price p is $\frac{1}{N}p[1-F(p)]$. Solving for the mixed strategy search equilibrium, we have

$$F(p) = 1 - \frac{\min(v, r')}{Np} \text{ for } p \in [\min(v, r')/N, \min(v, r')), \text{ and } Prob(P_1 = \min(v, r')) = 1/N. \quad (1)$$

$$G(p) = \frac{N}{N-1} - \frac{\min(v, r')}{(N-1)p} \quad p \in [\min(v, r')/N, \min(v, r')). \quad (2)$$

$$r' = \frac{N-1}{N-1-\ln N}c \quad (3)$$

The equilibrium payoff to the advertiser is $\min(v, r')/N$, the equilibrium payoff to the other sellers is $\min(v, r')/N^2$, r' is given by equation (3).

Q. E. D.

APPENDIX B: INSTRUCTIONS TO SUBJECTS

Instructions to Sellers

(Advertise-then-Price Game)

Welcome to this experiment which concerns market decision making. The funding for this study has been provided by several foundations. The instructions are simple, and by following them carefully you can earn money. Your total profits during the experiment will be calculated in experimental dollars. Each experimental dollar is worth 0.5 USD. At the end of the experiment, all your experimental dollars will be converted to USD and paid to you in cash. In addition, you will be paid 5 USD for showing up.

1. Short Overview

There are six participants in this experiment. All participants will act as SELLERS. There are also buyers in the market, but their behavior will be determined by computer program. These automated buyers follow the shopping rule explained in section III.

As a seller in this market, you can use your computer to sell units of the good. Your computer screen will display useful information about selling opportunities. Remember that the information on your computer screen is **private**. Please do not talk with other market participants during the experiment.

Each time for buying and selling is called a TRADING PERIOD. At the start of each period, sellers make two decisions. A seller chooses a PRICE for his or her units, *and* decides whether or not to REVEAL his or her price. (The exact rules of how this is done are spelt out in section II below.) A seller who chooses price p earns PROFIT p on every unit sold that period. The number of units sold depends on the choices by other sellers in the market, and on the buyers' shopping rule. Sellers who do not sell any units earn a profit of zero that period.

At the end of the trading period, the computer screen will display your profits for that period and your total profits over all periods so far. Then the new trading period will begin. Everyone has new opportunities to sell each period; old units do not carry over into the next period. 20 trading periods are scheduled in the experiment.

2. Choices of Sellers

At the beginning of each period, the computer **randomly** divides participants into two markets. There are 3 sellers and 3 automated buyers in each market. In other words, as a seller, you are equally likely to be put into either market 1 or market 2. **Transactions occur within the market only**.

As a seller, before choosing your price, you need to decide whether or not to reveal your price. If you decide to reveal your price, please click the button "Reveal." If you decide not to reveal your price, please click the button "Not to Reveal." Once all sellers in your market have decided whether or not to reveal price, it is time for sellers in the market to choose prices. The screen will show the number of sellers in your market (which is three), and the number of sellers in your market who decided to reveal price. Please put your chosen price in the given text box, and click the "OK" button to confirm. The computer will wait for all other sellers in your market to choose their prices.

3. The Shopping Rule for the Automated Buyers

After all sellers have chosen prices, the automated buyers will shop. Each automated buyer demands at most ONE unit of the good, and all of the automated buyers follow the **same** shopping rule. However, the automated buyers are **independent**. The price, or the prices, an automated buyer sees does not depend on any other automated buyers' choices. For each automated buyer, there are three possible scenarios:

(a) If **none** of the sellers in that automated buyer's market decide to reveal price, then the automated buyer will "search." This means that one of the three not revealed prices will be **randomly** picked and therefore observed by that automated buyer. If the observed price is less than or equal to $2.00, then the automated buyer will purchase at the observed price. If the observed price is higher than $2.00, then the automated buyer will not purchase any unit that period.

(b) If **some** sellers in the market decide to reveal prices but the rest of the sellers in the same market decide not to reveal prices, then the automated buyer observes the revealed prices. If the lowest revealed price in the market is less than or equal to $0.30, then the automated buyer will purchase at that lowest revealed price. If the lowest revealed price is higher than $0.30, then the automated buyer will "search." This means that one of the not revealed prices will be **randomly** picked and therefore observed by that automated buyer. (If there is only one not revealed price, then that price will be observed for sure.) If the newly observed price is less than or equal to $0.30, then the automated buyer will purchase at that observed price. If the newly observed price is higher than $0.30, then all of the three prices in the market will be shown to that automated buyer and the automated buyer follows the shopping rule described in (*c*).

(c) If **all** of the sellers in that automated buyer's market decide to reveal prices, then all of the prices are observed by the automated buyer. If the lowest price in the market is less than or equal to $2.00, then the automated buyer will purchase at the lowest price. If the lowest price is higher than $2.00, then the automated buyer will not purchase any unit that period.

Finally, if there is a tie at the price at which the automated buyer is going to purchase, then one of the tying prices will be **randomly** picked and the corresponding seller will sell one unit to that automated buyer.

Do you have any questions about the instructions or procedures? If you have a question, please raise your hands and one of us will come to your seat to answer it.

Instructions to Sellers

(Advertise-with-Price Game)

Welcome to this experiment which concerns market decision making. The funding for this study has been provided by several foundations. The instructions are simple, and by following them carefully you can earn money. Your total profits during the experiment will be calculated in USD. At the end of the experiment, all your profits will be paid to you in cash. In addition, you will be paid 5 USD for showing up.

1. Short Overview

[Exact same text as in Advertise-then-Price Game]

2. Choices of Sellers

At the beginning of each period, the computer **randomly** divides participants into two markets. There are 3 sellers and 3 automated buyers in each market. In other words, as a seller, you are equally likely to be put into either market 1 or market 2. **Transactions occur within the market only**.

As a seller, you need to make two decisions at the same time. You need to choose your price, and decide whether or not to reveal your price. If you decide to reveal your price, then all of the automated buyers in your market will observe your price. If you decide not to reveal your price, the automated buyer will not observe your price unless they "search," which will be explained in section III. Please put your chosen price in the given text box. If you decide to reveal your price, please click the option button "Reveal." If you decide not to reveal your price, please click the option button "Not to Reveal." Finally, please click the "OK" button to confirm.

3. The Shopping Rule for the Automated Buyers

After all sellers have chosen prices, the automated buyers will shop. Each automated buyer demands at most ONE unit of the good, and all of the automated buyers follow the **same** shopping rule. However, the automated buyers are **independent**. The price, or the prices, an automated buyer sees does not depend on any other automated buyers' choices. For each automated buyer, there are three possible scenarios:

(a) If **none** of the sellers in that automated buyer's market decide to reveal price, then the automated buyer will "search." This means that one of the three not revealed prices will be **randomly** picked and therefore observed by that automated buyer. If the observed price is less than or equal to $2.00, then the automated buyer will purchase at the observed price. If the observed price is higher than $2.00, then the automated buyer will not purchase any unit that period.

(b) If **at least one** of the sellers in the market decides to reveal prices, then the automated buyer observes the revealed prices. If the lowest revealed price in the market is less than or equal to $2.00, then the automated buyer will purchase at that lowest revealed price. If the lowest revealed price is higher than $2.00, then the automated buyer will not purchase any unit that period.

Finally, if there is a tie at the price at which the automated buyer is going to purchase, then one of the tying prices will be **randomly** picked and the corresponding seller will sell one unit to that automated buyer.

Do you have any questions about the instructions or procedures? If you have a question, please raise your hands and one of us will come to your seat to answer it.

REFERENCES

Anderson, S., A. de Palma and J-F. Thisse, 1992. Discrete Choice Theory of Product Differentiation. Cambridge: MIT Press.

Bagwell, K., and G. Ramey, 1994. Coordination Economies, Advertising and Search Behavior in Retail Markets. *American Economic Review* 84: 498-517

Bailey, J., 1998. *Intermediation and Electronic Markets: Aggregation and Pricing in Internet Commerce*. Ph.D Thesis, Department of Electrical Engineering and Computer Science, MIT, May 20, 1998.

Benham, L., 1972. The Effects of Advertising on the Prices of Eyeglasses. *Journal of Law and Economics* 15: 337-52

Brynjolfsson, E., and M. Smith, 2000. Frictionless Commerce? A Comparison of Internet and Conventional Retailers. *Management Science* 46: 563-585

Burdett, K., and K. Judd, 1983. Equilibrium Price Dispersion. *Econometrica* 51: 955-969

Butters, G. R., 1977. Equilibrium Distributions of Sales and Advertising Prices. *Review of Economic Studies* 44: 465-91

Cady, J. F., 1976. An Estimate of the Price Effects of Restrictions on Drug price Advertising. *Economic Inquiry* 14: 493-510

Capra, M., R. Gomez, C. Holt and J. Goeree, 2000. Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition. forthcoming in the *International Economic Review*.

Cason, T., and D. Friedman, 2003. Buyer Search and Price Dispersion: A Laboratory Study. *Journal of Economic Theory* 112: 232-260

Chevalier, J., and A. Goolsbee, 2003. Measuring prices and price competition online: Amazon and Barnes and Noble. *Quantitative Marketing and Economics*, forthcoming.

Clay, K., R. Krishnan, E. Wolff and D. Fernnades, 2002. Retail Strategies on the Web: Price and Non-Price Competition in the Online Book Industry. *Journal of Industrial Economics* 50: 351-367

Coursey, D., R. M. Isaac and V. L. Smith, 1984. Natural Monopoly and Contested Markets: Some Experimental Results. *Journal of Law and Economics* 27: 91-113

Cox, J. C., and R. L. Oaxaca, 1989. Laboratory Experiments with a Finite-Horizon Job-Search Model. *Journal of Risk and Uncertainty* 2: 301-330

Diamond, P., 1971. A Model of Price Adjustment. *Journal of Economic Theory* 3: 158-168

Friedman, D., 1989. Producers' Markets: A Model of Oligopoly with Sales Costs. *Journal Economic Behavior and Organization* 11: 381-398

Grossman, G. M., and C. Shapiro, 1984. Informative Advertising with Differentiated Products. *Review of Economic Studies* 51: 63-81

The Handbook of Experimental Economics, John H. Kagel and Alvin E. Roth, editors, Princeton University Press, 1995. Paperback edition, Fall 1997.

Harrison, G., and M. McKee, 1985. Monopoly Behavior, Decentralized Regulation, and Contestable Markets: An Experimental Evaluation. *The Rand Journal of Economics* 16: 51-69

Holt, C., 1994. Oligopoly Price Competition with Incomplete Information: Convergence to Mixed-Strategy Equilibria. Working Paper.

Houser, D., and J. Wooders, 2001. Reputation in Auctions: Theory, and Evidence from eBay. Working Paper.

Kreps, D., and J. Scheinkman, 1983. Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes. *The Bell Journal of Economics* 14: 326-337

Kwoka, J. E. Jr., 1984. Advertising and the Price and Quality of Optometric Services. *American Economic Review* 74: 211-16

Lee, H. G., 1997. Do Electronic Marketplaces Lower the Price of Goods? *Communications of the ACM* 41: 73-80

Milyo, J., and J. Waldfogel, 1999. The Effect of Price Advertising on Prices: Evidence in the Wake of *44 Liquormart*. *American Economic Review* 89: 1081-96

Morgan, J., H. Orzen and M. Sefton, 2003. An Experimental Study of Advertising and Price Competition. *European Economic Review*, forthcoming.

Morgan, J., H. Orzen and M. Sefton, 2001. An Experimental Study of Price Dispersion. *Games and Economic Behavior*, forthcoming.

Nerlove, M., 2002. Essays in Panel Data Econometrics, Cambridge University Press.

Quesenberry, C. P., and D. C. Hurst, 1964. Large Sample Simultaneous Confidence Intervals for Multinomial Proportions. *Technometrics* 6: 191-195

Robert, J., and D. O. Stahl, II, 1993. Informative Price Advertising in a Sequential Search Model *Econometrica* 61: 657-86

Smith, V. L., 2002. Method in Experiment: Rhetoric and Reality. *Experimental Economics* 5: 91-110

Snee, R., 1974. Graphical Display of Two-way Contingency Tables. *The American Statistician* 28: 9-12

Stahl, D., 1989. Oligopolistic Pricing with Sequential Consumer Search. *American Economic Review* 79: 700-712

Stigler, G. J., 1961. The Economics of Information. *Journal of Political Economy* 69: 213-225

Tirole, J., 1988. The Theory of Industrial Organization, Cambridge, MA: MIT Press.

Weitzman, M., 1979. Optimal Search for the Best Alternative. *Econometrica* 47: 641-654