Magic on the Internet: Experimental Tests of Auction Theory

by

David H.L. Reiley

A.B., Astrophysical Sciences, Princeton University, 1991

Submitted to the Department of Economics in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1996

Signature of Author...
Department of Economics
May 16, 1996

Certified by...
Glenn D. Ellison
Ford Career Development Associate Professor of Economics
Thesis Supervisor

Accepted by ....
Richard S. Eckaus
Ford International Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Magic on the Internet: Experimental Tests of Auction Theory

by

David H.L. Reiley

Submitted to the Department of Economics on May 16, 1996, in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Abstract

Chapter 1: I present experimental evidence on the effects of minimum bids in first-price, sealed-bid auctions. I make use of a methodological innovation in the experimental study of auctions: the ability to run auctions for real goods in a preexisting market on the Internet, rather than for cash payoffs in the laboratory. The auction experiments in this chapter yield data on the effects of reserve prices on both the auctioneer's expected revenue and on the number of bidders who choose to participate in the auction. The benchmark theoretical model considered here is the classical auction model described by Riley and Samuelson (1981), with an exogenous number of symmetric, risk-neutral bidders with independent private values. The data verify a number of the predictions of classical auction theory, despite the fact that at least one of its assumptions is violated, as bidder entry is demonstrated to be endogenous in this market. Because of the violation of the classical assumptions, some attention is also devoted to the implications of more recent theoretical auction models.

Chapter 2: Vickrey's revenue equivalence theorem is one of the oldest and most basic results in auction theory, predicting under certain circumstances that all four basic auction types (English, Dutch, first-price sealed bid, and second-price sealed bid) should produce equivalent revenues for the auctioneer. This chapter contains revenue comparisons of all four basic auction types using field, rather than laboratory, data. The data, generated in field experiments that auctioned collectible trading cards over the Internet, confirm some laboratory findings while disconfirming others in this market. An important confirmation is that first-price and Dutch auctions raise considerably higher revenues than do second-price and English auctions, perhaps due in part to bidder risk aversion. An interesting contrast with laboratory experiments is the finding that Dutch auctions yield higher revenues than first-price auctions in these field experiments, while laboratory experiments have consistently found the opposite revenue ranking.

Chapter 3: The introduction of Tagamet in the United States in 1977 represented both a revolution in ulcer therapy and the beginning of an important new industry. Today there are four prescription H2-antagonist drugs: Tagamet, Zantac, Pepcid, and Axid, and they comprise a multi-billion dollar market for the treatment of ulcers and other gastric acid conditions. In this chapter, we examine the determinants of sales in this market, using a carefully constructed data set made possible by IMS America. We concentrate particularly on the marketing of these drugs to physicians through detailing (i.e., direct visits to physicians by pharmaceutical sales representatives) and medical journal advertising, and we make an innovative attempt to distinguish between "industry-expanding" and "rivalrous" marketing efforts. We find that the impact of total marketing on the expansion of overall industry sales declines as the number of products on the market increases. In addition, we find that the stock of industry-expanding marketing depreciates at a near-zero rate, while the stock of marketing oriented towards rivalrous market share competition depreciates at a 40% annual rate. We also find that the products' sales are affected significantly by price, quality attributes (such as FDA-approved indications and adverse drug interactions), and order of entry into the market.

Chapter 4: We examine empirically the role of information in facilitating and explaining growth of the overall anti-ulcer drug market, as well as in shaping the changing market shares of the four patented H2-antagonist drugs. In addition to the detailing and journal advertising considered in the previous chapter, an additional source of information we consider here is the direct-to-consumer advertising by the pharmaceutical firms. Additional factors in the econometric model include pricing, product quality attributes, and order-of-entry effects. We find price elasticity estimates close to unity, as one would expect from the Lerner monopoly markup rule. It is worth noting that when marketing variables are omitted from the relative demand equations, price elasticity estimates fall to about half these values. We find that marketing information stocks positively affect sales, that the sales elasticity is largest for detailing, followed by journal pages, and is smallest for direct-to-consumer advertising.

Thesis Supervisor: Glenn D. Ellison
Title: Ford Career Development Associate Professor of Economics

Acknowledgments

I am grateful to all those people whose advice, encouragement, and support enabled me to complete this thesis. First, I wish to acknowledge those teachers who inspired me to pursue a career in economics, starting with Alan Blinder, who, as one of the best professors I've ever had, instilled in me a great intellectual curiosity about economic behavior. I would also like to single out Ernie Berndt, Eduardo Engel, Jim Poterba, Nancy Rose, and Klaus Schmidt for their efforts as outstanding classroom teachers.

Glenn Ellison has been a terrific adviser. His guidance enabled me to focus my research ideas, and his criticism enabled me to achieve more than I would ever have thought possible. I am grateful to him for his dedication and commitment.

Ernie Berndt has been an invaluable mentor and friend during my years at MIT; he has helped me through rough times and has taught me much about the process of applied economics research. I benefited greatly from his insights, as well as those of Linda Bui and Glen Urban, during the process of our joint research on the pharmaceutical industry. A number of other people helped to make this pharmaceutical research possible, and we have thanked them elsewhere, but I wish to give special thanks here to Michele Lombardi for his enthusiastic research assistance, especially under deadline pressure.

I have come to view experiments as a very important tool for learning about microeconomic behavior. I am grateful to Matthew Rabin for first introducing me to the ideas of experimental economics, and to Rachel Croson and Vernon Smith for serving as invaluable mentors as I learned about the process of conducting experiments for myself. Preston McAfee provided important encouragement in my research, and taught me most of what I know about auction theory.

I would never have happened on the idea of running field experiments on the Internet without my friend Skaff Elias, who introduced me both to the game *Magic: the Gathering*, and to the online auctions which the game cards had inspired. He also served as an important source of information about the business of manufacturing and marketing the game. James Kirtley helped me debug my auction rules before I actually presented them to the experimental subjects. Marius Hauser, my young research assistant, was a great help with the sometimes overwhelming task of managing my card inventory. Florian Zettelmeyer got me started with the computer skills I needed to take advantage of the Internet as a research tool.

My friends at MIT have provided me with valuable emotional support, and have helped me to think more clearly about my intellectual interests. My heartfelt thanks go to economist friends Meghan Busse, Kip King, Lucia Nixon, Kim Rueben, and Florian Zettelmeyer, as well as non-economists Timothy Chow, Arthur and Grace Chang Mateos, Myrna Snyder, and Carson Schütze. As always, my parents and sister have been a constant source of support and inspiration.

Most of all, I wish to thank my wife M.K. Lucking Reiley for her support throughout my graduate-school career. She has supported my work both directly, as an occasional research assistant, and indirectly through love, compassion, and devotion. She is responsible in large part for my success, and I dedicate this thesis to her.

To my wife and best friend,
Mary Katharine Lucking Reiley

Contents

Introduction

This thesis is made up of two distinct parts. The primary part, consisting of the first two chapters, contains the results of my independent research on experimental tests of auction theory. Chapter 1 is an investigation of the effects of minimum bids on auction revenue and bidding behavior, while Chapter 2 presents evidence on the effects of the choice of the auction mechanism: English-style auction, Dutch auction, first-price sealed-bid auction, or second-price sealed bid auction. Both of these chapters involve data collected in controlled experiments, conducted by running my own auctions on the Internet for cards from the game *Magic: the Gathering*. In addition to providing interesting evidence on the predictions of auction theory, they also demonstrate the potential of the Internet as a tool for conducting experimental research in economics, and they contain some reflections on the advantages and disadvantages of "field experiments" as contrasted with traditional laboratory experiments in economics.

The second part of this thesis, Chapters 3 and 4, contains the results of research on pricing and marketing in the anti-ulcer pharmaceutical industry, from the earlier. Both of these pharmaceutical chapters were co-authored with Ernie Berndt, Linda Bui, and Glen Urban. Chapter 3, presented at a conference sponsored by the National Bureau of Economic Research in 1994, appears in *The Economics of New Products*, edited by Robert Gordon and Timothy Bresnahan (University of Chicago Press, 1996), while Chapter 4 appears in the May 1995 edition of *The American Economic Review*, in the papers and proceedings of the 1995 annual meetings of the American Economic Association. Permission to reprint these chapters has been granted by the National Bureau of Economic Research and the American Economic Association, respectively.

Chapter 1, "What are the Effects of Reserve Prices in Auctions? Evidence from Field Experiments," contains experimental evidence on the effects of minimum bids in first-price, sealed-bid auctions. My auction experiments yield data on the effects of reserve prices on both the auctioneer's expected revenue and on the number of bidders who choose to participate in the auction. The benchmark theoretical model considered here is the classical auction model described by Riley and Samuelson (1981), with an exogenous number of symmetric, risk-neutral bidders with independent private values. The data verify a number of the predictions of classical auction theory, despite the fact that at least one of its assumptions is violated, as bidder entry is demonstrated to be endogenous in this market. Because of the violation of the classical assumptions, some attention is also devoted to the implications of more recent theoretical auction models.

Chapter 2, "Tests of Revenue Equivalence in Internet Magic Auctions," describes my experimental tests of Vickrey's revenue equivalence theorem, one of the oldest and most basic results in auction theory. The revenue equivalence theorem predicts that under certain circumstances, all four basic auction types (English, Dutch, first-price sealed bid, and second-price sealed bid) should produce equivalent revenues for the auctioneer. This chapter contains revenue comparisons of all four basic auction types, replicating a number of previous laboratory studies of the revenue equivalence theorem. My field data confirm some of the findings previously generated with laboratory data, but disconfirm other laboratory findings. An important confirmation is that first-price and Dutch auctions raise considerably higher revenues than do second-price and English auctions, perhaps due in part to bidder risk aversion. The most interesting contrast with laboratory results is the finding that Dutch auctions yield higher revenues than first-price auctions in these field experiments, while laboratory experiments have consistently found the opposite revenue ranking.

In chapter 3, "The Roles of Marketing, Product Quality, and Price Competition in the Growth and Composition of the U.S. Anti-Ulcer Drug Industry," my coauthors and I model the monthly sales of anti-ulcer drugs over a sixteen-year period, using a carefully constructed data set made possible by IMS America. The four prescription H2-antagonist drugs -- Tagamet, Zantac, Pepcid, and Axid -- together comprise a multi-billion dollar market for the treatment of ulcers and other gastric acid conditions. We examine the determinants of sales in this market, beginning with the introduction of Tagamet, the first H2-antagonist drug, in 1977. We concentrate particularly on the marketing of these drugs to physicians through detailing (i.e., direct visits to physicians by pharmaceutical sales representatives) and medical journal advertising, and we make an innovative attempt to distinguish between "industry-expanding" and "rivalrous" marketing efforts. We find that the impact of total marketing on the expansion of overall industry sales declines as the number of products on the market increases. In addition, we find that the stock of industry-expanding marketing depreciates at a near-zero rate, while the stock of marketing oriented towards rivalrous market share competition depreciates at a 40% annual rate. We also find that the products' sales are affected significantly by price, quality attributes (such as FDA-approved indications and adverse drug interactions), and order of entry into the market.

Chapter 4, "Information, Marketing, and Pricing in the U.S. Antiulcer Drug Industry," is a follow-up to the study in Chapter 3. My coauthors and I include updated data through 1994, and in addition to the detailing and journal advertising considered in the previous chapter, we add data on direct-to-consumer advertising by the pharmaceutical firms. We examine empirically the role of information in facilitating and explaining growth of the overall anti-ulcer drug market, as well as in shaping the changing market shares of the four patented H2-antagonist drugs. Additional factors in the econometric model include pricing, product quality attributes, and order-of-entry effects. We find price elasticity estimates close to unity, as one would expect from the Lerner monopoly markup rule. It is worth noting that when marketing variables are omitted from the relative demand equations, price elasticity estimates fall to about half these values. We find that marketing information stocks positively affect sales, that the sales elasticity is largest for detailing, followed by journal pages, and is smallest for direct-to-consumer advertising.

Chapter 1:

What are the Effects of Reserve Prices in Auctions? Evidence from Field Experiments

1 Introduction

Although the theory of auctions has received a considerable amount of empirical study, both in the laboratory and with field data, to date very little has been written about the effects of minimum bids, also known in the theoretical literature as reserve prices. In this chapter, I perform experimental tests of the theory of reserve prices in first-price, sealed-bid auctions. I make use of a methodological innovation in the experimental study of auctions: the ability to run "field experiments" in a preexisting market. The auction experiments in this chapter yield data which verify some of the predictions of Vickrey's classical auction theory, including the impacts of reserve prices on the number of submitted bids, the probability of sale, and the revenue earned on goods sold. I also find that bidders exhibit sophisticated strategic reactions to changes in reserve prices, just as predicted by classical theory. These confirmations of classical theory come in spite of the fact that at least one of its assumptions is clearly violated in this market: the assumption of an exogenously-determined number of bidders. This demonstrated violation also leads me to consider some recent theoretical models with endogenous bidder entry.

During the past eighteen months, a fascinating new market has sprung up on the Internet. It is a market for collectible cards from the game *Magic: the Gathering*, a game whose retail success is a story in itself. Launched in August 1993, this product has already grossed hundreds of millions of retail dollars, and now has over a million players worldwide. There are more than a thousand distinct types of cards which have been printed for use in this game, each of which has a slightly different role in game play. In the game scenario, players assume the roles of dueling wizards, each with their own libraries of magic spells (represented by decks of cards) that may potentially be used against the player's opponent. Cards are sold in random assortments, just like baseball cards, at retail stores ranging from small game and hobby shops to large chains such as Toys 'R' Us and Waldenbooks.

Perhaps the most interesting aspect of this product (to an economist, at least) is its interaction with the Internet to create a thriving, online, secondary exchange economy, in which each Magic card is a separate commodity. The Internet, with its convenient methods of transmitting messages, has facilitated new technologies for trading among individuals. For example, a Magic player with a few hundred unwanted cards can auction them off to the highest bidders online, in a process which is considerably easier even than the traditional practice of holding a garage sale for excess household goods. The Internet lowers transaction costs, enabling ordinary individuals to make trades without the assistance of retail or auctioning specialists. Transaction costs are particularly low for cards as opposed to, for example, computer hardware, because it is so easy to mail cards across the country, and thus this card market has flourished.

In this chapter, I test some of the predictions of auction theory, taking advantage of the opportunity to run "field experiments" in this marketplace. I run a series of auctions of my own in order to investigate the theory of reserve prices. Two distinct experimental designs are developed in this chapter. The first design examines the effects of a binary variable: whether or not minimum bids were used. By auctioning the same cards twice, once with and once without minimum bids, this design exploits within-card variation to find the effects of the treatment variable on bidding behavior. The second design investigates the effects of variation in the *level* of reserve prices, and it does so by assigning different reserve levels (as a proportion of card market value) to different cards. This second design allows for an analysis of potentially nonlinear effects of reserve prices on, for example, the number of participating bidders and the auctioneer's expected revenue.

Past empirical studies of auctions can be divided into two distinct segments: laboratory experiments, and empirical studies using field data. Dozens of laboratory experiments on auctions have been published, dating back at least to Smith (1967); for a comprehensive review, see Kagel (1995). The experimental literature includes tests of Vickrey's revenue equivalence theorem in an independent-private-values context (see, for example, Cox, Roberson, and Smith (1982), as well as investigations of the "winner's curse" phenomenon in common-value auctions (see, for example, Kagel and Levin (1986)). A few very recent papers, such as Levin and Smith (1995) and Cox, Dinkin, and Smith (1995) have also investigated auctions with endogenous bidder entry.

Empirical work using field data, such as that from government auctions for offshore oil rights, has been more limited in scope, because of data restrictions. Hendricks and Paarsch (1995) provide an excellent summary of this literature, which has included tests of the theory of common value auctions with asymmetric information (see, for example, Hendricks and Porter (1988)), as well as structural estimation of the underlying probability distributions in independent-private-value auctions (see, for example, Paarsch (1992), and Laffont, Ossard, and Vuong (1995)).

The present study's methodology is a hybrid between the laboratory approach and the approach of traditional field research; hence the term "field experiment." This methodology shares with laboratory experiments the important advantage of allowing the researcher to control certain economic variables of interest (in this case, the existence and level of reserve prices), rather than leaving the researcher subject to the vagaries of the actual marketplace, where agents do not often produce exactly the variation that the researcher might desire. It also shares with traditional field research the advantage of studying agents' behavior in a real-world environment, rather than in a more artificial laboratory setting.

To my knowledge, this is also the first empirical study, either in the laboratory or field literatures, to focus on the effects of minimum bid levels. I find strong confirmation of several basic predictions of classical auction theory; increasing the reserve price decreases the number of bids received and the probability of selling the good, and increases auction revenue for goods which are actually sold. I also find that bidders react in a strategically sophisticated way to the existence of reserve prices, just as predicted by the classical Bayesian Nash equilibrium theory. However, I do find some evidence of phenomena that the classical theory fails to predict, and this leads to an investigation of the implications of more recent auction theories where bidder entry is costly, and therefore the number of potential bidders is not exogenously determined.

The chapter is organized as follows. The next section describes the history and institutional details of the marketplace where these field experiments took place. Section 3 presents some background on the theory of reserve prices in auctions. This includes a look at classical auction theory, which began with Vickrey (1961) and was summarized by Riley and Samuelson (1981), as well as some of the more recent theoretical literature on auctions. The recent papers include that of Bulow and Klemperer (1995), which considers the effects of additional bidder, as well as some of the literature on endogenous-bidder-entry auction models. Special attention is paid to empirically verifiable predictions.

Section 4 describes in detail the two experimental designs used in this study, and also includes a demographic description of the experimental subject pool. Section 5 contains a description of the results of the experiments, and the chapter finishes with conclusions and suggestions for future research.

2 History and Structure of the Market

Soon after the introduction of Magic, Internet users formed a newsgroup (rec.games.deckmaster) devoted to the discussion of this new game. In addition to discussions about the rules of the game and strategies for constructing decks, many of the messages on this newsgroup were buy, sell, and trade offers for individual cards. The terms of trade were agreed to via electronic mail, and then the transactions were carried out through postal mail. Messages devoted to economic trades soon overwhelmed the discussion group, so within a few months a new group (rec.games.deckmaster.marketplace, later renamed rec.games.trading-cards.marketplace) was devoted exclusively to the trading of these cards.

Both the quantity and the variety of messages posted to this newsgroup are stunning. By the spring of 1995, there were nearly 6,000 messages being posted each week, making rec.games.trading-cards.marketplace the highest-volume newsgroup on the entire Internet. According to the April 1, 1995 edition of the monthly publication *Top 40 Newsgroups in Order by Traffic Volume*, this newsgroup edged out the second-place newsgroup, misc.jobs.offered, by 3,000 messages per month. Even the most active discussion groups on, say, political issues, have fewer than half as many messages as rec.games.trading-cards.marketplace.

Approximately 90% of the 26,000 messages per month are devoted to the trading of Magic cards, with the remaining 10% devoted to the trading of cards from other games which have followed Magic into the booming market. Interestingly, the messages utilize several different kinds of market mechanisms. Some people post cards they are willing to trade, along with "wish lists" of cards they would be willing to accept in return, and solicit responses by private electronic mail. Others post fixed prices at which they are willing to sell cards for cash. Many hopeful sellers conduct auctions of their unwanted cards, and there are a plethora of different auction mechanisms which coexist simultaneously. Some auctioneers hold English-style auctions, and a small number hold first-price sealed-bid auctions. I have even seen one example of a Dutch auction, in which prices fall over time until the winner finally submits the first bid.

Many Magic auctions are hybrids of the English auction and the first-price sealed-bid auction. The sellers post daily updates of the highest price offered on each card for a set period of days, and the highest bidder before midnight of the specified final day is named the winner. These auctions are similar to the English auction at first, as the bids rise continuously from one day to the next, but they end up as a sealed-bid auction, as no one has the opportunity to raise their bid after the specified ending date, and no bidder knows what the other bidders' final bids are.

Another observed difference among the auction mechanisms is that some auctioneers include minimum bids, while others do not. In addition, some auctioneers post a "shutout price" for each card, a bid level at which the bidder automatically wins the card, and no other bids are accepted, thus capping the realized revenue at a certain level. This has the flavor of many transactions which take place in newspaper classified ads, such as for used cars or musical instruments, in which the seller may advertise "$500 or best offer." Finally, auctioneers differ from each other in the shipping fees that they charge to customers: some have a fixed shipping cost of, say $.40 per card purchased, while others have decreasing marginal fees for shipping each additional card. This could represent a passthrough of the auctioneer's own shipping costs, or it might also represent a form of second-degree price discrimination, designed to get bidders to bid higher in order to maximize their chances of getting multiple cards in the same auction, and therefore to minimize shipping costs per card.

This marketplace is a real-world laboratory for an empirical economist interested in auctions. Laboratory experiments, which to date have provided the vast majority of data on bidders' behavior in auctions, have occasionally been criticized on the grounds that subjects' behavior in an artificial laboratory environment may not be exactly the same as their behavior would be in the "real world." The Magic card market provides us with an opportunity to run controlled experimental auctions in the field rather than in the laboratory. It is worth noting that since in any given week there are dozens of auctioneers holding Magic auctions on the Internet, an experimenter in this market can be a "small player" who does not significantly perturb the overall market.

This technique of running "field experiments" has a financial advantage over traditional laboratory techniques. Since real people actually demand the goods being auctioned in the experiments, the experimenter need not pay the traditional monetary inducement to get subjects to participate.[1]

[1] For example, Levin and Smith (1995), with financial support from the NSF and Resources for the Future, paid well over $2000 in cash inducements to 40 individuals at the University of Houston for their participation in a series of 7 experiments.) By contrast, for the first set of experiments discussed in this chapter, I purchased cards for approximately $1600 and resold them for $2000, realizing a 25% return.

The trading of Magic cards has become just as enthralling an activity for some people as the actual playing of the game. A University of Washington student, Jason Black (who goes by the name "Cloister"), wrote a computer program in his spare time that automatically searches the marketplace newsgroup for each instance of each card name (with some tolerance for misspellings) and gathers data on the prices posted next to each card name in the newsgroup messages. It then computes trimmed means, standard deviations, quantiles, and so on, and automatically places this data on the Internet for interested parties to read. The Cloister price list, as it is known, is recomputed on a weekly basis, with each week's list archived for public use. Many people have adopted the Cloister list as a standard price reference when trading, and I make considerable use of this data in my research project.

3 Theoretical Background

In this chapter, I examine experimentally the effects of minimum bids in auctions. This section reviews a small subset of the theoretical literature on auctions with reserve prices, focusing on predictions which may be tested by the experiments in this study. For more comprehensive reviews of the theoretical literature on auctions, see Wilson (1992) and McAfee and McMillan (1987a).

3.1 Classical Auction Theory

In classical auction theory, which dates back to Vickrey (1961), the effects of minimum bids (reserve prices) are relatively straightforward. The setting of an optimal reserve price allows the auctioneer to use his bargaining power to extract a bit of additional profit from the highest bidder, above and beyond the profits which would result merely from competition between bidders. In effect, a minimum bid represents a take-it-or-leave-it offer from the auctioneer to the highest bidder.[2]

Let us consider an example, using the symmetric, independent-private-values framework proposed by Vickrey (1961, 1962). Suppose that there are N bidders in a first-price sealed-bid auction, and that each bidder's valuation v_i for the good is drawn independently

[2] Incidentally, one reason why the somewhat jargony term "reserve price" is more useful for theorists than the more colloquial "minimum bid" is that in some auctions, instead of imposing a minimum bid from the beginning, the auctioneer may instead use an implicit, "secret reserve price," or choose to watch everyone make their bids before setting a "reserve price" for the final take-it-or-leave it offer to the highest bidder.

from the same distribution, with probability density function f(v) and cumulative distribution function F(v). The distribution is common knowledge to the bidders and to the auctioneer, but the realization of each bidder's valuation is private information. (Thus, bidders' values are both private and independent.) Also assume that bidders are risk-neutral. If bidder i's bid b_i is the highest in the auction, then he wins the good, pays the amount of his bid, and realizes profits equal to $v_i - b_i$. If he is not the highest bidder, he realizes zero profits.

There exists a unique symmetric Bayesian Nash equilibrium to this first-price auction game. Each bidder's equilibrium strategy is given by a bid function b(v), and the equilibrium condition turns out to be the following first-order differential equation[3] which the function b(v) must satisfy:

$$\frac{db}{dv} = (N-1)(v-b)\frac{f(v)}{F(v)} \tag{1}$$

The auctioneer's reserve price r determines the boundary condition for this differential equation:

$$b(r) = r \tag{2}$$

This just states that when a bidder's valuation is equal to the reserve price, he will bid at the reserve price. A bidder with a valuation less than r will not bid at all. (Note that if there is no reserve price, then the boundary condition is just b(0) = 0.)

The differential equation can be solved by rearranging it into the canonical form for first-order linear differential equations, multiplying by the integrating factor $F(v)^{N-1}$, integrating (by parts), and substituting in the boundary condition. The solution is:

$$b(v;r) = \begin{cases} v - \dfrac{1}{F(v)^{N-1}}\displaystyle\int_r^v F(u)^{N-1}du & v \geq r \\ 0, & v < r \end{cases} \tag{3}$$

[3] See Vickrey (1962), Riley and Samuelson (1981), or Wilson (1992) for more details.

What is the expected revenue R to the auctioneer? The revenue in an auction is equal to the bid of the highest bidder, so R can be computed as the expectation of the bid function over the maximum of the bidder's valuations. The CDF of the maximum of N independent draws from a distribution with CDF $F(v)$ is equal to $F(v)^N$, so the expected revenue R(r) can be written as:

$$R(r) = \int_r^{\infty} N \cdot b(v;r) F(v)^{N-1} f(v) dv \tag{4}$$

3.1.1 Example: Uniformly-Distributed Valuations

For a concrete example, suppose that bidders' valuation are distributed according to a uniform distribution on [0,1]. Then the bid function is:

$$b(v;r) = \begin{cases} \frac{N-1}{N} v + \frac{r^N}{N v^{N-1}}, & v \geq r \\ 0, & v < r \end{cases} \tag{5}$$

Graphs of the function b(v) are shown in Figure 1. The first graph displays the case N=2, while the second displays the case N=5. In each graph, there are two lines displayed: one for an absolute auction (r=0), and one with reserve price r equal to 0.5. (As will be seen in Figure 2 below, it turns out that r=0.5 yields the optimal revenue for the auctioneer.) With zero reserve price, the bid function is linear, and the slope is equal to (N-1)/N, which explains the fact that as the number of bidders increases, the slope of the equilibrium bid function increases. With a nonzero reserve price, the bid function stays at zero for valuations less than the reserve price, and then jumps up to a curve (this time, a nonlinear one) whose slope is again increasing in the number of bidders.

The auctioneer's expected revenue, from equation 4, is given by:

$$R(r) = \frac{N-1}{N+1} + \left(1 - \frac{2N}{N+1} r\right) r^N \tag{6}$$

Graphs of the function R(r) for N=2 and for N=5 are shown in Figure 2. (Since the variation in revenue is so slight for the N=5 curve, the figure also includes a close-up of the N=5 curve over a very small segment of the vertical axis, in order to demonstrate that a peak actually exists.) Four features of these curves are worth noting. First, each curve gradually

Figure 1: Bid functions for the uniform distribution:

For the case N=2:

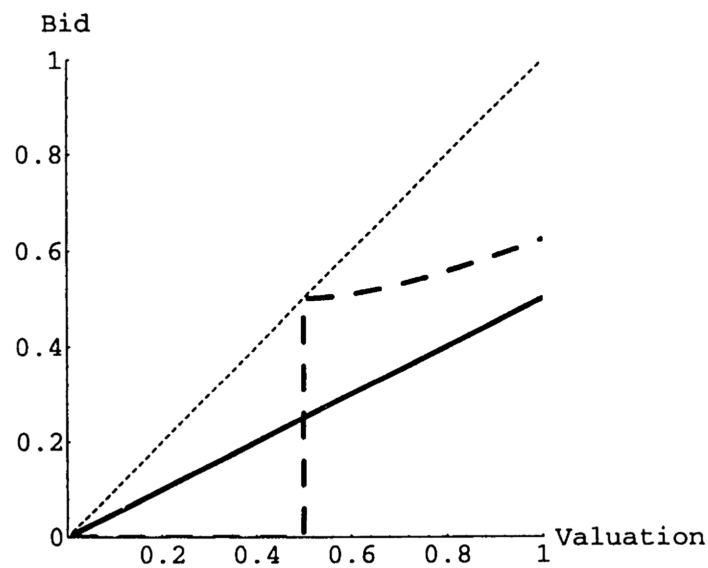

For the case N=5:

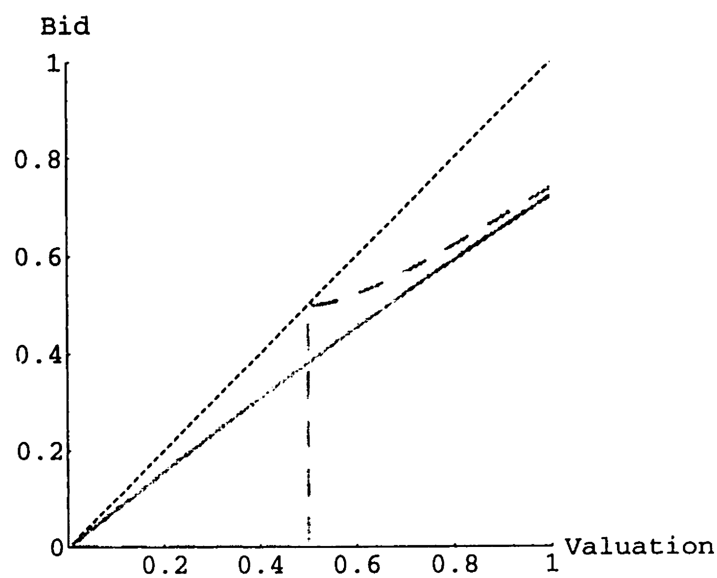

The solid lines represent the r=0 case, while the broken lines represent r=0.5.

Figure 2: Revenue curve for the uniform distribution.

The N=2 case is a black curve, while the N=5 case is a gray one:

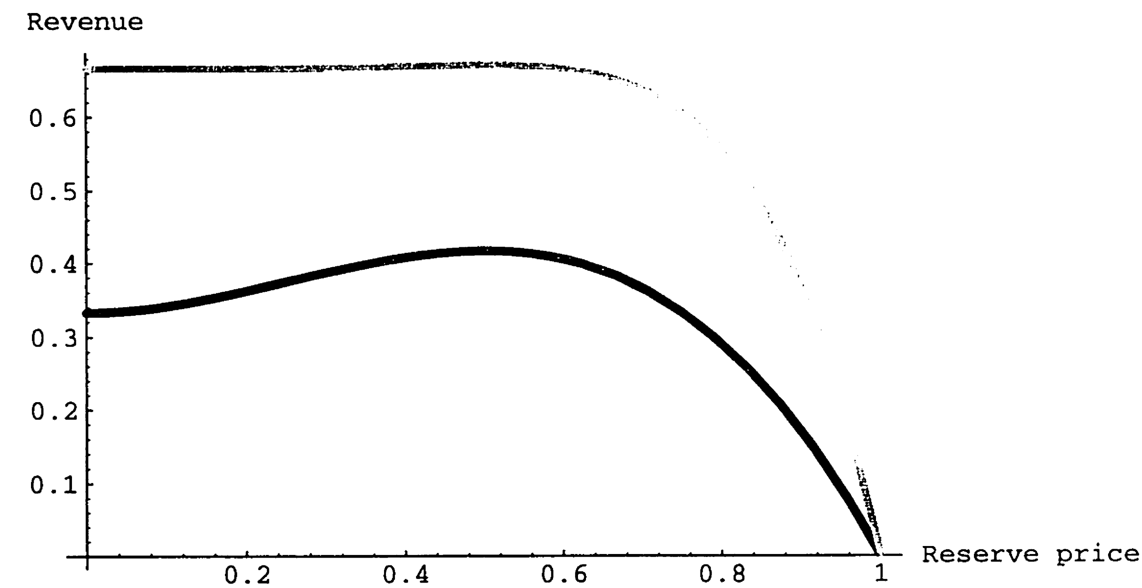

Vertical close-up of the optimum for N=5:

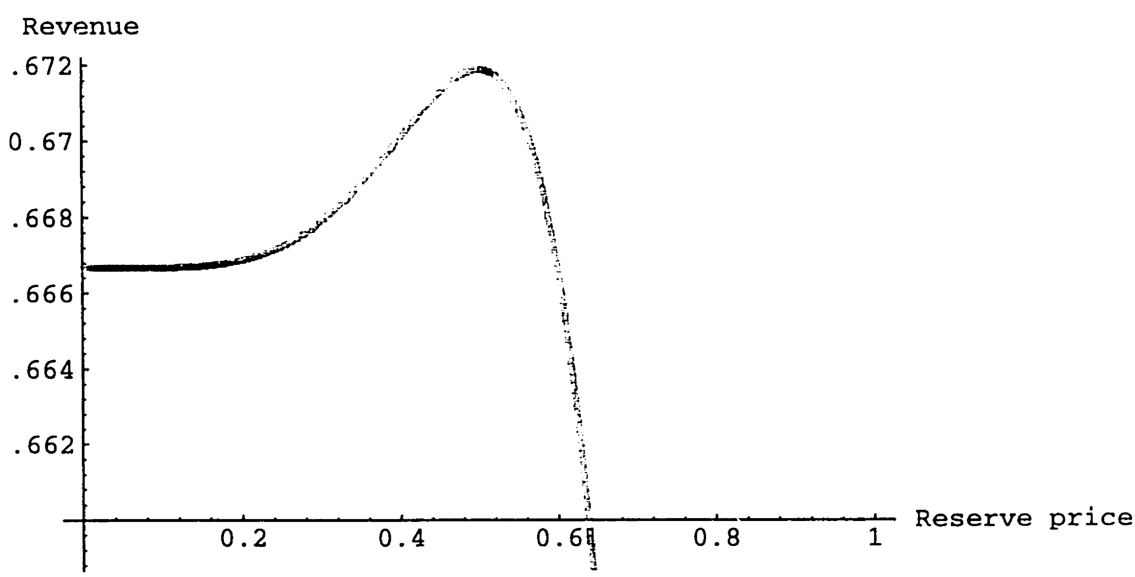

approaches its maximum at r=0.5, and then falls off relatively rapidly beyond the maximum. Second, the expected revenue gains from an optimal reserve price become very small as N increases. When the reserve price is increased from zero to its optimal value, auction revenue increases by 25 percent in the case N=2, but only by 0.78 percent in the case N=5. By the time the number of bidders reaches 10, the gains would be only 0.01 percent. (An intuitive explanation of this fact is that when the number of bidders is large (more than 5), reserve prices become superfluous to the revenue gained via competitive bidding pressures.) Third, the curve is flat at r=0, as R'(0)=0 for all N.[4] Finally, the curve stays flat much longer (to the left of the optimum) when N is large, becoming peaked only locally around the optimal reserve price.

3.1.2 Example: Exponentially-Distributed Valuations

The uniform distribution has finite support - in particular, no bidder in the above model could have a valuation greater than 1 - but in practice it might be that the distribution of valuations has no upper limit. For an example with unbounded support, then, consider bidder valuations drawn from an exponential distribution with exponential parameter $\beta=1$. In this case, the bid function still has a closed-form solution, but the functional form for general N is quite complicated and not terribly enlightening. The expected revenue R(r) also has a closed-form solution, but is similarly complicated. Rather than write out these equations, I will just display graphs of the functions. Figure 3 displays the bid functions for both

---

[4] It turns out that this result is not unique to the assumption of a uniform distribution of bidder valuations. Although I have not seen it proved previously, it is a general result of the independent-private-values model. The only assumption required is that there must not be a mass point at zero in the probability distribution of valuations: *F*(0)=0. The result can be proved by differentiating equation 4 with respect to *r*, using Leibniz' rule, which yields:

$$\frac{dR}{dr} = N \cdot F(r)^{N-1}[1 - F(r) - rf(r)]$$

Evaluated at r=0, this becomes:

$$\frac{dR}{dr} = N \cdot F(0)^{N-1}[1 - F(0)]$$

Therefore, $\frac{dR}{dr} \geq 0$ at r=0 always, and so long as F(0) is not equal to zero or one, the inequality is strict.

Also, assuming that there is a unique local maximum, it can be seen from the general expression for (dR/dr) that the revenue curve is either flat or increasing for all r less than the optimal r.

Figure 3: Bid functions for the exponential distribution.

For the case N=2:

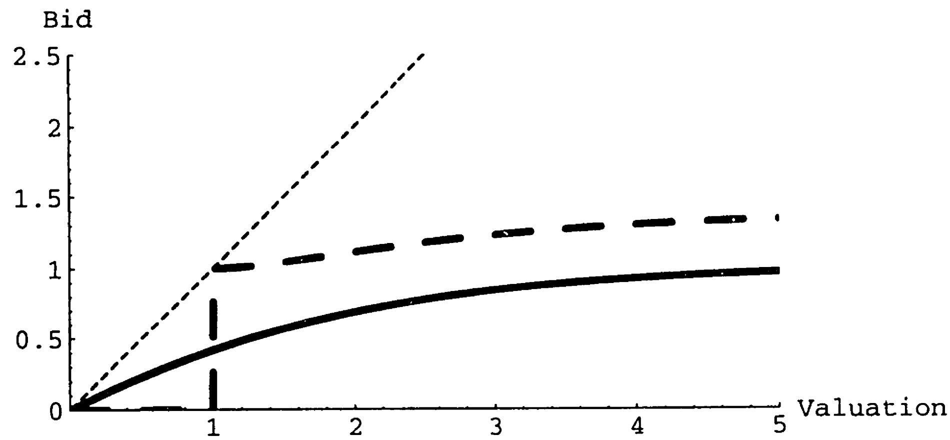

For the case N=5:

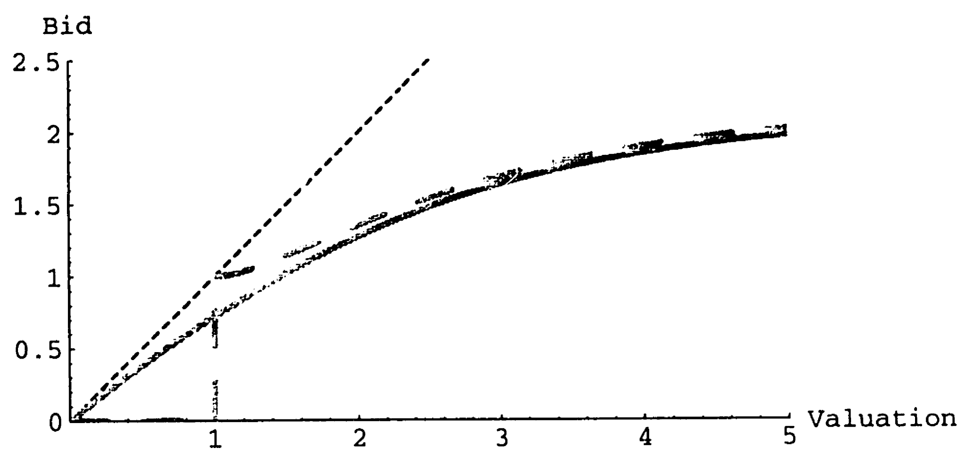

The solid lines represent the r=0 case, while the broken lines represent r=1.

N=2 and N=5, for the case with a zero reserve price as well as for the case where r=1 (which turns out to be the optimal reserve price for this example). Figure 4 displays the revenue curve R(r) for both N=2 and N=5.

Figure 4: Revenue curves for the exponential distribution.

The N=2 case is a black curve, while the N=5 case is a gray one:

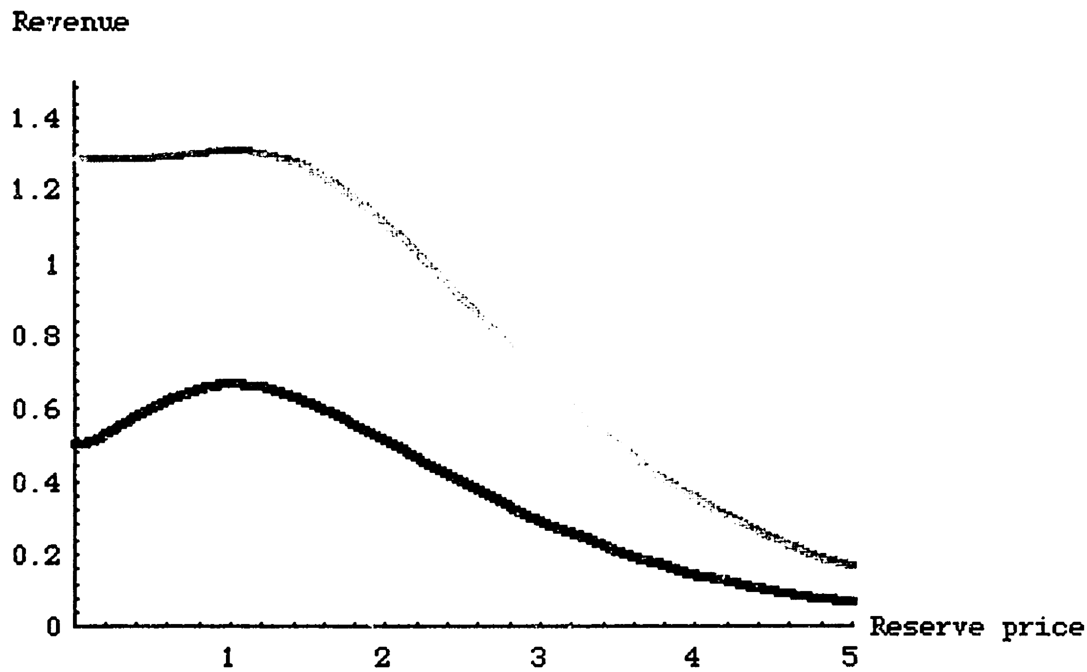

The same four properties noted for the uniform distribution still hold: the optimal reserve price (r=1) is independent of the number of bidders,[5] the expected revenue gains to the optimal reserve price fall off rapidly with larger N (from 33.6 percent at N=2 to 1.9 percent at N=5), the curve is flat at r=0, and the curve stays flat longer when N is larger.

Also worth noting are the differences between the uniform and exponential cases. First, the decline in revenue beyond the optimal reserve price is more gradual in the case of the exponential distribution. This is due in part to the fact that the support of the exponential distribution is infinite rather than finite - in the case of the uniform distribution, the revenue must go to zero at r=1, because there is zero probability that a bidder will value the good greater than 1. Second, the gains to be had with an optimal reserve price are more substantial (in percentage terms) under the exponential distribution, and the reduction in these gains with increased N is much less drastic.

[5] This turns out to be a general result of the IPV model. For any distribution F(v) of valuations, the optimal reserve price level is independent of the number of bidders. This was proved by Riley and Samuelson (1981). It follows directly from the first equation in footnote 4.

3.1.3 Testable Predictions of the Classical Model

The above examples illustrate several testable predictions about reserve prices which are true in general of the symmetric, independent-private-values auction model. The most basic predictions are that increasing the reserve price should decrease both the number of bidders who submit bids and the probability of selling the good, no matter what the distribution F(v) of valuations.

Another general prediction about bid levels is that $b(v;r+\Delta r) > b(v;r)$ when $v>r+\Delta r$ -- in other words, that a bidder whose valuation is high enough that he wants to participate in an auction with reserve price $r+\Delta r$ would submit a strictly higher bid when the reserve price is $r+\Delta r$ than when it is only r. This can be seen in Figures 1 and 3 from the fact that the dotted line (positive reserve price) is above the solid line (zero reserve price) for all valuations above the reserve price. It reflects the rational expectations of bidders in a Bayesian Nash equilibrium: if the reserve price is raised, then a bidder will realize that this will result in increased bid levels by the other bidders who choose to remain in the auction, and therefore his own optimal bid level will increase. Another statement of this prediction is that $\frac{\partial b}{\partial r} > 0$ for $v>r$.[6]

Classical auction theory also has predictions about the effects of reserve prices on auction revenue. One is that increasing the reserve price will increase the revenue earned on a good, conditional on its being sold. Several other properties can be identified by examining the auctioneer's overall expected revenue curve R(r). First, the optimal reserve price is independent of the number N of bidders. Second, the auctioneer's revenue curve is either flat or positively sloped for all reserve prices r less than the optimal reserve price, and the curve should be flat at $r=0$.[7] Beyond these general results which have been proved for all distribution functions F(v), there are also a couple of properties which were true of the uniform and exponential examples discussed above: as the number of bidders increases, an optimal reserve price should have less of an impact on expected revenue, and the revenue curve should become sharply negatively sloped for reserve prices greater than the optimal reserve.

---

[6] This can be proved by partial differentiation of equation 3.

[7] This is true unless there is a mass point of bidders with zero valuation, which seems unlikely. If a person valued the good at zero, then logically she should not be considered one of the N potential bidders.

3.2 Recent Theoretical Developments

Over the past fifteen years, there have been a number of important extensions to Vickrey's original IPV model of auctions. In this section, I single out two recent theoretical papers that have implications for the effects of reserve prices.

3.2.1 McAfee-Quan-Vincent: Reserve Prices with Endogenous Bidder Entry

A number of recent papers in auction theory[8] examine auctions where bidder entry is endogenous. In this literature, theorists recognize that the action of submitting a bid is not necessarily costless to the bidder, and therefore some potential bidders may choose not to enter the auction at all (even though these potential bidders value the good more than the seller does). An example of such a paper is that by McAfee, Quan, and Vincent (1995, hereafter MQV), in which the auctioneer chooses a reserve price and announces her auction, together with the level of her reserve price, to N potential bidders. Based on their knowledge of the reserve price, bidders can estimate the surplus they can hope to gain from bidding in the auction. They use this information to decide whether or not to incur the participation cost s. The equilibrium of the entry model is in mixed strategies, with each bidder's decision parametrized by a participation probability ρ which is identical across bidders. Finally, the auction is held among those bidders who decided to participate, and the good is awarded to the winner. If no bidder chooses to enter, or if no bidder who enters chooses to bid at least the reserve price, then the auctioneer keeps the good for herself and earns some outside option utility. In auction theory, this outside option value is commonly referred to as the good's "salvage value," as it is the value that the auctioneer would salvage from the good if she failed to sell it to any of the bidders.[9]

The main prediction of MQV is that the optimal reserve price is at least as high as the salvage value of the good. This is a testable prediction; raising the reserve price from some lower value to the expected salvage value of the good should raise revenues for the auctioneer.

---

[8] This literature begins with McAfee and McMillan (1987b).

[9] This value may come either from consuming the good herself, or from selling the good to someone who was not involved in the auction. It cannot be the amount that would be earned from offering the good a second time to any of the potential bidders from the original auction, because the theory of reserve prices assumes that the auctioneer is committed to keeping the good off the market if none of the bidders is willing to bid at or above the reserve price.

3.2.2 Bulow-Klemperer: The Benefits of an Additional Potential Bidder

Bulow and Klemperer (1995) show that under the general affiliated-values model of Milgrom and Weber (1982),[10] an absolute auction (no reserve price) with N+1 serious potential bidders will yield greater revenue to the auctioneer than an optimal auction (with an optimal reserve price announced only after the bid levels have been observed) with only N serious potential bidders. A "serious bidder," as defined by Bulow and Klemperer, is one with whom there would be gains from trade; i.e., the value of the good to the bidder is higher than the salvage value of the good to the auctioneer. Another way to state the Bulow-Klemperer theory is that gaining an additional serious bidder is more valuable to the auctioneer than the ability to replace an absolute auction with an optimal auction mechanism.

The theory of Bulow and Klemperer, together with my observations of other auctions in the market for Magic cards (many of which seemed to have set their minimum bid levels "too high"), have led me to propose my own theory. I hypothesize that when bidder entry is an important consideration, there may be an advertising value to absolute auctions, which causes them to generate more auction revenue than auctions with classically optimal minimum bid levels. I envision a three-stage model. First, the bidder notices the auction announcement and decides, based on a cursory inspection of both the good(s) available and the rules of the auction whether to bother with the auction at all. Second, if the bidder decides not to ignore the auction, she invests the time to gather additional information about her valuation for the good, and to decide on her optimal bid level. Third, assuming her optimal bid decision is above the reserve price, she submits her bid.

Suppose that at the first stage, the absence of reserve prices acts as a signal of potential bargains, inducing more bidders to commit to preparing bids. Then eliminating minimum bids altogether could be a profit-maximizing strategy for the auctioneer, by virtue of the additional serious bidders it would attract through a kind of "advertising effect."

In classical auction theory, the number of potential bidders is fixed at some number N, and the reserve price reduces the number of bidders by screening out only low-valuation bidders who wouldn't contribute much marginal revenue anyway. By contrast, I propose that minimum bids also have some probability of eliminating high-valuation bidders, at the stage where they make the decision whether to prepare a bid, and thus eliminating minimum bids altogether could be a profit-maximizing strategy for the auctioneer. A no-mini-

[10] The affiliated-values model developed by Milgrom and Weber (1982) is a generalization that includes as special cases both the private-values and common-values bidding paradigms.

mum strategy would open up the possibility of taking losses on some goods; however, if an auction with zero minimums is advertised, it is possible that more bidders will be induced to participate. Thus, although a few items may be sold at very low prices, they may serve as "loss leaders," similar to the goods advertised at deep discounts as weekly specials by supermarkets, enabling the auctioneer to collect higher revenues on average.

3.2.3 Testable Predictions from Recent Theory

In addition to tests of classical auction theory, I hope to provide tests of some of the extensions of classical theory. First, I wish to examine the main proposition of McAfee-Quan-Vincent, about the suboptimality of reserve prices lower than the salvage value of the good. Second, I hope to discover whether the absence of reserve prices carries an advertising value that would raise revenues by inducing entry by additional serious bidders.

4 Experimental Procedure

Two distinct experimental designs were developed for this chapter. The first design examines the effects of a binary variable: whether or not minimum bids were used. By auctioning the same cards twice, once with and once without minimum bids, it exploits within-card variation to find the effects of the treatment variable on bidding behavior. The second design investigates the effects of a continuous variable: the reserve price level (expressed as a fraction of the Cloister reference price). The across-card variation provides information that can be used to test predictions about the shape of the revenue curve R(r).

4.1 Within-Card Experiments

The first part of the data collection for this project consisted of two pairs of auctions which took place between February 24 and March 21, 1995. Each of the four auctions was a sealed-bid, first-price auction of several dozen individual Magic cards auctioned off individually. This simultaneous auction of many different goods at once, although not common in other economic environments,[11] is the norm for auctions of Magic cards on the In-

[11] Although simultaneous auctions are not traditional for familiar auctions, such as those of art, estate goods, or tulip bulbs, technological changes (such as the availability of computerized bidding) appear to be making them more common. In addition to the card auctions taking place on the Internet, a simultaneous-auction format was used for the recent FCC auctions of spectrum rights. See McMillan (1994) for details.

ternet. Running auctions in this simultaneous-auction format thus made the experiment as realistic and natural as possible for the bidders, who see many other similar auctions in the Internet marketplace for Magic cards.

Each auction lasted for one week, from the time the auction was announced to the deadline by which all bids had to be received. I announced each auction to potential bidders via two channels. First, I posted three announcements to the Internet newsgroup <rec.games.trading-cards.marketplace>, which is browsed every day by thousands of readers. For each auction, I posted a total of three newsgroup messages spaced evenly over the course of the week of the auction. Second, I solicited some bidders directly via email messages to their personal electronic mailboxes. My mailing list for direct solicitation was comprised of people who had already demonstrated their interest in participating in auctions for Magic cards by participation in previous ones.

The paired-auction experiment proceeded as follows. First, I held an absolute auction (no minimum bid) for 86 different cards (one of each card in the Antiquities expansion set), which took place over the course of a week. The subject line of the announcement read "Reiley's Auction #4: ANTIQUITIES, 5 Cent Minimum, Free Shipping!" so that potential bidders might be attracted by the unusually low minimum bid per card, essentially zero. (A 5-cent minimum is effectively no minimum, since the auction rules also required all bids to be in integer multiples of a nickel.) After the one-week deadline for submitting bids had passed, I computed the highest bid on each card. To each bidder who had won one or more cards, I mailed (electronically) a bill for the total amount owed.[12]

In addition to mailing bills to the individual winners, I also mailed a list of the winning bids to each bidder who had participated in the auction, whether or not they had won cards. This represented an effort to maintain my reputation as a credible auctioneer, demonstrating my truthfulness to those who had participated. Several bidders had specifically requested

---

[12] Although the standard practice in this marketplace is for auctioneers and other card sellers to charge buyers for postage and/or handling, I chose not to do this. I wanted bidders to bid independently, as much as possible, on each of the cards in which they were interested. Someone seriously interested in one card might decide to bid higher on a second card in the same auction than they would if the cards were auctioned independently, because they would like to spread out the postage costs per card by purchasing more than one card simultaneously from the same source. In addition, some of the cards I auctioned had rather low values, and I wanted to avoid having the card values be swamped by the cost of shipping. For example, if a bidder won a single card for 20 cents and then had to pay a fixed 50-cent shipping charge on top of that, the amount of useful information which could be derived from her bid would be rather suspect. Therefore, in the interests of keeping bid data as clean as possible, I decided to pay postage costs myself, and announced in advance that first-class shipping was included in the amount of each bid.

to see the overall results of the auction when it ended, so I was happy to perform this service for them in return for their help as participants in my experiment. I did not, however, give the bidders any explicit information about the number of people who had participated in the auction, or about the number of people who had received email invitations to participate.

After one additional week of buffer time after the end of the first auction, I ran the second auction in the paired experiment, this time with reasonably high minimum bid levels on each of the same 86 cards as before. The minimum bid levels were determined by consulting the standard (trimmed-mean) Cloister price list of Magic cards cited in section 2 of this chapter, and setting the minimum bid level for each card equal to 90% of the value of that card from the price list.

This contrast in minimum bid levels (zero versus 90% of the Cloister price list) was the only difference between the two auctions. Both auctions lasted exactly seven days, starting and ending on a Tuesday. The same 86 cards were up for bid in each auction. Each auction announcement was posted exactly three times to the marketplace newsgroup, and was emailed to primarily the same list of potential bidders.[13] Even the subject line of the announcements and mailings was kept identical, except that in the second auction, the words "5 Cent Minimum" were removed. By keeping all other conditions identical between the two auctions, I attempted to isolate the effects of minimum bids on potential bidders' behavior.

One condition that could not be kept identical, unfortunately, was the time period during which the auction took place. Because the two auctions took place two weeks apart, there were potential differences between the auctions that might have affected bidder behavior. First, the demands for the cards (or the supplies by other auctioneers) might have changed systematically over time, which is a realistic possibility in such a fast-changing market as this one.[14] Second, since the auctions shared many of the same bidders in common, the cards sold in the first auction may have affected the demand for the cards sold in the second auction.[15]

---

[13] The discrepancy consists of several people who indicated after receiving the first auction announcement that they were no longer interested in any more auctions, and wished to be removed from my mailing list. Since they demonstrated a complete lack of interest in participating, they were not really potential bidders after all. Thus, for purposes of the experiment, it does not really matter that they received mailings inviting them to the first auction but not the second.

[14] For example, certain cards from the Arabian Nights expansion set increased in value by a factor of ten during their first year out of print. It turns out that market prices for cards were actually rather stable during the month in which this experiment was conducted, but I did not know *a priori* what was going to happen to card prices.

To control for such potential variations in conditions over time, I simultaneously ran the same experiment in reverse order, using a different sample of cards. This second pair of auctions each featured the 78 cards in the Arabian Nights expansion set, with minimum bids present in the first auction but absent in the second. Just as before, minimum bids were set at ninety percent of the market price level from the Cloister price list. The first auction in this pair began three days after the start of the first auction in the previous pair, so that the auctions in the two experiments overlapped in time but were offset by three days. Also, I used a larger mailing list for my email announcement in this pair of auctions (232 people) than I had for the previous pair of auctions (50 people), with the first mailing list being a subset of the second mailing list. Otherwise, all other conditions were identical between the two pairs of auctions.

---

[15] For example, suppose that Bidder X is very anxious to obtain a Guardian Beast card for her deck, so that her valuation of the card is much higher than that of any of the other bidders in the experiment. She then wins the card in the first auction, and then has zero demand for that same card in the second auction, since she only really wants one copy. If this is generally the case for most cards, that the highest-value bidders in the sample are screened out in the first auction, then we might expect to see systematically lower revenues in the second auction.

The reader may be interested in the size of the stakes involved in this experiment. Figures 5 and 6 display graphically the distribution of minimum bid amounts (determined from

Figure 5: Distribution of minimum bids in Auction #8 (Antiquities).

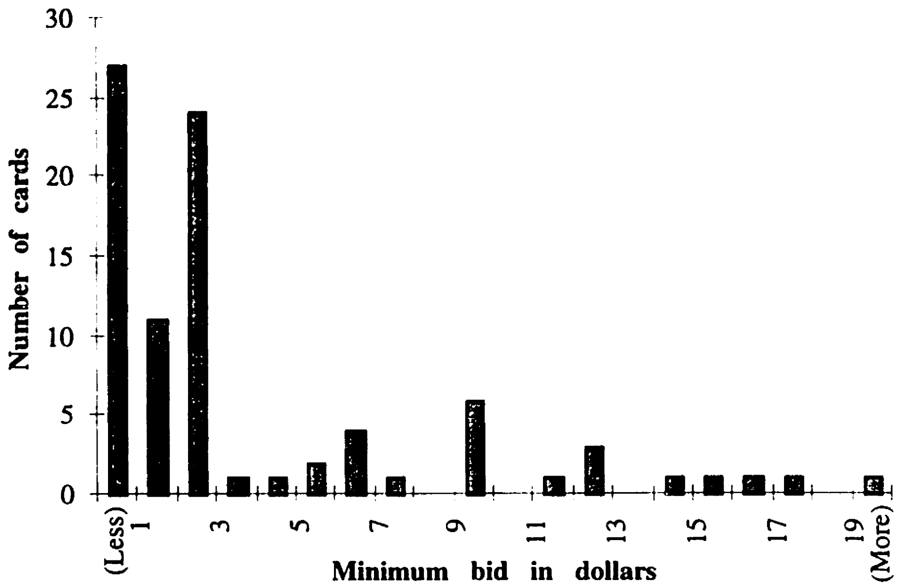

Figure 6: Distribution of minimum bids in Auction #5 (Arabian Nights).

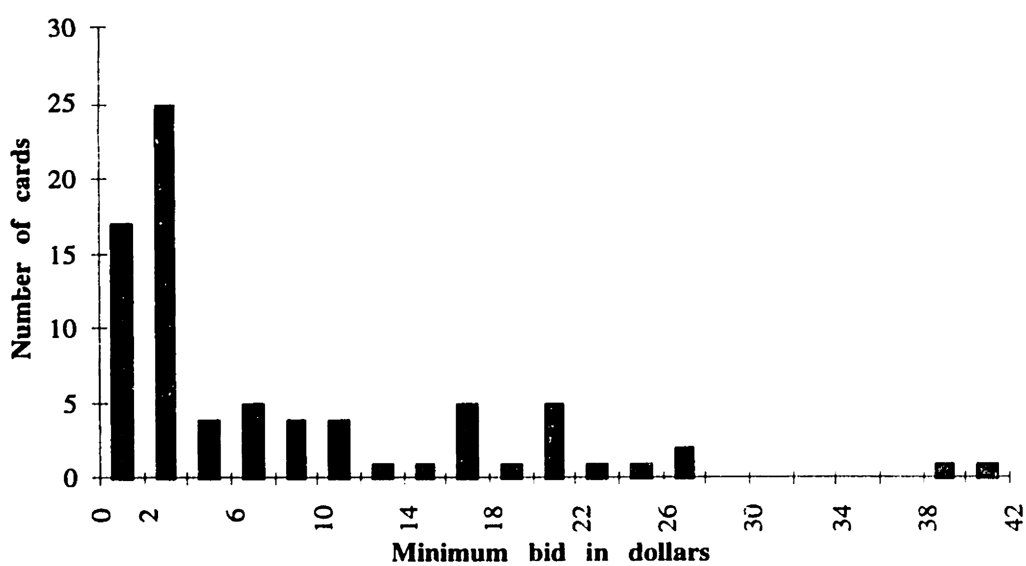

the aforementioned 90% rule) for each of the two sets of cards. The Arabian Nights cards tended to be more valuable, as they came from an earlier, more limited printing, and thus were in scarcer supply. The maximum posted reserve price in Auction #5 (Arabian Nights) was $40.10, while the maximum in Auction #8 (Antiquities) was only $19.85. Similarly, the median reserve price in Auction #5 was $3.65, compared with a median of $2.25 in Auction #8. While there were cards whose minimum bid values went as low as 30 cents, there were also dozens of cards valued at over ten dollars. The distributions of values illustrated in Figures 1 and 2 were quite standard for the Internet marketplace environment; in fact, they were probably higher than average.

A sample auction announcement, as it appeared to the potential bidders both in electronic mail and in the marketplace newsgroup, is displayed in the Appendix. For a review of the different auctions used in this experimental design, see Table 2 on page 41, which also includes summary statistics for each auction

4.2 Between-Card Experiments

A second set of experiments was designed to examine the effects of changes in the *level* of the reserve price, rather than merely changes in the *existence* of reserve prices. Four auctions (designated as Auctions #15, #16, #17, and #19) took place between October 3 and October 30, 1995. Again, each auction was a first-price, sealed-bid auction with a one-week timeframe for the submission of bids. Each was a simultaneous auction of 99 different items, this time with no overlap of items between auctions, and each card had a posted reserve price. Just as before, each auction was announced via three posts to the relevant newsgroup, as well as via email to a list of bidders.[16]

The reserve price for each card was chosen as a fraction of the current Cloister price of that card. In each of the first two auctions, nine cards were auctioned at a minimum bid of 10 percent of the Cloister price, nine at 20 percent, nine at 30 percent, and so on, up to a maximum of 110 percent of the Cloister price. After a quick analysis of the data from those auctions, I decided it was important to collect more data both at very low and at very high reserve price levels. Therefore, the third and fourth auctions were designed to have equal numbers of cards auctioned at reserve levels of 10, 20, 30, 40, 50, 100, 110, 120, 130, 140, and 150 percent of the Cloister price.[17]

[16] For this series of auctions, the bidder pool was quite a bit larger than before. 531 individuals were emailed to participate in Auction #15, and as some people specifically requested to be removed from my auction announcement mailing list, the list dropped to 489 individuals by the time Auction #19 began.

This variation in reserve price levels was designed to investigate how both bidder behavior and expected auction revenue would react to changes in the reserve price. It should also be possible to use this data to make a calculation of the optimal reserve price level. Normalizing by the Cloister price, since this is a standard reference price computed in the same way for all Magic cards, will make cross-card comparisons feasible.

Besides the exceptions noted above, all experimental protocols and bidder instructions were kept identical to those used in the auctions with reserve prices in the experimental design described in section 4.1.

4.3 Subject Pool Demographics

What are the demographics of the bidders who participated in these auctions? I did not ask bidders to fill out questionnaires or provide specific demographic information about themselves, because I did not want such information-gathering requirements to scare away bidders or to detract from the real-world-auction nature of the experiment. However, one can infer some demographic data from the electronic mail addresses of the participants. Primarily, one can tell whether the bidders were associated with a university or whether they were obtaining their access to the Internet through a commercial company. Table 1 displays this information for the first four auctions in this study, broken down by auction as well as aggregated over the entire bidder pool from the four auctions.

Table 1: Electronic-mail demographics of bidders in the within-card experiments.

	Auction #4	Auction #8	Auction #5	Auction #9	Total Bidder Pool
Educational (.edu)	10	3	17	33	55
Commercial (.com)	7	4	20	23	34
Other US (.gov,.mil, .org)	2	0	4	2	7
International	0	0	1	4	5
Total	19	7	42	62	101

[17] In my haste to run these follow-up auctions, I made some errors in computing the minimum bid levels, so these numbers were not quite equal in practice. Some cards were mistakenly assigned reserve levels of 60 percent and 90 percent in these auctions, but this does not compromise the integrity of the data.

In aggregate, there were 101 bidders who submitted bids in at least one of the first four auctions in this study. Of these, 55 bidders were students, faculty, or staff at American colleges and universities, as evidenced by the ".edu" domain name in their email addresses. Some of the universities represented in this experiment were the University of Chicago, the University of North Dakota, the University of Vermont, Penn State University, Stanford, Harvard, Berkeley, the University of Oklahoma, Lycoming College, and the University of Delaware.

An additional 34 participants were bidding from email addresses at commercial companies (".com"). These include employees using their email accounts at the companies they work for; I was able to identify bidders at AT&T, Motorola, and Rockwell, as well as at several smaller engineering and consulting companies. The ".com" designation also includes people who subscribe to commercial on-line services such as CompuServe, Delphi, and America Online, as well as smaller Internet service providers such as PrimeNet and Netcom. At least 25 of the 34 bidders in the commercial domain were subscribers to such services.

Another seven bidders sent their messages from the government (.gov), military (.mil), and nonprofit organization (.org) domains. These included employees of government agencies and national laboratories. Finally, there were also five participants in this experiment from other countries: Canada, France, and the Netherlands. I had actively tried to discourage international bidders from participating in my auctions, because of the complications involved in shipping cards outside the United States. I restricted the distribution of my mailing list to domestic email addresses, and promised free shipping only by first class mail within the United States, but still several enthusiastic foreign bidders submitted bids. I felt it would be inappropriate to refuse them, as I had not expressly prohibited international bids, and I noticed that international bidders were frequent participants in other Magic auctions on the Internet.

It is interesting to note that the bidders participated in this experiment from very diverse geographic areas. This is in marked contrast to most laboratory experiments, but very similar to many auctions for wines and fine art, where bids by fax and telephone have become common. Geographic locations cannot always be determined from email addresses, so I examined the postal return addresses on the payment envelopes mailed to me in Auction #5. This single auction included winners from Illinois, New York, Utah, Washington, Minnesota, northern and southern California, Virginia, Massachusetts, Texas, Oregon, Rhode Island, and Oklahoma.

One of the strengths of this study is its subject pool. The bidders in these experiments have wide diversity in their geographic and professional backgrounds, but share an intense interest in Magic cards. Their keen interest in the auctioned goods makes them very representative of people who bid in auctions in general, and indicates that the results obtained in this study will be applicable to auctions for other types of goods, from fine art to farm equipment.

5 Results

5.1 Within-Card Experiments

Table 2 shows a set of summary statistics for each of the four auctions in the within-card experiments.[18] The auctions are displayed in two pairs: first Auctions #4 and #8, for the 86 Antiquities card items, and then Auctions #5 and #9, for the 78 Arabian Nights card items. Auctions #4 and #9 were with no minimum bids, while Auctions #5 and #8 had sizable minimums (equal to 90% of the market price).

The table contains quite a bit of descriptive information about the auctions, including the number of participating bidders, the number of bids received, and the total payments received from winning bidders. I wish to single out a two key points. First, "real money" was involved in the auction transactions. Of the 73 different bills I sent to winning bidders over the course of the experiment, the median payment amount for each auction was between \$10 and \$24. A few individual payments even exceeded \$100.[19]

Second, in each auction there are multiple winners. The number of winners in each auction ranges from 6 to 25, and the fraction of bidders who win at least one card is between 40 percent and 86 percent. In each auction, the median number of cards won by each win-

[18] It may seem odd that the auctions are not numbered sequentially. The reason is that these four auctions were part of a series of auctions run for a larger research program, so the auctions numbered 1 through 4 and 6 through 7 were irrelevant to this project. This had two advantages where the experimental design is concerned. First, it helped avoid drawing bidders' attention to the point of my research. (For example, Auction #3 was an English auction, while Auction #7 was a second-price auction.) I feared that if they knew I was looking for the effects of reserve prices, it might distort their behavior (for example, they might consciously try to bid consistently from one auction to another). Second, it has the effect of making bidders unsure what I will do next. In particular, I didn't want bidders to expect that I would always auction the same card twice, for it might distort their behavior if they knew they would have a second chance to bid on the same card.

[19] For a graphical display of the distribution of winners according to the number of cards they won, and according to the amount of revenue they generated, see Figures 7 and 8. These express the total number of cards and the total amount paid, over all four auctions, by each individual.

Figure 7: Distribution of number of cards per winner, totalled over four auctions

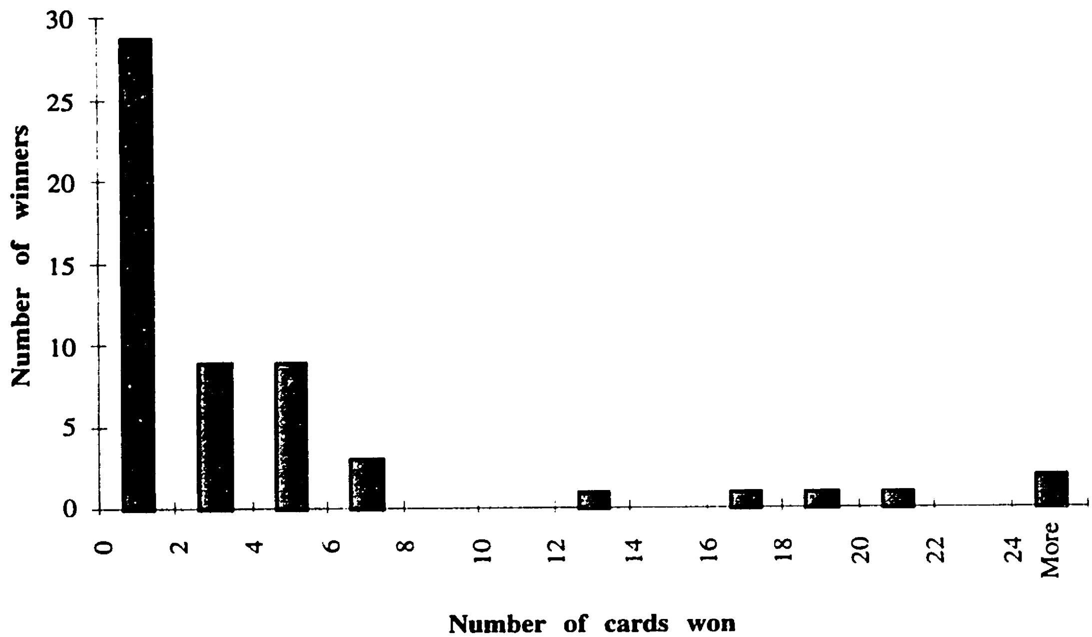

Figure 8: Distribution of revenue per winner, totalled over four auctions.

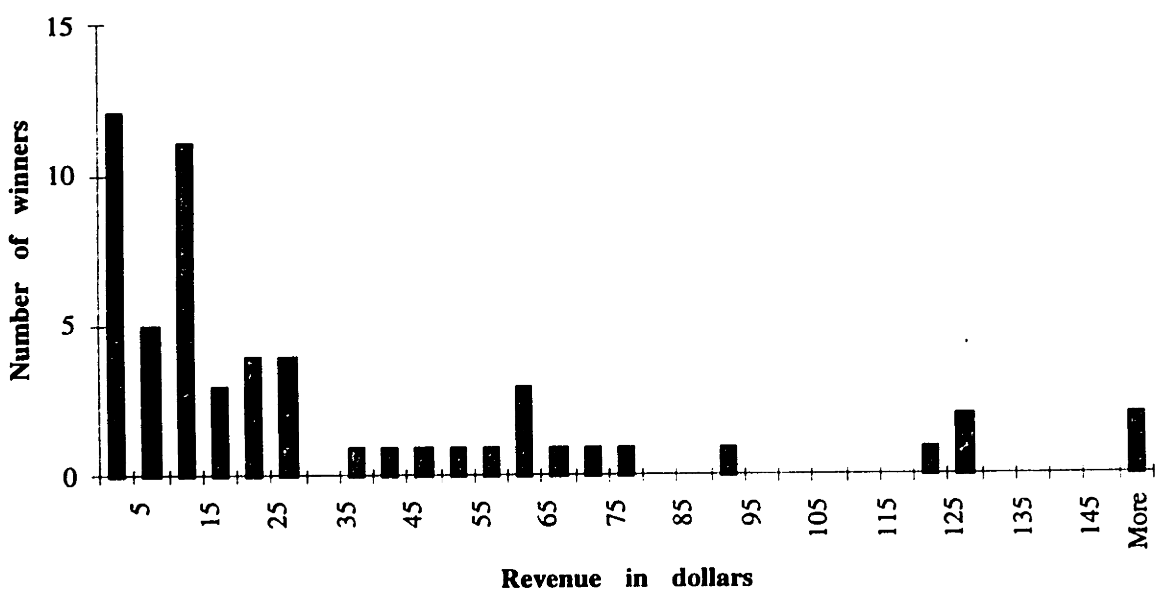

ner is between 2 and 3.5, while the maximum number of cards won by a single bidder ranges from 12 to 26. Except in Auction #8, no winner won more than 29 percent of the cards sold in any single auction. (In Auction #8, participation was very low: only 7 people submitted bids, 6 of whom won at least one card, and 39 of the cards went unsold.) The biggest

Table 2: **Summary statistics for within-card experiments.**

	Auction #4	Auction #8	Auction #5	Auction #9
Minimum bids?	No	Yes	Yes	No
Card set	Antiquities	Antiquities	Arabian Nights	Arabian Nights
Start date	Fri, 24 Feb	Fri, 10 Mar	Tue, 28 Feb	Tue, 14 Mar
End date	Fri, 3 Mar	Fri, 17 Mar	Tue, 7 Mar	Tue, 21 Mar
Number of items for auction	86	86	78	78
Number of items sold	86	47	74	78
Revenue from twice-sold cards	$189.90	$234.75	$783.80	$758.25
Total number of bids	565	71	238	1583
Total number of bidders	19	7	42	62
from email invitations	12	5	35	46
from newsgroup announcements	7	2	7	18
Number of email invitations sent	52	50	234	232
Total auction revenue	$292.40	$234.75	$783.80	$774.75
Revenue plus salvage	$292.40	$264.24	$787.76	$774.75
Number of winners	15	6	27	25
Winner/bidder ratio	78.9%	85.7%	64.3%	40.3%
Cards per winner:				
Max	25	26	18	12
as share of total	29.1%	55.3%	24.3%	15.4%
Min	1	1	1	1
Mean	5.7	7.8	2.7	3.1
Median	3	3.5	2	2
Payment per winner:				
Max	$70.00	$129.40	$128.00	$316.50
as share of total	23.9%	55.1%	16.3%	40.9%
Min	$3.00	$0.70	$2.55	$1.05
Mean	$19.49	$39.13	$29.03	$30.99
Median	$10.50	$23.68	$13.00	$13.15

spender in any of the auctions won cards totalling $316.50 of the total revenue of $758.25 in Auction #9, generating 41 percent of the revenue despite winning no more than 15 percent of the cards - evidently, she was particularly interested in high-value cards. Thus, it is not the case that some people are the highest bidders on all cards in an auction, suggesting that a given bidder's valuations for different cards are at least somewhat independent. Later in this chapter, I will report results of regressions in which each individual card will be taken as an independent observation; the evidence here indicates that this is not a terribly inappropriate assumption.

5.1.1 Entry Costs Are Relevant

Statistics on the number of card bids per participating bidder are shown in Table 3. As expected, individual bidders tend to submit fewer bids in the presence of minimums than they do in the absence of minimums. In the auctions with minimums, no single bidder submitted bids on even half of the cards; the maximum number of bids by a single bidder was 30. By contrast, there were bidders in both of the no-minimum auctions who submitted individual bids on every single card.

Table 3: Bids received in the within-card auctions.

	Auction 4	Auction 8	Auction 5	Auction 9
Minimum bids?	No	Yes	Yes	No
Card set	Antiquities	Antiquities	Arabian Nights	Arabian Nights
Number of bidders	19	7	42	62
Number of items for auction	86	86	78	78
Number of bids per bidder:				
Mean	29.7	10.1	5.7	25.5
Median	13.0	4.0	4.0	14.0
Max	86.0	29.0	30.0	78.0
Min	1.0	1.0	1.0	1.0

Interestingly, though, relatively few bidders followed this strategy of bidding on every single card in the absolute (no-minimum) auctions. Only one out of 19 bidders bid on every single item in Auction #4, and only six of 62 bidders bid on every single item in Auction #9. These statistics indicate that the cost of submitting a bid (the participation cost) is high enough to affect bidder behavior, and thus this experimental environment is appropriate for exploring endogenous-entry bidding models such as MQV. If there were no cost to submitting a bid, then one would expect to see all of the participating bidders submitting bids on every card (as low as a nickel, say), since every card does have some positive resale value even to people who get no consumption utility from it. My conclusion is that people deem the probability of a bidder getting a bargain (and thus a resale profit) on a card they're not terribly interested in is low enough that the expected profit from bidding does not always outweigh the cost of having to decide on a bid amount and to type the approximately ten characters required to submit another card bid. Indeed, the median number of card bids submitted by a single bidder was only 13 (of a possible 87) in Auction #4, and 14 (of a possible 78) in Auction #9, even though these auctions had no minimum bids.

5.1.2 Number of Participating Bidders

The imposition of minimum bids clearly had a measurable impact on bidder participation. For the Antiquities auctions, the imposition of nonzero minimum bids reduced the number of bidders from 19 to 7, and the number of individual card bids from 565 to 71. Similarly, the imposition of minimum bids in the Arabian Nights auctions lowered the number of participating bidders from 62 to 42, and the number of submitted bids from 1583 to 238. This is a confirmation of a prediction of classical auction theory: the number of participating bidders n(r) is a decreasing function of the reserve price.

5.1.3 Probability of Sale

The "thin" bidder pool for the Antiquities card auctions coincides with a high number of unsold cards in the minimum-bid auction, as shown in Table 2 on page 41. Only 47 of 86 cards were sold in the Antiquities auction with minimum bids, by contrast with 74 of 78 cards in the thicker Arabians market. As can be seen in the table, all cards were sold in the auctions without minimum bids. Thus, reserve prices decrease the probability of sale of each card, just as would be predicted by auction theory. Further, the effect on the probability of sale seems to be greatest when the number of potential bidders is small.

5.1.4 Revenue on Sold Cards

When it is not known whether reserve prices were set optimally, classical auction theory has an ambiguous prediction about the effect of minimum bids on overall expected revenue. However, the theory does have a very clear prediction for the revenue earned on sold goods: if the good is actually sold, reserve prices must have a non-negative effect on revenue. Thus, Table 2 also contains a column entitled "Revenue from twice-sold cards" which gives the total revenue earned on only those cards which were sold in both auctions. In both experiments, the cards sold under reserve prices earned more than the same cards did in the absence of reserve prices. In the Antiquities experiment the difference was $44.85, while in the Arabian Nights experiment it was $25.55. Thus, the aggregate data yield a confirmation of another of the basic predictions of classical auction theory: conditional on a good actually being sold in the auction, reserve prices have a non-negative effect on revenue.

5.1.5 Overall Auction Revenue

The middle of Table 2 contains two rows of data on overall revenue in each auction, counting both sold and unsold cards. Unfortunately, the evidence here on the shape of the expected revenue curve R(r) is more ambiguous than the evidence in the previous subsections.

In the first of the two experiments, the minimum bids had a negative effect on revenue. Without minimum bids, auction revenue was $292.40, while revenue was only $234.75 with minimum bids. This last figure does not take into account any salvage value of the 39 unsold cards, however.

This calls for an estimate of my salvage value for the goods I auctioned. I computed my salvage value by asking my local card dealer what he would pay me for my unsold cards. Looking at my batch of unsold cards, he told me that on average, he would pay me approximately 20 percent of their Cloister price. Thus, for each card, I will use a salvage value of 20 percent of its Cloister price.[20]

After adding in the appropriate salvage values, this figure becomes $264.24, which is still 9.6 percent lower than the amount of revenue received in the absence of minimums.

While minimum bids decrease revenue in the first experiment, they increase revenue in the second experiment. When minimum bids were imposed, revenue rose from $774.75 to $783.80. After including the salvage value of the unsold cards, the comparison changes to $774.75 versus $787.76. Thus, the auction with minimum bids raised $13.01 more revenue, a difference of 1.6 percent.

[20] I might have been able to shop around for a better price with a different card dealer, but this represents my best estimate of a salvage value, which by definition should be net of all administrative costs, including search costs.

The evidence on revenue from the aggregate data, then, is mixed. The inconsistency between the two auctions may be due to effects of the ordering of the auctions, as discussed earlier. Note that in each experiment, the revenue was higher in the first auction in the pair, suggesting that, for example, demand for the second copy of any given card may be lower than the demand for the first copy. In the next section, a disaggregated analysis, examining the data on a card-by-card and a bidder-by-bidder basis, will attempt to disentangle the confounding effects of time order from the effects of the minimum bid levels.

5.1.6 Disaggregated Evidence on Overall Auction Revenue

In this section, I analyze an econometric model which considers the extent to which the amount of revenue earned depends on the existence of minimum bids. In this specification, each card auction is a separate observation (N=328), and the dependent variable is the amount of revenue earned (REV), which is set to equal the salvage value of the card (20% of Cloister value) in those cases where the card went unsold. The explanatory variables are VAL, the Cloister value of the card, MIN, an indicator variable which indicates whether minimum bids were present in the auction, SECOND, an indicator for whether the auction was the second (rather than first) auction I ran for the that card, and AN, an indicator for whether the card was in the Arabian Nights experiment as opposed to the Antiquities one. The primary variable of interest is, of course, MIN.

I chose a log-log specification for the dependence of REV on VAL, so that the two dummy variables would have multiplicative (fixed-proportion) effects on revenue. The results of the ordinary-least-squares regression were as follows (standard errors are in parentheses). The number of observations was 328, and the R^2 was 0.853:

$$\ln \mathrm{REV} = \underset{(0.0629)}{-0.3073} \underset{(0.0274)}{+1.0251} \ln \mathrm{VAL} \underset{(0.0594)}{- 0.2177} \mathrm{SECOND} \underset{(0.0633)}{- 0.4527} \mathrm{AN} + \underset{(0.0594)}{0.1819} \mathrm{MIN}$$

The elasticity of auction revenue with respect to Cloister value is 1.02, which is within one standard deviation of unity. This means that auction revenue is directly proportional to the Cloister value of a card, which indicates that the Cloister values are indeed good measures of cards' market values. The coefficient on SECOND is negative and statistically significant, indicating that auctioning the same card for a second time with the same bidder pool has a tendency to reduce revenue (by about 20 percent). A potential explanation for this is that the top bidder in the first auction may have her demand satiated by winning the card, and therefore she would not bid for the same card in the second auction in the pair. The magnitude of the effect, however, was surprisingly large, in my opinion. Next, the coefficient on AN is statistically significantly greater than zero, indicating that the Arabian Nights auctions fetched higher revenue premiums over the Cloister price than did the Antiquities auctions. There are two reasons for this effect: first, the Arabian Nights cards were rarer and more expensive as a set than were the Antiquities cards, thus attracting greater interest, and second, the Arabian Nights auctions were advertised to a much higher number of potential bidders. As can be seen from the theoretical revenue curves in Figure 2 on page 24 and Figure 4 on page 27, the number of potential bidders is expected in classical auction theory to have a significant positive impact on auction revenue.

Finally, the coefficient on MIN is statistically significantly negative at -0.1819, which indicates that the imposition of reserve prices decreases the amount of revenue earned on a card by 17 percent, on average. This gives some support to my proposed theory of the advertising value of auctions without reserve prices: *reserve prices scare away serious bidders, and thus auction revenues decline.*

A caveat is in order here. One might argue that these data are merely consistent with the results of classical auction theory, and that my advertising theory is not required to explain such results. Since the parameters about the bidders (distribution of valuations for each good, number of bidders who would actually read my message, etc.) required to compute the classically optimal reserve price for each card were unknown to me, it is possible that I set the minimum bids suboptimally.[21] In particular, suppose that the reserve prices were set a bit too high. Then imposing the reserve prices could actually decrease revenue, by causing too many of the cards to go unsold. (When reserve prices are set too high, they screen out even the highest-valuation bidder in an auction.) This would account for the lower revenue raised under reserve prices.

Under this alternative explanation, many of the bids I classified as "serious" may nevertheless have been screened out by the too-high reserve prices. Indeed, the data does give some additional support for this explanation. Figures 9 through 12 show histograms of the levels of the bids received in the four different auctions. This detailed data reveals that when reserve prices are removed (Auctions #4 and #9), although there are more "serious bids" in the range between 20% and 90% of the Cloister value, there are actually fewer "very serious

[21] Indeed, after spending months observing this market environment and after running auctions myself, it is hard for me to imagine how an auctioneer in a real-word environment could ever have enough information to choose precisely the optimal reserve price.

Figure 9: Distribution of Auction #4 bids.

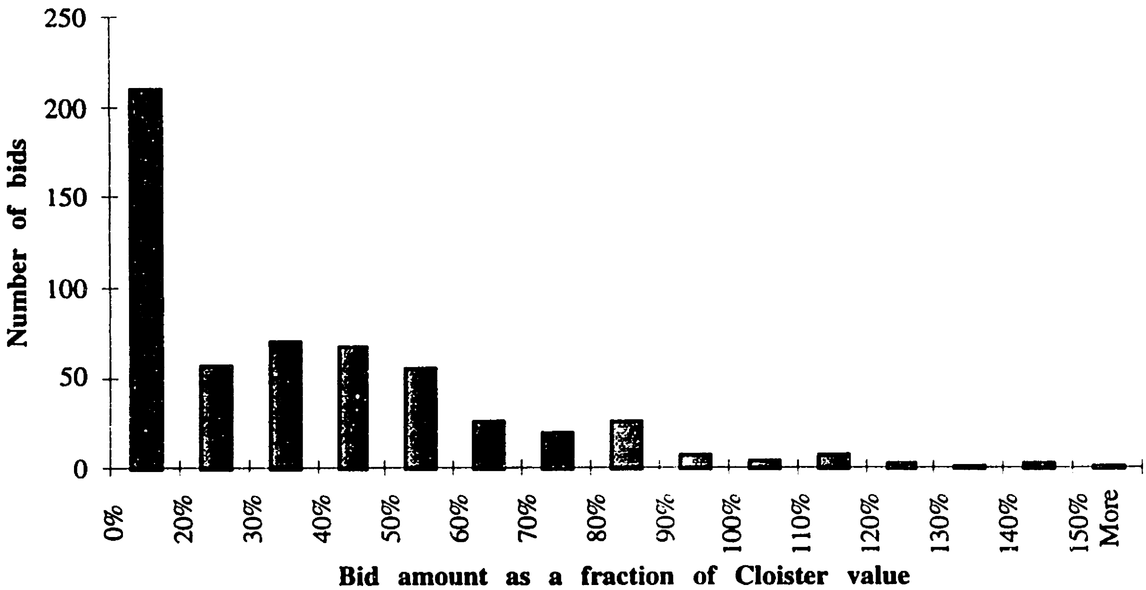

Figure 10: Distribution of Auction #8 bids

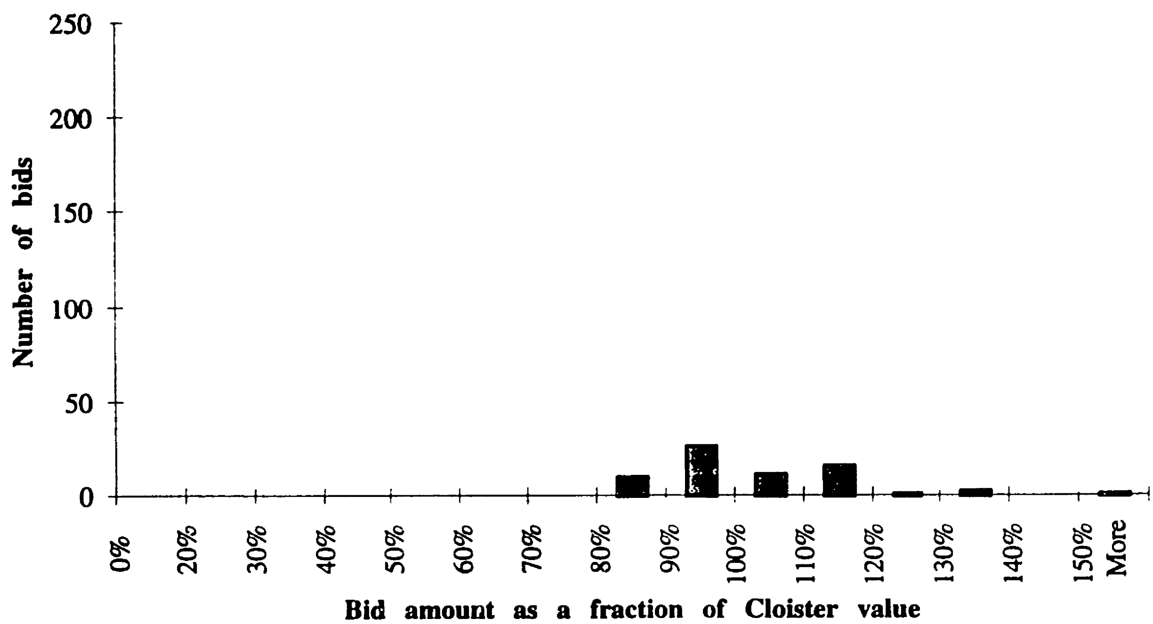

bids" in the range of 90% or higher. As classical theory would have it, those bidders whose valuations lay below 90% were forced not to bid at all, and those bidders whose valuations lay above 90% but who would bid below 90% in an absolute auction were forced to raise their bids to at least 90%. The revenue raised on sold cards was higher in the presence of reserve prices (as supported by the "revenue from twice-sold cards" figure in Table 2 on

Figure 11: Distribution of Auction #9 bids.

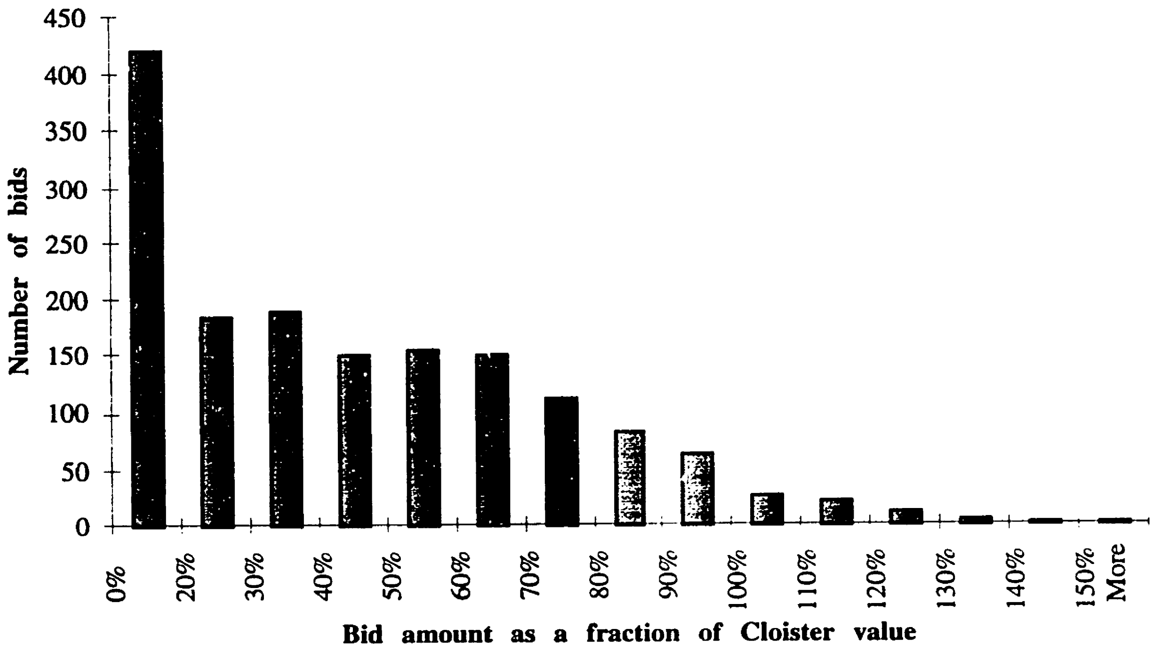

Figure 12: Distribution of Auction #5 bids.

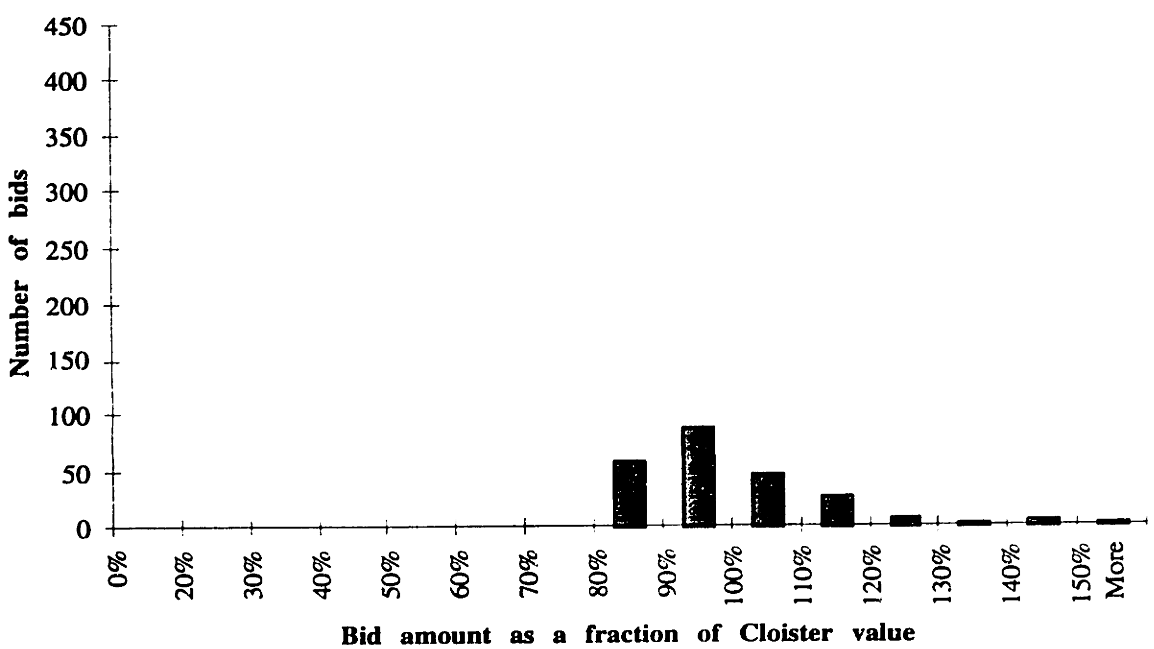

page 41), and the reason that reserve prices lowered auction revenue on average could be that they were set too high, causing too many cards to go unsold. Unfortunately, this experimental data cannot distinguish between an advertising effect which adds additional serious bidders and the classical effect of suboptimally high reserve prices.

5.1.7 Evidence on Bid Levels

Recall that in classical theory, $\frac{\partial b}{\partial r} > 0$ for a bidder with v>r. In this experiment, the implication is that if a bidder's valuation of a card is high enough that he chooses to participate in both auctions for the same card, he will (for strategic reasons) bid strictly higher in the presence of reserve prices than he will in the absence of the reserve prices. Returning to Figures 9 through 12, one can see that this prediction appears to be substantiated. Comparing Auction #5 to Auction #9, or Auction #8 to Auction #4, one can see that imposing reserve prices does not raise bid levels merely to the level of the reserve - it raises them still higher! The existence of reserve prices at the 90% level consistently increases the number of bids placed at levels 100%, 110%, 120%, and even 130% of Cloister value.

A more quantitative look at this data appears in Table 4. A total of 2,457 individual bids were placed on the cards in the four auctions. Of these, 2,148 occurred in the no-minimum auctions, while only 309 occurred in the auctions with reserve prices. Although the total number of bids was lower in the presence of reserve prices, it turns out that the number of bids greater than or equal to 90% of Cloister value was *greater* in the presence of reserve prices. Thus, reserve prices appear to have caused some bidders to increase their bids over the levels they would have chosen in an absolute auction.

Table 4: Comparison of submitted bid levels to the minimum bid levels.

Number of bids in:	Below the minimum	Equal to the minimum	Strictly above the minimum	Total
Auction 4	531	5	29	565
Auction 8	0	10	61	71
No-minimum auctions	1979	7	162	2148
Auction 5	0	59	179	238
Auction 9	1448	2	133	1583
With-minimum auctions	0	69	240	309
Total	1979	76	402	2457

The number of bids equal to the reserve price level increased from 7 to 69 when reserve prices were added, while the number strictly greater than the reserve price level increased from 162 to 240. This increase in the number of bids strictly greater than the reserve price is even stronger evidence in favor of the classical theory's prediction that $\frac{\partial b}{\partial r} > 0$ for bidders whose valuation is greater than the reserve price. To see this, consider a bidder who values a card at more than 90% of its Cloister price, but who in an absolute auction decides that his optimal bid is less than 90% of Cloister value. In the change from an absolute auction to an auction with minimum bids, a naive bidder might choose to raise his bid exactly to the level of the reserve price. A rational bidder, however, will take into account the strategic consideration that there is some probability that other bidders might also raise their bids in the presence of reserve prices, and will thus raise his bid to a level strictly greater than the reserve price. This experimental observation, that minimum bids cause a large increase in the number of bids strictly greater than the minimum, therefore demonstrates that the strategic considerations predicted by classical theory are indeed relevant to real-world bidding behavior.

Although the strategic considerations of classical theory have a measurable impact, there also seem to be some bidders who behave more naively, by bidding exactly at the reserve price. The classical bid equation (1) and its boundary condition (2) together imply that only a bidder with valuation v=r should bid exactly at the reserve price, which is an event of probability zero if the distribution of valuations is continuous. Therefore, the increase in the number of bids exactly equal to the reserve level, from 7 in the absolute auctions to 69 in the auctions with reserve prices, is somewhat contradictory of the classical theory. There are two potential interpretations of this finding. One is that there are some bidders in the sample who behave naively regarding reserve prices, while others behave as predicted. The other explanation is that the increase in the number of bids exactly equal to the reserve level is due to rounding error, caused by the fact that my auctions allowed only for discrete bid amounts in even multiples of a nickel. Therefore, a bidder who wished to bid at one penny over the minimum would have been forced instead to bid exactly at the minimum.

To examine in more detail the question of the effect of reserve prices on bid levels, I have assembled data on those bidders who bid on the same card twice, once in the presence of a reserve price, and once without. It proved to be a fairly small sample. There were 4 bidders submitting a total of 22 such bids in the Antiquities experiment, and 17 bidders submitting 88 such bids on cards in the Arabian Nights experiment. The data is displayed in

Figure 13. Of the 110 observations where the same bidder bid on the same card in both auctions, 70 of them had higher bids under reserve prices,[22] while 35 were lower under reserve prices and 5 were exactly equal between the two auctions. (The bids that were lower in the presence of reserve prices are a bit difficult to explain. Some of it appears random, but several bidders had a systematic tendency to behave this way: 3 of the 21 bidders accounted for 21 of these 35 aberrant observations.) The right tail of the distribution in Figure 13 contains much more mass than does the left tail, indicating once again that bids are often considerably higher in the presence of reserve prices.

Figure 13: Distribution of changes in bids on the same cards by the same bidders.

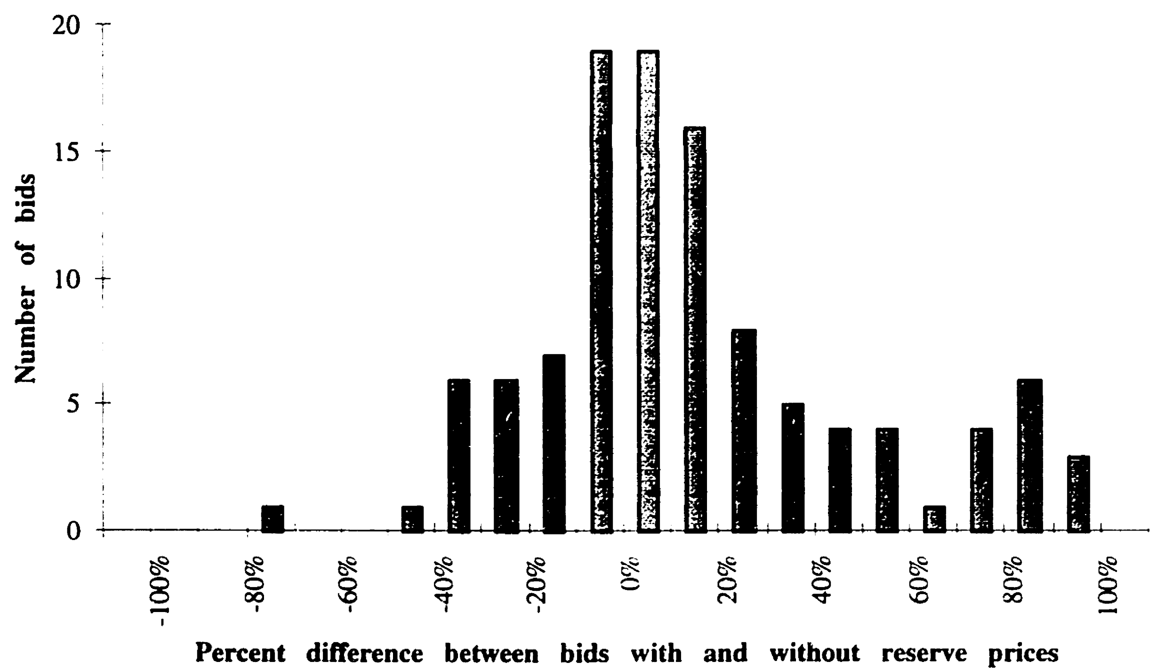

5.1.8 Is the Entry Equilibrium in Pure or Mixed Strategies?

In this section, I wish to consider one basic issue in endogenous-entry auction theory, the issue of whether the entry equilibrium is in pure strategies or in mixed strategies. The first endogenous-entry theories (see, for example, McAfee & McMillan (1987b)) modeled the bidder entry as a pure-strategy, asymmetric Nash equilibrium, in which exactly n bid-

[22] Of the 70 increases described above, it turns out that 28 increased up to the level of the reserve price, while 42 increased to a level strictly greater than the reserve price, which further corroborates classical bidding theory.

ders enter the auction (out of a total of N>n potential bidders), and n is determined endogenously from the other parameters of the model (the auction format, the degree of affiliation of bidder values, the cost of entry, and so on). More recently, mixed-strategy, symmetric entry equilibrium models have been proposed (see Levin and Smith (1994), MQV (1995)). In the mixed-strategy models, bidders enter with probability ρ, and stay out with probability $1-\rho$, where ρ is determined endogenously. Levin and Smith (1994) point out that the difference between pure-strategy (deterministic) models and mixed-strategy (stochastic) ones has implications for social welfare: if entry is stochastic, then expected social surplus is decreasing in the number N of potential bidders, because the variance of the number n of actual entrants is increasing in N, and such variance is costly. In common-value auctions, then, it turns out that auctioneers can increase both social welfare and their own profits by using reserve prices to discourage entry.

In Levin and Smith (1995), the authors perform an experiment in which they attempt to determine whether entry by bidders is stochastic or deterministic, and they find evidence in favor of their stochastic hypothesis. Their experiment, however, is very artificial. In fact, their experiment doesn't actually involve any auctions, but rather assigns simulated auction payoffs by a lottery procedure.[23] Thus, the experiment has little to do with actual auctions, but is instead a more abstract study of people's behavior in general entry games.[24]

---

[23] Subjects made the decision whether or not to incur the cost c to enter. After the entry outcome was observed, each of the n entrants had a 1/n chance of winning the payoff for that round of the experiment.

[24] In addition, I believe that this design biased the experiment towards finding a mixed-strategy entry equilibrium. Since the simple yes/no entry decision was the only opportunity a subject had to influence his payoff in the experiment, subjects might have decided to mix up their entry decisions at random merely to relieve the boredom of the experiment. If the auctions had actually been held, then perhaps the subjects would have learned something about the other subjects' bidding behavior. Then skillful bidders might have learned to continue entering consistently in subsequent rounds, while unskillful bidders might have learned to stay out; such behavior has been observed in experiments by Cox, Dinkin, and Smith (1996).

Levin and Smith argue that two effects actually bias their experiment *away* from finding a mixed-strategy entry equilibrium. One is that heterogeneity in risk preferences would cause more risk-averse bidders deterministically to stay out of the auction, while less risk-averse bidders would deterministically stay in. I would argue to the extent that risk-preference heterogeneity matters, it is also likely to matter in real-world auctions, and therefore this does not bias this laboratory experiment as compared with other auction environments. Their other argument is that repeated play could allow bidders to signal each other and coordinate to avoid the inefficiencies associated with stochastic entry. However, such efficiency gains are rather small, I believe they are not likely to be sufficient to induce coordination among bidders who may also have tastes to avoid boredom.

Therefore, the question of Levin and Smith warrants a second investigation in a more realistic environment. What is my evidence on deterministic versus stochastic entry equilibria? Very few bidders bid on a card both times it was offered, despite the fact that the same people were invited each time. In Table 2 on page 41, I reported that 19 and 7 bidders, respectively, bid in the two Antiquities auctions, but only 4 people overlapped between the two auctions. In the Arabian Nights auctions, there were 42 and 62 bidders, but only 17 of the bidders overlapped between the two. Thus, in each pair of auctions, there were a proportionally large number people who entered the first auction but not the second, and other people who entered the second auction but not the first. This argues in favor of a stochastic equilibrium, as the most natural kind of deterministic equilibrium is one in which the same bidders enter each time.

Two objections might be raised to this evidence. First, it might be the case that people enter one auction but not the other because the latter auction has reserve prices which are higher than they are willing to pay. However, this screening-out explanation cannot account for the bidders who bid in the presence of reserve prices but fail to bid in the absence of reserve prices; there were 3 such bidders in the Antiquities auctions, and 25 such bidders in the Arabian Nights auctions. The second potential objection is that bidders may have bid in the chronologically first auction, but not the second, in a pair because they had already bought the cards by the time the second auction occurred. This complaint potentially affects the 25 Arabian Nights bidders just cited, who bid in Auction #5 but not in Auction #9. Indeed, three of these bidders each placed a bid on a single card in Auction #5 and won it, so there would be no reason to expect them to bid in the second auction. However, none of the remaining 22 bidders won all the cards they bid on in Auction #5: ten did not win any cards at all, while the remaining twelve won an average of 50 percent of the cards they bid on. It is still possible that these bidders managed to purchase the rest of the cards they were interested in from someone else during the week that passed between my two auctions, but I can at least say that they did not buy them from me. Thus, the evidence is fairly strong that bidders in these auctions followed mixed strategies in their entry decisions: faced with the same auction opportunity, the same person might sometimes enter and sometimes fail to enter.

5.2 Between-Card Experiments

While the first set of auctions was designed to examine the effects of the *existence* of a reserve price, this second set of experiments was designed to investigate the effects of changes of the *level* of the reserve price. Summary statistics for these auctions are given in Table 5.

Table 5: Summary statistics for the between-card experiments.

	Auction #15	Auction #16	Auction #17	Auction #19
Card set	Artifacts	Black	White	Blue
Start date	Tue, 3 Oct	Fri, 6 Oct	Fri, 20 Oct	Mon, 23 Oct
End date	Tue, 10 Oct	Fri, 13 Oct	Fri, 27 Oct	Mon, 30 Oct
Number of items for auction	99	99	99	99
Number of items sold	98	92	77	78
Mean reserve level	60%	60%	85%	81%
Total number of bids	798	652	366	401
Total number of bidders	57	55	46	38
Number of email invitations sent	532	523	512	489
Total Cloister value	345.83	271.55	285.87	224.89
Total auction revenue	338.45	282.65	260.95	219.25
Revenue plus salvage	343.94	283.65	269.48	224.52

Each auction was a simultaneous auction of 99 different cards at different reserve price levels, each expressed as some fraction of that card's Cloister price. Some reserve prices were as low as 10% of the Cloister value of the card, while others were as high as 150% of the Cloister value. The average reserve price level varied slightly from auction to auction, from 60% to 85%.

As can be seen in the table, each auction had dozens of bidders and hundreds of bids on individual cards. The number of people receiving email invitations to participate declined with each successive auction, but only due to recipients asking to be removed from my mailing list, so the changes in the mailing list should not have affected the number of interested participants. It should be noted that the data gathered in this set of auctions is not directly comparable to the data from the within-card experiments, because the size and composition of the pool of participating bidders changed considerably during the intervening six months. Very few bidders overlapped between the two experiments; most of the bidders in the between-card experiment were new recruits.

The table also displays some aggregate statistics on revenue, including the total Cloister value of all the cards in each auction, the total revenue earned on cards which were sold, and a grand-total revenue figure which also includes the salvage value of the unsold cards. The auction revenue in each case was reasonably close to the total Cloister value of the cards; in Auction #16 I earned revenue greater than the total Cloister value, while in the three others I earned slightly less.

The following subsections describe the implications of this experiment for the theoretical questions outlined in section 3 of the chapter.

5.2.1 Bidder Participation

The first question, just as in section 5.1, is whether the number of bids submitted is a decreasing function n(r) of the reserve level. Because this second set of auctions provided data at a variety of different reserve levels (normalized the Cloister value for each card), I can now estimate the functional form of the function n(r).

The estimation proceeds as follows. When assigning reserve prices to each of the cards up for auction, I chose some to be at 10% of Cloister value, some at 20%, and so on up to 150% of Cloister value. We may think of the cards as having been assigned to one of fifteen discrete "bins," each of which represents a different reserve level.[25] Because there are at least 16 observations in each bin, it is thus possible to estimate n(r) as a nonparametric function of r, by assigning indicator variables to each bin and using them as regressors in a least-squares regression. In the absence of other regressors, this is equivalent to computing the mean number of bidders n(r) separately for each of the fifteen bins.

The reason I put the estimation problem into a regression context is that I also wanted to account for other potential determinants of the number of bidders. In particular, the Cloister value of the card might have an independent effect; for example, the $10 cards might attract more bidder interest than the fifty-cent cards. I also wanted to include auction dummy variables because of the simultaneous-auction format. For example, an auction

[25] When rounding was necessary in order to satisfy my requirement that all minimum bids be in even multiples of 5 cents, I always chose to round down to the nearest nickel. Thus, in practice the 10% bin contains cards whose reserve prices were less than or equal to 10% of Cloister value, the 20% bin contains cards whose reserve prices were between 10% and 20% of Cloister value, and so on.

with ten $10 cards might attract more bidders than an auction with only six $10 cards, and therefore the number of bids on the same fifty-cent card might attract more bidders if it were included in the former auction than in the latter.

The results of the regression analysis are reported in Table 6. The dependent variable is NUMBIDS, the number of bids received on a given card. The explanatory variables are the bin indicators RESBIN1 through RESBIN15, which correspond to reserve prices at 10% of Cloister value through 150% of Cloister value, as well as the card's Cloister value (CLOISTER), and auction dummy variables AUCT15, AUCT16, and AUCT17. (The excluded dummy variable was a dummy for Auction #19.) Two different specifications are reported, a base regression plus a regression that also included interactions between CLOISTER and the bin indicators.

The results confirm convincingly that the number of bids is a monotonically decreasing function of the reserve price level, at least at reserve prices between 10% and 70% of Cloister value. The coefficients decline from 9.8 at a reserve price of 10% down to 1.0 at a reserve price of 70%,[26] beyond which point all bin coefficients have point estimates less than 1 and are insignificantly different from each other. This confirms the prediction of classical auction theory, that the slope of n(r) is negative.

The Cloister value of the card also turns out to be positive and significant, indicating that more bidder interest is generated by more expensive cards. This makes sense in this market environment, where the transaction cost of purchasing an expensive card is a negligibly small fraction of the card value, but where the 32-cent cost of mailing a check for a 50-cent card is more likely to discourage bidder interest. In addition, two of the three auction dummies were positive and statistically significant, indicating that Auctions #15 and #16 were somewhat more popular than Auctions #17 and #19. This can perhaps be explained by the fact that the average reserve prices in the first two auctions were lower than in the latter two auctions, as shown in Table 5. Under this interpretation, bidders were more likely to bid in an auction if there were more potential bargains available in the set of cards up for auction. Thus, it seems that the individual card auctions were not entirely independent; the simultaneous nature of the auction format does have some effects on bidder behavior that would not be present in sequential auctions of each individual good.

[26] When CLOISTER is evaluated at its mean value, this indicates that in the base auction (#19), there are an average of 10.7 bidders at reserve levels of 10%, and 2.1 bidders at reserve levels of 70%.

Table 6: Least-squares results for NUMBIDS, the number of submitted bids per card.

Variable	Specification 1		Specification 2	
	Coefficient	Std. Error	Coefficient	Std. Error
RESBIN1	9.8043	(0.8411)	7.5440	(0.8510)
RESBIN2	8.0201	(0.8318)	4.9905	(0.9002)
RESBIN3	5.8973	(0.8390)	3.9177	(0.9398)
RESBIN4	6.0878	(0.9411)	3.0210	(1.0898)
RESBIN5	3.3543	(0.8433)	2.5734	(0.8328)
RESBIN6	2.8973	(1.2013)	1.9526	(1.4787)
RESBIN7	1.0484	(1.2039)	2.2154	(1.1647)
RESBIN8	-0.2924	(1.1799)	1.7709	(1.0728)
RESBIN9	-0.0082	(1.0153)	0.9068	(0.9893)
RESBIN10	-0.1592	(0.8430)	1.0387	(0.8342)
RESBIN11	0.2704	(0.8389)	0.6541	(0.9391)
RESBIN12	0.3449	(1.0775)	1.5417	(1.1762)
RESBIN13	0.4758	(1.1379)	1.8499	(1.4288)
RESBIN14	0.3371	(1.0762)	1.0199	(1.3071)
RESBIN15	0.5464	(1.0770)	1.3246	(1.2965)
AUCT15	3.3535	(0.6898)	3.2359	(0.6015)
AUCT16	2.1689	(0.6887)	1.5279	(0.6042)
AUCT17	-0.1215	(0.6285)	-0.3042	(0.5490)
CLOISTER	0.3867	(0.0501)		
CLOISTER*RESBIN1			1.3146	(0.1733)
CLOISTER*RESBIN2			1.7296	(0.2318)
CLOISTER*RESBIN3			1.3184	(0.2500)
CLOISTER*RESBIN4			1.5447	(0.2598)
CLOISTER*RESBIN5			0.7264	(0.1388)
CLOISTER*RESBIN6			0.8679	(0.3820)
CLOISTER*RESBIN7			0.1422	(0.1627)
CLOISTER*RESBIN8			0.0560	(0.0740)
CLOISTER*RESBIN9			0.2358	(0.1178)
CLOISTER*RESBIN10			0.0608	(0.1409)
CLOISTER*RESBIN11			0.3234	(0.2536)
CLOISTER*RESBIN12			-0.0690	(0.3036)
CLOISTER*RESBIN13			-0.0984	(0.3920)
CLOISTER*RESBIN14			0.1127	(0.4260)
CLOISTER*RESBIN15			0.0926	(0.3812)
R^2	0.4935		0.6322	
Number of observations	396		396	

The second specification, which adds interaction terms between the reserve price bins and the Cloister value, shares all the noted features of the first specification: the number of bidders is a decreasing function of the reserve price, and the first two auction dummies are positive and statistically significant. This regression adds one other result of interest: it demonstrates that the intrinsic Cloister value of a card significantly affects the number of bidders only when the reserve price level is low. That is, the interaction coefficients are positive and statistically significant only for the first six bins, representing reserve prices from 10 to 60 percent of Cloister value, but not significantly different from zero in the rest of the bins, from 70 to 150 percent of Cloister value.[27] This empirical observation has the following interpretation: not only do proportionally low reserve prices attract more bidder interest, but this effect is greatest when the card is intrinsically more valuable. A valuable card at a low reserve price attracts a lot more attention than a cheap card at a low reserve price, but the same is not necessarily true at high reserve prices.

Table 7 shows a final piece of analysis on the number of bids submitted on each card, disaggregating the data according to the size of each bid. Each row of the table displays observations from a different reserve price bin (0.1 through 1.5); the number of cards per bin is between 16 and 39, depending on the bin. Each column displays the number of bids (averaged over all card observations) at amounts greater than or equal to a certain level. For example, the first cell of the table shows that for cards with a reserve price level of 0.1 (that is, 10% of Cloister value), there were an average of 11.4 bids received per card auctioned. The cell immediately to the right excludes bids which were less than 20% of Cloister value, and the average number of bids falls to 9.7; this implies that approximately 2 of 11 bids on average were between 10% and 20% of Cloister value. The next cell on the right shows that the number of bids per card falls to 6.8 if bids less than 30% of Cloister value are excluded, and so on. In general, cells on the diagonal of the table show the total number of bids received which were greater than or equal to the reserve price, while cells to the right of the diagonal show the number of bids received which were at even higher price levels.

What is most interesting about Table 7 is a comparison of the numbers down each column. Within any particular column, the number of bids received appears to be an increasing function of the reserve price employed in the auction. For example, in the "bids >=0.4" column, we see that the number of bids received at levels greater than or equal to 40% of Cloister price increases as the reserve price increases. There are 4.4 such bids received per

27 The lone exception is the 90% bin, which has a coefficient of 0.11 that is just barely significantly positive at the 5% significance level.

Table 7: Number of bids received, as a function of the bid amount.

This reports the mean number of bids received across cards in each reserve-price bin, with std. dev. of the mean in italics.

Reserve Price	**Observed bids at amounts greater than or equal to a price level of:**														
	0.1	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**	**1.1**	**1.2**	**1.3**	**1.4**	**1.5**
0.1	11.436	9.667	6.769	4.410	3.231	1.949	1.256	0.923	0.590	0.436	0.333	0.282	0.256	0.256	0.231
	1.349	*1.206*	*0.893*	*0.632*	*0.529*	*0.382*	*0.246*	*0.206*	*0.136*	*0.115*	*0.106*	*0.090*	*0.088*	*0.088*	*0.078*
0.2		10.111	7.694	4.778	3.472	2.528	1.833	1.111	0.667	0.583	0.417	0.361	0.333	0.278	0.250
		1.187	*1.023*	*0.733*	*0.525*	*0.476*	*0.407*	*0.232*	*0.191*	*0.175*	*0.115*	*0.114*	*0.105*	*0.086*	*0.083*
0.3			8.314	6.200	3.343	2.200	1.514	0.943	0.686	0.543	0.286	0.257	0.114	0.086	0.086
			0.965	*0.901*	*0.545*	*0.395*	*0.308*	*0.201*	*0.191*	*0.161*	*0.120*	*0.111*	*0.055*	*0.048*	*0.048*
0.4				9.115	6.962	4.154	2.692	1.808	1.115	0.692	0.385	0.115	0.077	0.077	0.077
				1.380	*1.262*	*1.015*	*0.695*	*0.571*	*0.361*	*0.227*	*0.167*	*0.085*	*0.053*	*0.053*	*0.053*
0.5					6.118	4.882	2.706	1.529	1.059	0.794	0.500	0.441	0.324	0.265	0.206
					0.896	*0.794*	*0.471*	*0.344*	*0.257*	*0.218*	*0.195*	*0.190*	*0.162*	*0.136*	*0.125*
0.6						6.150	4.250	1.950	1.050	0.700	0.500	0.350	0.350	0.250	0.200
						1.047	*0.852*	*0.432*	*0.285*	*0.219*	*0.185*	*0.109*	*0.109*	*0.099*	*0.092*
0.7							5.438	3.313	1.750	0.625	0.375	0.250	0.188	0.188	0.125
							0.671	*0.546*	*0.470*	*0.301*	*0.180*	*0.112*	*0.101*	*0.101*	*0.085*
0.8								3.650	2.100	1.100	0.500	0.350	0.300	0.150	0.150
								0.519	*0.390*	*0.216*	*0.154*	*0.109*	*0.105*	*0.082*	*0.082*
0.9									3.440	2.400	1.000	0.400	0.320	0.160	0.120
									0.566	*0.436*	*0.271*	*0.141*	*0.111*	*0.075*	*0.066*
1.0										2.256	1.179	0.564	0.256	0.179	0.128
										0.293	*0.232*	*0.183*	*0.080*	*0.072*	*0.066*
1.1											2.606	1.485	0.485	0.424	0.212
											0.351	*0.258*	*0.180*	*0.151*	*0.084*
1.2												2.053	1.421	0.947	0.842
												0.883	*0.777*	*0.789*	*0.735*
1.3													1.800	0.900	0.600
													0.639	*0.464*	*0.373*
1.4														1.278	0.722
														0.311	*0.211*
1.5															1.375
															0.554

card auctioned at a reserve price of 0.1, 4.8 bids per card at a reserve price of 0.2, 6.2 bids per card at a reserve price of 0.3, and 9.1 bids per card at a reserve price of 0.4. This is consistent with the sort of strategic behavior predicted by classical auction theory: the reserve price can cause bidders to raise their bids even when the reserve is not binding. (See Figure 1 on page 23 and Figure 3 on page 26 for graphical illustrations of this property.)

For example, raising the reserve price from 0.2 to 0.3 increases not only the number of bids submitted at prices of 0.3, but also the number of bids at prices of 0.4 (in this case, from 4.8 bids to 6.2 bids). This phenomenon is generally true down each of the columns of the table, although as we might expect, the effect appears only in the last several rows of each column. That is, for example, increasing the reserve price from 0.2 to 0.3 has very little effect on the number of bids submitted at levels of 1.5 or more, just as changes in the reserve price have very little effect on the rightmost sections of the bid functions in Figures 1 and 3. Thus, Table 7 is yet another confirmation of the type of strategic bidding behavior predicted by classical auction theory.

5.2.2 Probability of Sale

The next question is whether the probability of selling a card is also a decreasing function of the reserve price, as would be predicted by the classical theory of first-price auctions. To examine this question, I employed two different specifications, both of which are reported in Table 8.

The first specification is a probit regression, reported in the first column of the table, with SOLD (zero if the card was unsold, one if sold) as the dependent variable. The binning technique could not be used for this estimation, because a number of the bins contained no unsold cards - thus, the corresponding indicator variables were perfectly correlated with SOLD and could not be included as regressors in probit estimation. Instead, the level of the reserve price (as a fraction of Cloister value) was used as a single regressor, RESLEVEL. The Cloister value and the Auction #15 dummy both had positive and significant coefficients, while the other two auction dummies were statistically insignificant. The reserve price has a significantly negative impact on the probability of sale: the regression model predicts that a card with a 100% reserve price is approximately 21 percent less likely to sell than is a card with a zero reserve price.

In the second column, I report the results of a linear probability model, estimated via ordinary least squares. Although a linear probability model generally has some undesirable econometric properties, it has the virtue in this case of allowing the probability of sale to

Table 8: Probability of sale. Dependent variable = SOLD.

Variable	Probit Coefficient	Probit Std. Error	OLS Coefficient	OLS Std. Error
RESBIN1			0.9564	(0.0585)
RESBIN2			0.9592	(0.0578)
RESBIN3			0.9032	(0.0583)
RESBIN4			0.9444	(0.0654)
RESBIN5			0.8430	(0.0586)
RESBIN6			0.8727	(0.0835)
RESBIN7			0.9247	(0.0837)
RESBIN8			0.8602	(0.0820)
RESBIN9			0.8647	(0.0706)
RESBIN10			0.7599	(0.0586)
RESBIN11			0.8478	(0.0583)
RESBIN12			0.5415	(0.0749)
RESBIN13			0.5468	(0.0791)
RESBIN14			0.4880	(0.0748)
RESBIN15			0.6533	(0.0748)
RESLEVEL	-1.4294	(0.2627)		
CLOISTER	0.2066	(0.0647)	0.0075	(0.0035)
AUCT15	1.2616	(0.4611)	0.0784	(0.0479)
AUCT16	0.2702	(0.2738)	0.0235	(0.0479)
AUCT17	-0.0951	(0.2262)	-0.0085	(0.0437)
CONSTANT	1.7750	(0.3123)		
R^2	0.2164		0.2249	
Number of observations	396		396	

vary nonparametrically with the reserve price. As in section 5.2.1, the regressors were RESBIN1 through RESBIN15, along with the Cloister value and auction dummy variables. The table shows in a more detailed format that the probability of sale is indeed decreasing in the reserve price, from 0.95 for the 10% reserve bin to 0.49 for the 140% one. Although the coefficient point estimates do not decline monotonically with the reserve price, the standard errors are large enough that one cannot reject for any pairwise comparison of bins the hypothesis that the larger reserve price bin has a coefficient less than or equal to the coefficient for the bin at the smaller reserve price.[28] The coefficients on the first four bins *are*

28 The exception is bin 15, which despite having extremely high reserve prices of 150% of Cloister value, had a large number of sold cards. It is likely that this is merely random noise.

pairwise significantly greater than the coefficients on the last four bins, indicating a decline in the probability of sale as the reserve price increases. Thus, both specifications confirm that *the probability of sale is a decreasing function of the reserve price.*

5.2.3 Revenue on Sold Cards

Table 9 displays the results of a nonparametric (binned) regression designed to test whether the revenue earned on a card, conditional on selling it, is an increasing function of the reserve price. The dependent variable is REV2CLO, the auction revenue normalized by the Cloister value of that card, and the sample is restricted to those 345 cards (out of a total of 396) which were actually sold in the auctions. The regressors are the same as those used in the previous two subsections.

Table 9: Revenue on sold cards. Dependent variable = REV2CLO.

Variable	Coefficient	Std. Error
RESBIN1	0.9433	(0.0964)
RESBIN2	0.8709	(0.0955)
RESBIN3	0.8890	(0.0977)
RESBIN4	0.8324	(0.1066)
RESBIN5	0.9552	(0.1012)
RESBIN6	1.1133	(0.1397)
RESBIN7	0.9357	(0.1372)
RESBIN8	1.2035	(0.1370)
RESBIN9	1.1365	(0.1205)
RESBIN10	1.2489	(0.1074)
RESBIN11	1.2421	(0.1025)
RESBIN12	1.3264	(0.1575)
RESBIN13	1.5364	(0.1648)
RESBIN14	1.6498	(0.1683)
RESBIN15	1.7700	(0.1463)
CLOISTER	-0.0060	(0.0057)
AUCT15	0.0598	(0.0815)
AUCT16	-0.0045	(0.0824)
AUCT17	-0.0125	(0.0799)
R^2	0.1974	
Number of observations	345	

The results confirm that the reserve price does indeed increase the revenue earned on sold cards, just as predicted by the classical theory. The point estimates do not increase monotonically, but the standard errors are such that for any pair of bins, one cannot reject the hypothesis that the bin with the larger reserve price has a larger revenue coefficient than the bin with the smaller reserve price. By contrast, one can reject the hypotheses that any of the four largest reserve price bins have coefficients less than or equal to the coefficients on any of the four smallest reserve price bins.

Besides the reserve prices, the only other statistically significant regressor was the dummy variable for Auction #15, which had a positive coefficient.

5.2.4 Overall Auction Revenue

Reserve prices do increase auction revenue on cards which actually sell, but what are their effects on overall expected revenue? To answer this question, I assign a salvage value of 20% of Cloister value to those cards which failed to sell at auction. Then I regress the normalized revenue REV2CLO per card on the same independent variables as in the previous four subsections. The results are shown in Table 10. Two specifications are reported: a base regression, and the base regression with added terms that capture interactions between the reserve price level and the Cloister value.

The major features of the results are that the expected revenue is approximately 80% of Cloister value for low (20% to 50%) reserve price levels, increases to approximately 100% of Cloister value for intermediate (80% to 110%) reserve price levels, and then declines back down to approximately 90% of Cloister value for very high (120% to 140%) reserve price levels. The additional control variables (Cloister value, auction dummies) all turn out to be statistically insignificant, as do the Cloister value-reserve price interaction terms. The expected revenue curve R(r), as given by the reserve price bin coefficients in the base regression, is plotted in Figure 14.

How does this curve compare with the expected revenue curves predicted by classical auction theory? (Recall the examples plotted in Figures 2 and 4 on page 24 and page 27.) The general prediction which should be testable with this data is that the revenue curve should be flat at r=0, and nondecreasing for all reserve prices less than the optimal one. The data suggest that this prediction may be violated, as the point estimates of expected revenue actually *decline* quite a bit (especially in the second specification) from reserve prices of 10% to reserve prices of 40%. This revenue decline between the first and fourth bins is statistically significant, at least in the second regression specification. My proposed explanation for the observed negative slope of the revenue curve at very low reserve price levels is that negligibly low reserve price levels of, say, 10% may have some advertising value, just as I posited for zero reserve prices back in section 3.2.2. Thus, the additional bidders ob-

Figure 14: The empirical expected revenue curve, R(r).

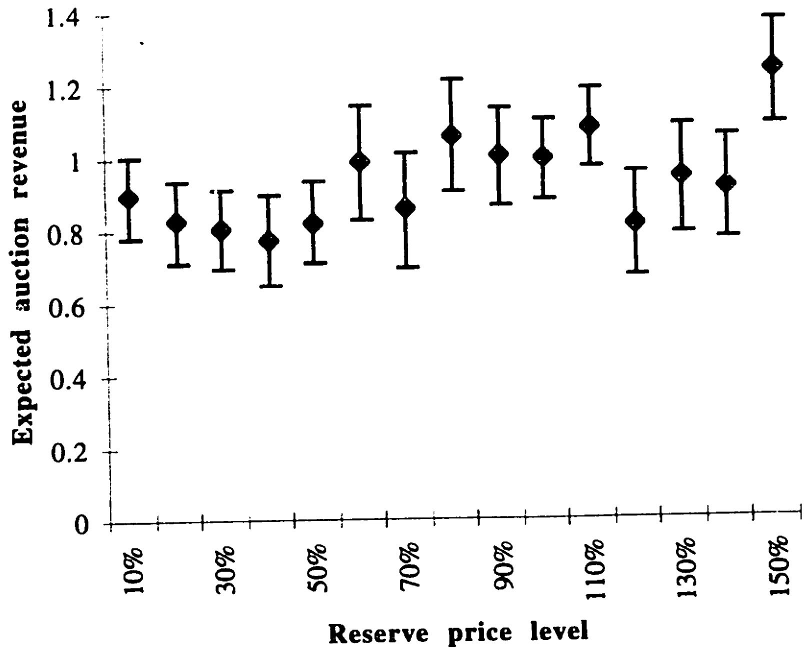

served at low reserve price levels in section 5.2.1 of this chapter may consist not only of bidders with low valuations (who would be screened out by higher reserve prices), but also of bidders with relatively high valuations (who, after having been attracted to the card auction by the very low reserve price levels at the stage of deciding whether to bother preparing a bid, decide after investing time in the auction to place optimal, and thus relatively high, bids).

A caution about this finding is in order. The observed negative slope of the revenue curve over the first four bins is not terribly large compared to the standard errors of the estimates, and only marginally statistically significant. Also, note that the coefficient estimate for the fifteenth bin is very high, significantly higher than the coefficient for the fourteenth bin. This would indicate that the revenue curve R(r) has two local maxima, and that the optimal reserve price is actually at 150% or more of the Cloister value. However, I believe that this is very difficult to accept, given that the Cloister value is an average transaction price for cards traded on the Internet. My interpretation, given the data in this section and in section 5.2.2, is that the high number of sold cards in bin 15 is a statistical fluke. It may also have been affected by uncertainty in the Cloister prices, which are computed as

Table 10: The expected revenue curve, R(r). Dependent variable = REV2CLO.

Variable	Specification 1		Specification 2	
	Coefficient	Std. Error	Coefficient	Std. Error
RESBIN1	0.8975	(0.1105)	0.9922	(0.1298)
RESBIN2	0.8281	(0.1093)	0.8141	(0.1373)
RESBIN3	0.8078	(0.1103)	0.7754	(0.1434)
RESBIN4	0.7748	(0.1237)	0.6452	(0.1663)
RESBIN5	0.8242	(0.1108)	0.8237	(0.1271)
RESBIN6	0.9891	(0.1579)	1.0989	(0.2256)
RESBIN7	0.8581	(0.1582)	0.8828	(0.1777)
RESBIN8	1.0606	(0.1551)	1.0893	(0.1637)
RESBIN9	1.0045	(0.1334)	1.0011	(0.1509)
RESBIN10	0.9976	(0.1108)	1.0026	(0.1273)
RESBIN11	1.0846	(0.1103)	0.9815	(0.1433)
RESBIN12	0.8179	(0.1416)	0.7598	(0.1794)
RESBIN13	0.9427	(0.1496)	0.8544	(0.2180)
RESBIN14	0.9182	(0.1415)	0.7374	(0.1994)
RESBIN15	1.2366	(0.1416)	1.1394	(0.1978)
AUCT15	0.1333	(0.0907)	0.1433	(0.0918)
AUCT16	0.0249	(0.0905)	0.0166	(0.0922)
AUCT17	-0.0198	(0.0826)	-0.0167	(0.0838)
CLOISTER	0.0021	(0.0066)		
CLOISTER*RESBIN1			-0.0335	(0.0264)
CLOISTER*RESBIN2			0.0074	(0.0354)
CLOISTER*RESBIN3			0.0153	(0.0381)
CLOISTER*RESBIN4			0.0471	(0.0396)
CLOISTER*RESBIN5			0.0019	(0.0212)
CLOISTER*RESBIN6			-0.0381	(0.0583)
CLOISTER*RESBIN7			-0.0058	(0.0248)
CLOISTER*RESBIN8			-0.0036	(0.0113)
CLOISTER*RESBIN9			0.0026	(0.0180)
CLOISTER*RESBIN10			0.0000	(0.0215)
CLOISTER*RESBIN11			0.0456	(0.0387)
CLOISTER*RESBIN12			0.0254	(0.0463)
CLOISTER*RESBIN13			0.0349	(0.0598)
CLOISTER*RESBIN14			0.0851	(0.0650)
CLOISTER*RESBIN15			0.0431	(0.0582)
R^2	0.0526		0.0731	
Number of observations	396		396	

trimmed means of price distributions with very high standard deviations. Therefore, given my reluctance to accept the revenue estimate for bin 15, I also should express caution about the intriguing finding that the slope of the revenue curve is decreasing in bins 1 through 4.[29]

5.2.5 A Test of McAfee-Quan-Vincent

Recall that the main prediction of the MQV paper is that raising the reserve price from some low value to the salvage value of the good will increase expected auction revenue, even in an endogenous-entry context. In this context, then, their prediction implies that raising reserve prices from 10% of Cloister value to 20% of Cloister value (which is my estimate of salvage value for unsold cards) should increase auction revenue. In other words, in the regression reported in section 5.2.4, the coefficient on bin 2 should be greater than the coefficient on bin 1. The point estimates give just the opposite result, although one cannot reject the hypothesis that difference between the two coefficients is positive. Therefore, the evidence is inconclusive, although it points mildly against the MQV theory. Note that if very low minimum bids have some advertising value that can attract additional serious bidders, as I proposed in section 3.2.2, this could account for reserve prices below the auctioneer's salvage value yielding higher revenue than reserve prices at the salvage value.

6 Conclusions

This study presents the results of controlled experimental auctions performed in a field environment. By auctioning real goods in a preexisting, natural auction market, I have obtained data in a manner that is intermediate between laboratory experiments and standard empirical IO studies using field data. Not every variable could be observed (for example, I could not assign "valuations" for each good to each bidder, as a laboratory experimentalist might), but I believe that I was able to hold constant all of the relevant variables in the environment except for the treatment variable, which in this case was the existence and level of reserve prices. By giving up the ability to observe and manipulate some of variables that laboratory experimenters can control, I gained a very realistic environment, akin to the field data traditionally used by empirical economists. The transactions which took place in this experiment were unquestionably real economic transactions.

29 In future research, I hope to resolve this issue both by collecting more experimental data, and by gathering more information about card characteristics (rarity, etc.) as additional variables that may help explain anomalies in the data.

The first result of this chapter is that despite its ignorance of endogenous-entry concerns, classical auction theory still accurately predicts a number of important features of the data. Holding all else constant, implementing reserve prices (1) reduces the number of bidders, (2) increases the frequency with which goods go unsold, and (3) increases the revenues received on the goods conditional on their having been sold. A more subtle and interesting result is that bidders appear to behave strategically in the presence of reserve prices, just as predicted by classical theory. A high-valuation bidder raises his bid in the presence of a minimum bid - not merely up to the level of the minimum, as a naive bidder might do, but even higher than the minimum, as would a bidder who correctly anticipates that other bidders will also raise their bids in the presence of posted minimums. In other words, I find confirmation of the classical prediction that $\frac{\partial b}{\partial r} > 0$ for bidders with valuations greater than the minimum bid level.

The second result is that entry costs are indeed an important feature of real-world auction markets, thus confirming the central hypothesis of endogenous-entry auction theory. The costs in the Magic-card market are not nearly as dramatic as those postulated in other markets (the auction market for offshore oil rights is a good example). Here, the cost of acquiring information about individual cards is quite small, but even the cost of typing in a bid amount appears to have observable effects.

Next, I find evidence that suggests that negligibly low reserve prices may have an advertising value that attracts additional serious bidders, which is potentially important in cases where bidder entry is endogenous. The additional bidders may raise expected auction revenue, outweighing the fact that reserve prices below salvage value may expose the auctioneer to potential losses.

Finally, this chapter sheds light on the question of which equilibrium model is the more realistic description of auction entry equilibrium: a pure-strategy deterministic equilibrium, or a mixed-strategy stochastic equilibrium. When the same cards were auctioned twice in rapid succession, very different sets of people decided to submit bids, despite the fact that the same superset of people were invited to participate both times. This can be interpreted as evidence in favor of the stochastic entry equilibrium model.

7 Future Research

This chapter suggests a number of interesting lines of research. First of all, it would be interesting to reproduce these experiments in a laboratory setting, to see whether the classical predictions about reserve prices can also be verified there. Such experiments could be undertaken both with fixed-N and endogenous-entry environments. Second, I would like to investigate whether reserve prices have different effects in thin versus thick markets. In this vein, it would be useful to perform an experiment in which both the reserve prices and the goods offered are held constant, but the number of invited bidders is manipulated as a treatment variable.

Another topic for future research is the exploration of the effects of reserve prices under different auction formats. The first-price, sealed-bid auction format was used here, because it was the easiest type of auction to administer, and the easiest for the bidders to understand. However, do the results found here also hold up under all four of the basic auction types (including second-price, English, and Dutch mechanisms)? I have already begun to experiment with conducting second-price, English, and Dutch auctions in this marketplace, as reported in Chapter 2.

In addition, it is worth investigating whether the simultaneous auction of K different goods has the same properties as K different sequential auctions, one for each good. Recent developments in information technology have made simultaneous-auction formats possible in a way that they never were before;[30] however, to date there has been very little research, either theoretical or experimental, on the properties of simultaneous versus sequential auctions.

Finally, I am optimistic that the "field experiment" methodology described in this thesis will be useful in other areas of research. Rather than relying solely on field data on behavior, which typically does not have the variation needed to test the economic theories of interest, economists can run their own businesses with the intent of collecting data on the effects of different decisions (such as the decision to employ reserve prices in auctions). With the new development of electronic markets for goods and services (such as those which have sprung up on the World Wide Web), detailed data recording in some businesses can be relatively straightforward. Therefore, it may now be possible for economists to collect data by running businesses for research purposes instead of (or as well as) for profit.

30 Indeed, the recent FCC auctions for communications spectrum were held in a simultaneous format. See McMillan (1994) for details.

References

Ashenfelter, Orley, "How Auctions Work for Wine and Art." *Journal of Economic Perspectives*, 1989, vol. 3, no. 3, pp. 23-36.

Black, Jason. *Cloister's Magic Card Price List*, <http://www.hhhh.org/cloister/pricelists/>, various weekly issues.

Bulow, Jeremy, and Paul Klemperer. "Auctions vs. Negotiations." *American Economic Review*, 1996, vol. 86, no. 1, pp. 180-194.

Cox, James C., Samuel H. Dinkin, and Vernon L. Smith. "Winner's Curse or Blessing: Conditions Under Which They Occur," Mimeo, University of Arizona, 1996.

Cox, James C., Bruce Roberson, and Vernon L. Smith, "Theory and Behavior of Single Object Auctions," in *Research in Experimental Economics*, Vernon L. Smith, ed., Greenwich, Conn.: JAI Press, 1982.

Hendricks, Kenneth, and Harry J. Paarsch, "A Survey of Recent Empirical Work Concerning Auctions." *Canadian Journal of Economics*, 1995, vol. 28, no. 2, pp. 315-338.

Hendricks, Kenneth, and Robert Porter, "An empirical study of an auction with asymmetric information," *American Economic Review*, 1988, vol. 78, pp. 865-883.

Kagel, John H. "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, J. Kagel and A. Roth, eds. Princeton: Princeton University Press, 1995, pp.501-585.

Kagel, John H., and Dan Levin, "The winner's curse and public information in common value auctions." *American Economic Review*, 1986, vol. 76, pp. 894-920.

Laffont, Jean-Jacques, Hervé Ossard, and Quang Vuong, "Econometrics of First-Price Auctions." *Econometrica*, 1995, vol. 63, no. 4, pp. 953-980.

Levin, Dan, and James L. Smith. "Equilibrium in Auctions with Entry." *American Economic Review*, 1994, vol. 84, no. 3, pp. 585-599.

Levin, Dan, and James L. Smith. "Auctions with Entry: An Experimental Investigation." Mimeo, University of Houston, 1995.

McAfee, R. Preston, and John McMillan. "Auctions and Bidding." *Journal of Economic Literature*, 1987a, vol. 25, no. 2, pp. 699-738.

McAfee, R. Preston, and John McMillan. "Auctions with Entry." *Economics Letters*, 1987b, vol. 23, no. 4, pp. 343-347.

McAfee, R. Preston, Daniel Quan, and Daniel Vincent. "How to Set Optimal Minimum Bids, with an Application to Real Estate." Mimeo, University of Texas, 1995.

McMillan, John, "Selling Spectrum Rights," *Journal of Economic Perspectives*, 1994, vol. 8, no. 3, pp. 145-162.

Paarsch, Harry, "Deciding between the common and private value paradigms in empirical models of auctions." *Journal of Econometrics*, vol. 51, pp. 191-215.

Riley, John G., and William Samuelson, "Optimal Auctions," *American Economic Review*, 1981, vol. 71, no. 3, pp. 381-392.

Smith, Vernon L. "Experimental Studies of Discrimination Versus Competition in Sealed-Bid Auction Markets." *Journal of Business*, 1967, vol. 40, pp. 56-84.

Smith, Vernon L. "Microeconomic Systems as an Experimental Science." *American Economic Review*, 1982, vol. 72, no. 5, pp. 923-955.

Top 40 Newsgroups in Order by Traffic Volume, `<ftp://rtfm.mit.edu/pub/usenet/news.lists/>`, April 1995.

Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 1961, vol. 16, no.1, pp. 8-37.

Vickrey, William. "Auction and Bidding Games," in *Recent Advances in Game Theory*, O. Morganstern and A. Tucker, eds. Princeton: Princeton University Press, 1962.

Wilson, Robert. "Strategic Analysis of Auctions," in *The Handbook of Game Theory*, R.J. Aumann and S. Hart, eds. New York: North-Holland, 1992, pp. 227-279.

Appendix. A Sample Auction Announcement.

Date: Tue, 28 Feb 1995 18:19:10 -0500
To: reiley@MIT.EDU (David Reiley)
From: reiley@MIT.EDU (David Reiley)
Subject: Reiley's Auction #5: ARABIANS, Free Shipping!

Hi! As a participant in one of my previous auctions, I thought you might be interested in this NEW AUCTION opportunity.

Please read the rules of this auction carefully, as each auction I run typically has a different set of rules.

This will be a SEALED-BID, FIRST-PRICE AUCTION. Here's how it works:

I will accept all bids up until the deadline of NOON (Eastern Standard Time), next TUESDAY, March 7, 1995. All bids are "sealed" in the sense that I will not post updates or otherwise reveal information about the highest bid until the auction is over.

After the deadline for bids has passed, I will award each card to the highest bidder at the price of their bid. The exception is that if no one bid at least the posted minimum bid for some card, that card will not be sold.

Note that SHIPPING IS INCLUDED in the bid price. If you win, you mail me exactly the amount of your bid, with no extra charges. This is to encourage everyone to bid separately on each card they're interested in - no worrying about having to win multiple cards in order to make it worth your while.

Here are THE RULES:

1. Submit bids via email to <reiley@mit.edu>. Make sure that the subject line of your email contains the phrase "Auction #5". (Simply using the "reply" command on most mail and news programs should work just fine.) If your message does not contain this text in the subject line, your message will be discarded.

2. In your email message, please put each of your bids on a separate line
of text.
Each bid line should be in the following format: the 3-digit identification number
of the card you're bidding on, immediately followed by a right parenthesis, and

then the amount of your bid in dollars and cents (such as 1.00). For example:

```
305) 2.00
306) 0.65
```

Including extra information on the bid line is okay, too. Such extra information might include email quote marks (such as the greater-than symbol), the card name, condition, etc. You may include anything that makes bidding easier for you, EXCEPT that your bid amount should be the only price-formatted number which appears on that line. In other words, no other number containing a decimal point should appear on that line.

For example, the following are also perfectly valid bids:

```
> 305) Ghosts of the Damned          U1     Blk   M  2.00
> 306) Demonic Torment   0.65
```

Bids that do not conform to these rules will be discarded.
Here are examples of INVALID bids:

```
>305) Ghosts of the Damned      $2                      [no decimal point]
306 ) Demonic Torment 0.65  [space between the 6 and the right parenthesis]
```

3. All bids must be in integer multiples of a nickel ($0.05), in US currency.

5. The auction closes on Tuesday, March 7, 1995, at noon (EST). Any bids received after that time will be ignored. All cards receiving a bid of at least the posted minimum bid will be sold at that point to the highest bidder. In the case of a tie, the winner will be the person whose bid was received first.

6. The winning bidder will be notified by email, and will be asked to pay the amount of his/her bid via US check or money order.

7. This payment will include free shipping within the United States, via first class mail. The cards will be wrapped in plastic sheaths and packed in cardboard for protection. All cards will be shipped within one week after the receipt of payment.

8. While this is a real auction for real cards, you should know that I plan to use data on the bids in this auction for economic research. In no case will individual bidders be identified in this research; anonymity will be preserved. By bidding in this auction, you indicate your consent to have your bid be used in economic research. If you do not approve of this, you

have the right not to participate in this auction. Should you have any questions or concerns about the use of data from this auction in academic research, please contact the chair of the COUHES committee at MIT by phone at 617-253-6787.

That's it! Enjoy the auction. Good luck, and thanks for participating!

Here is the LIST OF CARDS:

ID	Card Name	Rarity	Color	Cond	Minimum
501)	Bazaar of Baghdad	U3	Lnd	M	8.90
502)	City of Brass	U3	Lnd	M	17.85
503)	Desert	C11	Lnd	M	3.55
504)	Diamond Valley	U2	Lnd	M	26.40
505)	Elephant Graveyard	U2	Lnd	M	18.30
506)	Island of Wak-Wak	U2	Lnd	M	20.10
507)	Library of Alexandria	U3	Lnd	M	21.00
508)	Mountain	C1	Lnd	M	2.45
509)	Oasis	U4	Lnd	M	5.65
510)	Aladdin's Lamp	U2	Art	M	3.10
511)	Aladdin's Ring	U2	Art	M	3.05
512)	Bottle of Suleiman	U2	Art	M	2.75
513)	Brass Man	U3	Art	M	0.90
514)	City in a Bottle	U2	Art	M	10.60
515)	Dancing Scimitar	U2	Art	M	2.55
516)	Ebony Horse	U2	Art	M	3.05
517)	Flying Carpet	U3	Art	M	2.95
518)	Jandor's Ring	U2	Art	M	2.70
519)	Jandor's Saddlebags	U2	Art	M	2.80
520)	Jeweled Bird	U3	Art	M	5.40
521)	Pyramids	U2	Art	M	16.20
522)	Ring of Ma'ruf	U2	Art	M	21.85
523)	Sandals of Abdallah	U3	Art	M	6.30
524)	Cuombajj Witches	C4	Blk	M	2.40
525)	El-Hajjaj	U2	Blk	M	2.85
526)	Erg Raiders	C5	Blk	M	0.50
527)	Guardian Beast	U2	Blk	M	40.10
528)	Hasran Ogress	C5	Blk	M	1.25
529)	Junun Efreet	U2	Blk	M	13.50
530)	Juzam Djinn	U2	Blk	M	27.00
531)	Khabal Ghoul	U3	Blk	M	21.45
532)	Oubliette	C4	Blk	M	3.75
533)	Sorceress Queen	U3	Blk	M	5.05

534) Stone-Throwing Devils	C4	Blk	M	2.40
535) Dandan	C4	Blu	M	1.50
536) Fishliver Oil	C4	Blu	M	1.55
537) Flying Men	C5	Blu	M	2.40
538) Giant Tortoise	C4	Blu	M	1.30
539) Island Fish Jasconius	U2	Blu	M	3.20
540) Merchant Ship	U3	Blu	M	6.15
541) Old Man of the Sea	U2	Blu	M	24.65
542) Serendib Djinn	U2	Blu	M	9.95
543) Serendib Efreet	U2	Blu	M	3.95
544) Sindbad	U3	Blu	M	7.30
545) Unstable Mutation	C5	Blu	M	0.50
546) Cyclone	U3	Gre	M	6.30
547) Desert Twister	U3	Gre	M	1.95
548) Drop of Honey	U2	Gre	M	17.50
549) Erhnam Djinn	U2	Gre	M	15.70
550) Ghazban Ogre	C4	Gre	M	1.40
551) Ifh-Biff Efreet	U2	Gre	M	16.10
552) Metamorphosis	C4	Gre	M	1.60
553) Naf's Asp	C4	Gre	M	1.35
554) Sandstorm	C5	Gre	M	1.85
555) Singing Tree	U2	Gre	M	20.20
556) Wyluli Wolf	C5	Gre	M	3.30
557) Aladdin	U2	Red	M	16.45
558) Ali Baba	U3	Red	M	9.25
559) Ali from Cairo	U2	Red	M	39.95
560) Bird Maiden	C4	Red	M	2.15
561) Desert Nomads	C5	Red	M	2.00
562) Hurr Jackal	C4	Red	M	1.80
563) Kird Ape	C5	Red	M	0.60
564) Magnetic Mountain	U3	Red	M	2.55
565) Mijae Djinn	U2	Red	M	3.90
566) Rukh Egg	C4	Red	M	5.85
567) Ydwen Efreet	U2	Red	M	6.60
568) Abu Ja Far	U3	Whi	M	11.10
569) Army of Allah	C4	Whi	M	3.25
570) Camel	C5	Whi	M	1.45
571) Eye for an Eye	U3	Whi	M	3.20
572) Jihad	U2	Whi	M	23.00
573) King Suleiman	U2	Whi	M	11.85
574) Moorish Cavalry	C5	Whi	M	2.75
575) Piety	C4	Whi	M	1.45
576) Repentant Blacksmith	U2	Whi	M	9.40
577) Shahrazad	U2	Whi	M	11.95

578) War Elephant C4 Whi M 3.50

END OF LIST

Key to Card Types and Conditions:

C1 = Common
C2 = Common (twice as common as C1, printed twice per common sheet)
C4 = Common (four times as common as C1, printed four times per common sheet)
C11 = Common (printed 11 times per common sheet)
U1 = Uncommon
U2 = Uncommon (twice as common as U1, printed twice per uncommon sheet)
R = Rare

Lnd = Land
Art = Artifact
Blk = Black
Blu = Blue
Gre = Green
Red = Red :)
Whi = White

M = Mint
NM = Near Mint (never played, but has tiny blemishes from handling)
E = Excellent (played a few times, has small scuff marks)

If you are unfamiliar with some of these cards, you can get information about any Magic card (spell type, power, toughness, artist's name, etc.) from the following Internet sources:

http://www.itis.com:80/deckmaster/magic/cardinfo/
ftp://marvin.macc.wisc.edu/pub/deckmaster/card.info/lists.w.spoilers/

Remember to send any bids, comments, or questions about the cards or the rules of this auction to <reiley@mit.edu>.

Thanks!

Chapter 2:

Tests of Revenue Equivalence in Internet Magic Auctions

1 Introduction

In recent years, auctions have steadily increased in importance, both because of their practical significance in the economy and because auction theory has been one of the most interesting and elegant applications of game theory. Auction theory has been a particularly productive area of economic theory in terms of generating empirically testable predictions.

One of the most basic theorems in auction theory is known as the revenue equivalence theorem, which dates back to Vickrey (1961). Comparing four different basic auction formats (English, Dutch, first-price, and second-price auctions), Vickrey showed that in a simple model of consumer preferences and Nash equilibrium bidding behavior, the expected revenue to be collected by the auctioneer will be the same no matter which of the four auction mechanisms she chooses. Since 1961, auction theorists have extended the revenue equivalence theorem to more general models of bidder preferences, and have pointed out circumstances in which revenue equivalence might fail to hold - where one auction mechanism might be expected to generate more revenue than another. In this chapter, I report the results of a set of field experiments designed to test the revenue equivalence theorem, comparing the revenues earned when selling the same goods under each of the four basic types of auction in a particular market.

During the past two years, a fascinating new market has sprung up on the Internet. It is a market for collectible cards from the game *Magic: the Gathering*, a game whose retail success is a story in itself. Launched in August 1993, this product has already grossed hundreds of millions of retail dollars, and now has over a million players worldwide. There are more than a thousand distinct types of cards which have been printed for use in this game, each of which has a slightly different role in game play. In the game scenario, players assume the roles of dueling wizards, each with their own libraries of magic spells (represented by decks of cards) that may potentially be used against the player's opponent. Cards are sold in random assortments, just like baseball cards, at retail stores ranging from small game and hobby shops to large chains such as Toys 'R' Us and Waldenbooks.

Perhaps the most interesting thing about this product is its interaction with the Internet to create a thriving, online (secondary) exchange economy, in which each Magic card is a separate commodity. The Internet, with its convenient methods of transmitting messages, has facilitated new technologies for trading among individuals. For example, a Magic player with a few hundred unwanted cards can auction them off to the highest bidders online, in a process which is considerably easier even than the traditional practice of holding a garage sale for excess household goods. The Internet lowers transaction costs, enabling ordinary individuals to make trades without the assistance of retail or auctioning specialists. Transaction costs are particularly low for cards as opposed to, for example, computer hardware, because it is so easy to mail cards across the country, and thus this card market has flourished.

In 1994 and 1995, I observed auctioneers in this market employing a variety of auction mechanisms, including English, first-price sealed-bid, and even Dutch auctions, along with other variants and hybrids of these three auction types. The only basic mechanism from auction theory that I did not witness was the second-price sealed-bid auction. Thus, when I began to run my own auctions with a variety of different rules, I felt I could expect bidders to treat me no differently than they would treat any other auctioneer in this market. This provides a unique opportunity to perform tests of the revenue equivalence theorem in auctions for real goods. For more details about the history and structure of the Internet market for Magic cards, see Chapter 1.

A number of experimenters have performed previous tests of the revenue equivalence theorem in the laboratory, and have found some surprising violations of the theory. A particularly strong example is that laboratory experimenters have found consistent evidence that first-price sealed-bid auctions raise more revenue on average than do Dutch auctions, despite the fact that these two auction types are theoretically equivalent. The present study attempts to replicate this and other laboratory findings in a real-world market. The most interesting finding is that in comparing first-price with Dutch auctions, the result is the exact opposite of that in the laboratory experiments: Dutch auctions raise more revenue on average than first-price auctions.

The chapter is organized as follows. In Section 2, I review the theory of revenue comparisons between the four basic auction mechanisms: English, Dutch, first-price sealed-bid, and second-price sealed-bid. I discuss situations of strategic equivalence for bidders as well as possible equivalences in overall expected revenue. Section 3 reviews the findings of previous empirical investigations of revenue equivalence, most of which have taken place in laboratory settings. Section 4 describes my experimental design: a series of auctions of matched pairs of Magic cards. This section includes discussion of some of the features which distinguish these field experiments from previous laboratory experiments, including the simultaneous rather than sequential nature of the auctions, and the existence of variable entry by bidders, and some potential time-order problems not usually encountered in the laboratory. Results of the data analysis can be found in Section 5, followed by concluding remarks in Section 6.

2 Theory of the Four Basic Auction Mechanisms

The four basic auction mechanisms, as outlined by Vickrey (1961), are the English, Dutch, first-price sealed-bid, and second-price sealed-bid auctions. Probably the most familiar auction format is the English auction, in which bidders continue to submit increasingly high bids until all have dropped out of the bidding except for the winning bidder, who receives the good by paying the amount of his high bid. Instead of the steadily increasing prices of the English auction, the Dutch auction involves decreasing prices: a public price clock starts out at some very high level, and the price falls until the point when the first participant finds the price low enough to submit a bid. The first bidder is declared the winner of the Dutch auction, and receives the good at the price at which he stopped the clock. Dutch auctions have been used regularly for sales of tulip bulbs in Holland, and a form of Dutch auction is also used by clothing stores with automatic markdown policies, such as Filene's Basement in downtown Boston. These two auction types can be grouped together as "real-time" auctions, to distinguish them from auctions in which bids are sealed and there is no real-time bidding process involved.

First-price sealed-bid auctions frequently take place in the economy, such as in bidding for construction contracts. Each bidder has the opportunity to submit a single bid by a particular deadline, and after the deadline expires, the bids are examined and the highest bidder wins the good at the price she bid. (Bidding for construction contracts typically takes place as a buyer's auction rather than a seller's auction, so the construction contract is actually awarded to the lowest bidder, but the concept is essentially the same.) The second-price sealed-bid auction follows essentially the same procedure as the first-price sealed-bid auction, except that the winning bidder has to pay not the amount of her own bid, but the amount bid by the second-highest bidder. This type of auction is rarely seen in practice; it was proposed by Vickrey (1961) because of its desirable theoretical properties, which will be discussed in the next section. The most visible use of a second-price auction in the real world was the New Zealand government's auction of communications spectrum rights in 1990.[1]

2.1 Strategic Equivalence

One strong type of equivalence which can be found between two auction types is *strategic* equivalence. If auctions A and B are strategically equivalent, then not only do they produce the same amount of revenue for the auctioneer, they also have the property that a bidder in auction A would follow the same strategy that she would in auction B.

This property is best understood by considering specific examples. For example, it turns out that in a *private-values* model, the second-price sealed-bid auction format is strategically equivalent to the English auction format. A private-values model is one in which each bidder knows with certainty her own valuation v_i for the good, so that if she wins the auction and pays an amount x for the good, she earns a surplus of v_i - x, while if she does not win the good she gets a payoff of 0.

To see the strategic equivalence of the English and second-price auctions, consider the optimal strategy for a bidder in an English auction. The optimal strategy for bidder i is to continue raising other people's bids by some small amount until the bid level exceeds her valuation v_i for the good. Clearly, she wouldn't want to drop out of the bidding any earlier than this, because then she might miss out on the opportunity to earn a positive surplus by winning the good at a price less than v_i. Nor would she continue bidding after the bid level exceeds her valuation, because she wouldn't gain anything from this, and would only suffer the risk of winning the good at a price high enough to give her a negative payoff. The case of the second-price auction turns out to be strategically equivalent: in a second-price auction, a bidder with valuation v_i should submit a bid equal to v_i in order to maximize her payoff. Why should she submit a bid equal to her valuation? Consider instead submitting a bid higher than her valuation. In cases where her highest rival's bid is less than v_i, this higher bid would make no difference to her payoff, but in cases where her highest rival's bid is more than v_i, she would be running some risk of exceeding her highest rival's bid and earning a negative surplus. So a bid higher than v_i makes her worse off. Now consider bidding

[1] See McMillan (1994), p. 148, for details.

lower than v_i. In this case, she runs the risk of losing the auction in cases where her rivals all bid less than v_i, but does not gain anything in cases where at least one of her rivals bids more than v_i (she would not win the auction in either case). So a bid lower than v_i also makes her worse off, and thus it is optimal to bid her valuation. Therefore, the English and second-price auctions are strategically equivalent: in both cases, it is a dominant strategy to bid one's valuation. (It also turns out that in both cases, the winning bidder will end up paying the amount of the second-highest bidder's valuation, assuming that bid raises in the English auction can be infinitesimally small, so that bidding can stop as soon as the second-highest bidder drops out.)

A second example of strategic equivalence is that in the private-values model, the Dutch auction format is strategically equivalent to the first-price sealed-bid auction format. To see this, suppose that there are N bidders with private valuations v_i $(i=1,2,3,...,N)$ drawn from some probability distribution $F(v_1,v_2, v_3,...,v_N)$.[2] Unlike in a second-price sealed-bid auction, it turns out that in a first-price sealed-bid auction there is no dominant strategy equilibrium. We suppose that there will be a Bayesian Nash equilibrium in which each bidder attempts to maximize her payoff given her beliefs about the probability distribution of the highest of her rivals' bids, and these beliefs about rivals' bidding behavior will turn out to be correct in equilibrium. There will be some equilibrium bid function $b(v)$ that gives the optimal bid level for each possible valuation. Generally, $b(v)$ will be strictly less than v, so that the winning bidder always earns a positive surplus, but the exact functional form of the Nash equilibrium bid function depends on the functional form of the distribution of valuations. The essential concept is that each bidder must decide on a bid level, given her knowledge of her own valuation, her knowledge of the probability distribution of rival bidders' valuations. Now consider the Dutch auction, in which each bidder must decide how long to wait before submitting a bid at the current clock price. Just as in the first-price auction, the winning bidder in a Dutch auction wins the good and pays the amount of her bid. Deciding how long to wait before bidding the clock price is equivalent to answering the following question: "At what price will I bid, assuming that no one else has already submitted a higher bid than me?" Thus, the decision in the Dutch auction boils down to the decision of what price to bid, which makes the decision exactly equivalent to the case of the first-

[2] In Vickrey's original (1961) paper on auctions, he demonstrated his results for the case of *independent* private values (IPV). It turns out that independence is not necessary for the strategic equivalences to hold; bidder's valuations can have some known correlation with each other, but as long as the private valuations are known with certainty to each individual bidder, the English-second and the Dutch-first pairs are each strategically equivalent.

price sealed-bid auction. One might at first think that the time process of the Dutch auction would make a difference, because a bidder in a Dutch auction could decide to condition her strategy on the information about how other bidders had behaved in the early part of the auction. But the only time a bidder has the opportunity to bid is in the case when no one else has bid so far in the auction, so the only kind of information that she can use to condition her bid strategy is the information that no one else has bid so far. Thus, a bidder could decide on her equilibrium strategy before the auction even started, and no relevant information is gained by waiting. Because the rules of the game and the strategy decision are the same as in the first-price auction, the equilibrium strategy must also be the same. A bidder would bid the same amount if the rules of the auction were Dutch as if they were first-price sealed-bid, and therefore these two auction formats are strategically equivalent.

2.2 Vickrey's Revenue Equivalence Theorem

In addition to demonstrating strategic equivalence between English and second-price and between Dutch and first-price auction, Vickrey (1961) also demonstrated something more: that in an auction with independent private values (IPV) and risk-neutral bidders, *all four basic auction types yield the same expected revenue to the auctioneer.*[3] As discussed earlier, he same bidder would submit a higher bid under second-price auction rules (a bid equal to her valuation) than she would under first-price auction rules (a bid strictly less than her valuation). Interestingly, it turns out that the expected amount that the winner bids in the first-price auction is equal to the expectation of the second-highest valuation, and thus the expected payment from the winner in a first-price auction is equal to the expected payment from the winner in a second-price auction.

To formalize this statement, consider an auction with N risk-neutral bidders, each of whom has a privately-certain valuation drawn from a continuous distribution with distribution function F(x). Now suppose bidder i with valuation v_i for the good, and let $b(v_i)$ be the amount that she would optimally choose to bid in a Bayesian Nash equilibrium. Supposing that bidder i will be the winning bidder (the winning bidder in this equilibrium model is always the bidder with the highest valuation), one can compute the expected value of the highest of her rivals' valuations conditional on the fact that v_i is the highest of the N draws from the probability distribution. The conditional expected value of the highest of the val-

[3] A particularly readable exposition of this topic is the one by Riley (1989). Other useful references, including some generalizations of the Revenue Equivalence Theorem, can be found in the survey articles by McAfee and McMillan (1987) and by Wilson (1992).

uations of bidder i's rivals, conditional on bidder i having the highest valuation, turns out to be precisely equal to $b(v_i)$, which is bidder i's bid in the first-price auction. Thus, when the winning bidder in a first-price auction has valuation v_i, her payment to the auctioneer will be an amount $b(v_i)$ equal to the expected value of the second-highest bidder valuation conditional on v_i being the highest valuation. Since this is true for every possible value of v_i, it must also be true when one takes an expectation over all possible highest valuations: the amount paid by the highest bidder in a first-price auction is equal to the amount paid by the highest bidder in a second-price auction. And because of the strategic equivalences discussed in section Section 2.1, it follows that the expected payment of the winning bidder is the same in *all four* of the basic auction formats.

2.3 Theoretical Violations of Revenue Equivalence

In the thirty-five years since the publication of Vickrey's paper on the symmetric, risk-neutral independent-private-values (IPV) model, auction theorists have made comparisons of the four basic auction types under a variety of different models of bidder preferences and information structures. Under these different theoretical assumptions, it has been found that the Revenue Equivalence Theorem does not always hold. From the perspective of the present study, two directions of this research are particularly notable: generalizations to models with bidder risk aversion and to models where bidders' valuations are affiliated rather than independent and private.

2.3.1 Bidder Risk Aversion

When bidders are risk-averse, the Revenue Equivalence Theorem does not hold: the first-price auction yields more revenue than does the second-price auction. Mathematical details can be found in Riley and Samuelson (1981) or Maskin and Riley (1984), but here I present some economic intuition for this result. Compared with a risk-neutral bidder, a risk-averse bidder would be willing to accept a lower expected surplus in exchange for an increased probability of getting the positive payoff associated with winning the auction. This means that a risk-averse bidder would be willing to submit a higher bid in a first-price auction than a risk-neutral bidder would, and therefore that a first-price auction with risk-averse bidders will raise more expected revenue than a first-price auction with risk-neutral bidders. By contrast, risk aversion does not affect the outcome of a second-price auction, because risk aversion does not change the fact that it is still a dominant strategy for a bidder to bid her valuation. (In a second price auction, raising one's bid above one's valuation does not raise the probability of winning except in cases where the payoff would be negative.) So increased risk aversion increases revenue in a first-price auction, but not in a second-price auction. This fact, combined with the Revenue Equivalence Theorem for risk-neutral bidders, implies that with risk-averse bidders, expected revenue will be higher in a first-price auction than in a second-price one. (The Dutch-first and English-second strategic isomorphisms still hold, however.)

2.3.2 Affiliated Valuations

The second kind of extension to the IPV model that I wish to consider is the *affiliated-values* model, developed by Milgrom and Weber (1982) as a generalization which includes as special cases both the IPV model and the common-values model.

By contrast with a private-value auction, a common-value auction is one in which the good being auctioned has a fixed value which is the same for all bidders, but none of the bidders knows exactly what this value is. Usually a common-value auction is modeled by supposing that each bidder gets a different imprecise signal of the true common value of the good. An example of a common-value auction is an auction for offshore oil drilling rights, in which there is some fixed but unknown quantity of oil reserves available, and each oil company interested in participating in the auction has their own estimate (a noisy signal) of the true market value of the oil reserves.

In common-value auctions, an interesting problem is that of the "winner's curse." That is, the winning bidder will be the bidder with the most optimistic estimate of the good's value, so he may end up paying too much for the good unless he is careful. Each bidder should in a sense treat his signal as if it were an overestimate of the good's true value, because the only times his bid will actually have to be paid are the times when his signal is the highest among all the bidders. A naive bidder who wins the auction may end up paying more for the good than it is actually worth; this is the potential phenomenon known as the "winner's curse." In common-value auctions, rational bidders should bid conservatively in order to avoid this sort of negative outcome.

Milgrom and Weber (1982) have developed a model called the *affiliated values* model, which is a generalization that includes both the private-values and the common-values models as special cases. In the affiliated-values model, a bidder's valuation may depend on both a privately-certain component and a privately-uncertain (common-value) component. An example of such an auction with both private-value and common-value aspects is an auction for a fine painting, in which each bidder has a different private value of being able to view the art in his home, but in which each bidder also values the investment potential of the painting, which is a speculative future resale value (a common value) for which each bidder may have a different estimate.

In their paper, Milgrom and Weber show that revenue equivalence fails in the general affiliated-values model. The revenue in an English auction is greater than the revenue in a second-price auction, which in turn is greater than the revenue in a first-price or Dutch auction. The strategic equivalence of the Dutch and first-price auctions still holds in this model, but the strategic equivalence of the English and second-price auctions fails. The reason for this is the winner's-curse effect, which occurs whenever bidders are uncertain of their own valuations for the good and they believe that their uncertain signals of value are positively correlated with the signals of rival bidders. As discussed above, rational bidders will bid conservatively in the presence of the winner's-curse effect. To encourage bidders to bid less conservatively, the auctioneer may provide them with more information about the uncertain value of the good, thus reducing a bidder's risk of making a terribly overoptimistic estimate of the good's value. Note that in comparison with a first-price auction, a second-price auction provides bidders with more of this type of insurance, because in a second-price auction the winner's payment depends not just on his own bid (and signal of value), but also on the second-highest bidder's signal of value. Thus, a second-price auction reduces the effects of the winner's curse, which allows rational bidders to bid less conservatively, and thus raises expected auction revenue over that of a first-price auction. Similarly, an English auction provides bidders with even more information about other bidders' signals of value, because in addition to observing the second-highest bidder's bid, the winner also may get to observe some of the bids by other bidders in the auction. Thus, by reducing even further the bidders' need for conservative bidding, an English auction format yields higher expected revenue than a second-price auction. The additional information revealed about other bidders in the real-time English auction (as compared with the second-price sealed-bid auction) turns out to be important in the case of affiliated values, even though it is irrelevant in the private-values case.

Another interesting special case of the Milgrom-Weber model is that of affiliated private values. In this case, we assume that values are private in the sense that they are known with certainty to each individual bidder, but they are affiliated in the sense that bidders' valuations are positively correlated with each other rather than being independent draws from a known distribution. For example, suppose that bidder valuations are known to be drawn from a uniform distribution $U[a,a+1]$, but where the lowest value a is itself uncertain, having a uniform distribution on [0,1]. Then bidder valuations can in principle take any value between 0 and 2. But if a bidder finds that her valuation is v_i=1.2, she will know that it is likely that other bidders' valuations are likely to be relatively high (in particular, valuations between 0 and 0.2 are impossible). The uncertainty about valuations is symmetric: bidder i's subjective probability $F(v_j|v_i)$ that bidder j has a valuation no greater than v_j, conditional on knowing her own valuation v_i, is the same for all bidders $j \neq i$. In such cases, the strategic equivalence between English and second-price auctions still holds, because both auctions have the same dominant strategy when private valuations are known with certainty. However, the general revenue equivalence theorem fails: expected revenue is strictly higher in a second-price auction than in a first-price auction.

The intuitive explanation for this phenomenon is perhaps not obvious, as this failure of revenue equivalence in affiliated-private-value auctions is clearly not due to winner's-curse effects (as explained above and most commonly cited in the literature). The proof of the revenue equivalence theorem in the IPV case requires showing that a bidder with valuation v in a first-price auction will optimally bid an amount equal to the expectation of the highest of her rivals' valuations, conditional on her own valuation v being the highest of all valuations. Thus, in the IPV case, a bidder will submit a first-price auction bid equal to the amount she would expect to pay if she were following her dominant strategy in a second-price auction. By contrast, consider the affiliated-private-values example discussed in the above paragraph. The bidder with valuation v_i=1.2 knows that the parameter a may take on any value between 0.2 and 1, and that all of these values are equally likely. If the value of a were common knowledge, it would be an IPV model of valuations, and her Nash equilibrium bid in a first-price auction would be given by the following IPV formula:

$$b(v_i, a) = a + \frac{N-1}{N}(v_i - a) \qquad (1)$$

This is also equal (by Vickrey's revenue equivalence theorem) to the expected revenue, conditional on bidder i having the highest valuation, in a second-price auction where the parameter a is known. When the parameter a is unknown, the expected revenue if bidder i wins in a second-price auction is equal to the expectation of (1) over all possible values of a; for example, for the bidder with v_i=1.2, the expectation should be taken over all a between 0.2 and 1. But it turns out that bidder i's Nash equilibrium bid in a first-price auction is less than this expectation. The intuitive reason is that states of the world where a happens to be low (around 0.2) are more important to bidder i than states of the world where a is high (around 1), because when a is relatively low she is relatively more likely to actually be in the running to win the auction (i.e., to be close to having the highest valuation). Since the IPV bid function (1) is increasing in a, and since bidder i is influenced more by states of the world where a is small than states where it is large, her bid when a is uncertain will be lower than the expectation of (1) over all possible values of a. Therefore, in a first-price auction with affiliated private values, bidding competition is softened by the fact a bidder knows that her bid matters more (in terms of probability of winning) in states of the world where her valuation is relatively high than in states of the world where her valuation is relatively low. In other words, when valuations are privately certain to bidders, competition is softer in first-price auctions when the distribution of valuations is an uncertain function. This softened competition in first-price auctions is what causes the violation of revenue equivalence, and makes expected revenue lower in a first-price auction than in a second-price auction.

2.3.3 Summary

By relaxing the assumptions of the Vickrey auction model (risk neutrality and independent private values), one can obtain situations in which revenue equivalence does not hold. When bidder risk neutrality is replaced by risk aversion, the strategic isomorphisms of the two pairs of auctions still hold, but the Dutch and first-price auctions yield greater expected revenue than do the English and second-price auctions. However, the opposite revenue ranking occurs if the independent-private values assumption is relaxed: under affiliated bidder values, the English and second-price auctions raise more expected revenue than do the Dutch and first-price auctions. In addition, if the affiliated values are not strictly private but instead contain some common-value component, the strategic equivalence between the English and second-price auctions disappears, and the full revenue ranking for the auctioneer is English followed by second-price followed by first-price. Only the equivalence between the Dutch and first-price auctions remains. In a model with both risk aversion and affiliated values, the auctions cannot be easily ranked in terms of revenue because of the competing effects.

3 Previous Empirical Studies of Revenue Equivalence

It is reasonably difficult to obtain field IO data on real-world auctions that would allow one to test the equivalences between the basic auction types. Auction houses such as Sotheby's typically do not run both a second-price auction and an English auction for the same good to see whether both formats yield the same amount of revenue, nor does the Treasury Department in its T-bill auctions. I am aware of only one set of field data which allows comparisons between auction formats: a set of data on English and first-price auctions run by the U.S. Forest Service for timber harvesting rights. Using this data, Mead (1967) found a tendency for the first-price auctions to raise more revenue than the English auctions. However, Hansen (1985) pointed out a selection bias in this data, caused by the way the Forest Service chose which auction to us for each timber lot; after correcting for this bias, he found that the increased revenue in English auctions no longer could be said to be statistically significant. Thus, despite the fact that revenue equivalence is one of the most fundamental questions in auction theory, there has not yet been any definitive evidence on this question from empirical studies of real-world auctions.[4]

On the other hand, there have been a number of interesting tests of the revenue equivalence theorem performed in laboratory experiments. In the laboratory, economists enjoy precisely the opportunity to auction the same good twice with different mechanisms. The laboratory experimenter also has some control over the structure of bidder preferences. A laboratory experiment is typically conducted as follows: the experimenter tells the experimental subjects that they will be bidders in an auction for a fictitious good, which is worth different amounts of money to different people. Each bidder learns the probability distribution of bidder valuations, which is typically a uniform distribution (the easiest type of probability distribution to explain to subjects), and then he learns his own valuation for that auction. Bidding takes place, with the winning bidder paying to the auctioneer the amount specified by the rules of the auction and then selling the good back to the experimenter for the amount of his valuation. Thus, the winning bidder earns a cash payoff equal to his valuation minus his winning payment for the auction. In a single experimental session, the same subject pool typically participates in a series of perhaps ten, twenty, or thirty different auctions.

In this section, I review the results from laboratory studies of revenue equivalence between different auction formats. Most of these studies have concentrated on the independent-private-values (IPV) model, which was the first auction model to be understood theoretically and is a straightforward experimental design both to implement and to explain to

---

[4] For a review of the auction topics which *have* been studied empirically with field data, see the review by Hendricks and Paarsch (1995).

subjects. Even with such a simple model, the results have been very interesting, and not always in agreement with the theory.

3.1 Dutch-First Strategic Equivalence

The first comprehensive revenue-equivalence study was that of Coppinger, Smith, and Titus (1980). They ran auctions of each of the four basic types, with some experimental sessions consisting of Dutch and English auctions, and other sessions consisting of first- and second-price sealed-bid auctions. In each case, since the experimenters knew the distribution of bidder valuations (because they assigned the valuations to the bidders), they could compare the outcome to the theoretically predicted risk-neutral Nash equilibrium outcome.

Although they did not run experiments specifically pairing the Dutch and first-price auction formats against each other for direct comparison, they did provide strong indirect evidence against the theory of the Dutch-first strategic equivalence. Revenue in first-price sealed-bid auctions was significantly higher than the risk-neutral Nash equilibrium (RNNE) prediction, while revenue in Dutch auctions was approximately equal to or slightly below the RNNE level. Prices above the RNNE level are consistent with the theory of risk-averse bidders in a first-price auction, and since risk aversion was not something that could be controlled or measured in this experiment, that could be an explanation for the higher revenue in the first-price auctions. However, it does not explain why the Dutch auction failed to produce as much revenue as the first-price auction, since those two auction types should be strategically isomorphic under any sort of bidder risk preferences.

To follow up on this topic, Cox, Roberson, and Smith (1982) conducted another experimental study of bidder behavior in Dutch, first-price, and second-price (but not English) auctions. This time, there were direct paired comparisons of the Dutch and first-price auctions both within and between experimental sessions. Each experimental session consisted of thirty auction rounds, with each bidder getting an independent draw from the common-knowledge (uniform) distribution of valuations for each of the thirty rounds. Some of the experimental sessions consisted of 10 first-price, then 10 Dutch, followed by 10 more first-price auctions (FDF format), while others consisted of 10 Dutch, then 10 first-price, followed by 10 more Dutch auctions (DFD format). The reason for the reversal of the ordering of the auction formats was to allow comparisons of Dutch versus first-price bidding behavior within a single experimental session, controlling for possible hysteresis effects which might cause behavior in one auction format to carry over to the other auction format. In addition, each of the FDF sessions was paired with one of the DFD sessions, with the exact same realization of bidder valuations duplicated between the two sessions, but with different subjects. This technique allowed for direct comparisons between Dutch and first-price auction formats with the set of bidder valuations controlled to be identical between the two formats. This second study also added several other methodological improvements over the 1980 paper, including more careful randomization of valuations and standardization of allowable bid increment levels to be equal across auction formats.

Cox, Roberson, and Smith (1982) found the same kind of violation as before of the strategic equivalence between Dutch and first-price auctions. Dutch auction revenues were consistently lower on average than first-price auction revenues, which was particularly clear when comparing two experimental sessions with different auction formats but the same realization of the set of bidder valuations. This time the Dutch auction revenues tended to be slightly higher than the predicted RNNE level, but the first-price auction revenues were still higher. The authors proposed two different explanations of the failure of the strategic equivalence: that bidders enjoyed extra utility from the "suspense of waiting" in the Dutch auction (which was suggested by some observed bidder comments), or that bidders in the Dutch auctions underestimated the probability of losing the auction when waiting for the next tick of the Dutch price clock (by miscalculating their Bayesian updates, as has been observed in other experimental situations).

Cox, Smith, and Walker (1983) attempted to discriminate between the two models of the nonisomorphism by using the overall level of the distribution of valuations as a treatment variable. Specifically, they ran additional experiments in which the valuations of bidders were triple the levels of valuations in the experiments for the 1982 paper. If the "utility of suspense" explanation were correct (and if the utility of suspense is not affected by the rescaling of bidder valuations), then the tripling of values should cause Dutch auction bids to be proportionately closer to the first-price auction bids. The authors found that this experimental treatment did *not* cause Dutch auction revenues to increase, and thus they rule out the "utility of suspense" theory in favor of some other explanation, such as the probability miscalculation theory.

In summary, previous results have shown that the Dutch and first-price auction formats are *not* strategically equivalent in laboratory experiments, and there has been some success in learning why this equivalence fails. In this chapter, I will attempt to see whether these results can be replicated in an environment outside the laboratory.

3.2 English-Second Strategic Equivalence

As noted above, Coppinger, Smith, and Titus (1980) presented the first (and only) study comparing all four of the basic auction formats in an environment with independent private values (IPV). In terms of the English-second strategic equivalence, they found that the English and second-price auctions yielded approximately equivalent revenue. English auctions usually produced a price slightly above the theoretical prediction, which, they note, is to be expected given that the minimum bid increment could not really be infinitesimal as in the Vickrey theory. Second-price auctions tended to produce a price slightly below the Nash equilibrium prediction, although prices increased in repeated sessions to approach the predicted value; the authors interpreted this as evidence of bidder learning about an unfamiliar mechanism. They argue convincingly that the logic of bidding one's valuation in a second-price auction is not immediately apparent to bidders, and that bidders may tend to underbid their valuations in second-price auctions because that is what one does in order to earn a surplus in other, more familiar, auction types. In the long run, the second-price auctions generate only slightly less revenue than the English auctions, so the authors conclude that the English and second-price auctions are approximately strategically equivalent.

Cox, Roberson, and Smith (1982) provided more data on bidding in second-price auctions, and again found evidence of subjects' bidding lower than their valuations. However, this study was biased towards underbidding, because the computer software employed imposed a constraint that would not allow experimental subjects to bid greater than their private valuations in any of the experimental auctions.

Kagel, Harstad, and Levin (1987) performed a similar investigation of first-price, second-price, and English auctions in the context of affiliated private values (APV).[5] They did not constrain bidders to bid less than their valuations, and found in this case that second-price auction revenues were consistently *higher* than the theoretical prediction. In second-price auctions the tendency was for bidders to bid higher than their valuations, while in English auctions they did follow their dominant strategy. The English and second-price auctions were not paired for direct comparison using the same distribution of valuations. Nevertheless, there is clear indirect evidence that the strategic equivalence of English and second-price auctions is violated: people tend to bid higher in second-price auctions than they do in English auctions. Each of the four experimental sessions consisted of up to 31 sepa-

[5] As noted earlier, affiliation of private values does not change the theoretical strategic equivalence between English and second-price auctions.

rate rounds of auctions, so there some interesting data on changes in bidder behavior over time. In the English auctions, there was some tendency to overbid in the early rounds, with rapid convergence down to the dominant strategy of bidding no higher than one's valuation. In the second-price auctions, there was usually underbidding in the first round, with convergence by round 7 to overbidding that persisted for the remainder of the experimental session.

More evidence of overbidding in second-price auctions comes from a study of first-, second-, and third-price auctions by Kagel and Levin (1993). In two experimental sessions of 24 and 35 second-price auctions, they report 62 percent of all bids being above valuations, 30 percent approximately equal to valuations, and only 8 percent below valuations. Thus, although there is some underbidding observed, overbidding seems to be the most common and most persistent strategy in second-price auctions.

One puzzle is why different experimenters have obtained opposite results on second-price auctions. The first two experiments cited here (by Smith and his colleagues) showed a tendency for subjects to bid less than or equal to their valuations, while the other two experiments (by Kagel and colleagues) show a tendency to overbid their valuations. It is clear why the paper by Cox, Roberson, and Smith (1982) found no overbidding, as the experiment constrained bidders to bid less than or equal to their valuations. This leaves the puzzle of why Coppinger, Smith, Titus (1980) also failed to find a consistent tendency to overbid. This 1980 experiment did not constrain bids to be less than or equ... to valuations, but the experimenters did supplement their written instructions to the subjects with examples of how to compute their profits, including an example which demonstrated that negative profits could be earned by overbidding one's valuation v if one's highest rival bid were also greater than v.[6] By contrast, the experiments reported in Kagel, Harstad and Levin (1987) and Kagel and Levin (1993) did *not* include such examples for the bidders.[7] This, then, is a potential cause of the discrepancy. Kagel, Harstad, and Levin (1987) argue that the observed overbidding "is likely based on the illusion that it improves the probability of winning with no real cost to the bidder as the second-high-bid price is paid," and that people who behave this way receive very little negative reinforcement that might cause them to learn over time, as the overall probability of losing money in their experiment was only 6 percent.

---

[6] Private communication with Vernon Smith, March 1996.
[7] Private communication with John Kagel, March 1996.

In summary, then, the results of past experimental studies show that the strategic equivalence between English and second-price auctions for private-value auctions is, in general, not satisfied. Experimental subjects appear to bid higher in second-price auctions than they are willing to bid in English auctions. However, these results depend somewhat on the amount of experience of the subjects as well as the types of examples given to the subjects in the instructions for the experiment. In particular, although in repeated trials second-price auction revenues tend to be consistently above English auction revenues, the reverse is usually true during the first few trials of any given experiment. In the first round of an experimental session, experimenters have typically seen bidding above valuations in an English auction and bidding below valuations in a second-price auction.[8] After learning by bidders has taken place, however, the general observation is that people bid their dominant strate in English auctions but overbid in second-price auctions.

In this chapter, I will attempt to replicate these laboratory experiments in a real-world market environment. Because of the logistical difficulties of doing so, I will not be able to include many repeated trials, so learning effects may not be observed. But because my experiments take place in a preexisting market with already-experienced bidders, and because my bidders have plenty of time to make their decisions (days, as compared with minutes or seconds in laboratory experiments), I hope that any learning effects will already be accounted for even in a single-round auction experiment.

3.3 Revenue Equivalence Among All Four Formats

We have already seen that in IPV auctions, the Dutch and first-price auction mechanisms fail to be strategically equivalent, as do the English and second-price auctions, so there cannot be much hope for Vickrey's theoretical prediction that all four auction types should yield the same expected revenue. Nevertheless, it remains interesting to compare the first-price and second-price auction revenues, to see how close they come to revenue equivalence.

Coppinger, Smith, and Titus (1980) find that the first-price auction mechanism produces the highest revenues, followed by the Dutch auction, and then the English and second-price auctions. These same rankings were borne out by the experiments in Cox, Roberson, and Smith (1982) for the first-price, Dutch, and second-price auctions. In addition, these

[8] One caveat: this generalization comes from only a limited number of observations. The results of only seven experimental sessions (with a total of 42 subjects) of English auctions have been published, for example: five by Coppinger, Smith and Titus (1980) and two by Kagel, Harstad, and Levin (1987).

authors point to risk aversion as a possible explanation for the high revenues in the first-price auctions, which consistently yielded revenues above the RNNE prediction. This risk-aversion hypothesis has received a great deal of attention and detailed modeling in subsequent papers, such as Cox, Smith, and Walker (1988, 1992), Harrison (1989, 1992), Kagel and Roth (1992), and Kagel and Levin (1993); for an overview of this literature, see Section I.G. of the review article by Kagel (1995). According to Kagel, the bottom line is that risk aversion "probably has some role to play in explaining bidding in private value auctions," but that risk aversion doesn't entirely explain the deviations from the predicted RNNE equilibrium.

Kagel and Levin (1993) perform another direct comparison of the first-price and second-price auction mechanisms in the IPV case, and they duplicate the finding that first-price expected revenues are higher than second-price expected revenues in auctions with five participating bidders. Additional evidence from third-price auctions suggests that at least part of this difference is attributable to risk aversion. They also find that this difference disappears when the number of bidders is increased from five to ten.

In the affiliated-private-values (APV) case, Kagel, Harstad, and Levin (1987) find some violations of the theoretical RNNE prediction that the English and second-price auctions should raise more revenue than a first-price auction. In a number of cases, the first-price auction revenues were higher than the English auction revenues, and sometimes also higher than the second-price auction revenues. Just as in the IPV case, the first-price auction revenues were consistently higher than the theoretical RNNE predicted level.

To summarize, it appears that the following general statements are true. The results of previous IPV auction experiments show that the four auction types can probably be ranked from highest to lowest expected revenues as follows: first-price, Dutch, second-price, English. In addition to the failures of the strategic equivalences discussed in the previous two sections, the revenue equivalence between first-price and second-price auctions also fails, as first-price auctions earn higher revenues than do second-price auctions. This finding is probably attributable in part to risk aversion on the part of bidders, which is known theoretically to raise first-price but not second-price auction revenues in auctions. The result of higher revenues in first-price versus second-price auctions may disappear when the number of bidders becomes large, however, as well as in cases when bidders' private values are affiliated rather than independent. In the field experiments reported in this chapter, I can observe neither bidder risk preferences nor their distribution of bidder valuations (although, I will argue, bidder valuations are likely to be private rather than having a common-value component). Therefore, my results will focus mainly on the two predicted strategic equivalences discussed in the previous two sections, because those two theoretical predictions do not depend on risk preferences. Nevertheless, I will also present what evidence I can from this auction market on the expected revenues in first-price versus second-price auctions.

4 Experimental Procedure

To test the revenue equivalence theorem, I purchased over $2,000 worth of Magic cards and resold them via auctions in the Internet marketplace. The basic procedure was to auction two copies of the same card via two different auction mechanisms, in order to be able to make direct comparisons of revenue earned between the two different auction formats. The Internet marketplace for Magic cards consists of messages on the Usenet newsgroup <rec.games.trading-cards.marketplace>, supplemented with private electronic mail messages. For more details on the history and institutions of this marketplace, as well as a discussion of the demographics of the typical bidders, see Chapter 1.

The experiments reported in this chapter consist of eight different auctions: two each of the English, Dutch, first-price, and second-price formats. I grouped these eight auctions into four paired comparisons: two pairs comparing the Dutch and first-price formats, and two pairs comparing the English and second-price formats.

4.1 Simultaneous, Rather Than Sequential

Each of these eight auctions was actually a simultaneous auction of between 85 and 100 individual goods (each one a different type of Magic card) with separate bidding for each good. Historically, *simultaneous* auctions have been used frequently in sealed-bid auctions, such as those for Treasury bills or offshore oil drilling rights. By contrast, almost all English auctions (such as those at Christie's or Sotheby's auction houses) and Dutch auctions (such as the traditional tulip bulb auctions) have been *sequential* auctions. In a sequential auction, an audience of bidders gathers for a period of several hours, during which time one item after another comes up for bid. The bidding starts on the next item only after the auctioneer has declared the previous item sold.

Previous laboratory experiments on revenue equivalence in auctions have all been conducted with sequential auctions: subjects participate in repeated "rounds" of auctions for single objects. One advantage this offers is that it allows the experimenter to observe the effects of learning as bidders gain experience with a particular auction format. Why, then, have I chosen to conduct my experiments as simultaneous auctions? There are several reasons. First, simultaneous auctions are the norm in the Internet marketplace for Magic cards; practically all Magic auctions on the Internet are simultaneous auctions by a single auctioneer for dozens or even hundreds of different cards. Running simultaneous auctions, then, helped me make my auctions seem as natural and familiar as possible to the bidders who would be participating. Second, simultaneous auctions made it possible for me to collect data more rapidly, as it enabled me to collect data on over 80 auctions at a time. The bidding technology (electronic mail) used here allows people all over the world to participate in the auction, without their having to come to a single location as has previously been necessary for English and Dutch auctions; the tradeoff is that these auctions may take longer to complete than would an auction in an auction house. In order to allow time for bids to arrive via electronic mail from separate locations, it is not feasible to have bidders respond to each others' bids much more frequently than once per day. Thus, a single English auction may take as much as a month to complete, and this would make data collection via sequential auction very difficult.

In fact, one nice advantage of Internet technology is that it increases the feasibility of running simultaneous English and Dutch auctions. When there are multiple goods auctioned simultaneously in real time, it can be a challenge to organize for bidders the data on the current status of each good; networked computers provide an elegant solution to this problem.[9] In my auctions bidders received electronic updates with each of the current high bids laid out for them on their computer screens, and they had time to reflect on their bids before submitting them. A second advantage of electronic-mail bidding is that it relaxes the requirement that bidders be in the same place at the same time for real-time auctions. The availability of this type of bidding technology is a major reason why the FCC chose to run a simultaneous English auction format for its recent, well-publicized auctions of personal communications spectrum. The FCC auctions (which took place during the same period of time as the auctions reported in this chapter) were very similar in format to my English auctions, as they were simultaneous auctions of approximately 100 different goods (licenses) with periodic electronic updates of the high bids on each good. It is very likely that simultaneous auctions will continue to rise in importance in the future because of the increased availability of this technology.

[9] Other solutions are also possible. For example, charity auctions conducted via telethon have long used a simultaneous real-time English auction format.

4.2 Bidder Entry

Another way in which my experiments differ from laboratory tests of the revenue equivalence theorem has to do with the number of participating bidders. In such experiments, as in Vickrey's original theoretical model, the number N of bidders is fixed and certain. Each experimental subject knows, when deciding how to bid, exactly how many other bidders will be participating (typically a number between three and ten). By contrast, these auctions for Magic cards featured entry by an uncertain number of bidders. I advertised each auction to a large number of people, but I never knew for certain how many bidders would decide to enter, neither for the bidding on any particular card nor for the auction as a whole.

I used two different methods to invite bidders to participate: posting advertisements on the newsgroup <rec.games.trading-cards.marketplace>, and sending individual invitations via email to a private mailing list of people whom I knew were interested in Magic auctions. This is standard practice for auctioneers in this marketplace (sending a direct email invitation as opposed to posting a newsgroup advertisement significantly increases the probability that a given bidder will read your message), but it does create an asymmetry between the sealed-bid and the real-time auction formats. I designed my second-price sealed bid auctions to last one week each, but the English auctions had an unpredictable stopping date, as they did not stop until every card listed had been sold. In addition, I posted daily updates of the current high bids in the English auction to the newsgroup, whereas no such updates existed for the sealed-bid auctions. Posting my daily updates was done in order to keep with another standard practice among other auctioneers; I posted them to the newsgroup in addition to emailing them directly to anyone who had already submitted a bid or otherwise indicated that they would like to receive daily updates. (Updates were not sent to the entire original mailing list to which I announced the auction, but only to people demonstrating active interest in that particular English auction.) Another potentially relevant feature of the design was that I never told the bidders how many people had been invited to the auction; any time I sent a group mailing (either an announcement or an update), I sent the message so that each individual recipient would be unable to tell who else had received it.

The asymmetry mentioned above has to do with the amount of advertising each auction received on the newsgroup. The English (and, similarly, Dutch) auctions had updates posted to the newsgroup each day, and typically lasted several weeks. By contrast, the sealed-bid auctions lasted only one week, and I advertised them each only three times on the newsgroup (which was the maximum number of times I felt I could post the exact same message without suffering recriminations for cluttering up the newsgroup with "wasted bandwidth"). Thus one might expect the number of bidders attracted by advertisements to be higher in the auctions with daily updates than in the sealed-bid auctions, and since revenue in an auction is typically increasing in the number of participating bidders, we might expect there to be an upward bias to the revenue in the English and Dutch auctions. In order to counterbalance this asymmetry to some extent, I made the decision to send out some "extra" direct email invitations for the second auction in a pair. For example, my first auction was an English auction for which I sent out 90 direct email invitations in addition to the newsgroup postings, and it attracted a total of 38 bidders: 18 from the email and 20 from the newsgroup. After this auction ended, I ran a paired second-price auction for which I sent direct email invitations to the 20 "new" bidders as well as to the original 90-person mailing list (minus a few people who specifically asked to be removed from the mailing list).

Thus, the sealed-bid auctions did receive slightly less advertising than their counterpart daily-update auctions, which could potentially have biased the revenue results to favor the real-time auctions. However, an attempt was made to even out the bidder pool by making sure anyone who had entered one auction via a newsgroup advertisement was also invited to the next auction via a direct email invitation. The advantages of using the newsgroup advertisements were that they provided a source of new subjects for the research program, and they made the auctions as similar as possible to other auctioneer's Magic auctions already taking place on the Internet. The disadvantage was some loss of control over the size of the bidding pool. In my analysis of the revenue results below, I will consider the effects of excluding those observations which may have been biased by such advertising effects.

These complications arising due to the uncertain entry of bidders in my auctions illustrate one of the tradeoffs faced in the use of field experiments. In laboratory experiments designed to test the revenue equivalence theorem, there are no such difficulties: one just brings in a fixed number N of subjects, tells them the probability distribution of bidders' valuations, and instructs all N of them to bid in each auction. By contrast, in a field experiment one gives up some of the controls (such as the number of participating bidders and the true distribution of valuations) in exchange for increased realism. It is usually the case in real-world sealed-bid auctions, for example, that the bidders do not know precisely how many other bidders there are at the time they choose to make their bids. If field results on revenue equivalence turn out to be different from laboratory results, then it may indicate that the simple theory may not have been a terribly useful approximation, thus yielding fruitful directions for future laboratory experiments and theoretical research aimed at bringing economic models more in line with economic reality.

4.3 Time-Ordering Effects

One concern I had in designing the paired experiments was that there might be time-ordering effects on revenue, which could confound the revenue comparisons I wished to make between the different auction formats. For example, suppose I ran a first-price auction and sold a Akron Legionnaire card to Bidder X for a very rather high price of, say, $11.00. Then, when I followed this auction up with a matching Dutch auction, I might find that since Bidder X's demand has already been satisfied, the best price I can get from the remaining bidders is only $8.00. This demand-saturation effect is related to, but distinct from, the "revenue decline anomaly" described by Ashenfelter (1989). Ashenfelter describes situations in which multiple identical units of an item (such as cases of wine) are to be offered sequentially for auction, and despite the fact that the number of units to be sold is known to bidders in advance, the later units sold tend to fetch less revenue than the first units sold. In my auctions, such declines in revenues would not necessarily be anomalous; the difference here is that bidders did not know for sure that I would ever offer a second copy of any card I auctioned. In fact, I actively tried to discourage such beliefs in two ways: by including singleton cards along with the paired cards so that two paired auctions were not necessarily identical, and by occasionally running other auctions consisting entirely of singleton cards (including some auctions designed for other experiments not included in this chapter). In any case, whether or not the bidders *expected* to see the same card up for auction twice, there is some reason to believe that the second auction would fetch less revenue than the first.

Another potential time-order effect in my auction pairs has to do with the way bidders were invited to participate via direct email solicitations as well as by newsgroup advertisements, as noted in the previous section. I typically sent email invitations to more bidders for the second auction in a pair than I did for the first. For example, if a matched pair consisted of auction A followed by auction B, my email invitation list for auction B consisted of list for auction A plus any bidders who participated in auction A after discovering the advertisement on the newsgroup. Therefore, the number of participants in the second auction might be expected to be higher than the number of participants in the first auction, which would cause the second auction to tend to have higher revenue.

In order to control for these potential time-order effects, I ran each paired experiment twice, reversing the order of the treatment effect between the two pairs. This was the reason for running eight auctions rather than just four. For example, after running first an English auction and then a second-price auction for one set of cards, I also ran a second-price auction followed by an English auction for another set of cards. This would allow me to ascertain whether differences in revenue, if any appeared, were really due to the auction mechanism rather than merely the order in which the auctions appeared.

4.4 Dutch and First-Price Auctions

Table 1 gives an overview of the Dutch and first-price auction experiments I conducted. The first pair (experiment FD) consisted of Auction FD1, a first-price auction, followed by Auction FD2, a Dutch auction. Auction FD1 offered 88 different cards and lasted one week. Auction FD2 offered 87 cards (I was unable to locate a second "Mana Drain" card), each of which matched one of the cards in Auction FD1, and it lasted until the last card was sold, twelve days after it started.

The first-price auction was fairly straightforward to execute: potential bidders were told that they had one week's time to submit their bids, at which point I would award each card to the highest bidder at the amount of their bid. In order to maintain as much comparability as possible to the Dutch auctions, I imposed some constraints on the acceptable bid amounts.[10] For example, any bid less than $1.00 was required to be an even multiple of a nickel ($0.05), and any bid greater than $20.00 was required to be an even multiple of a dollar ($1.00). These corresponded to the bid decrement amounts that I planned to use in the corresponding Dutch auction. Appendix 1 shows a sample announcement for one of my first-price auctions.

At the end of the week, I computed the results of the first-price auction and emailed the winning prices to each of the participating bidders. The mailing of the results did not identify any of the winners. Anyone who won cards also received a congratulatory message listing the cards they had won, along with their total bill.[11] After receiving their payment in the mail, I sent them the bundle of cards they had won.

For the Dutch auction, there were additional design complications. Not only did I have to decide on the Dutch clock decrement amounts (how much the prices would fall each day

[10] The acceptable bid amounts were spelled out in detail in the rules of the first-price auction (see Appendix 1). If a bidder mistakenly submitted an unacceptable bid amount (such as $0.98), I rounded it down to the nearest acceptable amount (such as $0.95).

Table 1: Overview of Dutch-First Experiments

	Auction FD1	Auction FD2	Auction DF1	Auction DF2
Auction format	First-price	Dutch	Dutch	First-price
Card type	Black/Blue	Black/Blue	Red/Green	Red/Green
Start date	Fri, 5 May	Wed, 17 May	Wed, 27 Sep	Tue, 31 Oct
End date	Fri, 12 May	Mon, 29 May	Tue, 24 Oct	Tue, 7 Nov
Total number of cards auctioned	88	87	88	88
Total Cloister value	550.42	533.67	308.24	308.24
Number of matched cards	87	87	86	86
Maximum selling price	27.00	26.00	22.00	25.00
Minimum selling price	0.10	0.15	0.20	0.10
Cloister value of matched cards	533.67	533.67	303.25	303.25
Total revenue on matched cards	431.25	446.35	348.45	327.05
Number of participating bidders	32	63	88	42
from newsgroup announcements	3	7	7	3
from email invitations	29	56	81	39
Number of email invitations sent	403	379	586	472
Number of winners	20	22	22	22
Maximum number of cards to a winner	18	15	19	16
Maximum payment by a winner	92.75	80.50	57.25	92.75

a good went unsold), I also had to determine the starting prices. If I set the starting prices too low, I would put a cap on the amount of revenue I could earn, thus introducing a potential downward bias on the Dutch auction revenue. On the other hand, if I set them too high, the auction would take an unreasonably long time to clear, and bidders might lose interest too early and/or become annoyed with me for wasting their time. As a compromise, I decided to take the winning bid amounts from the first-price Auction FD1, and increase them

[11] Although auctioneers on the Internet typically ask winning bidders to pay shipping costs in addition to their winning bid price, I decided not to follow this practice, as I wanted to limit the extent to which bidders tried to win multiple cards in order to reduce the shipping cost per card. In order to encourage bidders to bid independently on each card, I told them from the beginning that their bid price would include free shipping within the United States.)

by five days' worth of Dutch clock decrements to get the starting prices for Auction FD2, the Dutch auction. In other words, I set the starting prices so that if each card sold after exactly five days, the Dutch auction revenue would be exactly the same as it was in the first-price auction.

The decrement amounts for the Dutch auction were on the order of a five percent decrease per day, with the exact amount depending on the current price level. I did not tell bidders explicitly how much prices would fall each day in the Dutch auction; rather, I told them that prices would fall by a "small amount" each day. My reason for this was to help enforce the real-time nature of the Dutch auction; I didn't want people trying to submit a future-price bid such as, "I would like to bid $4.30 for this card when its price has fallen to this level in two days' time, assuming no one else has bought it by then." Not giving explicit decrement amounts was intended to force people to look at the current price before submitting a bid.

The initial announcement of the Dutch auction stated the style and rules of the auction, along with the list of cards and their starting prices. Bidders were told that if they wanted to receive daily updates via email, they had to send me a message to that effect. I mailed updates once per day to this list of interested bidders, as well as posting them to the appropriate newsgroup, indicating which cards had been sold and to whom (so that winning bidders would know for sure that their bids had been accepted). For an example of such an update message, including the Dutch auction rules as stated to the bidders, see Appendix 2.

To my surprise, 25 of the 87 cards in Auction FD2 sold on the very first day of the Dutch auction, at prices which were all at least fifteen percent higher than their corresponding prices in the first-price auction. This indicated that some of the Dutch auction revenues could potentially have been higher if the prices had begun at a higher amount, so there was effectively data censoring from above.

The second pair of auctions described in Table 1 are labeled Auction DF1 and Auction DF2, which were, respectively, a Dutch auction followed by a first-price auction for the same cards. This time it was even more difficult to decide on the appropriate starting prices for the Dutch auction, because I did not have any first-price auction results to use as a baseline. Instead, I resorted to a standard price list for Magic cards called the Cloister price list, which gives a kind of average price for each of the Magic cards traded on the Internet marketplace, computed weekly.[12] I began with these Cloister values as a baseline, and marked

---

[12] For a more detailed description of the Cloister price list, see Chapter 1. Technical details on the computation of the prices may be found in Black (1995).

up these prices by five days' worth of Dutch clock decrements to obtain the starting price. This proved to be a reasonable technique; fortunately, none of the cards sold on the first day of the auction, although several sold on the second day. All but one card sold within about two weeks of the auction, but the final card remained unsold, all by itself, for an extra eleven days. Several bidders became impatient, and asked to have their names removed from the update list, before this card was sold. One week after this Dutch auction (Auction DF1) finished, I started a first price auction (Auction DF2) for comparison, using the same set of 88 cards. (However, two of these cards were only "near mint" instead of the "mint" condition of their counterparts in Auction DF1, and thus only 86 of the cards are strictly comparable.)

Table 1 displays an overview of each of the four auctions in the experiments designed to test the Dutch-first strategic equivalence. It includes data on the total number of cards in each auction, along with their total Cloister value. Next there are a set of statistics on the matched cards, including the maximum and minimum selling price for cards within each auction, the total auction revenue, and the total Cloister value of the matched cards. Other statistics include the number of participating bidders (equal to the number of bidders in the first-price auctions, and to the number of people who asked to receive daily updates for the Dutch auctions) as well as the number of bidders who actually won cards. It is worth noting that the number of participants was somewhat higher in the second pair of auctions (DF1 and DF2) than in the first pair of auctions (FD1 and FD2), mainly due to a difference in size of my announcement mailing list between those two time periods.

4.5 English and Second-Price Auctions

Four other auctions in this study were designed for an experimental test of the strategic equivalence between English and second-price auctions. An overview of these auctions may be found in Table 2. As in the Dutch-first experiments, two pairs of auctions were involved, in order to control for time-order effects. Auctions ES1 and ES2 were an English auction followed by a second-price auction for one set of cards, while Auctions SE1 and SE2 were a second-price auction followed by an English auction for another set of cards.

The second-price auctions were run according to the same basic set of sealed-bid rules as the first-price auctions discussed above. As in the first price-auctions, the second-price auctions each lasted one week. The only difference was that I took care to make clear to

Table 2: Overview of English-Second Experiments

	Auction ES1	Auction ES2	Auction SE1	Auction SE2
Auction format	English	Second-price	Second-price	English
Card type	Legends	Legends	White/Gold	White/Gold
Start date	Thu, Feb 9	Thu, Mar 9	Sat, 6 May	Sat, 20 May
End date	Sat, Mar 4	Thu, Mar 16	Sat, 13 May	Wed, 7 Jun
Number of cards auctioned	85	85	99	99
Total Cloister value	214.64	210.14	828.63	828.63
Number of matched cards	66	66	98	98
Maximum selling price	12.00	10.00	23.00	21.00
Minimum selling price	0.10	0.25	0.05	0.05
Cloister value of matched cards	107.94	107.94	804.33	804.33
Total revenue on matched cards	79.50	85.50	517.05	600.40
Number of participating bidders	40	27	43	38
from newsgroup announcements	22	13	3	12
from email invitations	18	14	43	26
Number of email invitations sent	90	85	385	372
Number of winners	26	15	27	17
Maximum number of cards to a winner	11	39	15	15
Maximum payment by a winner	26.50	69.20	154.00	136.50

bidders that the winning bidder would pay the *second-highest* amount bid, not the amount of his or her own bid. For a sample second-price auction announcement, see Appendix 3.

The English auctions required some additional design decisions, although fewer than in the Dutch auction case. I decided to issue daily updates, as in the Dutch auction, this time with the current high bid on each card displayed, so that bidders would have an opportunity to raise the current high bid. Bidders received these updates only after submitting a bid or otherwise indicating that they wished to receive daily updates; other people only received the initial announcement mailing. I employed a "Going, Going, Gone!" technique for declaring a card sold: a card which lasted a full day without a bid raise received an exclamation mark as a warning on that day's update, a card which lasted two full days received two exclamation marks as a second warning, and a card which lasted three full days was declared "SOLD!" on the update mailing. For a sample English auction update, see Appendix 4.

I also put restrictions on the levels of acceptable bids, as experience had taught me that when bid increment amounts were allowed to be very small (say $0.01), some bidders could become annoyed that the auction was being unnecessarily prolonged by other bidders making insignificant bid raises. Thus in Auction ES1, all bids were constrained to be even multiples of a nickel, effectively requiring all bid raises to be in increments of at least five cents. In Auction SE2 (after more feedback from bidders), I put even larger constraints on minimum bid increments for the higher-priced cards, such as "all bids between $10 and $20 must be in even multiples of 50 cents." (In fact, the allowed bid levels for Auction SE2 were the same as those in the first-price auction rules shown in Appendix 1.) To facilitate comparisons, I imposed the same bid level constraints on the second-price Auction ES2 as on the English Auction ES1, and the same constraints on the second-price Auction SE1 as on the English Auction SE2.

The summary statistics in Table 2 for the English and second-price auctions are exactly analogous to those in Table 1 for the Dutch and first-price auctions. One major difference between the two pairs of auctions in Table 2 is that the number of directly-invited participants was much lower in this first pair (Auctions ES1 and ES2) than in the second (Auctions SE1 and SE2). The increased number of invitations resulted in an increased number of bidders in the second-price auction, but not in the English auction. Advertising in the form of daily updates posted to the newsgroup evidently had a positive affect on participation in the English auctions. Another interesting point is that the time-order effect on revenues is the opposite of that expected, as the revenue on matched cards is greatest for the second auction in each pair; clearly, demand-saturation effects are not the dominant factor in this case. It may be that this revenue phenomenon is caused by the asymmetric advertising effects discussed earlier, and thus I will make an attempt to control for such effects in my analysis below. Finally, experiment ES contained a rather large number of unmatched cards, including 13 unmatched card types and 6 unmatched card conditions out of the 85 cards in each auction. However, the unmatched cards were chosen to be comparable in value across auctions, to keep the auctions as similar as possible overall. All subsequent data analysis in this chapter will of course exclude observations that were not perfect matches.

4.6 Design Summary

To summarize, the experimental design consisted of four auctions designed to test revenue equivalence between the Dutch and first-price auction formats, plus four additional auctions designed to test revenue equivalence between the English and second-price formats. Each auction consisted of dozens of individual card auctions, including many matched pairs of cards designed to facilitate direct revenue comparisons. Efforts were made to hold constant as much as possible about the two auctions in each pair, such as the constraints on allowable bid amounts. However, the real-world nature of these field experiments caused some variables to be beyond the experimenter's control. This is the result of a tradeoff which sacrifices some of the experimental control of a laboratory experiment for the increase in realism of a field experiment. In order to control indirectly for the uncontrolled environmental variables, each experiment was repeated with the time order of the experimental treatments reversed.

The paired auctions were designed to produce direct comparisons between Dutch and first-price auction revenue, as well as between English and second-price auction revenues. In addition, further comparisons among all four auction types may be made possible by the use of Cloister prices, which are independent measures of the market value of each card.

5 Results

5.1 Dutch and First-Price Auctions

For convenience in referring to the experimental data, I shall introduce the following notation. Let FD ("first-price, then Dutch") refer to the paired experiment described in the first two columns of Table 1, where the pair consisted of a first-price auction followed by a Dutch auction. Similarly, let DF ("Dutch, then first-price") refer to the paired experiment described in the last two columns of Table 1 on page 101. Figures 1 and 2 show graphs of the distribution of the revenue outcomes for the 87 matched pairs of cards in experiment FD, and the 96 matched pairs of cards in experiment DF. Figure 1 displays the revenue difference, in dollars, between the Dutch revenue and the first-price revenue In order to view the revenue differences as percentages rather than as absolute differences, Figure 2 takes the same difference data and normalizes it by the Cloister value of each card. Clearly, there is a general tendency for the difference to be positive: Dutch auctions appear to raise more revenue than first-price auctions. This is a surprising result, for it conflicts with the

Figure 1: Dutch revenue minus first-price revenue

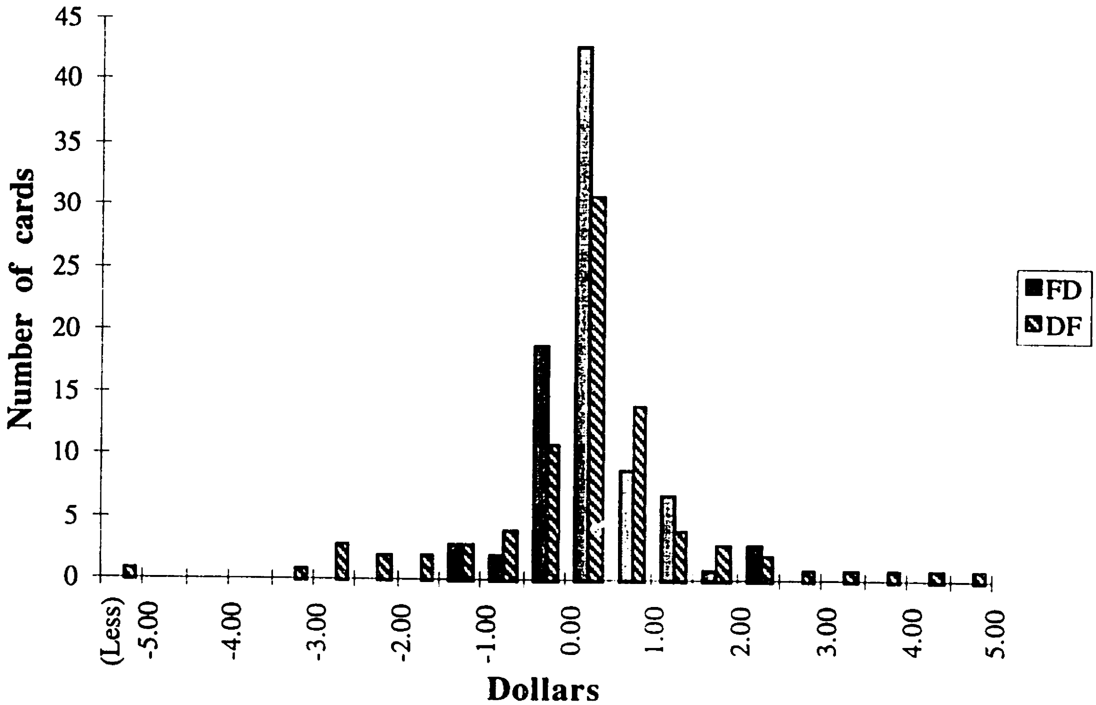

Figure 2: Dutch revenue minus first-price revenue

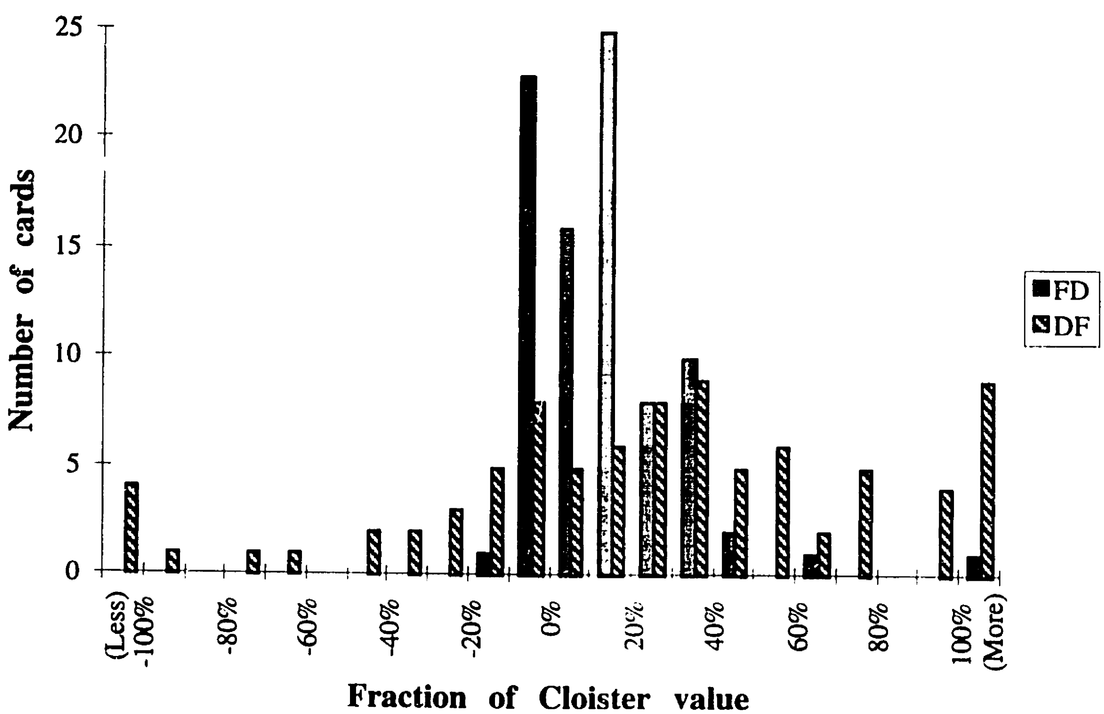

results from similar laboratory experiments, where there has been a systematic tendency for Dutch auctions to raise *less* revenue than first-price auctions.

For a statistical analysis of the data, see Table 3. The table is divided into two sections.

Table 3: The difference between Dutch and first-price revenues

	Experiment FD	Experiment DF	Combined totals
All observations	87	86	173
Dutch revenue higher	63	59	122
Equal revenue	12	5	17
First-price revenue higher	12	22	34
Mean Dutch-first difference (Std. deviation of mean)	0.38 (0.07)	0.25 (0.17)	0.32 (0.09)
Mean percentage difference (Std. deviation of mean)	14.9% (2.4%)	32.8% (13.7%)	23.8% (7.0%)
Wilcoxon statistic	5.66	2.78	5.78
Restricted sample	82	86	168
Dutch revenue higher	62	59	121
Equal revenue	8	5	13
First-price revenue higher	12	22	34
Mean Dutch-first difference (Std. deviation of mean)	0.42 (0.09)	0.25 (0.17)	0.33 (0.09)
Mean percentage difference (Std. deviation of mean)	-15.6% (0.3%)	32.8% (13.7%)	24.4% (0.6%)
Wilcoxon statistic	6.15	2.78	5.96

The top section displays the results for all observations of the matched pairs of cards; a matched pair here refers to a match on the card name as well as its advertised condition (mint, near mint, etc.). Of 87 such matched pairs in Experiment FD, the vast majority (63) were observations in which the revenue in the Dutch auction was greater than the revenue in the first-price auction. Similarly, 59 of 86 observations in Experiment DF had greater revenues in the Dutch auction.

Next is reported the average difference between the revenues in the Dutch and first-price auctions, which is the variable plotted in Figure 1. The mean difference was 38 cents in experiment FD, and 25 cents in experiment DF, favoring the Dutch auction in both experiments. A two-tailed t-test at the 95% confidence level rejects the hypothesis of a zero mean difference for experiment FD, but not for experiment DF. Pooling the observations from both experiments, one can also reject the hypothesis of revenue equivalence using a t-test, in favor of the hypothesis that the Dutch auction earns more revenue. The table also reports the mean *percentage* difference in revenue, with each difference normalized by the card's Cloister value, as in Figure 2. A two-tailed t-test on this variable results in a rejection of the revenue-equivalence hypothesis in each of the individual experiments.

With such controlled experimental data, it is also possible to perform an interesting nonparametric test of revenue equivalence. Rather than testing the mean difference, the Wilcoxon signed-rank test is a formal test of the hypothesis that the probability of obtaining a positive difference is equal to the probability of obtaining a negative difference.[13] The test statistic involves ranking the differences in revenue (Dutch revenue minus first-price revenue) according to their absolute value, and then computing the sum of the ranks of only those observations whose differences are positive.[14] This sum is a statistic which is distributed normally under the null hypothesis of equivalence, with known mean and variance that depend only upon the number of observations. Thus, by subtracting the mean and dividing by the standard deviation, we can obtain a test statistic which is distributed as a unit normal under the null hypothesis. I report this standard normal statistic as the "Wilcoxon statistic" in the table. For both of the individual experiments, as well as for the combined pool of data, the statistic is large enough to reject the null hypothesis at conventional significance levels. This indicates that the Dutch auction revenues have a significant tendency to be higher than the first-price auction revenues.

The bottom section of Table 3 contains results for a restricted sample of observations which attempts to exclude the potentially problematic advertising effects discussed in Section 4.2. (It includes only those observations for which the winning bidder in each auc-

---

[13] This test is nonparametric in the sense that it does not require the distributional assumptions of a t-test. Rather than assuming identically and independently distributed revenue differences, all one has to assume for the Wilcoxon test is that each observation has some constant probability p of yielding a positive rather than a negative revenue difference.

[14] Ties, meaning differences equal to zero, may be counted either as positive or negative, although this matters little in a situation such as this one where there are relatively few tie values. Just to be conservative, in computing my statistics I counted all ties as if they were negative differences, thus biasing the test against finding that Dutch revenues were higher than first-price revenues.

tion had been on the email invitation lists for both auctions, thus ruling out any cards won by people who were attracted to the auctions by newsgroup advertisements.) This turned out to exclude only five observations in this case, and only from the FD experiment. As can be seen in the table, the effect of this restriction on the sample was only to increase the statistical significance of the Dutch auction format's superior revenues.[15]

This finding of higher revenues in the Dutch auction format is unpredicted by the standard theory. In order to get a hint as to why this occurs, it is worth investigating whether the amount of the revenue violation varies with different kinds of goods. There were two card characteristics which were reasonable candidates as explanatory variables: the Cloister value of each card (CLOISTER), and the number of bidders who expressed specific interest in a particular card, as measured by the number of bids received for that card in the first-price auction (NUMBIDS).[16] The following equation reports the results of a regression of the amount of difference between Dutch and first-price revenues on a given card (DIFF) on three explanatory variables: CLOISTER and NUMBIDS, as well as a dummy variable DF which distinguishes Experiment DF from Experiment FD.

$$\begin{array}{lllll} \text{DIFF} = & 0.4536 & 0.1937\ \text{DF} & +0.0275\ \text{CLOISTER} & -0.0366\ \text{NUMBIDS} \\ & (0.1706) & (0.2770) & (0.0257) & (0.0282) \end{array}$$

None of the variables were statistically significant in this regression, nor were they significant in other specifications, such as those excluding one or more of the regressors, or those where the dependent variable was a percentage difference in revenue. Unfortunately, the available characteristics of the cards do not yield any insights as to the determinants of the observed tendency for Dutch revenues to be higher.

What are the possible explanations, then, for this violation of revenue equivalence between Dutch and first-price auctions, especially since this violation is exactly opposite the violation observed in laboratory experiments? One difference between these field experiments and the laboratory experiments is that the number of participating bidders has been exogenously fixed in laboratory experiments, but not in these field experiments: bidders in

---

[15] I also tried ranking normalized revenue differences, as in Figure 2 (dividing the difference by the Cloister value of each card), instead of dollar-valued revenue differences, in order to give more importance to observations whose revenue differences were large in percentage terms. However, this also resulted only in increased statistical significance of the Wilcoxon statistic.

[16] The amount of bidder interest on a card-by-card basis cannot be so easily observed in the Dutch auction, because in the Dutch auction only the highest bid(s) can be observed.

the Magic card auctions decided whether or not to participate after having seen the announcements of the auctions. As can be seen in Table 1, the number of participating bidders was higher in the Dutch auctions than in their corresponding first-price auctions (63 versus 32 bidders, and 88 versus 42 bidders, respectively). These increased numbers of bidders may have been the reason for the increased revenues in the Dutch auctions, assuming that these additional bidders might occasionally bid fairly high.[17] Thus, one plausible explanation for the revenue difference is that the real-time nature of the Dutch auction attracted more bidder attention than did the first-price auction, and the additional bidders yielded additional revenue. Under this explanation, there need not be any violations of the strategic equivalence of Dutch and first-price auctions in order to produce the violation of revenue equivalence. Nevertheless, it is still interesting to ask whether bidders are observed to bid higher on the same good under one auction format versus the other.

In addition to the card-level revenue data, I also observe individual bid-level data which may yield additional insight on this question. Unfortunately, the Dutch auction does not generate much data of this sort: only the winning bids (plus equally high bids submitted by losing bidders who lost because their bids weren't first) are actually observed --- by contrast, in a first-price auction the auctioneer observes bids submitted by each bidder on all cards which interest him or her. Despite the relative dearth of Dutch auction bid data, a few statistics could be generated, as follows.

There are two different ways a bidding anomaly (that is, a violation of strategic equivalence) could be observed. First, one could observe someone passing up the opportunity to bid in a Dutch auction at a price less than or equal to the amount he was willing to bid in a first-price auction. Such an anomaly violates strategic equivalence in favor of the first-price auction. Second, one could observe someone bidding in the Dutch auction at a price greater than the amount they bid in the matching first-price auction, which would be an anomaly in favor of the Dutch auction. Even if such anomalies exist, they will be difficult to observe, because of data censoring in the Dutch auction: the Dutch price level never falls any lower than the amount bid by the highest bidder, so one can learn very little about bidders with low valuations.

---

[17] On the other hand, it is worth noting that "participation" in the Dutch auction may have required somewhat less effort than participation in the first-price auction: to be counted as a "participant" in the Dutch auction, one did not necessarily have to submit any bids, but merely had to ask to receive updates of the current Dutch prices. It may be possible, then, that this measure of "participation" is somehow not measuring the same degree of participation as the number of bidders in the first-price auction.

There were a total of 257 bids observed in the two Dutch auctions. Of these, 38 bids (24 in FD2 and 14 in DF1), from 10 different bidders, matched first-price bids made by the same bidder on the same card in the paired first-price auction. These 38 observations represent opportunity to look for anomalies favoring the Dutch auction. Of these 38 observations, a full 30 of them did actually favor the Dutch auction, while 4 were equal across auctions and 4 favored the first-price auction. The 30 bid differences favoring the Dutch auction ranged in size from \$0.25 to \$16.00, with a mean of \$2.52, while the 4 differences favoring the first-price auction ranged only from \$0.10 to \$0.50. If there were no systematic tendency for a bidder to favor one auction format over another, then we would expect the number of differences in each direction to be approximately the same. The fact that the vast majority of differences favored the Dutch auction indicates that bidders may actually have a systematic tendency to bid higher in Dutch auctions.[18]

To look for bid anomalies favoring the first-price auction, I focus attention on the pool of all bidders who participated in both auctions in a pair (that is, who submitted bids in a first-price auction and also asked to receive updates in the matching Dutch auction). Such bidders submitted a total of 351 first-price auction bids in Experiment FD and 538 first-price auction bids in Experiment DF. I compare these first-price auction bids to the corresponding Dutch auction winning bids, to see whether a bidder ever passed up the opportunity to bid in a Dutch auction at a price less than or equal to the amount he was willing to bid in a first-price auction. Such an observation would be a violation of strategic equivalence in the sense that the bidder was willing to bid more in the first-price than in the Dutch auction. A search of the bid data indicates that such violations are rare. Of the 889 selected first-price bids, there were only 24 that were higher than, and 11 that were equal to, the corresponding winning bids in the Dutch auction. These were all the cases in which a bidder passed up the opportunity to buy a card in a Dutch auction at the price they demonstrated they were willing to bid in the first-price auction. The 24 violations favoring the first-price auction ranged in size from \$0.10 to \$3.00, with a mean of \$0.71. The fact that only 24 violations resulted from 889 possibilities appears to indicate that bid violations favoring the first-price auction format are very infrequent, but this is somewhat misleading. Even if all

[18] One caveat: Suppose bidders make random bidding errors that are equally likely to favor the Dutch auction as the first-price auction. Then, because individual Dutch bids are only observed when they are at the winning bid level, one would be much more likely to observe Dutch bids which have randomly been made too high than those which have randomly been made too low. In such a model of bidding errors, the above sample of Dutch bids is biased towards finding bid differences that favor the Dutch auction.

bidders had a tendency to bid higher in a first-price than a Dutch auction, it would have been impossible to detect this for any low-valuation bidders, because the Dutch price never fell far enough to be near the first-price bids of low-valuation bidders.

Unfortunately, the incomplete data-generating process of the Dutch auction makes it impossible to do a comprehensive comparison of the bid strategies by each bidder across the two different auction formats. All I can say is that I observed 30 cases of violations of strategic equivalence favoring the Dutch auction, and 24 cases favoring the first-price auction format. Nine different bidders generated the 30 bid violations favoring the Dutch auction, while eleven different bidders generated the 24 bid violations favoring the first-price auction, and two of these bidders overlapped between the two groups. Interestingly, the magnitude of the violations favoring the Dutch format was considerably greater on average than that of the violations favoring the first-price format. This suggests that some of the increased revenues observed in the Dutch auctions may be due to violations of strategic equivalence in individual bidding behavior, but the evidence is not comprehensive.

To summarize, the experimental data reveal a significant violation of the theoretical revenue equivalence between first-price and Dutch auctions. Even more interesting is the fact that the violation discovered is the exact opposite of that found in previous laboratory studies: I find that the revenue raised in Dutch auctions is higher than the revenue raised in first-price auctions, by an average of 30 cents, or approximately 25% of card value on average. At least part of the explanation is likely to be that in my auctions, bidder entry was endogenously determined, and that the Dutch auction format appears to have generated more participation among bidders. The individual bid data suggest, albeit inconclusively because the Dutch bid data are so sparse, that the revenue difference may also be due in part to individual bidders violating the predicted strategic equivalence between Dutch and first-price auctions. If bidders do indeed tend to bid higher in Dutch than in first-price auctions, one would need to explain why this is the opposite of the effect found in laboratory experiments. Perhaps it is the case that when bidders are bidding for real goods whose values are more difficult to articulate than the explicit cash payoffs used in the laboratory, they are more willing to accept high prices suggested to them, as in the Dutch auction, than to volunteer a similarly high price in a first-price auction bid. However, the "real goods" aspect of these auctions are not the only difference between these field experiments and the traditional laboratory experiments. Other differences include the simultaneous rather than sequential nature of the auctions, the fact that the number of other bidders was uncertain to each bidder, and the fact that the Dutch clock speed was measured in days rather than in seconds. All of these differences should be considered to be potential determinants of the different results, but the fact that bidder entry was greater in the Dutch auctions appears at the moment to be the most obvious cause of the difference.

5.2 English and Second-Price Auctions

A graphical overview of the English and second-price auction experiments can be found in Figures 3 and 4. Here I adopt notation analogous to the notation of the previous subsec-

Figure 3: English revenue minus second-price revenue

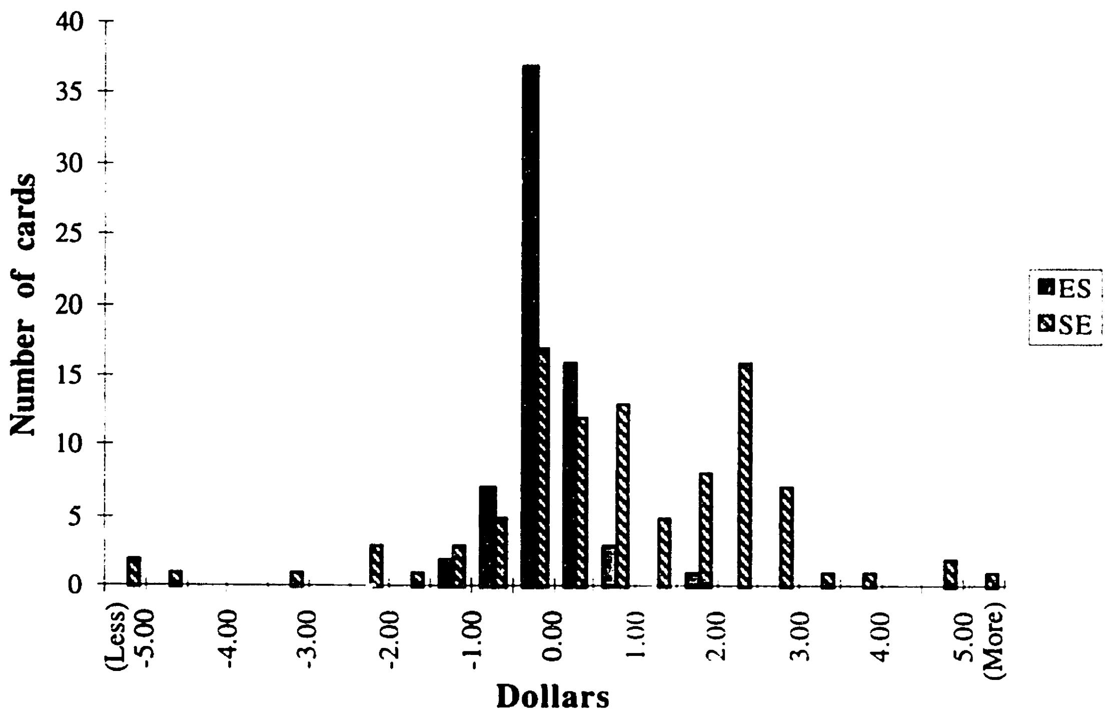

tion: Experiment ES consists of an English auction followed by a second-price auction (Auctions ES1 and ES2 in Table 2), while Experiment SE is the same type of experiment with the order reversed (Auctions SE1 and SE2 in Table 2). There are a total of 66 matched pairs in the ES experiment,[19] and a total of 98 in the SE experiment. Figure 3 shows the

[19] Experiment ES consisted of two of the earliest auctions in my research program; the reason for the relatively low number of observations is that I included a number of unmatched cards in this auctions in an effort to ensure that subjects would not come to expect every card I auctioned to have a matched pair in a later auction. Later, I realized I could achieve this same end by running occasional "decoy" auctions with no matches whatsoever, to offset the existence of pairs of auctions in which every card matched.

Figure 4: English-second revenue difference, normalized by Cloister value

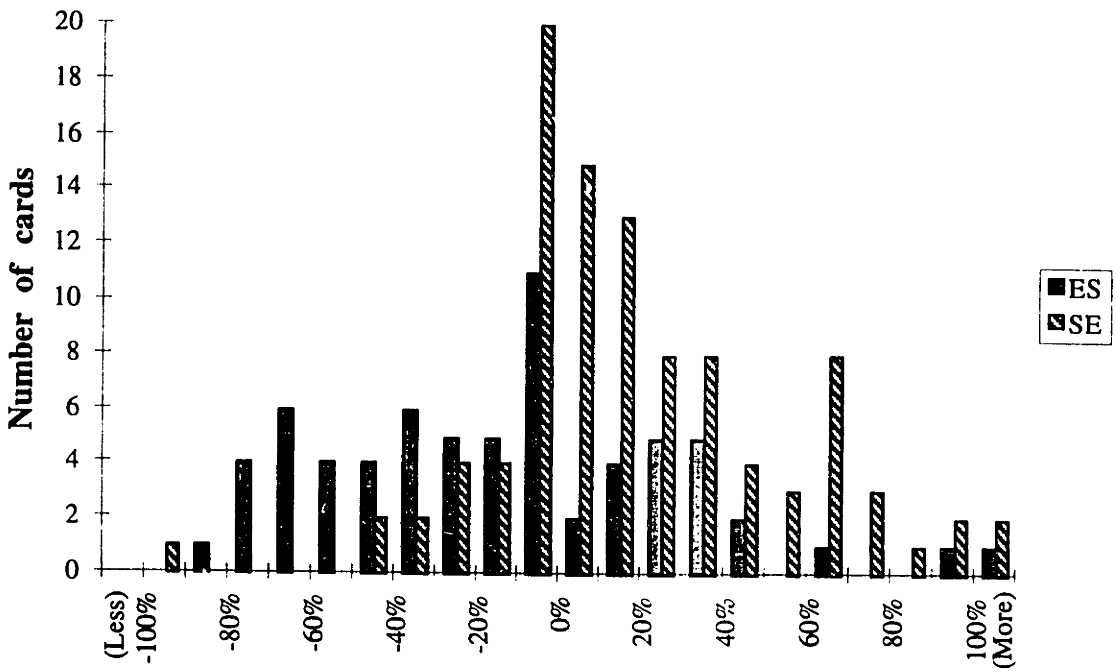

distribution of the revenue differences between the English and the second-price auctions for each card. It shows that there is a tendency for this difference to be positive in experiment SE, and a slight tendency for it to be negative in experiment ES. In other words, there is a time-order effect which appears to be greater than any effects due to the auction format treatment variable: whichever auction took place first in the pair tended to bring in less revenue than the one which took place second. Figure 4 takes a look at the normalized difference in revenue, in order to take into account the value of the card, but the results are qualitatively the same.

Table 4 displays numerical statistics for the two experiments. Again, the upper half of the table shows results for the full set of matched pairs of cards. A majority of observations (38 of 66) had higher English revenue in experiment ES, but only a small minority (24 of 98) had this feature in experiment SE. The Wilcoxon signed-rank test produced opposite results in the two different experiments: the test statistic was positive and significant for experiment SE, but negative and significant for experiment ES. This confirms my earlier suspicion that time-order effects appear to outweigh the intended experimental treatment effects.

Table 4: The difference between English and second-price revenues

	Experiment ES	Experiment SE
All observations	66	98
English revenue higher	20	65
Equal revenue	8	9
Second-price revenue higher	38	24
Mean Dutch-first difference (Std. deviation of mean)	-0.09 (0.06)	0.85 (0.19)
Mean percentage difference (Std. deviation of mean)	-14.9%% (4.8%)	18.7% (3.8%)
Wilcoxon statistic	-2.44	4.78
Restricted sample	22	73
English revenue higher	6	44
Equal revenue	2	7
Second-price revenue higher	14	22
Mean Dutch-first difference (Std. deviation of mean)	-0.15 (0.08)	0.79 (0.22)
Mean percentage difference (Std. deviation of mean)	-14.5% (9.3%)	18.1% (4.4%)
Wilcoxon statistic	-1.74	3.56

As noted, the effect here seems to be that the second auction in each pair tends to earn the most revenue, regardless of which auction was English and which one second-price. This could be partially due to my policy of sending an effectively greater quantity of direct email invitations to the second auction in each pair (as the second auction's invitation list included everyone on the first auction's email list, plus bidders who joined the first auction via a newsgroup advertisement). This argues for looking at the restricted sample of observations considered in the bottom half of Table 4, as the restricted sample excludes any observations which might have been tainted by such an effect. It includes only those cards whose winning bidders were invited via email to both auctions in the pair.

Restricting the sample according to this criterion does filter out a lot of observations, but it does not change the qualitative results. Again the majority of cards (14 of 22) in experiment SE had higher second-price than English revenues, while the majority of cards (44 of 73) in experiment ES had higher English than second-price revenues. The mean difference in revenue, as well as the Wilcoxon signed-rank test statistic, remained opposite in sign between the two experiments, although the statistical significance was reduced by the sample restriction. In particular, for experiment ES the Wilcoxon test can no longer reject the null hypothesis of revenue equivalence at the 95% confidence level. This may indicate that revenue equivalence really holds on this "clean" set of data, or it may just be due to the fact that the power of the test is reduced when the number of observations (only 22) becomes small.

Table 4 does not contain a column that pools the data for both auction experiments, as did Table 3, because the number of observations is much higher in one experiment than in the other. Simple pooling would give an inappropriate amount of weight to Experiment SE, which was the auction with more observations. Combining the restricted samples of data by giving equal weights to the two experiments (that is, computing a simple average) results in an overall mean difference of \$0.32 (±\$0.15) or a mean percentage difference of 1.8% (±6.8%). Thus, it seems that overall there is a tendency for the English auction revenues to be higher on average than the second-price auction revenues, but this difference is not necessarily statistically significant.

In general, the card-level data is somewhat inconclusive on the question of revenue equivalence between English and second-price auctions. Experiment SE, which is the experiment with more and better data, indicates that English revenue may be higher than second-price revenue, but the results of experiment ES indicate that this result may be just a time-order effect. On the other hand, this so-called "time-order effect" is opposite in sign to the order effect expected: it indicates that more revenue is earned by the second auction in any pair. In section 4.3, I discussed potential advertising effects that might cause the second auction to raise more revenue than the first, but the result remained even after my attempts to exclude observations which might include such advertising effects. Seeing the second auction in a pair earn more revenue than the first conflicts both with the anticipated demand-saturation effect discussed in section 4.3 and with the empirical findings in section 5.1 above and in Chapter 1 which document a time-ordering effect that favors the first auction in a pair. This causes me to be skeptical of the card-level data in these English and second-price auction experiments. In particular, Experiment ES, which was one of the earliest experiments in my research program, had a rather small amount of matched cards for auction (especially in dollar terms), and derived most of its bidders from newsgroup announcements rather than from email invitations, as shown in Table 2. Such factors warrant some suspicion of the data from the ES experiment. I am tempted to discount this data and rely on the data from Experiment SE, which in spite of a potential time-order bias which might have favored the second-price auction, indicate that the English auction revenues tend to be higher. In order to reach a more definite conclusion, it would be desirable to repeat Experiment ES to collect more data.

Next, I turn to data at the level of the individual bidder, to see if comparisons can be made between bids submitted by a single bidder on the same card in the two different auctions. If strategic equivalence holds, such pairs of bids should be approximately equal to each other. Starting with the set of all bidders who bid in both auctions in either pair, I restricted attention to those instances where the same bidder submitted both a second-price auction bid and an English auction bid on the same card.[20] This resulted in a total of 57 observations of matched pairs of bids submitted by a set of 7 bidders in experiment ES, and a similar set of 156 observations generated by 10 bidders in experiment SE.[21]

Unlike the Dutch auction, the English auction format allows me to collect data from each bidder on each card that they were interested in. However, there are still some limitations on the data which prevent me from making direct comparisons between the second-price auction bid and the highest English bid submitted by a particular bidder on a particular card. There are two reasons why, even if a bidder employed equivalent bid strategies, we would observe her highest English bid to be lower than his submitted second-price auction bid. First, the bidder may have *won* the English auction because the second-highest bidder dropped out, and thus her observed English bid would not be as high as her dominant strategy bid in the second-price auction. The second reason has to do with the possibility of "jump bids." Suppose a bidder was willing to bid \$5.00 either in a second-price or an English auction. In the second-price auction, I would actually observe his bid of \$5.00. Now suppose that in the English auction, his bid of \$4.50 was raised immediately to \$6.00 by

---

[20] One might argue that I should also have considered instances where, despite participating in both auctions in a pair, a bidder submitted a bid for card X in one auction but not in the other auction. (There were, in fact, dozens of such instances.) I could have included such observations by taking the nonbid to be a bid equal to zero, for example. However, I believe that in such cases, the decision not to bid on a card at all is probably due to changes in the bidder's demand for that card during the time between auctions (for example, the bidder obtains the card in some other way), rather than being caused by the choice of auction format. Therefore, I restrict my attention to those cases where the same bidder submits two bids on the same card in two different auctions. My assumption is that changes in bid levels may be due to the auction format, while discrete decisions whether or not to bid on a particular card are more likely to be due to demand shifts.

[21] One bidder overlapped between these two sets.

someone else. Then our bidder would never have the opportunity to submit a bid of $5.00, and thus his bid of $4.50 would mistakenly suggest that he was willing to bid less in the English auction than in the second-price auction. Thus, it is not possible to do a direct comparison of the bids in all cases to see whether strategic equivalence holds. However, as in the Dutch-first case, it is still possible to search for violations of strategic equivalence, and in this case the data is at least more complete.

A violation of strategic equivalence in favor of the English auction merely involves finding a bidder's English auction bid to be higher than her second-price auction bid on the same card. In order to discuss violations of the opposite sort, I shall first define the "English dropout price" (for a particular bidder and a particular card) as the highest bid level which that bidder failed to raise. (If that bidder was the winner on a card, then there was no "dropout price," so I define it as equal to zero.) For example, if bidder X had the current high bid of $1.50 on Tuesday, then bidder Y raised the price to $1.60 on Wednesday, and bidder X never again was observed to submit a bid on that card, then the "English dropout price" was $1.60 for bidder X. Now, a violation of strategic equivalence in the sense that the second-price bid was "too high" can be identified as a case in which the bidder's English bid was less than or equal to her English dropout price, which in turn was strictly less than her second-price bid. This then avoids counting as violations those cases where the bidder's English bid was too low because she either won the card or faced a jump bid that she was unwilling to raise.

Of the 213 matched-bid observations, 29 were violations of strategic equivalence in favor of the second-price auction format, ranging from $0.05 to $11.00, with a mean of $1.57 or 15.5% of Cloister value. There were 75 violations in favor of the English format, ranging from $0.20 to $8.00, with a mean of $1.70 or 17.8% of Cloister value. Of the remaining 109 observations, eight had exactly equal second-price and English bid strategies, while 101 were indeterminate (because the English dropout price did not rise high enough, or because of a jump bid, either of which would have prevented me from observing the bidder's English bid strategy precisely enough to compare it to her second-price bid.) Note that the categories used here have been designed to avoid introducing a sample-selection bias on the observations of bid violations. The violations favoring the English auction format outnumber the violations favoring the second-price format by a margin of nearly three to one,[22] and they are also larger in magnitude on average. This indicates that strategic equivalence fails to hold in this auction environment, and that bidders's strategies instead tend to involve higher amounts in an English than in a second-price auction.

Figure 5 displays the distribution of the difference in the bid level across the 112 non-indeterminate observations, while Figure 6 shows a similar graph of this difference as a percentage of each card's Cloister value. Observations to the right of center are observations

Figure 5: English-second differences observed in matched bids.

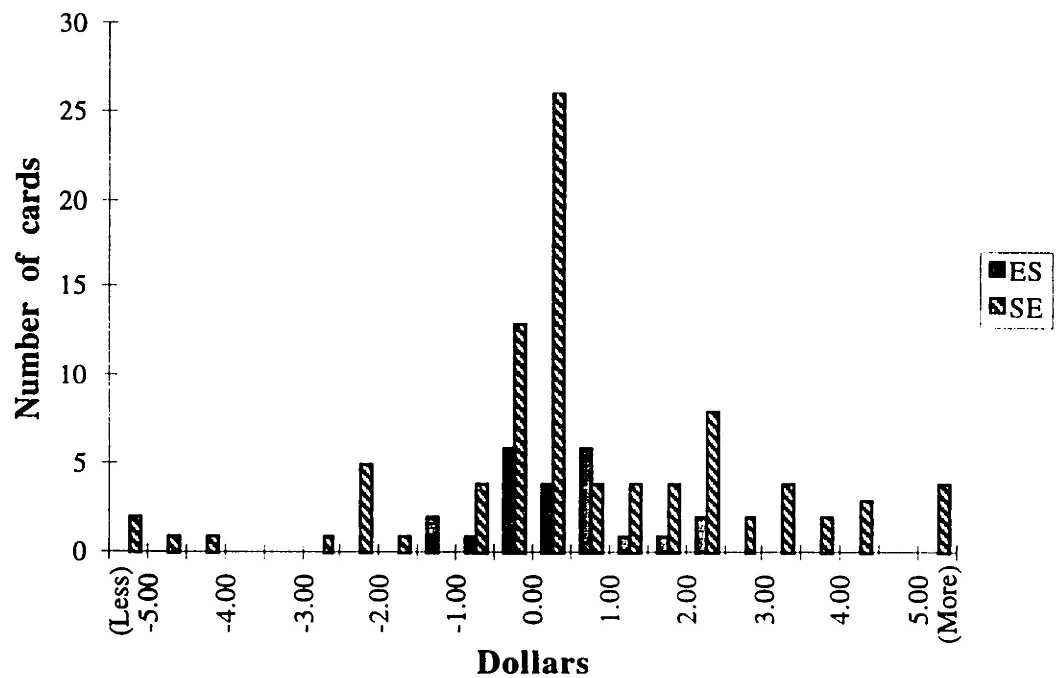

for which the English auction bid level was higher, while observations to the left are observations for which the second-price bid level was higher.

Breaking down these 112 matched-bid observations by individual bidder, it becomes apparent that some bidders behaved rather differently than others. Table 5 shows the mean difference (normalized by Cloister value) between the two bids submitted, with separate means computed for each bidder. Of the 16 different bidders for whom I had matched bid observations, ten of them tended to bid higher in the English auctions (that is, more of their

[22] Note that the category of "indeterminate" observations was developed specifically to avoid sample-selection bias problems in the data. For example, suppose that the price in an English auction never rose as high as bidder X's second-price auction bid on the same card. Then, although it might at first appear that his English bid were lower than his second-price bid, this is because the auction never gave him the opportunity to raise his English bid to be greater than or equal to his second-price bid. If bidders all were equally as likely to have higher bid strategies as lower bid strategies in English versus second-price auctions, then this "indeterminate" observation would be equally likely to go one direction as the other. I have some confidence, then, that the non-indeterminate data used are essentially free of sample-selection bias.

Figure 6: Bid differences normalized by Cloister value

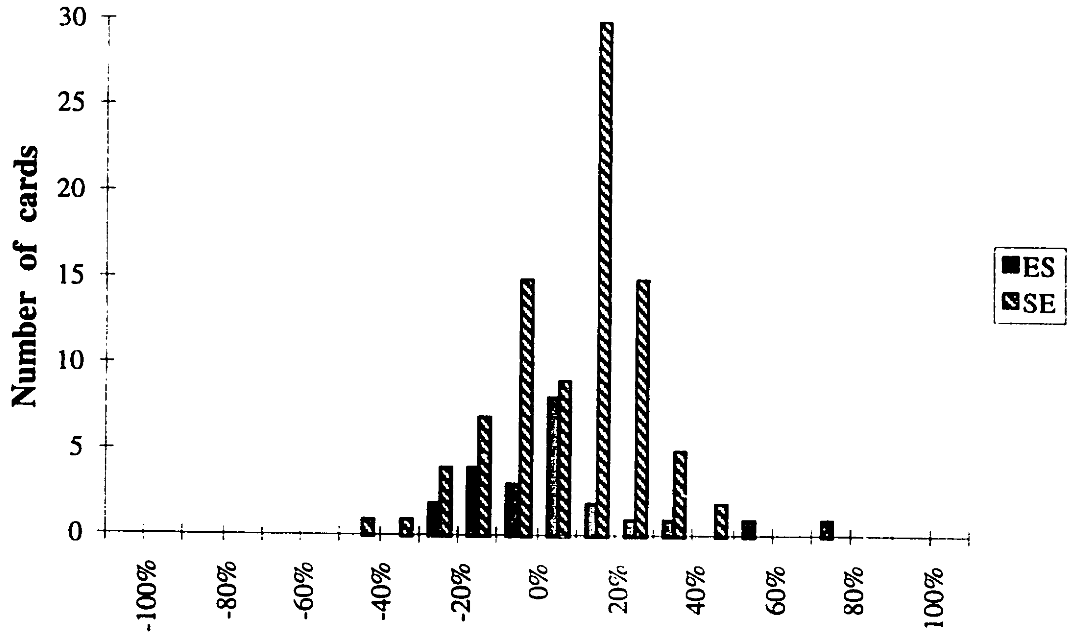

observed bid differences were positive than were negative), while five of them tended to bid higher in the second-price auctions. Of these, only bidder numbers 1, 6, and 9 seem convincingly and consistently to bid differently than strategic equivalence would suggest. (Each of these bidders yielded at least five different determinate observations, and for each the vast majority of observations favored the second-price auction.) Three bidders consistently seemed to favor the English auction, while one consistently favored the second-price auction. Thus, although the bidders taken as a whole tend to bid higher in the English rather than the second-price auction format, there is some heterogeneity by bidder.

To summarize, the card-level evidence on English and second-price auctions is somewhat inconclusive on the question of revenue equivalence, as Experiment ES failed to generate enough data to help distinguish potential time-order effects. Experiment SE, taken by itself, indicates that English revenues may be slightly higher than second-price revenues. This suspicion is confirmed somewhat by the bid-level data, in which there was a marked tendency (although slightly heterogenous across bidders) for bidders to violate strategic equivalence in favor of the English auction format. Taking the average across the bidder data displayed in Table 3, the bidders exhibited bids that were 3.0% higher on average in the English auction. This is somewhat surprising, as laboratory experiments have typically

Table 5: **Difference between English bid and second-price bid, as a percentage of the card's Cloister value**

Bidder	Experiment	Observations				Mean	Std. Dev.
		Positive	Zero	Negative	Index		
Bidder 1	SE	14	0	0	1	23.1%	9.8%
Bidder 2	ES	1	0	0	0	19.3%	
Bidder 3	ES	3	0	0	1	6.1%	2.0%
Bidder 4	SE	0	0	1	0	-2.7%	
Bidder 5	SE	8	1	6	5	3.6%	15.7%
Bidder 6	ES	0	0	5	17	-18.2%	11.3%
Bidder 7	SE	0	0	2	16	-11.7%	1.5%
Bidder 8	SE	1	0	0	0	17.2%	
Bidder 9	Both	36	3	10	34	9.7%	20.6%
Bidder 10	ES	0	0	1	0	-14.3%	
Bidder 11	SE	3	2	2	11	3.6%	9.7%
Bidder 12	ES	2	1	0	10	4.9%	4.6%
Bidder 13	SE	4	0	0	2	6.8%	5.2%
Bidder 14	SE	0	0	1	2	-23.8%	
Bidder 15	ES	1	0	1	2	10.8%	34.6%
Bidder 16	SE	2	1	0	0	14.1%	17.0%

observed overbidding in English versus second-price auctions with private values. A possible theoretical explanation is that bidders' valuations for Magic cards contain some common-value component, and this causes English bids to be higher for the reasons explained by Milgrom and Weber (1982). I feel, however, that the common-value component (which would be a speculative resale component to bidders' valuations) is very small in this market, so that valuations are approximately private. This makes the proposed theoretical explanation seem rather weak, but the experimental data are not strong enough to identify an alternative explanation to the potential puzzle.

5.3 A Revenue Comparison of All Four Auction Types

The previous two sections have investigated the strategic equivalence between Dutch and first-price auctions, and between English and second-price auctions. The experimental results thus far have demonstrated that Dutch auctions tend to raise more revenue than first-price auctions, and they have been unable to demonstrate a systematic difference between revenues in English and second-price auctions. This section attempts to address general revenue rankings of the four auctions: whether, for example, a first-price auction produces more or less expected revenue than a second-price auction.

The data in this chapter do not facilitate direct paired comparisons other than those comparisons already made in the previous two sections; I did not auction the same card twice in both first-price and second-price auctions, for example. However, it is still possible to use this data to make indirect comparisons, using the Cloister values of the auctioned cards as baselines for comparison. Unfortunately, not all of the data will be useful for such purposes, because there were significant differences between some of the auctions, such as the time period of the auction and the number of invited bidders, whose effects on revenue could be confounded with the effects of the auction format. Such considerations exclude four of the eight auctions undertaken in this study, but the remaining four (Auctions FD1, SE1, FD2, and SE2) all took place during the same month and with the same basic set of invited bidders[23] (see Tables 1 and 2 for details). The revenues earned on each of the cards in these auctions will allow additional revenue comparisons between auction formats. Because of potential time-order effects, I will not attempt to use this data to provide rankings of all four auction types; instead, I will group the first-price and Dutch auctions (Auctions FD1 and FD2) as one single unit and the English and second-price auctions (SE1 and SE2) as another unit, using the data to make comparisons between the two units. For convenience, I will label the first unit as FD and the second unit as ES.

As Section 3.3 indicated, the theoretical prediction in this case is somewhat ambiguous: if bidder risk aversion is important, FD may generate more revenue than ES, while if affiliation of bidder values is important, ES may generate more revenue than FD. If neither of these two effects are present, then the two units are predicted to generate equal expected revenue. In this market, my guess would have been that neither risk aversion nor affiliation

---

[23] Small changes in the numbers of invited bidders are due to certain bidders asking specifically to be removed from my mailing list; removing such bidders should not have any affect on revenue, because they wouldn't have bid in the auctions anyway.

of values were important. The stakes are relatively small, which might indicate approximate risk neutrality. The goods were purchased for private consumption (rather than resale) by people whose alternative options for buying the same cards (such as the availability of the card their local card dealer, or the cost of spending time looking for a good deal elsewhere on the Internet) seem unlikely to be correlated between bidders, which would indicate the absence of affiliation. Thus, to the extent that the theory accurately describes this market, I might have expected revenues to be equal between the FD and ES pairs. The degrees of bidder risk aversion and affiliation are impossible to know for certain, however (this would unfortunately be true in just about any real-world market), so I cannot strictly test the received theory. All I can say is which theoretical situation most nearly matches the experimental data.

To test whether FD auctions and ES auctions yield the same expected revenue, I estimate an econometric specification in which the auction revenue (REV) of each card depends on the card's Cloister value (CLOISTER), as well as on an intercept term which varies with the auction format (FD is an indicator variable that equals one only if the card was auctioned either via the first-price or via the Dutch format). I utilize a log-log specification, which causes any difference between the FD and ES formats to be estimated as a fixed percentage of card value. The estimation, with 370 observations and an R^2 of 0.805, gives the following results:

$$\begin{array}{llll} \ln \text{REV} = & -1.1890 & +0.2597\ \text{FD} & +1.2785 \ln \text{CLOISTER} \\ & (0.0752) & (0.0727) & (0.0332) \end{array}$$

As a specification check, note that the coefficient on Cloister value is close to unity, which indicates that auction revenue is nearly proportional to the Cloister value. (The fact that the coefficient is somewhat greater than unity indicates that auctions tend to generate proportionately slightly more revenue for higher-priced cards; this may indicate that higher-priced cards attract more attention in Internet auctions than do lower-price cards.[24]) Of primary interest is the fact that the coefficient on FD is positive and significant at the 95% confidence level. Its magnitude indicates that cards in FD auctions earn nearly 30% more revenue than cards in ES auctions. This is consistent with a theory of bidder risk aversion in auctions, but not with a theory of risk neutrality and affiliated bidder values.

[24] A similar regression in Chapter 1, for a different experiment, found a coefficient on log-Cloister that was almost exactly equal to unity, so the coefficient reported here is somewhat surprisingly large.

Table 6 reports the results of additional regression analysis on this data. The first spec-

Table 6: Regression analysis of FD versus ES revenues. The dependent variable is the natural logarithm of REVENUE.

Variable	Specification 1	Specification 2	Specification 3
C	-1.1890 (0.0752)	-0.7067 (0.0529)	-1.4058 (0.0975)
FD	0.2597 (0.0727)	0.1160 (0.0771)	0.5983 (0.1222)
ln(CLOISTER)	1.2785 (0.0332)	1.0000	1.4037 (0.0491)
FD*ln(CLOISTER)			-0.2254 (0.0659)

ification is the specification reported earlier, while the second specification constrains the coefficient on ln(CLOISTER) to be equal to one (so that revenues are exactly proportional to Cloister values). Under this specification, the coefficient on FD remains positive, but is no longer statistically significant at the 95% confidence level; the p-value for a two-tailed test becomes 0.867 in this case. Finally, I consider adding an interaction term between the log Cloister value and the indicator variable for the auction format; this is reported as the third specification in the table. The interaction term is negative and significant, which indicates that the difference between FD revenue and ES revenue declines as the Cloister value of the auctioned card increases. In fact, for Cloister values greater than approximately $13.00, this specification predicts that the ES auction formats would bring in greater revenues than the FD auction formats. There were rather few cards of such high values in the data, however, which make such extrapolations rather tenuous. In general, then, the data indicate that FD auctions raise more revenues than ES revenues for the types of goods sold in these auctions, but that the revenue difference is lower for the higher-valued goods.

Combining this result with the results of the previous two sections, I can say with some confidence that Dutch auctions raise the most revenue in this market, followed by Dutch auctions, English auctions, and second-price auctions. The revenue difference between the last two auction types is the most difficult to distinguish with the present data.

6 Concluding Remarks

During the 1980s, the question revenue equivalence between different auction formats was a subject of intense study in laboratory experiments. Their findings indicated that revenue equivalence was violated, and that the four auction formats could be ranked as first-price, then Dutch, then second-price and English.

The purpose of this chapter is to examine the extent to which these laboratory results are representative of auctions in the real world. In the market examined here, the revenue rankings of the auction formats are Dutch, then first-price, then English and second-price.

The field experiments in this chapter sacrificed some of the controls enjoyed by traditional laboratory experimenters (such as the ability to induce and observe bidder valuations) in exchange for greater realism: the auctions took place for real goods rather than cash payoffs, and they took place in a preexisting auction environment rather than a laboratory. Thus, this work is complementary to laboratory research. Vernon Smith, one of the pioneers of experimental economics, puts it this way: "Control is not everything, and if the lab stuff is relevant to the field we ought to be able to get consistency tests. I have never considered the lab to be a substitute for field empirical work."[25]

In a recent survey of auction theory, Wilson (1992) also discussed the tradeoff between field studies and experimental research. By comparison with laboratory research, Wilson pointed out that field IO studies of auctions "must contend with less complete data, and few controls on the auction environment are possible. On the other hand, they have the advantage that the data pertain to practical situations in which the stakes are often large and the participants are skilled and experienced." The field experiments in this chapter are intermediate between these two extremes. They share with laboratory experiments some, but not all, of the ability to control the relevant economic variables, and with field studies the advantage that subjects tended to be familiar with the auction environment. The stakes in my field experiments are no higher than those in most laboratory experiments, but these field experiments are more realistic in at least two important senses: bidders were bidding for real goods rather than being given an explicit cash payoff structure, and bidders were not explicitly told the distribution of others' valuations (such information is never known perfectly in practice).

---

[25] Private communication with Vernon Smith, April 1996.

The principal findings of this chapter are threefold. First, Dutch auctions tend to raise more revenue than first-price auctions by approximately five percent on average. Second, bidders tend to bid higher in English auctions than in second-price auctions by approximately three percent on average, although there is quite a bit of heterogeneity among bidders, and the overall revenue comparisons between English and second-price auctions were inconclusive. Finally, Dutch and first-price auctions tend to raise more revenue than English and second-price auctions, by nearly thirty percent on average.

The first finding contrasts with results found in the laboratory, where Dutch auctions consistently raise *less* revenue than first-price auctions. The most likely explanation for this different result is the fact that the number of bidders was not fixed in these experiments, as it has been in laboratory experiments. Indeed, bidder participation was greater in the Dutch auctions than in the corresponding first-price auctions in this study. An interesting hypothesis is whether these additional bidders might have been attracted by the novelty of participating in a Dutch auction, and whether this effect might disappear after many repeated trials. However, it is not clear that Dutch auctions were more novel than sealed-bid auctions in this market, as the vast majority of Magic auctions on the Internet during this period were English auctions.[26] Other differences between these field experiments and the more traditional laboratory experiments about revenue equivalence include: (1) the use of real goods rather than explicit cash payoffs, (2) much slower speed of the Dutch clock, and (3) simultaneous rather than sequential auction format. In bid-level data, there were a few more observations of bidders violating strategic equivalence in favor of the Dutch format than in favor of the first-price format, so I speculate that using real goods instead of laboratory cash payoffs may introduce a sort of anchoring effect of the Dutch clock price, which is absent in the first-price format.

The second finding, that of finding bidders violating strategic equivalence in favor of the English auction over the second-price auction format, also is somewhat contradictory of laboratory data, which tend to find overbidding in second-price auctions. A possible explanation of this discrepancy is that bidder valuations in these field experiments may have some common-value component which is affiliated across bidders, whereas laboratory experiments on revenue equivalence have been able to impose private values. If this is true, it is somewhat surprising, as there are good arguments (such as the fact that speculative buy-

[26] Dutch auctions were not entirely novel to this market, either. One bidder wrote me a message to explain that he did not want to participate in my Dutch auction, because, as he said, "I participated in one previous Dutch Auction on the Net, and I found it to be very very frustrating."

ing behavior is uncommon among these bidders) why bidder valuations in this market might be reasonably well approximated by private values. If affiliation of values is important in this market, then affiliation is probably even more important in other real-world markets with more possibility for speculative resale of the good.

The third finding, of higher revenues in Dutch and first-price auctions than in English and second-price auctions, is consistent with laboratory studies. It is also consistent with a theory of bidder risk aversion, in which bidders are reluctant to shade their bids very far below their valuations in a Dutch or first-price auction. One unexplained feature of the data is that the revenue difference between auction formats appears to decline with the value of the good, suggesting (if one takes risk aversion to be the source of the revenue difference) that risk aversion is more important for cards worth \$1.00 than for cards worth \$10.00. It may be that, as has been demonstrated in laboratory experiments, risk aversion cannot explain the entire difference.

These field experiments have provided, for the first time, evidence from real-world auctions that allow revenue comparisons between each of the four basic auction formats. However, a number of questions about revenue equivalence still remain. Just as it is important to check whether laboratory results extend to real-world auction markets, it is also important to see whether the results obtained in this market can be replicated in other real-world markets. It would be very interesting to attempt similar field experiments in a completely different market, such as the market for used farm equipment, to see if these results can be replicated. The findings in this chapter also point out useful directions for laboratory research, such as adding the potential for endogenous bidder entry to laboratory studies of revenue equivalence. Another important direction, especially given recent advances in communications technology which have led both to card auctions on the Internet and simultaneous English-style auctions for FCC communications spectrum, is towards more careful theoretical and experimental explorations of the properties of simultaneous (as opposed to sequential) auction formats.

References

Ashenfelter, Orley, "How Auctions Work for Wine and Art." *Journal of Economic Perspectives*, 1989, vol. 3, no. 3, pp. 23-36.

Black, Jason. *Cloister's Magic Card Price List*, `<http://www.hhhh.org/cloister/pricelists/>`, various weekly issues, 1995.

Coppinger, Vicki M., Vernon L. Smith, and Jon A. Titus, "Incentives and Behavior in English, Dutch and Sealed-Bid Auctions," *Economic Inquiry*, 1980, vol. 18, pp. 1-22.

Cox, James C., Bruce Roberson, and Vernon L. Smith, "Theory and Behavior of Single Object Auctions," in *Research in Experimental Economics*, Vernon L. Smith, ed., Greenwich, Conn.: JAI Press, 1982, pp. 1-43.

Cox, James C., Vernon L. Smith, and James M. Walker, "A Test that Discriminates Between Two Models of the Dutch-First Non-Isomorphism," *Journal of Economic Behavior and Organization*, 1983, vol. 4, pp. 205-219.

Cox, James C., Vernon L. Smith, and James M. Walker, "Theory and Individual Behavior of First-price Auctions," 1988, *Journal of Risk and Uncertainty*, vol. 1, pp. 61-99.

Cox, James C., Vernon L. Smith, and James M. Walker, "Theory and Misbehavior of First-price Auctions: Comment," *American Economic Review*, 1992, vol. 82, no. 5, pp. 1392-1412.

Hansen, Robert G. "Sealed-Bid Versus Open Auctions: The Evidence," *Economic Inquiry*, 1986, vol. 24, no. 1, pp. 125-142.

Harrison, Glenn W. "Theory and Misbehavior of First-price Auctions," *American Economic Review*, 1989, vol. 79, no. 4, pp. 749-762.

Harrison, Glenn W. "Theory and Misbehavior of First-price Auctions: Reply," *American Economic Review*, 1992, vol. 82, no. 5, pp. 1426-1443.

Hendricks, Kenneth, and Harry J. Paarsch, "A Survey of Recent Empirical Work Concerning Auctions." *Canadian Journal of Economics*, 1995, vol. 28, no. 2, pp. 315-338.

Kagel, John H. "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, J. Kagel and A. Roth, eds. Princeton: Princeton University Press, 1995, pp.501-585.

Kagel, John H., Ronald M. Harstad, and Dan Levin, "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 1987, vol. 55, no. 6, pp. 1275-1304.

Kagel, John H., and Dan Levin, "Independent Private Value Auctions: Bidder Behaviour in First-, Second-, and Third-Price Auctions with Varying Numbers of Bidders," *Economic Journal*, 1993, vol. 103, pp. 868-879.

Kagel, John H., and Alvin E. Roth, "Theory and Misbehavior of First-price Auctions: Comment," *American Economic Review*, 1992, vol. 82, no. 5, pp. 1379-1391.

Magic: the Gathering home page, <http://www.itis.com/deckmaster/index.html>.

McAfee, R. Preston, and John McMillan. "Auctions and Bidding." *Journal of Economic Literature*, 1987, vol. 25, no. 2, pp. 699-738.

Mead, Walter J. "Natural Resource Disposal by Oral Auction Versus Sealed Bids," *Natural Journal*, 1985, vol. 7, no. 2, pp. 195-224.

McMillan, John, "Selling Spectrum Rights," *Journal of Economic Perspectives*, 1994, vol. 8, no. 3, pp. 145-162.

Riley, John G., and William Samuelson, "Optimal Auctions," *American Economic Review*, 1981, vol. 71, no. 3, pp. 381-392.

Rothkopf, Michael H., and Ronald M. Harstad, "Modeling Competitive Bidding: A Critical Essay," Management Science, 1994, vol. 40, no. 3, pp. 364-384.

Smith, Vernon L. "Experimental Studies of Discrimination versus Competition in Sealed-Bid Auction Markets," *Journal of Business*, 1967, vol. 40, pp. 58-84.

Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 1961, vol. 16, no.1, pp. 8-37.

Wilson, Robert. "Strategic Analysis of Auctions," in *The Handbook of Game Theory*, R.J. Aumann and S. Hart, eds. New York: North-Holland, 1992, pp. 227-279.

Appendix 1. Sample first-price auction announcement.

Date: Fri, 05 May 1995 16:16:01 -0400
From: reiley@MIT.EDU (David Reiley)
Subject: Reiley's Auction #10: LG Black/Blue, 5 Cent Min, Free Shipping!
Newsgroups: rec.games.trading-cards.marketplace

** Please read the rules of this auction carefully, as each auction I run typically has a different set of rules. **

This will be a SEALED-BID, FIRST-PRICE AUCTION. Here's how it works:

I will accept all bids up until the deadline of NOON (Eastern Daylight Time), next FRIDAY, May 12, 1995. All bids are "sealed" in the sense that I will not post updates or otherwise reveal information about the highest bid until the auction is over.

After the deadline for bids has passed, I will award each card to the highest bidder at the price of their bid.

Note that SHIPPING IS INCLUDED in the bid price. If you win, you mail me exactly the amount of your bid, with no extra charges. This is to encourage everyone to bid separately on each card they're interested in - no worrying about having to win multiple cards in order to make it worth your while.

Here are THE RULES:

1. Submit bids via email to <reiley@mit.edu>. Make sure that the subject line of your email contains the phrase "Auction #10". (Simply using the "reply" command on most mail and news programs should work just fine.) If your message does not contain this text in the subject line, your message will be discarded.

2. In your email message, please put each of your bids on a separate line of text. Each bid line should be in the following format: the 3-digit identification number of the card you're bidding on, immediately followed by a right parenthesis, and then the amount of your bid in dollars and cents (such as 1.00). For example:

305) 2.00
306) 0.65

Including extra information on the bid line is okay, too. Such extra information might include email quote marks (such as the greater-than symbol), the card name, condition, etc. You may include anything that makes bidding easier for you, EXCEPT that your bid amount should be the only price-formatted number which appears on that line. In other words, no other number containing a decimal point should appear on that line.

For example, the following are also perfectly valid bids:

```
> 305) Ghosts of the Damned            U1     Blk   M  2.00
> 306) Demonic Torment   0.65
```

Bids that do not conform to these rules will be discarded.
Here are examples of INVALID bids:

```
>305) Ghosts of the Damned      $2                      [no decimal point]
306 ) Demonic Torment 0.65  [space between the 6 and the right parenthesis]
```

3. All bids must be in US currency, and must also satisfy the following:

   a. Bids under 1 dollar must be even multiples of a nickel ($0.05).
   b. Bids from 1 to 5 dollars must be even multiples of a dime ($0.10).
   c. Bids from 5 to 10 dollars must be even multiples of a quarter ($0.25).
   d. Bids from 10 to 20 dollars must be even multiples of a 50 cents ($0.50).
   e. Bids over 20 dollars must be even multiples of a dollar ($1.00).

4. The auction closes on Friday, May 12, 1995, at noon (EDT). Any bids received after that time will be ignored. All cards receiving a bid of at least five cents will be sold at that point to the highest bidder. In the case of a tie, the winner will be the person whose bid was received first.

5. The winning bidder will be notified by email, and will be asked to pay the amount of his/her bid via US check or money order.

6. This payment will include free shipping within the United States, via first class mail. The cards will be wrapped in plastic sheaths and packed in cardboard for protection. All cards will be shipped within one week after the receipt of payment.

7. While this is a real auction for real cards, you should know that I plan to use data on the bids in this auction for economic research. In no case will individual bidders be identified in this research; anonymity will be preserved. By bidding in this auction, you indicate your consent to have your bid be used in economic research. If you do not approve of this, you have the right not to participate in this auction. Should you have any questions or concerns about the use of data from this auction in academic

research, please contact the chair of the COUHES committee at MIT by phone at 617-253-6787.

That's it! Enjoy the auction. Good luck, and thanks for participating!

Here is the LIST OF CARDS for auction. All are from the LEGENDS expansion:

ID	Card Name	Rarity	Color	Cond
001)	Abomination	U1	Blk	M
002)	All Hallow's Eve	R	Blk	M
003)	Blight	U1	Blk	M
004)	Carrion Ants	R	Blk	M
005)	Chains of Mephistopheles	R	Blk	M
006)	Cosmic Horror	R	Blk	M
007)	Cyclopean Mummy	C2	Blk	M
008)	Darkness	C1	Blk	M
009)	Demonic Torment	U1	Blk	M
010)	Evil Eye of Orms-By-Gore	U1	Blk	M
011)	Fallen Angel	U1	Blk	M
012)	Ghosts of the Damned	C2	Blk	M
013)	Giant Slug	C2	Blk	M
014)	Glyph of Doom	C2	Blk	M
015)	Greed	R	Blk	M
016)	Headless Horseman	C1	Blk	M
017)	Hell Swarm	C1	Blk	M
018)	Hell's Caretaker	R	Blk	M
019)	Hellfire	R	Blk	M
020)	Horror of Horrors	U1	Blk	M
021)	Imprison	R	Blk	M
022)	Infernal Medusa	U1	Blk	M
023)	Jovial Evil	R	Blk	M
024)	Lesser Werewolf	U1	Blk	M
025)	Lost Soul	C2	Blk	M
026)	Mold Demon	R	Blk	M
027)	Nether Void	R	Blk	M
028)	Pit Scorpion	C2	Blk	M
029)	Quagmire	U1	Blk	M
030)	Shimian Night Stalker	U1	Blk	M
031)	Spirit Shackle	C1	Blk	M
032)	Syphon Soul	C2	Blk	M
033)	Takklemaggot	U1	Blk	M
034)	The Abyss	R	Blk	M
035)	The Wretched	R	Blk	M

036) Touch of Darkness	U1	Blk	M
037) Transmutation	C1	Blk	M
038) Underworld Dreams	U1	Blk	M
039) Vampire Bats	C2	Blk	M
040) Walking Dead	C1	Blk	M
041) Wall of Putrid Flesh	U1	Blk	M
042) Wall of Shadows	C2	Blk	M
043) Wall of Tombstones	U1	Blk	M
044) Black Common Set (15 cards)	C	Blk	M
045) Acid Rain	R	Blu	M
046) Anti-Magic Aura	C1	Blu	M
047) Azure Drake	U1	Blu	M
048) Backfire	U1	Blu	M
049) Boomerang	C2	Blu	M
050) Brine Hag	U1	Blu	M
051) Devouring Deep	C2	Blu	M
052) Dream Coat	U1	Blu	M
053) Elder Spawn	R	Blu	M
054) Enchantment Alteration	C1	Blu	M
055) Energy Tap	C2	Blu	M
056) Field of Dreams	R	Blu	M
057) Flash Counter	C2	Blu	M
058) Flash Flood	C2	Blu	M
059) Force Spike	C2	Blu	M
060) Gaseous Form	C1	Blu	M
061) Glyph of Delusion	C1	Blu	M
062) In the Eye of Chaos	R	Blu	M
063) Invoke Prejudice	R	Blu	M
064) Juxtapose	R	Blu	M
065) Land Equilibrium	R	Blu	M
066) Mana Drain	U1	Blu	M
067) Part Water	U1	Blu	M
068) Psionic Entity	R	Blu	M
069) Psychic Purge	C1	Blu	M
070) Puppet Master	U1	Blu	M
071) Recall	R	Blu	M
072) Relic Bind	U1	Blu	M
073) Remove Soul	C2	Blu	M
074) Reset	U1	Blu	M
075) Reverberation	R	Blu	M
076) Sea King's Blessing	U1	Blu	M
077) Segovian Leviathan	U1	Blu	M
078) Silhouette	U1	Blu	M
079) Spectral Cloak	U1	Blu	M

080) Telekinesis	R	Blu	M
081) Teleport	R	Blu	M
082) Time Elemental	R	Blu	M
083) Undertow	U1	Blu	M
084) Venarian Gold	C1	Blu	M
085) Wall of Vapor	C2	Blu	M
086) Wall of Wonder	U1	Blu	M
087) Zephyr Falcon	C2	Blu	M
088) Blue Common Set (15 cards)	C	Blk	M

END OF LIST

Key to Card Types and Conditions:

C = Common
C1 = Common (printed once per sheet of commons)
C2 = Common (twice as common as C1, printed twice per common sheet)
C4 = Common (four times as common as C1)
C11= Common (eleven times as common as C1)
U = Uncommon
U1 = Uncommon (printed once per sheet of uncommons)
U2 = Uncommon (twice as common as U1)
U3 = Uncommon (three times as common as U1)
R = Rare

Lnd = Land
Art = Artifact
Blk = Black
Blu = Blue
Gre = Green
Red = Red :)
Whi = White

M = Mint
NM = Near Mint (never played, but has tiny blemishes from handling)
E = Excellent (played a few times, has small scuff marks)

If you are unfamiliar with some of these cards, you can get information about any Magic card (spell type, power, toughness, artist's name, etc.) from the following Internet sources:

http://www.itis.com:80/deckmaster/magic/cardinfo/
ftp://marvin.macc.wisc.edu/pub/deckmaster/card.info/lists.w.spoilers/

Remember to send any bids, comments, or questions about the cards or the rules of this auction to <reiley@mit.edu>.

Thanks!

Appendix 2. Sample Dutch auction update.

Date: Fri, 19 May 1995 14:36:45 -0400
To: reiley@MIT.EDU (Dave Reiley)
From: reiley@MIT.EDU (Dave Reiley)
Subject: Reiley's Auction #12: Black/Blue Legends - Update #2

This is a DUTCH auction, which is the auction format traditionally used to sell tulip bulbs in Holland. Here's how it works:

By contrast with an ordinary (English-style) auction, in which the prices of the cards rise over time, this is an auction in which the prices of the cards FALL over time. The price of the card begins rather high and falls each day until someone is willing to bid. The first person to bid gets the card at the price that was posted that day.

Note that SHIPPING IS INCLUDED in the bid price. If you win, you mail me exactly the amount of your bid, with no extra charges. This is to encourage everyone to bid separately on each card they're interested in - no worrying about having to win multiple cards in order to make it worth your while.

If you are interested in participating in this auction, PLEASE EMAIL ME to let me know! I will then put you on the list of people to receive daily email updates of the price reductions. If I don't hear from you, this is the last email you'll receive about this auction.

Here are THE RULES. Please read them carefully:

1. Prices on unsold cards in this auction will be reduced once per day, Monday through Saturday, typically between noon and 1pm (EDT). The new prices will be announced through two channels: a daily post to the newsgroup <rec.games.trading-cards.marketplace>, and a daily email update to a list of interested bidders. If you are interested in receiving these updates via email, please request via email to <reiley@mit.edu> to be added to the mailing list.

2. When the price has reached the level at which you are willing to bid for the card, submit your bid(s) via email to <reiley@mit.edu>. Make sure that the subject line of your email contains the phrase "Auction #12". (Simply using the "reply" command on most mail and news programs should work just fine.) If your message does not contain this text in the subject line, your message will be discarded.

3. In your email message, please put each of your bids on a separate line of text. Each bid line should be in the following format: the 3-digit identification number of the card you're bidding on, immediately followed by a right parenthesis, and then the amount of the bid in dollars-and-cents format. (The bid amount is fixed by the current day's prices, but I ask you to include it as a cross-check.) For example:

```
305) 2.00
306) 0.65
```

Including extra information on the bid line is okay, too. Such extra information might include email quote marks (such as the greater-than symbol), the card name, condition, etc. You may include anything that makes bidding easier for you; just make sure that each bid is on a separate line.

For example, the following are also perfectly valid bids:

```
> 305) Ghosts of the Damned          U1    Blk   M  2.00
> 306) Demonic Torment   0.65
```

Bids that do not conform to these rules will be discarded.
Here are examples of INVALID bids:

```
>305) Ghosts of the Damned     $2                         [no decimal point]
306 ) Demonic Torment 0.65  [space between the 6 and the right parenthesis]
```

4. Any submitted bids which do not contain that day's correct price will be ignored, as will any submitted bids for cards which have already been identified as sold.

5. In subsequent updates, the winner of each card will be identified, so that bidders will know whether they were in fact the first people to bid for their respective cards. To preserve anonymity, each person will be identified only by the part of her email address which appears to the left of the @ sign. If you wish me to keep even that information confidential, please send me a nickname by which I can refer to you in updates.

6. The first bidder to submit a valid bid on any card will win that card. The winner will be asked to pay the amount of her bid, in the form of a US check or money order.

7. This payment will include free shipping within the United States, via first class mail. The cards will be wrapped in plastic sheaths and packed in

cardboard for protection. All cards will be shipped within one week after the receipt of payment.

8. While this is a real auction for real cards, you should know that I plan to use data on the bids in this auction for economic research. In no case will individual bidders be identified in this research; anonymity will be preserved. By bidding in this auction, you indicate your consent to have your bid be used in economic research. If you do not approve of this, you have the right not to participate in this auction. Should you have any questions or concerns about the use of data from this auction in academic research, please contact the chair of the COUHES committee at MIT by phone at 617-253-6787.

That's it! Enjoy the auction. Again, I'm quite serious about the fact that every card will be sold to the highest bidder, no matter how small the bid.

Good luck, and thanks for participating!

Here is the LIST OF CARDS for auction. All are from the LEGENDS expansion:

[Note: any cards marked with a "Winner" have already been sold.]

ID	Card Name	Rarity	Color	Cond	Price	Winner
201)	Abomination	U1	Blk	M	2.00	Matt.Gomes
202)	All Hallow's Eve	R	Blk	M	30.00	
203)	Blight	U1	Blk	M	3.00	wneale
204)	Carrion Ants	R	Blk	M	20.00	michele
205)	Chains of Mephistopheles	R	Blk	M	10.50	
206)	Cosmic Horror	R	Blk	M	7.75	michele
207)	Cyclopean Mummy	C2	Blk	M	0.30	
208)	Darkness	C1	Blk	M	0.75	michele
209)	Demonic Torment	U1	Blk	M	1.40	ent
210)	Evil Eye of Orms-By-Gore	U1	Blk	M	2.60	Matt.Gomes
211)	Fallen Angel	U1	Blk	M	10.50	manatee
212)	Ghosts of the Damned	C2	Blk	M	0.35	gbwhite
213)	Giant Slug	C2	Blk	M	0.35	
214)	Glyph of Doom	C2	Blk	M	0.35	gbwhite
215)	Greed	R	Blk	M	5.00	kmiyake
216)	Headless Horseman	C1	Blk	M	0.45	psharkey
217)	Hell Swarm	C1	Blk	M	0.55	gbwhite
218)	Hell's Caretaker	R	Blk	M	21.00	
219)	Hellfire	R	Blk	M	12.50	ent

220)	Horror of Horrors	U1	Blk	M	3.00	michele
... *(more cards)* ...						
286)	Zephyr Falcon	C2	Blu	M	1.00	gbwhite
287)	Blue Common Set (15 cards)	C	Blk	M	8.25	snag

END OF LIST

Key to Card Types and Conditions:

C = Common
C1 = Common (printed once per sheet of commons)
C2 = Common (twice as common as C1, printed twice per common sheet)
C4 = Common (four times as common as C1)
C11= Common (eleven times as common as C1)
U = Uncommon
U1 = Uncommon (printed once per sheet of uncommons)
U2 = Uncommon (twice as common as U1)
U3 = Uncommon (three times as common as U1)
R = Rare

Lnd = Land
Art = Artifact
Blk = Black
Blu = Blue
Gre = Green
Red = Red :)
Whi = White

M = Mint
NM = Near Mint (never played, but has tiny blemishes from handling)
E = Excellent (played a few times, has small scuff marks)

If you are unfamiliar with some of these cards, you can get information about any Magic card (spell type, power, toughness, artist's name, etc.) from the following Internet sources:

http://www.itis.com:80/deckmaster/magic/cardinfo/
ftp://marvin.macc.wisc.edu/pub/deckmaster/card.info/lists.w.spoilers/

Remember to send any bids, comments, or questions about the cards or the rules of this auction to <reiley@mit.edu>.

Thanks!

Appendix 3. Sample second-price auction announcement.

Date: Thu, 09 Mar 1995 22:41:45 -0500
From: reiley@MIT.EDU (David Reiley)
Subject: Reiley's Auction #7: LEGENDS, 5 Cent Minimum, Free Shipping!
Newsgroups: rec.games.trading-cards.marketplace
Organization: MIT Dept. of Economics

Here's another opportunity to get cards at great prices!

Please read the rules of this auction carefully, as each auction I run typically has a different set of rules.

This will be a sealed-bid, SECOND-PRICE auction. Here's how it works:

I will accept all bids up until the deadline of NOON (Eastern Standard Time), next THURSDAY, March 16,1995. All bids are "sealed" in the sense that I will not post updates or otherwise reveal information about the highest bid until the auction is over.

After the deadline for bids has passed, I will award each card to the highest bidder, but at the price submitted in the bid of the SECOND-HIGHEST bidder.

For example, suppose that I receive 20 different bids on a Demonic Hordes, and that the two highest bids are from Emily Dickinson and Nathaniel Hawthorne. Suppose that Emily bids $8.00 while Nathaniel bids $5.50. Then Emily wins the card, and the price she pays is only $5.50.

One interesting result of this type of scheme is that if you are the only bidder for a particular card, then you get it for FREE! Of course, I'm hoping this won't happen, but I guarantee that I will honor the terms of this auction, and not withdraw a card if the bids are too low.

Also, note that SHIPPING IS INCLUDED in the bid price. If you win, you mail me exactly the amount of the second-highest bid, with no extra shipping charges. This is to encourage everyone to bid separately on each card you're interested in - no worrying about having to win multiple cards in order to make it worth your while.

Here are THE RULES:

1. Submit bids via email to <reiley@mit.edu>. Make sure that the subject line of your email contains the phrase "Auction #7". (Simply using the "reply" command on most mail and news programs should work just fine.) If your message does not contain this text in the subject line, your message will be discarded.

2. In your email message, please put each of your bids on a separate line of text. Each bid line should be in the following format: the 3-digit identification number of the card you're bidding on, immediately followed by a right parenthesis, and then the amount of your bid in dollars and cents (such as 1.00). For example:

```
 305) 2.00
 306) 0.65
```

Including extra information on the bid line is okay, too. Such extra information might include email quote marks (such as the greater-than symbol), the card name, condition, etc. You may include anything that makes bidding easier for you, EXCEPT that your bid amount should be the only price-formatted number which appears on that line. In other words, no other number containing a decimal point should appear on that line.

For example, the following are also perfectly valid bids:

```
 > 305) Ghosts of the Damned            U1     Blk   M  2.00
 > 306) Demonic Torment   0.65
```

Bids that do not conform to these rules will be discarded.
Here are examples of INVALID bids:

```
 >305) Ghosts of the Damned       $2                        [no decimal point]
 306 ) Demonic Torment 0.65  [space between the 6 and the right parenthesis]
```

3. All bids must be multiples of a nickel ($0.05), in US currency.

4. The auction closes on Thursday, March 16 at noon (EST). Any bids received after that time will be ignored. All cards receiving a bid of at least one cent will be sold at that point to the highest bidder. In the case of a tie, the winner will be the first person to have posted a bid of the high amount.

5. The winning bidder will be notified by email, and will be asked to pay not

the amount of her own bid, but the amount of the second-highest bid. Payment should be made by US check or money order.

6. The winning bidder will be notified by email, and will be asked to pay the amount of his/her bid via US check or money order.

7. This payment will include free shipping within the United States, via first class mail. The cards will be wrapped in plastic sheaths and packed in cardboard for protection. All cards will be shipped within one week after the receipt of payment.

8. While this is a real auction for real cards, you should know that I plan to use data on the bids in this auction for economic research. In no case will individual bidders be identified in this research; anonymity will be preserved. By bidding in this auction, you indicate your consent to have your bid be used in economic research. If you do not approve of this, you have the right not to participate in this auction. Should you have any questions or concerns about the use of data from this auction in academic research, please contact the chair of the COUHES committee at MIT by phone at 617-253-6737.

That's it! Enjoy the auction. Again, I'm quite serious about the fact that every card will be sold to the highest bidder, no matter how small the bid.

Good luck, and thanks for participating!

Here is the LIST OF CARDS:

ID	Card Name	Rarity	Color	Cond
701)	Adventurers' Guildhouse	U1	Lnd	NM
702)	Cathedral of Serra	U1	Lnd	NM
703)	Seafarers' Quay	U1	Lnd	NM
704)	The Tabernacle at Pendrell	R	Lnd	M
...				
784)	Sol'kanar the Swamp King	R	Gld	M
785)	Tetsuo Umezawa	R	Gld	M

END OF LIST

Appendix 4. Sample English auction update.

Date: Sat, 03 Jun 1995 17:13:10 -0400
From: reiley@MIT.EDU (David Reiley)
To: reiley@MIT.EDU (David Reiley)
Subject: Reiley's Auction #13: Whi/Gld LEGENDS, 5 Cent Min, Free Shipping!
Newsgroups: rec.games.trading-cards.marketplace

[Note: If you are done bidding in this auction, you may email me to ask me to remove you from the mailing list. If you've won cards, I will also send you your final bill.]

Please read the rules of this auction carefully, as each auction I run typically has a different set of rules.

This will be a standard, English-style ("Going, Going, Gone!") auction.

There will be updates posted daily, but they will be mailed only to people who bid or who ask specifically to be added to the mailing list for this particular auction.

Note that SHIPPING IS INCLUDED in the bid price. If you win, you mail me exactly the amount of your bid, with no extra charges. You don't have to worry about having to win multiple cards in order to make the auction worth your while.

Here are THE RULES:

1. Submit bids via email to <reiley@mit.edu>. Make sure that the subject line of your email contains the phrase "Auction #13". (Simply using the "reply" command on most mail and news programs should work just fine.) If your message does not contain this text in the subject line, your message will be discarded.

2. In your email message, please put each of your bids on a separate line of text. Each bid line should be in the following format: the 3-digit identification number of the card you're bidding on, immediately followed by a right parenthesis, and then the amount of your bid in dollars and cents (such as 1.00). For example:

```
 305) 2.00
 306) 0.65
```

Including extra information on the bid line is okay, too. Such extra information might include email quote marks (such as the greater-than symbol), the card name, condition, etc. You may include anything that makes bidding easier for you, EXCEPT that your bid amount should be the only price-formatted number which appears on that line. This means, for example, that you should make sure to delete the previous high bid, so that your own bid is the only number containing a decimal point on that line.

For example, the following are also perfectly valid bids:

```
> 305) Ghosts of the Damned              U1     Blk   M  2.00
> 306) Demonic Torment   0.65
```

Bids that do not conform to these rules will be discarded.
Here are examples of INVALID bids:

```
>305) Ghosts of the Damned      $2                      [no decimal point]
306 ) Demonic Torment 0.65  [space between the 6 and the right parenthesis]
```

3. All bids must be in US currency, and must also satisfy the following:

   a. Bids under 1 dollar must be even multiples of a nickel ($0.05).
   b. Bids from 1 to 5 dollars must be even multiples of a dime ($0.10).
   c. Bids from 5 to 10 dollars must be even multiples of a quarter ($0.25).
   d. Bids from 10 to 20 dollars must be even multiples of a 50 cents ($0.50).
   e. Bids over 20 dollars must be even multiples of a dollar ($1.00).

4. If you wish me to keep your name confidential, please send me a nickname by which I can refer to you in updates when you are the highest bidder on a card. You need only do this for me once, and I will use the same nickname in subsequent auctions I hold. If you do not send me a nickname, I will identify you in updates by your name but not your full email address.

5. I will post an update of the high bids once per day to the newsgroup <rec.games.trading-cards.marketplace>. In addition, I will email a copy of this update to each participating bidder.

6. If a card's bid has not been raised in the past day, I will mark it Going Once (!) in the update. If it has not been raised in the past 2 days, I will mark it Going Twice (!!). If it has not been raised in the past 3 days, the card will be marked SOLD! In the case of tie bids, the winner will be the first person to have mailed the winning bid.

7. The winning bidder of each card will be notified by email, and will be asked to pay the amount due via US check or money order.

8. Payment of the winning bid will include free shipping within the United States, via first class mail. The cards will be wrapped in plastic sheaths and packed in cardboard for protection. All cards will be shipped within one week after receipt of payment.

9. While this is a real auction for real cards, you should know that I plan to use data on the bids in this auction for economic research. In no case will individual bidders be identified in this research; anonymity will be preserved. By bidding in this auction, you indicate your consent to have your bid be used in economic research. If you do not approve of this, you have the right not to participate in this auction. Should you have any questions or concerns about the use of data from this auction in academic research, please contact the chair of the COUHES committee at MIT by phone at 617-253-6787.

That's it! Enjoy the auction. Thanks for participating!

Here is the LIST OF CARDS. All are from the LEGENDS expansion set:

	Card Name	Rarity	Color	Cond	Bid	Status	Bidder
304)	Angelic Voices	R	Whi	M	14.50		michele
328)	Moat	R	Whi	M	17.50		michele
330)	Petra Sphinx	R	Whi	M	7.50		michele
338)	Spiritual Sanctuary	R	Whi	M	7.00	!	cdaveb
351)	Bartel Runeaxe	R	Gld	M	8.00		michele
371)	Livonya Silone	R	Gld	M	8.25	!	michele
374)	Nebuchadnezzar	R	Gld	M	13.00		michele
301)	Akron Legionnaire	R	Whi	M	7.50	SOLD!	mehl
302)	Alabaster Potion	C2	Whi	M	0.10	SOLD!	ristvan
303)	Amrou Kithkin	C2	Whi	M	0.05	SOLD!	jbb
305)	Cleanse	R	Whi	M	11.00	SOLD!	ent
306)	Clergy of the Holy Nimbus	C2	Whi	M	0.15	SOLD!	jbb
307)	D'Avenant Archer	C2	Whi	M	0.50	SOLD!	winston
308)	Divine Intervention	R	Whi	M	5.00	SOLD!	ent
309)	Divine Offering	C2	Whi	M	0.50	SOLD!	daleg
310)	Divine Transformation	R	Whi	M	5.50	SOLD!	gkearney
311)	Elder Land Wurm	R	Whi	M	5.00	SOLD!	snook
312)	Enchanted Being	C1	Whi	M	0.50	SOLD!	winston
313)	Equinox	C1	Whi	M	1.50	SOLD!	daleg
314)	Fortified Area	U1	Whi	M	0.80	SOLD!	ristvan

315)	Glyph of Life	C2	Whi	M	0.15	SOLD!	mabehr
316)	Great Defender	U1	Whi	M	1.00	SOLD!	hunt
317)	Great Wall	U1	Whi	M	0.50	SOLD!	ristvan
318)	Greater Realm of Preservat	U1	Whi	M	6.00	SOLD!	daleg
319)	Heaven's Gate	U1	Whi	M	1.50	SOLD!	gkearney
320)	Holy Day	C1	Whi	M	0.75	SOLD!	winston
321)	Indestructible Aura	C2	Whi	M	0.15	SOLD!	hunt
322)	Infinite Authority	R	Whi	M	7.25	SOLD!	michele
323)	Ivory Guardians	U1	Whi	M	2.10	SOLD!	ristvan
324)	Keepers of the Faith	C2	Whi	M	0.25	SOLD!	jbb
325)	Kismet	U1	Whi	M	3.50	SOLD!	daleg
326)	Land Tax	U1	Whi	M	5.50	SOLD!	daleg
327)	Lifeblood	R	Whi	M	9.00	SOLD!	mehl
329)	Osai Vultures	C1	Whi	M	0.25	SOLD!	hunt
331)	Presence of the Master	U1	Whi	M	3.00	SOLD!	gkearney
332)	Rapid Fire	R	Whi	M	5.00	SOLD!	ristvan
333)	Remove Enchantments	C1	Whi	M	2.00	SOLD!	daleg
334)	Righteous Avengers	U1	Whi	M	1.60	SOLD!	gkearney
335)	Seeker	U1	Whi	M	1.50	SOLD!	hunt
336)	Shield Wall	U1	Whi	M	1.00	SOLD!	ent
337)	Spirit Link	U1	Whi	M	5.00	SOLD!	daleg
339)	Thunder Spirit	R	Whi	M	19.50	SOLD!	michele
340)	Tundra Wolves	C2	Whi	M	0.10	SOLD!	ristvan
341)	Visions	U1	Whi	M	1.60	SOLD!	gkearney
342)	Wall of Caltrops	C1	Whi	M	0.20	SOLD!	ristvan
343)	Wall of Light	U1	Whi	M	1.60	SOLD!	ristvan
344)	White Common Set (15 cards)	C	Whi	M	3.80	SOLD!	jbb
345)	Adun Oakenshield	R	Gld	M	9.00	SOLD!	gkearney
346)	Angus Mackenzie	R	Gld	M	8.00	SOLD!	gkearney
347)	Arcades Sabboth	R	Gld	M	18.00	SOLD!	ent
348)	Axelrod Gunnarson	R	Gld	M	7.75	SOLD!	cdaveb
349)	Ayesha Tanaka	R	Gld	M	6.00	SOLD!	ent
350)	Barktooth Warbeard	U1	Gld	M	3.50	SOLD!	daleg
352)	Boris Devilboon	R	Gld	M	8.75	SOLD!	jbb
353)	Chromium	R	Gld	M	20.00	SOLD!	ent
354)	Dakkon Blackblade	R	Gld	M	21.00	SOLD!	cdaveb
355)	Gabriel Angelfire	R	Gld	M	9.00	SOLD!	cdaveb
356)	Gosta Dirk	R	Gld	M	6.50	SOLD!	ent
357)	Gwendlyn Di Corci	R	Gld	M	9.00	SOLD!	ent
358)	Halfdane	R	Gld	M	10.00	SOLD!	mehl
359)	Hazezon Tamar	R	Gld	M	10.00	SOLD!	ent
360)	Hunding Gjornersen	U1	Gld	M	3.50	SOLD!	daleg
361)	Jacques le Vert	R	Gld	M	8.25	SOLD!	inahsohn
362)	Jasmine Boreal	U1	Gld	M	3.00	SOLD!	winston

363)	Jedit Ojanen	U1	Gld	M	3.00	SOLD!	winston
364)	Jerrard of the Closed Fist	U1	Gld	M	3.00	SOLD!	winston
365)	Johan	R	Gld	M	15.00	SOLD!	mehl
366)	Kasimir the Lone Wolf	U1	Gld	M	3.00	SOLD!	winston
367)	Kei Takahashi	R	Gld	M	7.00	SOLD!	ent
368)	Lady Caleria	R	Gld	M	7.50	SOLD!	hunt
369)	Lady Evangela	R	Gld	M	7.00	SOLD!	ent
370)	Lady Orca	U1	Gld	M	3.00	SOLD!	winston
372)	Lord Magnus	U1	Gld	M	3.00	SOLD!	winston
373)	Marhault Elsdragon	U1	Gld	M	3.00	SOLD!	winston
375)	Nicol Bolas	R	Gld	M	21.00	SOLD!	gkearney
376)	Palladia-Mors	R	Gld	M	20.00	SOLD!	ent
377)	Pavel Maliki	U1	Gld	M	3.00	SOLD!	winston
378)	Princess Lucrezia	U1	Gld	M	3.00	SOLD!	winston
379)	Ragnar	R	Gld	M	7.00	SOLD!	ent
380)	Ramirez DePietro	U1	Gld	M	3.00	SOLD!	ent
381)	Ramses Overdark	R	Gld	M	9.50	SOLD!	michele
382)	Rasputin Dreamweaver	R	Gld	M	7.25	SOLD!	gkearney
383)	Riven Turnbull	U1	Gld	M	3.50	SOLD!	daleg
384)	Rohgahh of Kher Keep	R	Gld	M	8.50	SOLD!	cdaveb
385)	Rubinia Soulsinger	R	Gld	M	13.00	SOLD!	cdaveb
386)	Sir Shandlar of Eberyn	U1	Gld	M	3.50	SOLD!	daleg
387)	Sivitri Scarzam	U1	Gld	M	8.00	SOLD!	winston
388)	Sol'kanar the Swamp King	R	Gld	M	15.00	SOLD!	matthew.sojka
389)	Stangg	R	Gld	M	8.00	SOLD!	mehl
390)	Sunastian Falconer	U1	Gld	M	4.00	SOLD!	winston
391)	Tetsuo Umezawa	R	Gld	M	9.50	SOLD!	tfgiord
392)	The Lady of the Mountain	U1	Gld	M	3.50	SOLD!	daleg
393)	Tobias Andrion	U1	Gld	M	3.50	SOLD!	daleg
394)	Tor Wauki	U1	Gld	M	3.50	SOLD!	daleg
395)	Torsten von Ursus	U1	Gld	M	3.50	SOLD!	daleg
396)	Tuknir Deathlock	R	Gld	M	8.00	SOLD!	hutch
397)	Ur-Drago	R	Gld	M	6.00	SOLD!	ent
398)	Vaevictis Asmadi	R	Gld	NM	21.00	SOLD!	jbb
399)	Xira Arien	R	Gld	M	8.00	SOLD!	cdaveb

END OF LIST

Key to Card Types and Conditions:

C = Common
C1 = Common (printed once per sheet of commons)
C2 = Common (twice as common as C1, printed twice per common sheet)

U = Uncommon
U1 = Uncommon (printed once per sheet of uncommons)
U2 = Uncommon (twice as common as U1)
R = Rare

Lnd = Land
Art = Artifact
Blk = Black
Blu = Blue
Gre = Green
Red = Red   :)
Whi = White

M = Mint
NM = Near Mint (never played, but has tiny blemishes from handling)
E = Excellent (played a few times, has small scuff marks)

If you are unfamiliar with some of these cards, you can get information about any Magic card (spell type, power, toughness, artist's name, etc.) from the following Internet sources:

http://www.itis.com:80/deckmaster/magic/cardinfo/
ftp://ftp.itis.com/pub/deckmaster/card.info/lists.w.spoilers/

Remember to send any bids, comments, or questions about the cards or the rules of this auction to <reiley@mit.edu>.

Thanks!

Chapter 3:

The Roles of Marketing, Product Quality, and Price Competition in the Growth and Composition of the U.S. Anti-Ulcer Drug Industry

(co-authored with Ernst R. Berndt, Linda T. Bui, and Glen L. Urban)

1 Introduction

The introduction of Tagamet into the U.S. market in 1977 marked the beginning of a revolutionary treatment for ulcers, and the emergence of a new industry. What distinguished the products of this new industry was their ability to heal ulcers and treat pre-ulcer conditions pharmacologically on an outpatient basis, thereby substituting for traditional, and costly, hospital admissions and surgeries. Tagamet, known medically as an H_2-receptor antagonist, promoted the healing of ulcers by reducing the secretion of acid by the stomach.

A striking feature of the anti-ulcer market is that it has sustained growth in sales (quantity, not just revenue) for over fifteen years, and still shows no sign of slowing. New prescribing habits have clearly diffused to an ever-increasing number of physicians. Today there are a total of four H_2-receptor antagonists: Tagamet, Zantac, Pepcid and Axid. Zantac is now the United States' (and the world's) largest selling prescription drug, having estimated worldwide sales in 1992 of about \$3.5 billion. Moreover, both Zantac and Tagamet are among the ten top selling prescription drugs in the U.S.[1]

In this paper we attempt to explain the growth and changing composition of the anti-ulcer drug market. Although we examine the impacts of pricing and product quality, we devote particular attention to the role of firms' marketing efforts. We distinguish between two types of marketing: (i) that which concentrates on bringing new consumers into the market ("industry-expanding" advertising), and (ii) that which concentrates on competing for market shares from these consumers ("rivalrous" advertising). Note that of these two types, the market-expanding advertising has particular economic importance in a new market, because no matter how potentially beneficial is the new

[1] *MedAdNews*, "One Hundred Powerhouse Drugs", Special Supplement, May 1993, Vol. 12, No. 6. Incidentally, Tagamet is seventh, Pepcid ranks 17, Prilosec 25 and Axid 61 in terms of US sales. In terms of world sales, Tagamet is 7, Pepcid is 22, Prilosec is 49 and Axid is 67.

product, it can generate no consumer surplus until consumers have been informed about the new product and have been induced to experiment with it.

As others have done, we estimate the effects of industry-expanding advertising on sales. However, we also examine how the effectiveness of this socially beneficial type of advertising varies with market structure. We exploit two facts. First, in the earliest years of the market when Tagamet was a monopoly product, by definition, all of the Tagamet advertising was market-enhancing. Second, the timing of entry is largely exogenous in this industry, for patent protection ensures that firms cannot enter until their research laboratories develop a new molecule that has the desired impact and approval for use is given by the Food and Drug Administration (FDA).

We also analyze factors affecting the market shares earned by the limited number of firms in this market. A principal theme is that the patent and pioneer advantages to Tagamet were overcome by Zantac, the second entrant, through costly but effective marketing efforts, especially efforts that interacted with the apparent existence of favorable adverse drug interaction profiles relative to Tagamet. Moreover, Zantac's relative price, although higher than Tagamet's, declined substantially over time. Thus, evidence from this industry suggests that while the barriers to entry from patent and first-mover advantages are considerable, they are not insurmountable.

Our empirical analysis is based on an unusually rich and detailed data set. Beginning with the introduction of Tagamet in July 1977, we have obtained monthly data, for each of the products in this market, on quantity and average price of sales (separately for the retail drug and hospital markets), marketing efforts (minutes of detailing by sales representatives to physicians, and professional medical journal advertising), as well as product quality information including side effect profiles, adverse drug interactions, efficacy, dosage forms, and indications for which the product had received approval from the U.S. Food and Drug Administration.

We begin in Section 2 by providing background information on ulcers and ulcer treatments. Then in Section 3 we present an overview of data trends. We describe the growth of the anti-ulcer market, as well as the pricing and marketing behavior of the various market participants. We move on in Section 4 to develop an econometric framework for modeling the growth of the anti-ulcer industry. In particular, we examine the effects of "informative" or market-expanding marketing efforts on industry sales. In Section 5 we report findings from an analogous attempt to model factors affecting market shares earned by the various products in this industry. Here we examine in particular the roles of rivalrous marketing, product quality, order of entry and price competition. Finally, in Section 6 we offer some concluding observations and suggestions for future research. The paper also includes a data appendix.

2 Background on Ulcer Treatments[2]

Peptic ulcer disease occurs in 10-15% of the U.S. population. Ulcers located in the stomach proper are termed gastric ulcers (GU), while those in the duodenum (the bulb connecting the stomach to the small intestine) are called duodenal ulcers (DU). A related non-ulcerous condition is gastroesophageal reflux disease (GERD), which occurs in the esophagus. What the three conditions have in common is that they involve inflammation of tissue in the digestive tract that is exacerbated by the presence of the body's naturally occurring gastric acid. GERD and duodenal ulcers have roughly the same rates of occurrence in the U.S. population, whereas gastric ulcers are about one-fourth as likely. The incidence of ulcers in adult males is about twice that in adult females, and appears to be most common in individuals twenty to fifty years old.

Ulcers have a long history of clinical treatment. There is evidence that already in the first century A.D., coral powder (calcium carbonate, an antacid) was used to relieve symptoms of dyspepsia.[3]Early in the 20th century, conventional medical wisdom conformed to the notion "no acid, no ulcer." As a result, until the 1970's recommended treatments sought to neutralize gastric acid, and often consisted of hourly feedings of milk and/or antacids, as well as a dietary reduction of acidic food and drink. If ulcers persisted, surgery was undertaken. It is worth noting that while antacids such as Maalox and Mylanta neutralize gastric acid, they do not decrease the rate of gastric secretions (they may in fact increase them). Moreover, the required dosages of antacids are typically quite large, side effects can be considerable, and adverse interactions with other drugs are not uncommon. As a result, with antacids patient compliance can be problematic.

An alternative ulcer treatment involves acid suppression with anticholinergics, such as Pro-Banthine and Atropine. Anticholinergic agents decrease acid secretion by inhibiting receptors for the hormone acetylcholine in the acid-producing cells of the stomach lining. However, these agents have considerable unpleasant and adverse reactions, since acetylcholine is involved in a number of biochemical processes other than the secretion of gastric acid, and anticholinergics tend to be nonselective. The side effects of dry mouth, blurred vision, urinary retention, abnormally rapid heartbeat, and drying of bronchial secretions are particularly frequent.

In 1977 a revolutionary form of anti-ulcer drugs was introduced in the U.S., known as an H_2-receptor antagonist.[4]The H_2-receptor antagonists act by blocking the histamine-2 (H_2) receptor on parietal cells in the lining of the stomach -- cells that produce gastric acid. Histamine-2 is one of three "messenger molecules" (along with gastrine and acetylcholine) that can stimulate the

---

[2] The material that follows is taken in large part from Scouler [1993] and the references cited therein. Also see Fine et al. [1988] and McKenzie et al. [1990].

[3] See Fine et al. [1988].

[4] Tagamet was introduced into the U.K. one year earlier, in 1976.

production of acid by the parietal cells. By blocking the receptor for H_2 (and, unlike the anticholinergic drugs, avoiding any interference with other biochemical processes), an H_2-antagonist (henceforth referred to as an H_2) can decrease overall acid concentration in the stomach. H_2-antagonist healing rates are very high. A four- to six-week treatment period, for example, is associated with a healing rate of 70-80% for patients suffering from a duodenal ulcer.

SmithKline was the first pharmaceutical company to introduce an H_2-antagonist into the U.S. market (August 1977), and they dubbed it Tagamet (its chemical name is cimetidine). Thereafter three companies followed -- Glaxo with Zantac (ranitidine) in June 1983, Merck with Pepcid (famotidine) in October 1986, and Lilly with Axid (nizatidine) in April 1988. Each of these four H_2-antagonists is a slightly different chemical entity; Tagamet's patent protection could not prevent entry by such therapeutic substitutes.

Zantac was marketed very aggressively by Glaxo, in partnership with Hoffmann-LaRoche, and was also priced at a premium over Tagamet. Detailers (sales representatives who call on physicians) emphasized that unlike Tagamet whose original dosage required it to be taken four times daily, Zantac needed to be taken only twice per day. Moreover, Zantac detailers highlighted side effect profiles that had accumulated with Tagamet -- nausea, diarrhea, drowsiness, decreased sperm count, gynecomastia (swelling of the breasts in males) and adverse drug interactions.[5] Within eighteen months Tagamet responded to Zantac by introducing a twice per day version of its drug, but it continued to find itself on the defensive in terms of alleged side effect and adverse interaction profiles. A prolonged rivalry then ensued, first between Tagamet and Zantac in the form of new versions whose dosages were but once per day (thereby facilitating patient compliance even further), and later including additional competition from the newly entered Pepcid and Axid, each available with a once-daily dosage regimen.

In addition to side effect profiles and frequency of dosage, another form of rivalry among the four H_2-antagonists involved FDA-approved treatments (indications). Since several distinct types of ulcerous conditions exist, similar drug products can compete on the basis of efficacy for different indications. In the U.S., before a drug can be introduced into the market, the Food and Drug Administration must grant approval for at least one indication. When Tagamet was originally introduced into the U.S. market in August 1977, its approval was for duodenal ulcers; Tagamet was also the first to be approved for duodenal ulcer maintenance treatment (to prevent recurrence of a newly-healed duodenal ulcer, in April 1980) and gastric ulcers (December 1982). However, Zantac was the first to obtain approval for the GERD indication (May 1986),[6]and it was not until March 1991 that Tagamet obtained FDA approval for GERD. It is worth noting that, once FDA approval

[5] By June 1983, Tagamet had registered ten adverse interactions at the FDA. Zantac recorded its first adverse interaction in January 1992.

[6] Discussions with industry officials suggest that Glaxo actually invented the GERD indication at the FDA.

for an indication is granted, the manufacturer is permitted to provide promotional and marketing material *only* for approved indications. Thus, even though Tagamet had very similar clinical effects to Zantac, suggesting that it would probably be effective in the treatment of GERD, Tagamet promotions were not permitted to mention GERD until 1991. Although physicians often prescribe drugs for indications not approved by the FDA (called off-label prescribing), not having FDA approval for an indication which is held by a competitive product may constitute a significant disadvantage in the marketplace. Hence, even though Tagamet pioneered in the three anti-ulcer indications, that it lagged Zantac in the relatively populous GERD market was of considerable importance.

Today the four H_2-antagonist drugs are frequently viewed as being "...equally efficacious in their ability to suppress acid secretion,"[7] but different in their pharmacological profiles. McKenzie et al. [1990, p. 58] note that Tagamet is "the H_2-antagonist implicated with the most side effects and drug interactions," and that such adverse impacts occur "to a lesser extent" with Zantac. The third and fourth entrants -- Pepcid and Axid -- appear to have even less drug interactions and side effects.[8] What is not yet clear, however, is the extent to which apparent differences in side effect profiles simply reflect differential lengths of time over which the various drugs have been able to accumulate medical experience.

Modern ulcer medicines are not necessarily restricted to H_2-antagonists. One alternative therapy is Carafate (sucralfate), introduced into the U.S. by Marion Labs in August 1981. Instead of inhibiting acid secretion, Carafate acts by forming a protective coating over the ulcer that in turn promotes healing. While it is relatively free from side effects, Carafate has problems of convenience and compliance, since it must be taken four times per day, always on an empty stomach (before meals). It also acts more slowly than the acid inhibitors in relieving pain. For these reasons, Carafate serves a market niche, being used predominantly for older patients, and patients in intensive care.

Another entrant in the anti-ulcer market is Cytotec (misoprostol), introduced in December 1988. Cytotec has been targeted at ulcers associated with the use of non-steroidal anti-inflammatory drug (NSAID) therapy (pain relievers, such as Motrin). Its rather small market niche consists of patients who take NSAIDs chronically and are at greater risk for the development of peptic ulcer disease, or complications from peptic ulcers -- particularly the elderly, those with previous ulcers, concomitant debilitating diseases, and/or patients who smoke. A common side effect of Cytotec, however, is diarrhea, although it can often be mitigated by adjusting dosage.

The most recent treatment innovation to enter the anti-ulcer market is Prilosec (omeprazole), introduced into the U.S. by Merck, Sharp and Dohme in September 1989.[9] Prilosec is a powerful

---

7 McKenzie et al. [1990], p. 58.

8 *Ibid.*

new drug known as a proton pump inhibitor. It acts by directly blocking the action of the proton pump, which is the biochemical mechanism that actually produces the acid in the stomach. Initially approved for only the GERD indication, in June 1991 Prilosec was approved by the FDA for duodenal ulcer treatment. Although it is considerably more potent than the H_2-antagonists, there is some evidence from long-term studies on rats that Prilosec is associated with carcinoid tumors; hence it is approved only for short-term use. Dosing for Prilosec is unique in that it is supplied in a timed-release capsule, thus reducing dosage to once per day but yielding continuous levels of the drug within the body throughout the day.

With this brief overview on ulcer drugs and ulcer treatments as background, we now move on to a discussion of the pricing and marketing behavior of the manufacturers, the sales and market shares they attained, and the data sources underlying these statistics.

3 Overview of the Data

Most of the data used in this study originated with IMS America, a Philadelphia-based firm that independently collects data on the sales and marketing of pharmaceutical products. IMS sells its data to pharmaceutical manufacturers, among others, for their use in formulating marketing strategy.[10] IMS sales data track prescription pharmaceutical purchases made by hospitals and by retailers; market segments not monitored by IMS include food stores, dispensing physicians, HMO's, mail order, nursing homes, and clinics. IMS estimates that its drugstore audit covers 67% of the U.S. pharmaceutical market, and that its hospital audit encompasses an additional 16%.[11]

The level of aggregation of the IMS purchase data is at the presentational form, e.g., bottles of 30 tablets of a 150mg pill. For each presentational form, we compute average price as dollar purchases divided by number of units. We also convert these price and quantity measures into patient days and price per patient day, using the recommended daily dosage for duodenal ulcer treatment as the transformation factor. These monthly data series begin in August 1977 and continue through May 1993.

In addition to price and quantity data on drug purchases, we employ IMS data on marketing efforts from their Personal Selling Audit, earlier called the IMS National Detailing Audit. Based on a panel of about 3500 physicians who report the number of visits and minutes spent with detailers discussing particular drug products, IMS computes monthly detailing efforts by drug.[12] Using an estimated cost per detailing visit, IMS also estimates total detailing expenditures. Medical jour-

[9] Merck obtained the rights to market Prilosec in the US from AB Astra of Sweden. Prilosec was originally named Losec; however, its name was changed because of confusion surrounding the similarity of the name Losec to that of Lasix, a common diuretic.

[10] IMS America, 660 W. Germantown Pike, Plymouth Meeting, Pennsylvania 19462 (215-834-5000).

[11] Information on IMS is taken from the *IMS Pharmaceutical Database Manual.*

nal advertising expenditures are estimated by IMS in their National Journal Audit. Based on the number of square inches and pages of advertisements in about 300 major medical journals, as well as features such as the number of colors of each advertisement, IMS uses standard rate sheets to estimate total dollars of journal advertising, monthly, by product. We convert these current dollar expenditures into constant dollar magnitudes using the BLS Producer Price Index for Advertising in Professional Journals.

Discussions with industry personnel suggest that while these detailing and journal advertising expenditures likely understate total promotion costs (booths and promotions at conferences are not included, for example), there is no reason to suspect that the proportions differ across products, and thus we are led to believe that the relative expenditure data series are likely to be reasonably accurate. It is worth noting, incidentally, that according to one observer, in the early 1990's in the U.S. pharmaceutical industry, approximately \$3.1 billion was spent on detailing, about \$700 million was spent annually on journal advertising and direct-mail promotions, medical education expenses accounted for about \$400 million, and other forms of media and communication amounted to approximately \$300 million annually.[13]

Finally, data on recommended daily dosages and product-specific attribute information are taken from *Physician's Desk Reference*, annual issues from 1978 to 1993, and *U.S. Pharmacopeia Convention, Dispensing Information*. Further details regarding data sources and transformations are presented in the Data Appendix.

With this as background regarding data sources, we now present an overview of data trends. In Figure 1 we plot the quantity of U.S. sales (number of patient days of therapy) over time, separately for the retail drugstore and hospital markets, disaggregated into the H_2-antagonist (Tagamet, Zantac, Pepcid and Axid) and other anti-ulcer drugs (Carafate, Cytotec and Prilosec). Starting from zero in August 1977, by May of 1993 total monthly sales were almost 130 million patient days; of this, approximately 93% is sold via retail drug hospitals. Broken down by drug type, the H_2-antagonist class accounts for approximately 84% of total sales, while the other anti-ulcer drugs make up the remaining 16%. Because of this market dominance, hereafter we confine our analysis to the H_2-antagonist drugstore market.

The growth of H_2-antagonist sales over time has been remarkably steady. For example, if one runs a simple regression of log sales on a constant and a time counter, one obtains:

$$ln(Q_{H2}) = 16.4 + 0.012t, \qquad R^2 = 0.82,$$

implying an average annual growth rate (AAGR) of about 15%.

---

12 This sample size has increased with time. The 3500 number refers to 1993. In the mid 1980's, the sample size was about 2800.

13 Cearnal [1992], p. 23.

Figure 1. H2 Antagonist and Total Industry Drug Store and Hospital Sales

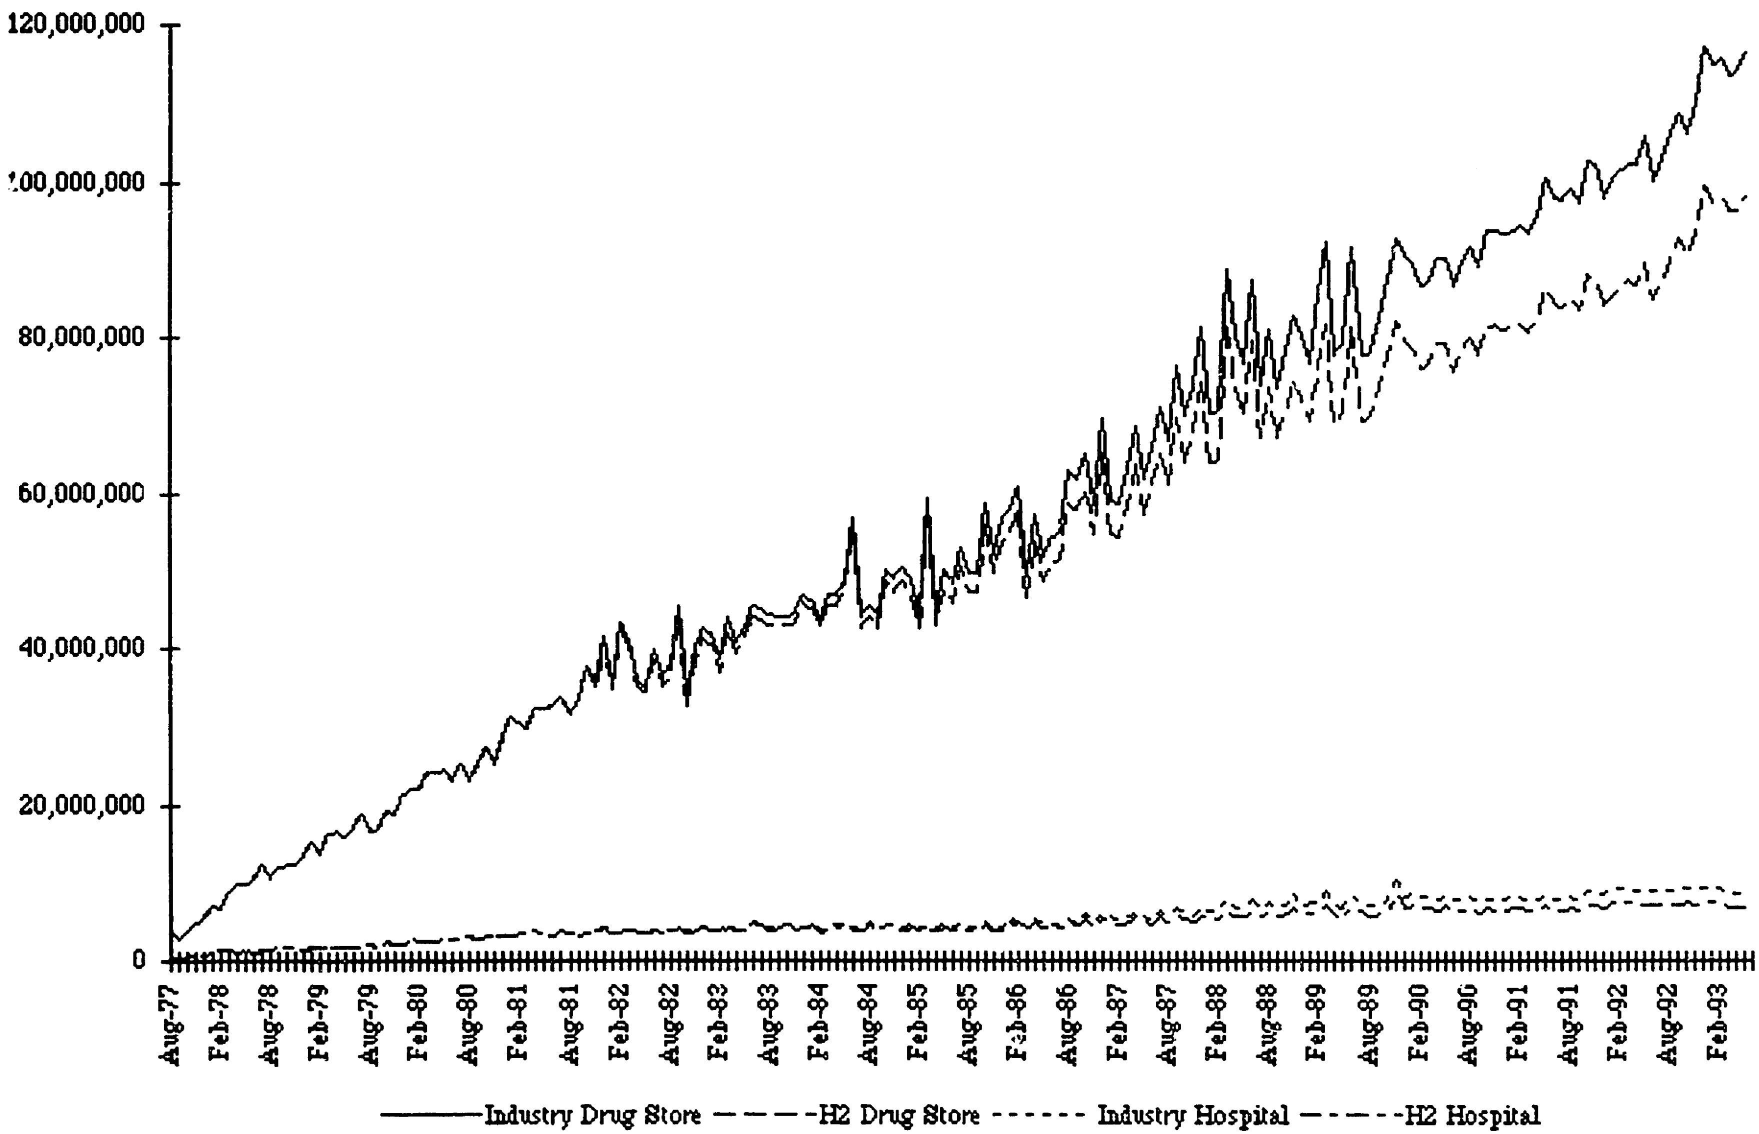

In Figure 2 we plot market shares of H_2-antagonist drugstore sales for the four H_2-antagonist drugs. Although Tagamet is the pioneer, Zantac enters in July 1983, and within one year it already captures about 25% of the total Tagamet-Zantac market. Tagamet's share continues to decline when Pepcid enters in October 1986, but Pepcid is less successful than Zantac; one year after entry, Pepcid has only approximately 8% market share. The sales of Zantac grow remarkably quickly and steadily, and by January 1988 Zantac sales overtake those of Tagamet. At about the same time (April 1988), Axid enters the market; as fourth entrant, however, Axid faces considerable competition, and after one year, its sales account for about only a 4% market share. By the end of our sample in May 1993, Zantac has captured about 55% of the quantity market share, Tagamet 21%, Pepcid 15%, and Axid 9%.

Although the entry of Zantac into the H_2-antagonist market increased total market sales, the sales of Tagamet fell. As seen in Figure 3, drugstore sales of Tagamet grew at a very rapid rate after entry in 1977, they began to level off a bit from 1981 to 1983, and although they peaked at about 46 million patient days in April 1984, after Zantac's entry in 1983 Tagamet's sales tended to decline. This general decline in sales continues to the end of our sample, when Tagamet monthly sales are less than half their peak -- about 21 million patient days. By contrast, sales of Zantac have generally increased over time, and by May 1993 Zantac accounted for about 54 million patient days per month. Although Zantac sales increase with time, as is seen in Figure 3, there is a modest decline in the growth slope beginning early 1988, coinciding with a slight rebound in Tagamet sales and the effects of entry by the fourth entrant, Axid. Although both Pepcid and Axid record considerable growth in sales, they clearly are dominated by the two earliest entrants, Tagamet and Zantac.

An interesting phenomenon occurs in the pricing behavior of the four products over this tumultuous time period. Price per day of duodenal therapy (based on recommended dosages, and adjusted for inflation using the overall Consumer Price Index with 1982-84 = 1.00) is displayed for the four products in Figure 4. After original entry until it faced competition from Zantac, Tagamet gradually decreased its real price from about \$1.00 to about \$.80 per day. When Zantac entered in late 1983, it charged a substantial premium (\$1.25 per day, a 56% premium). Thereafter, prices of <u>both</u> Zantac and Tagamet rose with time, although Tagamet's prices increased more rapidly. By the end of the sample, the Zantac price premium had narrowed from about 56% to 25%.

The third and fourth entrants, Pepcid and Axid, followed price policies that fell generally somewhere between that of Tagamet and Zantac. At the end of the sample period covered by our data (May 1993), the price per day of therapy ranged from a low of about \$1.41 per day for Pepcid, \$1.44 for Tagamet and \$1.62 for Axid, to a high of \$1.80 per day for Zantac. An interesting recent development is that in November 1993 (after the end of our sample), Tagamet announced a major change in its pricing policy, offering rebates directly to consumers.[14]

Finally, as is seen in Figure 4, there does not appear to be any substantial competitive pricing policy response by incumbents to the entry of new competitors in the H_2-antagonist market. Indeed, the only price trend break that coincides with entry is that for Tagamet upon entry by Zantac, which resulted in the incumbent Tagamet increasing rather than decreasing its price.[15] Note also that price trends do not show breaks around the time of entry by Pepcid and Axid.

Pricing policy, however, is not the only instrument for competitive rivals. In the U.S. pharmaceutical industry, marketing plays a very significant role. In Figure 5 we plot monthly detailing minutes for the two principal rivals, Tagamet and Zantac; cumulative detailing minutes since original product launch are plotted for each H_2-antagonist drug in Figure 6.

As seen in Figure 5, the launch of Tagamet coincided with a very substantial detailing effort --about 180,000 minutes in September 1977, after which detailing efforts gradually diminished. High levels of Tagamet detailing occurred in mid-1980 and early 1983, apparently in response to Tagamet receiving FDA approval for the new indications of duodenal ulcer maintenance (April 1980) and gastric ulcer therapy (December 1982). When Zantac entered with a very aggressive detailing effort in July 1983 (over 350,000 minutes), Tagamet responded with about a 50% increase in its own detailing efforts. More detailing peaks for both Tagamet and Zantac occurred in 1986, a year in which Pepcid entered and Zantac obtained FDA approval for the treatment of GERD. Both Tagamet and Zantac appear to have anticipated the entry of Axid in April 1988 by increasing their detailing in February 1988 (substantially for Tagamet, more modestly by Zantac), but both detailing levels declined again after Axid's entry.

Although month-to-month variations are apparent in Figure 5, there are definite trends in the intense Zantac-Tagamet detailing rivalry. As is seen in Figure 6 where cumulative detailing minutes are plotted for all four products, over its entire life Tagamet has out-detailed Zantac. However, in terms of detailing minutes per year, Zantac has considerably outpaced Tagamet. In part, Zantac has been able to do this because it has had two sales forces resulting from Glaxo's co-marketing agreement in the U.S. with Hoffmann-LaRoche. In terms of cumulative minutes of detailing through the end of our sample, the relative magnitudes are for every one minute of Axid detailing, there have been 3.21 minutes of detailing for Tagamet, 2.60 minutes for Zantac, and 0.88 for Pepcid.

According to Bond and Lean [1977], one way in which pioneering advantages occur in the pharmaceutical industry is in the effectiveness of advertising. Bond and Lean argue that to convince physicians to switch from an existing drug to a new one and thereby to overcome advantages

---

[14] See *New York Times* [November 9, 1993].

[15] For a discussion of the possible social welfare impacts of a pioneer raising its price in response to the introduction of a competitive product by a second entrant, see Perloff and Suslow [1994]. Related literature is found in Bresnahan-Reiss [1990], Cocks [1975], Cocks-Virts [1974], and Reekie [1978].

Figure 2. Drug Store Market Shares

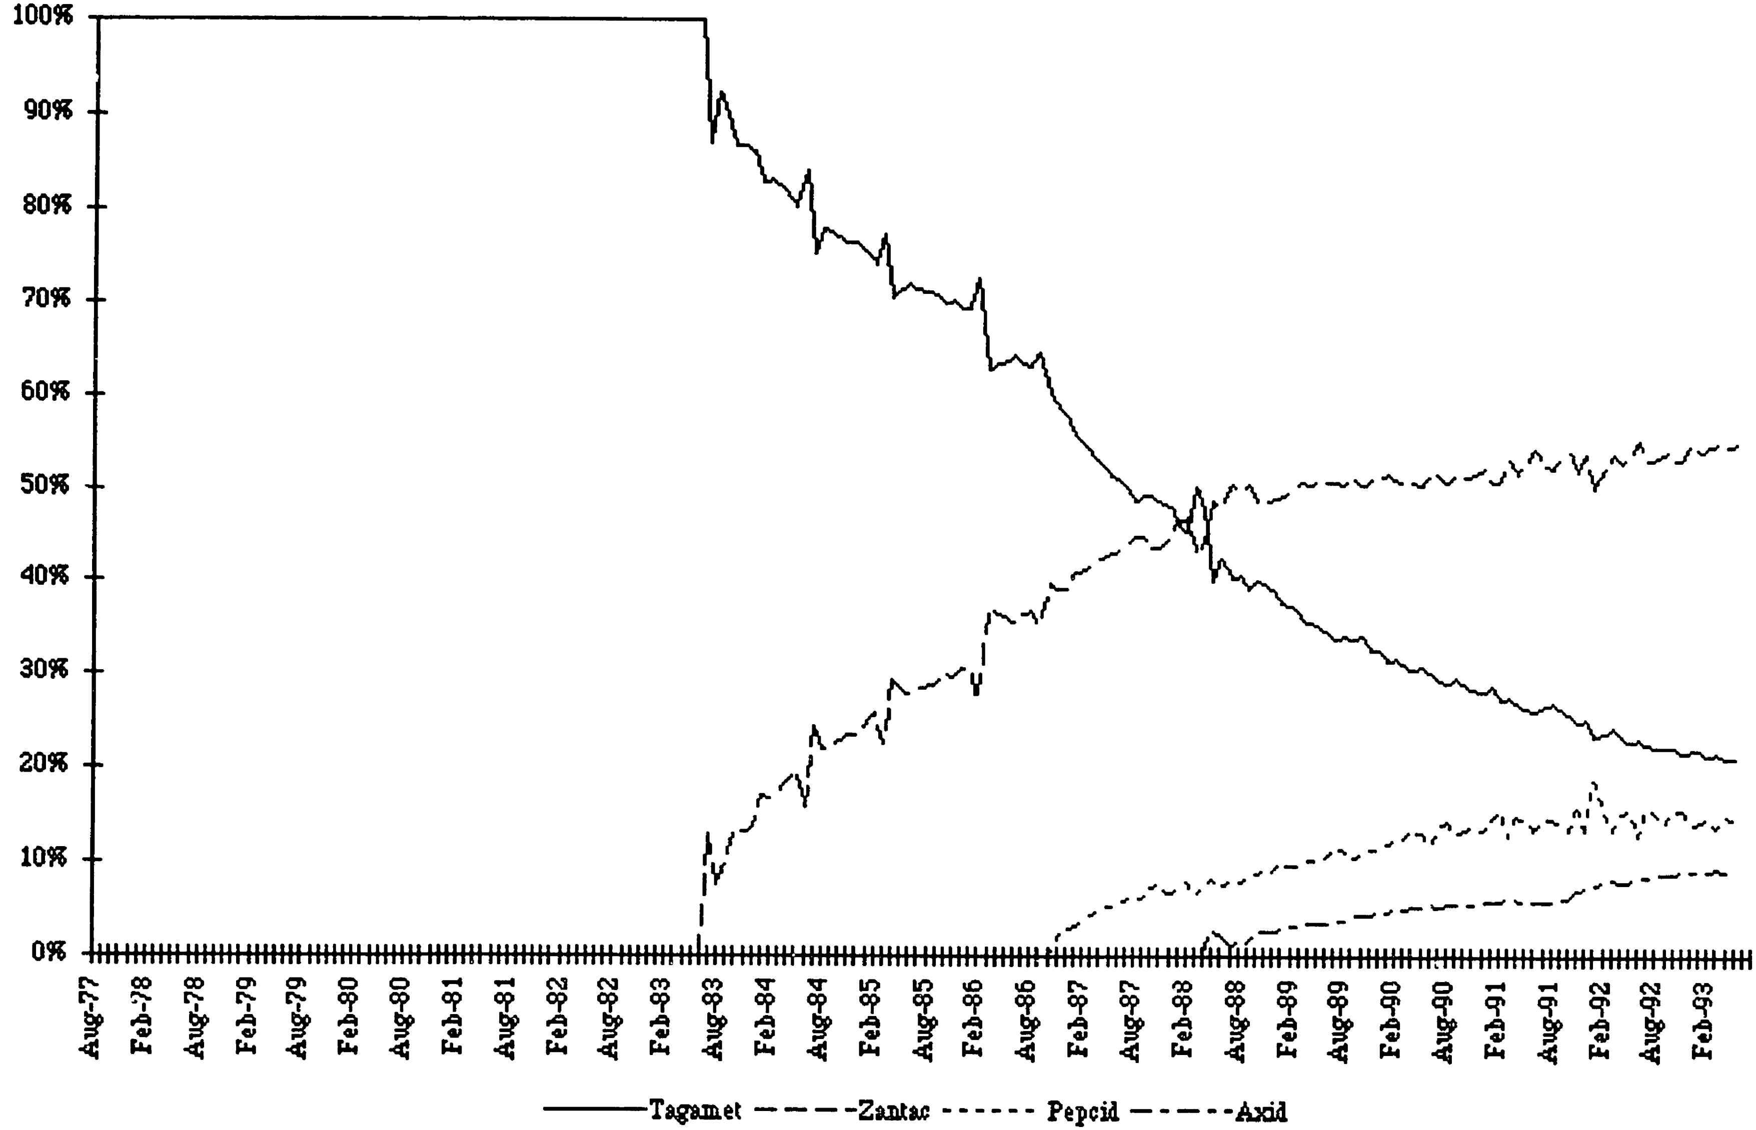

Figure 3. H2 Antagonist Drug Store Sales

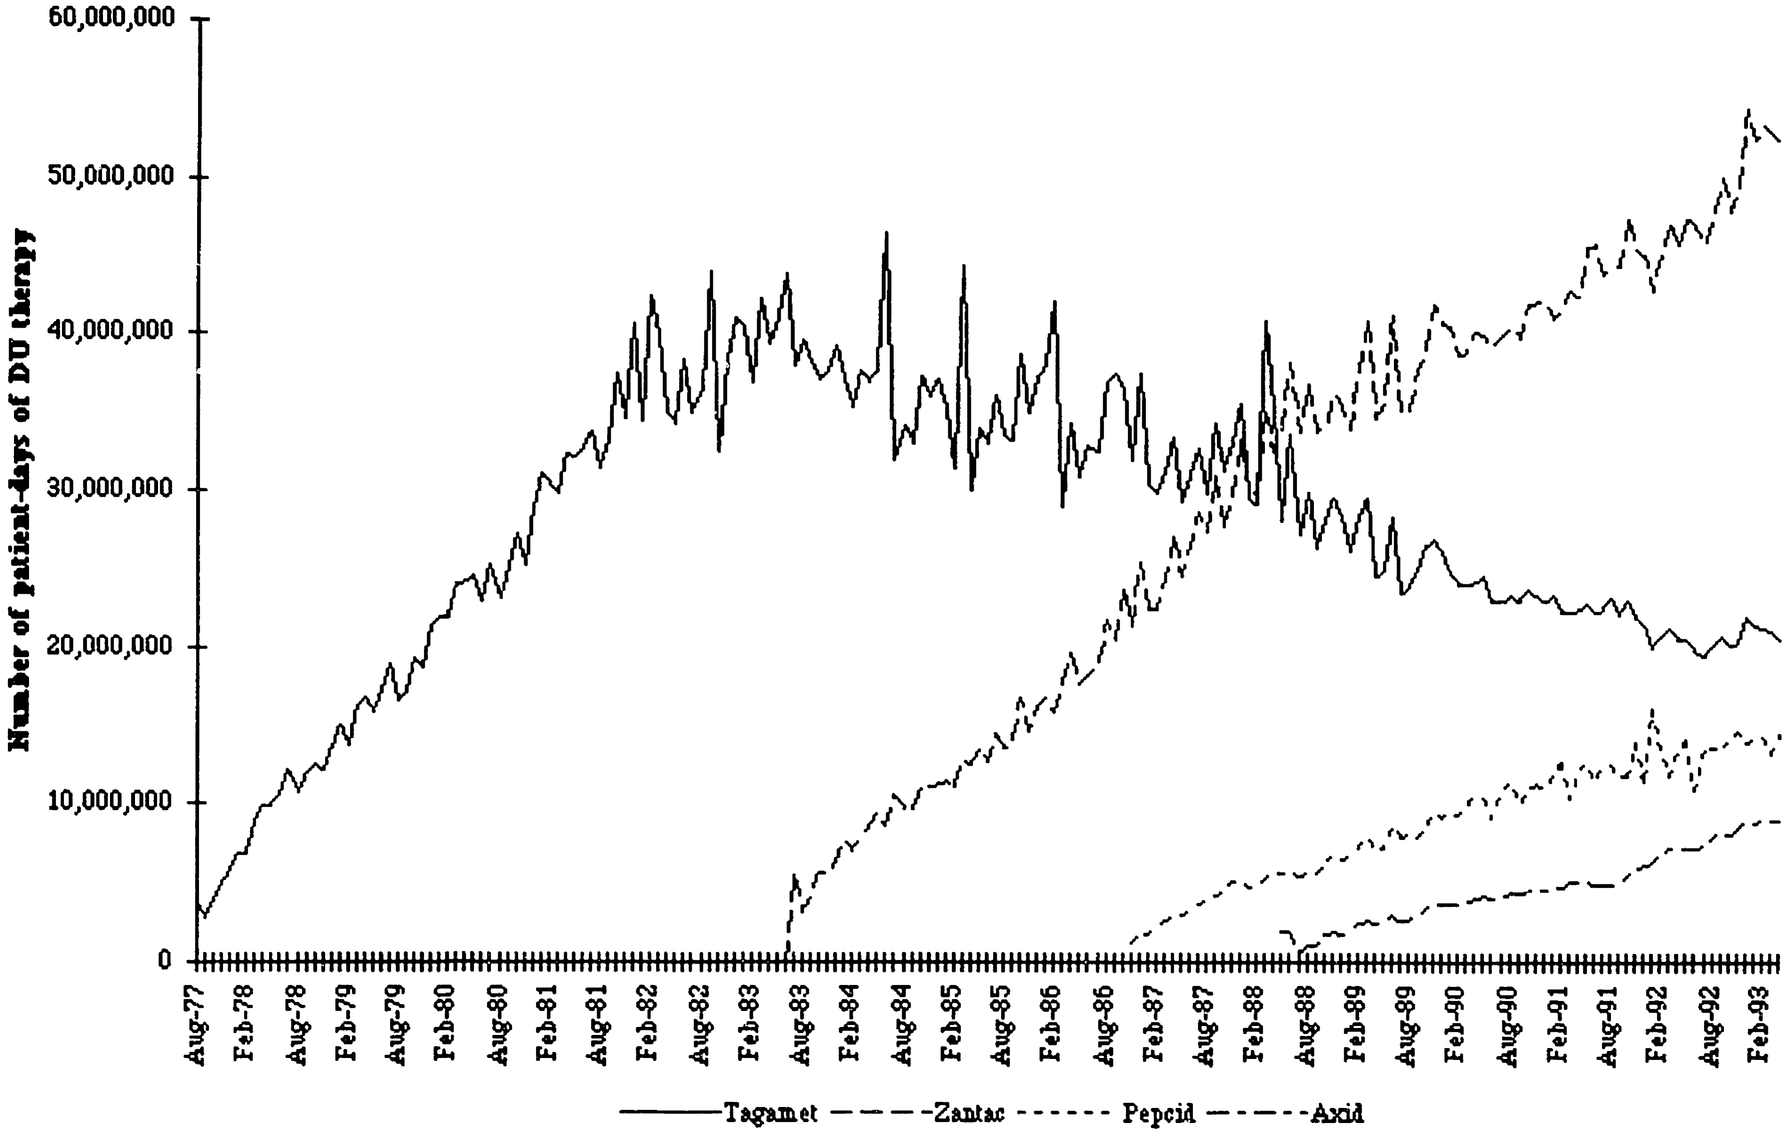

Figure 4. Real Drug Store Prices

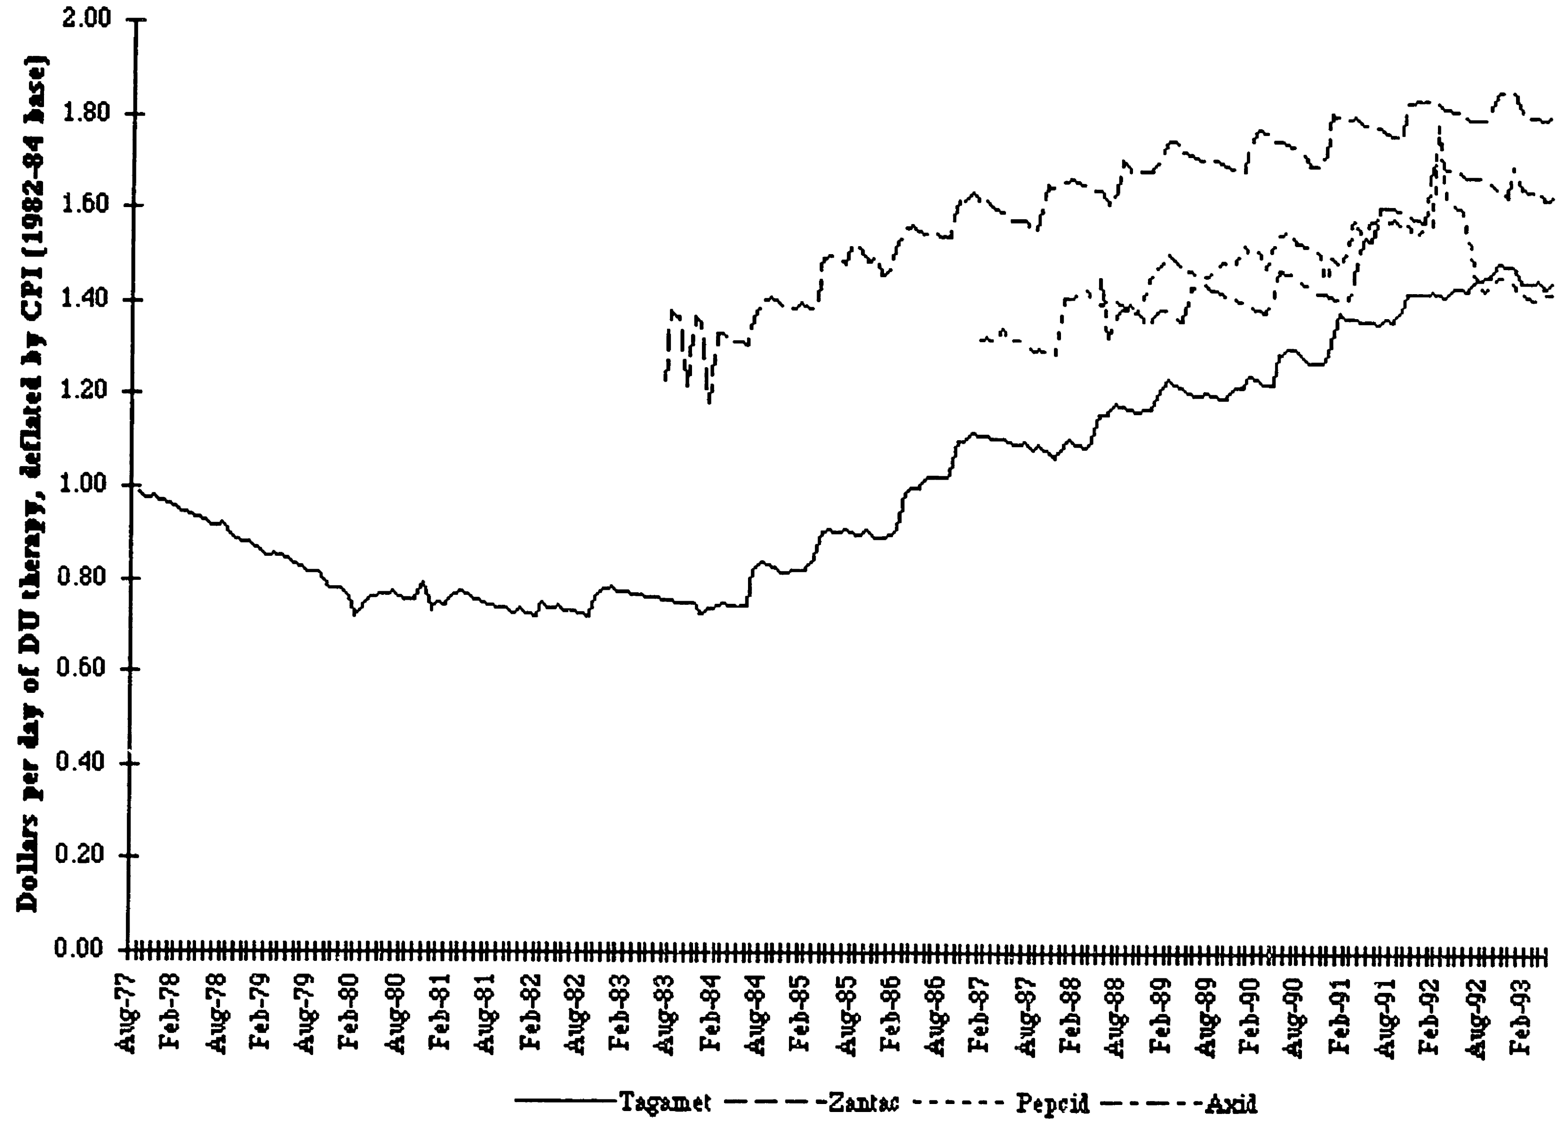

Figure 5. Tagamet and Zantac Detailing Minutes

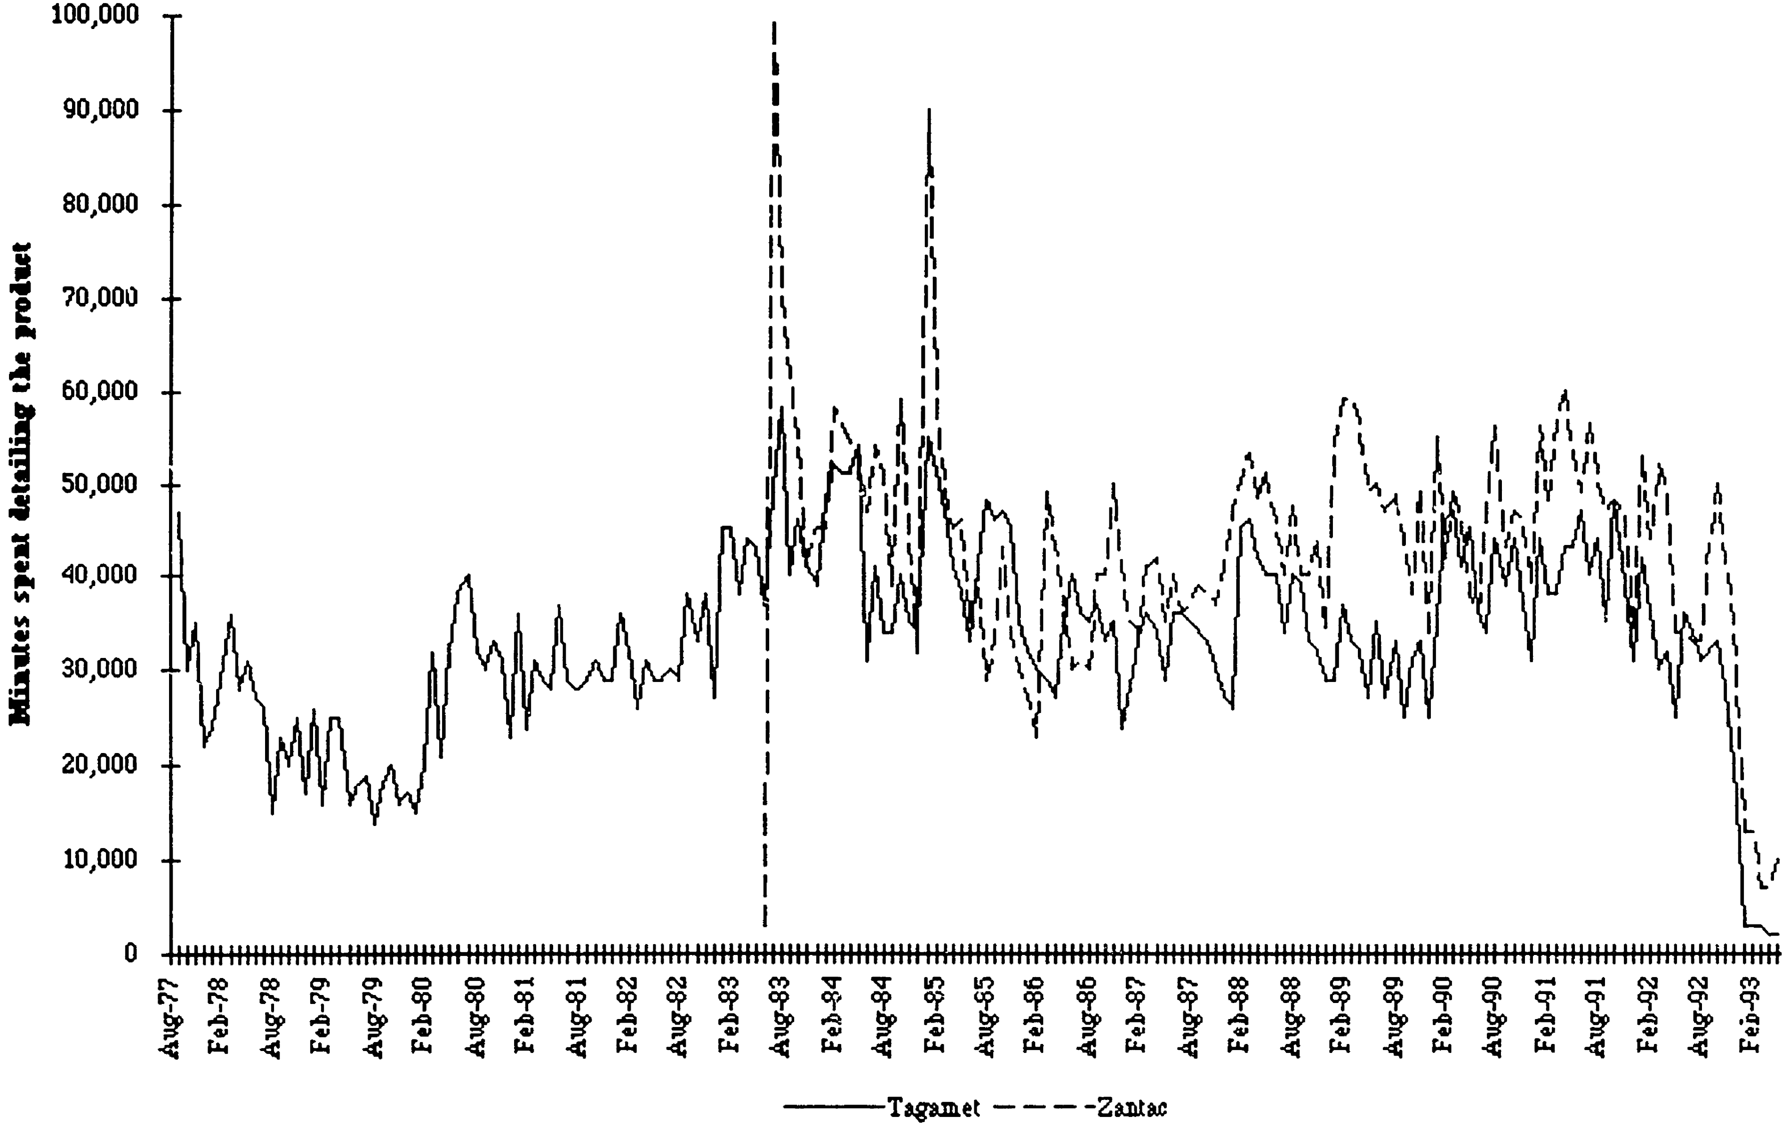

Figure 6. Cumulative Minutes of Detailing

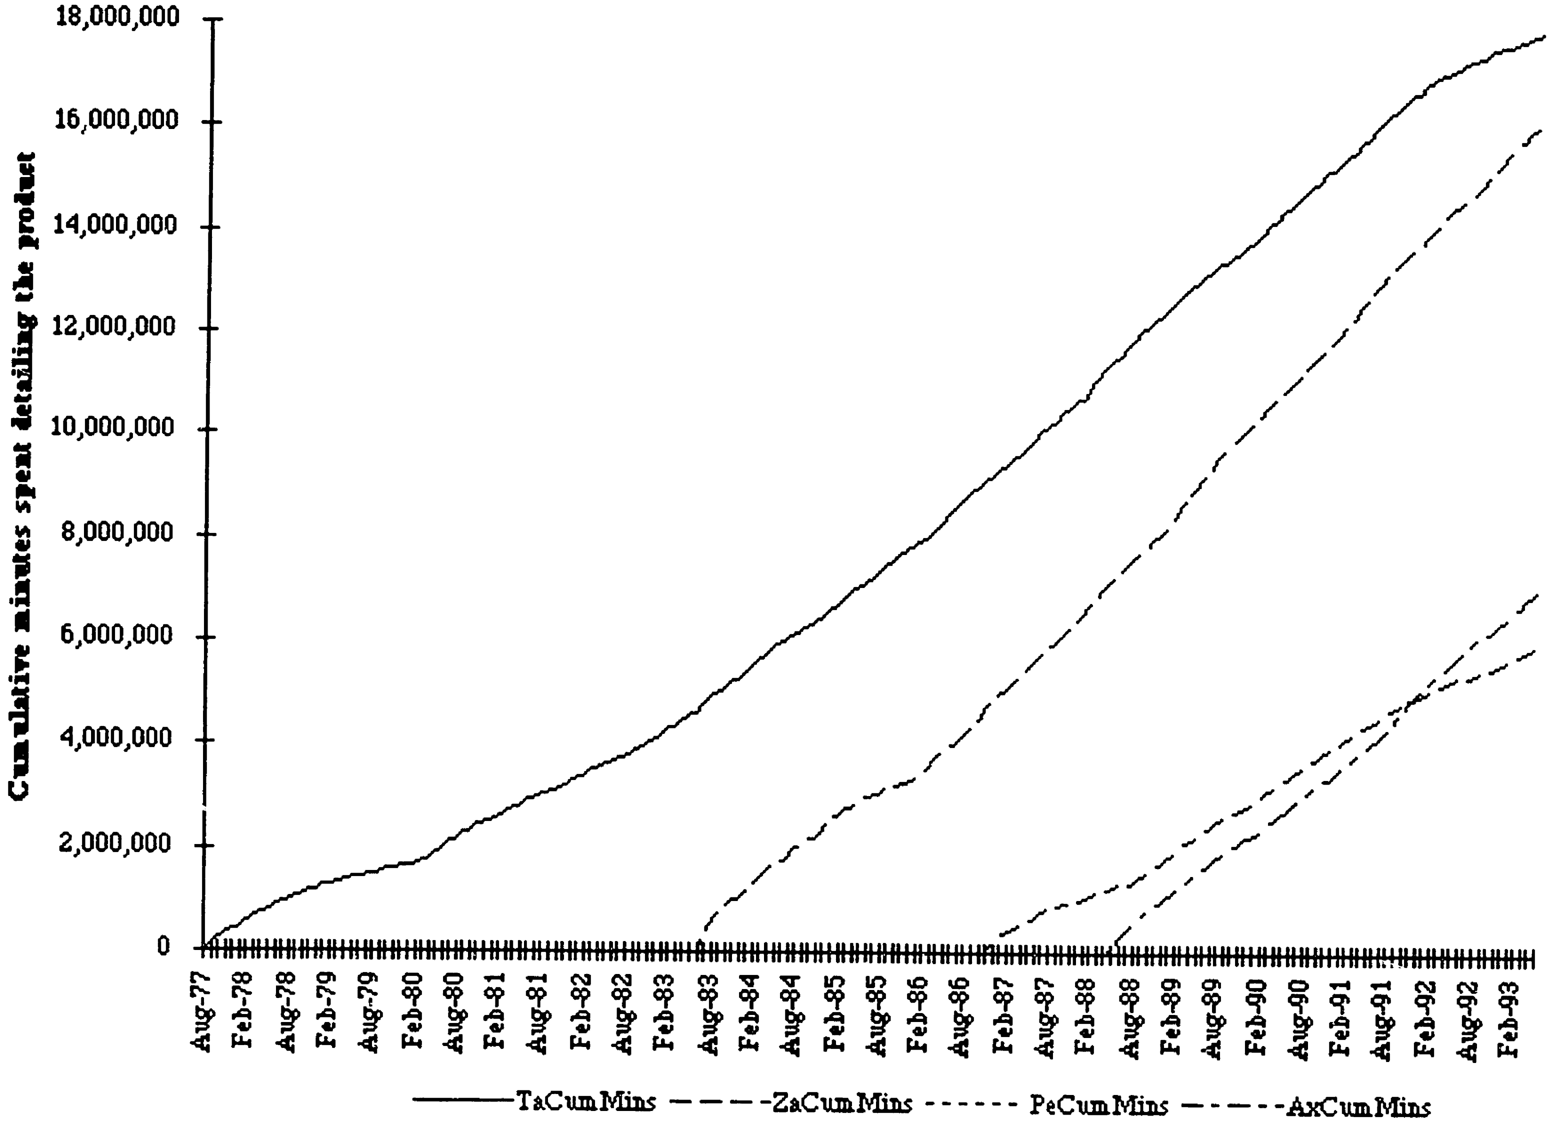

accruing to early entrants, the later entrant may be expected to offer either a lower price and/or a heavier promotion.[16]The Bond-Lean conjecture relates of course to the considerable theoretical and empirical literature in marketing and economics dealing with first mover advantages.[17] It is therefore of interest to examine whether this conjecture is consistent with the data from the H_2-antagonist drug market. Although we present econometric evidence on order of entry effects later in Section 5, in Figure 7 we display cumulative detailing/cumulative sales ratios as a function of order of entry after one year in the marketplace (the leftmost set of bars), after two years (the middle set), and after three years (the rightmost set). The results are striking. Given any duration of time, cumulative detailing/sales ratios are always lowest for the pioneer (Tagamet), are always larger for the second entrant (Zantac), always increase further for the third entrant (Pepcid), and are always highest for the final entrant (Axid). Moreover, since a disproportionate amount of detailing occurs immediately following product launch, for all four H_2-antagonist products the cumulative detailing/sales ratios decrease as the time interval since launch increases.

Detailing is not the only form of marketing rivalry, however. Another instrument for bringing product information to the attention of prescribing physicians is via medical journal advertising. It is worth mentioning that relative to detailing, estimated expenditures on journal advertising are rather modest; as observed earlier, expenditures on detailing are approximately four to five times as great as expenditures on journal advertising in the overall U.S. pharmaceutical industry, although substantial variations occur across products.

It might be noted that to convert nominal to real dollars, one must employ a deflator. We use the BLS price index for scientific and professional journals. Based on a preliminary analysis of advertising rates charged by two major medical journals, the *New England Journal of Medicine* and the *Journal of the American Medical Association*, however, we found that the BLS deflator appeared to rise less rapidly in the 1980's than did advertising rates in these journals. An alternative measure of real medical journal advertising involves a simple page count. This measure does not account well, of course, for variations in copy quality, or in journal circulation. Later in this paper we discuss these two measures further. For our current purposes, it is sufficient to note that the two measures are reasonably highly correlated. In Figure 8 we plot cumulative medical journal dollars spent for each of the four H_2- antagonist products, using the BLS deflator. Clearly the launch of Tagamet coincided with a considerable journal advertising campaign. Thereafter until receiving

[16] As Bond-Lean [1977, p. vi] state, "Neither heavy promotion nor low price appears to have been sufficient to persuade prescribing physicians to select in great volume the substitute brand of late entrants...When other things are equal, physicians appear to prefer the brands of existing sellers to those of new sellers."

[17] On first mover advantages, see, for example, the surveys and references in Kalyanaram-Urban [1992], Robinson [1988], Robinson-Fornell [1985], Robinson-Kalyanaram-Urban [1994], Samuelson-Zeckhauser [1988], Schmalensee [1982], and Urban-Carter-Gaskin-Mucha [1986]. For an alternative interpretation, see Golder-Tellis [1992].

Figure 7. Cumulative Detailing/Sales Ratios

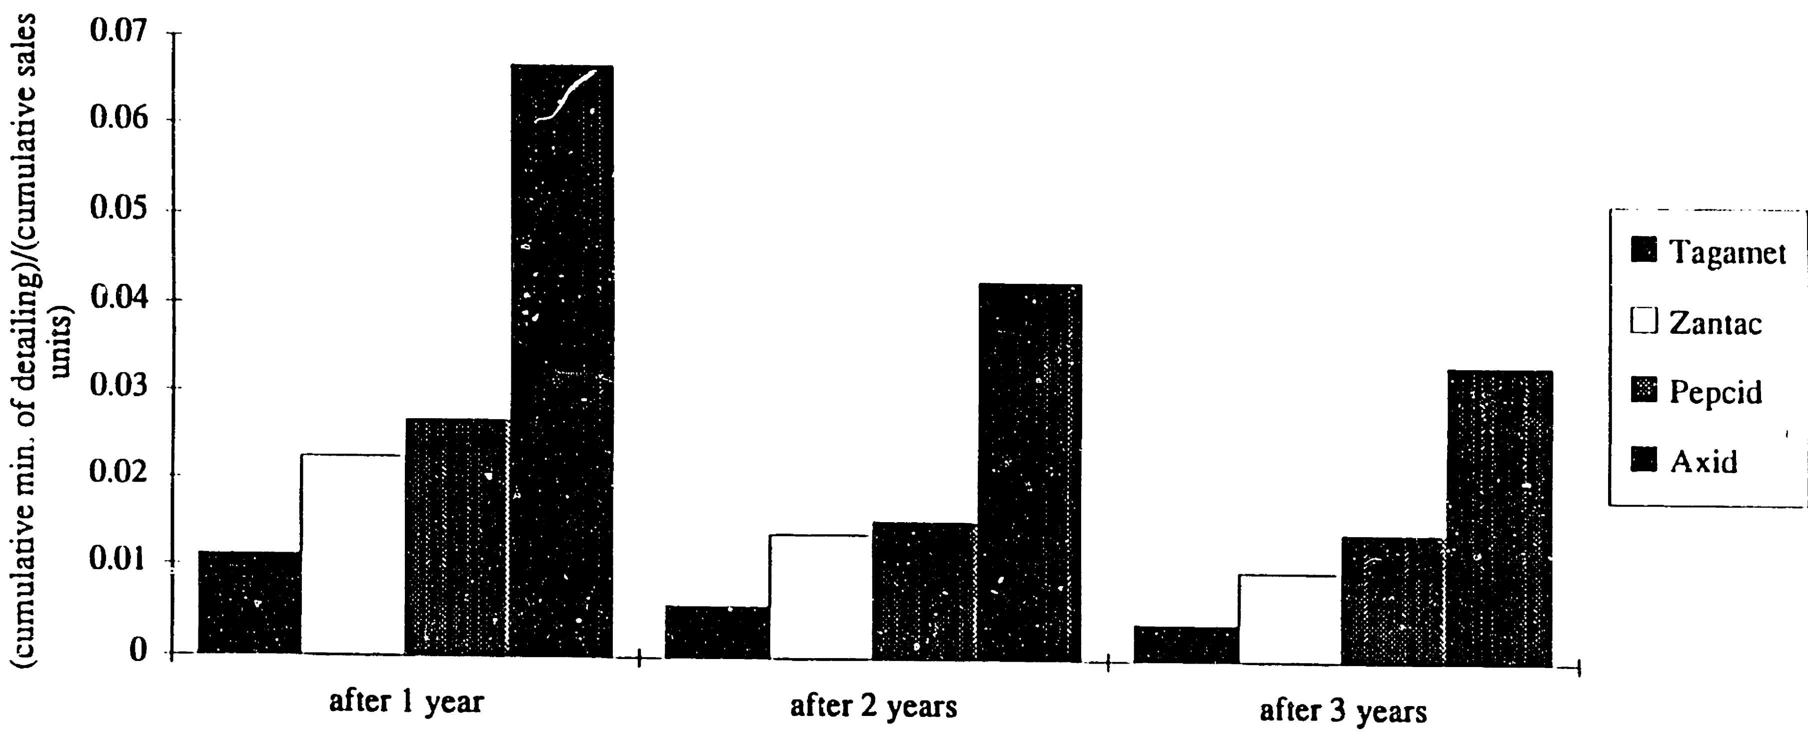

Figure 8. Cumulative Real Journal Advertising

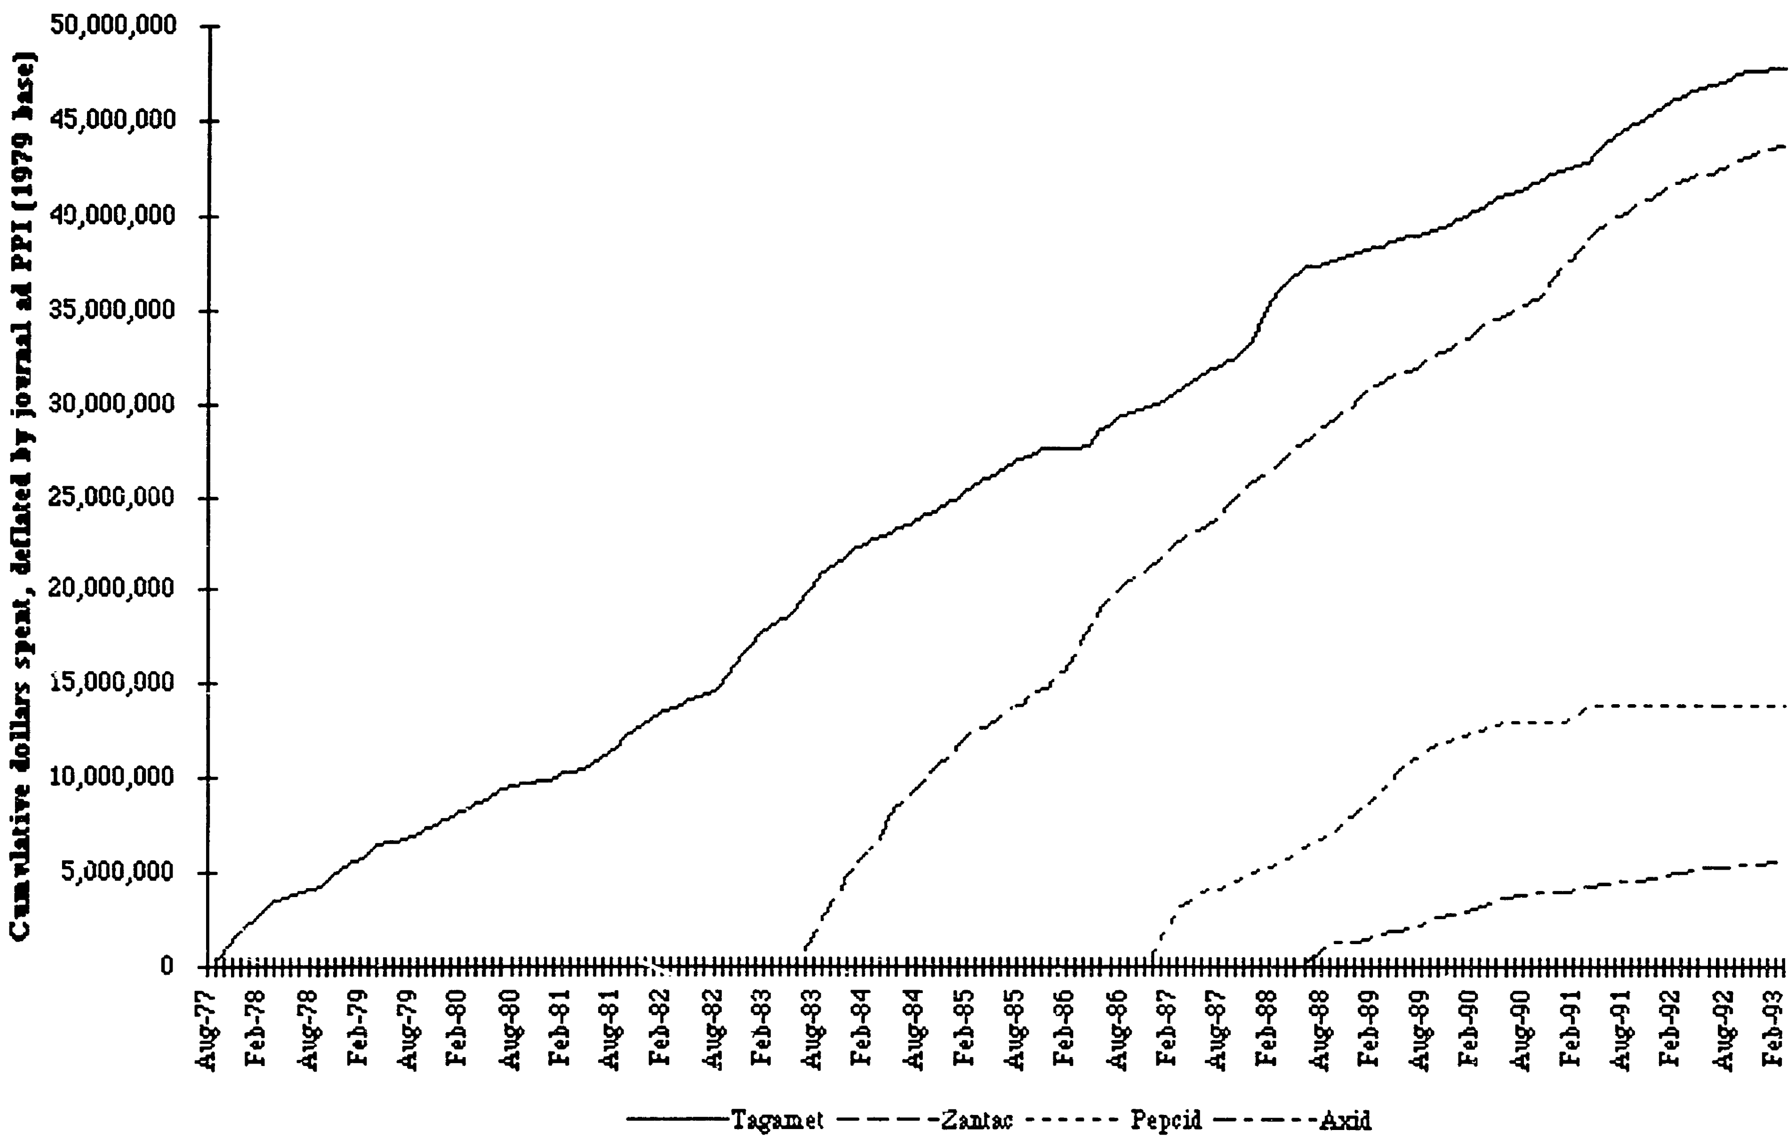

FDA approval for duodenal ulcer maintenance in April 1980, Tagamet journal advertising was relatively modest, with temporary increases around the time of FDA approval for gastric ulcer treatment (December 1982) and for GERD (March 1991). It is noteworthy that Tagamet journal advertising increased only moderately after the entry of Zantac in August 1983, and it did not respond aggressively when Pepcid entered in late 1986. In terms of its response with journal advertising to entry by Pepcid and Axid, Zantac was roughly similar to Tagamet. Spurts in Zantac journal advertising appear to follow closely the obtaining of FDA approval for gastric ulcer treatment (June 1985), and the simultaneous approval for duodenal ulcer maintenance and GERD (May 1986). Finally, a comparison of Figures 6 and 8 reveals that Pepcid and Axid differed considerably in their choice of marketing medium in the sense that Axid has relied much more heavily than Pepcid on detailing, and much less on medical journal advertising.[18]

With this overview of price, product quality and marketing competition data trends in the H_2-antagonist market, we now turn our attention to modeling the growth in overall industry sales, and to modelling changes in the shares earned by the various products. We begin in Section 4 with an analysis of overall industry growth, and then consider market shares in Section 5.

4 Econometric Analysis of Growth in Industry Sales

In this paper we consider the four H_2-antagonist products as constituting a distinct market or industry. However, since Tagamet and Zantac so clearly dominate the H_2-antagonist market, we shall also consider a separate, simpler market -- that consisting only of Tagamet and Zantac. We first digress to consider theory and measurement issues, and then present econometric results.

4.1 Theoretical and Econometric Considerations

The traditional approach to modeling demand for a product involves calling upon the economic theory of consumer demand, in which consumers are assumed to maximize utility given prices of products and an overall budget constraint; additional assumptions are then employed to aggregate up from the individual consumer to an overall industry demand. In the context of pharmaceutical products, this approach is unlikely to be useful, for the typical decisionmaker (the physician) is not the consumer (the patient) who actually pays for the prescription drug product. Moreover the marginal price paid by the patient often differs considerably from the price received by the dispensing pharmacy, due to the existence of third-party insurance and various co-payment schemes. While a discussion of such principal-agent problems is beyond the scope of this paper, we believe the ex-

18 Industry sources say that this is not only true for Axid, but for all of Lilly's products. Lilly's corporate strategy has been to use a much higher percentage of detailing over journal advertising in their marketing efforts. Lilly's mix of detailing to advertising is approximately 90%-10%, whereas the industry average is 75%-25%.

istence of these institutional arrangements clearly suggests that rigid adherence to the traditional neoclassical approach of demand analysis is unlikely to be useful here.

Although we eschew here the direct use of conventional utility-maximizing economic behavior, we still wish to incorporate the most important insights of demand analysis. Thus we specify that quantity demanded depends on the price of the product, product characteristics, and marketing efforts. We now discuss these three factors affecting demand in further detail.

In terms of price, economic theory suggests that quantity demanded depends on real rather than nominal price; since we employ time-series data, we deflate average product price by the Consumer Price Index (CPI). Also, although product-specific price data are available, for examining overall industry demand one must construct an industry price index. The important point here is that since we wish later on in this paper to investigate the extent of price-substitutability among drugs, when constructing an aggregate price index for the industry it is important we not implicitly assume a value for such substitutability. In particular, if one simply summed up patient-days of therapy across drugs, then summed up total revenue across drugs, and finally, calculated price as total industry revenue divided by total industry patient-days, one would implicitly be assuming that the various drugs are perfectly substitutable. To circumvent this problem, we employ the economic theory of price indexes, and calculate the industry price using the Fisher-Ideal price index.[19]

In terms of quality, to the extent that product quality characteristics affect the size of the potential market, they should be included in an overall industry demand equation. We would expect that the size of the potential patient market would depend on the specific indications for which the FDA has granted approval. We shall concentrate on one particular indication, GERD, which represented an especially large potential new market, and for which the H_2-antagonists first received FDA approval relatively late in the sample. Specifically, when the FDA granted approval to Glaxo's Zantac for GERD, Zantac detailers were permitted to provide specific information to physicians concerning the treatment of GERD. This was significant, for instead of being confined to detailing to gastroenterologists who saw ulcer patients, now Zantac detailers also made calls on general practitioners who commonly saw patients having GERD symptoms. This undoubtedly expanded the potential market.

Such reasoning suggests that a dummy variable, say, GERD (taking on the value of one following FDA approval) be employed in the overall industry demand equation. However, it is worth noting that information concerning efficacy of drugs for different indications typically diffuses prior to formal FDA approval. The medical community is often aware of results of clinical trials prior

---

[19] Specifically, the Fisher-Ideal price is the geometric mean of the Laspeyres and Paasche price indexes, where each of them is computed using updated weights. New products are incorporated as soon as feasible (i.e., in the second period of their existence, so that their first difference is calculated). For further details concerning the Fisher-Ideal price index, see W. Erwin Diewert [1981,1992].

to the FDA reviewing the clinical trial data and coming to a final decision concerning approval for a new indication. As a result, a great deal of prescribing is done off-label, prior to the FDA granting approval. Thus, it is not clear how reliable the GERD dummy variable will be in capturing major changes in the size of the potential patient base.

The third set of factors affecting industry demand involves marketing efforts. Earlier we noted that in this industry, the two principal forms of marketing efforts are minutes of detailing and either pages or deflated dollars of medical journal advertising. There are several important issues concerning the measurement of marketing efforts. First, since drug marketing is largely a matter of providing information about the existence and usefulness of the product, we expect its impact to be long-lived; once a physician has been informed, it is hard to see how such information might be destroyed. Indeed, precisely because of this durability, firms typically expend a particularly large amount of marketing effort in the early stages of a new product's life. Hence the impact of marketing on sales is likely better measured by a cumulative stock of marketing efforts since product launch, rather than simply by the flow of current monthly expenditures. We will also want to allow for the possibility that this stock of information depreciates or deteriorates over time, although we expect the depreciation rate to be quite low.

We therefore employ the well-known perpetual inventory method. Let M_t be the *stock* of marketing effort at the end of month t, (as measured by the stocks of journal advertising or detailing minutes), let δ be the monthly rate of depreciation of this stock, and let m_t be the flow of marketing effort during time period t. Define M_t as the depreciation adjusted stock of marketing effort carried over from the last month, $(1 - \delta)M_{t-1}$, plus new marketing efforts (m_t) during month t, i.e.,

$$\begin{aligned} M_t &= (1-\delta)M_{t-1} + m_t \qquad . \qquad\qquad (1) \\ &= \sum_{\tau=0}^{t} (1-\delta)^{t-\tau} m_{t-\tau} \end{aligned}$$

We construct separate stock measures for detailing and for journal advertising. Unlike the typical case for capital stock accounting, we have no problem with establishing benchmark or "starting values" since we know that prior to August 1977, the Tagamet journal (and detailing) stocks were zero. To implement Eq. (1), one must however assume rates of depreciation for each of these stocks. As discussed below, we will use the historical data on marketing and sales to estimate δ econometrically, rather than assume its value *a priori*.

The other major issue in measuring the effects of marketing efforts entails an innovation of this paper. Other authors have pointed out that advertising be modeled as having two simultaneous effects in the market: overall advertising by all firms affecting overall market demand, and relative levels of advertising among firms affecting the individual firms' market shares.[20] We take this modeling one step further here by hypothesizing that firms may choose to direct their marketing efforts to emphasize one of the two effects more than the other. Although the degree to which firms' marketing efforts are directed, say, at overall market expansion cannot be directly observed from data on quantities of marketing done by firms, we now propose a method for estimating this effect econometrically.

To clarify this concept, we now discuss it in the context of the anti-ulcer drug market. When SmithKline (SK) marketed Tagamet from its introduction in 1977 until the entry of Zantac in 1983, SK did not worry about competing for market share in the H_2-antagonist market, for patent status conferred on them a temporary monopoly position. In this monopoly position, the goal of marketing for SK was to convince more and more physicians of the utility of H_2-antagonists in treating ulcer patients. SK, and no other firm, reaped the rewards of having expended efforts on diffusing information on H_2-antagonist drugs to physicians, since SK held 100% market share. However, once Zantac entered the market, another SK marketing goal appeared: to preserve market share against Zantac among those doctors who had already adopted the H_2-antagonist technology. Similarly, although Zantac detailers could benefit somewhat from continuing to reach out to new doctors and patients still not converted to the H_2-antagonist technology, Zantac detailers also had strong incentives to persuade physicians already in the H_2-antagonist market to begin prescribing Zantac instead of Tagamet, emphasizing the alleged Zantac advantages of lower-frequency dosing and fewer adverse drug interactions. Unlike the monopoly case, in this duopoly situation the marketing efforts of firms may have both market-expanding and rivalrous (product-positioning) aspects.

Moreover, to the extent that Zantac would reap some of the benefits of Tagamet's market-expanding efforts to persuade physicians to adopt the H_2-antagonist drugs, and that Tagamet might similarly benefit somewhat from Zantac's market-expanding promotions, each firm's market-expanding promotional effort exerts a positive externality (spillover) on the other firm's sales. Similarly, we might consider rivalrous marketing to exert negative inter-firm externalities. When the number of products in the market increases, *ceteris paribus*, we would expect a decrease in firms' incentives to engage in market-expanding promotional efforts, and correspondingly greater incentives to engage in marketing with a more rivalrous content.[21] The practical implication of this hypothesis is that in a duopoly, *ceteris paribus*, one would expect the product marketing of the two participants to have a smaller impact on industry demand than would be the case if this adver-

---

[20] See, for example, Schmalensee [1972]. There is a considerable body of literature on a related, but distinct, approach that decomposes advertising into its "information" and "persuasive" components. For examples in the context of the pharmaceutical industry, see Leffler [1981], and Hurwitz and Caves [1988].

[21] This also implies that incentives to advertise, and perhaps the content of advertising messages, can be expected to vary with industry structure.

tising had occurred in a monopoly market structure, for some of the duopolists' advertising would primarily impact market share, not overall industry demand; similarly, *ceteris paribus*, for a given amount of cumulative marketing stocks, one might plausibly expect that in a triopoly the effects of marketing on industry demand would be less than in a duopoly.

In this paper we examine this hypothesis empirically by inferring econometrically the proportionate impact (relative to a monopolist) that marketing efforts have under varying market structures. To this end, we distinguish cumulative marketing efforts according to the market structure in which such expenditures originally occurred. Let $M_{1,t}$ be the marketing stock at end of month t that accumulated in the monopoly market environment, let $M_{2,t}$ be the marketing stock at end of month t that accumulated in the duopoly market environment, and let $M_{k,t}$ be the marketing stock at end of month t that accumulated in a market environment consisting of K products. Define the "effective industry marketing" stock M_t as the weighted sum of the cumulative marketing efforts distinguished by market structure, *i.e.*,

$$\overline{M}_t = \mu_1 M_{1,t} + \mu_2 M_{2,t} + \mu_3 M_{3,t} + \ldots + \mu_k M_{k,t} \tag{2}$$

where the $M_{k,t}$, $k = 1,\ldots,K$, are each defined as in (1). *Ceteris paribus*, we therefore might plausibly expect that

$$\mu_1 > \mu_2 > \mu_3 > \ldots > \mu_k, \tag{3}$$

reflecting the fact that in terms of affecting overall industry demand, participants' market-expanding effects decline as the number of products in the industry increases.[22] Since in a monopoly <u>all</u> marketing efforts affect industry demand, we normalize the μ_k's by setting $\mu_1 = 1$.

It is worth noting that two other hypotheses might be proposed involving the μ_k's. First, if the effectiveness of firms' marketing on industry sales is independent of market structure, then $\mu_2 = \mu_3 = \mu_4 = 1$. Second and alternatively, if $\mu_2 = \mu_3 = \mu_4 = 0$, then in the presence of any competition all marketing efforts are rivalrous and affect only market shares. Note that in such a case of possibly but not necessarily socially "wasteful" marketing, firms' marketing efforts generate a zero-sum change in industry sales. In our empirical analysis, we will estimate the remaining μ_k's in Eq. (2) and assess whether the evidence is consistent with any of these hypotheses.

We begin with some definitions of variables. Let Q_t be total units of sales for all products (a Fisher-Ideal quantity index), let PR_t be the corresponding real price index (deflated by the CPI), let $D_{k,t}$ be the stock of minutes detailed by product k at the end of time period t, let $J_{k,t}$ be the stock

---

[22] Note that the μ's do not deal at all with the effects of marketing stocks on the market *shares* garnered by the various firms in the market. We discuss determinants of market shares further in Section 5 below.

of pages advertised in medical journals by product k at time t, and let $GERD_t$ be the above-noted GERD dummy variable.

In terms of a mathematical formulation, we specify a traditional log-linear demand equation, where, however, the use of identities (1) and (2) necessitates estimation by nonlinear least squares (NLS) procedures. In particular, let

$$lnQ_t = \beta_0 + \beta_1 lnPR_t + \beta_2 ln\bar{D}_t + \beta_3 ln\bar{J}_t + \beta_4 GERD_t + \varepsilon_t, \qquad (4)$$

where ε_t is an identically normally distributed random error term, and where $ln\bar{D}_t$ and $ln\bar{J}_t$ are natural logarithms of the "effective industry marketing" stocks of number of minutes detailed and pages of medical journal advertisements,[23] respectively, defined as:

$$\bar{D}_t = D_{1,t} + D_{2,t} + D_{3,t} + D_{4,t} \qquad (5)$$

and

$$\bar{J}_t = J_{1,t} + J_{2,t} + J_{3,t} + J_{4,t}. \qquad (6)$$

In turn, following Eq. (1), define the effective stock of minutes at end of month t for a market structure consisting of k products as

$$D_{k,t} = (1 - \delta_M)D_{k,t-1} + MIN_{k,t} \qquad (7)$$

where δ_M is the constant rate of depreciation for the detailing minutes stock, and MIN is the number of minutes detailed during month t, where month t was one in which the market structure consisted of k products. The construction of effective stocks of journal pages $J_{k,t}$ by type of market structure is analogous to that in Eq. (7). Since Eqs. (5) - (7) are nonlinear in the μ's and δ's, for convenience we

will constrain and $\delta_M = \delta_J$, but of course the μ_k (equal for minutes and journal pages) will still be permitted to differ with industry structure k in order that the hypothesized inequality in Eq. (1) might emerge.

There is one other issue that merits attention. At the industry level, one would expect price to be simultaneously determined with quantity. Moreover, as has been emphasized by, among others,

---

[23] Two possible measures of medical journal advertising are current dollar expenditures divided by a BLS price index for advertising in professional journals, and the number of pages of medical journal advertising. An exploratory examination of the BLS price index suggested to us that in the 1980's and early 1990's it increased much less rapidly than advertising rates in the *New England Journal of Medicine* and the *Journal of the American Medical Association*. On the other hand, the page measure does not account well for variations in copy quality, or in journal circulation. Results from preliminary regression estimation suggested that the page measure provided more plausible parameter estimates.

Dorfman-Steiner [1954] and Schmalensee [1972], advertising efforts are also likely to be jointly determined with price and quantity. In terms of stochastic specification, therefore, it may well be the case that LNPR, LND and LNJ are correlated with E, in which case estimation by nonlinear least squares would provide inconsistent estimates of the parameters. In the next section, we therefore report results of a Hausman test for this possible endogeneity, and since we find the correlation to be significant, we also estimate and report results using the nonlinear two-stage least squares (NL-2SLS) estimator.

4.2 Results of Econometric Analysis

Our data set consists of 189 monthly observations beginning in September 1977. We proceed using two alternative definitions of the market, one comprised of the two dominant products, Zantac and Tagamet, and the other consisting of all four H_2-antagonists. In each case, we begin by setting the depreciation rate $\delta = 0$; we then examine and choose among several possible alternative specifications. Given reasonable regression equations, we perform a grid search for the best-fit value of δ by re-estimating the models assuming a variety of depreciation rates, where $0 \leq \delta \leq 1$. We choose as our final set of parameter estimates that value of δ and the other parameters for which the sum of squared residuals is minimized (the sample likelihood function is maximized). Our findings are summarized in Table 1; the top half is that for the two-product market, while the bottom is that of the four-product market.

Table 1: Parameter Estimates in the Two- and Four-Product Industry Models[a]

Method	NLS	NL-2SLS	NLS	NL-2SLS
Market	T-Z	T-Z	T-Z-P-A	T-Z-P-A
β_0	-6.574*	-5.165*	-7.291*	-7.110*
	(-0.46)	(-0.54)	(-0.58)	(-0.68)
β_1	-0.901*	-1.072*	-0.737*	-0.737*
	(-0.11)	(-0.14)	(-0.14)	(-0.20)
β_2	0.534*	0.413*	0.574*	0.574*
	(0.06)	(0.08)	(0.06)	(0.08)
β_3	0.210*	0.275*	0.166*	0.174*
	(0.06)	(0.08)	(0.06)	(0.07)
β_4	0.157*	0.164*	0.117*	0.118*
	(0.03)	(0.03)	(0.03)	(0.03)
μ_2	0.688*	0.892*	0.577*	0.600*
	(0.07)	(0.12)	(0.08)	(0.11)
μ_3			0.812*	0.848*
			(0.14)	(0.18)
μ_4			0.464*	0.491*
			(0.09)	(0.13)
δ	0.002	0.000	0.000	0.000
	(0.00)			
R^2	0.992	0.992	0.994	0.994
D-W	1.767	1.729	1.909	1.907
N	189	189	189	189

a. Notes: T-Z: Tagamet-Zantac. T-Z-P-A: Tagamet-Zantac-Pepcid-Axid. Standard errors reported in parentheses. * denotes significance at the 95 percent level.

First, as seen in the top row of Table 1, the iterative NLS procedure yielded an optimum when δ is very small (0.2% per month), and is not significantly different from zero.[24] While we expected a low value for this depreciation rate since knowledge and information about a product is very durable, that we obtained such a very low rate of depreciation is somewhat surprising. It is worth noting, however, that in an inter-industry productivity study estimating the depreciation rate of R&D capital (another good whose use involves potential spillovers, and for which information plays a central role), Griliches and Lichtenberg [1984] reported an estimated depreciation rate of zero.

---

24 The implicit standard error estimates in Table 1 are conditional on the value of δ. The t-statistic for δ was computed by comparing the likelihood function at $\delta = 0$ with that at $\delta = 0.0020$, and then computing the implied test statistic.

Second, the estimate of μ_2 is about 0.69, and with a standard error estimate of about 0.07, it is significantly different both from unity and from zero. Since μ_1 has been normalized to unity, this estimate of μ_2 implies that, *ceteris paribus*, observed marketing stocks of detailing minutes and journal pages are only about 70% as effective in changing *industry* sales when they occur in a duopoly (Tagamet and Zantac), relative to when they take place in a monopoly (Tagamet). This is a plausible result, for anecdotal evidence suggests to us that much of the Zantac-Tagamet duopoly era contained highly competitive marketing aimed at securing market share, rather than focused on increasing overall industry growth.[25]Nonetheless, as was shown in Figure 1, during this duopoly industry sales grew rapidly.

Third, in terms of marketing effectiveness, as is seen in the top row of Table 1, the elasticity of sales with respect to effective cumulative industry detailing minutes (LND) is slightly greater than 0.5, and is about two and one-half times as large as that for effective cumulative industry journal pages (LNJ), whose value is about 0.2.

Fourth, each of these two marketing elasticities is estimated to be considerably smaller in absolute magnitude than the market price elasticity, which is slightly less than unity (-0.90).

Finally, although we have some hesitations concerning its reliability in tracking physician awareness, the coefficient on GERD (a dummy variable equal to one during the time period in which the FDA approved an H_2-antagonist drug for the GERD indication) is positive and significant; the estimate implies that, *ceteris paribus*, FDA approval for GERD increased the market size by about 15%.

These NLS results are based on the assumption that the regressors are uncorrelated with the disturbance term (in our context, that the regressors are all exogenous). We have tested for this assumption using a Hausman specification test, based on instruments that will be discussed below. We find that the joint null hypothesis of no correlation between ε and LNPR, ε and LND, and ε and LNJ is soundly rejected:[26]the likelihood ratio test statistic is 49.2, while the 0.01 critical value for the five restrictions is 15.1.[27]This implies that NLS generates inconsistent parameter estimates, and suggests that we instead employ the NL-2SLS estimator.

We utilize two groups of exogenous variables to form the instruments. One group is common to both firms: the log of the producer price index for intermediate materials, the log of the wage rate for production workers in the pharmaceutical industry, the GERD dummy variable, and a time

25 For a journalist's account of Glaxo's marketing activities and their success in the marketplace, see Lynn [1991].

26 More precisely, the null hypothesis involves testing that the various component (monopoly, duopoly) stocks of MIN and PJL are uncorrelated with E. Hence under the alternative hypothesis there are five endogenous variables, monopoly stocks of MIN and PJL, duopoly stocks of MIN and PJL, and price.

27 Coefficients on each of the marketing stock variables, and on the price variable, were significantly different from zero as well.

counter. The other set incorporates firm-specific variation, but aggregates them to the industry level: the number of details by firms for their products other than those in the H_2-antagonist market, and the number of real dollars of medical journal advertisements for the firms' non H_2-antagonist products. To make these variables comparable to the components of the regressors LND and LNJ (see Eqs. (5) and (6) above), we construct stocks separately by type of industry structure, and then cumulate them assuming $\delta = 0$.

The results of the NL-2SLS estimation are presented in the second row of the top panel in Table 1. Relative to the NLS findings, a number of results are worth noting. First, with NL-2SLS the criterion function is optimized when $\delta = 0$. This estimate is low, but as noted above, it is not without precedent in a related context. Second, under NL-2SLS estimation, the estimate of μ_2 increases from 0.69 to about 0.89, and now is no longer significantly different from unity. It is, however, significantly different from zero. Third, the price elasticity estimate under NL-2SLS is slightly larger in absolute value (-1.07 vs. -0.90), unlike estimates of the detailing minute elasticity (0.41 vs. 0.53). Fourth, for the journal page elasticity, under NL-2SLS estimation the estimate increases from 0.21 to 0.28. Hence with NL-2SLS as well as NLS estimation, the estimates of the journal page and detailing elasticities are much smaller in absolute value than is the price elasticity. Finally, under either estimation method, the R^2 is above 0.99, and the Durbin-Watson test statistics are very close to 2.0.

We now turn to a discussion of findings obtained under a four-product market definition; results are given in the bottom panel of Table 1. As shown in the table, under either estimation method the goodness of fit is above 0.99, and the Durbin-Watson test statistic is again quite close to 2.0. For both NLS and NL-2SLS, the estimated δ at the optimum was 0.00. Hence, the very low depreciation estimate results for marketing efforts of detailing minutes and pages of medical journal carries over from the two-product to the four-product market context. Also, the Hausman test for exogeneity is again clearly rejected, although not as decisively as in the two-firm analysis; here the likelihood ratio test statistic is 41.3, while the 0.01 critical value for the nine restrictions is 21.7. This suggests again that the NLS estimates may be inconsistent, and that we instead employ the NL-2SLS estimator.

The NL-2SLS estimate for the market price elasticity in this four-product market is slightly smaller (in absolute value) than in the two-product case, around -0.74 vs. -1.07. The estimate of the sales elasticity with respect to cumulative detailing minutes is somewhat larger here (0.57 vs. 0.41), while that with respect to journal pages is slightly smaller (0.17 vs. 0.28). Moreover, with the larger four-firm market definition the GERD coefficient declines slightly, from about 15% to 12%.

Of particular interest, however, are the estimates of μ_2, μ_3 and μ_4. Recall from the discussion surrounding Eq. 3 that, *ceteris paribus*, we hypothesized that $1 > \mu_2 > \mu_3 > \mu_4$. As is seen in the bottom panel of Table 1, this pattern is largely, but not completely borne out; although less than unity, typical estimates of these three parameters are 0.6, 0.8 and 0.5, respectively. Why it is that marketing efforts in the triopoly epoch were more effective in generating industry sales than during the two- and four-product eras is an issue meriting further examination. Moreover, the joint null hypothesis that these μ's are all unity (that the effectiveness of marketing efforts on sales is independent of market structure) is decisively rejected, as is the joint hypothesis that $\mu_2 = \mu_3 = \mu_4 = 0$, the latter indicating that market-expansion spillovers do not entirely disappear when competition begins. While these spillovers are considerably lower in the duopoly period than would be the case in a monopoly, and are lower when there are four products on the market than two, in this market the relationship between μ_k and the number of products in the market is not completely monotonic.

5 Econometric Analysis of Factors Affecting Market Shares

To this point our analysis has focused on overall market demand, with alternative definitions of the market. We now report on an exploratory effort at modeling the factors that affect individual market shares earned by each of the products. The results reported here are those from our initial research; we intend to extend this analysis in future research. As in Section 4, we begin with a discussion of considerations drawn from economic theory, and then report on statistical findings.

5.1 Theoretical and Econometric Considerations

The specification of market share or relative demand functions traditionally draws on the economic theory of consumer behavior. As noted earlier, however, principal-agent problems and wedges between marginal relative prices paid and received imply that one cannot directly employ the economic theory framework of consumers maximizing utility, given prices and budget constraints.

Consistent with traditional economic specifications, however, we would expect that relative rather than level prices affect market shares. Moreover, within the marketing literature, there is n-ple precedent for specifying that relative values (ratios) of product characteristics, and relative marketing efforts, affect market shares. In addition, both the economic and marketing literatures suggest that order of entry can be expected to be a significant determinant of market shares.[28] Following Urban et al. [1986], we employ a market share specification of the general form:

$$\frac{Q_{jt}}{Q_{1t}} = f\left(\frac{p_{jt}}{p_{1t}}, \frac{MIN_{jt}}{MIN_{1t}}, \frac{PJL_{jt}}{PJL_{1t}}, X_{jt}, ENT_{jt}\right) \tag{8}$$

[28] See, for example, Schmalensee [1982], Kalyanaram-Urban [1992] and Urban et al. [1986].

where Q_{jt}/Q_{1t} is the sales of product j relative to product 1 (the first or pioneer entrant, in this case, Tagamet) in month t, P_{jt}/P_{1t} are the corresponding relative prices per day of therapy, MIN_{jt}/MIN_{1t} and PJL_{jt}/PJL_{1t} are relative cumulative stocks of minutes of product detailing and cumulative pages of medical journal advertisements (defined as in Eq. 1), X_{jt} are a set of s variables measuring the quality of product j relative to the pioneer (e.g., dosage frequency, number of (adverse) drug interactions reported to the FDA, whether product j has a GERD indication advantage relative to the pioneer, etc.), and ENT_{jt} is the order of entry of product j (i.e., 2 for Zantac, 3 for Pepcid and 4 for Axid.).

In our context, the pioneer product is Tagamet, and thus all variables in Eq. (8) are measured for product j relative to Tagamet. Since market shares are 100% for Tagamet during its monopoly epoch (September 1977 through July 1983), the data set for which market share analysis is appropriate commences in August 1983; data prior to this are not employed. In the case of a two-product market definition, the data therefore consist of Zantac/Tagamet relative quantities beginning with August 1983, a total of 118 observations. For the four-product (H_2-antagonist) market definition, however, the data set is expanded to incorporate relative Pepcid/Tagamet data points (December 1986 onwards), as well as relative Axid-Tagamet observations (beginning with June 1988), giving us a total of 255 observations. Note that in this four-product model the data take the form of an unbalanced panel.

Finally, in terms of econometric considerations, one would expect that relative market shares, relative marketing efforts, and relative prices are jointly determined. For this reason, we compare the OLS and NLS results with those based 2SLS and NL-2SLS.

In terms of mathematical formulation, we specify a relative demand equation as in Eq. (8), where variables are logarithmically transformed:

$$\begin{aligned} ln\left(\frac{Q_{jt}}{Q_{1t}}\right) &= \beta_1 ENTRY_{jt} + \beta_2 ln\left(\frac{p_{jt}}{p_{1t}}\right) + \beta_3 ln\left(\frac{MIN_{jt}}{MIN_{1t}}\right) + \beta_4 GERD_{jt} + \beta_5 ln\left(\frac{FREQ_{jt}}{FREQ_{1t}}\right) \\ &\quad + \beta_6 ln\left(\frac{INT_{jt}}{INT_{1t}}\right) + \beta_7 AGE_{jt} + \varepsilon_{jt} \end{aligned} \tag{9}$$

where $ENTRY_{jt}$ takes on the value 2 for all Zantac observations, 3 for Pepcid and 4 for Axid, $FREQ_{jt}$ is the recommended daily dosage frequency of drug j, INT_{jt} is the number of (adverse) drug indications of drug j reported to the FDA as of time t,[29] $DGERD_{jt}$ is a variable indicating whether product j has a GERD indication advantage relative to Tagamet (1 if an advantage, zero if no advantage, -1 if a disadvantage), and AGE_{jt} is the number of months product j has been in the marketplace. Notice that if the relative price, relative detailing minutes, relative adverse interaction

[29] Data on INT_{jt} are taken from annual issues of the *Physicians Desk Reference*.

and relative dosing frequency variables were all unity, and if the products had no GERD advantage, then at age zero, the relative quantities would depend solely on order of entry effects. Thus the coefficient on ENTRY reflects disadvantages confronting later entrants into the market, other things held equal. The coefficient on AGE reflects the impact of marketplace experience on sales, holding ENTRY (and other variables) fixed. *A priori*, we expect that $\beta_2 < 0$, $\beta_3 > 0$, $\beta_4 > 0$, $\beta_5 < 0$, $\beta_6 < 0$ and $\beta_7 > 0$.[30]

As noted earlier, the data set for this market share model begins when the Tagamet monopoly period ends and Zantac enters. To implement the model empirically, we must make an assumption concerning the "starting value" of the Tagamet stock of detailing minutes. Since the results of our industry analysis suggested depreciation rates for effective industry marketing stocks were zero, we begin the duopoly era using Tagamet's end of monopoly era value for MIN_t, assuming $\delta = 0$. However, we will permit δ, the depreciation rate for these stocks, to differ from zero now that competition has emerged, reflecting in part the fact that the content of marketing may now become more susceptible to counter-claims, and therefore, become less long-lived. Although we have not yet developed a formal model describing optimal behavior in this context, we would not be surprised if the depreciation rate δ in the rivalrous context were larger than it was in the industry-expanding environment.

5.2 Results of Econometric Analysis

We begin with a market share analysis for the two-product market, Tagamet and Zantac. Conditional on any given rate of depreciation, the market share model of Eq. (9) is linear in the parameters. We proceed by estimating parameters in Eq. (9) by ordinary least squares (OLS) under different rates of depreciation and then choose as our preferred model that set of δ and the other parameters that minimizes the sum of squared residuals (maximizes the sample likelihood function). Results from preliminary analysis suggested that it was difficult to obtain precise estimates of both marketing instruments -- minutes of details and pages of medical journal advertising, reflecting in part the fact that the simple correlation between MIN_{jt}/MIN_{1t} and PJL_{jt}/PJL_{1t} was 0.98. In the results presented in the top two rows of Table 2 below, the LNPJL variable was therefore deleted.

30 Note that if one insisted, this logarithmic functional form could be rationalized as deriving from the relative demand equations based on a constant elasticity of substitution (CES) indirect utility function augmented by marketing and product characteristic variables.

Table 2: Parameter Estimates in the Two- and Four-Product Market Share Models[a]

METHOD	OLS	OLS	2SLS	2SLS	OLS	2SLS
MARKET	T-Z	T-Z	T-Z	T-Z	T-Z-P-A	T-Z-P-A
β_1	-0.054 (-0.17)	-0.116 (-0.17)	-0.147 (-0.20)	-0.181 (-0.17)	-0.492* (-0.01)	-0.507* (0.04)
β_2	-0.862* (0.24)	-0.840* (-0.24)	-0.885* (0.24)	-0.886* (-0.24)	-0.643* (-0.06)	-0.693* (-0.07)
β_3	1.003* (0.06)	0.950* (0.05)	0.922* (0.10)	0.893* (0.06)	0.731* (0.02)	0.673* (0.04)
β_4	0.087* (0.02)	0.085* (0.02)	0.093* (0.02)	0.094* (0.02)	0.032 (0.02)	0.046* (0.02)
β_5	-0.066 (-0.05)		-0.023 (0.06)			
β_6	-0.090* (-0.04)	-0.093* (-0.04)	-0.097* (-0.04)	-0 099* (-0.04)	-0.251* (0.02)	-0.232* (0.03)
β_7	0.010* (0.00)	0.010* (0.00)	0.011* (0.00)	0.011* (0.00)	0.012* (0.00)	0.013* (0.00)
δ	0.000	0.000	0.000	0.000	0.039* (0.00)	0.042* (0.01)
R^2	0.993	0.993	0.993	0.993	0.990	0.989
N	118	118	118	118	255	255

a. Notes: T-Z: Tagamet-Zantac. T-Z-P-A: Tagamet-Zantac-Pepcid-Axid. Standard errors reported in parentheses. * denotes significance at the 95 percent level.

Several points are worth noting. First, as in the two-product industry equation, the likelihood function is maximized at the point where $\delta = 0$. A second somewhat unexpected result is that the coefficient on the relative frequency of dosage variable ($LNRFRQ_t \_ LN(FREQ_{jt}/FREQ_{1t})$), though negative, is insignificantly different from zero. We therefore set this parameter to zero, and re-estimate the model. As is seen in the second row of Table 2, the logarithm of the relative quantities of Zantac to Tagamet ($LN(Q_{jt}/Q_{1t})$, the dependent variable) is significantly negatively affected by relative price ($LNRPR \_ LN(P_{jt}/P_{1t})$) -- the own-price elasticity is about -0.8, and is very substantially affected by the relative stocks of cumulative detailing minutes ($LNRMIN \_ LN(MIN_{jt}/MIN_{1t})$) -- this elasticity estimate is about 1.0. As hypothesized, the coefficient on the GERD advantage variable is positive (0.08) and significant, while that on $LNRINT \_ LN[(INT_{jt}+1)/(INT_{1t} + 1)]$, where INT is the number of (adverse) drug interactions reported to the FDA, is negative (-0.09) and significant. Finally, while the order of entry coefficient (in this 2-firm model, essentially just the intercept term) is negative, its standard error is quite large. By contrast, the coefficient on the AGE variable is positive and highly significant.

We then perform a Hausman test to check for possible endogeneity of LNRPR and LNRMIN. The exogenous variables used here are the same as those noted in Section 4 above, except now the firm-specific number of details and dollars of medical journal advertising for products other than those in the H_2-antagonist market are employed, as are dummy variables for whether the product has received FDA approval for duodenal maintenance therapy, gastric ulcers, GERD, and stress ulcer prevention. These latter variables are particularly useful as instruments, since they represent "shocks" and new information for marketing efforts. The results of the Hausman test are not as clear as in the overall market analysis; here the likelihood ratio test for exogeneity of LNRPR and LNRMIN is 6.67, while the 0.01 chi-square critical value for the two restrictions is 6.63. As a sensitivity check, we proceed with 2SLS estimation. Our 2SLS results are presented in the bottom two rows of the top panel in Table 2.

With 2SLS estimation, the fitting optimum is again reached with the depreciation rate $\delta = 0$. Essentially, the results are the same as those obtained under OLS estimation. In particular, the own-price elasticity estimate is about -0.9, about the same in absolute value as the elasticity of sales with respect to cumulative detailing. The DGERD advantage is significant and equal to about 10%, AGE is significant and about 1% per month, while both ENTRY and LNRFRQ are negative but insignificantly different from zero. Finally, Zantac's relative market share is significantly negatively affected by its number of drug interactions relative to Tagamet. At the end of the sample, incidentally, values of INT are 12 for Tagamet and 1 for Zantac. Hence, the INT product quality variable is particularly important in explaining the growth in Zantac's market share and the corresponding decline of Tagamet.

In summary, in the two-product market, relative Zantac-Tagamet quantities demanded are systematically related to relative product prices, relative cumulative detailing efforts, relative product quality (relative adverse interactions and GERD, but not, apparently, by dosing frequency), and by the length of time the product has been on the market. Moreover, for both OLS and 2SLS estimation, the goodness of fit is above 0.99.[31]

We now turn to the broader market definition, one encompassing all four H_2-antagonist products (Tagamet, Zantac, Pepcid and Axid). The results of this analysis are given in the bottom panel of Table 2. First, we now uncover evidence suggesting that in the rivalrous market context, depreciation rates on detailing minutes differ substantially from zero. Specifically, with OLS estimation, the sample log-likelihood function is maximized when $\delta = 0.039$; this monthly rate of 3.9% corresponds with an annual rate of about 38%. Second, with this expanded market definition, order of entry effects (no longer just an intercept term) become very large and significant; the -0.492 estimate corresponds with about a 39% disadvantage accruing to each later entrant, *ceteris paribus*,

[31] Durbin-Watson test statistics in the two OLS equations are 1.646 and 1.624, while in the two 2SLS equations they equal 1.627 and 1.608.

and is remarkably close to the "consensus" estimate of order of entry effects (-0.5) in numerous other markets surveyed by Robinset et al. [1994]. Third, although the price elasticity estimate is slightly smaller in absolute value in this four-firm market than in the two-firm context (-0.7 vs. -0.9), the standard error estimates are much smaller, and the t-statistics are therefore larger. Further, the elasticity of relative sales with respect to relative cumulative detailing minutes is slightly larger in absolute value than the price elasticity (0.73 vs. -0.64), and is also highly significant. Finally, as hypothesized, the coefficient on the relative number of (adverse) drug interactions variable (LNRINT) is negative and significant (-0.25, t-statistic of 12), and that on GERD is positive (0.03), but the latter coefficient is of only marginal statistical significance (t-statistic of 1.9). The AGE coefficient is again slightly greater than 1%, indicating that length of time in the marketplace affects relative sales in a positive manner. Goodness of fit is again about 0.99.[32]

To check on the possible endogeneity of relative prices and relative detailing stocks, we again perform a Hausman specification test. The null hypothesis of exogeneity of LNRPR and LNRMIN is decisively rejected; the likelihood ratio test statistic is 8.53, while the 0.01 chi-squared critical value for the two restrictions is 6.63.

Parameter estimates under 2SLS estimation are given in the bottom row of Table 2. Several findings are of particular interest. First, the estimate of δ at the fitting optimum is 0.042, and is significantly different from zero; this monthly depreciation rate of 4.2% implies an annual rate of about 40%. Hence, these results suggest that in the four-product anti-ulcer market, relative detailing efforts have a long-lived rivalrous impact that depreciates at about 40% per year. Second, order of entry effects are very substantial and statistically significant (-0.51, t-statistic of 59), and again conform remarkably closely to the -0.5 consensus estimate reported by Robinson et al. [1994] for numerous other packaged goods type products. Third, the absolute values of the price and advertising elasticities are roughly the same -- 0.7, and each is significantly different from zero. Thus the evidence suggests that in the four-firm market, relative shares garnered by the four products vary systematically and significantly with order of entry, relative prices and relative cumulative detailing minutes. In terms of product quality variables, increases in the relative number of adverse drug interactions reported to the FDA negatively impact relative sales, whereas having a GERD approval advantage relative to Tagamet positively affects relative sales.

In summary, this exploratory four-firm market share analysis suggests that order of entry, pricing behavior, marketing behavior and product quality all affect relative sales quantities in the hypothesized manner. Moreover, rivalrous detailing appears to depreciate at about 40% per year.

Before leaving this discussion, however, we believe it is of interest to report estimates of total price elasticities. The price elasticity estimates reported in Table 2 focus only on relative quantities

[32] Since the data set now consists of an unbalanced panel, the traditional Durbin-Watson test statistic is no longer appropriate.

(market shares), but leave fixed the size of total industry demand at, say Q; denote these price elasticities by E^*_{jj}. A total price elasticity also captures the impact of a product's price change on total industry demand; denote such a price elasticity by E_{jj} (no asterisk). As has been shown by, *inter alia*, Berndt-Wood [1979], the relationship between E^*_{jj} and E_{jj} is as follows:

$$E_{jj} = E^*_{jj} + \left(\frac{\partial \ln Q_j}{\partial \ln Q}\right)\left(\frac{\partial \ln Q}{\partial \ln P}\right)\left(\frac{\partial \ln P}{\partial \ln P_j}\right) \tag{10}$$

where Q_j is quantity demanded of product j, Q is total industry demand, and P is industry price. The first partial derivative in Eq. (10) can be assumed to equal unity (other things equal, demand for product j grows equiproportionally with market demand, i.e., according to its market share), while the second partial derivative is the industry or market price elasticity (estimated values of which are given in Table 1). The last partial derivative in Eq. (10) indicates the impact of a change in product j's price on the overall industry price index; it can be approximated by the revenue share of product j in total industry revenues.

Alternative OLS and 2SLS estimates of the E^*_{jj} are given in Table 2, while NLS and NL-2SLS estimates of the industry price elasticity are presented in Table 1. For the two-product market, 1993 drug-store revenue shares for Tagamet and Zantac are approximately 0.25 and 0.75. For the four-product market, these shares are approximately 0.19 (Tagamet), 0.60 (Zantac), 0.12 (Pepcid) and 0.09 (Axid). Together, these relationships imply that in the two-market context, the 2SLS estimate of the total own-price demand elasticities for Tagamet and Zantac are approximately -1.154 and -1.690, respectively, while in the four-product market, the 2SLS estimated total own-price demand elasticity is -0.909 for Tagamet, -1.153 for Zantac, -0.820 for Pepcid, and -0.799 for Axid. Note that while these point estimates imply that some of the demand elasticities are less than one in absolute magnitude, the associated standard errors may well imply that reasonable confidence intervals include values of one and above (in absolute value).

6 Concluding Remarks

In this paper we have attempted to explain the phenomenal growth of the H_2-antagonist anti-ulcer drug industry in the U.S., as well as changes in the market shares garnered by the various products over time. Although we have examined the roles of product quality, order of entry and price, we have focussed particular attention on the role of various marketing efforts. Our framework and results can be summarized as follows.

First, marketing efforts such as detailing and medical journal advertising have long-lived impacts. Thus in explaining current period sales, a stock of cumulative detailing or cumulative medical journal advertising is a more appropriate measure of marketing impacts than is current monthly expenditures. In the context of industry demand, we distinguish investments of firms in these marketing activities by the industry structure prevailing when the expenditures originally occurred. In a monopoly market structure, all marketing expenditures are market-expanding, for the monopolist has 100% market share. In a market structure with k products, however, marketing activities become more rivalrous, and as k becomes large, we expect relatively little "spillover" of a firm's marketing efforts in affecting industry demand. We have hypothesized, therefore, that in terms of affecting *industry* demand, the relative effects of marketing expenditures originally made when k products were in the market will tend to decline as k increases. In other words, we hypothesize that the effectiveness of marketing in generating industry sales depends on market structure in a systematic manner.

In our empirical analysis of the anti-ulcer drug market, we obtained considerable but not quite unanimous support for this hypothesis. In particular, normalizing the impact of a monopolist's marketing investments on current sales to unity, we estimated the impact in a duopoly to be 0.6, in a three-product industry to be 0.8, and in a four-product market to be 0.5; these last three numbers are all statistically significantly different from unity (implying that we reject the hypothesis that the effectiveness of marketing efforts is independent of market structure), and from zero (indicating that we reject the hypothesis that once there is competition, the only impact of marketing is on market share, and none on overall market size). Thus our results suggest that in the anti-ulcer drug market there is clear evidence of spillovers, and that these spillovers are considerably less than 100% effective. Moreover, for the most part, these spillovers decline as the number of products in the industry increases.

Second, we find that at the industry level, both cumulative minutes of detailing and cumulative pages of medical journal advertising affect sales; typical estimates of these elasticities are 0.5 and 0.2, respectively. At the market share level, relative sales of products are also positively related to relative cumulative minutes of detailing; this elasticity is typically in the range of 0.7 to 0.9. Together these results imply that the marketing efforts of firms in the anti-ulcer drug market had substantial effects, both in terms of affecting market shares and the size of the overall industry.

Third, a somewhat unexpected result we obtained is that at the industry level, the rate of depreciation of stocks of both minutes of detailing and medical journal advertising was estimated to be zero. We believe that this result reflects the fact that market-expanding marketing primarily involves informing physicians about the usefulness of this class of drugs, and that once a physician begins prescribing these drugs, he/she is not likely to forget about their existence and stop prescribing them. By contrast, at the level of market shares a rather different picture emerges. In particular, in the four-product market (Tagamet, Zantac, Pepcid and Axid), we find that the market share impact of the stock of detailing minutes deteriorated at an annual rate of around 40%, reflecting perhaps a more rivalrous content of marketing efforts.

The remarkable growth in the market share of Zantac over time can be partially explained, then, by the very substantial marketing efforts undertaken by Glaxo. However, pricing policies also had an impact. Zantac increased share over Tagamet in part because the price premium commanded by Zantac declined from about 56% in 1983 to only about 25% in 1993. Our estimates of industry price elasticities range from about -0.7 to -0.9, while estimates of cross-price elasticities between any pair of the four products are about 0.7.

Another set of important factors affecting sales of anti-ulcer drugs concerns product quality attributes. At the industry level, the evidence suggests that the size of the market was enlarged considerably when the FDA granted approval for the GERD indication -- a condition that occurs in a relatively large population. At the market share level, we find that when a product had a GERD approval advantage relative to other products, its market share increased. Thus another reason why Zantac fared so well in the marketplace is that for quite some time it was the only product having received FDA approval for the treatment of GERD. Another variable affecting market share significantly is the number of adverse interactions with other drugs reported to the FDA. Relative to its competitors, on this account Tagamet fared relatively badly (by 1993, Tagamet had 12 drug interactions, Zantac and Axid had only one, and Pepcid had none). Thus Zantac also enjoyed advantages from this product quality characteristic. An unexpected result we obtained, however, was that dosing frequency did not appear to affect market shares in a statistically significant manner.

Finally, we found that, as in many other markets, order of entry effects are very substantial. In particular, holding price, marketing efforts and product quality constant, relative to the n^{th} product, the $(n+1)^{th}$ entrant can expect about 40% lower sales.

The results of this paper are of considerable interest in the current health care reform debate. Critics of the pharmaceutical industry have argued that much detailing is merely aimed at market share, and is socially wasteful. Some have suggested placing ceilings on the marketing activities of pharmaceutical firms, but our findings demonstrate that this could have negative social welfare impacts. The findings in this paper suggest that marketing efforts also play a very important role in the diffusion of information to physicians, although the degree to which this is true probably declines somewhat as the number of products in a market increases. Moreover, our results suggest that in order to overcome pioneer product advantages, later entrants have found it necessary to advertise more intensively. An implication of these results is that if all pharmaceutical firms were constrained in their marketing activities, it is possible that the benefits would accrue primarily to the pioneer firms, at the expense of later entrants who would be prevented from trying to overcome pioneer product advantages. Thus, such a policy could have anti-competitive impacts, although it would be consistent with a patent system that rewards innovation.

The research reported in this paper should be extended in a number of ways. First, although the industry and market share equations are plausible and provide important initial evidence on the roles of marketing, price and product quality competition in the anti-ulcer market, the underlying models could be modified in a number of useful ways. The most obvious extension is to reformulate the models within an explicitly dynamic diffusion framework, such as those involving the Gompertz, logistic or other more general diffusion curve formulations. In such a framework, marketing and pricing policies might not only affect the long-run or equilibrium level of demand, but they might also affect the speed at which a long-run equilibrium level is approached.

A second useful extension would involve incorporating data on direct-to-consumer marketing. In 1988 SKB experimented with a "Tommy Tummy" television advertising campaign that was aimed directly at consumers but did not mention Tagamet by name. More recently, Glaxo has advertised in magazines and on television, suggesting that patients with heartburn and acid discomfort should see their physician. These ads are sponsored by the Glaxo Research Institute, and, consistent with FDA regulations on direct-to-consumer advertising, do not mention the Zantac product by name unless the requisite warning and other product information is also fully disclosed. Since these advertisements typically do not mention product name, their impact is more likely to be on industry demand than on market share. Moreover, direct-to-consumer advertising may change the physician-patient information sharing relationship, and therefore could modify the diffusion process. It would be useful to examine whether such effects have actually occurred, and by extension, how effective is direct-to-consumer marketing in the anti-ulcer marketplace.

Third, and perhaps most importantly, the findings of this paper suggest interesting topics in the theory of industrial organization. What is the optimal marketing strategy for firms when there are spillovers and marketing activities have long-lived impacts? What is the correspondingly optimal pricing behavior? How does this optimal behavior vary with market structure? How is the optimal behavior affected by federal tax provisions that allow the expensing (rather than amortizing) of long-lived marketing investments? What are the implications for social welfare?

Obviously, much remains to be done. We believe we have demonstrated quite clearly that marketing efforts are very important in understanding the diffusion and economic success of new products. Product quality and pricing behavior have also been shown to play important roles in the diffusion process. We hope the results of this study contribute to this and other related research projects that enrich our understanding of the economics of new products.

References

Berndt, Ernst R. and David O. Wood [1979], "Engineering and Econometric Interpretations of Energy-Capital Complementarity," *American Economic Review*, Vol. 69, No. 3, June, 342-354.

Bresnahan, Timothy F. and Peter C. Reiss [1990], "Entry in Monopoly Markets," *Review of Economic Studies*, Vol. 57, October, 531-553.

Bond, Ronald S. and David F. Lean [1977], *Sales, Promotion and Product Differentiation in Two Prescription Drug Markets*, Staff Report to the Federal Trade Commission, Washington DC: Federal Trade Commission, Bureau of Economics, February.

Cearnal, Martin E. [1992], "Medical Marketing Communications Today: Use and Abuse," in Dev S. Pathak, Alan Escovitz and Suzan Kucukaslan, eds., *Promotion of Pharmaceuticals: Issues, Trends, Options*, Binghamton, NY: Haworth Press for Pharmaceutical Products Press, 23 -32.

Cocks, Douglas L. [1975], "Product Innovation and the Dynamic Elements of Competition in the Ethical Pharmaceutical Industry," in Robert B. Helms, ed., *Drug Development and Marketing*, Washington, DC: American Enterprise Institute for Public Policy.

Cocks, Douglas L. and John R. Virts [1974], "Pricing Behavior of the Ethical Pharmaceutical Industry," *Journal of Business*, Vol. 47, July, 349-362.

Diewert, W. Erwin [1981], "The Economic Theory of Index Numbers: A Survey," in Angus Deaton, ed., *Essays in the Theory and Measurement of Consumer Behavior in Honor of Sir Richard Stone*, Cambridge: Cambridge University Press, 163-208.

Diewert, W. Erwin [1992], "Fisher Ideal Output, Input, and Productivity Indexes Revisited," *Journal of Productivity Analysis*, Vol. 3, 211-248.

Dorfman, Robert and Peter O. Steiner [1954], "Optimal Advertising and Optimal Quality," *American Economic Review*, Vol. 44, No. 5, December, 826-836.

Fine, Steven N., Andrew J. Dannenberg and David Zakim [1988], "The Impact of Medical Therapy on Peptic Ulcer Disease," in David Zakim and Andrew J. Dannenberg, eds., *Peptic Ulcer Disease and Other Acid-Related Disorders*, New York: Academic Research Associates, Inc., 1-13.

Golder, Peter N. and Gerard J. Tellis [1992], "Do Pioneers Really Have Long-Term Advantages? A Historical Analysis," Cambridge, MA: Marketing Science Institute, September, Report No. 92-124.

Griliches, Zvi and Frank Lichtenberg [1984], "R&D and Productivity Growth At The Industry Level: Is There Still A Relationship?", in Zvi Griliches, ed., *R&D, Patents and Productivity*, Chicago: University of Chicago Press for the National Bureau of Economic Research, 465-502.

Hurwitz, Mark A. and Richard E. Caves [1988], "Persuasion or Information? Promotion and the Shares of Brand Name and Generic Pharmaceuticals," *Journal of Law and Economics*, Vol. 31, October, 299-320.

Kalyanaram, Gurumurthy and Glen L. Urban [1992], "Dynamic Effects of the Order of Entry on Market Share, Trial Penetration, and Repeat Purchases for Frequently Purchased Consumer Goods," *Marketing Science*, Vol. 11, No. 3, Summer, 235-250.

Leffler, Keith B. [1981], "Persuasion or Information? The Economics of Prescription Drug Advertising," *Journal of Law and Economics*, Vol. 24, April, 45-74.

Lynn, Matthew [1991], *The Billion Dollar Battle: Merck v. Glaxo*, London: Mandarin Paperbacks of Reed International Books, Ltd.

McKenzie, Constance A., Ellen S. Underwood, Kim Poinsett-Holmes, and Lynn Graham [1990], "Peptic Ulcer Disease: Therapeutic Options," *U.S. Pharmacist*, October, 53-64.

Montgomery, David B., and Alvin J. Silk [1972], "Estimating Dynamic Effects of Market Communications Expenditures," *Management Science*, Vol. 18, No. 2, June, 485-501.

New York Times [1993], "Drug Company Breaks Tradition With Promotion Focused on Price," *New York Times*, Tuesday, November 9, Page D2, Column 1.

Perloff, Jeff and Valerie Y. Suslow [1994], "Higher Prices from Entry: Pricing of Brand-Name Drugs," Ann Arbor, MI: University of Michigan, School of Business Administration, unpublished manuscript.

Reekie, W. D. [1978], "Price and Quality Competition in the United States Drug Industry," *Journal of Industrial Economics*, Vol. 26, March, 223-237.

Robinson, William T. [1988], "Sources of Market Pioneer Advantages: The Case of Industrial Goods Industries," *Journal of Marketing Research*, Vol. 25, February, 87-94.

Robinson, William T. and Claes Fornell [1985], "The Sources of Market Pioneer Advantages in Consumer Goods Industries," *Journal of Marketing Research*, Vol. 22, No. 2, August, 297-304.

Robinson, William T., Gurumurthy Kalyanaram and Glen L. Urban [1994], "First Mover Advantages for Pioneering New Products: A Survey of Empirical Evidence," *Review of Industrial Organization*, Vol. 9, No. 1, 1 -23.

Samuelson, William and Richard Zeckhauser [1988], "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, Vol. 1, March, 349-365.

Schmalensee, Richard L. [1972], *The Economics of Advertising*, Amsterdam: North-Holland Publishing Company.

Schmalensee, Richard L. [1982], "Product Differentiation Advantages of Pioneering Brands," *American Economic Review*, Vol. 27, 349-365.

Scouler, Bonnie Jean [1993], "A Segmentation Analysis of the Ulcer Drug Market," S.M. Thesis, Alfred P. Sloan School of Management, Massachusetts Institute of Technology, May.

Suslow, Valerie [1993], "Are There Better Ways to Spell Relief? A Hedonic Pricing Analysis of Ulcer Drugs," Ann Arbor, MI: University of Michigan, School of Business Administration. Revision, October.

Urban, Glen L., Theresa Carter, Steve Gaskin and Zofia Mucha [1986], "Market Share Rewards to Pioneering Brands: An Empirical Analysis and Strategic Implications," *Management Science*, Vol. 32, June, 645-659.

Data Appendix: Data Sources from IMS America

We hope that this discussion will serve as a useful reference for economists who will be using IMS sales data on pharmaceuticals in the future, as there are a number of important issues and quirks to the data which are not well documented in IMS literature.

A. U.S. Drugstores Audit (USD) and U.S. Hospitals Audit (USH):

A panel of pharmaceutical wholesalers report to IMS each month on the sales of each presentational form (unit dose syringes, bottles of 100 tablets, etc.) of each drug product (Tagamet, Zantac, etc.) to drugstores and hospitals in the United States. From the sales reports they obtain in this audit, IMS computes national projections of the number of units and the dollars of revenue of each presentational form of each product sold each month in the United States, separately for drugstores and for hospitals. In recent years, the panel has grown to encompass nearly the entire universe of pharmaceutical wholesalers, according to IMS, making the audit nearly a full census, and the projections therefore quite accurate.

One interesting feature of all of the IMS data used in this study is that although IMS has been collecting such data for decades, the company keeps computer records of only the immediate past six years, on a rolling basis. In order to have the opportunity to study the anti-ulcer market since its very inception, which dates back over fifteen years, we chose to type in numbers by hand from archived monthly IMS publications. The sales data from January 1986 through December 1981 comes directly from IMS computer records, but all other IMS data used in this study was hand-typed.

Because the sales data contained so many different numbers (quantities and revenues each month for each presentational form of each drug for a total of over 5200 hand-typed numbers in the fifteen-year sample, above and beyond the 8000 numbers provided by IMS in computer format), and because the original copies of the published data were often very difficult to read (often the numbers were available only on poor-quality microfilm, where a 3 was indistinguishable from an 8), we deemed the possibility for error to be very high. We therefore chose to invest several months in ensuring the integrity of the hand-typed data by carefully checking it for typographical errors. It turns out that there is a reasonable degree of variation from month to month in the sales quantities and revenues for each individual drug presentation (variation often on the order of 10% or more), but the prices of the drug presentations (IMS-reported revenues divided by IMS-reported units) are relatively stable. Therefore, our method of error correction was to sort the data by presentational form of each drug, and then print separate graphs of the drugstore and hospital prices of each presentation as a function of time. We were easily able to spot potential typographical errors as outliers on these graphs, at which point we were able to correct the errors by checking them against copies of the original published data. (Unfortunately, we had to make more than one trip back to Philadelphia in order to obtain copies of data pages which were missing from our collection! It was easy to lose a page, or miss it in the first place, because our data was obtained from dozens of three-inch thick monthly volumes of printed data, or their microfilm equivalents, in which the data of interest was contained on just a page or three in the middle of each hefty tome.) In all, we corrected a few dozen serious errors on the approximately 100 graphs printed, but as a result we are now quite confident of the reliability of the data, to the extent that it accurately matches the data collected by IMS.

Nevertheless, additional manipulations remained to be performed on this data set in order to put it into a form that would be useful for this study. Details of those manipulations follow.

First, as noted earlier, there were multiple presentational forms of each drug sold. To obtain a single number describing the quantity of each brand of drug (e.g. Tagamet) sold in a given month, we summed up the total number of milligrams of the chemical sold that month. For example, if in August 1979 SmithKline sold 6200 bottles of 100 Tagamet 300mg tablets and 1600 packages of 10 unit-dose containers of 10ml of Tagamet syrup at 5mg/ml concentration, then we would compute the total number of milligrams of Tagamet sold that month as:

$$(6200)(100)(300\text{mg}) + (1600)(10)(5\text{mg/ml}) = 1{,}868{,}000 \text{ mg}.$$

An alternative approach to constructing a single monthly sales series for each drug, and one which a number of other studies have adopted, would be merely to proxy a drug's total sales (units and revenues) by the sales of a single leading presentation. The advantage of this alternative approach is its computational simplicity; by contrast, our method required dozens of additional hours of data manipulation. However, there is a serious disadvantage to the simpler approach, especially in this ulcer market, which is that the leading presentation changes over time. For example, see Figure A1, which displays the sales of Tagamet over time, broken down by its four major product forms (note that even this is a simplification of the full sales data set, as, for example, the portion of the graph corresponding to Tagamet 800mg tablets represents a sum of two different presentational forms: bottles of 30 tablets, and unit-dose packages of 100 tablets). From this graph, we see that although originally Tagamet was sold only in the 300mg form, by 1992 the 400mg form had become the "leading form," considerably overtaking the 300mg sales. Other drugs in the sample present similar problems, having more than one presentation which hold significant shares of the drug's total sales.

Also, we chose to include only those presentational forms which were intended to be taken orally by patients: tablets, capsules, and oral liquids. Excluded were those forms packaged in vials, minibags, syringes, etc., for injection or intravenous administration. One reason for this decision is that we intended to concentrate our study mainly on the drugstore market, where the bulk of the

Figure A1. Tagamet Drug Store Sales

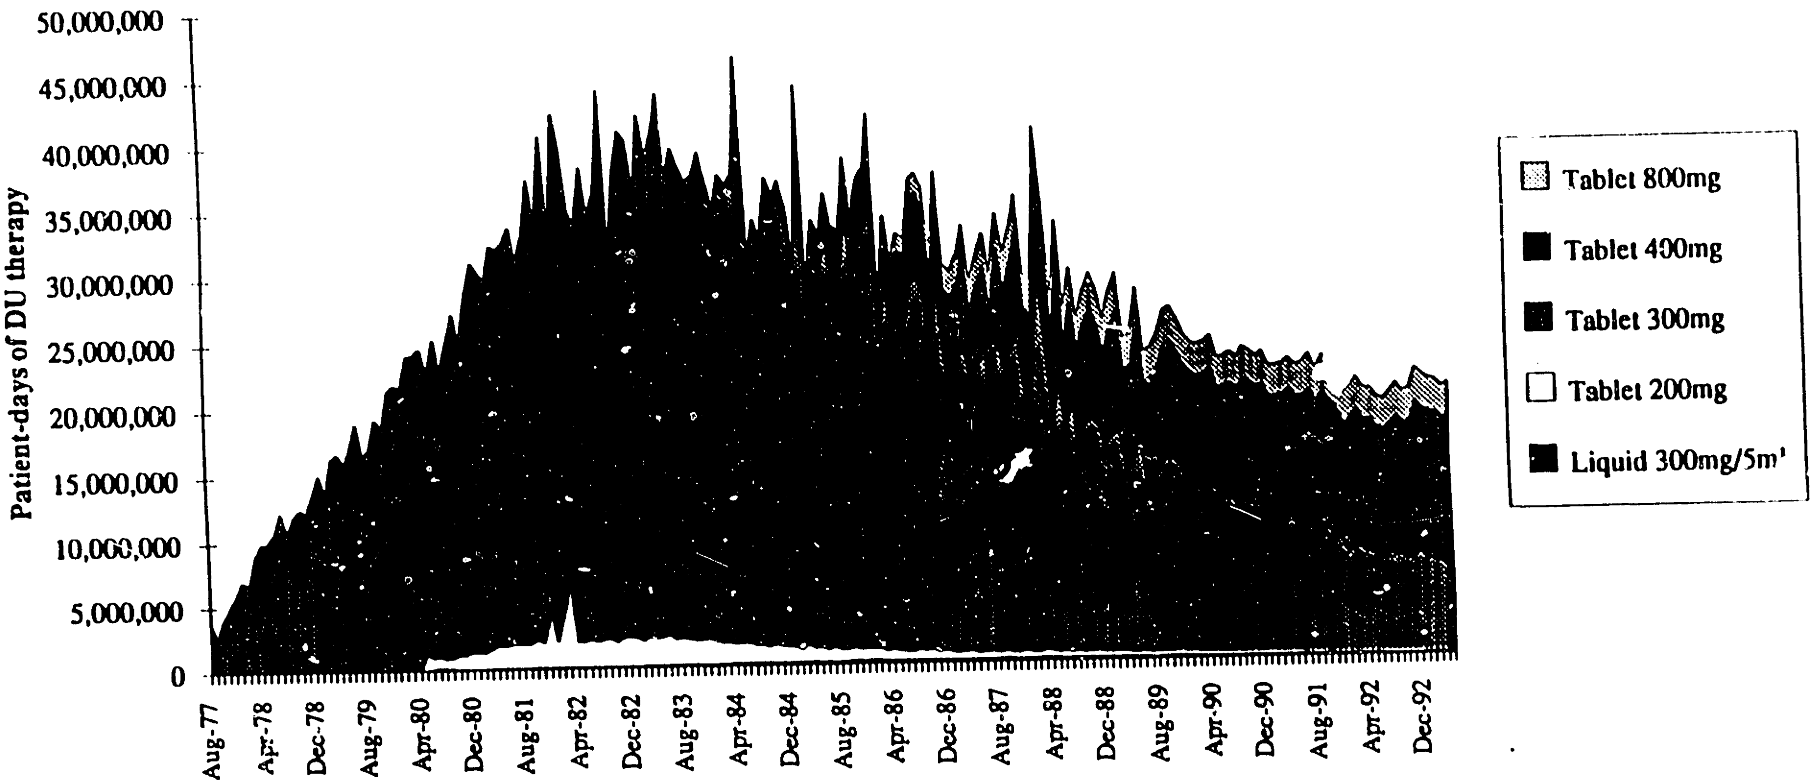

anti-ulcer sales occur, and where detailing to physicians is most salient. By contrast, the non-oral preparations are developed mainly for hospitals, although some non-oral sales also show up in the drugstore market, presumably with the end consumers being patients either in nursing homes or under hospice care. (For the twelve-month period ending in May 1993, drugstore sales revenues for non-oral presentations of H_2-antagonist s were less than one thousandth as much as revenues for oral presentations. Even in the hospital market, non-oral presentations brought in less than half as much revenue as the oral presentations during that time period.) A second, very substantive reason for including only the oral preparations is that we learned, from conversations with doctors and pharmaceutical marketing professionals, that the non-oral preparations are generally used for very different purposes: instead of healing painful ulcers in otherwise healthy people, as the tablets and capsules are intended, the intravenous administration of anti-ulcer medication is used mainly for the prevention of ulcers in emergency-room patients at risk for ulcers due to the increased acid secretion brought on by trauma, for example, and in patients who are at risk for ulcers due to regimens of large doses of non-steroidal painkillers. Anti-ulcer medication may also be injected as part of a complete anesthesiology for surgery. These uses require very different numbers of milligrams of drug than do the standard therapies (DU, GU, DU maintenance, and GERD) that are usually administered orally, and the price per milligram of drug tends to be an order of magnitude higher for the intravenous preparations (likely a combination of two effects: price discriminations, and the more complicated packaging and storage requirements of the IV preparations). So rather than confound the two types of uses, we have chosen to define our market of interest to be the orally-administered anti-ulcer drugs.

Next, we had to find a way to make the quantity units comparable across drugs. Milligrams were not an appropriate unit for comparison, because, for example, treatment of an active duodenal ulcer with Tagamet requires 800mg of drug to be ingested per day, but an equivalent therapy with Pepcid requires only 40mg of drug. Each drug is a different chemical entity, with different molecular weights, different rates of absorption, and different rates of binding to bioactive sites in the body, which combine to cause wide variation in the amounts of mass of drug that must be consumed to achieve the same desired effect. Because marginal manufacturing costs in the pharmaceutical industry are generally much lower than prices, we have chosen to concentrate on the demand side of the market in our choice of quantity units: patient-days of therapy. (This may have some concordance with the producer side as well, for although the different chemicals may not have the exact same marginal costs of synthesis, it is at least plausible to assume that packaging the drug into tablets, and the tablets into bottles, should have approximately the same marginal cost per tablet, regardless of the chemical being so packaged.) This choice of quantity units considers 800mg of Tagamet to be the same mount of drug as 40mg of Pepcid, for purposes of computing sales levels and market shares, since these quantities are therapeutically equivalent.

The quantity of patient-days of therapy of a drug sold in a given month is equal to the total number of milligrams sold divided by the number of milligrams per day of active duodenal ulcer therapy for that particular drug (in the case of Cytotec, which is not indicated for active DU therapy, we instead used the daily recommended dosage for NSAID-induced ulcer prevention). Thus, continuing our earlier example, we would find that in our hypothetical month, there were sold:

(1,868,000mg)/(800mg/day) = 233,500 patient-days of therapy.

The number of milligrams per day of therapy used for our quantity conversions were as follows for each drug:

Table A1: Number of milligrams per day of therapy, used for quantity conversions.

DRUG	MILLIGRAMS PER DAY OF THERAPY
TAGAMET	800
ZANTAC	300
PEPCID	40
AXID	300
PRILOSEC	20
CARAFATE	400
CYTOTEC	0.2

Defining our quantity unit to be the total number of milligrams divided by the standard dosage in milligrams per day is, unfortunately, not without problems. First, the same drug may be used for slightly different therapies, and it may be taken in different dosages for the different purposes. For example, Zantac may be prescribed at a dosage of 300mg per day (either 300mg once daily, or 150mg twice daily) for active duodenal ulcer, gastric ulcer, or GERD, but its recommended dosage for duodenal maintenance therapy is only 150mg per day, half of that required for the other therapies. Each of the H_2-antagonist s has a similar prescribing regimen for those four different indications. Therefore, our quantity measures are not literally the number of patient-days of therapy being consumed, but rather the number of patient-days of therapy which would be consumed of all of the sales were for treatment of active DU. Second, while we have assumed that the milligram dosage required for DU therapy remained constant over time, this was not the case for Tagamet. At the time of its introduction in 1977, the recommended dosage for DU therapy was 1200mg per day (300mg, for times daily), but subsequent experimentation showed that lower doses could be just as effective for ulcer healing, and by 1988 the recommended dosage was only 800mg per day (either 800mg once daily or 400mg twice daily). We have taken the approach that a milligram of Tagamet in 1977 is the same quantity as a milligram of Tagamet in 1990, despite the fact that people may have been consuming fewer milligrams on average in the later years for the same length of treatment. Since we have no way of knowing how many DU patients were taking 1200mg of Tagamet versus 800mg of Tagamet at any point in time (the choice between the two depended upon the vagaries of individual doctors' prescribing habits), we feel that we have chosen the most appropriate way to proceed.

A final modification which needed to be made to the sales data concerns the fact that the data collection from pharmaceutical warehouse invoices has, at different times during the sample, been rather lumpy. This problem introduces seasonal noise into the data, which can be eliminated by rescaling the sales and revenue figures. For purposes of rescaling, there are three distinct periods in our sample. Until December of 1980, the sales audit was actually conducted at a sample of pharmacies rather than at warehouses, and there was no lumpiness to the data, so no rescaling was required. From January 1981 to December 1989, the data were apparently (according to the best information we could obtain from IMS, whose data specialists are not accustomed to answering questions about historical data) reported from warehouses on the bases of full weeks, so some months could contain four weeks of data, while others contained five. This causes large month-to-month variations in the sales data, which is obviously inappropriate for a detailed monthly analysis of the competitive effects of price and advertising on sales. A lengthy investigation has failed to reveal an appropriate way to rescale the data to correct for these fluctuations. (Based on conversations with IMS representatives, we tried several possibilities, such as rescaling the data by the number of Wednesdays in each month, but none turned out to be correct.) Thus our best approximation to the truth is that the sales data for this period of time contain a component of stochastic measurement error. In the third period of the sample, from January 1990 to the present, the number of reporting weeks per month were standardized so that the first four weeks of the year were designated as January, the next four weeks as February, the next five weeks as March, and repeated in a 4-4-5 pattern each quarter. (The single exception is December 1991, which for accounting purposes - there are not exactly 52 weeks in each year - was designated as a month of 6 weeks rather than 5.) To rescale the data for our purposes, we divided the IMS sales figures in each month from January 1990 to the end of the sample by the number of reporting weeks in that month, and then multiplied by 4.33 in order to retain the same normalization of physical units as in the original IMS data.

Finally, in order to transform the nominal prices from the IMS data into real prices, we deflated by the Consumer Price Index (1982-84 base years).

B. National Detailing Audit (NDA)

The National Detailing Audit (which as of 1993 has been subsumed by the Integrated Promotional Services, Office Contact Report) is a service that collects data from a nationwide panel of doctors about the visits which have bee paid to them by pharmaceutical sales representatives. The doctors participating in the panel keep a log of the number of minutes they spend talking to detailers on each detail visit. If the detailer talks to the doctor about more than one product (for example, a Lilly detailer might discuss both Axid, an ulcer drug, and Ceclor, an antibiotic, with a family doctor), the physician makes an estimate of how many minutes were devoted to each product. From this panel, IMS then reports nationally projected estimates of the number of details and the number of minutes spent detailing each product, each month.

The detailing data series, unlike the sales data series, consist of only one observation per month, since detailing is performed at the level of the drug brand, rather than at the level of the presentational form. This fact made typographical errors much less of a problem than in the sales data, despite the fact that we had to manually enter the monthly detailing data for every month in the more than fifteen years of our sample. We collected monthly data on details and minutes for our seven drugs of interest, as well as for the total number of details and minutes done by each of the manufacturers producing these drugs (across all of their products) and for the total number of details and minutes in the entire U.S. pharmaceutical industry. These last two types of data are intended to be used as instruments for brand detailing, which is a potentially endogenous variable.

Beyond typographical errors, there were still some corrections to be made. In 1986, IMS expanded its panel of doctors from 1400 physicians reporting two weeks of every month to 2800 physicians reporting full months. Concurrently, they changed their projection methodology, and it turns out that a scaling factor of 0.74 must be applied to the data for all months prior to January 1986 in order to make it comparable to the data for January 1986 through December 1992. (In suggesting this scaling factor, IMS cautions that it is much more confident in its ability to measure the relative shares of detailing by different products than its ability to measure absolute levels. Nevertheless, we assume that after applying the recommended transformations, we can make reasonably accurate comparisons of the levels of detailing in different periods).

A second change occurred in January 1993, when IMS significantly increased the breadth of coverage of detailing data. Under the newly created Integrated Promotional Services, there exist a wide variety of reports, including the Office Contact Report, which is the most directly related to the now-defunct National Detailing Audit. In the Office Contact Report, there are now reported many more *types* of details, including sample drops, educational visits, service visits, and telephone calls, than it did before, so the data on details and minutes are not easily comparable.

We were able to construct a measure of the number of details for 1993 which would be comparable to the pre-1993 years, by looking at the new breakdown of details into the various new IMS categories, and counting as details only those visits which were either "full discussion" details or "brief mention" details, which is what IMS considers to be the "traditional" details that doctors were intended to include in their reports for the NDA prior to 1993. Although similarly disaggregated information on minutes of detailing are not readily available in IMS's printed reports, we were able to match up 1993 on minutes of detailing with the pre-1993 data on minutes by special arrangement with IMS, who provided us with computer-generated reports from their database on the number of minutes of detailing in 1993 devoted to full discussions and brief mentions.

C. National Journal Audit (NJA)

In this audit, IMS performs a complete census of advertising in medical journals. They subscribe to every known medical journal and examine every advertisement in every issue of each journal. They note the number of whole and partial pages, the number of colors used in printing the ads, the location of the advertisement in the journal (for example, if it was found at the very front, or if it was printed on the back cover, in either case getting more exposure than an ad buried in the middle of the publication), and other attributes that affect the cost of placing an advertisement. Then, using standard rate sheets, they compute the cost of each of the advertisements placed. Reported in the NJA monthly report are the total number of pages of advertising published for each product in that month (weighting all journals equally, regardless of circulation or professional influence), as well as the total estimated cost of all medical journal ads for that product.

We consider the cost figure to be the most accurate single measure of the amount of medical journal advertising done for a particular product, because (assuming that medical journal advertising is close to being a competitive industry), the prices of the ads reflect the reach of the advertisements, in terms of number of doctors reached, visual impact of the advertisement (through color, for example), etc. These cost figures are reported in nominal dollars, so to obtain a real measure of medical journal advertising effort, we deflate these series by the PPI for Advertising in Professional and Institutional Periodicals (BLS product code 2721-415).

As with the detailing data, the series we collected from the NJA include monthly series for each drug product in our sample, as well as monthly series on total monthly advertising for each manufacturer producing one of the products in our sample and the total monthly advertising by the pharmaceutical industry as a whole.

Chapter 4:

Information, Marketing, and Pricing in the US Anti-Ulcer Drug Market

(co-authored with Ernst R. Berndt, Linda T. Bui, and Glen L. Urban)

1 Background

There are two cost conditions that have considerable bearing on the structure and behavior of the pharmaceutical industry. First, sunk costs are very large. In particular, the costs of bringing a product to market -- doing basic research, winning patent approval, engaging in development, performing clinical trials, and obtaining final approval from the Food and Drug Administration (FDA) -- are currently estimated at about $360 million per drug. Second, for most traditional pharmaceutical products, the marginal costs of manufacturing are very small. Although appropriate cost data are not publicly available, it is not uncommon for generic drugs to sell at 25-30% of the pre-patent expiration price. Informal discussions with industry officials suggest that for the H_2 drugs, production costs are only about 10% of the price.

These cost conditions have implications for pricing. Patent protection gives firms the ability to influence price, and to the extent one is willing to use the Lerner markup relation as a pricing rule of thumb, one would expect price-marginal-cost conditions to approximate $(P - MC)/P = -1/\varepsilon$, where ε is the demand price elasticity. With manufacturing costs at 10-25% of price (markups 75-90%), the implied demand price elasticity would range from -1.1 to -1.3. However, elasticities of that size contrast with the common perception that demand for prescription drugs is extremely price inelastic. Peter Temin [1980, Ch. 5], for example, notes that traditionally physicians have been relatively unaware of drug prices. Other observers have suggested that moral hazard due in the form of third-party (insurance) payment practices also contributes to low price responsiveness. Very little econometric evidence on demand elasticities for drugs is available, in part because the traditional consumer demand paradigm (utility maximization, marginal rates of substitution equal to relative marginal prices, etc.) cannot be expected to describe behavior adequately in a market in which principal-agent and moral hazard problems (stemming from relationships among physicians, patients and insurers) are widespread.[1] In this paper we report elasticity estimates viewed from the vantage of the firm, not the "consumer" -- whoever that may be.

Since marginal production costs are so very small, enhancing revenues is essentially the same as increasing profits, and thus drug firms face strong incentives to shift out the demand curves. Thus it is not surprising that marketing-sales ratios are quite high in the pharmaceutical industry. The largest component (70-80%) of marketing has traditionally involved detailing to physicians; it consists of a company representative providing as much product information as possible to physicians, given the typical short time of the visit (3-10 minutes) and the content regulation enforced by the FDA. Medical journal advertising is also carried out, but is less extensive than detailing. Finally, in the last few years pharmaceutical companies have increasingly employed direct-to-consumer advertising in various media.

The information content of marketing efforts deals primarily with product differentiation and non-price aspects. In the H_2 market, five quality attributes are of particular importance.[2] First, the various H_2 drugs are viewed as being roughly similar in efficacy (the four- to six-week treatment healing rate is about 70-80% for duodenal ulcer patients), although there is some evidence suggesting that Zantac has a significantly lower relapse rate than does Tagamet for patients on duodenal maintenance treated at recommended dosages.[3] Second, less frequent dosages are thought to enhance patient compliance. When Zantac entered in 1983, its twice daily product attributes were compared negatively with the four times a day regimen recommended for Tagamet. Tagamet responded with a twice-a-day version in late 1984, after which considerable rivalry ensued; today all four H_2 drugs have a once-a-day version. A third quality attribute involves adverse interactions with other drugs. Here Tagamet has been on the defensive, for early on it was discovered that Tagamet interacted with the liver system in a way that could affect the metabolism of other drugs. As of 1994, Tagamet had reported to the FDA significant drug interactions with ten other drugs, whereas Zantac and Axid had only one reported drug interaction, and Pepcid had none. A fourth quality characteristic involves side effects. Here again Tagamet has been somewhat on the defensive, for conditions such as dizziness in the elderly and gynecomastia (breast swelling) for males are apparently not as prevalent with Zantac, Pepcid and Axid.

---

[1] See, however, Baye, Maness and Wiggins [1994].
[2] For more extensive discussion, see Chapter 3.
[3] See K.R. Gough et al. [1984].

Finally, the four products compete in terms of medical conditions (indications) for which the FDA has granted treatment approval. Although Tagamet was the first to win approval for duodenal ulcers, duodenal ulcer maintenance, and gastric ulcers, in 1986 Zantac was the first to obtain approval for gastroesophageal reflux disease (GERD), a rather common condition that varies from modest heartburn and acid indigestion to being very painful. The FDA permits marketing only for approved indications. Although Tagamet obtained FDA approval for GERD in 1991, and even though Tagamet had very similar effects to Zantac suggesting it would likely also be effective in treating GERD, not having FDA approval for GERD whereas Zantac did may have constituted a significant marketplace disadvantage for Tagamet.

In terms of pricing, at entry Zantac was priced at a significant 80% premium over Tagamet, but by May 1994 this premium had gradually declined to 19%. In May 1994, the price per day's treatment (to drug stores) was \$2.61 for Zantac, \$2.56 for Axid, \$2.30 for Tagamet and \$2.17 for Pepcid; quantity shares for the four products were 49%, 12%, 22% and 17%, respectively.

To understand the roles of marketing, pricing and quality attributes in explaining the growth and changing composition of the H_2 market, we now outline an econometric model first for the H_2 industry as a whole, and then for the market shares garnered by the four H_2 drugs.

2 An Econometric Model of the H_2 Market

At the industry level, we expect the quantity demanded (number of patient days of duodenal ulcer therapy) to depend on price per treatment day, various marketing efforts, and quality attributes. Since marketing efforts provide long-lived information, it is important that cumulative information stocks be distinguished from current period new information flows. Define the cumulative marketing information stock S_t at end of month t as

$$
\begin{aligned}
S_t &= (1-\delta)S_{t-1} + F_t \qquad (1)\\
&= \sum_{\tau=0}^{t} (1-\delta)^{t-\tau} F_{t-\tau}
\end{aligned}
$$

where F_t is the flow of new marketing information efforts during month t, and δ is the monthly depreciation rate. Since is unknown, we estimate it econometrically. In terms of marketing efforts, we distinguish three instruments: the minutes of detailing to physicians (DET), the number of pages of medical journal advertising (PJL), and the target rating points of direct to consumer advertising (DCA).[4] It is worth noting that the H_2 DCA efforts did not mention any H_2 product by name, but only encouraged viewers to seek advice from their physician if they experience heartburn and acid indigestion.

Although such DCA advertising is plausibly intended to augment overall industry demand, when two or more products exist marketing efforts are often only focussed on a particular brand. During its monopoly era Tagamet recouped all the benefits of its marketing efforts (it had 100% market share). However, once Zantac entered, even though rivalry between Tagamet and Zantac was intense, some of Tagamet's marketing efforts might have spilled over to the benefit of Zantac, and vice versa. Similarly, once Pepcid and Axid entered, while marketing efforts were typically focused on specific brands, spillovers to Zantac and Tagamet might have occurred. To allow for marketing spillovers affecting industry (rather than just product-specific) demand, we define the effective industry marketing stock S^*_t as a weighted sum of the marketing information stocks originally formed in various market structures:

$$S^*_t = \mu_1 S_{1t} + \mu_2 S_{2t} + \mu_3 S_{3t} + \mu_4 S_{4t}, \tag{2}$$

where S_{1t} is the surviving marketing information stock at end of month t that originally accumulated in the Tagamet monopoly era, S_{2t} is the similar stock formed during the Tagamet-Zantac duopoly, S_{3t} is that from the Tagamet-Zantac-Pepcid triopoly, and S_{4t} is that from the Tagamet-Zantac-Pepcid-Axid rivalry. Since in a monopoly *all* marketing efforts affect industry demand, we normalize the μ's by setting $\mu_1 = 1$. Several interesting hypotheses involve the μ's.

First, if the effectiveness of firms' marketing on *industry* sales is independent of market structure, then $\mu_2 = \mu_3 = \mu_4 = 1$. Second, if in the presence of competition marketing efforts only affect market shares and have a zero-sum impact on industry demand, then $\mu_2 =$

---

4 Target rating points is defined as the target reach (the percent of the over age 35 population who view the message over the course of the ad campaign) times the frequency, where frequency is the number of times the average target individual views the message. For further discussion, see Philip Kotler [1991, pp. 606-608]. The proprietary DCA data were kindly provided us by Lowe & Partners/SMS in cooperation with Glaxo, Inc.

$\mu_3 = \mu_4 = 0$. Finally, if the industry sales-augmenting effects of firms' marketing decline as the number of products in the industry increases, then $1 > \mu_2 > \mu_3 > \mu_4 > 0$.

For our industry demand equation, we specify a log-log model, where Q_t is quantity, P_t is CPI-deflated price, DET^*_t, PJL^*_t, and DCA^*_t are the effective industry stocks defined in (1) and (2), and DGERD is a dummy variable taking on the value of one following FDA approval for GERD:

$$LNQ_t = \beta_0 + \beta_1 LNP_t + \beta_2 LNDET^*_t + \beta_3 LNDET^*_t + \beta_4 LNDCA^*_t + \beta_5 DGERD_t + \varepsilon_t. \quad (3)$$

Since the effective industry marketing stocks depend on the μ's and δ's in complex ways, and since marketing efforts, pricing and quantity demanded are likely to be jointly determined (see Richard Schmalensee [1972]), we estimate parameters in Eq. (3) by nonlinear two-stage least squares (NL-2SLS).[5]

Our econometric model of market shares follows Glen Urban et al. [1986] in specifying variables relative to the incumbent (Tagamet). In particular, using a log-log framework, we specify that in month t, demand quantities of product j relative to the incumbent [LNQJ1 = ln (Q_j/Q_1), with j = Zantac, Pepcid, or Axid] depend on relative prices LNPRJ1, relative detailing, journal pages and direct to consumer marketing stocks LNJPJ1, LNDTJ1, and LNDCJ1, the number of adverse drug interactions for product J relative to Tagamet (LNINTJ1),[6] a discrete variable DSGERD indicating whether product J has a GERD indication advantage relative to Tagamet (1 if an advantage, 0 if no advantage, -1 if a disadvantage), an order of entry variable ENTRY taking on the value of 2 for all Zantac observations, 3 for Pepcid and 4 for Axid, and an AGE variable indicating the number of months product J has been in the marketplace. Again, an instrumental variable procedure is employed to allow for simultaneity.

Our data sources are described more fully in Chapter 3.[7] The direct to consumer marketing data is that for a campaign begun by Glaxo (the manufacturer of Zantac) in June 1992, and it extends through May 1994.

---

[5] As instruments, we employ the PPI for intermediate goods, production worker wages in the pharmaceutical industry, cumulative marketing efforts by the four companies on non-H_2 products for each of the three instruments, and time.

[6] To accommodate zeros, 1.0 is added to both the DCA and the INT variables.

[7] Here we extend the data base to May 1994. Data on prices, quantities, details and journal pages are from IMS International.

3 Econometric Results

Based on 201 monthly observations from September 1977 through May 1994, we estimated parameters of the industry Eq. (3) by NL-2SLS. To be parsimonious in parameters, we constrained the μ's and δ's to be the same for the DET and PJL marketing stocks, but allowed to differ for DCA. The preferred model was chosen based on the lowest value of the NL-2SLS criterion function.

Our estimated H_2 *industry* price elasticity is -0.689 (t=3.80), while elasticity estimates for the DET, PJL and DCA surviving stocks are 0.553 (t=7.52), 0.198 (2.79) and 0.008 (2.67).[8] Hence, industry demand is positively affected by all three of the firms' marketing instruments, but DET is most effective; the sum of the three marketing elasticities is 0.759, suggesting decreasing returns to scale. In terms of spillovers, the estimates of μ_2, μ_3 and μ_4 are 0.601 (6.59), 0.924 (5.30) and 0.410 (4.00); these μ's are jointly significantly different from one, and from zero, indicating that marketing spillovers occur, and that the effectiveness of firms' marketing efforts on industry sales depends on market structure. However, the extent to which spillovers occur does not decline monotonically with the number of products in the market. The DGERD dummy variable coefficient is 0.104 (3.20), indicating that FDA approval for GERD increased the size of the H_2 market by about 10%. Finally, the NL-2SLS criterion function is optimized at the point where for the DET and PJL stocks is 0.00, while that for the DCA stock is 0.15, implying an annual DCA deterioration rate of about 80%. Although we are somewhat surprised that the information stocks of DCA and PJL marketing do not depreciate at all, we note that a similar = 0 finding in the context of R&D knowledge stocks has been reported by Zvi Griliches and Frank Lichtenberg [1984].

Turning now to econometric results based on the market share model, we obtained the following NL-2SLS results, based on 291 monthly observations:

$$
\begin{array}{lllll}
LNQJ1_t = & -0.427\,ENTRY & -0.667\,LNPRJ1_t & +\ 0.649\,LNDTJ1_t & +\ 0.167\,LNJPJ1 \\
 & (44.00) & (8.95) & (19.77) & (6.31) \\
 & +0.052\,DSGERD_t & -\ 0.252\,LNINTJ1 & +\ 0.010\,AGE & \\
 & (2.17) & (9.00) & (16.65) &
\end{array}
$$

[8] The equation R^2 is 0.995, and the Durbin-Watson is 1.912.

with an R^2 of 0.983. Order of entry effects are negative and strong, implying significant first-mover advantages, consistent with evidence from other markets (see Urban et al. [1986]). The within-H_2 price elasticity estimate is -0.67 and significant, while coefficients on relative stocks of surviving detailing (0.649) and journal pages (0.167) are positive and significant. The estimated monthly depreciation rate for the DET and PJL stocks is 0.030 (t = 13.77), implying that relative information marketing stocks deteriorate at about 30% per year. With respect to quality variables, the DSGERD coefficient is 0.05, while that on relative adverse drug interactions is -0.25, suggesting that Tagamet's market share was significantly negatively affected by its GERD and adverse drug interactions disadvantages in the H_2 market. Finally, the age coefficient is positive and significant, implying that, *ceteris paribus*, longevity in the marketplace positively affects market shares.[9]

4 Concluding Remarks

We have reported results on factors affecting the growth and composition of the H_2 drug market. With an H_2 industry own-price elasticity of -0.69 and between-H_2 price elasticities of -0.66, the implicit brand-specific own-price elasticities in May 1994 are -0.80 for Tagamet (standard error of 0.08), -1.03 (0.12) for Zantac, -0.76 (0.08) for Pepcid, and -0.74 (0.08) for Axid. Except for Zantac, these elasticity estimates are still slightly smaller than the -1.1 to -1.3 values one might expect based on the Lerner markup rule of thumb; nevertheless they are not far from one, and clearly differ from zero. It is worth noting that when marketing variables are omitted from the relative demand equations, price elasticity estimates fall to about half these values. We find that marketing information stocks positively affect sales, that the sales elasticity is largest for detailing, followed by journal pages, and is smallest for direct-to-consumer advertising. Marketing information appears to display overall decreasing returns to scale. We also find that order of entry effects are significant, as are quality attributes.

[9] Although DCA is arguably intended to affect industry demand rather than market shares, when the DCA information variable is added, shares of Tagamet and Axid were positively affected relative to those of Zantac and Pepcid.

References

Baye, Michael R., Robert Maness and Steven N. Wiggins, "Demand Systems and the 'True' Cost of Living for Pharmaceuticals," College Station, TX: Texas A&M University, Department of Economics, Working Paper, May 1994.

Berndt, Ernst R., Linda Bui, David Reiley and Glen Urban, "The Roles of Marketing, Product Quality and Price Competition in the Growth and Composition of the U.S. Antiulcer Drug Industry," Cambridge, MA: National Bureau of Economic Research, Working Paper 4904, October 1994.

Gough, K.R., M.G. Korman, K.D. Bardhan, F.I. Lee, J.P. Crowe, P.I. Reed and R. N. Smith, "Ranitidine and Cimetidine in Prevention of Duodenal Ulcer Relapse," *The Lancet*, September 22, 1984, vol. 2, pp. 659-662.

Griliches, Zvi and Frank Lichtenberg, "R&D and Productivity Growth at the Industry Level: Is There Still A Relationship?," in Zvi Griliches, ed., *R&D, Patents and Productivity*, Chicago: University of Chicago Press, 1984, pp. 465-502.

Kotler, Philip, *Marketing Management*, Seventh Edition, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1991.

Schmalensee, Richard L., *The Economics of Advertising*, Amsterdam: North-Holland Publishing Company, 1972.

Temin, Peter, *Taking Your Medicine: Drug Regulation in the United States*, Cambridge, MA: Harvard University Press, 1980.

Urban, Glen, L., Theresa Carter, Steve Gaskin and Zofia Mucha [1986], "Market Share Rewards to Pioneering Brands: An Empirical Analysis and Strategic Implications," *Management Science*, June 1986, vol. 32, pp. 645-659.