Collusion and Optimal Reserve Prices in Repeated Procurement Auctions

Charles J. Thomas[1]

July 30, 2001

[1]Federal Trade Commission. I thank Dave Balan, Ken Chay, Stephen Holland, Dan Hosken, Dave Reiffen, and Bart Wilson for helpful comments. Send correspondence to Federal Trade Commission, 600 Pennsylvania Avenue, NW, Washington, DC 20580, or to cjthomas@ftc.gov. This work represents the opinions of the author and does not necessarily represent the position of the Federal Trade Commission or any individual Commissioner.

JEL: C72, D44, L1
Keywords: auction, procurement, collusion, optimal reserve prices

Abstract

Previous research on collusion in procurement markets uses static mechanism design theory to address the limitations on collusive activity imposed by asymmetric information, but in most instances it does not address how to enforce the proposed mechanisms. This paper uses repeated game theory to examine the sustainability of two commonly reported collusive schemes that have been identified as optimal static mechanisms. If a buyer does not select its reserve price strategically, or if its value is large relative to the sellers' costs, then collusion may be sustainable for a wide range of plausible discount factors. However, even mildly sophisticated reserve price selection can dramatically shrink the set of discount factors for which collusion can be sustained. These findings provide a rationale for existing arguments that buyers are vulnerable to collusion, but suggest that buyers possess tools that may profitably induce sellers to act competitively. The analysis also reveals that collusion tends to be more easily sustained if the sellers' costs have a low mean or a high variance, or, in some instances, if the number of sellers increases.

1 Introduction

While economists have suggested that collusion occurs in many procurement settings,[1] there are few analyses of it compared to studies of collusion in other strategic environments. In particular, relatively little is known or has been proposed about either the form that collusion takes, or the likelihood that it can be sustained. These are important limitations to arguments that collusion is a pervasive problem in auction markets, because what does seem clear is that colluding sellers must overcome several obstacles to sustain supracompetitive prices.

One difficulty facing a cartel in a procurement market is that in many instances its members cannot credibly reveal their production costs to one another. Hence, the cartel cannot directly determine which member can most profitably fulfill a particular contract, and therefore the cartel cannot obtain the maximum possible surplus from each transaction. Economists have developed theoretical models of auction environments to examine means by which cartels might attempt to appropriate rents from buyers.[2] The predominant approach taken has cartels construct collusive bidding schemes as solutions to mechanism design problems. In the mechanism design approach, each cartel member reports (possibly falsely) its privately known production cost to the mechanism, which then specifies the price the firm is to offer to the buyer. Because the asymmetric information about costs limits the sellers' ability to collude, the cartel's primary goal is to design optimal mechanisms that induce truthful revelation of its members' production costs.[3] A common feature of these mechanisms is rigid-pricing, in which the seller selected to win sets its price at a specific level regardless of its production costs, while other sellers either offer higher prices or do not submit price offers. Treatments vary in whether or not the participants are able to make transfer payments to non-winning cartel firms. This distinction is important, because the use of transfer payments likely increases the probability that competition authorities can prove the existence of price-fixing agreements.

Another difficulty facing a cartel in *any* market is that its members' adherence to the collusive arrangement must be enforced. In most instances, the mechanism design approach to collusion in procurement markets abstracts away from concerns about enforcing the proposed mechanisms.[4]

[1] For example, several allegations of collusion are reported in Scherer and Ross [1990], Porter and Zona [1993], Baldwin, Marshall, and Richard [1997], and McAfee and McMillan [1992].

[2] There exists a small, but growing, empirical literature that attempts to expose collusive behavior using bidding data from past auctions. For example, see Porter and Zona [1993] and Baldwin, Marshall, and Richard [1997]. Some of the empirical papers use the theoretical papers as a starting point for their econometric modeling.

[3] Note that in this instance it is the cartel members who design the mechanism, in contrast to the typical mechanism design setting in which the auctioneer is the designer.

[4] The exceptions are particular mechanisms that are enforceable without considering repeated interaction.

This is not a trivial concern, because sellers may have incentives either to set prices other than the ones specified by the mechanism, or to not make agreed-upon transfer payments. That is, while truthfully revealing costs is made incentive compatible, abiding by the mechanism's rules is not. For example, suppose that two firms report their costs of 3 and 4 and are told to offer prices of 6 and 7, respectively. The firm with cost 4 has incentives to set its price just below 6 in order to win the contract. As the mechanism design approach is static, it does not address the possibility or prevention of such deviations.

In this paper I use repeated game theory to evaluate the ability of sellers in first-price procurement auctions to enforce adherence to two commonly reported rigid-pricing collusive schemes identified as being optimal static mechanisms.[5] Specifically, I determine the lowest value that sellers can place on future profits such that the short-term gain from defection is outweighed by the long-term loss from rivals' retaliation to that defection. The lower is the critical discount factor for which collusion can be sustained, the more likely it is that the sellers' discount factor in a particular market is high enough to sustain the collusive agreement.

The critical discount factor for which collusion can be sustained depends in a complicated manner on the number of firms, the distribution of their production costs, and the reserve price. Consequently, its theoretical determination offers limited insight about either the ease of sustaining collusion, or how sustaining collusion is affected by changes in the strategic environment. To more clearly identify the conditions under which collusion can occur in procurement markets, I therefore compute critical discount factors for a large number of cost distributions and market concentration levels, under various assumptions about the buyer's reserve price strategy. Next I fit to the generated data a simple function of the strategic environment's primitives. This technique usefully summarizes the data and permits straightforward numerical evaluation of comparative statics.

The computational results reveal that the reserve price imposed by the buyer has a substantial impact on the sellers' ability to collude, which is an interesting element that is absent from standard collusion models in which non-strategic buyers are embodied in the market demand curve. If the buyer cannot or does not credibly impose a reserve price, or simply has a large value relative to the sellers' costs, then collusion likely can be sustained for a wide range of discount factors. In contrast, if the buyer can credibly impose a strategically selected reserve price, then the set of discount factors for which collusion can be sustained may be quite small. With a well-chosen reserve price, say at

[5] In a first-price procurement auction, the firm offering the lowest price is designated the winner and is paid the price it offered.

or below the level that is optimal versus static Nash equilibrium price-setting, the critical discount factor for sustaining collusion is high because the short-term gain from defection is large relative to the per-period loss from rivals' retaliation to that defection. Hence, firms must highly value future profits in order for the loss of the foregone future collusive profits to outweigh the short-term gain from defection.

The computational results also reveal how the number of sellers and their cost distributions affect the sellers' ability to collude. First, in some instances collusion is more easily sustained as the number of sellers increases. This results from a widening of the gap between the per-period cooperative and noncooperative expected profits as the number of sellers increases from low levels. Second, collusion appears to be more easily sustained when sellers have low expected costs, all else equal. This result likely emerges because sellers with high expected costs have little to gain from collusion, which exacerbates a seller's temptation to deviate from the collusive agreement when it has unexpectedly low costs. Third, collusion appears to be more easily sustained when the sellers' costs are more variable.

To conclude the analysis, I evaluate the sustainability of collusion by comparing the computed critical discount factors to plausible values of firms' actual discount factors, which helps one to assess the economic significance of the static mechanism design approach to collusion in auctions. The comparison suggests that firms using the optimal static mechanism are unlikely to be able to collude against sophisticated buyers. However, unsophisticated buyers may be vulnerable to collusion.

The following three papers are representative of the mechanism design approach. McAfee and McMillan [1992] identify optimal static mechanisms in single-object first-price auctions with an exogenously determined reserve price and revelation of the winning bid, both with and without the ability to make transfer payments. Announcing the winning bid makes it immediately clear if a winning firm has deviated from its prescribed bid. The authors do not model the enforcement device, but appeal to the repeated game approach I examine in this paper. Without transfer payments, they find that the optimal mechanism specifies a price offer at the reserve price by all sellers whose production costs are below the reserve price. With transfer payments, they find that the low-cost seller wins at the reserve price and pays an equal amount to each losing seller. I examine both of these mechanisms in this paper.

Graham and Marshall [1987] identify static mechanisms in single-object second-price and English auctions with an endogenously determined reserve price. Second-price and English auctions have the attractive feature that, once costs are truthfully announced, no firm can achieve a short-term gain by deviating from the collusive scheme. Therefore, there is no need to consider repeated interaction to enforce adherence to the mechanism, because, even in a static setting, abiding by the mechanism is incentive compatible. Consequently, the buyer cannot select the reserve price to deter collusion, in contrast to my finding in first-price auctions.[6]

Cramton and Palfrey [1990] examine collusive mechanisms in a quantity-setting oligopoly model, and they explicitly consider outside options that might arise if the sellers cannot agree on a mechanism. Like McAfee and McMillan [1992], they do not formally address the incentive to deviate from the mechanism, but appeal to repeated game concepts to justify the enforceability of the mechanism's quantity prescription after the sellers reveal their costs.

A new approach is taken by Athey, Bagwell, and Sanchirico [2000], who use the same basic structure as in the present paper, with the exceptions that all sellers' price offers become public information and that sellers cannot make explicit transfer payments. The authors show that implementing the optimal static mechanism without transfer payments is the optimal collusive device in the repeated game, provided that the sellers have sufficiently high discount factors. When sellers' discount factors are below the critical level, the authors partially characterize more profitable collusive mechanisms that do not use rigid-pricing. They do not address the two central goals of this paper, a systematic appraisal of the discount factors for which rigid-price collusion can occur, and the effect on such collusion of the buyer's reserve price selection and the sellers' cost distributions.

The rest of this paper is structured as follows. Section 2 presents a standard one-shot procurement auction and states known results relevant to the succeeding analysis. Section 3 uses the one-shot procurement auction as the stage game of an infinitely repeated game, and determines conditions necessary for the sustainability of two collusive schemes. Section 4 uses specific assumptions on the number of firms and their cost distributions to numerically determine how patient the sellers must be to successfully employ the proposed collusive schemes, under various assumptions about the buyer's strategic sophistication in selecting the reserve price. Section 5 briefly concludes,

[6] The susceptibility to collusion of English and second-price auctions is offered as one reason why first-price auctions are more frequently employed in procurement settings. However, a countervailing factor is that the identified mechanism in English and second-price auctions requires the sellers to communicate about each contract, which increases the risk of successful antitrust prosecution. Therefore, a buyer's use of the second-price format might make it easier to *prove* the existence of price-fixing agreements, even though it simultaneously increases the likelihood of the sellers' implementing an agreement. If the sellers chose to forego the communication necessary to implement collusion using the second-price mechanism, then they could employ the mechanism identified by McAfee and McMillan [1992]. However, that mechanism's specified bidding behavior would lead to outcomes identical to those in a first-price auction.

while the Appendix contains all proofs.

2 The One-Shot Procurement Auction

A prospective buyer of a product solicits price offers simultaneously from each of N sellers. Prior to the making of offers to the buyer, each seller i draws its production cost, c_i, independently from the cumulative distribution $F(c)$. Assume that F has a differentiable density f with support $[\underline{c}, \overline{c}]$. The buyer purchases from the low-priced seller at the offered price. In auction terminology, this is a symmetric independent private value (IPV) first-price auction. Assume further that the buyer and sellers all are risk neutral, the number of firms is exogenous, entry is blockaded, and it is costless for sellers both to learn their production cost and to participate in the procurement process. The buyer's next best supply alternative costs $c_B \geq \overline{c}$, and the buyer's profit from purchasing from one of the N sellers at price p is $c_B - p$. Seller i's profit from winning with price p_i is $p_i - c_i$. Prior to the submission of offers, the buyer imposes a commonly known reserve price, r, that is less than or equal to $\overline{c}$, and above which price offers will be rejected.[7] Following most of the auction literature, I assume that the buyer can commit to the procurement format.

The one-shot procurement process has been extensively discussed elsewhere, so here I simply state relevant known results. Proofs can be found in, for example, McAfee and McMillan [1992].[8] The first result describes the sellers' static Nash equilibrium price-setting strategies and expected profits, while the second describes the buyer's expected profit and optimal reserve price versus static Nash equilibrium price-setting.

Result 1 *In the symmetric IPV one-shot procurement auction with reserve price $r \leq \overline{c}$, for $c \in [\underline{c}, r]$ the sellers' static Nash equilibrium price-setting function is*

$$p^{NE}(c|r) = c + \int_c^r \left[\frac{1 - F(s)}{1 - F(c)}\right]^{N-1} ds.$$

A seller's interim expected profit in the static Nash equilibrium, when it has drawn cost $c \leq r$, is

$$\overline{\pi}_S^{NE}(c|r) = \int_c^r [1 - F(s)]^{N-1} \, ds,$$

[7] Technically speaking, the buyer might select a reserve price that exceeds $\overline{c}$, but in the symmetric equilibrium sellers will never set a price greater than $\overline{c}$.

[8] While their proofs all are for auctions where bidders are buyers rather than sellers, all of the proofs have analogues in the procurement model.

and its ex ante expected profit in the static Nash equilibrium, when it has not yet drawn its cost, is

$$\overline{\pi}_S^{NE}(r) = \int_{\underline{c}}^{r} F(c)\left[1 - F(c)\right]^{N-1} dc.$$

Result 1 illustrates that sellers shade their price above their cost by an amount that depends on the reserve price and the number of rivals. As the reserve price increases, the amount of shading and expected profits increase. As the number of rivals increases, the amount of shading and expected profits decrease.

Result 2 *In the symmetric IPV one-shot procurement auction, if the sellers employ their static Nash equilibrium price-setting strategies, then the buyer's expected profit from imposing reserve price r is*

$$\overline{\pi}_B^{NE}(r) = \int_{\underline{c}}^{r} \left[c_B - p^{NE}(c|r)\right] N \left[1 - F(c)\right]^{N-1} f(c) dc.$$

The buyer's optimal reserve price versus static Nash equilibrium price-setting, r^{NE}, solves

$$c_B = r^{NE} + \frac{F\left(r^{NE}\right)}{f\left(r^{NE}\right)},$$

provided the resulting $r^{NE} \leq \overline{c}$. Otherwise, $r^{NE} = \overline{c}$, and the reserve price is said to be non-binding.

The reserve price presented in Result 2 has some interesting features. First, the buyer's optimal reserve price versus static Nash equilibrium price-setting is independent of the number of sellers. Second, for a certain range of reserve prices, both the sellers' and the buyer's expected profit increase as r increases. This occurs because social welfare increases as the probability of a sale increases, so parties on both sides of the transaction can gain as r increases. In many instances, the buyer's expected profit eventually decreases in the reserve price, leading to a binding optimal reserve price strictly less than $\overline{c}$. However, if c_B is sufficiently large, then the buyer optimally sets r at $\overline{c}$, because its outside option is so costly relative to the most costly seller. Therefore, the cost to the buyer of possibly losing a purchase from one of the sellers outweighs the gain from leading the sellers to set their prices more aggressively.

3 The Repeated Procurement Auction

Suppose that the same buyer and sellers play the preceding one-shot game in each of an infinite number of discrete periods, with the sellers' costs redrawn each period and with future profits discounted by the factor $\delta \in [0, 1)$. While employing the static Nash equilibrium price-setting strategies each period is a subgame perfect equilibrium, of greater interest is the characterization of subgame perfect equilibria yielding the sellers higher expected profits. In what follows, I assume that the buyer does not reveal the winning price, as is common in private sector procurement.[9] Additionally, I assume that the winner in each period is revealed, and that the buyer commits at the outset to a reserve price to be used throughout the game.[10] Subject to the reserve price, the sellers are presumed to select the most collusive outcome possible.

I examine two rigid-pricing collusive schemes in which the seller selected to win the contract offers a price equal to the reserve price. Both schemes are all-inclusive, in the sense that all sellers are members of the cartel. The sellers not selected to win the current contract either offer higher prices or do not submit price offers. The first scheme assumes that the sellers cannot credibly reveal their costs to each other prior to submitting price offers, and that they cannot or do not make side payments. It uses a form of bid rotation in which the randomly selected winner has the option of declining its position. This scheme requires an initial meeting for the participants to agree on the signals to be employed, but this is not too troubling from the perspective of applications. As McCutcheon [1997] argues, existing competition policy might enhance collusion by promoting initial meetings but not subsequent ones. This bid rotation scheme is identical, in terms of the sellers' expected profit, to the optimal mechanism without transfer payments identified by McAfee and McMillan [1992]. Following their nomenclature, I refer to a cartel using this scheme as a **weak cartel**. Scherer and Ross [1990] report that a version of this scheme was used in the electrical manufacturers' conspiracy. While bid rotation is unappealing because the low-cost seller does not necessarily win the contract, McAfee and McMillan [1992] prove that earning supracompetitive profits via their mechanism is incompatible with the efficient allocation of contracts.

The second scheme assumes that the sellers cannot credibly reveal their costs to each other prior to submitting price offers, but that they can make side payments. McAfee and McMillan [1992] illustrate that the sellers can implement this scheme by holding their own first-price auction

[9] With slight modification, all of the results in this section hold if the buyer reveals the winning price, as is common in public sector procurement.

[10] Perhaps surprisingly, the buyer can more successfully deter collusion if he cannot commit to a single reserve price at the outset of play. Specifically, if the sellers revert to static Nash equilibrium price-setting following a deviation, and if the buyer switches from the optimal reserve price versus collusive price-setting to the optimal reserve price versus static Nash equilibrium price-setting, then the sellers' per-period ex ante expected profits will be higher than if the buyer had not changed the reserve price. This higher profit level in the punishment regime makes collusion more difficult to sustain. In fact, in some instances the sellers' ex ante expected profit in the punishment phase exceeds their ex ante expected profit in the collusive phase, implying that collusion cannot be sustained for any discount factor. Therefore, assuming a one-time commitment to the reserve price provides a bound on the sellers' ability to sustain collusion.

prior to the buyer's first-price auction. They refer to a cartel using this scheme as a **strong cartel**. Strong cartels require more frequent contact than do weak cartels, which increases the likelihood that the antitrust authorities may discover and be able to prove the existence of the agreement. However, this concern is balanced against the fact that the strong cartel finds collusion easier to sustain than does the weak cartel.[11]

Four potential stumbling blocks to successful collusion are the selection of the winning firm, the selection of a price, the detection of deviations, and the punishment of deviations. The schemes used by the weak cartel (WC) and the strong cartel (SC) clearly overcome the first three. The winner and the price are predetermined, and deviation can be recognized after the current round of offers and dealt with in the next period. To punish deviations, I employ perpetual reversion to the static Nash equilibrium. Though there exist more severe punishments, imposing them suffers from difficulties similar to those preventing efficient collusion. For example, suppose all firms $i \neq j$ are supposed to punish firm j in each period by setting $p_i = c_i$. Because no prices are revealed and the sellers' costs are private information, a punishing firm could deviate to a price $\widetilde{p}_i > c_i$ without being detected if it lost when it should have won, and earn a positive expected profit. Thus, it seems such a severe punishment could not be implemented credibly. In contrast, by definition Nash reversion can be implemented credibly.

Collusion can be sustained provided that the profit from colluding exceeds the profit from defecting today and facing punishment in the future, for any current cost realization. It is important that the incentive constraints be satisfied for all possible cost draws. If they were not, then the positive probability of the designated seller's being undercut by sellers with particular cost realizations will prompt the designated seller to undercut as well. Such a response precludes the existence of an equilibrium in which the selected seller offers a price of r. I define the critical discount factor, $\overline{\delta}^X(r)$, to be the smallest discount factor for which the incentive constraints are satisfied for all possible costs, for $X \in \{WC, SC\}$ and for reserve price r.

3.1 Weak Cartels

Suppose that the sellers do not know each other's costs before prices are offered and cannot make side payments, but that they are able to engage in some signaling after drawing their costs. Specifically, the sellers construct an ordering, say alphabetically by name, then use the outcome of

[11] Perhaps surprisingly, the strong cartel scheme is identical, in terms of expected profit, to a rigid-pricing scheme in which the sellers credibly can reveal their costs to each other prior to submitting price offers, without a need to meet or otherwise explicitly communicate, but cannot make side payments. Therefore, the theoretical and computational results regarding strong cartels apply equally well to the credible revelation setting.

a public randomization device to select a winner for each contract as it is offered. The designated winner competes unopposed and offers a price equal to the reserve price. The seller that is selected through the randomization process is able to costlessly decline the invitation to be the winning firm, which it will do if its cost exceeds the reserve price. The next seller on the list is then designated to be the winning firm, it also has the option of declining, and so on. The remaining sellers either do not submit price offers, or submit price offers that exceed the reserve price. This is a form of bid rotation, though the designated seller is determined randomly each period rather than selected sequentially from a predetermined ordering of the firms.

Lemma 1 in the Appendix proves that the outcomes using bid rotation are identical in expectation to those using the optimal mechanism, without transfer payments, determined by McAfee and McMillan [1992]. Their mechanism specifies that each seller with a cost less than the reserve price submits a price equal to the reserve price. Other sellers either do not bid or set a price equal to their cost.[12] However, for several reasons the bid rotation scheme I have constructed may be easier to implement than the optimal mechanism. First, while bid rotation does not require that the winning price be revealed, enforcing the optimal mechanism is difficult unless the buyer announces the winning price; otherwise, sellers can deviate by slightly undercutting the reserve price. Such a deviation cannot be detected immediately and dealt with in the next period, which impedes collusion. Second, bid rotation can be helpful in avoiding investigation by the antitrust authorities, as opposed to other schemes, such as the optimal mechanism, that have several firms submitting the same price offer.[13] Third, bid rotation is immune to the buyer's disrupting the scheme, say by using a non-random tie-breaking procedure.[14]

For a weak cartel to sustain supracompetitive prices in an auction, a seller with cost c designated to lose must not prefer to slightly undercut the designated winner's price and precipitate perpetual reversion to static Nash equilibrium price-setting. Therefore, it must be the case that for all c, for a given reserve price r,

$$\left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{WC}(r) \geq (r-c) + \left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{NE}(r). \tag{1}$$

[12]On a technical note, a weak cartel can profitably implement the optimal mechanism only if $[1-F(v)]/f(v)$ decreases in v. For a procurement auction, the equivalent condition is $-F(c)/f(c)$ decreases in c. This condition is equivalent to requiring that F be log-concave, a restriction that is employed in the simulations.

[13]However, this scheme still suffers from the weakness that the winning firm in each period sets a price of r. Under noncooperative bidding, that outcome occurs with probability zero. Despite this result, there exist many procurement settings in which the government has received several identical bids. See Scherer [1967], Comanor and Schankerman [1976], and McAfee and McMillan [1992] for details.

[14]See McAfee and McMillan [1992] for elaboration on this point.

$\overline{\pi}_S^{WC}(r)$ denotes a seller's ex ante expected profit per-period from participating in the weak cartel when the reserve price is r.[15] The left hand side of (1) is the net present value of future expected collusive profit, while the right hand side is the short-term gain from defecting plus the net present value of future expected static Nash equilibrium profit.

The incentive constraint can be rearranged to solve for the discount factor, $\overline{\delta}^{WC}(c,r)$, above which defection is prevented by a firm with cost c, yielding

$$\overline{\delta}^{WC}(c,r) = \frac{(r-c)}{\overline{\pi}_S^{WC}(r) - \overline{\pi}_S^{NE}(r) + (r-c)}. \tag{2}$$

As $\overline{\pi}_S^{WC}(r)$ and $\overline{\pi}_S^{NE}(r)$ are independent of a firm's current cost, it is straightforward to show that $\overline{\delta}^{WC}(c,r)$ is highest for a firm with the lowest cost. Therefore, $\overline{\delta}^{WC}(r) \equiv \overline{\delta}^{WC}(\underline{c},r)$ is the critical discount factor. It increases as the short-term gain from defection $(r-\underline{c})$ increases, and it increases as the per-period loss from defection $\left(\overline{\pi}_S^{WC}(r) - \overline{\pi}_S^{NE}(r)\right)$ decreases.

Proposition 1 *In a weak cartel using the symmetric bid rotation scheme with reserve price r, a seller's ex ante expected profit from participating is*

$$\overline{\pi}_S^{WC}(r) = \frac{1-[1-F(r)]^N}{NF(r)} \int_{\underline{c}}^{r} F(c)dc.$$

Defection by a seller with the lowest possible cost can be deterred if and only if $\delta \geq \overline{\delta}^{WC}(r)$, where

$$\overline{\delta}^{WC}(r) \equiv \frac{N(r-\underline{c})}{\int_{\underline{c}}^{r} F(c)\left[\frac{1-[1-F(r)]^N}{F(r)} - N[1-F(c)]^{N-1}\right]dc + N(r-\underline{c})}.$$

$\overline{\delta}^{WC}(r)$ may either increase or decrease as r increases.

3.2 Strong Cartels

Suppose that the sellers do not know each other's costs before prices are offered, but that they are able to make side payments to each other. McAfee and McMillan [1992] show that in the optimal scheme, the low-cost seller sets its price at the reserve price, pays each other seller an equal

[15] Note that the randomness in the assignment process creates for a firm incentives that are unaffected by whether the firm has recently been the selected winner of another project. This greatly simplifies the analysis by reducing the number of constraints that must be satisfied.

fraction of the cartel's surplus, $\frac{r-T(c)}{N-1}$, and receives $T(c)$, where

$$T(c) = r - \frac{\int_c^r (r-s)(N-1)[1-F(s)]^{N-1} f(s)ds}{[1-F(c)]^N}.$$

Thus, a firm can receive a payment even if its cost exceeds the reserve price. The side payment scheme can be implemented by the sellers' holding their own first-price auction prior to the buyer's first-price auction, which Graham and Marshall [1987] refer to as a pre-auction knockout (PAKT). I assume that the strong cartel attempts to collude in this fashion, with the seller winning the PAKT making transfer payments to the losing sellers before the buyer's procurement auction.

Lemma 2 in the Appendix proves that a seller's expected profit from using the strong cartel scheme is identical to its expected profit from using a rigid-price scheme in a setting in which the sellers can credibly reveal their costs to one another but cannot make side payments. The latter scheme allocates contracts efficiently and avoids the asymmetric information that prevents the cartel members from knowing which seller can most efficiently fulfill a particular contract. Considering the setting in which costs can be credibly revealed does more than provide a baseline for determining how asymmetric information across sellers affects their ability to collude. While the assumption of credible cost revelation is challenged in the mechanism design approach, it has been used in the literature on information sharing in oligopoly.[16] Moreover, this assumption may be valid in certain settings, such as when transport costs are relatively high and colluding firms create territorial restrictions. For example, in the Addyston Pipe conspiracy[17] and various conspiracies in lumber contracting,[18] the designated winners were local firms and presumably were the low-cost providers.

It is somewhat more complicated for a strong cartel to sustain collusion than it is for a weak cartel, because with the strong cartel's side payment scheme there are more instances in which a seller potentially has a profitable deviation. As in a weak cartel, in a strong cartel a seller with cost c that is designated to lose the buyer's auction must be deterred from slightly undercutting the reserve price to beat the designated winner. That is, it must be the case that

$$\left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{SC}(r) \geq (r-c) + \left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{NE}(r). \tag{3}$$

[16] See Gal-Or [1986], Shapiro [1986], and Vives [1984].

[17] *Addyston Pipe and Steel Company v. United States*, 175 U.S. 211 (1899).

[18] See Baldwin, Marshall, and Richard [1997] for details.

Note that the side payment does not come into play in (3), because at the moment the losing seller is deciding whether or not to deviate from the collusive arrangement, it already has received the side payment.

The low-cost seller with cost c in the current period, who should win the PAKT, must be deterred from the following two deviations. First, after the PAKT but before the buyer's auction, the low-cost seller may elect not to make the agreed-upon side payments. As the PAKT reveals all sellers' costs, the low-cost seller knows precisely how much profit it will earn in the buyer's auction, assuming that the deviating seller will win by an arbitrarily small amount less than the second-lowest cost. Regardless of the low-cost seller's cost, the profit from this deviation is maximized when all other sellers' costs exceed r. Thus, to prevent such a deviation it must be the case that

$$T(c) + \left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{SC}(r) \geq (r-c) + \left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{NE}(r). \tag{4}$$

Second, the low-cost seller may elect to overstate its cost in the PAKT, receive a side payment, and then slightly undercut the designated winner in the buyer's auction. To prevent such a deviation, it must be the case that

$$T(c) + \left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{SC}(r) \geq \frac{E\left[r - T(c_{(2)})|c_{(1)} = c\right]}{N-1} + (r-c) + \left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{NE}(r), \tag{5}$$

where $c_{(i)}$ denotes the ith lowest order statistic from the cost distribution $F(c)$.

By themselves, (4) and (5) look like they may cause substantial difficulty in calculating the critical discount factor, relative to the calculation with a weak cartel. However, as $T(c) > 0$, then (4) holds whenever (3) holds. Therefore, (4) is superfluous. Similarly, if (3) holds, then the only way that (5) can fail to hold is if

$$T(c) < \frac{E\left[r - T(c_{(2)})|c_{(1)} = c\right]}{N-1}.$$

However, because bidding truthfully is made incentive compatible in the optimal static mechanism, a seller always prefers winning and making side payments over losing and receiving a side payment when it should have won. Therefore, the preceding inequality cannot hold, and (5) holds whenever (3) holds. Consequently, the only incentive constraint that requires checking for the strong cartel is (3), which can be rearranged to yield

$$\overline{\delta}^{SC}(c,r) = \frac{(r-c)}{\overline{\pi}_S^{SC}(r) - \overline{\pi}_S^{NE}(r) + (r-c)}. \tag{6}$$

(6) is quite simple and is analogous to the single incentive constraint that must be satisfied for the weak cartel to be able to sustain collusion. As was determined for the weak cartel, $\overline{\delta}^{SC}(c,r)$ is highest for a seller with the lowest cost, $\underline{c}$. Hence, the critical discount factor is $\overline{\delta}^{SC}(r) \equiv \overline{\delta}^{SC}(\underline{c},r)$.

Proposition 2 *In a strong cartel using the symmetric side payment scheme with reserve price r, a seller's ex ante expected profit from participating is*

$$\overline{\pi}_S^{SC}(r) = \frac{r-\underline{c}}{N} - \frac{1}{N}\int_{\underline{c}}^{r} [1-F(c)]^N \, dc.$$

Defection by a seller with the lowest possible cost can be deterred if and only if $\delta \geq \overline{\delta}^{SC}(r)$, where

$$\overline{\delta}^{SC}(r) \equiv \frac{N(r-\underline{c})}{(N+1)(r-\underline{c}) - \int_{\underline{c}}^{r} [(N-1)F(s)+1][1-F(s)]^{N-1} ds}.$$

Collusion is more easily sustained as the reserve price increases. That is, $\overline{\delta}^{SC}(r)$ decreases as r increases.

Note that expected profits in the punishment regimes of the schemes used by the weak and strong cartels are identical. Therefore, Propositions 1 and 2 indicate that the difference between the critical discount factors for weak and strong cartels is caused by the difference in expected collusive profits. It can be shown that expected profits are larger for the strong cartel. Hence, for a given reserve price a strong cartel's critical discount factor is always lower than a weak cartel's. However, the buyer's optimal reserve price may depend upon the type of cartel faced, which in turn will affect the critical discount factor. The buyer's response to the potential for collusive behavior is the topic discussed next.

3.3 The Buyer's Response

In the analysis so far, the reserve price has been taken as given. However, if the buyer has discretion in setting the reserve price, then the buyer may be able to substantially affect the sellers' ability to collude.

For any reserve price, the buyer's expected profit is identical when facing any two rigid-price collusive schemes in which the seller selected to win sets its price at the reserve price, provided that its cost is less than the reserve price. Thus, the optimal reserve price versus such collusive price-setting is the same regardless of the collusive scheme.

Result 3 *Facing the collusive scheme X in which a single seller offers the reserve price, the buyer pays the reserve price if at least one seller's cost is below the reserve price. Thus, the buyer's expected profit as a function of the reserve price, r, is*

$$\overline{\pi}_B^X(r) = (c_B - r)\left(1 - [1 - F(r)]^N\right).$$

The buyer's optimal reserve price, r^X, solves

$$c_B = r^X + \frac{1 - \left[1 - F(r^X)\right]^N}{N\left[1 - F(r^X)\right]^{N-1} f(r^X)},$$

provided that $r^X \leq \overline{c}$. Otherwise, $r^X = \overline{c}$ and the reserve price is non-binding. If $N \geq 2$, then the reserve price is always binding. Moreover, the optimal reserve price weakly decreases as N increases.

When the N sellers collude by having the firm selected to win set its price at the reserve price, the buyer essentially is facing a single seller whose costs are drawn from the distribution $G(r) = 1 - [1 - F(r)]^N$. $G(r)$ is the distribution of the minimum of N draws from the distribution $F(r)$. Thus, Result 3's formula for the optimal reserve price in essence is identical to the one presented in Result 2. Moreover, because the buyer's optimal reserve price versus the collusive scheme decreases in the number of sellers, the buyer's optimal reserve price when facing such a collusive scheme is less than the optimal reserve price when facing static Nash equilibrium price-setting.[19]

While Result 3 determines the buyer's optimal reserve price when it knows it is facing a rigid-pricing scheme by a strong or a weak cartel, that reserve price is not necessarily the one that would be selected by a buyer concerned about collusion. Instead, the buyer may be able to deter collusion by selecting a different reserve price, which in turn affects the critical discount factor.

I focus on three reserve prices available to the buyer. First, the **Nash reserve price**, presented in Result 2, is the reserve price that is optimal versus static Nash equilibrium price-setting. Second, the **naive reserve price**, equal to the upper support of the cost distribution ($\overline{c}$), is the reserve price

[19]This result follows because the optimal reserve price versus static Nash equilibrium price-setting is identical for all N. Thus, the optimal reserve price versus collusive price-setting with $N = 1$ is the same as the optimal reserve price versus static Nash equilibrium price-setting. The argument is completed by using the fact that the optimal reserve price versus collusive price-setting decreases as N increases.

that would be set if the buyer were unsophisticated, could not credibly impose a lower reserve price, or simply had a large value relative to the sellers' costs.[20] Each of these reserve prices determines a critical discount factor that the sellers' discount factor must exceed in order for collusion to be sustainable.

Third, the **minimal deterrent reserve price** is the lowest reserve price for which the buyer's expected profit versus static Nash equilibrium price-setting, using that reserve price, is equal to the buyer's expected profit versus collusive price-setting, using the optimal reserve price versus collusive price-setting (r^X from Result 3). As the buyer always can let the sellers collude and earn $\overline{\pi}_B^X(r^X)$, for $X \in \{WC, SC\}$, then the buyer will never try so hard to induce static Nash equilibrium price-setting that it actually earns less than that amount. The minimal deterrent reserve price determines the **critical deterrent discount factor**, which is the lowest discount factor the sellers can have in order to sustain collusion against a strategic buyer.

Because for any reserve price the buyer's expected profit versus static Nash equilibrium price-setting exceeds its expected profit versus collusive price-setting, there always exists a set of discount factors for which the buyer finds it profitable to deter collusion through its selection of the reserve price. The differences the buyer can create in the critical discount factor may not appear to be substantial, say if the discount factors are distributed uniformly over the interval [0,1]. However, if the true discount factors are distributed uniformly over [0.95,1], then increasing the critical discount factor from, say, 0.98 to 0.991 decreases the probability that the sellers can collude by 55 percent.[21] The next section examines more carefully the extent of the buyer's ability to deter collusion through its selection of the reserve price.

4 Simulation Results

The theoretical results in Section 3 analytically determine the minimal discount factor for which collusion is sustainable, but they do not provide much insight into the actual value of the critical discount factor, or how that value is affected by the number of sellers, their cost distributions, or the level of the reserve price. These quantitative issues are important for several reasons. First, concern about collusion via these mechanisms is warranted only if the critical discount factors

[20] I am being somewhat imprecise. If the buyer's value exceeded the sellers' maximum cost, and the buyer could not credibly impose a reserve price below its value, then the buyer would be even more susceptible to collusion than the results I present will show. This is the case because, while colluding, the designated winning seller will offer a price equal to $c_B > \overline{c}$, but if punishment is entered into, then the highest price set will be $\overline{c}$.

[21] There exists a 40 percent chance the sellers' discount factor is greater than 0.98, and an 18 percent chance it is greater than 0.991. Hence, increasing the critical discount factor decreases the probability the sellers can collude by $(40 - 18)/40 = 55$ percent.

correspond to discount factors that firms might actually use to value future profits. Second, it is useful to know how the critical discount factor changes as a function of the primitives of the strategic environment. For example, if there exist conditions under which collusion is more easily sustained as the number of sellers increases, then under those conditions the buyer may wish to restrict the number of sellers from which it accepts price offers. Similarly, if collusion is more easily sustained as the variance of the sellers' costs falls, then there may exist incentives for the sellers to adopt production technologies that reduce that variance.

Researchers have turned to computation to obtain both qualitative and quantitative insights from models that are analytically cumbersome or intractable. For example, Quirmbach [1993] provides computational evidence on the relationship between the incentives to perform R&D and the extent of post-innovation collusive behavior. He finds that while the social welfare ranking of Bertrand and Cournot competition can vary, both tend to outperform perfect collusion. Prior to this finding, it was unclear whether the prospect of collusive profits tended to lead to higher welfare by promoting more intensive R&D. Haubrich [1994] provides computational evidence that helps bridge the gap between the theory and the reality of executive compensation. He finds that empirical regularities of compensation packages are consistent with reasonable risk aversion parameters in standard principal-agent models, despite the widely held preexisting belief that such models could not realistically explain observed compensation schemes.

This section complements Section 3's theoretical results by computing critical discount factors for a varying number of sellers and for a wide range of specific functional forms of the sellers' cost distributions. For the sake of space, I compute critical discount factors only when $c_B = \overline{c}$, for which the Nash reserve price is binding. I evaluate the data by using a second-order polynomial to approximate the nonlinear function that determines the critical discount factor. This technique usefully summarizes how the critical discount factor is influenced by the number of firms (N) and by the mean (μ) and variance (σ^2) of the sellers' cost distribution, three parameters whose effect one typically might wish to evaluate.

I compute the discount factor for each (N, μ, σ^2) triplet in a grid of the three-dimensional parameter space. The number of firms varies from 2 to 30. The cost distributions all are from the family of Beta distributions, which has support $[0, 1]$, includes the uniform distribution as a special case, and varies smoothly in its strictly positive parameters p and q.[22] By varying the parameters, the Beta distribution can take on a variety of shapes and can be made to look like many other

[22] Moving the support to $[X, X + 1]$, while increasing c_B by X, has no effect on the quantitative results to follow.

distributions. The mean and the variance uniquely determine the two Beta parameters, and hence uniquely determine the distribution. One can show that in order for the Beta parameters to be strictly positive, the mean and the variance must satisfy $\mu(1-\mu)-\sigma^2>0$. Moreover, I further restrict the mean and variance so that the Beta distribution is log-concave.[23] This condition is required for the optimal mechanism without transfer payments to yield supracompetitive profits.[24] The permissible set of means and variances is illustrated in Figure 1, with a grid size of 0.01. Subject to the parameter restrictions and the grid size, I compute the critical discount factor for 478 distributions. Given that the number of firms varies from 2 to 30, for each combination of the two cartel schemes and the three reserve price strategies, I generate 13,862 critical discount factors (478 distributions$\times$29 numbers of firms).

Figure 1 Here

I use the least squares criterion to fit the computed discount factors to the (mostly) second-order polynomial equation

$$\delta=\alpha_0+\alpha_1 N+\alpha_2\mu+\alpha_3\sigma^2+\alpha_4 N^2+\alpha_5 N\mu+\alpha_6 N\sigma^2+\alpha_7\mu^2+\alpha_8\mu\sigma^2+\alpha_9\sigma^4+\alpha_{10}N^3+\alpha_{11}N^4.$$

I add the N^3 and N^4 terms to the otherwise second-order polynomial because some of the generated data have a pattern requiring fourth-order terms. The results of the regressions for each combination of the two cartel schemes and the three reserve price strategies are reported in Table 1. As is evident from the adjusted R^2 values, the approximations provide a very good fit of the generated data. The importance of the second-order terms is evident from the reduction in the adjusted R^2 measures of on average about 0.245 points when only first-order terms are used.

Table 1 Here

The worst fits predictably involves the deterrent reserve price strategy. With the naive reserve price strategy, the reserve price does not change with N, μ, or σ^2. With the Nash reserve price strategy, the reserve price changes with μ and σ^2, but not with N. Because the critical discount factor is a function of the reserve price and of N, μ, and σ^2, there should exist greater variability

[23]One can show that the Beta distribution is log-concave provided that $q\geq 1$. The proof involves solving for the conditions under which either the distribution itself is log-concave, or the density is log-concave (which implies the distribution is log-concave).

[24]See Theorem 1 in McAfee and McMillan [1992].

in the critical discount factors with the Nash reserve price strategy than with the naive reserve price strategy. With the deterrent reserve price strategy, the reserve price changes with N, μ, and σ^2. Hence, the critical discount factors versus this strategy should exhibit the most variation, and consequently the fit should be worse than with the other two reserve price strategies.

The following three examples use the regression results in Table 1 to show how the critical discount factor is affected by changing the three parameters. Figure 2 illustrates the effect of changing the mean of the sellers' cost distribution for a strong duopoly cartel whose cost distribution's variance equals the average level of the variance in the simulations ($\overline{\sigma}^2 = 0.0382$). The evaluated mean costs are those from the parameter-grid that induce log-concavity of the cost distribution, given the assumed variance. For each reserve price strategy in this example, collusion becomes more difficult to sustain as the distribution's mean increases. This effect is so pronounced that, against a buyer using the deterrent reserve price strategy (panel (c)), collusion is extremely difficult to sustain once the mean cost exceeds 0.6. In fact, once the duopolists' expected costs exceed 0.5, their discount factor must be at least 0.983 to possibly be able to sustain collusion versus the deterrent reserve price strategy. In contrast, for the same range of expected costs the duopolists' discount factor can be as low as 0.93 to sustain collusion versus the naive reserve price strategy.

Figure 2 Here

Figure 3 illustrates the effect of changing the variance of the sellers' cost distribution for a strong duopoly cartel whose cost distribution's mean equals the average level of the mean in the simulations ($\overline{\mu} = 0.4221$). The evaluated variances are those from the parameter-grid that induce log-concavity of the cost distribution, given the assumed mean cost. For the naive and Nash reserve price strategies, the actual data show a slight increase in the critical discount factor as the variance increases, while the predictions show a slight decrease. For the deterrent reserve price strategy, the predicted and actual data match closely and show a slight decrease in the critical discount factor as the variance increases. The sellers' discount factor must always exceed 0.973 for a strong duopoly cartel to sustain collusion against a strategic buyer using the deterrent reserve price strategy.

Figure 3 Here

Figures 4 illustrates the effect of changing the number of sellers in a strong cartel when the cost distribution's mean and variance equal their average levels in the simulations. The predicted and actual values match extremely closely, with one significant difference for the Nash reserve price strategy shown in panel (b). The actual data show the decrease in the critical discount factor as the number of firms increases from two. Though the critical discount factor does increase in N with three or more sellers, in this example collusion is more difficult to sustain with two sellers than it is with up to nine. This result is not uncommon in the data, though even the quartic terms fail to capture it fully in the predictions. This failure exists even in panel (c), for which the actual data illustrate a small but real decline in the critical discount factor as the number of sellers increases from two to three.

Figure 4 Here

To examine the generality of the examples just presented, I use the regression results in Table 1 to provide a comprehensive numerical evaluation of comparative statics. The effect of increasing the mean is predicted by the regression equation to be

$$\frac{d\delta}{d\mu} = \alpha_2 + \alpha_5 N + 2\alpha_7 \mu + \alpha_8 \sigma^2.$$

Similarly, the effect of increasing the variance is predicted to be

$$\frac{d\delta}{d\sigma^2} = \alpha_3 + \alpha_6 N + \alpha_8 \mu + 2\alpha_9 \sigma^2.$$

These derivatives can be calculated for any (N, μ, σ^2) triplet in the grid of feasible parameters. General patterns in the signs of these derivatives provide evidence about the influence of changes in the primitives on the ease of sustaining collusion.

Using the regression results in Table 1, the first row of Table 2 reports the fraction of feasible (N, μ, σ^2) triplets for which the derivative of the critical discount factor, with respect to the mean cost, is strictly positive. The second row similarly reports the fraction of strictly positive derivatives with respect to the variance. Because one might be more curious about the derivatives solely in more concentrated markets, rows 3 and 4 report about the same derivatives as rows 1 and 2, limited to those (N, μ, σ^2) triplets in the feasible parameter grid with ten or fewer sellers.

Table 2 Here

The results in Table 2 suggest that collusion typically becomes more difficult to sustain as the mean cost increases, particularly in more concentrated markets. While this result potentially is due partly to the change in the reserve price as μ changes, the fractions reported for the naive reserve price, which is unchanged as μ changes, indicate that the result is largely due to the increase in the mean cost causing a decrease in the difference between the per-period cooperative and noncooperative payoffs. With the exception of a weak cartel facing the naive reserve price strategy, the results in Table 2 provide even stronger evidence that collusion typically becomes easier to sustain as the variance of the sellers' costs increases. This result is caused by increasing the difference between the per-period cooperative and noncooperative payoffs at σ^2 increases. With no a priori reasoning about how the difference between the per-period cooperative and noncooperative payoffs changes as the mean or the variance change, the results in Table 2 provide useful insights into the determinants of sellers' ability to collude.

Examining the level of the critical discount factor is just as valuable as is examining the preceding comparative statics, because the critical discount factors that emerge from the simulations are quite high. For example, the average critical discount factors for duopoly facing the deterrent reserve price strategy are approximately 0.966 for a strong cartel and 0.979 for a weak cartel. In contrast, the critical discount factor for successful collusion in a homogeneous product Bertrand duopoly is $\frac{1}{2}$ when firms have constant identical marginal costs and split the market each period. If the Bertrand duopolists alternate production, which is more comparable to the procurement setting examined here, then the critical discount factor is $\frac{2}{3}$. While the difference between the Bertrand games and the procurement setting is striking, it is not clear whether the critical discount factors for strong and weak cartels are high from a practical perspective.

Relating the discount factors to plausible real-world values helps one to evaluate the likelihood of tacit collusion using the two proposed schemes. By definition, the one period discount factor has a one to one relationship to the one period discount rate, R, which is given by $\delta = \frac{1}{1+R}$. If a firm's discount rate were known, then one could determine the firm's discount factor. This discount factor then could be compared to the critical discount factors derived in the simulations.

One measure of a firm's discount rate is its opportunity cost of capital. Another measure of a firm's discount rate is its "hurdle" rate, which is the minimum annual return on investment necessary for an investment opportunity to be undertaken. While many factors are considered by firms selecting investment opportunities, it seems reasonable to use a firm's hurdle rate as its annual discount rate.

If the contracts considered in the repeated game model are not solicited annually, then one must adjust the hurdle rate to calculate the correct discount factor. If contracts are offered every t years, where t years constitute one period from the repeated game perspective, and R is the hurdle rate, then the one period discount factor is $\delta = \frac{1}{(1+R)^t}$. If contracts are offered more than once per year, then t is less than one.[25]

Table 3 presents the per-period discount factors associated with four hurdle rates over six different time periods, as measured by the number of contracts offered per year. For hurdle rates of 15 percent and higher, the resulting discount factors are below those reported in Table 4 as being necessary for successful collusion by a weak cartel versus a sophisticated buyer, even with a small number of sellers. With a 10 percent hurdle rate, the same conclusion holds if six or fewer contracts are offered per year. The critical discount factors in Table 4 illustrate that a strong cartel also faces difficulty when the buyer uses the deterrent reserve price strategy. Hence, in many instances sellers with costs drawn from the "average" Beta distribution examined in Table 4 could not successfully collude against a sophisticated buyer by using the optimal mechanism with or without side payments. In contrast, if the buyer cannot credibly impose a reserve price, then collusion can be sustained by strong and weak cartels for several values of the hurdle rate, as shown by the critical discount factors in Table 4 associated with the naive reserve price strategy.

Table 3 Here

Table 4 Here

Even if the sellers' discount factor is such that they can collude, the buyer can decrease the sellers' discount factor by decreasing the frequency with which contracts are let. As Tables 3 and 4 show, a sophisticated buyer can eliminate the possibility of collusion by a strong cartel with a 15 percent hurdle rate by offering contracts every two months rather than every month. Note that the change in frequency does not change the critical discount factor, but does change the sellers' true discount factor.

Finally, the buyer may prefer to reduce the number of sellers from which it accepts price offers. As the computational results revealed, in some instances collusion may be easier to sustain with more firms. Of course, the desire to eliminate collusion must be balanced against the expected profit in both cases.

5 Conclusion

[25] One could envision a procurement setting in which there are multiple buyers within an industry, which in many respects would appear to be equivalent to the single buyer setting with a further increase in the number of contract offerings. One potential difficulty is that individual buyers may have incentives to game the system by raising their reserve price without inducing collusion, because the sellers are unable to collude against the remaining buyers who chose their reserve prices for deterrent purposes.

This paper uses tools from repeated game theory to examine the ability of sellers in repeated first-price procurement auctions to sustain two collusive bidding schemes that are optimal from a static mechanism design perspective. Specifically, for different numbers of firms, different reserve price strategies, and a large number of cost distributions, I numerically determine the value that firms must place on future profits such that abiding by the collusive scheme is more profitable than is defecting from it and inciting retaliation. Computing the critical discount factors necessary for collusion to be sustainable illuminates the analytically complex effects of the various parameters of the strategic environment. Moreover, relating the computed discount factors to plausible real-world discount factors helps one to assess the practical relevance of the static mechanism design approach to collusion in auctions.

The computational results reveal that the buyer's choice of the reserve price has a large impact on the sellers' ability to collude. If the buyer is sophisticated in its choice of the reserve price, then collusion tends be sustainable only for extremely high discount factors that correspond to what appear to be unreasonably low hurdle rates within the firm. Successful collusion requires the sellers to place such high value on future profits because the short-term gain from cheating on the collusive agreement is large relative to the per-period loss of collusive profits. Thus, sellers must value those foregone profits highly in order to resist their temptation to cheat. The necessity of such extremely high discount factors suggests it may be unlikely that tacit collusion using the previously identified bidding schemes can be supported as a subgame perfect equilibrium of the repeated game. This conclusion is consistent with the assertion by Graham and Marshall [1987] that cartels in first-price auctions in the United States are unstable. Firms in such auctions may attempt to collude, only to find that their arrangement cannot be sustained.

The results also provide strong evidence that collusion is more difficult to sustain when the sellers' costs have a high mean or a low variance. Surprisingly, collusion in highly concentrated markets might be more easily sustained as the number of firms increases, which is not a feature of traditional oligopoly models. The likelihood of collusion initially increases in the number of sellers if the difference between the expected per-period collusive and punishment profits increases as the number of colluding sellers increases, because increasing that difference decreases the critical discount factor. The buyer may be able to use this phenomenon to eliminate collusion by reducing the number of sellers from which it accepts price offers. As a complementary strategy, the buyer also can deter collusion through the more familiar method of increasing the duration of contracts, which affects sellers' per-period discounting.

The strong conclusion that the probability of rigid-price collusion is low versus strategic buyers in an independent private values environment raises the issue of how the many reported instances of collusion in auction and procurement markets were supported. Relaxing three assumptions used in the present analysis generates three possible explanations of how collusion might more readily be sustained. First, buyers may not be very sophisticated in their reserve price selection, they may not be able to credibly commit to the reserve prices necessary to severely limit collusion, or they may have large values relative to the sellers' costs. If so, then the results for naive buyers apply, and they indicated that collusion was more easily sustained.

Second, collusion might be more profitable in common value environments, and it is possible that the reported instances of collusion tended to occur in procurement settings in which common value elements were more important than private value elements. Collusion in common value environments generates better information about the contract's true value and permits sellers to avoid the winner's curse. If the value of collusion in common value environments exceeds its value in private value environments, then it may be easier to sustain.

Third, the two schemes I examine may be less effective than other collusive schemes one could design. In particular, colluding firms might coordinate on a price less than the reserve price, in order to make deviation less attractive. However, coordinating on a lower price reduces the value of collusion, which reduces the per-period loss from defection and makes collusion more difficult to sustain. Moreover, if sellers coordinate on a price less than, but affected by, the reserve price, then one must account for the possibility that the buyer will select its reserve price differently. Clearly, determining more effective means of colluding in procurement markets is a worthwhile area for future research.

Appendix

Lemma 1 *Consider a bid rotation scheme in which each of N firms is equally likely to be selected to win the contract at price r. If a firm is selected but its cost exceeds r, then one of the remaining firms is randomly selected to win the contract at price r. This firm can also decline, and so on. A firm's ex ante expected profit using this scheme is*

$$\overline{\pi}_S^{BR}(r) = \frac{1 - [1 - F(r)]^N}{NF(r)} \int_{\underline{c}}^{r} F(c)dc.$$

Proof of Lemma 1: Suppose that a firm has cost less than r and is selected to win, either initially or subsequent to other firms being selected but having costs exceeding r. The firm's expected profit, conditional on having cost less than r, is

$$\frac{\int_{\underline{c}}^{r}(r-c)\,f(c)dc}{F(r)}=\frac{\int_{\underline{c}}^{r}F(c)dc}{F(r)},$$

where the term on the right hand side following from integrating the left hand side by parts. The probability the firm's cost is less than r is $F(r)$, and the probability the firm is selected from the N firms is

$$\frac{1}{N}+\frac{1}{N}\left[1-F(r)\right]+\frac{1}{N}\left[1-F(r)\right]^{2}+\cdots+\frac{1}{N}\left[1-F(r)\right]^{N-1}=\frac{1-\left[1-F(r)\right]^{N}}{NF(r)}.$$

Therefore, the firm's ex ante expected profit is

$$\overline{\pi}_{S}^{BR}(r)=\frac{1-\left[1-F(r)\right]^{N}}{NF(r)}\times F(r)\times\frac{\int_{\underline{c}}^{r}(r-c)\,f(c)dc}{F(r)}=\frac{1-\left[1-F(r)\right]^{N}}{NF(r)}\int_{\underline{c}}^{r}F(c)dc,$$

which is the desired result. ■

Proof of Proposition 1: The determination of ex ante expected profit is made in McAfee and McMillan [1992]. The following incentive constraint must be satisfied to prevent defection by a firm with cost $\underline{c}$ that is not selected to win:

$$\left(\frac{\delta}{1-\delta}\right)\overline{\pi}_{S}^{WC}(r)\geq r-\underline{c}+\left(\frac{\delta}{1-\delta}\right)\overline{\pi}_{S}^{NE}(r).$$

The smallest δ for which this is satisfied is

$$\overline{\delta}^{WC}(r)\equiv\frac{N\left(r-\underline{c}\right)}{\int_{\underline{c}}^{r}F(c)\left[\frac{1-\left[1-F(r)\right]^{N}}{F(r)}-N\left[1-F(c)\right]^{N-1}\right]dc+N\left(r-\underline{c}\right)},$$

which is the desired result. ■

Proof of Proposition 2: A seller's interim expected profit is

$$\overline{\pi}_{S}^{SC}(c|r)=\begin{cases}\left[T(c)-c\right]\left[1-F(c)\right]^{N-1}+\left[1-\left[1-F(c)\right]^{N-1}\right]\int_{\underline{c}}^{c}\left[\frac{r-T(s)}{N-1}\right]\left[\frac{(N-1)\left[1-F(s)\right]^{N-1}}{1-\left[1-F(c)\right]^{N-1}}\right]f(s)ds \text{ if } c\leq r\\ \left[1-\left[1-F(r)\right]^{N-1}\right]\int_{\underline{c}}^{r}\left[\frac{r-T(s)}{N-1}\right]\left[\frac{(N-1)\left[1-F(s)\right]^{N-1}}{1-\left[1-F(r)\right]^{N-1}}\right]f(s)ds \text{ if } c>r\end{cases}$$

Through tedious integration of the interim expected profit, one can show that the seller's ex ante expected profit is

$$\int_{\underline{c}}^{\overline{c}}\overline{\pi}_{S}^{SC}(c|r)f(c)dc=\frac{r-\underline{c}}{N}-\frac{1}{N}\int_{\underline{c}}^{r}\left[1-F(c)\right]^{N}dc.$$

As shown in the text, only the following incentive constraint must be satisfied to prevent defection by a firm with cost $\underline{c}$ that is not the selected winner :

$$\left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{SC}(r) \geq r-\underline{c}+\left(\frac{\delta}{1-\delta}\right)\overline{\pi}_S^{NE}(r).$$

The smallest δ for which this is satisfied is

$$\overline{\delta}^{SC}(r) \equiv \frac{N\left(r-\underline{c}\right)}{\left(N+1\right)\left(r-\underline{c}\right)-\int_{\underline{c}}^{r}\left[\left(N-1\right)F(c)+1\right]\left[1-F(c)\right]^{N-1}dc}.$$

Finally, it is straightforward to show that the derivative of $\overline{\delta}^{SC}(r)$ with respect to r is negative. ■

Lemma 2 *Consider a setting in which the sellers can credibly reveal their costs to one another, and in which the low-cost seller is designated to win by setting a price equal to the reserve price r. There are no side payments, so firms designated to lose earn nothing. A seller's ex ante expected profit from participating in the credible revelation scheme is*

$$\overline{\pi}_S^{CR}(r) = \frac{r-\underline{c}}{N} - \frac{1}{N}\int_{\underline{c}}^{r}\left[1-F(c)\right]^{N}dc.$$

Proof of Lemma 2: A seller's interim expected profit is

$$\overline{\pi}_S^{CR}(c|r) = (r-c)\left[1-F(c)\right]^{N-1}.$$

The seller's ex ante expected profit is

$$\int_{\underline{c}}^{r}\overline{\pi}_S^{CR}(c|r)f(c)dc = \frac{r-\underline{c}}{N} - \frac{1}{N}\int_{\underline{c}}^{r}\left[1-F(c)\right]^{N}dc,$$

which is the desired result. ■

References

[1] Athey, S., Bagwell, K., and Sanchirico, C. "Collusion and Price Rigidity" unpublished manuscript (2000).

[2] Baldwin, L., Marshall, R., and Richard, J., "Bidder Collusion an Forest Service Timber Sales." *Journal of Political Economy*, Vol. 105, No. 4 (1997), pp. 657-699.

[3] Comanor, W., and Schankerman, M., "Identical Bids and Cartel Behavior." *Bell Journal of Economics*, Vol. 7 (1976), pp. 281-286.

[4] Cramton, P., and Palfrey, T., "Cartel Enforcement with Uncertainty About Costs." *International Economic Review*, Vol. 31, No. 1 (1990), pp. 17-47.

[5] Gal-Or, E. "Information Transmission-Cournot and Bertrand Equilibria." *Review of Economic Studies*, Vol. 53 (1986), pp. 85-92.

[6] Graham, D., and Marshall, R., "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions." *Journal of Political Economy*, Vol. 95, No. 6 (1987), pp. 1217-1239.

[7] Haubrich, J., "Risk Aversion, Performance Pay, and the Principal-Agent Problem." *Journal of Political Economy*, Vol. 102, No. 2 (1994), pp. 258-276.

[8] McAfee, R.P., and McMillan, J. "Bidding Rings." *American Economic Review*, Vol. 82, No. 3 (1992), pp. 579-599.

[9] McCutcheon, B., "Do Meetings in Smoke-Filled Rooms Facilitate Collusion?" *Journal of Political Economy*, Vol. 105, No. 2 (1997), pp. 330-350.

[10] Porter, R., and Zona, J.D., "Detection of Bid Rigging in Procurement Auctions." *Journal of Political Economy*, Vol. 101, No. 3 (1993), pp. 518-538.

[11] Quirmbach, H., "R&D: Competition, Risk, and Performance." *RAND Journal of Economics*, Vol. 24, No. 2 (1993), pp. 157-197.

[12] Scherer, F.M., "Focal Point Pricing and Conscious Parallelism." *Antitrust Bulletin*, Summer (1967), pp. 495-503.

[13] Scherer, F.M., and Ross, D. *Industrial Market Structure and Economic Performance*. Boston: Houghton Mifflin (1990).

[14] Shapiro, C. "Exchange of Cost Information in Oligopoly." *Review of Economic Studies*, Vol. 53 (1986), pp. 433-446.

[15] Vives, X. "Duopoly Information Equilibrium: Cournot and Bertrand." *Journal of Economic Theory*, Vol. 34 (1984), pp.71-94.

Table 1
Regression Results by Cartel Type and Reserve Price Strategy for

$$\delta = \alpha_0 + \alpha_1 N + \alpha_2 \mu + \alpha_3 \sigma^2 + \alpha_4 N^2 + \alpha_5 N\mu + \alpha_6 N\sigma^2 + \alpha_7 \mu^2 + \alpha_8 \mu\sigma^2 + \alpha_9 \sigma^4 + \alpha_{10} N^3 + \alpha_{11} N^4$$

	Weak Cartel Naive RP	Weak Cartel Nash RP	Weak Cartel Deterrent RP	Strong Cartel Naive RP	Strong Cartel Nash RP	Strong Cartel Deterrent RP
Intercept	0.69379149	0.86431262	0.91699998	0.68188789	0.84017995	0.89139666
Number of Firms	0.03704046	0.00248617	0.00267516	0.03874445	-0.00017688	0.00342790
Mean	0.21043461	0.29331887	0.26129209	0.17315251	0.25770755	0.29180795
Variance	0.24967026	-0.40937132	-0.51774748	-0.16285821	-0.42003522	-0.57022403
Number of Firms^2	-0.00250669	0.00036730	-0.00000426	-0.00253982	0.00085835	0.00003237
Number of Firms*Mean	-0.00792688	-0.00569608	-0.00300236	-0.00711913	-0.00671399	-0.00378106
Number of Firms*Variance	-0.01335749	-0.00272237	-0.00146551	0.00159331	-0.00248975	-0.00237113
Mean^2	0.00772701	-0.15093804	-0.19797418	0.02759866	-0.07577537	-0.20118931
Mean*Variance	0.33454753	0.38034936	0.46419345	-0.18393295	0.03053952	0.47785869
Variance^2	-1.12080580	0.67125324	1.45184525	0.84536447	2.30518151	1.67026818
Number of Firms^3	0.00008605	-0.00002024	-0.00000176	0.00008394	-0.00004206	-0.00000427
Number of Firms^4	-0.00000109	0.00000032	0.00000004	-0.00000103	0.00000063	0.00000008
Data Points	13,862	13,862	13,862	13,862	13,862	13,862
Adjusted R^2	0.94958884	0.90662891	0.78330459	0.96408683	0.93118992	0.87697972
Adjusted R^2 (1st order only)	0.67843531	0.68311733	0.48413183	0.71895946	0.75560037	0.61968881

Table 2
Percentage of Feasible (N,μ,σ^2) Triplets with Strictly Positive Derivatives

	Weak Cartel Naive RP	Weak Cartel Nash RP	Weak Cartel Deterrent RP	Strong Cartel Naive RP	Strong Cartel Nash RP	Strong Cartel Deterrent RP
$\frac{d\delta}{d\mu}(N \le 30)$	94.75%	86.57%	76.32%	86.62%	90.92%	80.41%
$\frac{d\delta}{d\sigma^2}(N \le 30)$	71.98%	0%	0%	0%	0.02%	0%
$\frac{d\delta}{d\mu}(N \le 10)$	100%	99.33%	85.61%	100%	100%	90.66%
$\frac{d\delta}{d\sigma^2}(N \le 10)$	100%	0%	0%	0%	0.07%	0%

Table 3
Per-Period Discount Factors Associated with Various Hurdle Rates and Time Horizons

Hurdle Rate	12 contracts/year	6 contracts/year	4 contracts/year	3 contracts/year	2 contracts/year	1 contract/year
5%	0.996	0.992	0.988	0.984	0.976	0.952
10%	0.992	0.984	0.976	0.969	0.953	0.909
15%	0.988	0.977	0.966	0.954	0.933	0.870
20%	0.985	0.970	0.955	0.941	0.913	0.833

Table 4
Critical Discount Factors for Weak and Strong Cartels
Versus Naive and Deterrent Reserve Price Strategies
(μ and σ^2 at average levels from parameter set)

Number of Firms	Weak Cartel Naive RP	Weak Cartel Deterrent RP	Strong Cartel Naive RP	Strong Cartel Deterrent RP
2	0.849492	0.990717	0.810981	0.977497
3	0.877683	0.990965	0.837246	0.977051
4	0.899341	0.991818	0.861076	0.978607
5	0.914829	0.992628	0.879408	0.980327
6	0.926262	0.993319	0.893588	0.981893

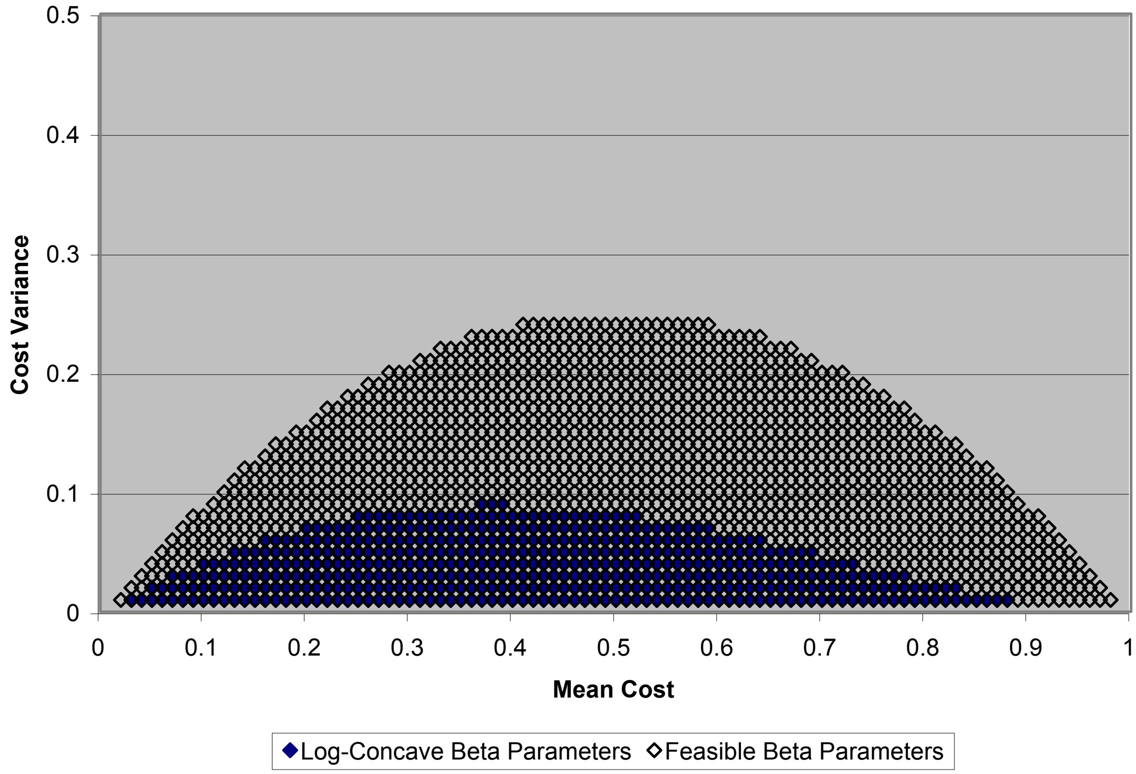

Figure 1
Feasible and Log-Concave Beta Parameters
(Grid size = 0.01)

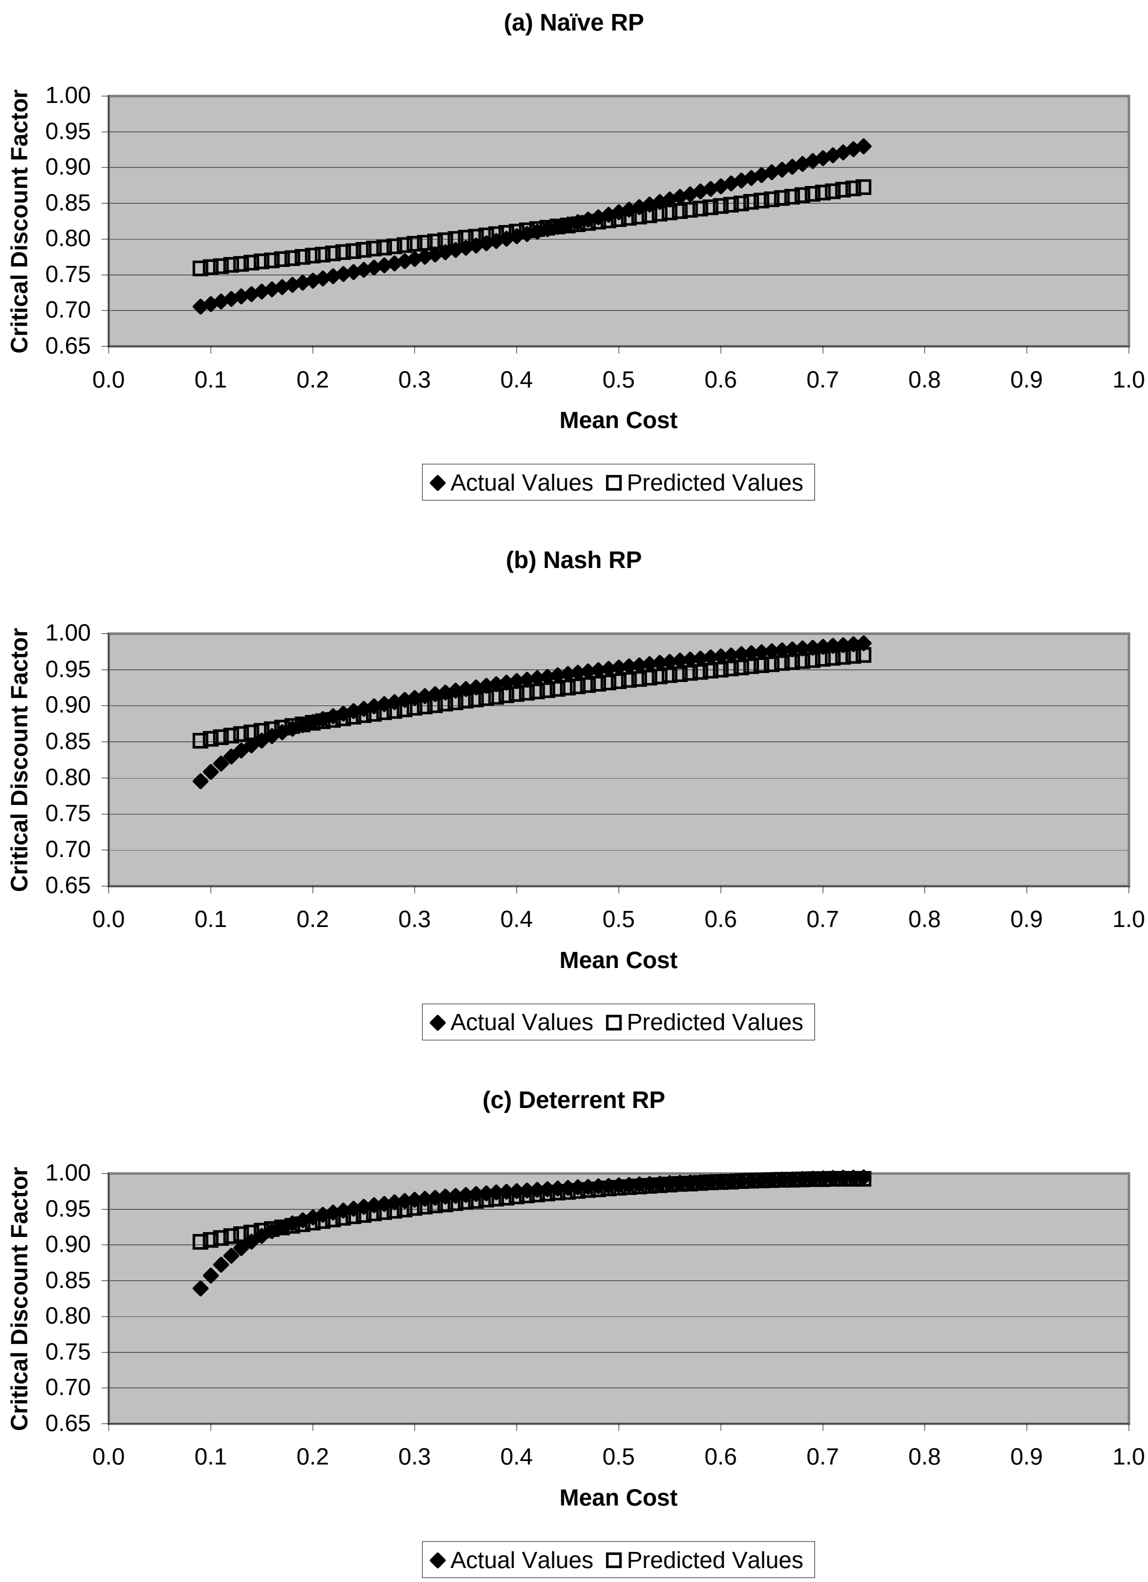

Figure 2

Critical Discount Factors for Strong Duopoly Cartel with σ^2 at Average Level

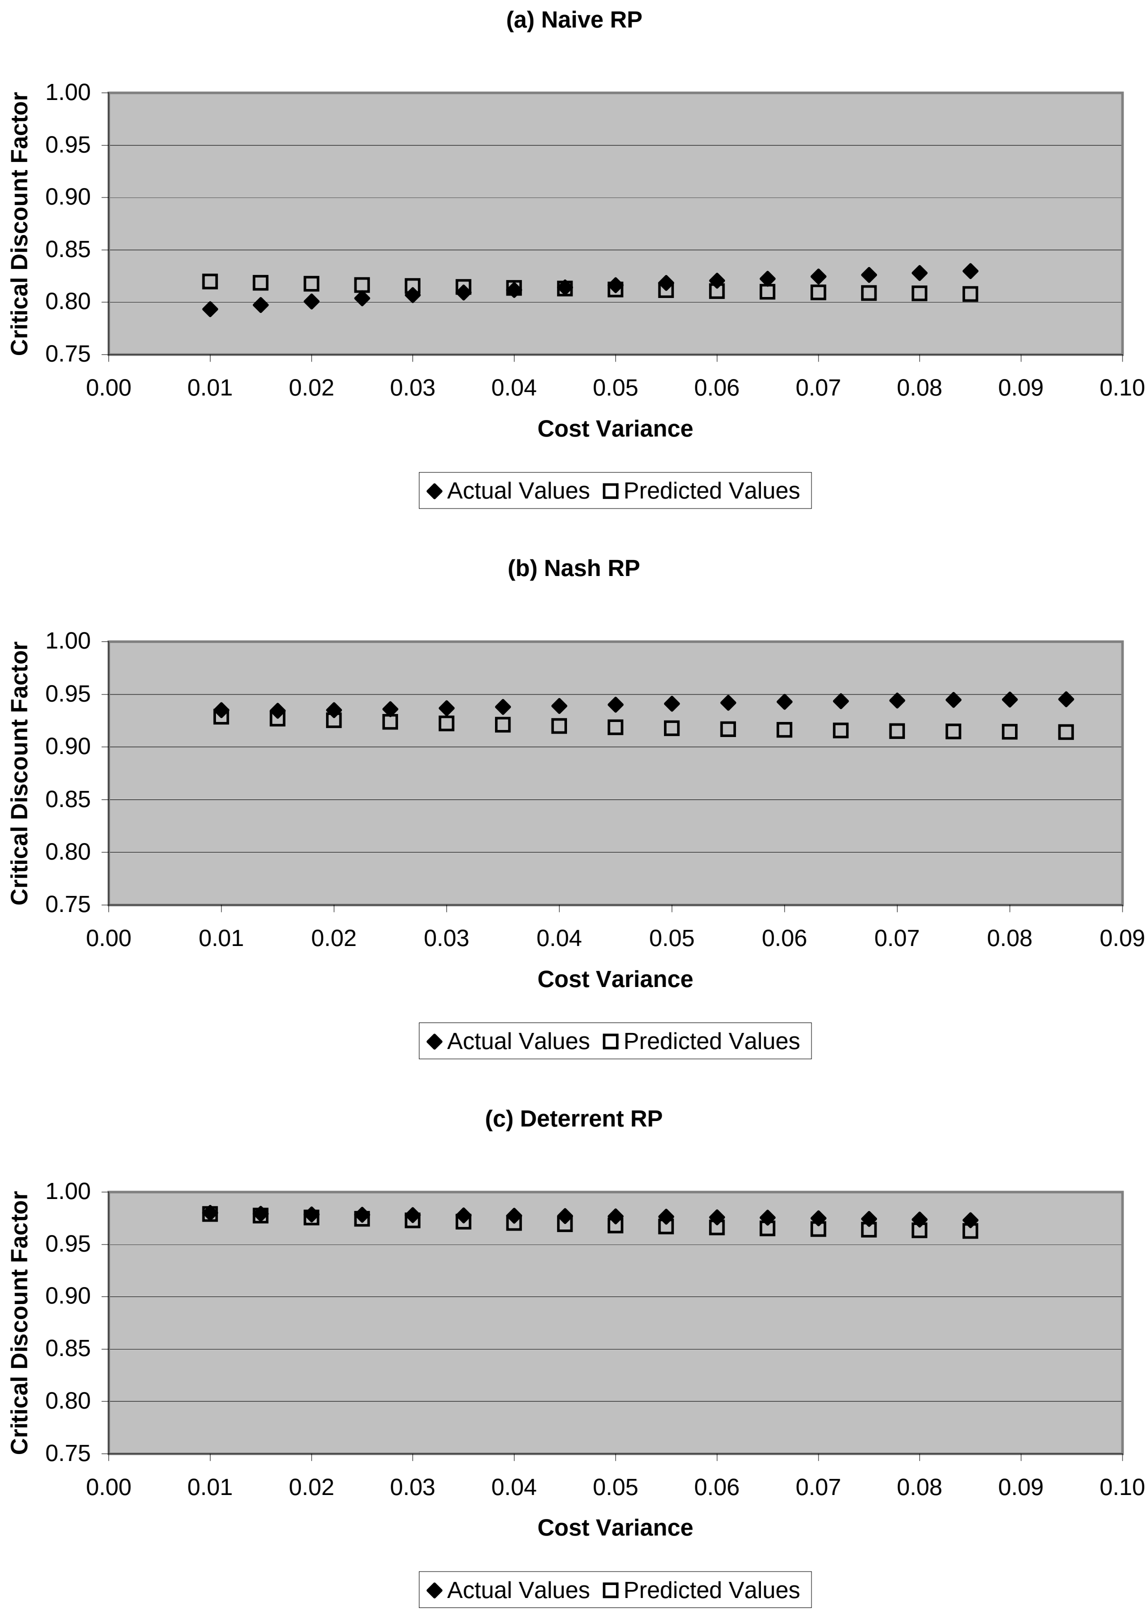

Figure 3
Critical Discount Factors for Strong Duopoly Cartel with μ at Average Level

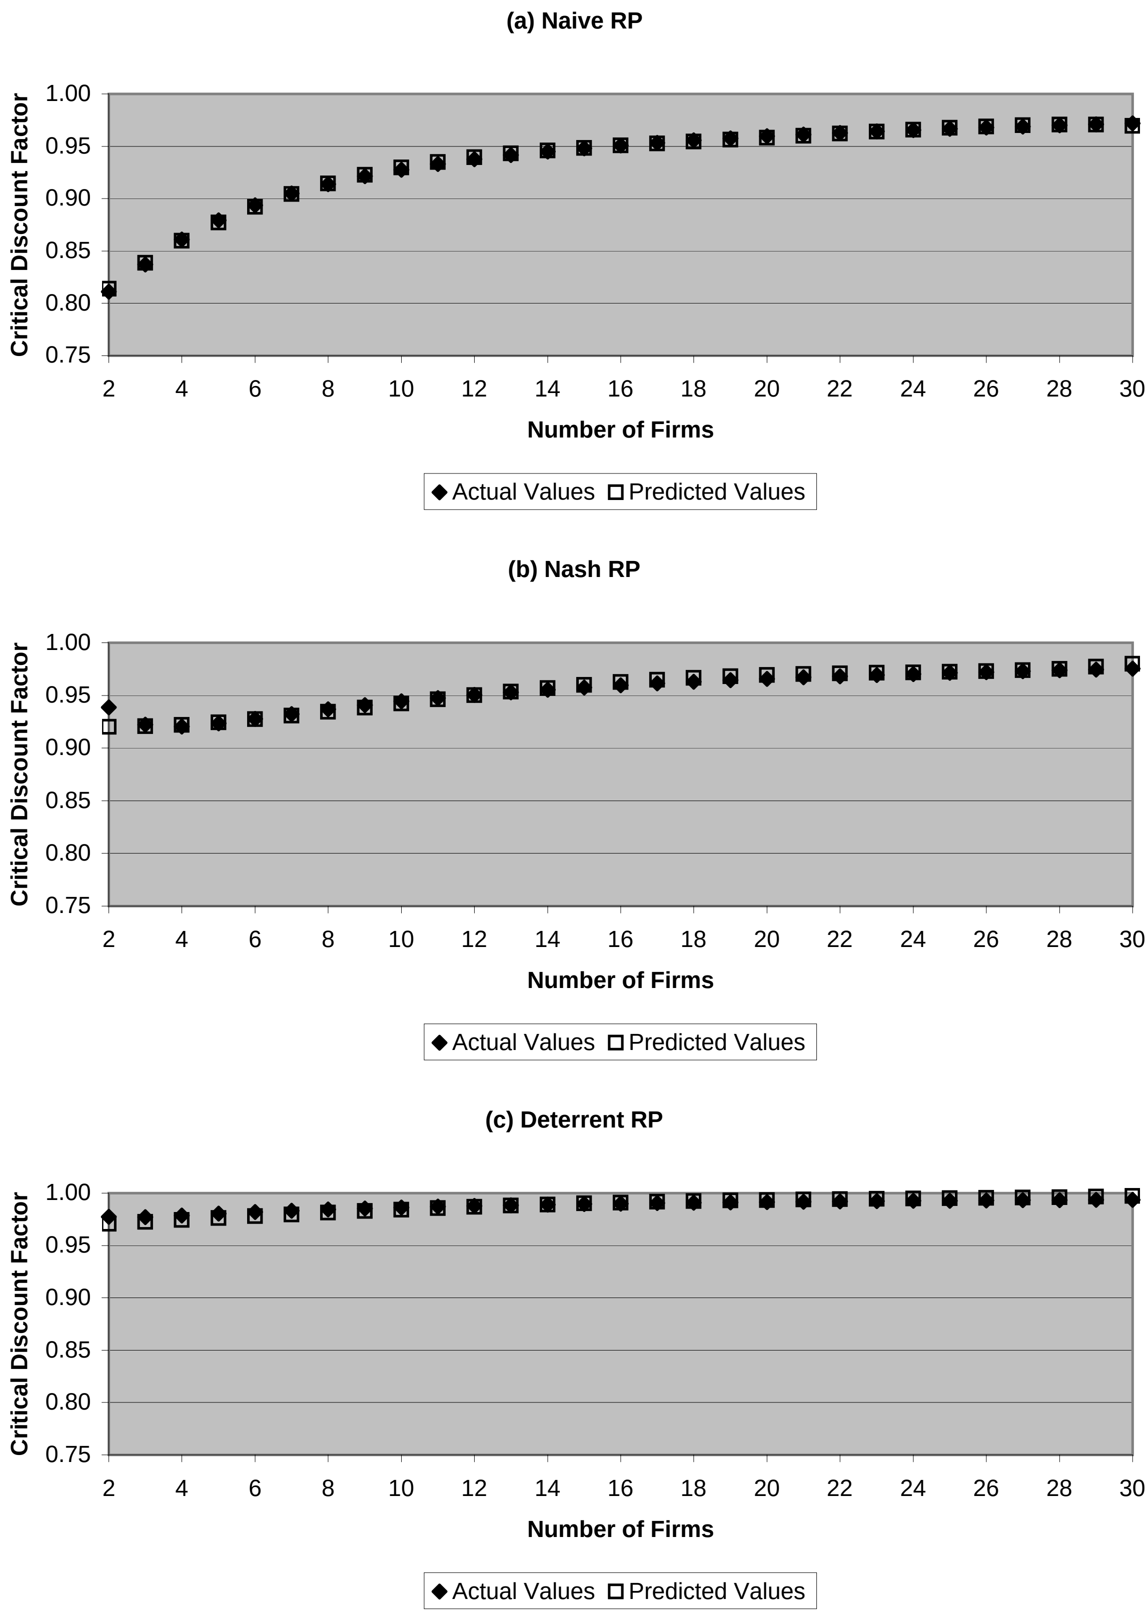

Figure 4

Critical Discount Factors for Strong Cartel with μ and σ^2 at Average Levels