The Pennsylvania State University

The Graduate School

Department of Economics

ESSAYS IN AUCTIONS AND COLLUSION

A Thesis in

Economics

by

Steven P. Schulenberg

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

December 2003

We approve the thesis of Steven P. Schulenberg.

Date of Signature

Robert C. Marshall
Professor of Economics
Head of the Department of Economics
Thesis Advisor
Chair of Committee

Herman J. Bierens
Professor of Economics

Ronald L. Filippelli
Professor of Labor Studies and
Industrial Relations

J. Tomas Sjöström
Professor of Economics

ABSTRACT

This thesis is composed of four chapters.

The first chapter is titled "The Impact of Cost Synergies on Bidding in the Georgia School Milk Market," co-authored with Robert C. Marshall, Matthew E. Raiff, and Jean-Francois Richard. Each summer milk processors around the country participate in sealed bid procurements for the right to provide public schools with milk throughout the subsequent academic year. School district contracts are an important part of vehicle routing problems that milk processors solve on an ongoing basis. There are allegedly substantial cost savings for a milk processor from servicing a district that is directly adjacent to one they already service. In this paper, following the work of Krishna and Rosenthal (1996), we construct a procurement model allowing for cost synergies. The equilibrium bid function maps directly into an empirical specification. Using data from a time period when bidders were allegedly acting non-cooperatively, our structural parametric estimation results give significant support for the presence of cost synergies in the bidding. This essay demonstrates that a market outcome that is consistent with the division of the market among cartel members is also consistent non-cooperative behavior, indicating that a test for the presence of collusion in this market should consider these synergies or will risk “over-detecting” collusion.

The second chapter is titled "Numerical Analysis of Asymmetric Auctions with Optimal Reserve Prices," co-authored with Robert C. Marshall. We

numerically calculate equilibrium bids and optimal seller reserve prices for first price auctions within the independent private values framework where two bidders draw their valuations from heterogeneous distributions. In particular, we consider a type of bidder asymmetry that may result from collusion. In contrast to the case of a zero reserve price, once optimal reserve prices are introduced first price auctions no longer dominate second price auctions in terms of expected revenue for the seller. In fact, for the cases examined herein, the opposite is true. There are two main conclusions that we draw from this. First, most of the literature has focused on the case of a zero reserve price. Our paper shows that results that hold with a zero reserve price may not also hold when optimal reserve prices are introduced. Second, there may be situations where auctioneers facing collusive bidders will maximize revenue when using a second price auction. This is contrary to the intuition that first price auctions are less susceptible to cartels than first price auctions, and are therefore preferable.

The third chapter is titled "Optimal Reserve Prices at First and Second Price Auctions with Endogenous Coalition Formation." This extends the second essay to consider two coalitions of bidders that form endogenously. In general, the incentive to form a cartel given an auction type will change as the reserve price changes. This essay first shows how the incentives to form a cartel change with the reserve price. Then, it is shown that when coalitions of bidders form endogenously as predicted by the recursive-core, a more sophisticated choice of the reserve price by the auctioneer (relative to the choice in the second essay) can affect the coalitions that result in such a way to again make first price auctions revenue dominate second price auctions even when there is a relatively small number of bidders.

The fourth chapter is titled "Improving the Corporate Leniency Policy." The United States Department of Justice (DOJ) has instituted the Corporate Leniency Policy, which allows a single member of a cartel to receive amnesty from criminal punishment in exchange for substantial evidence against the remainder of the cartel. This has been heralded as the most effective weapon against collusion, and has been involved with most recent antitrust litigation. Using the Davidson and Deneckere (1990) model and outcomes as a baseline, an antitrust policy without leniency results in more collusion. The addition of leniency still results in more collusion than in Davidson and Deneckere, but not as much as with just the antitrust policy. However, including a bounty with leniency can improve the situation (from the view of the DOJ) if the criminal fines are sufficiently large.

TABLE OF CONTENTS

LIST OF FIGURES vii

LIST OF TABLES ix

CHAPTER 1: THE IMPACT OF COST SYNERGIES ON BIDDING IN THE GEORGIA SCHOOL MILK MARKET 1

CHAPTER 1: REFERENCES 53

CHAPTER 2: NUMERICAL ANALYSIS OF ASYMMETRIC AUCTIONS WITH OPTIMAL RESERVE PRICES 84

CHAPTER 2: REFERENCES 101

CHAPTER 3: OPTIMAL RESERVE PRICES AT FIRST AND SECOND PRICE AUCTIONS WITH ENDOGENOUS COALITION FORMATION ... 105

CHAPTER 3: REFERENCES 134

CHAPTER 4: IMPROVING THE CORPORATE LENIENCY POLICY 143

CHAPTER 4: REFERENCES 166

LIST OF FIGURES

Chapter 1

FIGURE 1: GEORGIA SCHOOL DISTRICTS, COLOR-CODED BY 1989 WINNER 70

FIGURE 2: HISTOGRAMS OF WINNING BIDS AND NON-WINNING BIDS FOR ALL YEARS........71

FIGURE 3: AVERAGE WINNING BID BY YEAR 72

FIGURE 4: DIFFERENCE OF THE AVERAGE WINNING BID: THREE BIDDERS VS. TWO BIDDERS AND FOUR BIDDERS VS. TWO BIDDERS 72

FIGURE 5: AVERAGE WINNING BID BY ENROLLMENT CATEGORY........73

FIGURE 6: PERCENTAGE OF PROCUREMENTS IN EACH MONTH BY ENROLLMENT CATEGORY: 1978-1991 73

FIGURE 7: DIFFERENCE OF THE AVERAGE WINNING BID: JULY VS. JUNE AND AUGUST VS. JUNE 74

FIGURE 8: DIFFERENCE IN THE PERCENTAGE OF THE TOTAL ENROLLMENT WON FROM THE PREVIOUS YEAR FOR EACH OF THE FOUR LARGEST FIRMS........76

FIGURE 9: WINNING BID VS. SECOND LOWEST VOLUME RATIO, 1989........76

FIGURE 10: COST SYNERGIES THROUGH ADJACENCY........77

FIGURE 11: ESTIMATED COST AND BID DENSITIES FOR CASE 1........80

FIGURE 12: ESTIMATED BID FUNCTION FOR CASE 1........80

FIGURE 13: ESTIMATED COST AND BID DENSITIES FOR CASE 2........81

FIGURE 14: ESTIMATED BID FUNCTION FOR CASE 2........81

FIGURE 15: HISTOGRAMS OF ALL BIDS FOR TWO, THREE, OR FOUR (POTENTIAL) BIDDERS........83

Chapter 2

FIGURE 1:COALITION OF 1 VERSUS COALITION OF 10 103

FIGURE 2:COALITION OF 4 VERSUS COALITION OF 7 104

Chapter 3

FIGURE 1: OPTIMAL RESERVE AND THE CORRESPONDING EXPECTED REVENUE AND EXPECTED SURPLUS AT A SECOND PRICE AUCTION WITH AN ALL-INCLUSIVE CARTEL OF SIZE N 137

FIGURE 2: PER CAPITA EXPECTED SURPLUS AT A FIRST PRICE AUCTION WITH FIVE BIDDERS 138

FIGURE 3: DIFFERENCE BETWEEN PER CAPITA EXPECTED SURPLUS OF AN ALL-INCLUSIVE CARTEL OF FIVE AND OF A COALITION OF ONE AGAINST A COALITION OF FOUR...................... 139

FIGURE 4: EXPECTED REVENUE AT A FIRST PRICE AUCTION WITH VARIOUS COALITIONS OF FIVE BIDDERS 140

FIGURE 5: EXPECTED REVENUE AT A FIRST PRICE AUCTION WITH VARIOUS COALITIONS OF FIVE BIDDERS 141

FIGURE 6: EXPECTED REVENUE AT A FIRST PRICE AUCTION WITH FIVE NON-COOPERATIVE BIDDERS 142

LIST OF TABLES

Chapter 1

TABLE 1: SUMMARY STATISTICS: 1978-1991 71

TABLE 2: AVERAGE WINNING BIDS IN NON-COLLUSIVE YEARS: 1978-1992, 1989-1991 (NO. OF OBS.)....................74

TABLE 3: LOCATION CATEGORIES 75

TABLE 4: AVERAGE WINNING BIDS IN NON-COLLUSIVE YEARS BY LOCATION CATEGORY AND NUMBER OF BIDDERS (NO. OF OBS.) 75

TABLE 5: INCUMBENCY CATEGORIES....................75

TABLE 6: AVERAGE WINNING BIDS IN NON-COLLUSIVE YEARS BY INCUMBENCY CATEGORY AND NUMBER OF BIDDERS (NO. OF OBS.) 75

TABLE 7: AVERAGE WINNING BIDS IN 1989 FOR ALPHA AND NON-ALPHA DISTRICTS, BY NUMBER OF BIDDERS (NO. OF OBS.) 77

TABLE 8: REGRESSIONS OF LN OF BIDS FOR 1989 PROCUREMENTS....................78

TABLE 9: STRUCTURAL ESTIMATES....................79

TABLE 10: MEAN AND STANDARD DEVIATION OF ESTIMATED COST AND BID DISTRIBUTIONS (3 BIDDERS)....................80

TABLE 11: WEIGHTING FUNCTIONS....................82

TABLE 12: ESCALATION MULTIPLIERS (TOP) AND DEFLATING INDEX (BOTTOM, NORMALIZED TO MAY 1978) 82

Chapter 2

TABLE 1: AUCTIONEER'S OPTIMAL RESERVE, EXPECTED REVENUE, STANDARD ERROR, AND BIDDERS' EXPECTED SURPLUS (PER CAPITA) - N = 5....................95

TABLE 2: AUCTIONEER'S OPTIMAL RESERVE, EXPECTED REVENUE, STANDARD ERROR, AND BIDDERS' EXPECTED SURPLUS (PER CAPITA) - N = 11 96

TABLE 3: AUCTIONEER'S EXPECTED REVENUE AT FIRST AND SECOND PRICE AUCTIONS, MEAN REVENUE DIFFERENTIAL, AND STANDARD ERROR 97

Chapter 3

TABLE 1: FIRST PRICE AUCTION WITH THREE BIDDERS 126

Chapter 1

The Impact of Cost Synergies on Bidding in the Georgia School Milk Market

Robert C. Marshall, Matthew E. Raiff, Jean-Francois Richard,
and Steven P. Schulenberg

1. Introduction

School districts throughout the country conduct procurements during the summer months for the provision of milk throughout the subsequent academic year. These procurements are invariably first price sealed bid. The bidders are milk processors who have either processing plants or distribution centers within a reasonable proximity of the school districts conducting the lets. In the late 1980's and early 1990's there were numerous antitrust cases brought by states against bidders for rigging bids at these procurements.

In this paper our objective is more modest than to model collusion at these procurements. Instead we try to obtain a reasonable characterization of bidder behavior during a time period when collusion was supposedly not a pervasive phenomenon. Such a characterization is an important first step for determining how one should then proceed with the modeling of collusion. Specifically, if non-cooperative behavior is badly mis-characterized then collusion may be “detected” when, in fact, the collusion just better captures the true non-cooperative behavior of bidders.

Milk processors deliver fluid milk each day to hundreds of grocers, restaurants, convenience stores, hospitals, schools, and many other consumers. The delivery of milk to many locations is done by medium sized refrigerated trucks known as peddle trucks. Tens of trucks leave a processor each morning and return empty later in the day. As a processor loses and wins contracts they solve new vehicle routing problems. Ideally, trucks leave each morning at full capacity and have routes which allow them to deliver all milk within a normal working day.

Processors obtain their raw milk from dairy cooperatives. The agreements between processors and dairy cooperatives are long standing relationships. If there is any sudden significant change in quantities purchased from the cooperative, up or down, then the processor is required to pay large premiums or penalties.

These two features of the milk processing industry mean that milk processors want to sell a specific volume of fluid milk and, in addition, they want to sell the volume so as to best solve their vehicle routing problems. The latter phenomenon implies the potential existence of cost synergies in the market. Specifically, milk processors may prefer to have contracts for delivery that are geographically adjacent than to have contracts that are geographically dispersed. Figure 1 shows the contracts won by processors in Georgia in 1989 – it is evident from this colored map that processors acquire school districts in geographic groupings.

Milk processors know quite a bit about each other. Production technologies and input costs are common knowledge. But, they do not know how another processor will solve it's vehicle routing problem and, in particular, they do not know how a specific school milk contract assists with a competitor's solution. This part of a processor's costs is private information.

In this paper we attack the following question – are synergies important in understanding non-cooperative behavior in this market and, if so, how important? Although our ultimate interest is in understanding the collusion that occurred, this question is an essential precursor to address for several reasons. First, the non-cooperative behavior serves as a benchmark against which the collusive behavior can be measured. Suppose vehicle routing cost synergies are indeed present for several school districts for more than one milk processor. Suppose there are other districts where the vehicle routing cost synergies are not relevant or relevant for only one producer. The former districts will have low prices while the latter will have high. Without properly understanding this non-cooperative source for price dispersion we may falsely find collusion[1] in districts where cost synergies were not strong or entirely absent.

Second, understanding non-cooperative behavior will allow a more accurate characterization of collusive behavior. If synergy is not considered, then examination of the winners of each school district might result in the conclusion, perhaps inaccurate, that collusion was implemented via a geographic allocation scheme.

There are many factors in this market that may lead to differences in prices between procurements, such as size of the school district, time of year of the procurement, distance between the processor and the district, number of bidders, incumbency, location, and synergy. We first investigate each of these potential factors using reduced form measures that are suggestive of their existence and potential importance. Using a time period during which collusion is unlikely in the market[2], we find some of these factors (number of bidders as well as synergy) to be significant while several are not.

As a theoretical foundation for our structural estimation we use the work of Krishna and Rosenthal (1996, subsequently KR) who provide a model of non-cooperative bidder behavior in the presence of synergistic values. We adapt this

---

[1] Of course, we may also inaccurately measure the impact of collusion for the same reasons.

[2] Per plea periods and by our analysis.

model to our setting – first price procurements. Factors found to be significant in the reduced form estimation can be included directly into the structural model. In particular, a wonderful feature of the KR model is that there is a single synergy parameter, and if that parameter is zero then the model produces a bid function that is identical to the bid function for a single object first price sealed bid procurement. In other words, the standard first price sealed bid model is fully nested within the KR synergy model. This greatly simplifies inference since inter-model comparisons rely on the significance of a single parameter.

There are separate contributions in the literature regarding the estimation of structural procurement models[3], reduced form estimation of cost synergies[4], and reduced form empirical analyses of collusion in school milk procurements[5]. The

---

[3] In first price procurements costs are not directly observed. Instead, bids are observed which are a function of the unobserved costs. The major difficulty of this stems from the fact that this function will depend on the unknown cost distribution, and hence is also unknown to the researcher. The literature has focused on this latency and the numerous econometric issues that arise as a result. A subset of important work in this vein includes Donald and Paarsch (1993,1996), Laffont, Ossard and Voung (1997), Sareen (1999), Guerre, Perrigne, and Vuong (2000), and Qi (1996).

Baldwin, Marshall and Richard (1997) explore whether price variations in the Forest Service timber sales in the Pacific Northwest are better explained by collusion or variations in timber supply conditions, but their data consists exclusively of English auctions.

[4] Ausubel, Cramton, McAfee, and McMillan (1997) examine prices from the FCC auctions of spectrum for evidence of geographically based synergistic effects. They find that prices are relatively high for a given piece of spectrum if the last bidder to drop out of the bidding for that piece of spectrum already held licenses in adjacent areas.

[5] Hewitt, McClave and Sibley (1993) examine the Texas school milk market and find evidence of a complementary bidding scheme that the authors believe to be only consistent with collusive behavior.

Pesendorfer (1994a, 1994b, 2000) examines the Florida and Texas school milk market. He marginal contribution of this paper is the use of structural estimation to measure cost synergies in school milk procurements during a period of time when behavior was allegedly non-cooperative.

The paper is organized as follows. In Section 2 we describe Georgia school milk procurements. The relevant theory as adapted from KR is presented in Section 3. The data is discussed in Section 4. Section 5 discusses the effect of the market and procurement characteristics on the bidding by the milk suppliers. The reduced form model and the estimation results are discussed in Section 6. An enumeration of our most important maintained hypotheses appears in Section 7. In Section 8 we present our empirical model and estimation results. Concluding remarks are offered in Section 9.

---

finds that the variance in the share of contracts per cartel member is smaller for a cartel that is unable to make side payments. Consistent with results from the asymmetric auction literature he fids that the cartel bids less aggressively than the competitive fringe and that the cartel's bids first order stochastically dominate the distribution of non-cartel bids.

Raiff (1997) examines the Georgia milk data used in this paper. He employs an exogenous switching regression model to determine if a particular school district was receiving collusive bids. Raiff finds that the cartel's winning bids was less sensitive to cost measures than non-cooperative winning bids.

Porter and Zona (1999) examine the Ohio school milk market. An argument is put forth that competitive bids in the milk processing industry should be a monotonically increasing function of the distance from the school district to the processing plant. The reasoning is based on the assertion that firms would not be able to charge high prices for districts nearby, because they are less costly to service, unless they entered into a collusive agreement with their competitors. Porter and Zona find that the defendant's bids are lower on districts further away and they interpret this as evidence that these firms established territorial monopolies as part of a collusive agreement among the firms.

2. Georgia School Milk Market

A detailed discussion of the milk processing industry is contained in Appendix A.

2.1. Procurement

The Board of Education in every school district solicits bids, and awards one-year school milk contracts to the lowest responsible bidder. Bids are typically solicited for four types of milk: whole white, low-fat chocolate, low-fat white and skim. Between two and four weeks before the let, the Board of Education of the school district distributes to potential vendors a detailed description of the contract for their district containing the list of schools to be supplied, the contract period, the delivery times and estimated quantities for the individual milk items. The schools may increase or decrease quantities of any or all items listed in the specification. All costs of delivery are included in the per item price. Each year these procurements begin in late April and continue until the end of August. There are 181 Public School Districts in the State of Georgia – 159 county districts and 22 independent districts.

The contract usually allows for both firm and escalating bids. A firm bid is a fixed price per half pint for the life of the contract. An escalating bid contains an escalation clause which allows the firm/school district the right to request price increases/decreases during the year. The typical escalation clause in Georgia calls for a one-tenth cent increase/decrease per fifteen cent increase/decrease in the FMO price of raw milk[6] per hundredweight where 186 half pints of milk equal approximately one hundredweight. The scoring function we develop, which is discussed in Section 4.1 and Appendix C, accounts for the differences in these two types of bids.

3. Relevant Theory: Synergistic Bidding

The following model is adapted from Krishna and Rosenthal (1996) and will be used in the structural estimation. The model characterizes bidding behavior at a first-price procurement where the bidders experience cost synergies.

There are N districts holding procurements and these are assumed to be arranged on a circle. Each bidder is interested in M districts. For instance, with $M = 3$ bidder i is interested in procurements at adjacent locations $i - 1$, i and $i + 1$. Note that this implies that each procurement has 3 (of more generally M)

---

[6] The major factor input for a milk processing plant, raw milk, is regulated by Federal Milk Marketing Orders. The FMO equals the Minnesota / Wisconsin price plus a fixed differential typically increasing in distance from the upper Midwest. Various discussions with industry experts have confirmed that, on average, the price of raw milk accounts for 65% of the cost of producing a half-pint of milk.

bidders.

Again, consider the case of $M = 3$ (in which bidder i effectively competes with bidders $i - 2$, $i - 1$, $i + 1$, and $i + 2$). Bidder i receives a private signal for c which is his cost if he wins only one contract. Signals are i.i.d. and distributed according to the distribution function F with density f. If the bidder wins two objects, the total cost is $2c - \alpha$ and if he wins three then the total cost is $3c - 2\alpha$.

The objects are sold using separate and simultaneous first-price procurements. We are searching for a symmetric and increasing equilibrium strategy $B(\cdot)$ with inverse $B^{-1}(\cdot)$. Note that this implies a bidder submits the same bid at each procurement. Bidder i's expected payoff from bidding b with a cost of c is:

$$\begin{aligned}\Pi(b,c) &= (1-F)^4(3b+2\alpha-3c)+2(1-F)^3F(2b+\alpha-2c)\\ &\quad +2(1-F)^2F(b-c)+(1-F)^2F^2(b-c)\end{aligned}$$

where $F \equiv F(B^{-1}(b))$. The first term captures the expected payoff from winning all three procurements (and thus beating all four other bidders with interest in these contracts). The second term captures the expected payoff from winning exactly 2 contracts (which, by symmetry, must be either $\{i-1, i\}$ or $\{i, i+1\}$ and involves beating all but one competitor who is either bidder $i-2$ or $i+2$.)

The third term captures the expected payoff from winning only the $i - 1$ or $i + 1$ contract. (To win only the $i - 1$ contract bidder i must beat both bidders $i - 1$ and $i - 2$ but lose to bidder $i + 1$. The bid of $i + 2$ is not relevant to the case. The same logic applies to the $i + 1$ contract.) The fourth term captures the expected payoff from winning only the i-th contract.

The above expression simplifies to:

$$\Pi(b,c) = \left(1 - F\left(B^{-1}(b)\right)\right)^2 \left(3b + 2\alpha - 3c - 2\alpha F\left(B^{-1}(b)\right)\right).$$

The first order condition for an optimum is:

$$\begin{aligned} 0 = \; & -2\left(1 - F\left(B^{-1}(b)\right)\right) f\left(B^{-1}(b)\right) \left(B^{-1}(b)\right)' \left(3b + 2\alpha - 3c - 2\alpha F\left(B^{-1}(b)\right)\right) \\ & + \left(1 - F\left(B^{-1}(b)\right)\right)^2 \left(3 - 2\alpha f\left(B^{-1}(b)\right) \left(B^{-1}(b)\right)'\right) \end{aligned}$$

Imposing symmetry and simplifying reduces this to:

$$-2f(1 - F)b + (1 - F)^2 b' = 2\alpha f(1 - F)^2 - 2cf(1 - F).$$

Let $G(c) \equiv (1 - F(c))^2$ and $g(c) \equiv G'(c)$ so that the differential equation becomes

$$\frac{d}{dc}[G(c)B(c)] = g(c)c - \frac{2}{3}\alpha\frac{d(1 - F(c))^3}{dc}$$

and therefore the symmetric equilibrium bid function is

$$B(c) = c + \frac{1}{G(c)}\int_{y=c}^{1} G(y)\,y\,dy - \frac{2}{3}\alpha(1 - F(c)).$$

One can demonstrate that if each bidder is interested in a general number M of districts (with M bidders at each procurement), and with the cost distribution F as a function of parameter vector θ, the equilibrium bid function is[7]

$$B(c; M, \theta, \alpha) = c + \frac{1}{G(c; M, \theta)}\int_{c}^{1} G(u; M, \theta)\,du - \frac{M-1}{M}\alpha(1 - F(c; \theta)) \quad (3.1)$$

where $G(c; M, \theta) \equiv (1 - F(c; \theta))^{M-1}$. This allows a simple adjustment for different numbers of bidders between procurements in the estimation procedure.

A remarkable feature of this equilibrium is that if $\alpha = 0$ then the symmetric equilibrium for the synergistic bidding model is simply the standard single object

---

[7] In order to preserve symmetry, we assume that if M is an even integer then each bidder's interests must be skewed in the same way. For example, with $M = 4$ bidder i is interested in $i - 1$, i, $i + 1$, $i + 2$.

bid function. This feature allows for a simple empirical test between the two models since the synergy model nests the single object model.

4. Data Selection and Description

The original data set for the State of Georgia contains 6,947 bids on a total of 2,368 contracts. The data set contains observations as early as 1973 and as late as 1992. When available the following information is given for each contract: the name of the school district, the date of the let, the identity of the bidding firms and the identity of the individual who signed the winning bid, the bids for the individual milk items, the total number of half pints sold, whether a bid was firm or contained an escalation clause and a notes field for any peculiarities regarding the procurement. Firms may submit two bids for the same contract: one with and one without an escalation clause.

As usual when working with primary data, many observations were highly problematic. Our data set reduction steps are precisely enumerated in Appendix B. In the end we were left with a total of 1,531 procurements, generally with 2 to 5 bidders at each, and a total of 4,829 bids. The most sparse year of data was 1978 with 45 procurements, and the most complete was 1989 with 163 procurements.

4.1. Scoring Function

Our data provides bids, per half pint, for each of the four milk products. The winner is identified in the data and we often have bids of the losers. However, we do not have the precise scoring function used by each school district in each year that maps the multi-dimensional bids into a single number and thus determines the winner. Nevertheless, we needed to score these bids in order to undertake our empirical analyses. The scoring function also adjusts bids for inflation, using a price index based on the price of raw milk and the motor fuel price index. Appendix C contains a detailed discussion of our scoring.

4.2. Summary Statistics

Summary statistics are provided for our data set in Table 1. The variables in the tables are all bids, the winning bid, non-winning bids, and the price of raw milk, all in cents per half pint[8]. The winning bids are on a tighter range (as well as being 0.35 cents lower on average) than all bids, as standard procurement theory would suggest. Histograms of winning bids and also non-winning bids are shown

---

[8] Variation in the price of raw milk may seem surprising given that there are Federal price supports for milk. In fact, the price of raw milk changes quite often and by substantial amounts. First, there is a strong seasonal component in the price which is related to the reproductive cycle of cows. Second, climate conditions have a very big impact on the milk productivity of cows. A drought in Georgia may cause a sudden increase in the price in that state relative to other locations.

in Figure 2.

5. Procurement Characteristics and Bids

In the construction of our structural model, we will be using several maintained hypotheses. In order to gain insight into what the appropriate assumptions are, we will first use a reduced form model. Such a reduced form analysis serves another essential purpose. There are potentially important market and procurement characteristics which could be problematic for our subsequent structural form analysis such as collusion, the size of the school districts, the within-summer timing of the procurement, location of the processing plants and distribution centers, incumbency, volume, and synergy. In the absence of operational structural models which can account for such complications, their potential relevance is investigated within our reduced form analysis.

5.1. Collusion

The majority of the milk suppliers that appear in our data were legally determined to have participated in collusive bidding during the years 1983 to 1988. Much of the evidence was based on the confessions of some of the suppliers involved. The period prior to 1983 fell outside of the statute of limitations for bid rigging cases, and hence it is not clear if this period was truly non-cooperative or was just not investigated by law enforcement authorities. The investigation began between the 1988 and 1989 procurements. We will treat bids in 1989 as if they reflect non-cooperative behavior.

In Figure 3 we provide a time series plot which shows the average winning (scored) bid in each year. Winning bids increase and remain relatively high for the periods 1983 to 1988 compared to other years. This is consistent with our expectations given the guilty pleas for this period.

5.1.1. Non-Cooperative Bidding and the Number of Bidders

Although we have complete information on winning bids we have incomplete information on losing bids, so we cannot simply use the number of bids shown as the number of potential bidders[9] at a procurement. We also used the bidders in the surrounding time periods to determine this number on the assumption that if a firm submitted a bid in the previous and following year then they were also potential bidders in the current year.[10]

---

[9] The relevant number of bidders is the number which the bidding firms expect to participate in a procurement, rather than the actual number that did submit bids. Those bidders that are expected to submit bids are what are referred to as potential bidders.

[10] For example, our data occasionally show the same three firms submitting bids over the length of our data except for one year, and in that year only one firm submitted a bid. We do not believe that the other two firms did not bid, but rather that they simply were not recorded

Standard competitive procurement theory predicts that bids, and in particular the winning bid, should be decreasing in the number of bidders. Figure 4 shows the difference between the average winning bid for three bidder versus two bidder procurements as well as four bidder versus two bidder procurements. Note that this difference is generally negative, as theory would predict, for all years except the collusive period. Collusive bidding during the plea period may have been largely unrelated to costs. Histograms of all bids classified by the number of bidders are provided as Figure 15.[11] This figure shows, again as standard theory would predict and consistent with the above, that the distribution of bids with three or four bidders is shifted towards higher bids relative to the distribution with two bidders. This figure also provides some evidence that the spread of the bid distributions is relatively consistent across differing numbers of bidders. One implication of these observations is that it is a reasonable assumption that the underlying cost distribution is not conditional on the number of bidders.[12]

---

in our data. We therefore treat this as a three bidder procurement. Similar adjustments were made where appropriate for each procurement.

[11] Only auctions with two, three, or four (potential) bidders have been included, due to a relatively small number of observations for other numbers of bidders.

[12] This assumption is made in the estimation below, as described in Section 8.

5.2. Size of District

The size of the school district can be measured by the total enrollment in all of the schools in the district. This has several potential effects on the underlying costs of supplying milk to a district and hence on the bids submitted.

1. A high enrollment corresponds to a more populated district[13]. Such districts would typically have more and better developed roads, which might increase flexibility in solving a vehicle routing problem.

2. Delivering a small amount of milk to a rural school might require relatively high labor costs per half pint, while delivering to relatively large districts might require the use of so much of the capacity of a peddle truck that it cannot be part of any route involving another school.

3. There might be efficiencies involved with unloading large quantities of milk at a given school.

4. Finally, winning a procurement associated with a high enrollment allows a processor to place a large proportion of its volume. This could make large

---

[13] The correlation between the total enrollment of a school district and its total population is 0.94.

schools more attractive by reducing the chance that a firm will be left with excess volume at the end of the summer.

We separated the districts in our data into four separate enrollment categories. 63% of the districts fall in the 0-4000 enrollment category, 24% in the 4001-8000 category, 10% in the 8001-20000 category, and 3% in the 20000+ category[14]. For 1989, our most complete year of data, there are 108, 33, 15, and 7 procurements in the categories respectively, out of a total of 163. Figure 5 shows the average winning bid for districts in each of these categories by year. This shows that while bids for all districts with enrollments under 20000 are similar, the bids at the largest districts are significantly lower (10.06 cents on average for the largest districts compared to 10.29 cents on average for all other districts). But not only are there very few observations for the 20,000+ districts but these lets have more bidders than lets for smaller districts. This suggests that the effect of enrollment on the bidding may be significant, but must be investigated with a more complete model (such as the reduced form model presented below) that would control for other relevant factors.

---

[14] Enrollment for each district did change from year to year, but typically by a small amount. Each district is in the same enrollment category for all years.

5.3. Timing

The procurements take place throughout the entire summer. Of the 1531 scorable procurements in our data set, 274 did not have a date recorded. Of those that did, approximately one half of the procurements take place in August, one quarter in July, and the remainder throughout May and June. Districts typically hold their procurements at roughly the same time each summer. This implies that some districts choose the beginning of the summer, and some choose the end.

Figure 6 shows the proportion of the procurements that occurred in each month, both overall and for each of the enrollment categories as described above. This shows that larger schools hold their procurements earlier, moving from an average of July 12 for the largest schools to August 1 for the smallest. Figure 7 shows the difference in average winning bids for August versus June as well as July versus June. For the majority of the non-collusive years, winning bids decrease from June to July and from July to August, while the reverse is true in the collusive years. Since the large districts tend to have earlier procurements, this is in conflict with the observation from the previous section that winning bids tend to decrease with enrollment.

In order to investigate this further, Table 2 shows the average winning bids in each month in each of the above enrollment categories, for the non-cooperative years – 1978-1982 and 1989 to 1991. The number of procurements that are in each category are shown in parentheses. From this Table it is unclear if larger or earlier schools receive lower winning bids. Winning bids decrease over the summer for the midsize schools, but not the smallest schools. Winning bids decrease with increasing enrollment in July, but not June. This suggests that the date of the procurement has little direct effect on the bids, although this will also be investigated more fully in the next section.

5.4. Location of the Processing Plants and Distribution Centers

It seems intuitively obvious that the distance between a milk processing plant and a school district would have a bearing on costs and thus bids. But any processor can set up a remote distribution center at low cost. In fact, a large parked semitruck can be a distribution center provided that there is an available power source to refrigerate the milk.[15] Furthermore, a processor may be operating numerous delivery routes throughout diverse geographic areas. A processor may have numerous retail contracts in a grouping of counties over 100 miles away from the processing plant. The marginal cost of servicing a school district in one of these counties may be quite low as a consequence of the existing retail contracts.

---

[15] In what follows the term "distribution center" refers to both a processing facility and a non-processing distribution facility, such as a refrigerated semitruck.

There is another more subtle issue stemming from the nature of non-cooperative asymmetric procurements that must be considered before attempting to measure the impact of location on price. Consider a situation with two otherwise identical firms, one that is close to a given school district and one that is far away. The close firm will have low costs but a relatively high bid function, and the far firm will have high costs but a relatively low bid function. The result of this is that the bid distribution for both firms will look very similar (although not identical), even though they have distinct cost distributions.

It is possible to find comparisons that do allow for potential asymmetries, like distance, to be measured. Consider two otherwise identical firms that are at differing distances from a given district. The winning bid will typically be lower when both firms are close than when one is close and one is far. Similarly, the winning bid will typically be lower when one is close and one is far than when both are far. Separating districts into these categories (where the meanings of "close" and "far" are chosen appropriately) and comparing the winning bids across them allows the effect of distance to be seen.

We divided the Georgia school districts into five disjoint groups: districts containing one firm's distribution center, districts with two or more firms' distribution centers, districts adjacent to those with one firm's distribution center, districts adjacent to those with two or more firms' distribution center, and finally districts that are not in the above categories. This final group of districts are those not near any distribution centers. These location categories are identified in Table 3. The average winning bids for non-collusive years (78-82, 89-91) in each of these categories, controlling for the number of bidders, is shown in Table 4.

Since the distribution centers tend to be located near areas with higher populations, it is natural to expect that districts closer to distribution centers will have more bidders. The evidence is mixed. It suggests that the distance from a bidder's distribution center to a district may not have a significant direct impact on its bid, but more importantly it suggests that there may also not be a direct effect on costs.

5.5. Incumbency

Our data set contains 694 scorable procurements during the non-collusive years. Of these, the identity of the winner in the previous year was not known for 94 procurements. This occurs either for the first year a district appears in the data or for years following missing observations. For the 600 procurements for which the incumbent is known, 394 were won by this incumbent. This suggests the possibility that the incumbents generally submit lower bids (and hence win more often) than the non-incumbents due to a cost advantage related directly to incumbency.[16]

If it is significant, incumbency naturally causes a cost asymmetry between firms. We know from the equilibrium of an asymmetric procurement that although the incumbent has a cost advantage, the bids of all of the bidders would be similar. The incumbent would not bid out the entire cost advantage, since the other bidders do not have it. The non-incumbents would have to bid more aggressively since they are facing one bidder with a typically lower cost. As a result, it is misleading to compare bids by an incumbent to competing bids by non-incumbents and then claim that the difference accurately measures an incumbency effect. A stronger implication of this is that even if incumbency does provide a significant cost advantage for one firm, using bidder specific incumbency information would bias a conclusion about incumbency towards insignificance since the bids would not reflect this difference in costs.

We separated our data into four categories (shown in Table 5). These are procurements where the incumbent won, where the incumbent did not win but bid, where the incumbent did not bid, and where the incumbent was not known.

---

[16] One potential reason for this could be that the delivery options become more flexible as the cafeteria managers become more familiar with the peddle truck drivers. Initially, the drivers might have only a particular half hour window to drop off the milk. This places restrictions on the routes that may be used. After some time, the cafeteria managers often give keys to the drivers that allow delivery at any time, even before school employees get there. This increases flexibility, and allows for better solutions to the vehicle routing problem.

If firms are bidding as the equilibrium of an asymmetric procurement model suggests, than bids in the first two categories should be equal on average. If, in addition, incumbency does give one firm a significant cost advantage, then bids in the third category should be higher than that of the first two. Note that this allows incumbency to be investigated using district specific data rather than bidder specific data, thus avoiding the asymmetry problems discussed above.

Table 6 shows the average winning bids in each of these categories, both overall and separated by the number of bidders. For every number of bidders, procurements where the incumbent wins have lower average winning bids than those procurements where the incumbent does not participate. This suggests that incumbency is significant. However, it also is the case that for two and three bidder procurements, which contain most of the data, the bids are roughly a third of a cent lower when the incumbent wins than when a non-incumbent beats the incumbent. This is contrary to expectations and hence suggests that incumbency is perhaps not the direct cause of the difference in observed bids.[17] This is inves-

---

[17] Given the nature of the incumbency advantage in this market, it is not unreasonable to expect it to have an insignificant effect on costs and bidding. The procurement is for contracts that last for the entire school year. While it may be the case that the advantages discussed above (keys and experience with routing) are significant at first, a new supplier would likely be able to obtain these advantages well before the contract ran out. This would give the incumbent a very temporary advantage which, when spread throughout the life of a contract, may not realistically amount to much of a cost savings.

tigated further in the reduced form model.

5.6. Volume

Through contracts with dairy cooperatives, milk processors have a fixed supply of class A raw milk arrive at their facility each day. The amount of raw milk that a supplier will purchase during a school year is generally determined well before the procurements for that school district occur. There are large financial penalties associated with increasing or decreasing one's raw milk purchases from the cooperative. This gives potentially strong incentive for the suppliers to win enough procurements each summer so that the total enrollment to be served will account for all of the raw milk to be purchased (while maintaining a small reserve to allow some flexibility).

For each of the four largest milk processors, Figure 8 shows the difference from the previous year in the percentage of the total enrollment won. Firms tend to win generally the same proportion of the Georgia enrollment each year. There are a few things to note regarding this figure. The data is only reasonably complete from 1982 to 1990.[18] The variability shown in the tables decreases during the first few years of the collusion, remains quite small for the remainder of the collusive

---

[18] The lack of data greatly increases the variability seen in the percentages in years outside this range.

period, and then increases again as the collusion breaks down (especially 1990). This suggests that one main accomplishment of the collusive behavior may have been to stabilize the volume placed in Georgia schools by the cartel members.

Through a given procurement season, if a firm has already placed a relatively large amount of its volume, then the incentive to win any particular remaining procurement is lessened. Conversely, if a firm has already placed a relatively small amount of its volume, then the incentive to win each remaining procurement is increased. In this light, it is the firm with the second lowest percentage of its volume placed that has the largest effect on the winning bid. As the volume already placed of this firm increases the winning bid should tend to increase if volume considerations affect bidding. (The volume placed of the other firms would also affect the winning bids, but in a less direct way.) We constructed a district specific variable that reflects the relative volume placed of this second firm. We used the most complete competitive year of the data, which is 1989. For each firm, we determined the proportion of the total enrollment for the entire previous year (1988) that was won by a given firm. Call this Z_i where i denotes a given firm. Given the long term nature of the contracts with dairy farms, the proportion won by each firm should be the same in 1989. For each procurement in 1989, we aggregated the enrollments for all other procurements that had already been decided up to that point in time. In addition, we did this aggregation for each given firm. Then, for firm i we took the ratio of the latter to the former. Call this Y_i. The variable $\frac{Y_i}{Z_i}$ thus captures where firm i stands in 1989 relative to 1988 in terms of its placement of volume. A ratio of exactly one indicates that the firm in question is on pace to serve the same enrollment in 1989 as it did in 1988. As the ratio decreases towards zero the firm falls farther behind this pace and hence has more incentive to bid aggressively. The district specific variable we used was then the second lowest such ratio (second lowest of the $\frac{Y_i}{Z_i}$) at each procurement.

Figure 9 shows a scatter plot of the winning bid against the second lowest volume ratio for each 1989 scorable procurement. The solid horizontal line shows the mean winning bid. The figure indicates that there is little correlation between the winning bid and the second lowest volume ratio.

5.7. Synergy

As discussed above, the delivery costs are a substantial part of the overall costs of supplying milk to public schools, and the majority of the delivery costs are incurred specifically while the half-pints are being transported by peddle truck from distribution centers to the schools. Delivery to a single school typically requires only a portion of the capacity of a peddle truck. Ideally, delivery routes should be chosen so that a peddle truck leaves the distribution center full, drops off milk at multiple schools and retail establishments, and finishes a standard working day empty. Winning adjacent school districts increases the flexibility of a processor for solving vehicle routing problems. Again, we use the terminology “cost synergy” to refer to cost savings realized from winning adjacent school districts.

As an augmentation to Figure 1, we offer the observation that milk suppliers often win clumps of school districts that are located a distance from their distribution centers while not winning the districts that contain the distribution center. For example, the pink supplier has a distribution center on the border of the Stephens and Franklin school districts, while no other firm has a distribution center nearby. Despite this, the pink supplier wins neither of these districts in 1989. An even stronger example can be seen by comparing pink’s districts won in 1989 and 1990 (both non-collusive years). In 1990, pink does not win the Oconee-Clarke-Oglethorpe clump, and instead wins a clump consisting of Hancock, Taliaferro, and Warren. These districts are significantly far away from the distribution center, and there are many closer school districts on which this supplier did not bid aggressively enough to win.

Although Figure 1 is compelling, the sequential nature of the school milk procurements in any given summer provides another look at synergy. If in year

t a district is won by a different firm from that in $t-1$ then the subsequent procurements in surrounding districts should also tend to be won by different firms in period t. For each district for 1978 to 1991 we found the proportion of the total enrollment of the surrounding districts that was to be served by a different firm than the previous year. For districts where the incumbent won, this proportion was 35.4% while for districts where the incumbent did not win, the proportion was only 25.7% – this difference provides support for the synergy effect. Note that the total enrollment served in our data is approximately 11.5 million so a difference of 10% represents a large amount of milk.

From Figure 1 it is clear that processors buy districts in clusters. Within these clusters only one firm will realize a cost synergy. However, on the boundaries of the clusters, multiple firms will have the potential for cost synergy. It is these districts that where the cost synergy would have the most impact on the bids.

Exposition of our definition of a boundary district is facilitated by referring to Figure 10. Consider period t and look back at the ownership of the 16 districts in Figure 10 in the periods $t-2$ and $t-1$. We consider four milk processors denoted A, B, C, and D. Each cell of the table is a school district. The letter above the diagonal in each cell denotes the winner in period $t-2$ while the letter below denotes the winner in period $t-1$.

It is unlikely that synergy effects will be measurably observed for the four upper left districts and the four lower right districts. It appears that there is no second firm seriously contesting for ownership of these blocks.

The four lower left districts are each stand-alone districts. We would expect there to be no cost synergies realized from ownership of these.

All measurable cost synergy effects will be in the upper right four districts. Firm A is extending its adjacent holdings to district $(1,3)$ as well as $(2,3)$. Firm B is extending to $(2,3)$ and $(2,4)$. Furthermore, although $(1,4)$ does not switch it seems clear that there is enough potential synergistic bidding pressure on $(1,4)$ from A, B, and D to include it as a "boundary district" for the purposes of measuring synergy.

In Figure 10 districts $(1,3)$, $(2,3)$ and $(2,4)$ are of a particular type – districts that had a change in the winning firm in the previous two years when at least one of the winners had won a neighboring district in both years. 42 of the 133 scorable procurements are of this type in 1989.

District $(1,4)$ is a different type – a district won by the same firm in the two previous years, not adjacent to a like district, where the potential for synergistic bidding is strong since at least one competitor has been in a position for at least one year to enjoy the potential benefits of adjacent ownership by acquiring the district. 5 of the 133 scorable procurements were of this type in 1989.

Both types of districts are hereafter referred to as "alpha districts".

The average winning bids for alpha and non-alpha procurements in 1989 are shown in Table 7 for each number of bidders and overall, with the number of procurements in each category shown in parentheses. With two and four bidders the winning bids are roughly the same, but are lower in the alpha districts with three bidders and, to a lesser extent, overall. This is not inconsistent with synergy.

6. Reduced Form Estimation

Our data analysis, from this point forward, focuses exclusively on the 133 scorable procurements for 1989. The variables used in our regression analysis are below.

The dependent variable is the natural log of the winning bid and lowest losing bid for each procurement, deflated as described in Section 4.1 and Appendix C.3.[19]

Independent variables:

1. The natural log of the price of raw milk.[20]

2. Dummy variables for three bidders and also for four or more bidders. We

---

[19] All reduced form estimates were also run separately for the winners and lowest losers. These two estimates were quite similar. We only report the pooled estimates.

[20] Recall that the bids are deflated by the price of raw milk in the month of the procurement.

expect this coefficient to be negative, and more negative with more bidders.

3. Enrollment and enrollment squared. This will control for the size of each district. We expect these to be jointly insignificant, although if they are significant we expect the overall effect of enrollment to be negative. The square is included since we do not expect enrollment to have a linear effect. Enrollment is measured in ten thousands of students.

4. A dummy variable coded 1 if the procurement was known to have happened early in the year, where early is considered as no later than July 15, and 0 otherwise. This had a value of 1 for 27 of the 133 procurements which corresponds to 44% of the total enrollment. If the date was not known, which was the case for 13 procurements and 6% of the enrollment, the dummy for the procurement was chosen to be 0.

5. A dummy variable to control for the location of the processing plants and distribution centers. This is coded 1 if the district contained at least one processing plant or distribution center, and 0 otherwise. This had a value of 1 for 43 of the 133 procurements. If location is significant, the coefficient on this dummy variable should be negative.

6. A dummy variable coded 1 if the incumbent submitted a bid at the pro-

curement and 0 otherwise. If incumbency does provide a significant cost advantage, then this dummy should have a positive coefficient. However, we expect this coefficient to be insignificant.

7. The second lowest volume ratio, as described in the volume discussion above. Since higher values indicate more volume has been previously allocated and hence that placing volume is less important for the relevant bidders, this should have a positive coefficient if it is significant.

8. A dummy variable coded 1 if the district was in an alpha district, as described in Section 5.7, and 0 otherwise. The had a value of 1 for 47 of the 133 procurements. If synergy is significant, this should have a negative coefficient.

The result of this regression is shown in Table 8.[21] Standard errors (using White's heteroskedasticity correction) are shown in parentheses. The most notable result is that only the bidder and alpha dummy variables are significant. As expected, the dummy variables for the number of bidders are both negative. These imply that, all else equal, the bids decrease as the number of bidders increases.

---

[21] The regression (and the subsequent tests) were performed using EasyReg International, an econometric program provided by H. J. Bierens. The EasyReg homepage is available at http://econ.la.psu.edu/~hbierens/EASYREG.HTM.

The negative and significant coefficient on the alpha dummy implies that cost synergy is indeed important, and that the milk suppliers account for this in their bidding.

To test if the remaining variables (other than number of bidders and synergy) are jointly insignificant, the F test, the Wald test, and the Kiefer-Vogelsang-Bunzel (KVB) test may each be used. The F-test null distribution is F(7,255) which gives a 10% critical value of 1.74. The F-test statistic is 1.61, and therefore the null of joint insignificance is accepted. The Wald test 10% critical value is 12.02. The corresponding test statistic is 11.29, so again the null of joint insignificance can not be rejected. Finally, the KVB test statistic is 50.37 and the 10% critical value is 67.37, so the null of joint significance can not be rejected with this test either.[22]

However, as discussed above, the synergy corresponding to a school district may be affected by the size of the district. To investigate this, we ran a second regression. This was identical to the first with two exceptions. The synergy dummy variable was interacted with enrollment and enrollment squared, and these were included as additional regressors. The results of this regression are shown in Table 8. The two additional variables are jointly insignificant. This suggests

---

[22] Note that although the KVB test is designed to allow hypotheis testing in models with heteroskedasticity, performing a Breusch-Pagan test leads to the conclusion that the errors are in fact homoskedastic. Also note that while the KVB test may not be invariant to the order of the observations, different orders produced similar results.

that synergy is not dependent on the size of the school district. An alternative explanation is that synergy is indeed dependent on the size of a district, but also dependent on the size of the surrounding districts. This would lead to all districts in a given area having similar synergy values, hence also motivating a constant synergy.

The final regression eliminates all insignificant regressors, leaving only a constant and the bidder and alpha dummy variables. Not surprisingly, all remaining variables are significant. The coefficients and standard errors are not significantly different from either of the other regressions. This regression supports the maintained hypotheses that are introduced next.

7. Maintained Hypotheses

Several of the maintained hypotheses used for the structural estimation are enumerated below.

7.1. Symmetric Independent Private Values

All underlying assumptions of the symmetric independent private values model are assumed to be valid. Firms are risk neutral, some part of cost information is privately known to each firm, they each draw the privately known part of their cost from a common underlying distribution for each given contract, and these privately known cost draws are independent between firms. All firms act non-cooperatively, which appears to be reasonable for 1989.

From discussions with industry participants it is clear that each milk processor knows everything of relevance regarding every other milk processor with one exception – the cost implications in terms of solving the vehicle routing problem from owning a given school district is private information to the firm.

The regression in Table 8 indicates that the price of raw milk, the size of the district, the location of processing plants and distribution centers, incumbency, and volume concerns all do not have a significant effect on the bids (deflated) submitted by the milk suppliers. This suggests that, except for synergy, symmetry may reasonably be imposed across districts and bidders in a structural model.

7.2. Simultaneous Procurements

Our structural model has all procurements occurring simultaneously. This is not consistent with the Georgia milk market, since the procurements occur throughout the summer. However, the reduced form estimation suggests that this is an acceptable simplifying assumption. The timing of the procurement was found to have no significant effect on the bidding. Similarly, the volume already placed by each firm was also found to be insignificant.

7.3. Number of Bidders Known and Unaffected by Synergy

Our observations consist of school districts where we presume to know the number of bidders. This is important since, for instance, M=3 means that we assume there are exactly three potential bidders for these lets. The danger of this maintained is that we may find large synergies when, in fact, bidders were simply anticipating more bidders than our data records as participants. To avoid this problem, as noted earlier, we used the surrounding years to help determine the number of potential bidders.

It seems reasonable that as synergy increases, the district becomes more attractive to win and hence more milk suppliers will submit bids for that district. However, our model does not consider an endogenous number of bidders. We assume that the number of bidders is the same regardless of the synergy value of any given district.

8. Empirical Specification and Structural Estimation

Using the equilibrium bid functions that result from the above model, the distribution of bids is fully determined by the distribution of costs and the value of α.

We chose to use a modified Weibull distribution as the cost distribution. While the lognormal distribution is a more traditional choice, it should be noted that for any lognormal distribution there is a Weibull that exactly matches in the first two moments.

We presume that costs are drawn from a distribution that has a cumulative given by

$$F(c;\beta,\gamma,\nu) = 1 - \exp\left[-\left(\frac{c-\nu}{\beta}\right)^{\frac{1}{\gamma}}\right]$$

for $c \in [\nu, \infty)$ where ν is a boundary parameter.[23] Note that $\left(\frac{c-\nu}{\beta}\right)^{\frac{1}{\gamma}}$ is distributed as an exponential(1), and hence that $\left(\frac{c-\nu}{\beta}\right)^{\frac{1}{\gamma}} \in [0, \infty)$. The mean of the cost distribution is given by

$$\beta\Gamma(\gamma+1)+\nu$$

and the standard deviation is given by

$$\beta\sqrt{[\Gamma(2\gamma+1)-\Gamma^2(\gamma+1)]}$$

---

[23] Note that this assumes that the cost distribution is independent of the number of bidders. As discussed in Section 5.1, our data suggest that this is a reasonable assumption.

where Γ is the standard gamma function given by

$$\Gamma(\gamma+1)=\int_0^{\infty} x^{\gamma}\exp(-x)\,dx.$$

8.1. Operational Reformulation of the Bid Function

A convenient version of the bid function from Section 3 is straightforward to obtain. Transformations of the cost and the bid are taken as

$$\phi \equiv (M-1)\left(\frac{c-\nu}{\beta}\right)^{\frac{1}{\gamma}}$$

and

$$y \equiv (M-1)\left(\frac{B(c)-\nu}{\beta}\right)^{\frac{1}{\gamma}}.$$

Note that ϕ is distributed as an exponential($\frac{1}{M-1}$).

Using these, a transformed version of the bid function is given by

$$y^{\gamma}=\phi^{\gamma}+\exp(\phi)\,\Gamma(\gamma+1)\,[1-\Psi(\phi,\gamma)]-\frac{(M-1)^{\gamma+1}}{M}\frac{\alpha}{\beta}\exp\left(-\frac{\phi}{M-1}\right)$$

where Ψ is given by

$$\Psi(\phi, \gamma) = \frac{1}{\Gamma(\gamma)} \int_0^{\phi} x^{\gamma - 1} \exp(-x)\, dx.$$

The use of the transformations considerably simplifies computer coding, as there exist preprogrammed numerical procedures to evaluate Γ and Ψ.

8.2. Boundary Parameter

It is now well-known that boundary conditions play a critical role in the estimation of structural auction models - see e.g. Donald and Paarsch (1996) or Florens, Protopescu and Richard (2001). Problems originate from the fact that the bid function (1) is essentially horizontal for costs drawn from the left tail of their distribution. This creates immediate invertibility problems for low bids, as well as poor identification of the left boundary of the cost distribution. In the present paper, these two problems are taken care of in the following way: (i) we use a boundary estimate for ν, the lower bound of the cost distribution; and (ii) we introduce an auxiliary penalty function which restricts ν to an a priori reasonable region[24]. The combination of these two modifications produces a numerically

---

[24] Bayesians would refer to that function as to a prior density and would, accordingly, reinterpret our estimates as a posterior mode.

stable algorithm. They leave asymptotic distributions of the estimators unaffected (since, in particular, boundary estimates are super-efficient, converging at speed n instead of $\sqrt{n}$). Obviously, they do affect small sample distributions which is why we choose to compute finite sample standard deviations and critical values by Monte Carlo simulations. Finally, while the estimates of β and γ are obviously affected by the selection of the penalty function, we shall verify that our estimate of α, the key parameter of interest, is robust across a broad range of 'sensible' specifications.

Let J be the number of procurements where there are I_j bidders at procurement j. Let ϕ_S denote this expected minimum value of ϕ, conditional on J and $(I_1, ..., I_J)$, and let ϕ_r be the lowest observed value of ϕ in the data.

$$\phi_r = \min_{j \in J} \min_{i \in I_j} \phi_{ij}$$

Also let ϕ_{1j} be the lowest ϕ of all the bidders at procurement j (i.e. that of the winning bidder).

$$\phi_{1j} = \min_{i \in I_j} \phi_{ij}$$

Since ϕ_{1j} is the lowest of M_j draws, it is distributed as an exponential $\left(\frac{M_j}{M_j - 1}\right)$.

Then,

$$\begin{aligned}
\Pr(\phi_r \geq a) &= \prod_{j=1} \Pr\left(\phi_{1j} \geq a\right) \\
&= \prod_{j=1} \Pr\left(\left(\frac{c_j - \nu}{\beta}\right)^{\frac{1}{\gamma}} \geq a \left(\frac{M_j}{M_j - 1}\right)\right)
\end{aligned}$$

Recall that $\left(\frac{c_j - \nu}{\beta}\right)^{\frac{1}{\gamma}}$ is distributed as an exponential(1). Therefore,

$$\Pr(\phi_r \geq a) = \exp\left[-a \sum_{j=1} \left(\frac{M_j}{M_j - 1}\right)\right]$$

indicating that ϕ_r is distributed as an exponential $\left(\left[\sum_{j=1}^{J} \left(\frac{M_j}{M_j - 1}\right)\right]^{-1}\right)$. Hence,

$$\phi_S = E(\phi_r) = \left[\sum_{j=1} \left(\frac{M_j}{M_j - 1}\right)\right]^{-1}.$$

The expected lowest observed transformed bid is then given by

$$y^{\gamma} \geq \phi_S^{\gamma} + \exp(\phi_S) - (\gamma + 1)\left[1 - \Gamma(\phi_S, \gamma)\right] - \frac{(M - 1)^{\gamma+1}}{M} \frac{\alpha}{\beta} \exp\left(-\frac{\phi_S}{M - 1}\right).$$

Substituting in for y and isolating ν suggests the estimator

$$\hat{\nu}_S = \min_{i,j} \left[\begin{array}{c} B(c_{ij}) - \frac{\beta}{(M_j - 1)^\gamma} \left(\phi_S^\gamma + \exp(\phi_S) - (\gamma + 1)\left[1 - \Gamma(\phi_S, \gamma)\right]\right) \\ - \frac{M_j - 1}{M_j} \alpha_j \exp\left(-\frac{\phi_S}{M_j - 1}\right) \end{array} \right].$$

This was used in our estimation.

8.3. Posterior Density

In the estimation, we used the same data as in the reduced form estimation. This was the lowest two bids at each procurement (the winning bid as well as the lowest losing bid), where these procurements have different numbers of bidders. The posterior density of bids is then the joint density of the two smallest of M bids multiplied by a penalty function for ν.

The joint distribution of the two smallest bids is fully described given the joint distribution of the two smallest costs and the bid function. For notational convenience, define $\theta \equiv (\beta, \gamma, \nu)$. Let $H_{12}(b_1, b_2; M, \theta, \alpha)$ be the joint cumulative distribution function of the two smallest bids. Since the bid function is increasing and with the inverse bid function given by $B^{-1}(\cdot; M, \theta, \alpha)$,

$$\begin{aligned} H(b_1, b_2; M, \theta, \alpha) \;=\; & \sum_{j=2} \sum_{i=1} \frac{M!}{i!\,(j-i)!\,(n-j)!} F^i\left[B^{-1}(b_1; M, \theta, \alpha); \theta\right] \\ & \times \left[F\left[B^{-1}(b_2; M, \theta, \alpha); \theta\right] - F\left[B^{-1}(b_1; M, \theta, \alpha); \theta\right]\right]^{j-i} \\ & \times \left[1 - F\left[B^{-1}(b_2; M, \theta, \alpha); \theta\right]\right]^{M-j} \end{aligned}$$

Taking derivatives and simplifying gives the following.

$$\begin{aligned} h(b_1, b_2; M, \theta, \alpha) \;=\; & n\,(n-1)\left[1 - F\left[B^{-1}(b_2; M, \theta, \alpha); \theta\right]\right]^{M-2} \\ & \times \left[\frac{f\left(B^{-1}(b_1; M, \theta, \alpha); \theta\right)}{\frac{\partial}{\partial c_1} B(c_1; M, \theta, \alpha)\big|_{c_1 = B^{-1}(b_1; M, \theta, \alpha)}}\right] \\ & \times \left[\frac{f\left(B^{-1}(b_2; M, \theta, \alpha); \theta\right)}{\frac{\partial}{\partial c_2} B(c_2; M, \theta, \alpha)\big|_{c_2 = B^{-1}(b_2; M, \theta, \alpha)}}\right] \end{aligned}$$

The derivative of the bid function is

$$\frac{\partial}{\partial c} B(c; M, \theta, \alpha) = (M-1) f(c; \theta) \left[\frac{1}{(1 - F(c; \theta))^M} \int_c^1 G(u; M, \theta)\, du - \frac{\alpha}{M}\right]$$

Substituting in then gives the appropriate joint density.

$$
\begin{aligned}
h(b_1, b_2; r, M, \theta, \alpha) &= \frac{n}{n-1}\left[1 - F\left(B^{-1}(b_2; M, \theta, \alpha); \theta\right)\right]^{M-2} \\
&\times \frac{1}{\left(1 - F\left(B^{-1}(b_1; M, \theta, \alpha); \theta\right)\right)^M}\left(\int_{u=B^{-1}(b_1;M,\theta,\alpha)}^{1} G(u; M, \theta)\, du - \frac{\alpha}{M}\right)^{-1} \\
&\times \frac{1}{\left(1 - F\left(B^{-1}(b_2; M, \theta, \alpha); \theta\right)\right)^M}\left(\int_{u=B^{-1}(b_2;M,\theta,\alpha)}^{1} G(u; M, \theta)\, du - \frac{\alpha}{M}\right)^{-1}
\end{aligned}
$$

For the boundary estimator penalty function, we assume that ν comes from a normal distribution with a mean of 7.0 and a standard deviation of 0.15[25]. Our posterior density is then the above joint density of bids multiplied by the penalty function for ν.

Define

$$
z_i \equiv \left(\frac{B^{-1}(b_i; M, \theta, \alpha) - \nu}{\beta}\right)^{\frac{1}{\gamma}}.
$$

Then our objective function becomes the sum, over all procurements, of the log

---

[25] As previously noted, estimates of α were highly robust to changes in this mean and standard deviation.

of the posterior density $\ln p(\beta, \gamma, \alpha; M)$, where

$$\begin{aligned}
\ln p(\beta, \gamma, \alpha; M) &= \ln\left(\frac{M}{M-1}\right) - (M-2)\,z_2 \\
&\quad - \ln\left[\exp(Mz_1)\,\frac{\beta\gamma}{(n-1)^{\gamma}}\,\Gamma(\gamma)\left[1-\Psi\left((M-1)\,z_1,\gamma\right)\right]+\frac{\alpha}{M}\right] \\
&\quad - \ln\left[\exp(Mz_2)\,\frac{\beta\gamma}{(n-1)^{\gamma}}\,\Gamma(\gamma)\left[1-\Psi\left((M-1)\,z_2,\gamma\right)\right]+\frac{\alpha}{M}\right] \\
&\quad + \ln\left(0.15\sqrt{2\pi}\right) - \frac{1}{2}\left(\frac{\nu-7.0}{0.15}\right)^2
\end{aligned}$$

Parameter estimates are obtained by maximizing this function. This maximization only must be done over β, γ, and α, since ν is estimated using these as described above.

8.4. Data and Covariates

The data set is the same as in Section 6 – the winning and lowest losing bids at 133 procurements in 1989, for a total of 266 observations. We did not parameterize β or γ since the reduced form estimation implied that any district specific effects were negligible. The synergy parameter is defined as a coefficient multiplied by the alpha dummy variable used in the reduced form estimation. More specifically, let $d_j = 1$ if procurement j is for an alpha district. Then the synergy parameter for procurement j is given by $\alpha_j = d_j\alpha$. The synergy parameter is constrained to be non-negative during the estimation.

8.5. Estimation Results

Several numerical issues arose in the estimation. These are discussed in Appendix D.

8.5.1. Case 1: No synergy ($\alpha = 0$)

We first estimated with the restrictive assumption that there was no synergy ($\alpha = 0$). As a result of the features of the model, this meant we were simply estimating a standard single object first price procurement. Maximum likelihood estimates are shown in Table 9. The value of the objective function at the estimates is -69.92. Standard errors are shown in parentheses immediately below the estimated parameters. These were calculated by simulating bids and estimating the parameters a total of one thousand times, and then calculating the standard deviation of the estimates across trials. The cost density for these estimates is shown in Figure 11 along with the corresponding bid density for the case of 3 bidders[26]. Note that the bid density is similar to the histogram shown in Figure

---

[26] Only the results for 3 bidders are shown since results for other numbers of bidders are similar, this is the most common number of bidders, and this is closest to the average number of bidders (2.81) observed in the 1989 data used in the estimation. There were 52 procurements with 2 bidders, 55 with 3 bidders, and 26 with 4 or more bidders in this data.

2. The mean and standard deviation of the cost and bid distributions (with 3 bidders) are shown in Table 10.

The bid function for the case of 3 bidders using these estimates is shown in Figure 12 (along with the 45 degree line for reference). These estimates show that, with the average cost of 9.43, the bid markup over this cost is 0.53 cents. From the shape of the bid function, the markup is higher as cost decreases (to approximately 1.8 cents when cost is 7 cents), and goes to zero as the cost increases.

8.5.2. Case 2: Synergy ($\alpha \geq 0$)

This was the same estimation as above with the exception that α is constrained to be non-negative rather than simply zero. The previous results along with a small positive value for α were convenient starting values for the optimization. The parameter estimates obtained are shown in Table 9, along with standard errors calculated as described above[27]. The value of the objective function is -44.61. The cost density for these estimates is shown in Figure 13 along with the corresponding bid density for the case of M=3. Note that this bid density is also similar to the histogram shown in Figure 2. The mean and standard deviation of

---

[27] We investigated the robustness of our parameter estimates to economically reasonable changes in the prior for ν. The estimated parameters were found to be remarkably stable with respect to the choice of a prior for ν. This was especially true for our parameter of interest α.

the cost and bid distributions are also shown in Table 10. The bid function for the case of M=3 using these estimates is shown in Figure 14.

To test the significance of the estimated synergy parameter, we simulated the difference between the objective functions in both cases and created a distribution based on the results. This was done by simulating bids under the null hypothesis (i.e. $H_0 : \alpha = 0,$ case 1 parameter estimates). Then the parameters were estimated under both cases, and the difference of the objective functions was calculated. This was repeated one thousand times. The maximum difference from all these trials was 10.9. Since our estimated difference is 25.3 we reject the null that the true value of α is zero.

In order to test the significance of only allowing the alpha districts to realize synergy, we also estimated this same case while giving each district the synergy (i.e. $d_j = 1, \forall j$). The resulting value of the objective function was not significantly different from that of Case 1.

Interpretation of α The estimate of alpha implies that firms realize a cost advantage of approximately 0.18 cents per half pint when they win adjacent districts. This is significant, but is it big or small? The average cost draw is approximately 9.5 cents. Recall that the average price of raw milk is 7.5 cents. This leaves 2 cents for processing and distribution costs. The price of raw milk and processing technologies are common across milk processors. Differentials between the profits of firms largely come from how they solve their distribution and delivery problems. As noted in Appendix A.2[28], "on the most common types of wholesale milk distribution routes, distribution costs will be between 11 cents and 17 cents per gallon delivered". Distribution costs are then approximately .7 cents to 1.1 cents per half pint. The estimated synergy parameter then implies that owning adjacent school districts could reduce distribution costs for a processor by 15% to 25%.

Recall that our estimation assigned zero synergy to all firms in non-alpha districts. These districts were typically those where only one firm had significant potential to realize the benefits of synergy. Even when only one firm enjoys a synergy there will be an impact on the bidding as non-synergy firms realize they need to be aggressive when bidding against a strong rival. Our estimation of α is based on the difference between the bids at the procurements for alpha and non-alpha districts. By treating the non-alpha districts as if there was no synergy at all, we have effectively narrowed this difference and hence have almost surely

---

[28] Source is University of Minnesota's Agricultural Experiment Station Bulletin number 530-1979.

estimated a lower bound for α.

9. Conclusion

The vehicle routing problem is a significant part of the optimization problem that confronts milk processors. Geographically adjacent contracts can be highly desirable since the marginal cost of servicing a contract adjacent to an existing one is potentially lower than obtaining a similar size contract in an area where the given firm has no existing business. In this paper we account for the possible cost synergies associated with adjacency. Estimation of a reduced form model shows that synergy is significant in this market, while other potentially important market characteristics do not affect the bidding. The Krishna and Rosenthal (1996) synergy framework is extended to a first price sealed bid procurement. Estimation of this structural model reveals significant and substantial synergies. These synergies may account for as much as 25% of distribution costs.

This result implies that any potential model of the collusion that occurred in this market must take into account the affect that such synergies would have on the bidding. This is important for two reasons. First, a major goal of the collusion will likely be to arrange winners of each district to collectively maximize the synergy that the cartel members obtain. Ignoring this will result in a mischaracterization of the collusion. Second, if this synergy is not considered, then a methodology that attempts to distinguish between collusive and non-cooperative behavior may over-detect collusion. Our next goal in this research program is to understand how synergies would be incorporated into a collusive mechanism and to use this to understand what we observe in the collusive years.

References

1. Ausubel, Lawrence M., Peter Cramton, R. Preston McAfee, and John McMillan, 1997, “Synergies in Wireless Telephony: Evidence from the Broadband PCS Auctions,” Journal of Economics and Management Strategy, Vol. 6, No. 3, pp. 497-527.

2. Baldwin, Laura, Robert C. Marshall and Jean-Francois Richard, 1997, “Bidder Collusion at Forest Timber Sales,” Journal of Political Economy, Vol. 105, pp. 657-699.

3. Bierens, H. J. (2002), "EasyReg International", Department of Economics, Pennsylvania State University, University Park, PA

4. Donald, Stephen and Harry Paarsch, 1993, “Piecewise Pseudo-Maximum Likelihood Estimation in Empirical Models of Auctions,” International Economic Review, Vol. 34, pp. 121-148.

5. Donald, Stephen and Harry Paarsch, 1996, “Identi…cation, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm”, Econometric Theory, Vol. 12, No. 3, pp. 517-67.

6. Fischer, Martin, Jerome Hammond and Wallace Hardie, 1979, “Fluid Milk

Processing and Distribution Costs," Station Bulletin 530, Agricultural Experiment Station, University of Minnesota.

7. Florens, Jean-Pierre, C. Protopopescu and Jean-Francois Richard, 2001, "Inference in a Class of Game Theory Models," University of Pittsburgh, Department of Economics, Working Paper.
8. Guerre,Emmanuel, Isabelle Perrigne, and Quang Vuong, 2000, "Optimal Nonparametric Estimation of First-Price Auctions", Econometrica, Vol. 68, No. 3, pp. 525-74.
9. Hewitt, Cynthia, James McClave and David Sibley, 1993, "Incumbency and Bidding Behavior in the Dallas-Ft. Worth School Market," mimeo, University of Texas at Austin.
10. Krishna, Vijay and Robert Rosenthal, 1996, "Simultaneous Auctions with Synergies," Games and Economic Behavior, Vol. 17, pp. 1-31.
11. Laffont, Jean-Jacques, Herve Ossard, and Quang Vuong, 1995, "Econometrics of First-Price Auctions," Econometrica, Vol. 63, pp. 953-908.
12. Levin, Jonathan, 1997, "An Optimal Auction for Complements", Games-and-Economic-Behavior, Vol. 18, No. 2, pp. 176-92.

13. Pesendorfer, Martin, 1994a, “A Study of Cartel Agreements in First Price Auctions. Part I: Cartel Agreements,” mimeo, Northwestern University.

14. Pesendorfer, Martin, 1994b, “A Study of Cartel Agreements in First Price Auctions. Part II: Cartel Bidding Behavior,” mimeo, Northwestern Univeristy.

15. Pesendorfer, Martin, 2000, “A Study of Collusion in First-Price Auctions”, Review of Economic Studies, Vol. 67, No. 3, pp. 381-411.

16. Porter, Robert and Douglas Zona, 1999, “Ohio School Milk Market: An Analysis of Bidding,” RAND Journal of Economics, Vol. 30, No. 2, pp. 263-88.

17. Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, Numerical Recipes in Fortran, Second Edition, Cambridge University Press, 1992

18. Raiff, Matthew, 1997, “Identifying Collusive Behavior in the Georgia School Milk Market,” mimeo, Penn State University.

19. Qi, Junfeng, 1997, “Structural Analysis of Asymmetric Auctions”, mimeo, Duke University.

20. Sareen, Samita, 1999, "Posterior Odds Comparison of a Symmetric Low-Price, Sealed-Bid Auction within the Common-Value and the Independent-Private-Values Paradigms", Journal of Applied Econometrics, Vol. 14(6), pp. 651-76.

Appendices

A. Georgia School Milk Market

A.1. Technology

The production of fluid milk consists of pasteurizing it, removing all the butterfat from the raw milk and combining it in varying proportions, along with flavorings, vitamins and other ingredients, to produce the four types of milk sold in the school milk market: whole white, low-fat chocolate, low-fat white and skim. The various types of half-pints of milk are then packaged and delivered.

A.2. Cost

Our primary source for information regarding the cost structure facing a fluid milk processing facility is the University of Minnesota's Agricultural Experiment Station Bulletin number 530-1979. This study focuses on the estimation of fluid milk processing and distribution costs. Processing costs are estimated using the "economic-engineering method" for alternative sizes of processing plants. Estimates of milk distribution costs for various types of distribution routes are presented. For our purposes the following analysis is most important. The authors write, “on the most common types of wholesale milk distribution routes, distribution costs will be between 11 cents and 17 cents per gallon delivered. In contrast, processing costs minus packaging materials expense were 9.94 cents to 16.19 cents per gallon. Distribution costs are an important determinant of the total costs incurred by milk processors. For many milk processing firms, distribution costs exceed the costs of processing milk.” This study strongly suggests the importance of synergies. The authors demonstrate how vital the vehicle routing problem is in determining the profitability of a processing plant.

The major factor input for a milk processing plant, raw milk, is regulated by Federal Milk Marketing Orders. The FMO equals the Minnesota / Wisconsin price plus a fixed differential typically increasing in distance from the upper Midwest. Various discussions with industry experts have confirmed that, on average, the price of raw milk accounts for 65% of the cost of producing a half-pint of milk.

A.3. Supply of Raw Milk

A processing plant’s source for raw milk, either from a dairy cooperative or an independent farmer, is extremely inflexible to changes in the quantity purchased by the processing plant. Typically, if a processing plant wished to increase or decrease production by 10% they would have to notify their suppliers 6 months in advance unless they were willing to pay a significant premium or penalty for the milk. An increase or decrease of 20% would require approximately a one year notice to the suppliers. This property of the fluid milk market makes it very reasonable to assume that each bidder only wants a subset of the school milk contracts that are available at any one time. An increase or a decrease in the quantity of raw milk a processor purchases can happen in real time but it will require significant premiums.

A.4. Milk Processor's Vehicle Routing Problem

This subsection highlights the general characteristics of the vehicle routing problem faced by fluid milk processing plants. The problem is most easily understood by breaking it down into the four key components: the nature of demand, the information on demand, capital and labor requirements, and scheduling requirements.

A.4.1. Nature of Demand

The primary activity is delivery of fluid milk but the milk cases are recycled so the peddle truck driver is also making a pick-up at most locations. The processor provides a portfolio of products ranging from ice tea and fruit drinks to ice cream and butter. The size of the delivery is monitored and controlled by the peddle truck driver. It is imperative that all customers never run out of milk. Because of this concern a priority list for customers can arise when a customer places an order due to unexpected sales at their establishment.

A.4.2. Information on Demand

All demand is not known in advance. The peddle truck driver is solving an inventory control problem at each of the locations. Demand is very predictable but only after a significant amount of time will the driver become very efficient at solving this problem. The changing landscape of the retail market due to entry and exit will mean that sales at a particular location will be changing and must be accounted for by the driver. The customer base is very stable. Brand loyalty is very important in the industry and reputable establishments do not shop their dairy contract. Retail outlets and schools establish routines with the peddle truck drivers. These consist of the time of day, number of deliveries per week and standard quantity. A real-time inflow of demand problem is very common with schools because the cafeteria managers are notoriously poor at monitoring their inventory needs.

A.4.3. Capital and Labor Requirements

The 18 ft. flat-bed truck was the standard peddle truck during the 1980s. It is common for the peddle truck driver to load their own trucks. The drivers place their orders the night before, the product is pulled and put onto the loading dock and the driver then loads the truck while keeping in mind the route he will be driving during the course of the day. Because milk is extremely heavy an efficient load on a peddle truck is very important for timeliness in the delivery process. The union drivers work on flat rate plus 3% commissions with the average salary being $37,500 in 1997. Turnover for drivers is very uncommon because processors recognize the substantial learning curve facing a driver. Schools frequently require the processor to provide coolers for the milk. The coolers are a non-trivial capital investment for the processor.

A.4.4. Scheduling Requirements

Customers are assigned certain days of the week. But because schools operate on an erratic schedule the scheduling process for a school district is difficult. The window times for delivery are somewhat hard for a new account but as the driver establishes a reputation with the customer these windows become softer. A typical school district is part of three to seven different routes. A dedicated school route would require a very dense and highly populated school district such as Atlanta. Normally, school contracts are added onto existing retail routes.

B. Selection of Districts

The data originally contained information on procurements for institutions other than public schools, such as hospitals and prisons. Removing these left 1,986 procurements. The data was then restricted to be from 1978 to 1991, leaving 1,809 procurements. We then removed all observations that contained a potential scoring problem. Information on the bids submitted was not available for 47 observations. Of the remaining, 94 procurements resulted in two or more milk suppliers servicing the schools in each district. Contracts were negotiated in 47 instances, usually directly as extensions of the previous year. In 16 procurements, the lowest bid was not accepted due to past service problems. Finally, 4 remaining procurements were for a length of time other than the standard year. Removing all procurements with these problems left 1,601. The final criteria implemented was that there must be a determinable number of bidders for each district. This left a total of 1,531 procurements, generally with 2 to 5 bidders at each, and a total of 4,829 bids. The most sparse year of data was 1978 with 45 procurements, and the most complete was 1989 with 163 procurements.

C. Scoring Function

The scoring function is an imperfect substitute for the methods employed by the procurement officials throughout the State of Georgia. Three important issues are addressed by the scoring function. First, the function is designed to handle how usage rates across milk items affect the scoring of a bid. Second, the function accounts for the differences in a bid submitted with an escalation clause and one submitted without an escalation clause. Third, the function will deflate the bids to account for the differences throughout time. Throughout the paper, the term "bids" refers to actual bids that have been adjusted in all three ways.

C.1. Relative Weights

Each firm submitted bids on the individual milk items[29]. These bids generally were identical for each of the items, but often were not. To map these multi-dimensional bids into one dimension, we used a weighted average. The weights were constructed using a data set provided to us by Info-Tech Inc. The data set contained the ex-post usage rates across milk items. These usage rates were obtained from the invoices from 26 school districts. All observations with a usage rate for a particular type of milk exceeding 80%, except when the only two items

---

[29] There are four separate milk items: low-fat chocolate, low-fat white, skim, and whole white.

were low-fat white and whole white, were deleted from the sample because these observations were most likely recorded improperly by the school district. For example, we do not believe that in 1984 Bryan County bought 96% chocolate milk and only 4% whole white. Using ex-post information to score a bid poses a series of potential problems. We do not know for certain how closely this information relates to the scoring procedure used by an individual school district. Also, the accuracy of the ex-post usage rate numbers is a serious problem for most school milk data sets. A careful examination of the ex-post numbers in the Texas data set leads one to believe these numbers are largely suspect.

Average usage rates were constructed for the remaining data set. Four different weighting functions were developed to account for the four different combinations of milk items seen at procurements in our data. They are presented in Table 11.

C.2. Escalation Clause

Many firms submitted two bids for a single contract, one with an escalation clause and one without. We made the assumption that any such bid pair would be, on average, scored as equivalent bids. Using this, we calculated a multiplier for each month that, on average, scaled the escalated bids in these pairs to be the same as the bid without an escalation clause. All bids containing an escalation clause were scaled by the appropriate multiplier. The multipliers are presented in Table 12.

The missing values are for those months in which an escalated bid did not appear in our final data set. The ”other” value was used when the specific month and day of the procurement was not known, which was the case for 274 of our 1531 observations.

Note that the multiplier is higher (and hence that bids with an escalating clause are relatively lower) for the years with higher inflation. This is consistent with what one would expect. When inflation is higher, an escalating clause places more risk on the school districts, therefore making such bids less desirable. Hence, to make bids with an escalating clause competitive to those without, they must be relatively lower with higher inflation.

C.3. Deflating

The time period we are using has a significant amount of inflation. We constructed a price index that was more appropriate than the CPI for our data. As noted above, the price of raw milk and the distribution costs are two of the major factors in determining the cost of a firm to produce and supply milk. Our index was a combination of the price of raw milk and the motor fuel price index for the

South[30], using weights of 0.8 and 0.2 as suggested by the University of Minnesota's study. The weighted and escalated bids were divided by this index, and these were used as our scored bids. Using May 1978 as the base month and year, the values of the index are presented in Table 12.

D. Numerical Considerations

All of the code was written in Fortran 77, and was compiled and run on a RS/6000 AIX machine. All of the floating point numbers used were double precision. All integers were declared as 4 byte variables. No implicit data types were used to ensure that values were stored properly. Each individual part of the code was tested separately. At several points in the code the relevant variables were checked to ensure they were neither too small or too large.

D.1. Estimation of Boundary Parameter

Recall that the estimator of the boundary parameter is given by

$$b_S = \min_{i,j}\left[\begin{array}{c} B(c_{ij}) - \frac{\beta}{(M_j-1)^{\gamma}}\left(\phi_S^{\gamma} + \exp(\phi_S) - (\gamma+1)\left[1 - \Psi(\phi_S,\gamma)\right]\right) \\ - \frac{M_j-1}{M_j}\alpha_j \exp\left(-\frac{\phi_S}{M_j-1}\right) \end{array}\right].$$

---

[30] This is available from http://www.economagic.com/em-cgi/data.exe/blscu/CUUR0300SETB.

The major numerical concerns involved with this are the exponential and the functions ¡ and ª. The exponential was evaluated with the implicit function DEXP. ¡ and ª are evaluated directly with the IMSL functions DGAMMA and DGAMDF respectively.

D.2. Inversion of the Bid Function

The bid function must be inverted to recover a cost from an observed bid, given a full set of parameters. Since an analytic solution for the inverse bid function is not available, the inversion was done through bisection. The bid minus the bid function creates an expression that has a value of zero at the desired cost. The bisection was done with an double precision adaptation of the function "rtbis" as described in Press (1992, pages 346-7). This was chosen for its simplicity but mainly because it guarantees that the bisection will succeed (up to a specified tolerance) as long as the root is initially bracketed. The bisection is accurate up to 10^{i15}.

D.3. Evaluation of the Objective Function

This step has the same concerns as the estimation of the boundary parameter, and were addressed in the same way. In addition, the objective function uses factorials.

These were evaluated using the implicit function DFAC. All of the arguments to this function were small integers, so there was no issue of inaccuracy.

D.4. Maximization of the Objective Function

The objective function was maximized with respect to the parameters of the cost distribution, $\theta = (\beta, \gamma, \nu)$, and the synergy parameter α. All of the optimizers used throughout this project were inherently minimizers, so the objective function was simply multiplied by -1. This minimization was initially attempted using the IMSL routines DUMPOL and DUMINF. These proved to not be very robust, and failed several minimization tests. The minimizer that was used to produce the estimates reported below was written by Jean-Francois Richard, and was an adaptation of the multidimensional simplex routine "amoeba" described in Press (1992, pages 402-6). This proved to be extremely robust, passed all accuracy tests, and was not noticeably slower than the IMSL routines. Another advantage of this method is that only function evaluations are required, so the (analytical or numerical) gradient was not ever required to be calculated.

This minimizer searches over the parameter space, starting with some initial range, until the decrease in the function value achieved by each step is less than a specified tolerance. Then the minimizer begins the search again around the parameter values of this initial convergence, and again continues until the specified tolerance is reached. This restarting is done to help ensure that the search does not end on a local minimum. The specified tolerance was 10^{-12}. This particular minimizer is extremely robust (relative to other minimizers) in the sense that there is typically great flexibility in the starting ranges that can be chosen which, at termination, result in the same minimum point.

Figure 1: Georgia School Districts, Color-Coded by 1989 Winner

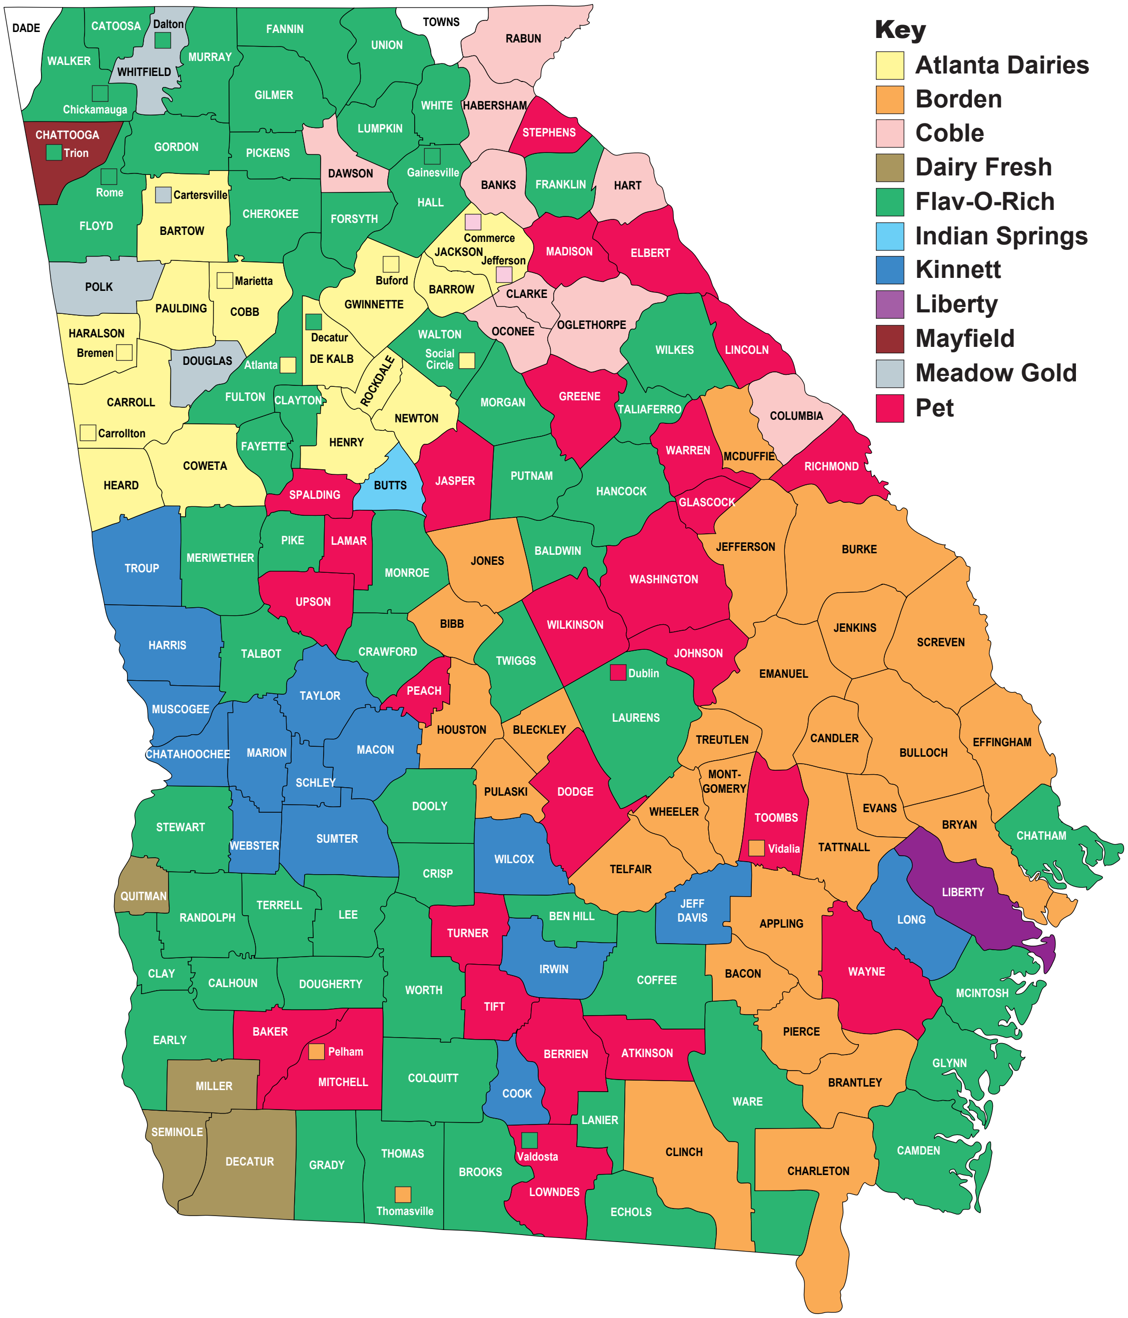

Table 1: Summary Statistics[*]: 1978-1991				
	All Bids (Cents)	**Winning Bids (Cents)**	**Non-Winning Bids (Cents)**	**Price of Raw Milk (Cents)**
Mean	10.63	10.28	10.79	7.53
Median	10.70	10.37	10.89	7.57
Std. Dev.	1.20	1.08	1.21	0.56
Min.	7.61	7.61	7.81	6.12
Max.	15.99	14.85	15.99	8.80
Count	4829	1531	3298	

[*] Throughout the paper, bids are standardized as described in Section 4.1 and Appendix C.

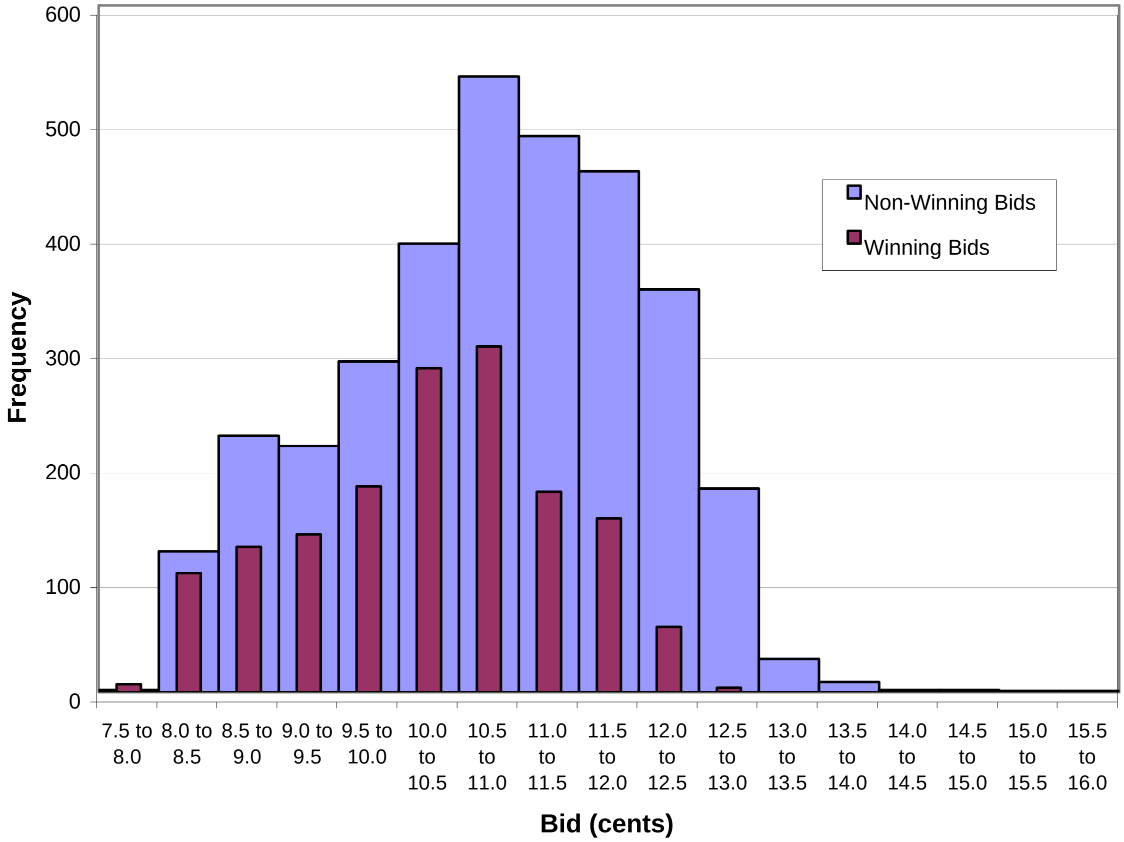

Figure 3: Average Winning Bid by Year

Columns give the number of observations by the number of bidders (number provided if > 9, with total shown at top)

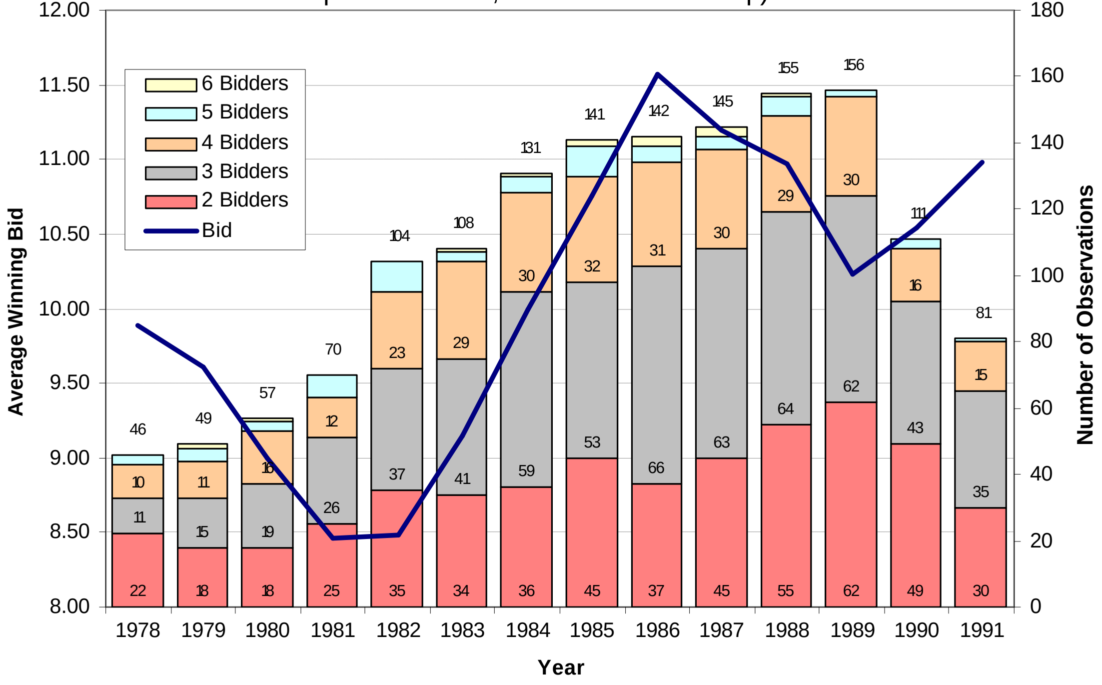

Figure 4: Difference of the Average Winning Bid: Three Bidders vs. Two Bidders and Four Bidders vs. Two Bidders

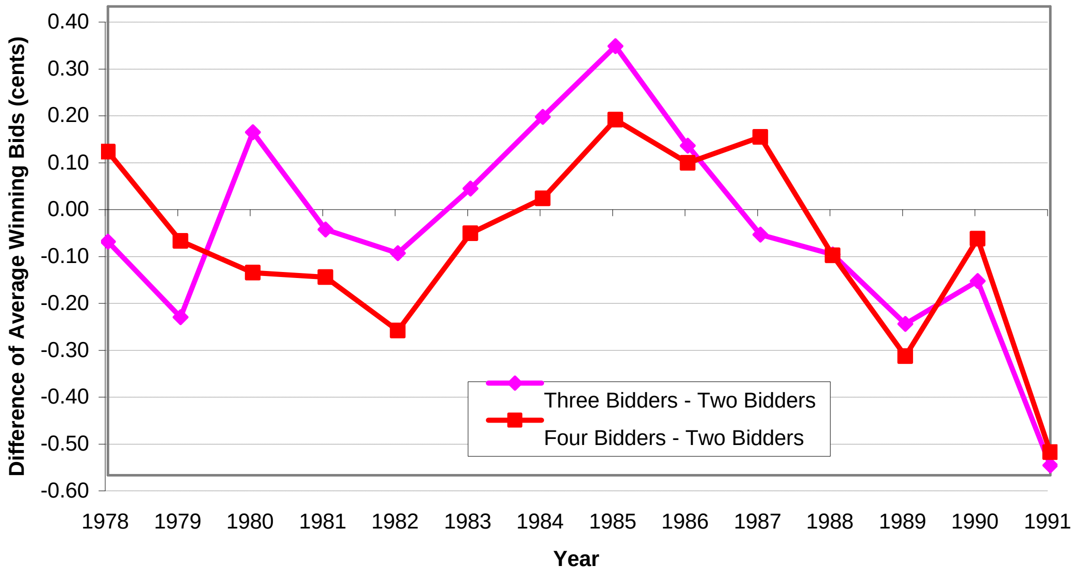

Figure 5: Average Winning Bid by Enrollment Category

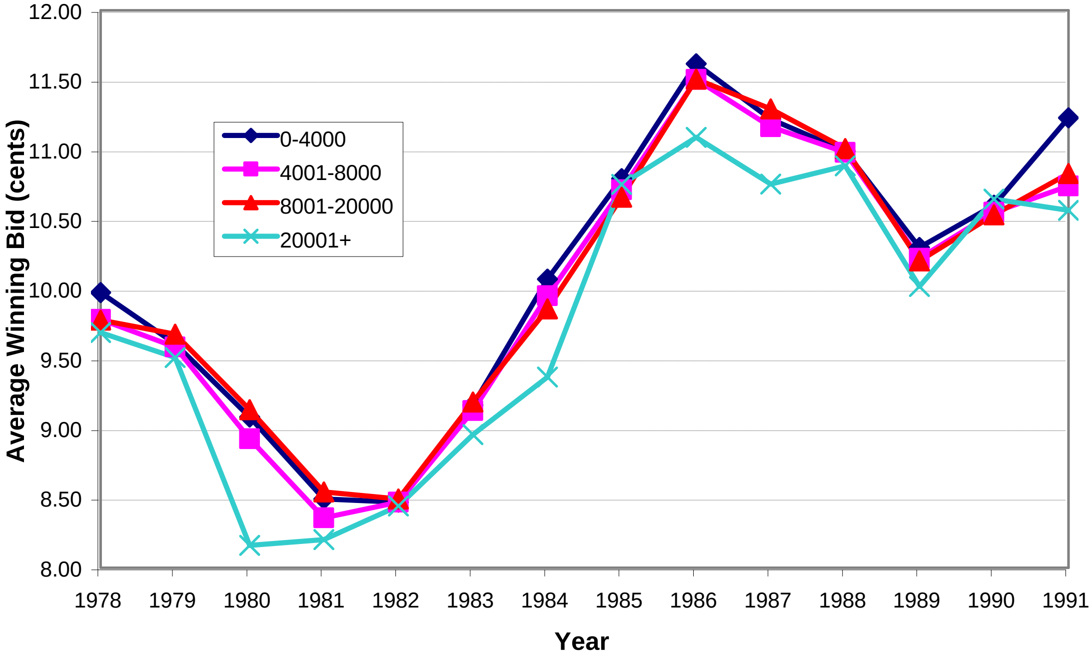

Figure 6: Percentage of Procurements in Each Month by Enrollment Category: 1978-1991

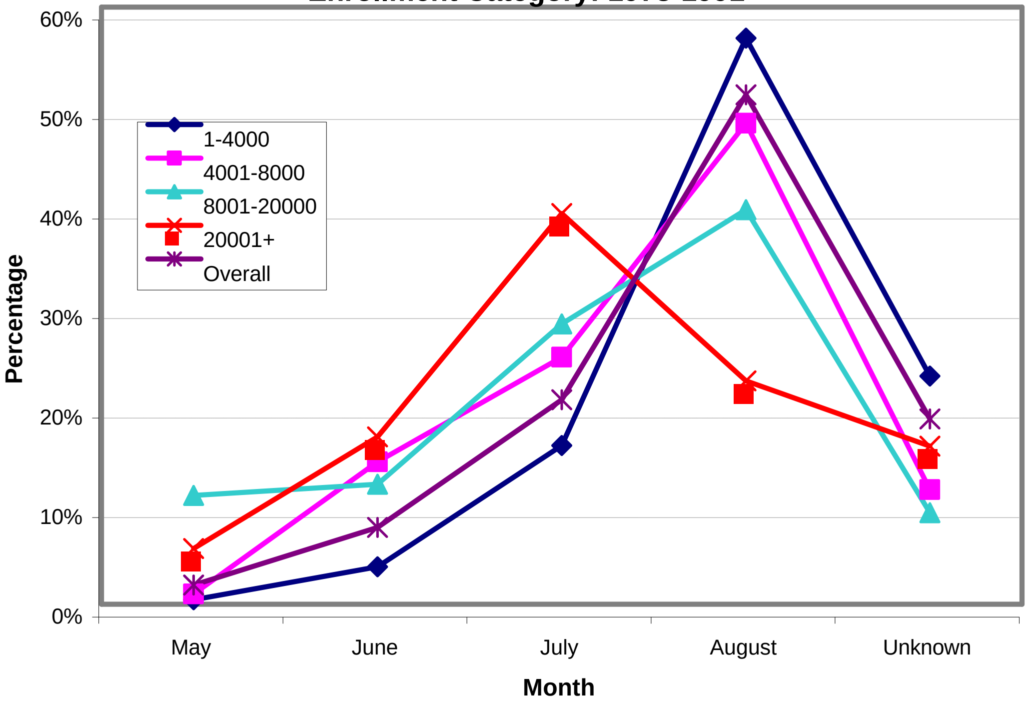

Figure 7: Difference of the Average Winning Bid: July vs. June and August vs. June

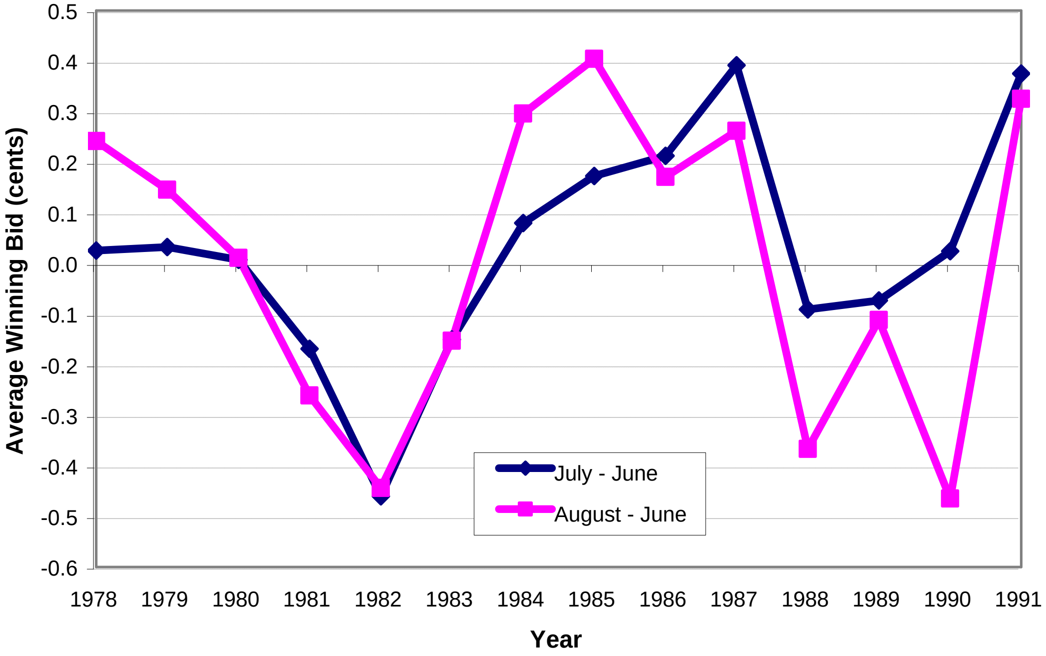

Table 2: Average Winning Bids in Non-Collusive Years: 1978-1982, 1989-1991 (No. of Obs.)				
Enrollment Category	**June**	**July**	**August**	**Month Unknown**
1-4000	9.77 (16)	9.89 (84)	9.99 (255)	9.25 (70)
4001-8000	9.77 (24)	9.89 (47)	9.99 (78)	9.25 (15)
8001-20000	10.26 (8)	9.69 (22)	9.67 (26)	9.11 (4)
20001+	10.18 (5)	9.68 (14)	10.30 (4)	8.12 (1)
Overall	9.87 (53)	9.82 (167)	9.90 (363)	9.23 (90)

Table 3: Location Categories		
Category	**Description**	**No. of Obs.**
1	Contains 1 Distribution Center	83
2	Contains 2+ Distribution Centers	38
3	Adjacent to 1 Distribution Center	121
4	Adjacent to 2+ Distribution Centers	247
5	No Close Distribution Centers	167

Table 4: Average Winning Bids in Non-Collusive Years by Location Category and Number of Bidders (No. of Obs.)

Location Category	**2 Bidders**	**3 Bidders**	**4 Bidders**	**5 Bidders**
1	9.57 (21)	9.63 (33)	9.62 (26)	8.58 (3)
2	9.82 (12)	9.53 (15)	9.22 (10)	9.66 (1)
3	10.37 (42)	9.59 (46)	9.70 (28)	8.31 (4)
4	9.76 (97)	9.82 (91)	9.61 (41)	9.17 (18)
5	10.05 (83)	9.80 (55)	9.44 (23)	9.74 (6)

Table 5: Incumbency Categories		
Category	**Description**	**No. of Obs.**
1	Incumbent Won	377
2	Incumbent Bid but Lost	150
3	Incumbent did not Bid	52
4	Incumbent Unknown	90

Table 6: Average Winning Bids in Non-Collusive Years by Incumbency Category and Number of Bidders (No. of Obs.)

Incumbency Category	**2 Bidders**	**3 Bidders**	**4 Bidders**	**5 Bidders**
1	9.95 (166)	9.66 (133)	9.75 (65)	8.86 (13)
2	10.35 (32)	9.95 (73)	9.61 (35)	9.43 (10)
3	10.57 (24)	10.07 (14)	9.81 (8)	9.20 (6)
4	9.14 (35)	9.24 (28)	8.91 (24)	9.14 (3)

Figure 8: Difference in the Percentage of the Total Enrollment Won from the Previous Year for Each of the Four Largest Firms

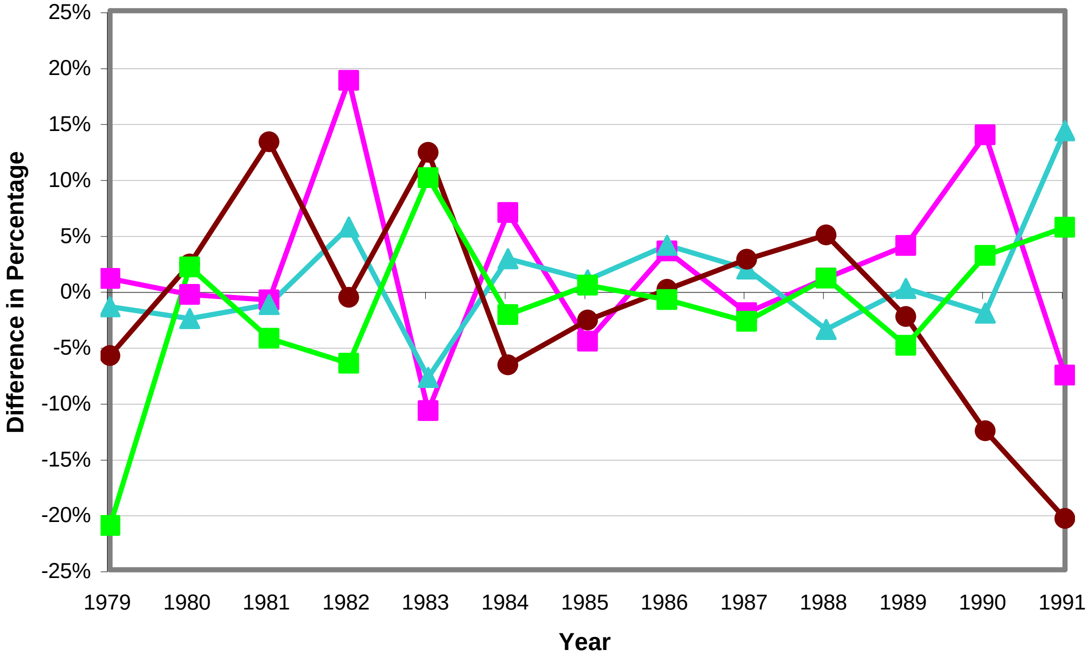

Figure 9: Winning Bid vs. Second Lowest Volume Ratio, 1989

Volume ratio is calculated as the proportion of the total enrollment in 1988 won by a firm divided by the proportion of the total enrollment let to date in 1989 won by the firm

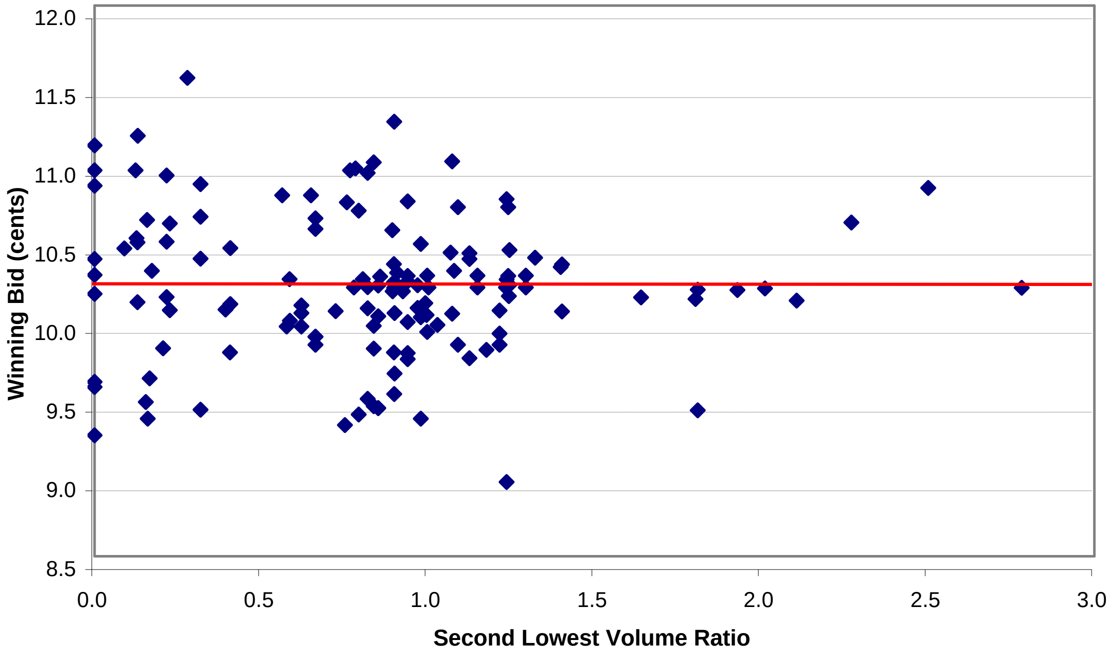

Figure 10: Cost Synergies Through Adjacency

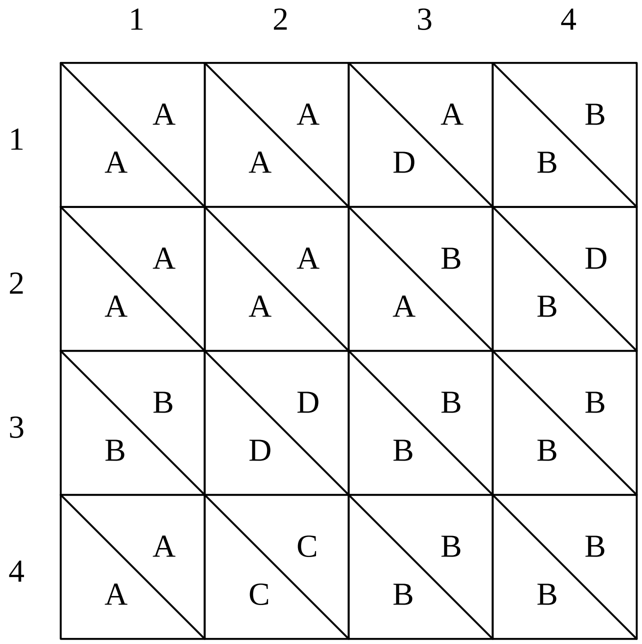

16 hypothetical school districts in year *t*, 4 firms, winners denoted in year *t-2* and *t-1*.

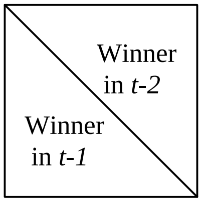

Table 7: Average Winning Bids in 1989 for Alpha and Non-Alpha Districts, by Number of Bidders (No. of Obs.)

	Alpha	**Non-Alpha**
Overall	10.11 (47)	10.23 (86)
2 Bidders	10.41 (18)	10.36 (34)
3 Bidders	9.91 (41)	10.23 (31)
4 Bidders	10.03 (5)	10.00 (20)

Table 8: Regressions of ln of Bids for 1989 Procurements

	Regression 1	**Regression 2**	**Regression 3**
Variable	**Coefficient (Std. Err.)**	**Coefficient (Std. Err.)**	**Coefficient (Std. Err.)**
Constant	3.18^{**} (1.05)	3.27^{**} (1.05)	2.36^{**} (0.00468)
ln of Price of Raw Milk	-0.396 (0.516)	-0.442 (0.515)	-
Dummy for 3 bidders	-0.0278^{**} (0.00622)	-0.0277^{**} (0.00624)	-0.0291^{**} (0.00603)
Dummy for 4+ bidders	-0.0458^{**} (0.00803)	-0.0463^{**} (0.00806)	-0.0427^{**} (0.00778)
Enrollment	$-8.98\bullet10^{-3}$ $(7.44\bullet10^{-3})$	$-6.10\bullet10^{-3}$ $(8.38\bullet10^{-3})$	-
Enrollment Squared	$1.26\bullet10^{-3}$ $(1.04\bullet10^{-3})$	$9.52\bullet10^{-4}$ $(1.15\bullet10^{-3})$	-
Early Dummy	$9.27\bullet10^{-3}$ $(1.06\bullet10^{-2})$	$-7.30\bullet10^{-3}$ $(1.08\bullet10^{-2})$	-
Location Dummy	$-1.47\bullet10^{-4}$ $(5.75\bullet10^{-3})$	$-6.39\bullet10^{-4}$ $(5.78\bullet10^{-3})$	-
Incumbency Dummy	0.0134 (0.0110)	0.0135 (0.0110)	-
2nd Lowest Volume Ratio	-0.0116 (0.00883)	-0.0127 (0.00895)	-
Alpha Dummy	-0.0107^{*} (0.00581)	-0.0111 (0.0101)	-0.0100^{*} (0.00570)
Alpha Dummy · Enrollment	-	$-5.39\bullet10^{-3}$ $(2.41\bullet10^{-2})$	-
Alpha Dummy · Enrollment Squared	-	$-4.21\bullet10^{-3}$ $(6.98\bullet10^{-3})$	-
R^2	0.171	0.174	0.135
F	5.28	4.43	13.60
No. of Obs.	266	266	266

*Significance of 10%
**Significance of 1%

Table 9: Structural Estimates

Parameters	**Case 1** **(a = 0)**	**Case 2** **(a ³ 0)**
b	2.748 (0.0185)	2.841 (0.0227)
g	0.414 (0.0104)	0.377 (0.0116)
n	6.996 (0.000623)	6.996 (0.000558)
a	-	0.179 (0.0532)
Objective Function	-69.92	-44.61

Table 10: Mean and Standard Deviation of Estimated Cost and Bid Distributions (3 Bidders)

	Statistic	**Cost Dist.**	**Bid Dist.**
Case 1 **(a = 0)**	**Mean**	9.43	10.04
	Standard Deviation	1.07	0.83
Case 2 **(a ³ 0)**	**Mean**	9.52	10.08
	Standard Deviation	1.02	0.79

Figure 11: Estimated Cost and Bid Densities for Case 1

The bid density is based on 3 bidders

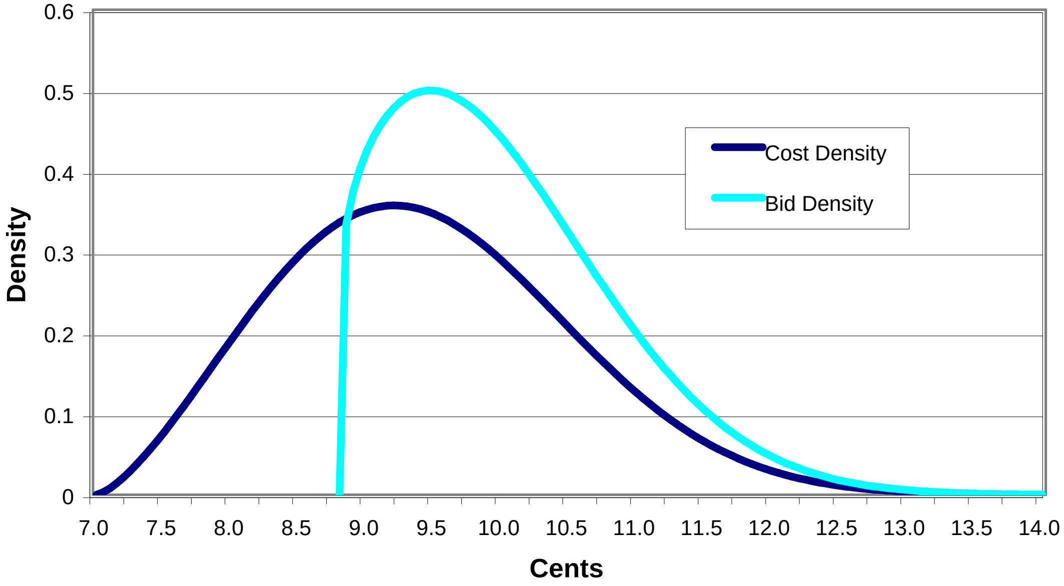

Figure 12: Estimated Bid Function for Case 1

The bid function is based on 3 bidders

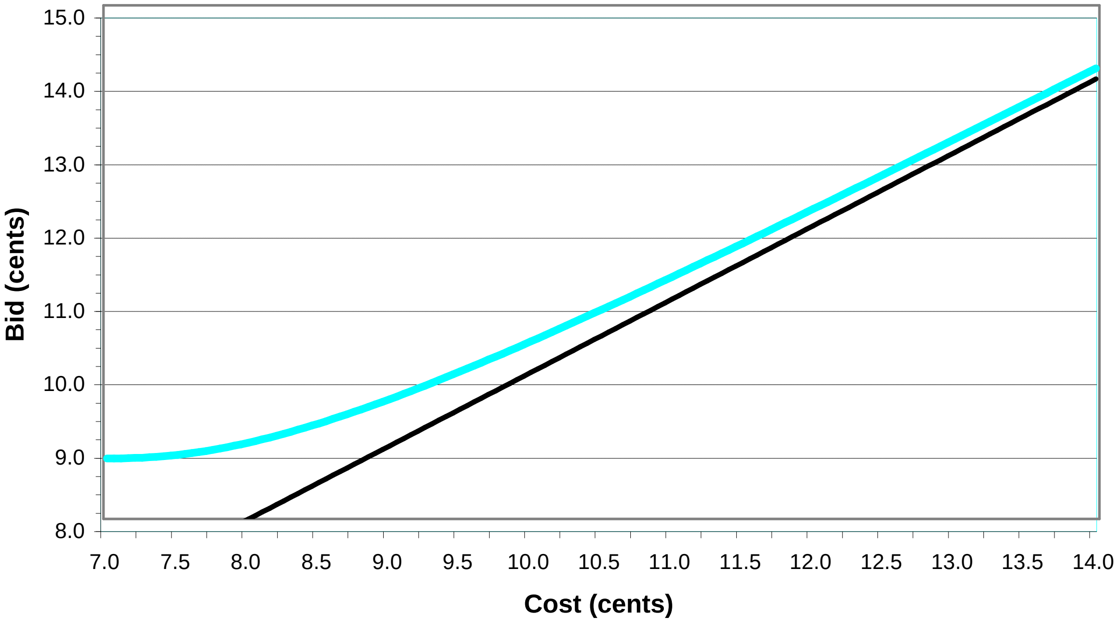

Figure 13: Estimated Cost and Bid Densities for Case 2

The bid density is based on 3 bidders

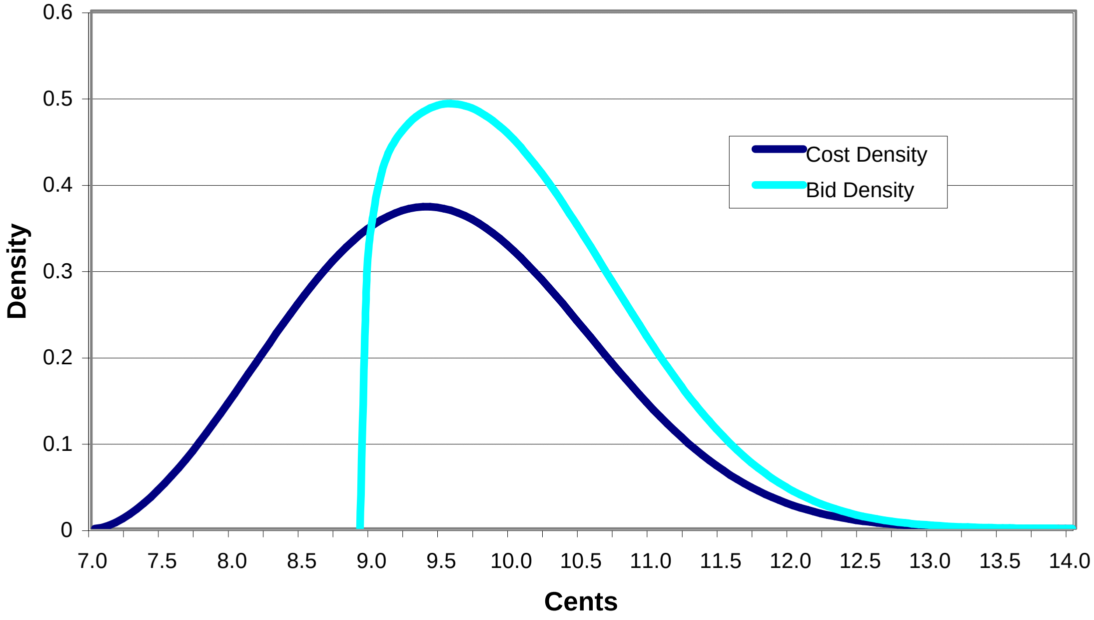

Figure 14: Estimated Bid Function for Case 2

The bid density is based on 3 bidders

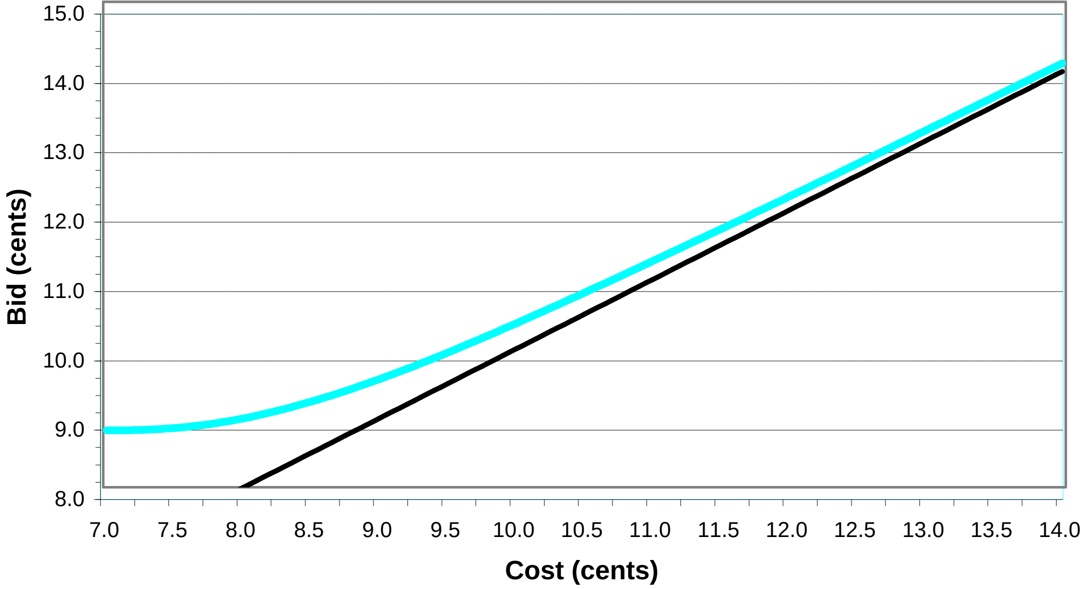

Table 11: Weighting Function			
	LFC[1]	**LFW[2]**	**WW[3]**
Category 1[4]	0.520	0.085	0.395
Category 2	-	0.087	0.913
Category 3	0.471	-	0.529
Category 4	0.57	0.43	-

Table 12: Escalation Multipliers (top) and Deflating Index (bottom, normalized to May 1978)					
	May	**June**	**July**	**August**	**Other**
1978	1.039 1.000	- 1.013	1.042 1.021	1.058 1.028	- 1.016
1979	- 1.212	1.067 1.260	1.055 1.302	1.054 1.337	1.048 1.278
1980	- -	- 1.602	1.066 1.602	1.054 1.601	1.045 1.600
1981	- 1.751	- 1.744	1.036 1.740	1.037 1.735	1.036 1.743
1982	- 1.621	1.035 1.679	1.026 1.699	1.016 1.694	1.019 1.673
1983	1.016 1.444	1.018 1.446	1.019 1.447	1.022 1.446	1.028 1.446
1984	- 1.392	1.014 1.401	1.016 1.401	1.043 1.402	1.039 1.401
1985	1.018 -	1.021 1.379	1.023 1.369	1.020 1.349	1.022 1.375
1986	1.019 1.286	1.035 1.293	1.028 1.280	1.025 1.267	1.023 1.281
1987	1.015 1.304	1.015 1.312	1.018 1.317	1.024 1.329	1.025 1.316
1988	1.026 1.269	1.019 1.262	1.018 1.268	1.011 1.276	1.018 1.269
1989	- 1.370	1.015 1.373	1.024 1.369	1.019 1.371	1.015 1.370
1990	1.109 1.429	1.076 1.455	1.077 1.485	1.069 1.551	1.067 1.480
1991	1.040 1.312	1.038 1.322	1.056 1.326	1.084 1.354	1.058 1.330

---

[1] LFC is low-fat chocolate.

[2] LFW is low-fat white.

[3] WW is whole white.

[4] Category 1 bids were for low-fat chocolate, low-fat white, and whole white. Category 2 bids were for low-fat chocolate and whole white. Category 3 bids were for low-fat chocolate and low-fat white. Category 4 bids were for low-fat white and whole white.

Figure 15: Histograms of All Bids for Two, Three, or Four (Potential) Bidders

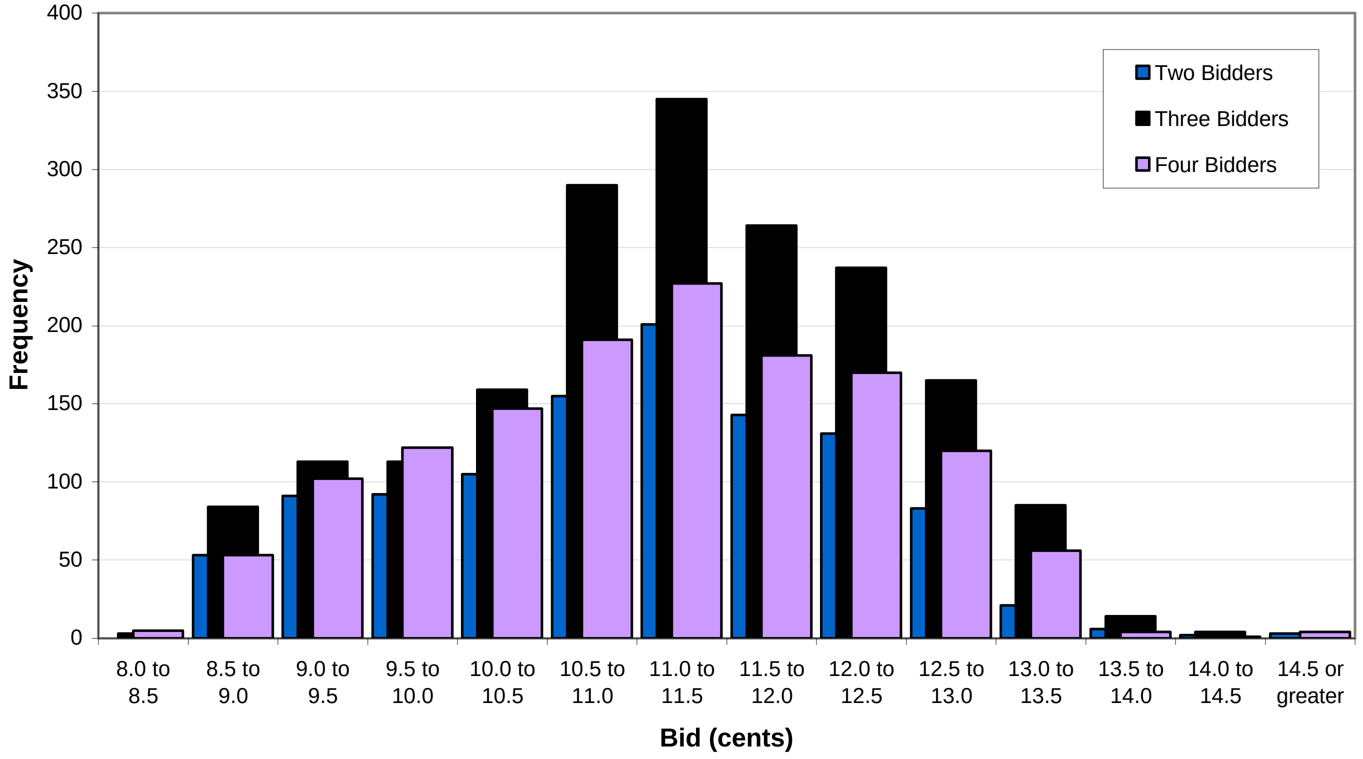

Chapter 2

Numerical Analysis of Asymmetric Auctions with Optimal Reserve Prices

Robert C. Marshall and Steven P. Schulenberg

1. Introduction

The benchmark result of the theory of auctions is the revenue equivalence theorem (RET). The RET is derived within the independent private values (IPV) framework. One of the underlying assumptions of the IPV model is that all bidders draw valuations from the same common underlying distribution. If this assumption is violated then the RET no longer holds. The first price and second price auctions will yield different expected revenue for the seller if bidders draw values from different distributions (excepting pathological cases).

Bidder asymmetry is natural to investigate for two related reasons. First, in many practical circumstances bidders are very different, ex ante. In many cases these differentials can be palatably proxied via distributional heterogeneity. Second, if bidders are ex ante symmetric then collusion by any strict subset will produce bidder asymmetry. In this note we focus our attention on a kind of asymmetry that might arise from collusion by symmetric bidders.

Lebrun (1996, 1998a, 1998b, 1999) and Maskin and Riley (2000a, 2000b) have obtained numerous analytic results for the case of bidder value asymmetry. For a particular kind of bidder asymmetry that might reasonably arise from collusion by bidders whose valuations are distributed continuously uniformly, it is well known that first price auctions generate higher expected revenue for the seller than do second price auctions.[1] This result is derived in the absence of any reserve price policy by the auctioneer. This note addresses the question of whether this revenue ranking will reverse when the seller acts strategically by setting an optimal reserve price.[2]

Li and Riley (1999) have also performed some numerical analysis of asymmetric auctions with reserve prices. Despite using the optimal second price reserve for both auction types, rather than the optimal reserve for each type, they found that the first price auction provided a higher expected revenue in the examples that they considered. These examples were of a different type of asymmetry than we consider. Their results indicate that, for some types of asymmetry, the revenue ranking is as expected.

---

[1] Marshall, Meurer, Richard and Stromquist (1994) provided numerical analysis for collusive-like asymmetries. Lebrun (1998b) subsequently provided the general analytic results.

[2] This was inspired by a conjecture made in an earlier version of Maskin and Riley (2000a).

In this note we numerically calculate the optimal reserve price for the first price auction and then the corresponding expected revenue. We then compare this to the expected revenue for a second price auction with an optimal reserve. We find evidence in support of the reversal of the revenue ranking when an optimal reserve price is used – second price auctions generate higher expected revenues than first price auctions. This has significant practical implications. First, the literature concerning numerical solutions to first price auctions has devoted attention almost exclusively to the zero reserve case. In the case of distributional symmetry a zero reserve changes no qualitative results and, in fact, often simplifies exposition of key points. However, our results suggest that, to date, the literature on distributional heterogeneity may be focused on a special case (zero reserve).

Second, it is widely accepted that first price auctions are less susceptible to bidder collusion than second price auctions. Corroboration of the reversal in revenue ranking raises a tension – second price auctions are more susceptible to collusion but, with optimal reserves, they may generate higher expected revenue than first price auctions for collusive-like bidder asymmetries.

In Section II we present the model and discuss the numerical analysis. Section III contains the numerical results. A brief discussion is offered in Section IV.

2. Model and Numerical Analysis

The model and notation here parallels that in Marshall, Meurer, Richard, and Stromquist (1994, subsequently referred to as MMRS).

A single object is to be sold. We present results for both a first price and second price auction. The group of potential bidders comprise n risk neutral individuals who all draw their valuations independently from a uniform distribution on $[0, 1]$.[3] A subgroup of $k_1 < n$ bidders form a coalition. We assume this coalition acts as one bidder who draws a valuation from the cumulative distribution x^{k_1} where x 2 $[0, 1]$.

Consider next the $k_2 = n$ ¡ k_1 remaining bidders. We assume these k_2 bidders form a (counter) coalition which acts as one bidder that draws a valuation from the cumulative distribution x^{k_2} where x 2 $[0, 1]$.

The auctioneer is the sole owner of the item. He sets a reserve price, r, to maximize expected revenue. The reserve price for the second price auction is analytically derived from the auctioneer's optimization problem. The fact that a reserve price does not impact the dominance of incentive compatible strategies makes the analytics straightforward.

---

[3] The uniform distribution has become a benchmark for analysis of auction problems. Our use of it here allows immediate contrast to MMRS in which all cases were based on the uniform.

Bid functions are denoted by ϕ, appropriately subscripted (1 for the k_1-coalition, 2 for the k_2-coalition). Lebrun (1999) has shown that these bid functions are strictly monotone increasing. Inverse bid functions are denoted by λ.

An equilibrium of this model was shown to exist by Athey (2001) and, in addition, Lebrun (1999), Maskin and Riley (2000b), and Bajari (2001) have shown that an equilibrium exists and is unique. An obvious necessary condition for (λ_1, λ_2) to be a pair of Nash equilibrium strategies is that they have a common support in the form of an interval $[0, t_*]$, where t_* is the bid associated with a unit valuation. The (numerical) determination of this unique t_* is a critical component of the problem to be solved.

3. Coalition versus Coalition

3.1. The Differential Equations

Let $t = \phi_1(v)$ denote the Nash equilibrium bid submitted by coalition 1 when its highest valuation is v. Hence t is given by

$$t = Argmax(v - t)[\lambda_2(t)]^{k_2}. \tag{3.1}$$

The first-order condition generates the following differential equation:

$$k_2[\lambda_1(t) - t]\lambda_2'(t) = \lambda_2(t). \tag{3.2}$$

The corresponding equation for coalition 2 is given by

$$k_1[\lambda_2(t) - t]\lambda_1'(t) = \lambda_1(t). \tag{3.3}$$

Since the bidders are faced with a known reserve price r, the initial conditions are

$$\lambda_1(r) = \lambda_2(r) = r. \tag{3.4}$$

The terminal condition requires the existence of a number $t_* \in (0, 1)$ such that

$$\lambda_1(t_*) = \lambda_2(t_*) = 1. \tag{3.5}$$

3.2. Numerical Solution

MMRS have shown that for the special case of $r = 0$, the bid functions as $v \to 0^+$ have finite slope. Let l_i denote the (right-) derivative of λ_i at the origin,

$$l_i = \lim_{t \to 0^+} \lambda_i'(t). \tag{3.6}$$

If $\lim_{t \to 0^+} \lambda_1'(t)$ and $\lim_{t \to 0^+} \lambda_2'(t)$ both exist and are different from 1, then they are equal to l_1 and l_2 as given by:[4]

$$l_1 = 1 + \frac{1}{k_2}, \qquad l_2 = 1 + \frac{1}{k_1}. \tag{3.7}$$

This implies that a forward solution to Eqs. (2) and (3) will produce a linear solution given by

$$\lambda_i(t) = l_i \cdot t, \tag{3.8}$$

which does not satisfy the terminal condition. Therefore, MMRS suggest the use of a backward "shooting" algorithm for systems of differential equations such as Eqs. (2) and (3), and this method was used again here.

---

[4] A proof of the differentiability of $\lambda_1$and λ_2 at $t = 0$ and that they equal l_1 and l_2, respectively, can be found in Lebrun (1998b).

In the case where $r \neq 0$,

$$\lim_{t \to r^+} \lambda_i^0(t) = 1 \tag{3.9}$$

The solution to the system was found using the fourth-order Runge-Kutta method.

3.3. The Algorithm

We were unable to derive a general analytical solution for t_* and, therefore, our algorithm included an iterative search for the unique equilibrium value of t_*. With $t > t_*$, solutions to Eqs. (2) and (3), near the reserve price, showed a tendency towards zero or were no longer strictly monotone increasing. With $t < t_*$, solution to Eqs. (2) and (3) violated the initial condition in an orderly way. These properties were used to find t_* (to approximately 10 significant digits, as determined through experimentation with $r = 0$ and comparison to the analytical results that are available as described in Appendix A of MMRS). Li and Riley (1999) have also used these properties when performing their numerical analysis[5].

A single run of computation requires initializing certain parameters, evaluating the corresponding numerical solution, and then determining if another run is

---

[5] Li and Riley (1999), pages 20-21.

necessary.

1. Initialization. The parameters to be initialized are below.
    1. r: The currently conjectured value for the optimal reserve.
    2. t_*: The appropriate t_* for the conjectured reserve.
    3. N: The number of (equal length) subintervals of $(0, t_*)$ to be considered. These subintervals are of the form (t_{j-1}, t_j) with $t_0 = 0$ and $t_{N+1} = t_*$.
    4. MC: The number of Monte Carlo repetitions used in the determination of expected revenue.
    5. ε_r: A small positive number to be used in the maximizing of expected revenue. This will be discussed further below.
    6. ε_t: A small positive number to be used in the determination of t_* for each conjectured reserve.
2. Numerical Evaluation. Expected revenue as a function of the reserve is assumed to be quasi-concave.[6] This allows the method used for finding the reserve that maximizes the expected revenue to be a fairly simple "hill-climbing" method. The method chosen was the Golden Section Search

---

[6] We produced expected revenues for numerous reserve values on a grid and, in each case, found support for this conjecture.

as described in section 10.1 of Press, Flannery, Teukolsky, and Vetterling (1986).

3. Convergence Criterion

    1. Evaluation of t_* The equilibrium value of t_* was found by continually narrowing the range $(\underline{t_*}, \overline{t_*})$, starting with $\underline{t_*} = r$ and $\overline{t_*} = 1$, until the following inequality was satisfied:

    $$\overline{t_*} - \underline{t_*} \le \varepsilon_t \tag{3.10}$$

    Adjustments were made in t_* in the direction that reduced the quantity $(\lambda_1(r) - r)^2 + (\lambda_2(r) - r)^2$. Experimentation showed that values of ε_t on the order of 10^{-10} quickly provided equilibrium values of t_* that were highly accurate.

    2. Maximization of Expected Revenue. For each conjectured reserve, the bid functions were calculated and the revenue was computed by Monte Carlo. The interval $(r, 1)$ was divided into N subintervals of equal length and the bids corresponding to the separation points were computed and stored. The optimal reserve was calculated by continually

narrowing the range $(\underline{r}, \overline{r})$, starting with $\underline{r} = 0$ and $\overline{r} = 1$, until the following inequality was satisfied:

$$\overline{r} - \underline{r} \leq \varepsilon_r \tag{3.11}$$

Since maximization was being performed on a function that was evaluated by approximation, great care had to be taken to maintain as much accuracy as possible. The following values of (N, MC, ε_r) were chosen as a compromise between computing time and accuracy:

$N = 20{,}000$

$MC = 1{,}000{,}000$

ε_r on the order of 10^{-7}.

4. Comparison to Second Price

The results are presented in Table 1 and Table 2 while the bid functions for the ($k_1 = 1$, $k_2 = 10$) and ($k_1 = 4$, $k_2 = 7$) cases are provided in Figures 1 and 2.[7] We have included some of Tables III and IV from MMRS for comparison purposes

---

[7] Although Figure 1 seems to indicate a near linear bid function for the coalition of 1, in fact at the reserve price the slope of the bid function is zero.

(zero reserve price entries).

TABLE 1

Auctioneer's Optimal Reserve, Expected Revenue, Standard Error,

and Bidders' Expected Surplus (Per Capita) - n = 5

	k_1	k_2	Res.	Rev.	Std. Err.	ES(k_1)	ES(k_2)
Second	1	4	.65387	.61315	.00024	.02240	.03317
Price	1	4	0	.4667		.0333	.0833
	2	3	.61308	.63360	.00022	.02841	.03292
	2	3	0	.5833		.0417	.0556
First	1	4	.66020	.60867	.00023	.03646	.02876
Price	1	4	0	.5057		.0860	.0567
	2	3	.61581	.63329	.00020	.03217	.02993
	2	3	0	.5875		.0523	.0467

Based on 1 million drawings.

TABLE 2

Auctioneer's Optimal Reserve, Expected Revenue, Standard Error, and Bidders' Expected Surplus (Per Capita) - n = 11

	k_1	k_2	Res.	Rev.	Std. Err.	ES(k_1)	ES(k_2)
Second	1	10	.78181	.74774	.00020	.00584	.01154
Price	2	9	.76392	.76266	.00019	.00657	.01164
	3	8	.74464	.77557	.00017	.00744	.01154
	4	7	.72606	.78507	.00016	.00844	.01116
	5	6	.71338	.79047	.00015	.00950	.01046
First	1	10	.78557	.74158	.00020	.01662	.01029
Price	2	9	.76117	.75896	.00017	.01415	.01003
	3	8	.74357	.77448	.00016	.01212	.00963
	4	7	.71988	.78506	.00014	.01114	.00974
	5	6	.71184	.79045	.00013	.01031	.00985

Based on 1 million drawings

For each case, the mean difference in revenue between auction types and the associated standard error are shown in Table 3 below.

TABLE 3

Auctioneer's Expected Revenue at First and Second Price Auctions, Mean Revenue Differential, and Standard Error

k_1	k_2	SP Rev.	FP Rev.	Mean Diff.
1	4	.61315	.60867	.00448
				(.00008)
2	3	.63360	.63329	.00031
				(.00009)
1	10	.74774	.74158	.00616
				(.00006)
2	9	.76266	.75896	.00370
				(.00006)
3	8	.77557	.77448	.00108
				(.00007)
4	7	.78507	.78506	.00001
				(.00008)
5	6	.79047	.79045	.00003
				(.00008)

Based on 1 million drawings.

Several points are worthy of note.[8]

1. Table 3, Column “Mean Diff.”

   When using an optimal reserve price, for a given k_1 and k_2, the expected revenue with a second price auction is strictly larger than with a first price auction.

2. Tables 1 and 2, Column “Res.”

   1. For a given k_1 and k_2 the difference in the optimal reserve between first price and second price auctions appears to change sign as k_1 and k_2 change.
   2. Revenue differentials between first and second price auctions shrink when optimal reserves are used relative to a reserve of zero, and become more insignificant as the asymmetry between the coalitions (i.e., the difference between k_1 and k_2) decreases.[9]

3. Tables 1 and 2, Column “ES(k_1)” and “ES(k_2)”

---

[8] In the two cases studied here $k_1 < k_2$. One can think of k_1 as the "weaker" type, or, smaller coalition (sometimes a singleton) while k_2 is the "stronger" type, or, larger coalition.

[9] Small revenue differentials were also found in Li and Riley (1999) when optimal reserve prices were considered, as discussed on pages 27-29. Their observation was made when using the optimal second price reserve when calculating expected revenue for both the first and second price auctions.

1. As expected, the use of an optimal reserve price reduces the per capita payoff to both k_1 and k_2 types (relative to no reserve).

2. For the first price auction the per capita payoff gain from being in k_1 rather than k_2 dramatically shrinks with an optimal reserve as opposed to the absence of a reserve. (In the $k_1 = 1$ and $k_2 = 4$ case for the first price auction the percentage gain to being the k_1 type is 27% with an optimal reserve while it is 52% without a reserve.) For the second price auction the optimal reserve dramatically reduces the relative gain from being a k_2 type.

3. If we interpret the k_1 and k_2 types as opposing coalitions, then the reserve price is changing the benefit to defection from a coalition at first price auctions. For the $n = 5$ case, a member of the coalition $k_2 = 4$ who is considering defecting to coalition $k_1 = 1$ would find such a move profitable with the optimal reserve ($.03217 > .02876$) but not profitable with a zero reserve ($.0523 < .0567$).

5. Discussion

The remarkable finding is that, with this particular bidder asymmetry, the use of an optimal reserve seems to provide the seller with higher expected revenue from a second price auction than a first price auction. It was our prior belief that this would not prove to be true. At this time we have no intuition for the result.[10]

There seem to be two very important implications of this finding. First, the previous analysis in the literature of asymmetric bidders has focused on a special case – zero reserves. This implies that with bidder asymmetry, conclusions about different auctions may not in general carry over directly from the case of zero reserve prices to positive reserves, unlike with symmetric bidders. This suggests that previous results should be carefully re-examined for nonzero reserve prices. Second, the revenue superiority of second price auctions when optimal reserves are used raises the possibility that an auctioneer may prefer a second price auction to a first price auction when bidders are known to be colluding.

---

[10] It is intriguing to note that for the zero reserve case the bid functions have positive slope at values of zero but for any strictly positive reserve the bid functions have zero slope at values equal to the reserve.

References

Athey, Susan, 2002, “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information,” Econometrica, 69, pp. 861-890.

Bajari, Patrick, 2001, "Comparing Competition and Collusion in Procurement Auctions: A Numerical Approach", Economic Theory, 18, pp. 187-205.

Lebrun, Bernard, 1996, ”Existence of an Equilibrium in First Price Auctions”, Economic Theory, 7, pp. 421-443.

_________, 1998a, ”Comparative Statics in First Price Auctions”, Games and Economic Behavior, 25, pp. 97-110.

_________, 1998b, ”The First and Second Price Auctions in a Class of Asymmetric Examples”, mimeo.

_________, 1999, ”First Price Auctions in the Asymmetric N Bidder Case”, International Economic Review, 40, 1, pp. 125-142.

Li, Huagang, and John G. Riley, 1999, "Auction Choice", UCLA Working Paper.

Marshall, Robert C., Michael J. Meurer, Jean-Francois Richard, and Walter Stromquist, 1994, ”Numerical Analysis of Asymmetric First Price Auctions”, Games and Economic Behavior, 7, pp. 193-220.

Maskin, Eric and John Riley, 2000a, ”Asymmetric Auctions”, Review of Economic Studies, 67, pp. 413-438.

_____________, 2000b, ”Equilibrium in Sealed High Bid Auctions”, Review of Economic Studies, 67, pp. 439-454.

Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling, 1986, Numerical Recipes, Cambridge University Press: New York.

Figure 1: Coalition of 1 versus Coalition of 10

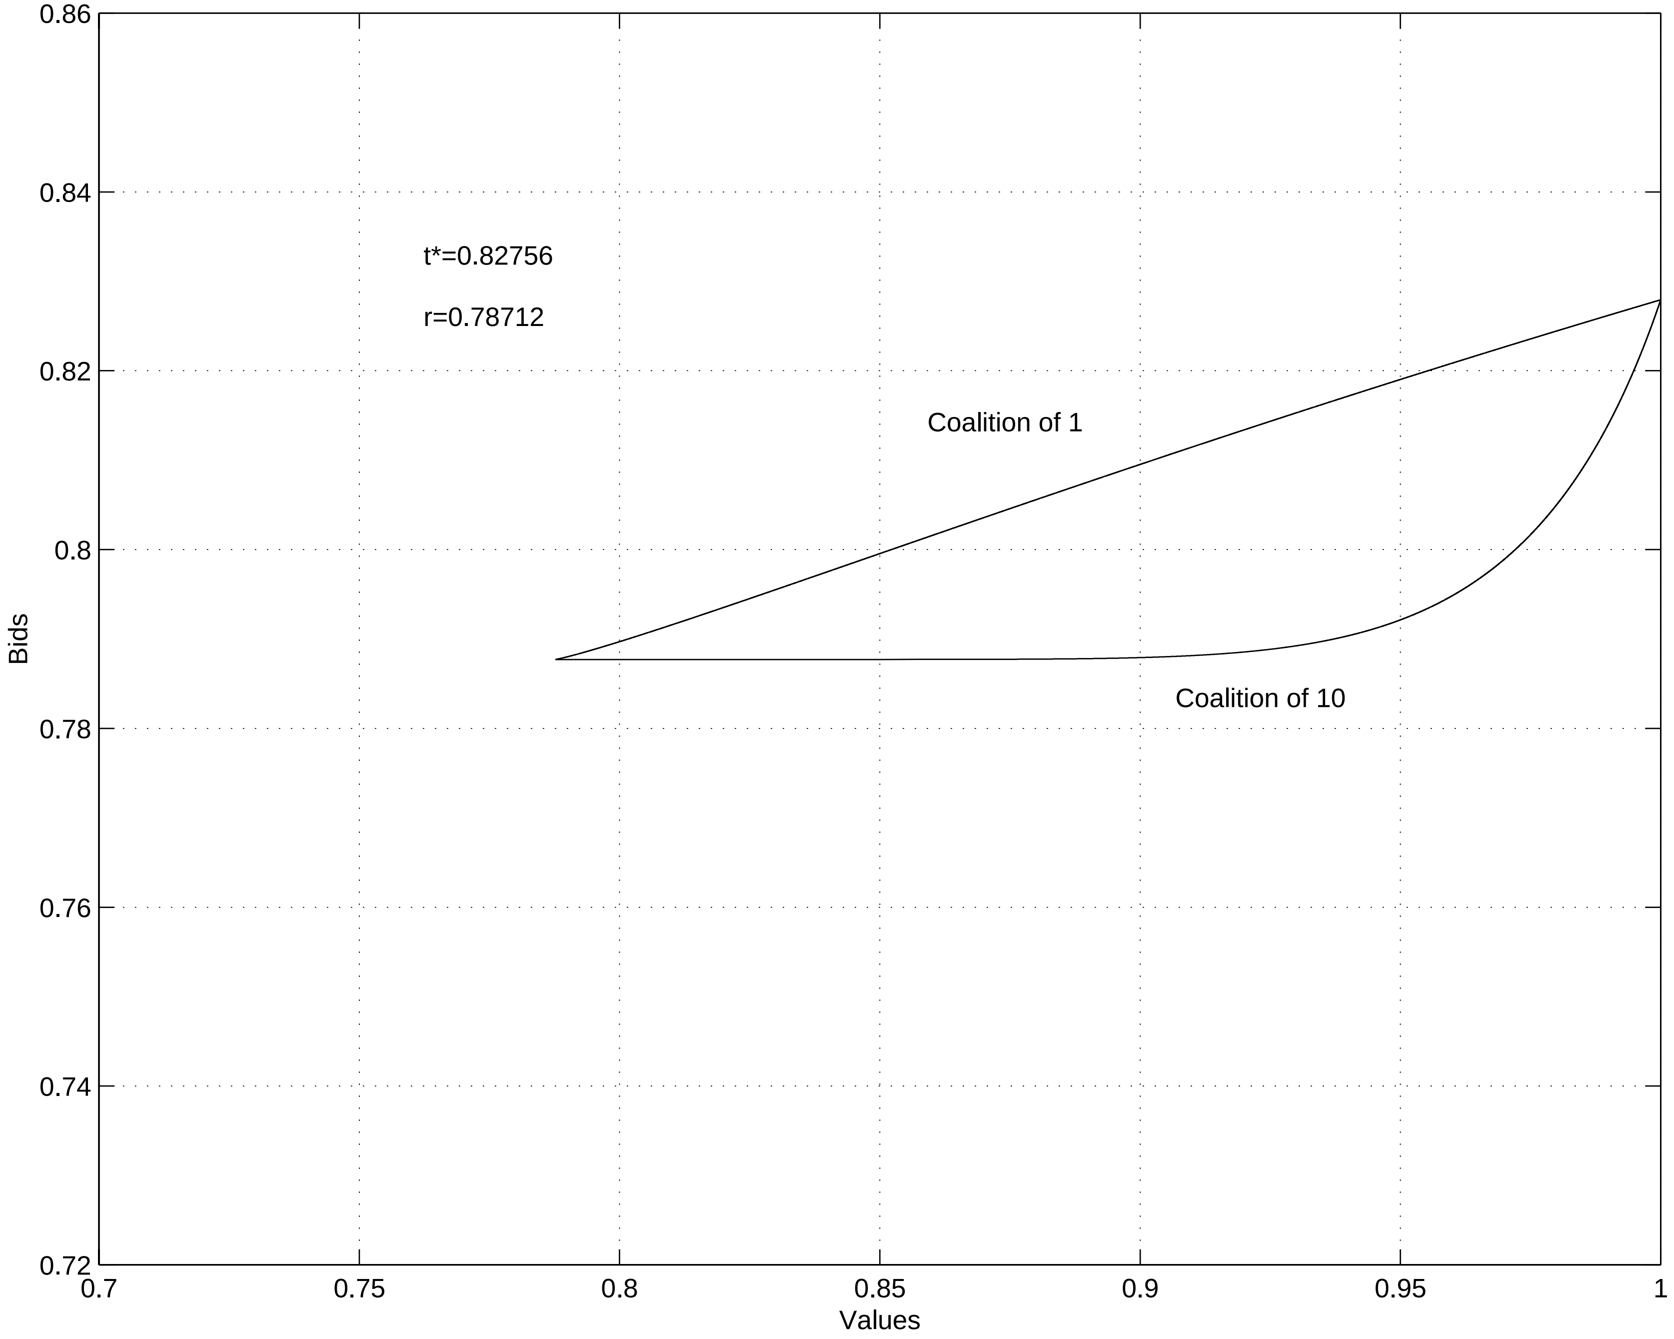

Figure 2: Coalition of 4 versus Coalition of 7

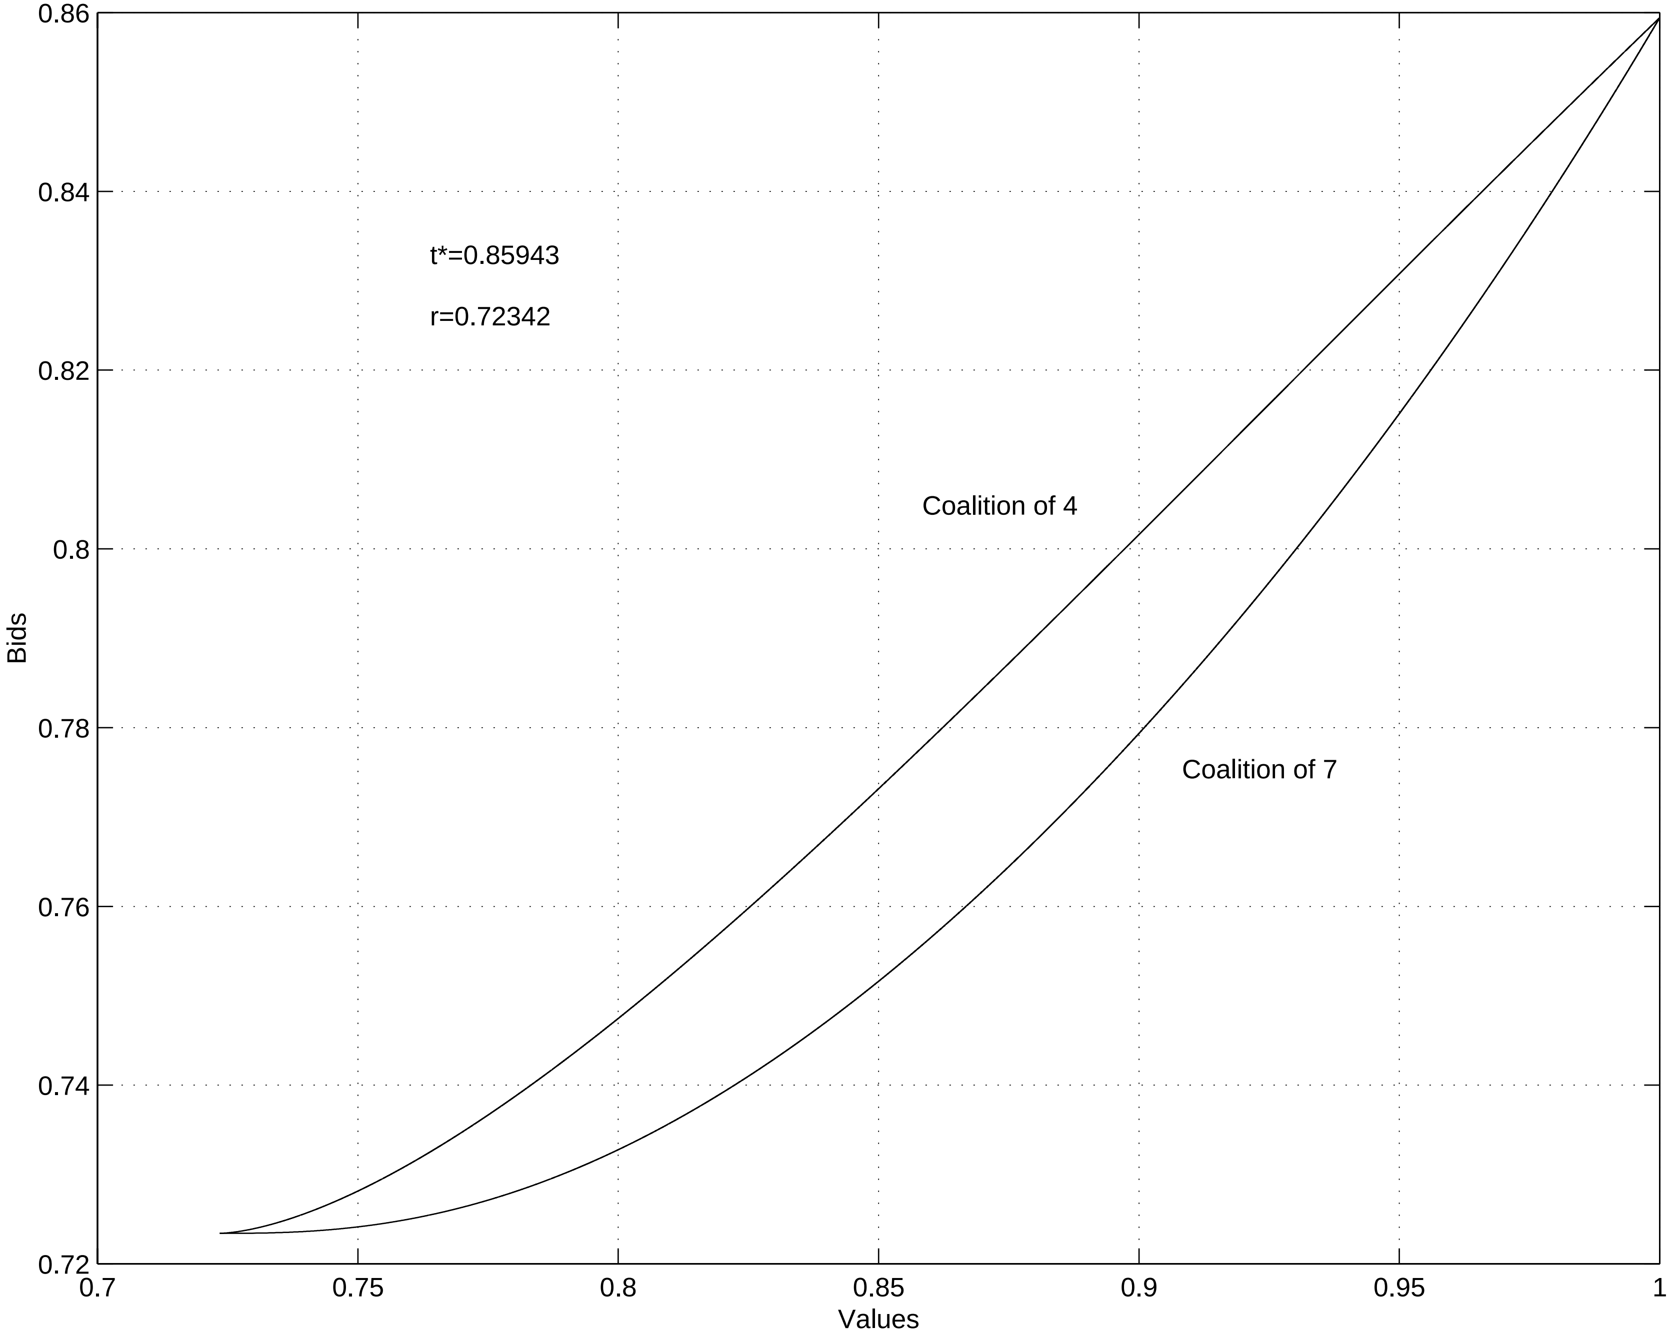

Chapter 3

Optimal Reserve Prices at First and Second Price Auctions with Endogenous Coalition Formation

Steven P. Schulenberg

1. Introduction

It is commonly believed that first price auctions are less susceptible to collusion than first price auctions. The intuition behind this is clear. Consider a second price auction, with zero reserve price for simplicity, where the only two participating bidders are colluding. One bidder should submit a high bid while the second bidder submits a bid of zero, thus allowing the cartel to win with a price paid of zero. The second bidder will not have incentive to deviate by bidding higher than the first bidder if the first bidder's bid is sufficiently high. Now consider a first price auction. One bidder should submit a very low bid while the second bidder again submits a bid of zero. Now, however, the second bidder can deviate and win the object with a low price paid. Therefore, the incentive to cheat on the collusive agreement is larger at a first price auction than a second price auction. Since collusion will reduce the price paid, this suggests that an auctioneer facing colluding bidders should choose a first price auction rather than a second price auction.

If a first price auction and a second price auction have different expected revenues, then the assumptions of the revenue equivalence theorem (RET), the benchmark result of the theory of auctions, cannot be true. The RET is derived within the independent private values (IPV) framework. One of the underlying assumptions of the IPV model is that all bidders draw valuations from the same common underlying distribution. If this assumption is violated then the RET no longer holds, and first price and second price auctions will yield different expected revenue for the seller if bidders draw values from different distributions (excepting pathological cases).

If bidders are ex ante symmetric then collusion by any strict subset will produce bidder asymmetry. Given this, it is natural to focus attention on a kind of asymmetry that might arise from collusion by symmetric bidders. Lebrun (1996, 1998a, 1998b, 1999) and Maskin and Riley (2000a, 2000b) have obtained numerous analytic results for the case of bidder value asymmetry. For a particular kind of bidder asymmetry that might reasonably arise from collusion by bidders whose valuations are distributed continuously uniformly, it is well known that first price auctions generate higher expected revenue for the seller than do second price auctions in the absence of reserve prices. Marshall, Meurer, Richard and Stromquist (1994, subsequently referred to as MMRS) provided numerical analysis for collusive-like asymmetries. Lebrun (1998b) subsequently provided the general analytic results.

These results are dependent upon the auctioneer not using a reserve price (i.e., using a reserve price of zero.) However, by using an appropriate reserve price, an auctioneer can significantly increase the expected revenue at any auction. Marshall and Schulenberg (2003, subsequently referred to as MS) found that, when using expected revenue maximizing reserve prices, second price auctions typically generate higher expected revenues than first price auctions with the type of asymmetry considered in MMRS. This suggests that, in contrast to the suggestion of the first paragraph, auctioneers may face bidder collusion and yet prefer a second price auction over a first price auction. In this analysis, the coalitions of bidders were assumed to have formed exogenously, allowing the auctioneer to have no impact on the decision of the bidders to collude or not.

The main goal of this paper is to relax this assumption and investigate the role that the auctioneer can have on the formation of coalitions. We consider bidders who must decide what coalitions they wish to form after learning of the reserve price. This allows the auctioneer to affect the degree of collusion, since increasing the reserve price acts to decrease the expected surplus of a coalition. We find that when there are enough bidders, the auctioneer may be able to raise the reserve above the value considered by MMRS and force smaller coalitions to form. If the reserve is chosen correctly, the expected revenue will increase as a result. This may again cause first price auctions to perform better than second price auctions in the presence of bidder collusion.

An auction is considered in three stages. First, the auctioneer announces an auction type (first or second price) and a reserve price. Bidders then play a cooperative game in the second stage to form any coalitions. Finally, the auction occurs. The auctioneer announces a reserve price, the bidders learn their values, the bidders submit their bids, and then the object is exchanged for the appropriate payment. The solution concept employed for the first step is that of the recursive-core, as presented in Huang and Sjostrom (2001). This involves finding the characteristic function that gives the worth of each coalition. An auction necessarily involved externalities, where the worth of any coalition depends on how the remaining players outside the coalition choose to behave. Huang and Sjostrom (2001) makes the very reasonable suggestion that the same solution concept used to solve the game, such as the core, is also used to predict the behavior of the players outside any given coalition when determining the worth of that coalition.

In this analysis, the issues of collusive mechanisms and incentives to cheat on the coalition at the auction are ignored. Instead, it is assumed that there is some suitable enforcement strategy so that once an individual commits to a coalition, he will truthfully report her value and bid according to the coalition's recommendation.[1] Then the only two concerns regarding an individual's inclusion in a coalition are if the coalition can offer the individual more than any other coalition, and if the coalition gains from the individual's inclusion net of the amount offered. If these conditions are not met, than there is no mechanism that will cause the individual to join to coalition. The analysis here is then looking at the maximum possible size of the coalitions, and in a sense the maximum possible amount of collusion.

It is also assumed that the bidders must decide what coalition (if any) to join prior to learning their values. There are situations where this is not a desirable assumption. However, this reduces the complexity of the analysis enormously. Without this assumption, not only would all of the possible coalitions have to

---

[1] Huang and Sjostrom (2001) assumes that "utility is transferable and binding contracts can be written within coalitions." They further assume that a coalition may be viewed as a "composite player" who is solely concerned with maximizing the joint payoff of the members of that coalition. These assumptions follow Ichiishi (1991) and Ray and Vohro (1997).

be considered, but also for each all of the (infinite) sets of possible values should be considered. Given the extreme asymmetry, this could almost assuredly not be done analytically, meaning that a subset of the possible values would have to be considered numerically and the remainder interpolated in some fashion. Hence without this assumption, any analysis would be extremely cumbersome and at best incomplete.

The paper proceeds as follows. Section II contains a brief overview of the recursive-core as defined in Huang and Sjostrom (2001). In Section III, the model is presented. Section IV discusses the third stage of the game (the auction) as well as the numerical analysis. The first and second stages, the choice of auction type and reserve price followed by coalition formation, are then presented in Section V. Results to the complete game and a brief discussion are offered in Section VI.

2. Brief Overview of the Recursive-Core

Huang and Sjostrom (2001) considers the problem of assigning values to each coalition in a cooperative game when the values depend on the actions of other coalitions. The members of a (possible) coalition must account for these externalities by predicting the behavior of each other coalition before concluding the magnitude of their value. This is clearly a relevant concern for auctions. The bidding strategy of and hence payoff to any coalition of bidders will be determined by factors such as the number of other bidders, the value distributions of these other bidders, how these other bidders have arranged themselves into coalitions, and the size of the reserve price.[2]

For example, consider an auction with four bidders (A, B, C, and D) initially formed into the all-inclusive cartel. If bidders A, B, and C considered a joint defection, then they do so knowing that bidder D must then act as an individual regardless of the solution concept in consideration. Suppose bidders A and B considered a joint defection. The remaining two bidders C and D may stay together, resulting in a coalition of two against a coalition of two, or they may instead choose to act individually, resulting in a coalition of two against two individual bidders. To the extent that the payoff to the (potentially) defecting coalition differs between outcomes, bidders A and B must predict the choice of bidders C and D when making a proper decision to defect or not. A similar, but slightly more complex, problem confronts a single bidder considering defection.

Any coalition smaller than the all-inclusive cartel must use some solution concept, such as the core, to determine its behavior. Huang and Sjostrom (2001)

---

[2] As will be discussed in more detail below, these factors have a greater impact at first price auctions than second price auctions.

make the natural assumption that, when determining its "worth", the coalition assumes the remaining subset of players will behave as predicted by the same solution concept. They provide several examples and illustrate how their theory differs from the most similar alternatives.

3. Model

The model and notation here parallels that in MS, which extends the model in MMRS by considering a non-zero reserve price.

A single object is to be sold. We present results for both a first price and second price auction. The auctioneer is the sole owner of the object. He sets a reserve price, r, to maximize expected revenue.

The group of potential bidders is composed of N risk neutral individuals who all draw their valuations independently from a uniform distribution on $[0,1]$.[3] These bidders form M disjoint coalitions, with n_i bidders in coalition i. Clearly, $\sum_{i=1}^{M} n_i = N$. We assume each coalition acts as one bidder who draws a valuation from the cumulative distribution x^{n_i} where $x \in [0,1]$.

Bid functions are denoted by ϕ, appropriately subscripted. Lebrun (1997) has shown that these bid functions are strictly monotone increasing. Inverse bid

---

[3] The uniform distribution has become a benchmark for analysis of auction problems.

functions are denoted by λ. An obvious necessary condition for $(\lambda_1, ..., \lambda_M)$ to be a set of Nash equilibrium strategies is that they have a common support in the form of an interval $[r, t_¤]$, where $t_¤$ is the bid associated with a unit valuation. The (numerical) determination of $t_¤$ is a critical component of the problem to be solved.

There are three stages of the game played by the auctioneer and bidders. In stage one, the auctioneer announces the auction type and reserve price. In stage two, the bidders arrange themselves into coalitions. In stage three, the bidders (and hence coalitions) observe their values, and the auction occurs.[4]

4. Stage Three: Auction Subgame

Consider the (generally) asymmetric auctions that will occur in stage three. The solution algorithm is nearly identical to that used in MS, and a more detailed discussion can be found in that paper. The notable difference is that the focus in that paper was on the single reserve which resulted in the largest expected revenue for the seller at each auction type. Here, we are concerned with the entire range of potential reserves, and the resulting expected revenues and surpluses. As will

---

[4] The game considered in MS differs by reversing the order of stages 1 and 2, and MS also begins analysis after the coalitions have been formed.

be seen in later sections, the optimal reserve as determined in MS will typically not be the optimal reserve of the larger game considered here.

The third stage begins with the auction type, reserve price, and coalitions all known. Given these, we can find the expected revenue and expected surplus for both a first price and second price auction. Since second price auctions are solvable analytically, these are discussed first.

4.1. Second Price

Given the assumptions on coalition behavior, a coalition can be thought of as a single bidder with the appropriate distribution. Hence, it is a dominant strategy for the coalition to bid its value. Then given the coalitions and the auctioneer's choice of a reserve price, the expected revenue to the auctioneer can be written as

$$ER = \frac{1}{N+1}\left(1 - r^{N+1}\right)\sum_{i=1}\left[\frac{n_i(N-n_i)}{N-n_i+1}\right] + \sum_{i=1}\left(\frac{1}{N-n_i+1}r^{N-n_i+1}(1-r^{n_i})\right). \tag{4.1}$$

For example, with $N = 2$, $r = \frac{1}{2}$ and $n_i = 1$ for all i, this gives the expected revenue as $\frac{5}{12}$.

Note that by maximizing this with respect to the reserve, an expression for the optimal reserve given the coalition structure can be found. This reserve is the value of $r \in [0,1]$ that makes the following equation hold:

$$\sum_{i=1} \left(r^{-n_i} - n_i - 1 \right) = 0 \tag{4.2}$$

Note that for $n_i = 1$ for all i, the left side reduces to $N(r^{-1} - 1 - 1) = N\left(\frac{1}{r} - 2\right)$ giving the well known result that the optimal reserve is equal to $\frac{1}{2}$.

The expression for the expected surplus for coalition i is given by

$$ES_i = \frac{1}{N+1}\frac{n_i}{N - n_i + 1} - \frac{1}{N - n_i + 1} r^{N - n_i + 1} + \frac{1}{N+1} r^{N+1}. \tag{4.3}$$

A convenient aspect of this is that the expected surplus to a coalition only depends on the number of bidders overall and in that particular coalition. The arrangement of the remaining bidders in coalitions is irrelevant. For example, with $N = 2$, $r = \frac{1}{2}$ and $n_i = 1$, the expected surplus for coalition (i.e., bidder) i is $\frac{1}{12}$.

4.2. First Price

First price auctions in this environment are generally asymmetric and not analytically solvable. This makes the analysis much more difficult, since it requires all evaluations to be done numerically. The exceptions to this are special cases when

$n_i = n_j$ for all i,j.[5] In such cases, revenue equivalence holds and hence the second price analytical results are applicable for the first price case.[6] For all other cases, the expected revenue and surplus was calculated as described below.

4.2.1. The Differential Equations

Let $t = \phi_i(v_i)$ denote the Nash equilibrium bid submitted by coalition i when its highest valuation is v_i. Hence t is given by

$$t = Arg\max\,(v_i - t)\prod_{j\neq i}\lambda_j(t)^{n_j}$$

The first order condition, imposing symmetry so that $v_i = \lambda_i(t)$, is then

$$(\lambda_i(t) - t)\left[\sum_{j\neq i} n_j\lambda_j'(t)\lambda_j^{n_j-1}(t)\prod_{k\neq i,j}\lambda_k^{n_k}(t)\right] - \prod_{j\neq i}\lambda_i^{n_j}(t) = 0.$$

---

[5] Note that this includes the case of an all-inclusive cartel (i.e., $i = j = 1$), and the case of each coalition with only a single member (i.e., $n_i = n_j = 1$).

[6] The conditions that result in expected revenue equivalence also result in expected surplus equivalence. For example, at either a first or a second price auction, an all-inclusive cartel will either bid zero (if its value is less than the reserve) or the reserve (if its value is at least the reserve). Since the resulting payments are also equivalent in each case (i.e., either zero or the reserve), the expressions derived for expected surplus and expected revenue are identical in each case.

Simplifying and rearranging gives rise to the following differential equations:

$$\lambda_i'(t) = \frac{\lambda_i(t)}{(M-1)\,n_i}\left[\sum_{j\neq i}\left(\frac{1}{\lambda_j(t)-t}\right) - \frac{M-2}{\lambda_i(t)-t}\right], \; i \in \{1,...,M\}. \tag{4.4}$$

Note that with $M = 2$, this reduces to

$$\lambda_1'(t) = \frac{\lambda_1(t)}{n_1\,(\lambda_2(t)-t)}, \; \lambda_2'(t) = \frac{\lambda_2(t)}{n_2\,(\lambda_1(t)-t)}$$

which corresponds to the system of differential equations in MMRS.

Since the bidders are faced with a known reserve price r, the initial conditions are

$$\lambda_i(r) = r, \; i \in \{1,...,M\}. \tag{4.5}$$

The terminal condition requires the existence of a number $t_* \in [0,1]$ such that

$$\lambda_i(t_*) = 1, \; i \in \{1,...,M\}. \tag{4.6}$$

An equilibrium of this model was shown to exist by Athey (2001), and Lebrun (1999), Maskin and Riley (2000b), and Bajari (2001) have shown that an equilibrium exists and is unique.

4.2.2. Numerical Solution and Algorithm

For each considered set of coalitions and reserve price, we had to determine the (unique) value of t_*, numerically solve the bid function, and then calculate the expected revenue and surpluses. To find the optimal reserve price, we searched over all possible reserve prices and simply chose the reserve that resulted in the highest expected revenue.

MMRS have shown that for the special case of $r = 0$, the bid functions as $v \to 0^+$ have finite slope. Let l_i denote the (right-) derivative of λ_i at the origin

$$l_i = \lim_{t \to 0^+} \lambda_i'(t).$$

This implies that a forward solution to the system of differential equations will produce a linear solution given by

$$\lambda_i(t) = l_i \cdot t,$$

which does not satisfy the terminal condition. Therefore, MMRS suggest the use of a backward "shooting" algorithm for systems of differential equations such as above, and this method was used again here.[7]

The solution to the system was approximated using the fourth-order Runge-Kutta method. This proved to be accurate as long as the number of bidders in each coalition were not large.[8] The relevant value range $[r, t^*]$ was divided into $20,000$ subintervals of equal length. These subintervals are of the form $[t_{j-1}, t_j]$ with $t_0 = 0$ and $t_{N+1} = t^*$. Starting at the (conjectured) value of t^*, the bid functions were approximated at the ends of each subinterval.

We were unable to derive a general analytical solution for t^* and, therefore, our algorithm included an iterative search for the equilibrium value of t^*. Beginning with $[\underline{t^*}, \overline{t^*}] = [r, 1]$, the range $[\underline{t^*}, \overline{t^*}]$ was continually narrowed until $\overline{t^*} - \underline{t^*} \leq \varepsilon_t$. Solutions to the differential equations showed a tendency towards $-\infty$ for values near the reserve value or were no longer strictly monotone increasing when a guessed value of t^* was too high. This, along with straightforward comparison to the initial conditions, made it clear whether a conjectured t^* was too high or too low. The value of ε_t used was on the order of 10^{-10}, and this quickly provided equilibrium values of t^* to approximately 10 significant digits.[9]

---

[7] In the case where $r \neq 0$, $\lim_{t \to r^+} \lambda_i^0(t) = \infty$.

[8] For instance, the accuracy was very poor with $M = 2$, $n_1 = 1$, and $n_2 = 100$.

[9] This was determined through experimentation with $r = 0$ and comparison to the analytical results that are available as described in Appendix A of MMRS.

The expected revenue and surpluses, given a set of coalitions and reserve price, were computed by Monte Carlo. For each Monte Carlo repetition, values for each coalition were drawn. The corresponding bids (if the values were above the reserve price) were found from the approximated bid functions, using linear interpolation within a subinterval. The winning bid (if any) was determined, and then the revenue and surpluses were found and recorded. Expected revenue and surplus were calculated as the mean over all of the Monte Carlo repetitions. One million repetitions were used in each case.

5. Stages One and Two: Coalition Formation and Optimal Reserve Price

Using the results obtained as above for each relevant auction, the coalitions that will result from any announced reserve can be determined. The important comparisons for this are then between the expected per capita surplus (abbreviated as PCES) of the bidders in the different possible coalitions. Consider the potential defection of a bidder from a coalition (A) to another coalition (B).[10] The most that coalition B is willing to offer to a bidder in coalition A is the increase in

---

[10] Coalition B may consist of only that defecting bidder.

expected surplus to coalition B if the bidder switched. If such a defection is profitable, then it must be that the increase in the expected surplus is larger than the PCES earned by a bidder in coalition A. If so, there is no possible division of the surplus for coalition A that would remove the incentive for all bidders to defect. Therefore, a bidder will wish to defect to a coalition and this coalition will wish to accept this bidder only if the increase in expected (total) surplus to the coalition due to the addition of the bidder is greater than the PCES the bidder earned in his original coalition. The relevant comparison is then of the PCES that results from each potential set of coalitions.

The expected surplus to each possible coalition for different numbers of bidders was calculated using the methods described above. This allowed the following question to be answered: given a reserve price, what coalitions of bidders will result?

5.1. Second Price

5.1.1. Stage Two: Coalition Formation

Graham and Marshall (1987) and Mailith and Zemsky (1991) have shown that, for second price auctions, the PCES increases as coalitions get larger for any reserve price. This is straightforward to see in our environment. From the above equation for expected surplus, the PCES is given by

$$\begin{aligned} PCES_i &= \frac{1}{n_i} ES_i \\ &= \frac{1}{N+1}\frac{1}{N-n_i+1} - \frac{1}{n_i}\frac{1}{N-n_i+1} r^{N-n_i+1} + \frac{1}{n_i}\frac{1}{N+1} r^{N+1} \end{aligned}$$

Then

$$\begin{aligned} \frac{\partial PCES_i}{\partial n_i} &= \frac{1}{N+1}\frac{1}{(N-n_i+1)^2} + \frac{1}{n_i^2 (N-n_i+1)} r^{N-n_i+1} \\ &\quad - \frac{1}{n_i (N-n_i+1)^2} r^{N-n_i+1} \\ &\quad + \frac{1}{n_i (N-n_i+1)} r^{N-n_i+1} \ln r - \frac{1}{n_i^2 (N+1)} r^{N+1} \end{aligned}$$

Note that

$$\begin{aligned} \frac{\partial \left(\frac{\partial PCES_i}{\partial n_i} \right)}{\partial r} &= \frac{1}{n_i^2}\frac{r^{N-n_i+1}}{r} + \frac{1}{n_i}\frac{r^{N-n_i+1}}{r} \ln r - \frac{1}{n_i^2}\frac{r^{N+1}}{r} \\ &= \frac{1}{n_i^2} r^N \left[r^{-n_i} - 1 \right] + \frac{1}{n_i} r^{N-n_i} \ln r \end{aligned}$$

Since $r - n_i > 1$, $\frac{\partial \left(\frac{\partial PCES_i}{\partial n_i} \right)}{\partial r} > 0$. This implies that $\frac{\partial PCES_i}{\partial n_i}$ is smallest when the reserve is smallest (i.e. $r = 0$). Substituting $r = 0$ in $\frac{\partial PCES_i}{\partial n_i}$ gives

$$\frac{\partial PCES_i}{\partial n_i} \Big|_{r=0} = \frac{1}{N+1} \frac{1}{(N - n_i + 1)^2} > 0.$$

Hence, the PCES is increasing as the coalition grows larger given any reserve price.

This immediately implies that the bidders always have incentive to form the all-inclusive cartel, regardless of the chosen reserve price, and hence this is the unique r-core. Recall from above that the expected surplus of a coalition depends only on the number of bidders in that coalition and the total number of bidders in other coalitions, and does not depend on how these others bidders are split into coalitions. There is no need for the coalition in question to predict the behavior of the other bidders, since there are no externalities that will result, and the r-core is then equivalent to the core.

5.1.2. Stage One: Optimal Reserve Price

Since the auctioneer can always expect to face an all-inclusive cartel, the best response by the auctioneer is then to choose a reserve that maximizes revenue when facing such a coalition. Recall from above that the optimal reserve satisfies

$$\sum_{i=1}^{M} \left[r^{-n_i} - n_i - 1 \right] = 0.$$

For the all-inclusive cartel, $M = 1$ and $n_1 = N$. Hence, the optimal reserve at a second price auction in given by

$$r^{*}_{SP} = (N + 1)^{-\frac{1}{N}}.$$

Substituting this into the expression for expected revenue gives

$$ER^{*}_{SP} = \frac{N}{N + 1} (N + 1)^{-\frac{1}{N}}.$$

Similarly, expected surplus is given by

$$ES^{*}_{SP} = \frac{N}{N + 1} - (N + 1)^{-\frac{1}{N}} + (N + 1)^{-\frac{2N+1}{N}}.$$

Figure 1 shows, for various numbers of bidders, the optimal reserve price and also the expected revenue and expected surplus at this reserve. For example, with five bidders the expected revenue is 0.582.

5.2. First Price

Unlike the second price auctions, first price auctions do have externalities. That is, a coalition of three bidders will bid differently and have a different expected surplus when bidding against either a coalition of two bidders or two bidders acting individually. The r-core will therefore not necessarily be identical to the core.

5.2.1. Simple Example: Three Bidders

To illustrate the r-core in this context, consider a first price auction with three bidders. There are three possible coalition structures: (a) the all-inclusive coalition, (b) two coalitions, one with a single bidder and another with two bidders, and (c) three coalitions each with a single bidder.[11] The expected revenue and surpluses are simple to calculate in the first and last cases, while the middle case requires numerical determination.

First, suppose that the auctioneer has chosen a reserve price of zero. The expected revenues and surpluses are shown in the table below.

---

[11] There are actually three ways that the case of two coalitions can result, since each of the three bidders could be the single bidder coalition. Since the bidders are identical, these situations are identical except for the names of the bidders. To avoid unnecessary complication in the exposition, these three cases will not be discussed separately. This simplification will be similarly applied for the remainder of the paper.

Table 1: First Price Auction with Three Bidders

Coalition	Competing Coalitions	Exp. Rev.	Exp. Surp.	Per Capita Exp. Surp.
3	None	0.000	0.750	0.250
2	1	0.417	0.250	0.125
1	2	0.417	0.083	0.083
1	1,1	0.500	0.083	0.083

Consider the all-inclusive cartel. The per capita expected surplus is 0.25. Suppose that one of the bidders was considering a defection. In order to know whether this was profitable, this bidder would have to make a prediction about the behavior of the remaining two bidders given a defection. The remaining bidders would have a per capita expected surplus of 0.125 if they joined into a single coalition, and only 0.083 each if they bid separately. Hence, the core predicts that these bidders would form a single coalition. The potential defector then would then have an expected surplus of 0.083 if acting alone, compared to a per capita surplus of 0.25 as part of the all-inclusive cartel. Hence, a single bidder would not choose to defect.[12] Similarly, a coalition of two bidders would not defect, since the per capita surplus after defecting is 0.125 compared with 0.25 as part of the

---

[12] This conclusion is the same regardless of the behavior of the two remaining bidders, since the surplus of the defecting bidder is the same in each possibility. However, this is not generally true so the discussion here reflects the standard argument that would be necessary.

all-inclusive cartel. Therefore, the unique element of the r-core is the all-inclusive cartel.

This result depends only on the following facts: (a) the PCES is greater for a coalition of two facing a single other bidder than for a single bidder facing two other bidders acting separately, (b) the PCES is greater for the all-inclusive cartel than for a single bidder facing a coalition of two bidders, and (c) the PCES is greater for the all-inclusive cartel than for a coalition of two bidders facing a single bidder. In this example, these facts are all true for every reserve price.[13] The unique r-core is then the all-inclusive cartel, regardless of the reserve price. The optimal reserve price chosen by the auctioneer will then be identical to that for a second price auction, which is 0.63.[14] The auctioneer is therefore indifferent to the auction type with three bidders.

5.2.2. Five Bidders

It is more generally true that a sufficient condition for the unique element of the r-core to be the all-inclusive cartel is that the PCES for the all-inclusive cartel is greater than that of each smaller coalition in all possible coalition structures.

---

[13] This not strictly correct at $r = 1$, where the expected surplus is always zero. However, the expected revenue is also zero, so this reserve is never an optimal choice for an auctioneer and hence is omitted from the analysis.

[14] See Figure 1.

This condition holds for second price auctions generally, as discussed above, and also for first price auctions with three bidders. One immediate corollary to this is that for the case of a first price auction with a zero reserve price, the r-core consists only of the all-inclusive cartel. However, with sufficiently many bidders, this condition does not hold for all reserve prices.

Suppose that there are five bidders at a first price auction. There are two important coalition structures in this case: (a) the all-inclusive cartel, and (b) two cartels, one with a single bidder and one with the remaining four bidders. The PCES of the cartel with four bidders is greater than that of any finer coalition for all relevant reserve prices, and hence bidders have no incentive to defect.[15] The r-core then has only the above two coalition structures as possible elements.

The PCES given a reserve price is shown for both coalition structures in Figure 2. Consider first the all-inclusive cartel, and suppose that a single firm is considering a defection. This potentially defecting firm would form a coalition of 1, leaving behind the coalition of 4. Against an exogenously formed all-inclusive

---

[15] In other words, there is no coalition structure with coalitions of no more than three bidders that provides a PCES greater than that of a coalition of four bidders competing with a single other bidder. This implies that there is no profitable defection from the coalition of four. This fact has been verified for reserve prices except possibly those very close to one. Even if this is not true for some reserve price sufficiently high, such a reserve price will never be optimally chosen (since the resulting expected revenue will be extremely small) and hence the conclusions of this paper remain unchanged.

cartel of size 5, the optimal reserve price is 0.699.[16] At this reserve and for all smaller reserves, the PCES is greater for the all-inclusive cartel than it is for a bidder who defects. There is no incentive to defect, so the unique element in the r-core for all such reserve prices is the all-inclusive cartel.

However, now consider a higher reserve price. Figure 3 shows the difference between the PCES with an all-inclusive cartel and the PCES for a coalition of one against a coalition of four. At all reserves greater than 0.735, the PCES is higher for a single defecting firm than it is in the all-inclusive cartel. For these reserves, there is incentive for one bidder to defect from the all-inclusive cartel. Consider a reserve price of 0.80. At this, the PCES of the all-inclusive cartel is 0.0154, while it is 0.0156 for a defecting firm. The remaining coalition of size four has a PCES of 0.0134.[17]

Therefore, the r-core always contains only one element. If the auctioneer sets a reserve price of less than 0.735, the r-core consists of the all-inclusive cartel. If instead the auctioneer sets a reserve price of 0.735 or higher, the r-core consists of two coalitions of size four and one respectively. The choice of reserve price will depend on the expected revenues for the auctioneer, which will in turn depend on

---

[16] See Figure 1.

[17] As noted above, there is no further incentive for firms to defect from the coalition of four. The PCES of each of two single bidders against a coalition of three is less than the PCES of a coalition of four against a single bidder, for all relevant reserve prices.

the resulting coalition structure.

Figure 4 shows the expected revenue against both sets of coalitions by reserve price. Against an all-inclusive cartel, the optimal reserve is 0.699 giving an expected revenue of 0.582. Recall that against coalitions of size four and one, the optimal reserve (shown in Table 1) is 0.659 giving an expected revenue of 0.608. However, with this reserve, such coalitions will not form.

Figure 5 shows the same expected revenues but only over the region of interest. The horizontal line shows the expected revenue earned by conceding to an all-inclusive cartel and choosing a reserve of 0.699. The vertical line is placed at 0.735, so for all reserves greater than this two coalitions will form. Then clearly any reserve that corresponds to the expected reserve curve to the upper right of the intersection of these lines will give the auctioneer higher expected revenue than is possible from an all-inclusive cartel. Since expected revenue in this region is declining in the reserve, the auctioneer should choose the smallest reserve so that one bidder will defect from the all-inclusive cartel. This would provide an expected revenue of 0.593. Note that this is larger than the expected revenue of 0.582 that resulted from a second price auction, and hence in equilibrium the auctioneer will choose a first price auction with the above reserve.

6. Discussion

The result of the analysis of the five bidder case is that for a first price auction, the auctioneer may benefit from choosing a reserve that is not optimal in the sense that it maximizes revenue conditional on a set of coalitions, but is optimal in the sense that it maximizes revenue conditional on the coalitions that will form given the reserve price. By raising the reserve, the auctioneer increases the relative attractiveness of defecting from the all-inclusive cartel, and then benefits from the induced competition. This appears to be true for more than five bidders.[18]

Recall that MS show results with exogenous coalitions where, in some cases, a second price auction with an optimal reserve would revenue dominate a first price auction with an optimal reserve. This was contrary to the belief that first price auctions are less susceptible to collusion than second price auctions. However, their analysis did not allow the auctioneer to act in the more sophisticated manner considered here. By allowing the auctioneer to choose a reserve price that will affect the coalitions formed, with sufficiently many bidders first price auctions again dominate second price by making the most extensive collusion not

---

[18] The same steps as above were taken with both seven and eleven bidders. The numerical analysis and comparisons show that the auctioneer has incentive to choose a first price auction with a reserve price that is sufficiently high to force a defection from the all-inclusive cartel, resulting in a higher expected revenue than at a second price auction.

sustainable.[19]

This analysis shows that if the bidders consider the reserve price when colluding, then an appropriate choice of the reserve price may prevent some coalitions from forming to the benefit of the auctioneer. This suggests that observed auction types and reserve prices that may appear non-optimal given the coalition structure at an auction may in fact be optimal. This also explicitly shows one way in which first price auctions are less susceptible to collusion than second price auctions.

This potential action of the auctioneer to combat collusion requires the auctioneer to understand the nature of the bidders at the relevant auction. Figure 6 shows the expected revenue at an auction with 5 bidders acting non-cooperatively. This, as well as the other plots of expected revenue in this paper, show that the expected revenue drops towards zero very quickly as the reserve becomes high enough. A careful balance must be maintained between choosing a high enough reserve to discourage collusion and choosing a reserve that is not too high as to dramatically decrease the expected revenue, especially if there is a reasonable possibility that an all-inclusive cartel will not form even if the reserve price is low.

This research also gives some insight into a stable first price collusive mech-

---

[19] As intuition behind why this is not true for a small number of bidders, note that the incentive to defect from an all-inclusive cartel increases as the number of bidders increases. More specifically, the PCES for an all-inclusive cartel decreases much faster with an increasing number of bidders than the expected surplus of a single defector.

anism. While a second price mechanism will be able to support an all-inclusive cartel, the above results suggest that in some situations there is no first price mechanism that can do the same. Any research on such mechanisms should then not focus entirely on supporting full collusion.

7. References

Athey, Susan, 2002, “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information,” Econometrica, 69, pp. 861-890.

Bajari, Patrick, 2001, "Comparing Competition and Collusion in Procurement Auctions: A Numerical Approach", Economic Theory, 18, pp. 187-205.

Graham, Daniel A., and Robert C. Marshall, 1987, “Collusive Bidder Behavior at Single-Object Second-Price and English Auctions”, Journal of Political Economy, 95(6), pp. 1217-39.

Huang, Chen-Ying, and Tomas Sjostrom, 2002, “Consistent Solutions for Cooperative Games with Externalities”, Games and Economic Behavior, forthcoming.

Ichiishi, Tatsuro, 1981, “A Social Coalitional Equilibrium Existence Lemma”, Econometrica, 49(2), pp. 369-77.

Lebrun, Bernard, 1996, ”Existence of an Equilibrium in First Price Auctions”, Economic Theory, 7, pp. 421-443.

________, 1998a, ”Comparative Statics in First Price Auctions”, Games and Economic Behavior, 25, pp. 97-110.

________, 1998b, ”The First and Second Price Auctions in a Class of Asymmetric Examples”, mimeo.

________, 1999, ”First Price Auctions in the Asymmetric N Bidder Case”, International Economic Review, 40(1), pp. 125-142.

Li, Huagang, and John G. Riley, 1999, "Auction Choice", UCLA Working Paper.

Mailath, George J., and Peter Zemsky, 1991, “Collusion in Second Price Auctions with Heterogeneous Bidders”, Games and Economic Behavior, 3(4), pp. 467-86.

Marshall, Robert C., Michael J. Meurer, Jean-Francois Richard, and Walter Stromquist, 1994, ”Numerical Analysis of Asymmetric First Price Auctions”, Games and Economic Behavior, 7, pp. 193-220.

Marshall, Robert C., and Steven P. Schulenberg, 2003, “Numerical Analysis of Asymmetric Auctions with Optimal Reserve Prices”, working paper.

Maskin, Eric and John Riley, 2000a, ”Asymmetric Auctions”, Review of Economic Studies, 67(3), pp. 413-38.

_____________, 2000b, ”Equilibrium in Sealed High Bid Auctions”, Review of Economic Studies, 67(3), pp. 439-54.

Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T.

Vetterling, 1986, Numerical Recipes, Cambridge University Press: New York.

Ray, Debraj, and Rajiv Vohra, 1997, “Equilibrium Binding Agreements”, Journal of Economic Theory, 73(1), pp. 30-78.

Figure 1: Optimal Reserve and the Corresponding Expected Revenue and Expected Surplus at a Second Price Auction with an All-Inclusive Cartel of Size N

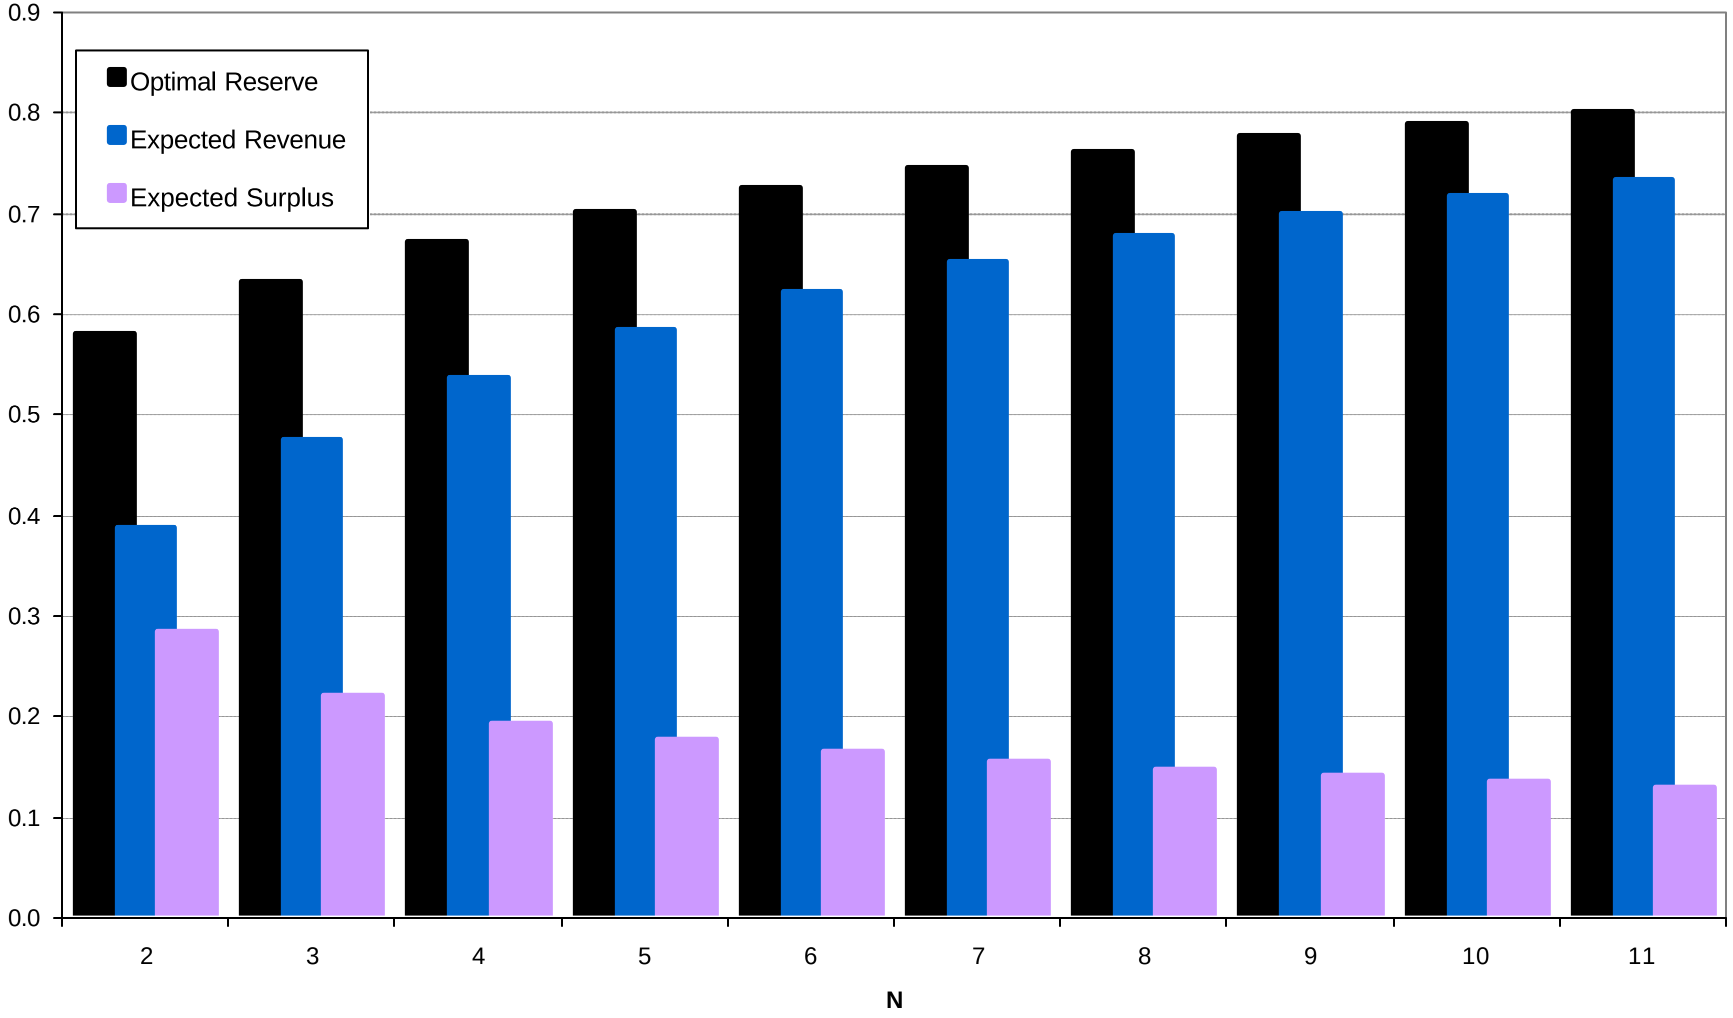

Figure 2: Per Capita Expected Surplus at a First Price Auction with Five Bidders

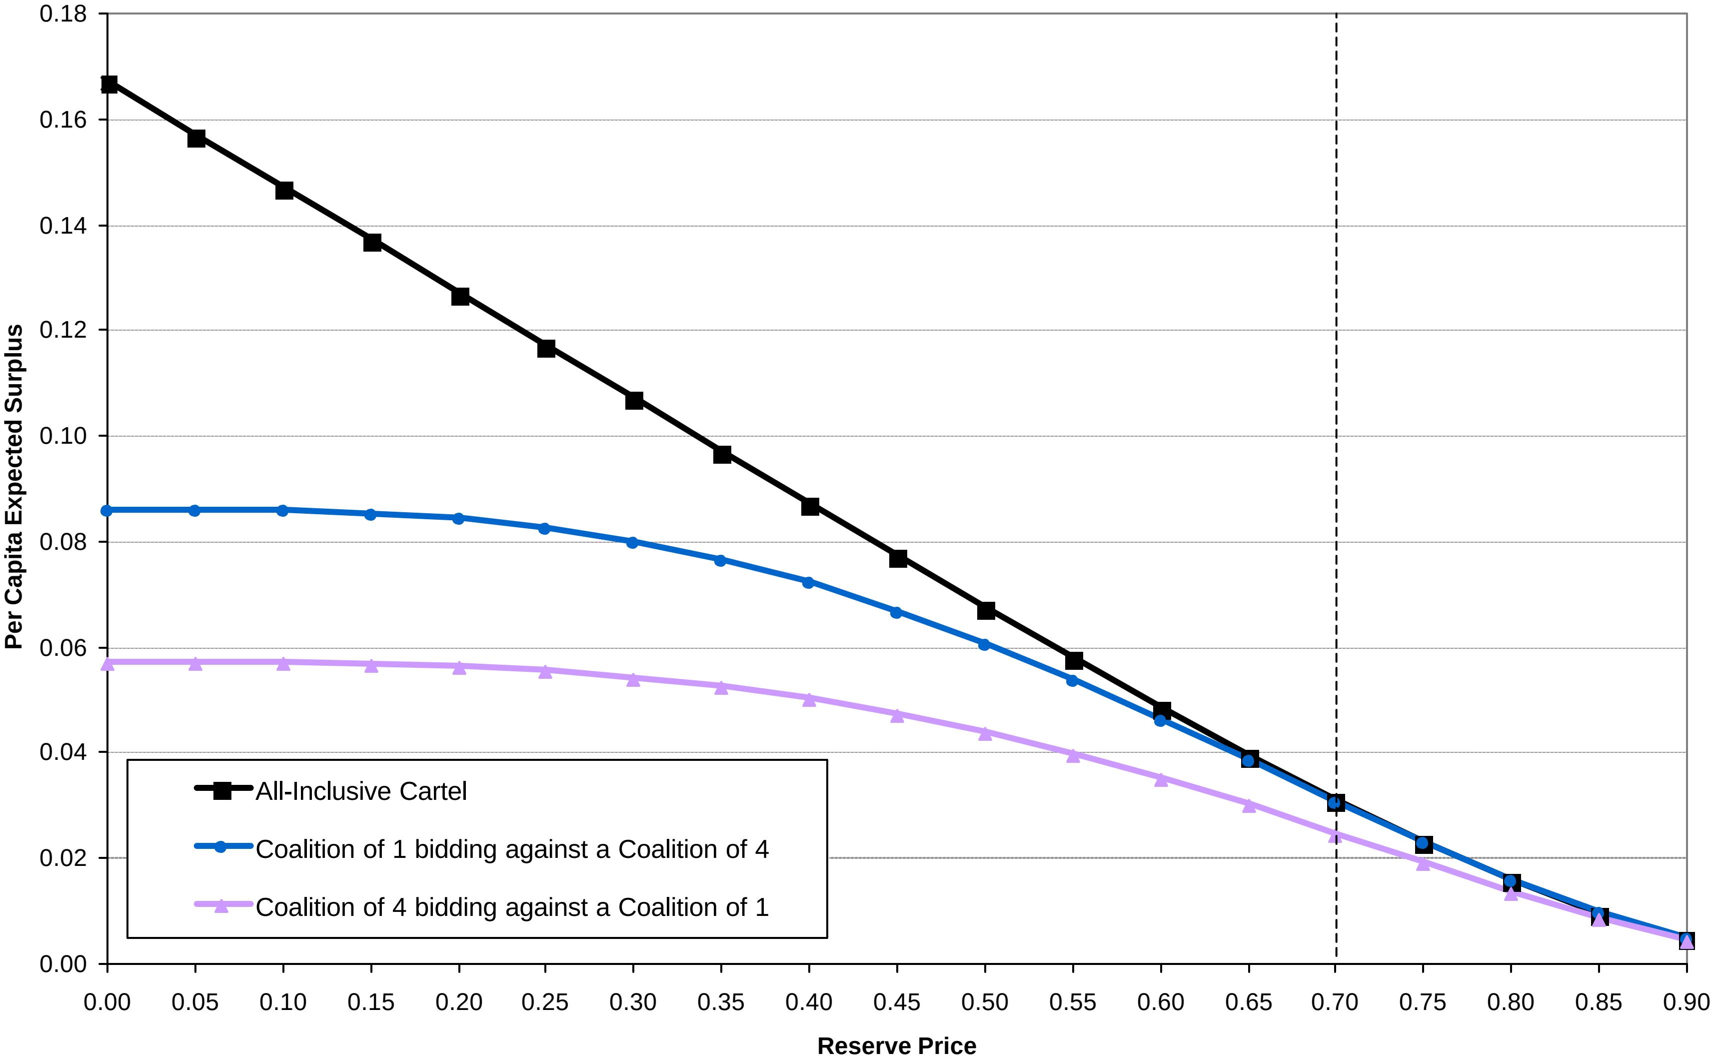

Figure 3: Difference between Per Capita Expected Surplus of an All-Inclusive Cartel of Five and of a Coalition of One against a Coalition of Four

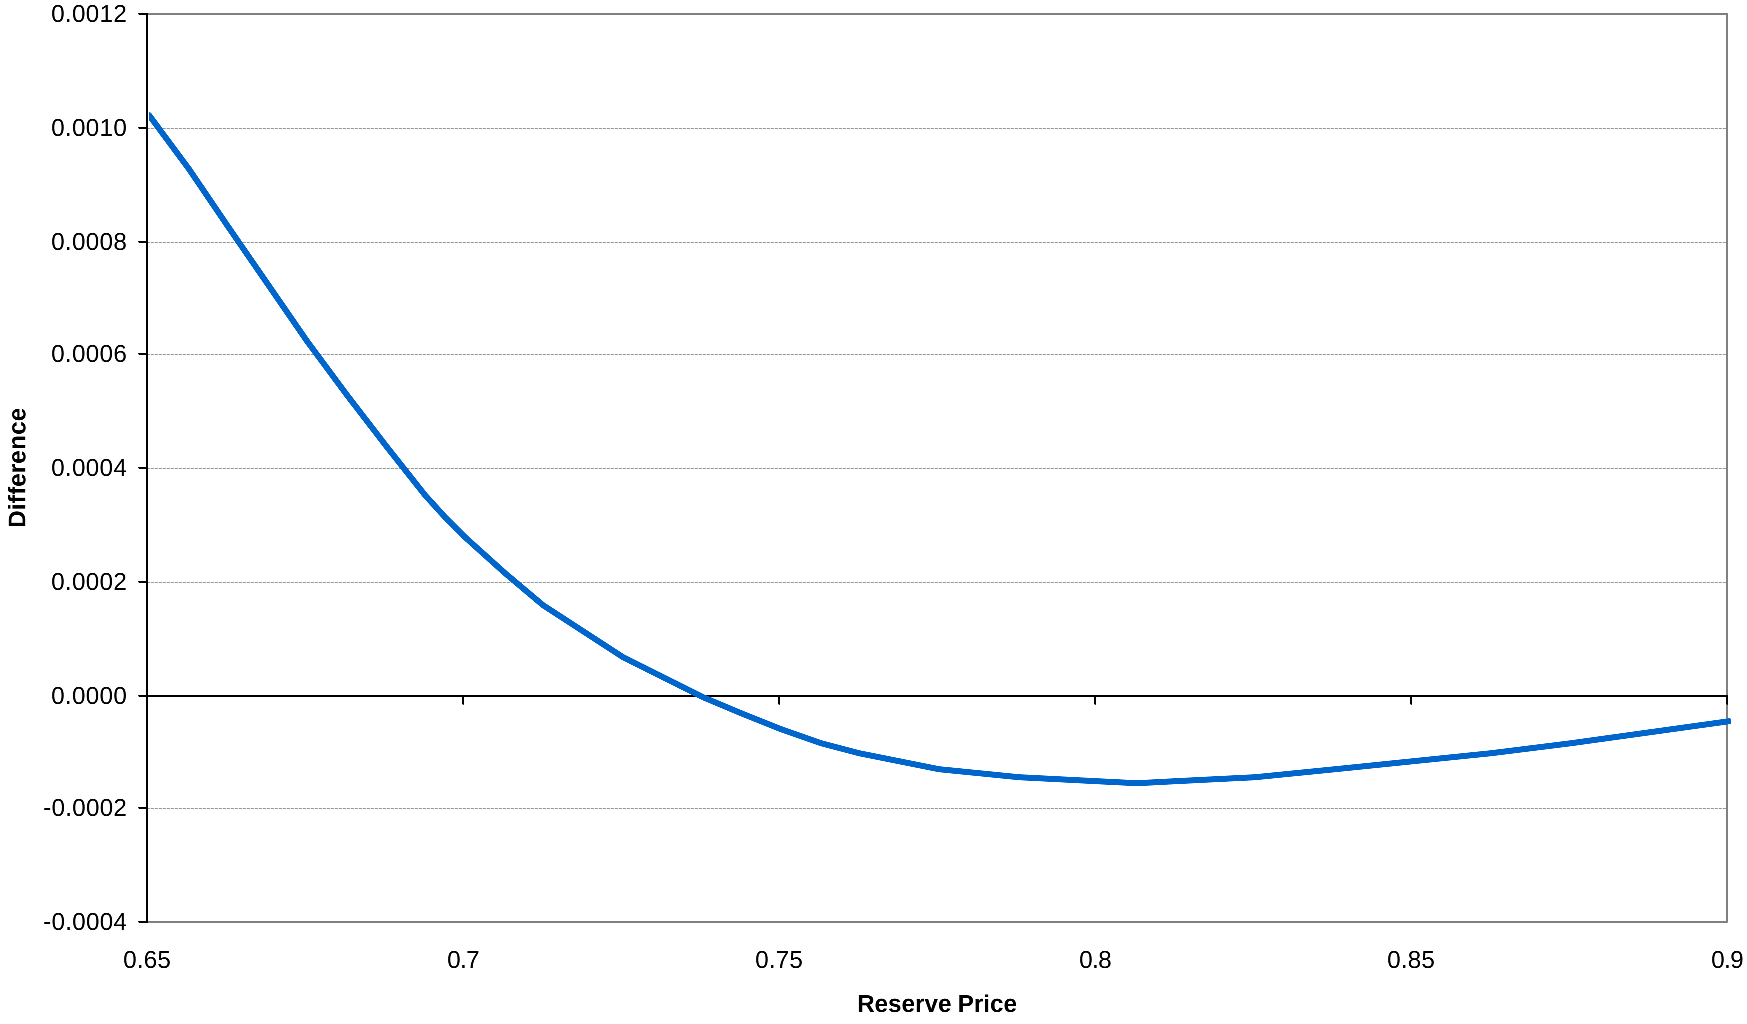

Figure 4: Expected Revenue at a First Price Auction with Various Coalitions of Five Bidders

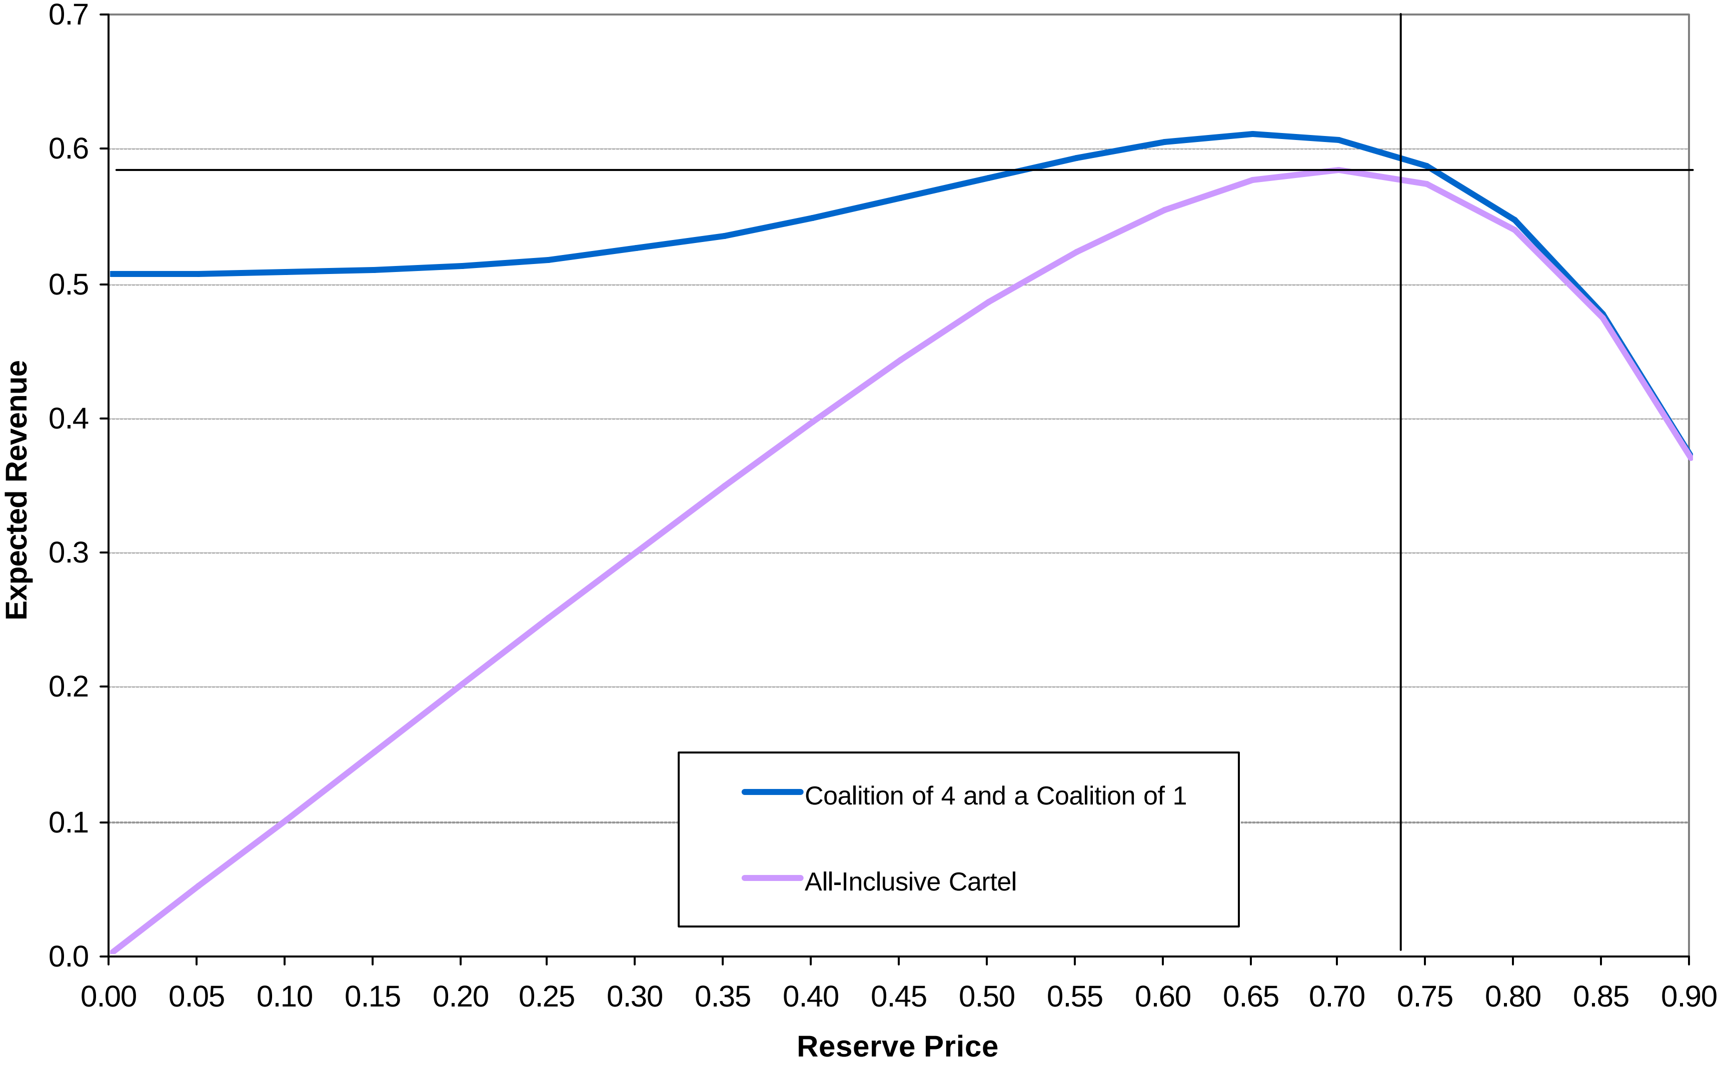

Figure 5: Expected Revenue at a First Price Auction with Various Coalitions of Five Bidders

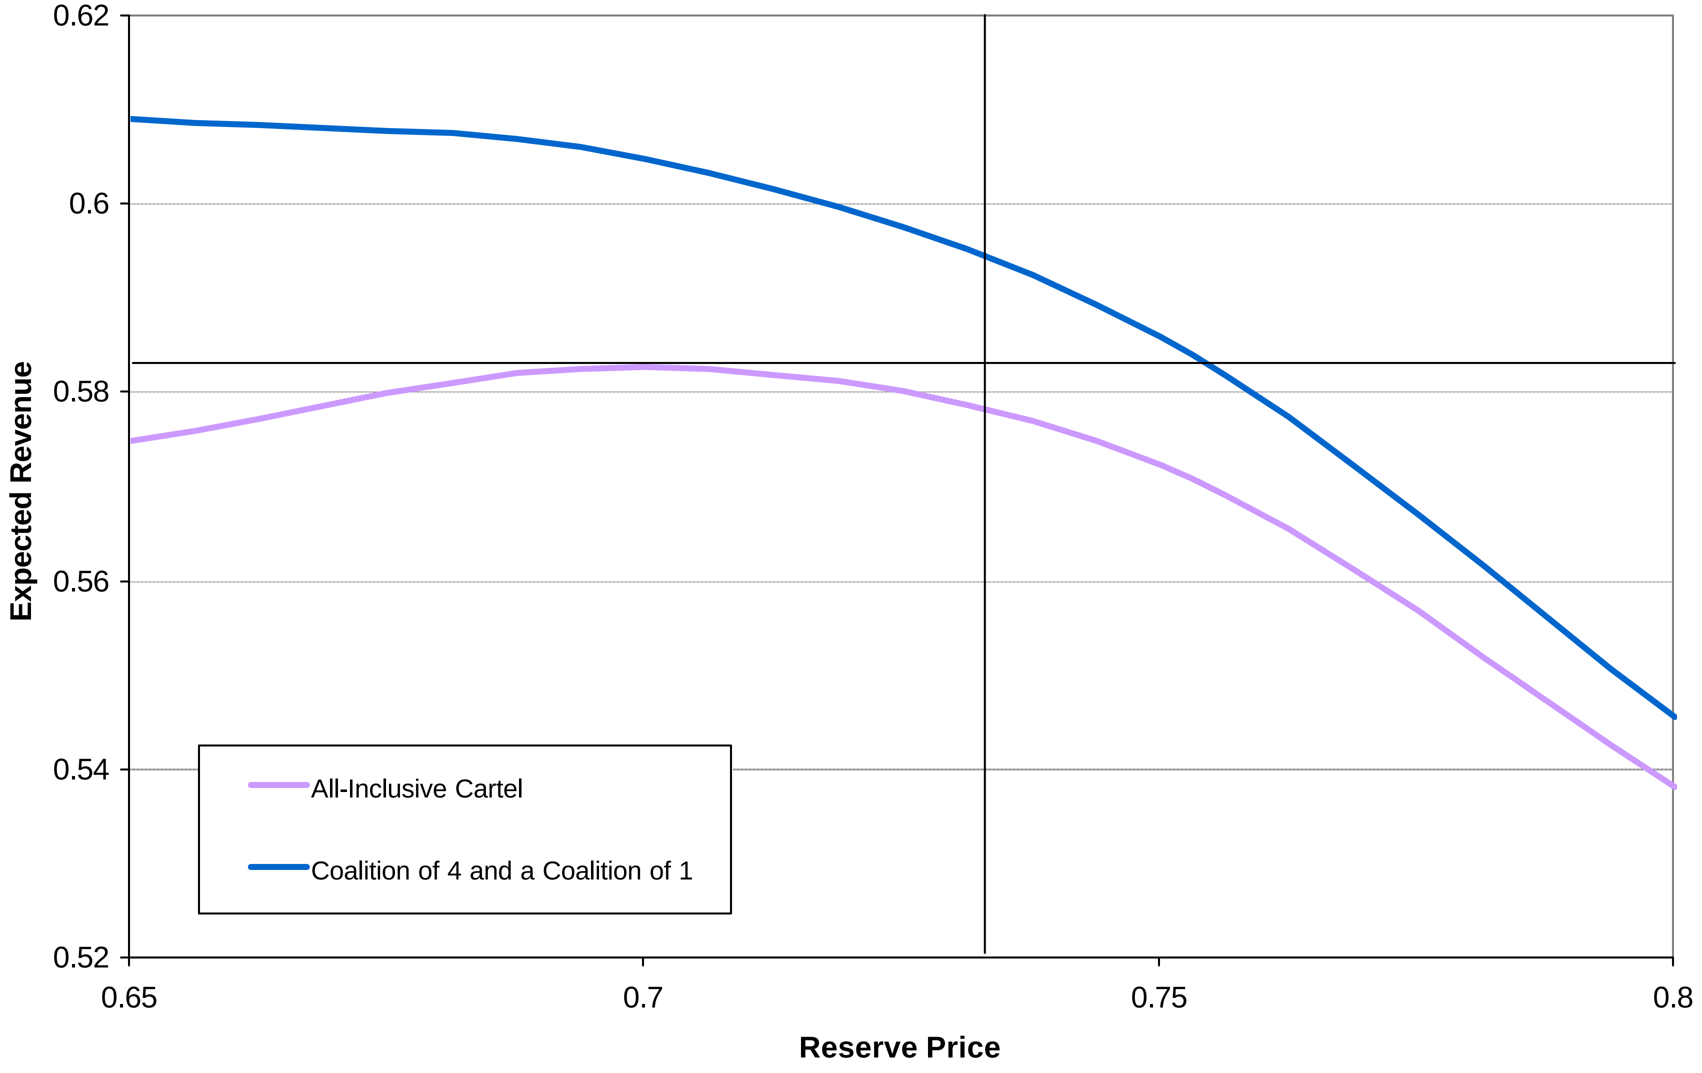

Figure 6: Expected Revenue at a First Price Auction with Five Non-Cooperative Bidders

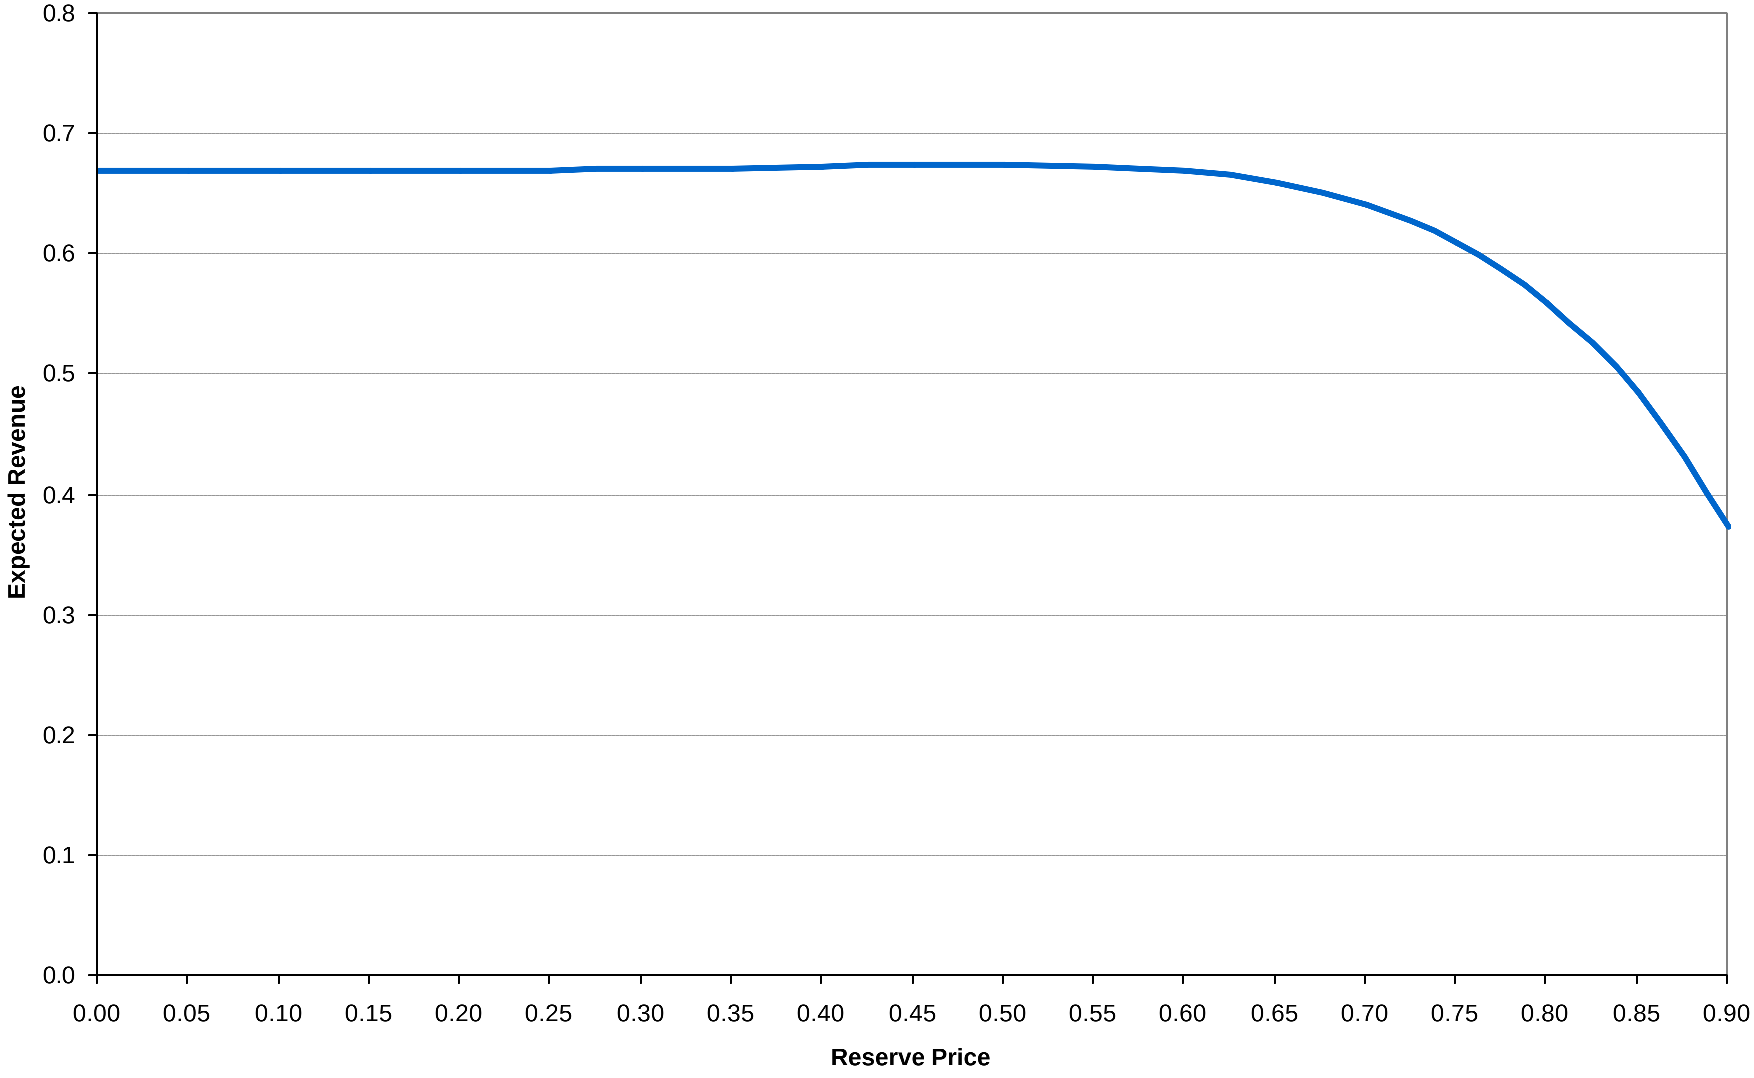

Chapter 4

Improving the Corporate Leniency Policy

Steven P. Schulenberg

1. Introduction

The Corporate Leniency Policy has recently been declared by the Department of Justice (DOJ) to be the single greatest weapon against explicit collusion. There is significant anecdotal evidence to back this assertion. The majority of the antitrust cases currently in the courts involve leniency in some manner. Since the leniency policy was modified in 1986 to make it much more usable, the amount of collusion that the DOJ has knowledge about has increased dramatically.

The leniency policy provides the opportunity for firms participating in a cartel to come forward to the DOJ. If they are able to either provide evidence that leads to a new investigation or that significantly strengthens the position of the DOJ in an existing investigation, then the firm can receive amnesty and avoid some or all of the criminal fines. In addition, the individuals involved may also avoid criminal fines and prison sentences. This opportunity is available to exactly one member of the cartel, which is that firm which comes forward and provides such evidence first.

The fundamental concern of cartels is the incentive constraint, which balances the incentive for firms to continue the collusive agreement with the incentive for firms to cheat (and thus earn short term high payments but also be punished). This constraint determines what degree of collusion may be maintained. Theory has established some factors that affect the incentive constraint, such as the interest rate and the cost of capacity. The antitrust policy of the DOJ affects the constraint by adding punishments (criminal penalties and civil fines) which may only be imposed if a firm colludes. The leniency policy modifies this addition by allowing the firm to avoid (at least part of) these punishments. The DOJ has established the current regulations under the premise that the punishments are sufficiently high to discourage cartels from forming and that the leniency policy promotes the discovery of cartels that would otherwise go undetected. In other words, it is the belief of the DOJ that both the fines and the potential for leniency affect the constraint in a manner that causes collusion to be more difficult to sustain.

The purpose of this paper is to show that this may not be the case. Using Davidson and Deneckere (1990) as a basis, it is first shown that an antitrust policy without leniency can actually encourage collusion, by providing an additional punishment mechanism for the cartel to use to dissuade deviations. It is then shown that the addition of the leniency policy reduces this undesirable effect of the antitrust policy some, but not entirely. This suggests the unappealing result that the current antitrust policy of the United States increases the amount of collusion above that which would occur with no government oversight. However, a modification of the leniency policy (motivated by a precedent set by the False Claims Act) can reverse this. By providing a bounty to the firm which comes forward, the DOJ can cause members of a cartel to have more of an incentive to deviate than with no government oversight.

The paper proceeds as follows. Section 2 gives a description of the Corporate Leniency Policy. The relevant literature is presented in Section 3. Section 4 describes the basis model used in this paper, followed by the equilibrium, as determined in Davidson and Deneckere (1990), in Section 5. The extensions to an antitrust policy without leniency and then with leniency are discussed in Sections 6 and 7 respectively. Section 8 describes the extension when a bounty is given to the firms that come forward. Concluding remarks are given in Section 9.

2. The Corporate Leniency Policy

The Corporate Leniency Policy allows firms, providing they meet the necessary conditions, to avoid most or all criminal penalties involved with a successful antitrust prosecution. This implies protection from both criminal fines for the corporation, and also protection from criminal fines and prison sentences for the individuals involved. Amnesty does not, however, affect the restitution that the firm is required to pay to any damaged parties, and hence it is still likely to be subject to civil penalties.

Corporations are eligible for amnesty if they meet several conditions. Most importantly, the corporation must be the first member of the cartel to come forward before a DOJ investigation has begun. If an investigation is underway, then amnesty is still available if the DOJ does not yet have enough evidence to make a strong case against the cartel. In addition, the corporation must have terminated the collusion promptly and must fully cooperate with the DOJ throughout the investigation, the confession must have been a corporate act rather than just by a few individuals, and the corporation must not have been the leader in or originator of the cartel.[1] If all of these conditions are met, then the corporation

---

[1] There is some discretion involved in whether or not a corporation has met all these conditions. However, the DOJ has apparently made a point to make all decisions as transparent as

automatically receives amnesty (along with individuals within the corporation).

The Corporate Leniency Policy was first instituted in 1978. At first it was used sparingly, with approximately one application for amnesty filed per year. In 1993, three major changes were made to the policy. Prior to 1993, amnesty was not available if an investigation was already underway. Additionally, if an investigation was not underway, the first corporation coming forward was only eligible for amnesty rather than automatically receiving it. This allowed the DOJ to retain some discretion, but it also required firms to gamble when coming forward. Finally, the criminal protection did not originally extend to individuals. The overall effect of these changes was to make amnesty more beneficial, make it applicable to more cartels, and reduce the uncertainty involved.

These changes did not have an immediate impact, presumably because it took some time for corporations to become confident that the changes would be applied as written. However, it did not take long for the number of amnesty applications to dramatically increase, to a rate of approximately two per month by 1998. This rate was further bolstered by the formal introduction of the "Amnesty Plus" provision. This allows firms under investigation in one industry to apply for amnesty in an unrelated industry, even if they do not directly qualify for amnesty in the initial

---

possible to remove any uncertainty firms may have.

market.[2] In addition to the standard amnesty benefits in the second market, the corporation will often receive a large reduction in penalties for the first through a plea agreement. Over half of the most recent DOJ investigations began due to this provision.

3. Relevant Literature

The types of indirect evidence that may be used to infer that firms were explicitly price-fixing are described in Kovacic (1993)[3]. Making such an inference is a difficult proposition, since most collusive market outcomes may also be loosely justified as also consistent with non-cooperative behavior. At the end of this article, it is noted that one way to ease this problem is by generating more direct evidence of collusion. Kovacic proposes that a "bounty system that focused on horizontal price fixing might increase the rate at which express, covert collusive agreements are detected and punished." This was motivated by the observed suc-

---

[2] For instance, if the firm was not "first in" or if the initial investigation was already too far along.

[3] Proving that antitrust violations have occured is done using evidence that is either direct or indirect. Examples of direct evidence include detailed meeting notes, copies of the written cartel agreement, or videotapes of cartel members discussing prices. As cartels are unlikely to make such materials readily available, the majority of cases involve mostly indirect evidence. This category includes the outcomes of a collusive agreement which, when observed, allow an inference on the existence of such an agreement to be drawn. Examples of this type of evidence include unexplained rises in price, increasingly stable market shares, anti-competitive transactions between rival firms, and certain patterns of public price announcements.

cess of the False Claims Act, which allows a bounty to be paid to citizens who provide evidence of fraud in government contracts. As shown here, this is a very reasonable conjecture.

The Corporate Leniency Policy has also been analyzed in Spagnolo (2000). This considers a standard Bertrand situation, and shows that a leniency program allows collusion to occur when it otherwise would not. The addition of the leniency program to the antitrust policy is the increment that allows collusion to be supported. This distinctly differs from the current paper, where the leniency program mitigates the incentives to collude provided by the standard antitrust policy. Spagnolo (2000) also discusses one-shot first price auctions, claiming that the leniency policy makes collusion possible where it otherwise is not.

Other authors have noted that there may be effects of government policies that are contrary to those intended. Bernheim (1984) examines the effects of several policy instruments that have the goal of reducing the concentration of suppliers in a given market, such as subsidization of entry or the forced breakup of dominant firms. It is found that such policies have both positive and negative effects on concentration, with the net effect ambiguous. Examples are provided where the policies in fact result in an increase in concentration.

In a more related setting, McCutcheon (1997) has suggested that the antitrust policy (not considering the leniency program) can have the unexpected result of increasing the ability of firms to collude. As will be discussed below, the punishments that may occur if the firms are found to be explicitly colluding can be incorporated into the mechanism of the cartel. As a result, the extent of collusion that can be sustained will increase. This occurs in McCutcheon (1997) when renegotiation proof equilibria are considered. As discussed in Farrell and Maskin (1989), it is unreasonable to expect that firms will carry out punishments if all of the firms involved would benefit by instead renegotiating a new agreement. In some situations, firms are able to collude effectively when renegotiation is not possible but cannot collude at all when renegotiation is available and costless. This suggests that if renegotiation is costly enough, then firms will choose not to take this option and will instead simply carry out the punishments in the original agreement. McCutcheon (1997) discusses renegotiation as occurring at a meeting that carries with it a nontrivial cost resulting from the possibility of detection and subsequent punishment, and provides some evidence that this cost might be sufficient to prevent renegotiation. As a result, antitrust policy leads to increased punishments being available to cartels and thus a greater extent of sustainable collusion.

4. Basis Model

The basis model used in this paper is that of Davidson and Deneckere (1990). They considered a two-stage infinite horizon game with two firms producing a homogeneous product. There is no possibility of entry. At time zero, the firms simultaneously and non-cooperatively choose capacity levels K_i at a cost of c per unit. In the second stage, from time one and on, the firms choose prices and costlessly produce the resulting output up to their capacity. Firms may not change capacity after time zero. The (constant) interest rate r is used by the firms to discount the per-period payoffs.

The market demand curve is assumed to be

$$D(p) = 1 - p$$

with the appropriate inverse $P\left(x\right)$. The firm with the lower price sells a quantity that is the smaller of the total demanded and its capacity. The residual demand for the firm with the higher price is assumed to be the additional amount demanded at its price beyond the capacity of the other firm.[4] Hence,

$$\begin{aligned} q\left(p_j\right) &= D(p_j) && \text{if } p_j < p_i \\ &= \max\left\{0, D\left(p_j\right) - K_i\right\} && \text{if } p_j > p_i \end{aligned}$$

When firms choose an identical price, the firms are assumed to share the market using the sharing rule

$$S_i = \frac{\min\left\{K_i, 1\right\}}{\min\left\{K_1, 1\right\} + \min\left\{K_2, 1\right\}}.$$

This allows firms to share the quantity proportionally with capacity, until the point when capacities reach the size of the market (which is 1, since $P(1) \equiv 0$).

Firms are assumed to collude in the second stage but not in the first.[5] The collusion is enforced with grim trigger strategies, which have firms choosing collusive prices until a deviation occurs, after which they revert to static Nash equilibrium prices forever. In equilibrium, the firms will then choose capacities and then prices

---

[4] This assumption requires that the customers with the highest reservation prices are served first, leaving the smallest possible residual demand to the higher priced firm. This is also assumed in Kreps and Scheinkman (1983). A different assumption is that the lower priced firms serves customers proportionally, which leads to the different residual demand seen in Davidson and Deneckere (1986).

[5] The authors do not attempt to explain this choice beyond referring to “widespread observance” and arguments such as that in Fershtman and Muller (1986).

such that there is no incentive to deviate.

Equilibria are found by first considering the price subgames for each set of capacities, and then finally considering the full game. For each set of capacities and prices, there are three relevant profits. Let the profit for firm i be π_i^N in the static Nash equilibrium, π_i^c when firms are colluding, and π_i^{ch} when cheating optimally. The obvious relationship between these profits is $\pi_i^N \leq \pi_i^c \leq \pi_i^{ch}$.

5. No Antitrust Policy (Davidson and Deneckere Outcomes)

Consider a firm that is currently colluding at given capacities and prices, and is considering the benefits of cheating. By not cheating, the firm will earn π_i^c each period. The total discounted payoff is then

$$\pi_i^c + \frac{1}{1+r}\pi_i^c + \left(\frac{1}{1+r}\right)^2 \pi_i^c + \left(\frac{1}{1+r}\right)^3 \pi_i^c + ... = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \pi_i^c.$$

If the firm instead cheated, it would receive π_i^{ch} in the current period and then π_i^N in each period after. The total discounted payoff in this case is then

$$\pi_i^{ch} + \frac{1}{1+r}\pi_i^N + \left(\frac{1}{1+r}\right)^2 \pi_i^N + \left(\frac{1}{1+r}\right)^3 \pi_i^N + ... = \pi_i^{ch} + \sum_{t=1}^{\infty} \left(\frac{1}{1+r}\right)^t \pi_i^N.$$

Recall that

$$\sum_{n=1}^{\infty}\left(\frac{1}{1+x}\right)^{n}=\frac{1}{x}.$$

We can then write the payoff when not cheating as

$$\pi_i^c+\frac{1}{r}\pi_i^c$$

and the payoff when cheating as

$$\pi_i^{ch}+\frac{1}{r}\pi_i^N.$$

The net gain from cheating is then

$$\left(\pi_i^{ch}+\frac{1}{r}\pi_i^N\right)-\left(\pi_i^c+\frac{1}{r}\pi_i^c\right)=\left(\pi_i^{ch}-\pi_i^c\right)-\frac{1}{r}\left(\pi_i^c-\pi_i^N\right)$$

Define

$$Z_i=\left(\pi_i^{ch}-\pi_i^c\right)-\frac{1}{r}\left(\pi_i^c-\pi_i^N\right).$$

Firm i therefore has no incentive to cheat if $Z_i \leq 0$.

For each set of capacities (K_1, K_2), each set of chosen prices (p_1, p_2) gives rise to a set of profits $\left\{(\pi_i^N, \pi_i^c, \pi_i^{ch}), i = 1, 2\right\}$ and therefore (Z_1, Z_2). Define the set of prices which can be supported (i.e. no cheating), given (K_1, K_2), as

$$\Omega = \{(p_1, p_2) : Z_1 \leq 0 \text{ and } Z_2 \leq 0\}$$

Davidson and Deneckere (1990) show that there are three categories of equilibria that result from the capacity and price game described above.

- 1. Unconstrained Semi-Collusive Equilibria (USE): Firms charge the (single period) monopoly price. These have high capacity and high excess capacity.
  2. Constrained Semi-Collusive Equilibria (CSE): Firms charge less than the monopoly price but above the static Nash price. These have moderate capacity and moderate excess capacity.
  3. Noncollusive Equilibria (NCE): Firms charge the static Nash price, since no higher price is supportable. These have low capacity and no excess capacity.

Note that more collusion is accompanied with more capacity and more excess capacity. The excess capacity is needed to give firms the ability to punish.

USE occur with small interest rates and/or costs of capacity, while NCE occur with large interest rates and/or costs of capacity. CSE may or may not occur at intermediate values. Small interest rates mean that the future is not discounted much, making punishments more meaningful. Therefore, more collusion may be supported as the interest rate falls. As the cost of capacity rises, the cost of carrying excess capacity also rises. This makes having excess capacity to use when punishing less attractive, and hence the level of collusion that may be supported will fall.

6. With Antitrust Policy but no Leniency

Now consider the same situation with the addition of an antitrust policy. If a cartel is discovered, then each participating firm is administered an exogenously determined criminal fine $F^{cr} > 0$ and a civil fine $F^{ci} > 0$ which must be paid in the following period.[6] It is assumed that the probability of an existing cartel being detected and punished is only non-zero immediately following a meeting where renegotiation occurs. It is also assumed, following McCutcheon (1997), that the expected cost of such a meeting is greater than the cost to a punishing firm of actually carrying out the punishment. As a result, the cartel will never rationally choose to renegotiate and hence there is zero chance that the collusion will be

---

[6] The fines in general would not realistically be exogenous. However, there is also not a clear mechanism that would lead to certain amounts due to, for instance, pleas and settlements. There would be an upper bound on any fines due to a bankruptcy constraint.

detected through some type of auditing mechanism.[7] Therefore the fines will only be imposed if the government discovers the cartel through the explicit choice of a participating firm.

Firms may, in any period after period zero, come forward and present evidence of their collusion. There is no direct benefit to coming forward, and doing so guarantees that all firms (including the firm coming forward) will pay both fines. Firms may then use this as part of their strategy. Specifically, firms use a grim trigger strategy that includes the criminal and civil fines. If a firm cheats, then both firms revert to the static Nash equilibrium and in addition the discovering firm comes forward to the DOJ.

For a given set of capacities and prices, the payoffs to colluding remain as without an antitrust policy. However, the payoffs to cheating are now lower since there are fines that must additionally be paid. Specifically, the payoff when cheating is

$$\pi_i^{ch} + \frac{1}{r}\pi_i^N - \frac{1}{1+r}\left(F^{cr} + F^{ci}\right)$$

---

[7] This is a reasonable assumption given the difficulty of successful antitrust prosecution based solely on indirect evidence. This is also a good assumption if the colluding firms truly believe that they will not be caught, since that is the relevant belief for cartel stability.

The net gains from cheating are then

$$Z_i^A = \left(\pi_i^{ch} - \pi_i^c\right) - \frac{1}{r}\left(\pi_i^c - \pi_i^N\right) - \frac{1}{1+r}\left(F^{cr} + F^{ci}\right)$$

Note that $Z_i^A < Z_i$.

The set of supportable prices is

$$\Omega^A = \left\{(p_1, p_2) : Z_1^A \leq 0 \text{ and } Z_2^A \leq 0\right\}.$$

Since the net gain from cheating is decreased, there are, for given capacities, prices at which cheating was beneficial with no antitrust policy but is not beneficial with an such a policy. Therefore, $\Omega \subset \Omega^A$. Hence more collusion and higher prices may be supported.

The main result of this section is that for a given interest rate and cost of capacity, in general more collusion and higher prices will result with DOJ oversight and an antitrust policy (as described above), since firms may use it as an additional method of punishment. This will also reduce the amount of excess capacity that is needed, since its role in punishment is not needed as much.

7. With Antitrust Policy and Leniency

In addition to the leniency policy of the previous section, suppose that there is also a leniency policy. Firms may, in any period after period zero, choose to come forward and apply for leniency simultaneously with choosing a price. If firm i comes forward, then only firm j must pay a criminal fine of F^{cr} in the next period while both firms must pay a civil fine of F^{ci} next period. Therefore, both firms pay a fine, but the firm that came forward pays a smaller overall fine. Also, the firms must choose the static Nash price in all following periods.

Firms still use a grim trigger strategy, where they revert to the static Nash price in all periods after a firm has cheated. In addition, the non-cheating firm will come forward immediately after the cheating occurs and therefore receive amnesty. However, the cheating firm will anticipate this. Optimal cheating then has a firm choosing the same cheating price as without leniency, but in the same period coming forward and applying for leniency.[8] This is done to avoid the criminal fines, since the civil fines must be paid regardless.

Since the cheating firm only now faces the civil fine, the net gains from cheating

---

[8] Recall that one requirement of receiving amnesty is that the firm must stop colluding immediately, which is consistent with this outcome.

are

$$Z_i^L = \left(\pi_i^{ch} - \pi_i^c\right) - \frac{1}{r}\left(\pi_i^c - \pi_i^N\right) - \frac{1}{1+r}\left(F^{ci}\right)$$

Note that $Z_i^A < Z_i^L < Z_i$.

The set of supportable prices is

$$\Omega^L = \left\{(p_1, p_2) : Z_1^L \leq 0 \text{ and } Z_2^L \leq 0\right\}$$

Since the net gain from cheating is decreased, $\Omega \subset \Omega^L \subset \Omega^A$. Hence in general less collusion and lower prices may be supported with the addition of leniency on top of the antitrust policy than with just the antitrust policy. This shows that, as advertised, the leniency policy does reduce the ability of firms to collude. However, it is also the case that there is still more collusion than would result if there was no antitrust policy.

8. Policy Recommendation: Bounty

The goal of the DOJ should be to make as little collusion supportable as possible at any given interest rate and capacity cost. This is not accomplished through the basic antitrust policy, and, as shown above, in fact more collusion is made possible. The addition of leniency as described above alleviates this some, but not completely. However, the benefit of a firm applying for leniency can be changed to make the DOJ improve over the Davidson and Deneckere (1990) outcomes.

Suppose again that firms may come forward and apply for leniency. If firm i comes forward, firm j must pay both F^{cr} and F^{ci}. Firm i does not have to pay F^{cr}, but still must pay F^{ci}. This is unchanged from above. In addition, firm i receives a bounty of $B > 0$. This is a payment made directly to the firm that comes forward in exchange for evidence of the cartel.

Using the same punishment strategy as above, the net gains from cheating are now:

$$Z_i^B = \left(\pi_i^{ch} - \pi_i^c\right) - \frac{1}{r}\left(\pi_i^c - \pi_i^N\right) - \frac{1}{1+r}\left(F^{ci} - B\right)$$

The cheating firm will again come forward. This firm will still have to pay a civil fine, but will be compensated with the bounty. Note that Z_i^B is increasing in B,which implies that a larger bounty gives firms a greater incentive to cheat. Also note that if the bounty is larger than the civil fine ($F^{ci} - B < 0$) then $Z_i^A < Z_i^L < Z < Z_i^B$.

The set of supportable prices is

$$\Omega^B = \left\{(p_1, p_2) : Z_1^B \leq 0 \text{ and } Z_2^B \leq 0\right\}$$

Again if the bounty is larger than the civil fine, $\Omega^B \subset \Omega \subset \Omega^L \subset \Omega^A$. For given capacities, the set of prices that can be supported by collusion is smaller. The highest prices that can be supported are therefore smaller with a bounty than without. In some instances the addition of the bounty will cause collusion to go from possible to impossible. In other instances collusion is supportable in either case, but the addition of bounty gives a lower resulting price. Therefore the antitrust policy with leniency and a bounty allows less collusion and lower prices to be supported than in any of the above three situations.

It is reasonable that the DOJ cannot pay out more in an bounty than it takes in as (criminal) fines for a given market. Therefore, a reasonable constraint on the bounty paid is

$$B \leq F^{cr}.$$

Since Z_i^B is increasing in B, the DOJ should choose as large a bounty as possible, or $B = F^{cr}$. The bounty is then effective in reducing collusion from Davidson and Deneckere levels if the criminal fines are larger than the civil fines.

The above result suggests that in addition to leniency as it currently stands, the DOJ should pay firms that come forward a bounty equal to the criminal fine collected from the other firm. The DOJ should also increase the criminal fines to ensure that they surpass the civil fines. Then, the DOJ will be effective in reducing rather than increasing collusion.

9. Discussion

As noted above, there is precedent in the government paying a bounty in the False Claims Act. This allows a bounty to be paid to an individual that was not involved in creating the false claim. This is similar to the clause in the leniency policy stating that leniency may only be obtained if "[t]he corporation did not coerce another party to participate in the illegal activity and clearly was not the leader in, or the originator of, the activity."[9]

The DOJ has declared the leniency policy to be incredibly successful, based on the number of successful antitrust prosecutions that have involved leniency since the changes in the policy were implemented in 1986 and subsequently understood. This is consistent with the above results, since leniency is an improvement over fines alone. The number of firms coming forward could then be viewed as collusion that could no longer be supported, with the cartel not adjusting properly to prevent cheating from being profitable. While this is certainly a noteable achievement, it is certainly not sufficient reason to not consider additional elements of an

---

[9] See http://www.usdoj.gov/atr/public/guidelines/lencorp.htm.

antitrust policy that may provide even further success. The results of this paper suggest that including a bounty with leniency may be one method of improvement.

This paper does not attempt to claim that the entirety of the antitrust policy is detrimental to the DOJ's goals. Instead, similar to McCutcheon (1997), the goal here is only to show that there are policies that may not in every situation have the desired effect. It is left for future research to consider a more ambitious model of the antitrust policy to answer the larger question about its overall effect and how improvements may be reached, and such research will hopefully be in part motivated by the current paper.

One element of such a model that will constitute an improvement over the current state of the research is the consideration of renegotiation proof equilibria. This paper, as well as the precursors Davidson and Deneckere (1990) and Benoit and Krishna (1987), considers enforcement of collusive agreements through the use of grim trigger punishment strategies. As noted in the literature, such as by Farrell and Maskin (1989), it is reasonable to believe that firms will choose to renegotiate away from paths of actions that are commonly inferior for all firms. It would certainly be interesting and informative to reconsider the current topic with renegotiation in mind.

However, such an investigation is outside of the scope of this paper for several reasons. While it is straightforward process to consider renegotiation for the portion of the model that extends beyond Davidson and Deneckere (1990), proper implementation would also require renegotiation in the remainder as well. Rederiving the entirety of the results in Davidson and Deneckere (1990) with renegotiation, despite the fact that Farrell and Maskin (1989) provides several necessary parts of the analysis, is significantly beyond the simple point that this paper attempts to make. Beyond this, requiring firms to consider renegotiation may be inappropriate given the goals of this paper, one which is to note that in some scenarios the antitrust policy may make collusion more, rather than less, sustainable. Assuming that firms will renegotiate will reduce the extent of possible collusion and possibly also remove any undesirable effects of the antitrust policy. However, if in fact firms will not renegotiate, then it may very well be that these same undesirable effects are actually present. There exists both theoretical and anecdotal evidence that suggests that firms may not always renegotiate, such as discussed in McCutcheon (1997). This paper takes the conservative approach for the questions of interest.

References

1. Abreu, Dilip, 1986, "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, Vol. 39, pp. 191-225.

2. Abreu, Dilip, D. Pearce, and E. Stacchetti, 1986, "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, Vol. 39, pp. 251-69.

3. Benoit, J. P. and Vijay Krishna, 1985, "Finitely Repeated Games," *Econometrica*, Vol. 53, pp. 905-22.

4. Benoit, J. P. and Vijay Krishna, 1987, "Dynamic Duopoly: Prices and Quantities," *Review of Economic Studies*, Vol. 54, pp. 23-36.

5. Bernheim, B. Douglass, 1984, "Strategic Deterrence of Sequential Entry into an Industry," *The RAND Journal of Economics*, Vol. 15, pp. 1-11.

6. Blume, Andreas, 1994, "Intraplay Communication in Repeated Games," *Games and Ecomomic Behavior*, Vol. 6, pp. 181-211.

7. Brock, W. and J. Scheinkman, 1985, "Price-Setting Supergames with Capacity Constraints," *Review of Economic Studies*, Vol. 52, pp. 371-82.

8. Davidson, Carl and Raymond Deneckere, 1990, “Excess Capacity and Collusion,” *International Economic Review*, Vol. 31, No. 3, pp. 521-41.

9. Davidson, Carl and Raymond Deneckere, 1987, “Long-run Competition in Capacity, Short-run Competition in Price and the Cournot Model,” *The Rand Journal of Economics*, Vol. 17, pp. 404-15.

10. Farrell, Joseph and Eric Maskin, 1989, “Renegotiation in Repeated Games,” *Games and Economic Behavior*, Vol. 1, pp. 327-60.

11. Friedman, James W., 1971, “A Non-Cooperative Equilibrium for Supergames,” *Review of Economic Studies*, Vol. 38, pp. 1-12.

12. Green, E. and Robert Porter, 1984, “Non-cooperative Collusion Under Imperfect Price Information,” *Econometrica*, Vol. 52, pp. 87-100.

13. Marshall, Robert C. and Michael J. Meurer, 1995, "Should Big-Rigging Always be an Antitrust Violation?", working paper.

14. McCutcheon, Barbara, 1997, “Do Meetings in Smoke-Filled Rooms Facilitate Collusion?,” *Journal of Political Economy*, Vol. 105, pp. 330-50.

15. Osborne, M. J. and C. Pitchik, “Price Competition in a Capacity Constrained Duopoly,” *Journal of Economic Theory*, Vol. 38, pp. 238-60.

16. Spagnolo, Giancarlo, 2000, “Self-Defeating Antitrust Laws: How Leniency Programs Solve Bertrand’s Paradox and Enforce Collusion in Auctions,” Stockholm School of Economics working paper.

Steven P. Schulenberg

Bates White, LLC
2001 K St. NW, Suite 700
Washington, DC 20006

Telephone: 202.354.1189
E-mail: steven.schulenberg@bateswhite.com

Education:

- Ph.D., Economics, The Pennsylvania State University, expected December 2003
- B.S.E., Electrical Engineering and Economics, Duke University, May 1997

Fields of Interest:

- Econometrics, Applied Microeconomics, Industrial Organization, Game Theory, Auctions and Procurements

Working Papers:

- “Numerical Analysis of Asymmetric Auctions with Optimal Reserve Prices," (with Robert C. Marshall)
- “The Impact of Synergies on Bidding in the Georgia School Milk Market," (with Robert C. Marshall, Matthew E. Raiff, and Jean-Francois Richard)
- “Deterrence of Collusion Through The Corporate Leniency Policy"
- “Collusion at Second Price Auctions with Synergies"

Professional Experience:

- Manager, Bates White, LLC, August 2002 to present
- Senior Consultant, Bates White, LLC, July 2001 to August 2002
- Instructor for Intermediate Microeconomics, The Pennsylvania State University, Spring 2000 to Spring 2001
- Research Assistant for Dr. Robert C. Marshall, The Pennsylvania State University, Summer 1997 to Summer 2001

Awards and Distinctions:

- Charles S. And Grace Lipson Graduate Research Fund
- Aftoora Scholarship Fund
- University Graduate Fellowship