

CHAPTER 20

AUCTIONS AND BID RIGGING

KEN HENDRICKS, R. PRESTON MCAFEE, AND
MICHAEL A. WILLIAMS

20.1. INTRODUCTION AND OVERVIEW

20.1.1. Auctions

An auction is “a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants” (McAfee and McMillan, 1987).¹ Auctions are among the oldest market institutions; “marriage auctions” for brides were used in Babylon as early as 500 BC. Products commonly bought and sold using auctions include art and jewelry, US Treasury Bills, used cars, the assets of bankrupt companies, and radio spectrum. Auctions are also commonly used in the government procurement of goods and services. Online auctions comprise a significant portion of the rapidly growing electronic commerce marketplace.

Auction theory provides an explicit model of price formation. Traditional models of competition that furnish equilibrium prices given demand and supply characteristics, such as Cournot and Bertrand models of competition, do not explain the price-formation process in terms of buyer-seller interactions. In contrast, auction mechanisms explicitly lay out rules that govern the formation of prices. A typical auction consists of a single auctioneer, responsible for selling an object, and a number of bidders who wish to buy the object. The auctioneer may announce a reserve price, that is, the

¹ A number of survey papers on auction theory and empirical work have been published. See, e.g., McAfee and McMillan (1987), Wilson (1992), Klemperer (1999), and Hendricks and Porter (2007). In addition, Klemperer (2008) analyzes the use and misuse of bidding markets in antitrust economics.

lowest price at which the object would be sold. There are four basic types of single-item auctions analyzed in the economics literature: (1) the English auction, (2) the Dutch auction, (3) the first-price sealed-bid (FPSB) auction, and (4) the second-price sealed bid (SPSB) auction (also called a Vickrey auction).

These four types of auctions are characterized by the following two properties: (1) whether bids are open or closed (i.e., sealed) and (2) whether the winning bidder pays an amount equal to the winning bid or the highest losing bid. An English auction is an open-bid auction in which bidders sequentially raise the price of the item until no bidder is willing to raise it further. A Dutch auction is an open-bid auction that proceeds in the opposite direction: The seller begins the auction at a particular high price, and lowers the price until a bidder is willing to purchase the item. In both FPSB and Dutch auctions, the winning bidders pay an amount equal to the winning bid. In a closed auction, each bidder may submit only a single, sealed bid (that is, no bidder may observe another bidder's bid). The seller then considers all submitted bids simultaneously. In FPSB and SPSB auctions, bids are submitted in a sealed fashion and are considered simultaneously by the seller to determine the winning bidder. In a FPSB auction, the highest bidder wins the item and pays an amount equal to the highest bid, that is, the winning bid. In a SPSB auction, the highest bidder wins the item and pays an amount equal to the second-highest bid, that is, the highest losing bid. In all four auction formats, the seller does not sell the item if the price determined by the auction is below the reserve price.

In the theoretical literature on auctions, such as Milgrom and Weber (1982), English auctions are usually approximated by Japanese auctions for the sake of analytical convenience. In a Japanese auction, the seller starts the auction with the price of the item at a minimum level. Each bidder presses and holds down a button while the price of the item is raised continuously. Bidders drop out of the auction by releasing their button when the price of the item exceeds their willingness to pay, that is, their valuation. Each bidder can observe all of the bids. The auction ends when the second-highest bidder drops out, so the winning bidder is the last remaining bidder. The continuous increase in price in a Japanese auction simplifies the theoretical analysis and makes it transparent that the winning bid in an English auction is approximately the second-highest valuation.

20.1.2. Standard Auction Models

The economic literature on auctions typically distinguishes two extreme types of information environments: private value (PV) and common value (CV). In a PV auction, each bidder knows her own value but not those of other bidders. For example, an art auction with art collectors who do not intend to resell the item would be analyzed using the PV framework. Each bidder's valuation is a personal characteristic that does not depend on other bidders' valuations. In a CV auction, the value of the item up for auction is the same for all bidders, but that value is unknown to the bidders before the auction. Prior to

bidding, bidders receive different private signals about the true value of the object. For example, in an offshore oil tract auction, bidders conduct their own geologic research to form their private estimates of the value of an offshore tract.

The standard economic model of an auction consists of a single seller and n bidders, $i = 1, 2, \dots, n$, who bid at an auction to purchase a single indivisible good. Bidder i 's valuation for the good is V_i , a real-valued random variable distributed according to the cumulative probability distribution F_i . The value of V is known only to bidder i , but F_i is known to the seller and the remaining bidders. In the symmetric independent private value model (IPV), the valuations of bidders are assumed to be independent draws from the same distribution F . In the CV model, the valuations of bidders are informative signals about a common value V unobservable to the bidders. In this case, V_i is typically modeled as an independent draw from a conditional distribution, $H(\cdot | V)$, known to the bidders. Bidders in a common value auction are susceptible to the "winner's curse." A bidder tends to win when she overestimates V and wins rarely when she underestimates V . The PV and CV frameworks also have been merged to study auction models in which bidders' information consists of both private and common value components; see, for example, Milgrom and Weber (1982) or Goeree and Offerman (2003).

Auctions are generally modeled as a Bayesian game, and bidders have Bayesian-Nash equilibrium bidding strategies. In a FPSB IPV auction, the unique equilibrium with n risk-neutral, symmetric bidders is the collection of bidding functions $b_i(v_i) = v_i - [(\int_{-\infty}^x F(s)^{n-1} ds) / F(x)^{n-1}]$ for each $i = 1, 2, \dots, n$. So each bidder bids less than her value, and the markdown factor decreases with the number of bidders. In contrast, in a SPSB PV auction, each bidder has a weakly dominant strategy to bid her valuation, regardless of the number of bidders. In a FPSB CV auction, each bidder shades her bid relative to her expected value of the object conditional on winning. However, the shading factor typically increases with the number of bidders due to the winner's curse. In SPSB CV auctions, the equilibrium is not in dominant strategies and not unique without some refinement. One equilibrium involves bidding more than the average value (specifically bidding the value conditional on a second bidder having the same signal). In this case, the actual bid will exceed the willingness to pay, although the second-highest bid, the price, will not.

A fundamental (and remarkable) result in auction theory is the revenue equivalence theorem for IPV auctions. An auction mechanism (e.g., an English or Dutch auction) is efficient if the mechanism allocates the object to the bidder with the highest valuation (Dasgupta and Maskin, 2000). The revenue equivalence theorem essentially states that all efficient, symmetric IPV auctions yield the same expected revenue. Thus, for IPV auctions, the expected winning bid is the same whether the auction mechanism is English, Dutch, FPSB, or SPSB. A necessary assumption for the theorem to hold is that the bidder with the lowest feasible valuation expects zero surplus. A first version of this result was obtained by Vickrey (1961) and later generalized by Myerson (1981) and Riley and Samuelson (1981). The revenue equivalence result does not hold, in general, for common-value auctions.

20.1.3. Collusion in Auctions: Bid Rigging

In a bid-rigging scheme, bidders agree to collude to limit competition and obtain the good for a lower price (or higher in the case of procurement) than would result in the absence of such explicit collusion.² A group of bidders who collude in an auction is called a bidding ring. Bidding rings may be all-inclusive, that is, include every bidder participating in an auction, or incomplete, consisting only of a subset of the set of all bidders. As Marshall and Marx (2012, p. 175) summarize: “Colluding bidders suppress their rivalry through the elimination of meaningful bids by all colluding bidders except for the ring bidder with the highest value. All other details and logistics of bidder collusion flow from this foundational principle of the ring.” Bid rigging is generally illegal in the United States, the EU, China, and many other countries.³ Bid rigging is the most common form of explicit collusion in auction markets (the other, less common, form being explicit collusion between one or more bidders and the auctioneer). This should be expected since noncooperative behavior is not jointly optimal for bidders. Bidders are collectively better off colluding and transferring gains from trade from the seller to the ring.

Rings differ from cartels insofar as they focus on individual auctions rather than a broad cross-section of product and geographical markets. The focus simplifies the task of a ring: it needs to coordinate the bids submitted by bidders. However, in doing so, rings face some of the same challenges as cartels: detection by the antitrust agencies or by the seller, internal enforcement, entry, and private information about the gains from trade. The ability of rings to solve these problems, and the nature of the solution, depends on the type of auction the seller uses.

One solution is to have all bidders submit identical bids, preferably at or near the reserve price, and let the seller randomly select the winning bidder. This scheme wastes a lot of the potential surplus from collusion because the good is unlikely to go to the bidder with the highest valuation. It is also difficult to enforce. Individual bidders have a strong incentive to bid slightly more than the agreed-upon price and win the good for sure. It is also easy to detect. Not surprisingly, this simple form of collusive bidding has become much less common after the US Federal Trade Commission prosecuted numerous cases of identical bidding in the mid-twentieth century. A second solution is to have

² There is a substantial economic literature on noncooperative, tacit collusion in auctions. For theoretical analyses, see, e.g., Brusco and Lopomo (2002), Fabra (2003), and Blume and Heidhues (2006). For empirical analysis, see, e.g., Cramton and Schwartz (2000) and Ishii (2009). For experimental analyses, see, e.g., Sherstyuk (1999), Sherstyuk (2002), Kwasnica and Sherstyuk (2007), Li and Plott (2008), Sherstyuk and Dulatre (2008), Boone et al. (2009), Haan, Schoonbeek, and Winkel (2009), Potters (2009), Phillips and Menkhaus (2009), Brown, Plott, and Sullivan (2009), Hinloopen and Onderstal (2010), and Müller and Normann in volume 1 of this *Handbook*. In addition, there is also a substantial economic literature on corruption in auctions, i.e., bidders colluding with the auctioneer. See, e.g., Compte, Lambert-Mogiliansky, and Verdier (2005). We do not analyze these literatures here.

³ An interesting exception in the United States is the legality of explicit bid rigging in hostile takeovers of publicly traded companies. See McAfee et al. (1993).

bidders pay kickbacks to each other as payment for either refusing to bid or for submitting “phantom” bids, that is, intentionally losing bids. In order to avoid detection, bidders coordinate on the amount of the “serious” bid, and then “nonserious” bidders place phantom bids. This technique requires more communication and explicit discussion of prices, but makes collusion more difficult to detect by antitrust authorities using statistical methods (see discussion below).

The side payments among bidders can also solve the problem of private information. The gains from trade captured by the ring from the seller, called the collusive surplus, are maximized when the serious bid is submitted by the ring member with the highest valuation. But in order to make this assignment, the members of the ring need to be incentivized to reveal their valuations. In one-shot auctions, rings typically solve this issue by holding a knockout auction, either before or after the seller’s auction, for example, Marshall and Marx (2012). The bidder with the highest valuation typically wins the knockout auction. The ring members share a sum of money equal to the difference between the price in the ring’s knockout auction and the price in the original auction.

How the collusive surplus is allocated among ring members depends on the nature of the ring. If the ring members are symmetric, for example, their valuations are drawn from the same distribution, the ring may agree to allocate equal shares of the collusive surplus to each ring member. Alternatively, if some members of the ring are economically stronger than others, for example, have lower costs or higher valuations, the ring may allocate the collusive surplus based on those characteristics; see Graham, Marshall, and Richard (1990). For example, Asker (2010) analyzes how members of a bidding ring for the purchase of stamps in the United States allocated the collusive surplus among ring members depending on their bids in a knockout auction. Ring members shared each increment between bids in the knockout auction, provided their bids were above the auction price in the original auction. Half the increment was kept by the winner of the knockout auction, with the balance shared equally between those bidders who bid equal to or more than the incremental bid in the knockout auction.

In some auctions, for example, procurement contracts and corporate takeovers, a bid-rigging ring may allocate the collusive surplus by using subcontracts between ring members, by bringing “losing” ring bidders back into the deal as equity owners, or by splitting the takeover target into pieces (see McAfee et al., 1993).

In repeated auctions, the ring has more scope for meeting the challenges of collusion. The simplest assignment is for ring members simply to take turns, with each being the sole bidder in a given auction. More sophisticated bid rotation schemes involve bidders communicating with each other regarding items they would like to win. The ring can solve the incentive problem without knockout auctions or side payments by agreeing on future allocations (Athey and Bagwell, 2001). Bidders can keep a “tally sheet” recording each bidder’s winnings to ensure the bidders’ wins approximately balance out over time. This type of scheme is more efficient than simple bid rotation schemes, but requires more coordination and communication, thus increasing the likelihood the collusion is discovered and prosecuted by antitrust authorities.

The enforcement issue is also more easily solved by rings in repeated auctions. To deter bidders from cheating on a collusive agreement in repeated auctions, the ring can punish members who cheat by reversion to non-cooperative bidding. A “bidding war” ensues in which all firms pay high prices and earn low profits until the collusive agreement can be reestablished (see, e.g., Porter, 1983 and 2005 in the context of a railroad cartel).

In what follows, we will survey the literature on the theory and practice of bidding rings in one-shot auctions and in repeated auctions. The main theme is how the type of auction, whether it is first-price or second-price, sealed bid or oral, affects the incentive of bidders to collude and the way in which they collude.

20.2. BID RIGGING: THEORY

20.2.1. The Incentive to Collude in Different Auction Formats

Economic theory shows that the incentive of bidders to collude in an auction depends in part on the auction format (see, e.g., Robinson, 1985, Waehrer, 1999, and Marshall and Marx, 2007). In an ascending bid auction, for example, an English auction, the ring bids up to the highest valuation of its members. This bidding strategy is the same as if the ring member with the highest valuation had bid when there was no collusion. Thus, if in the presence of the bidding ring a non-ring member wins the auction, that non-ring member pays the same price that she would have paid in the absence of the ring. As Marshall and Marx (2012, p. 176) discuss, this has an important implication: “if a ring wins, and there are gains to their collusion, then the ring captures all of those gains; however, if the ring loses, there are no gains from the collusion for non-ring bidders. The collusion is self-contained in this sense.”

In contrast, at a sealed-bid auction, for example, a FPSB auction, the ring submits a bid lower than the amount that the ring member with the highest valuation would have bid in the absence of the collusion. This shading of the ring’s bid can sometimes result in a non-ring member winning the collusive auction when that non-ring member would not have won the auction in the absence of the collusion. As Marshall and Marx (2012, p. 176) note: “the extra bid shading by the highest-valuing colluding bidder opens the possibility that the ring does not capture all the gains to its collusive conduct. The non-colluding bidders are beneficiaries, in expected terms, from the collusion. . . . The leakage of some of the collusive gain to the outside bidders, which is absent at the ascending-bid auction, means the incentives for suppression of rivalry through collusion are typically weaker at the sealed-bid auction than at the ascending-bid auction.”

An interesting corollary to this discussion relates to the incentives of non-ring members to join a ring. In the case of ascending bid auctions, non-ring members have no

incentive to remain outside of the ring, since the ring captures all the collusive gains. In contrast, in the case of FPSB auctions, non-ring members may find it more profitable on average to remain outside the ring since the collusive gains leak out to non-ring members.

20.2.2. Bid Rigging in One-Shot versus Repeated Auctions

Given an auction format, what is the optimal, incentive-compatible collusive mechanism? Is the collusive mechanism efficient? Is it unique? Economic theory attempts to answer these questions. A primary issue in the theoretical analysis of bid rigging is stability. A collusive scheme cannot be sustained in the absence of an appropriate enforcement device, such as side payments or punishments. In the absence of such a device, bidders have an incentive to lie about their valuations or otherwise deviate from the collusive scheme. Theoretical studies of bid rigging in auctions fall into two broad areas: collusion in one-shot auctions and collusion in repeated auctions.

20.2.2.1. *Bid Rigging in One-Shot Auctions*

Graham and Marshall (1987) develop models of bidding rings in SPSB and English IPV auctions. They assume that all bidders are ex ante identical in the sense that their valuations are drawn from the same distribution. They describe a collusion mechanism in which, prior to the main auction, each ring member submits a sealed “reported bid” to a risk-neutral ring center. The ring center determines the two highest reported bids and selects the member with the highest reported bid to act as the sole bidder in the actual auction. The ring center specifies that the bid submitted in the actual auction should equal the highest reported bid submitted to the ring. Other ring members are instructed to bid zero or not submit a bid. If the ring member wins the main auction, she pays the auctioneer an amount equal to the second-highest bid of all bids submitted in the main auction. She also must pay the ring center the difference between (1) the second-highest reported bid to the ring and (2) the second-highest bid of all bids submitted in the main auction, if this amount is positive.

Graham and Marshall show that the auctioneer’s best response to the formation of a bidding ring is to set a reserve price that increases as a function of the number of bidders in the ring. For any given number of bidders in a ring, each ring member’s payoff to collusion decreases as the reserve price increases. Thus, for any given reserve price, the expected payoff to ring members increases as the number of bidders in the ring increases. Therefore, the Nash equilibrium is characterized by (1) a ring that includes all bidders and (2) a reserve price optimal for such a ring.

From an ex ante perspective, the ring center’s expected payments to ring members equal the center’s expected revenues from the ring member who submits the highest reported bid to the ring. In this sense, the ring center is ex ante budget balancing. However, a difficulty with the Graham and Marshall collusion mechanism is that the

ring center is not ex post budget balancing. That is, ex post the center's expected payouts exceed its expected revenues.

Mailath and Zemsky (1991) relax the identical-bidder assumption in Graham and Marshall (1987) and show that an ex post budget balancing efficient collusion can be achieved in a second-price auction. Furthermore, this outcome is possible even for a proper subset of bidders. They construct an explicit mechanism that implements this ex post efficient collusive result. With this efficient mechanism, the collusive surplus can always be divided up in such a way that, not only will every bidder wish to participate, but every subset of bidders will also wish to participate. As a consequence, no subset of bidders can do better by colluding among themselves and excluding the other bidders. That is, there exist allocations of the collusive surplus achieved by the ring coalition that make each subring better off than it could be in the absence of the ring.

In contrast to the study of bidding rings in SBSP and English auctions, McAfee and McMillan (1992) analyze bidding rings in a FPSB IPV game.⁴ They define two types of cartels: weak cartels in which cartel members cannot make side payments, and strong cartels in which side payments are permitted. For weak cartels, they prove that the optimal collusive scheme (for a large class of valuation distribution functions) is identical bidding. All bidders with valuations above the reserve price bid the reserve price, and those with valuations below the reserve price submit bids equal to zero. They show that the optimal collusive scheme is inefficient, since the winning bidder is picked at random by the auctioneer out of all the bidders submitting identical bids. Thus, the winning bidder is not necessarily the bidder with the highest valuation for the object.

For strong cartels, McAfee and McMillan (1992) demonstrate that an efficient, optimal mechanism exists for all-inclusive rings. They show this mechanism can be implemented by having the bidders hold a first "knockout" FPSB auction, and then having the winner of this knockout round participate in the actual auction if that winner's bid in the knockout round exceeds the reserve price. The winning cartel member then pays each of the losers an equal share of the difference between her bid in the prior, knockout auction and the reserve price.

McAfee and McMillan (1992) characterize optimal collusive mechanisms when members of the bidding ring report their valuations to a direct, incentive-compatible mechanism, for example, a "center." McAfee and McMillan assume the center specifies and enforces the bids. Marshall and Marx (2007) and Lopomo, Marx, and Sun (2011) extend the model of McAfee and McMillan (1992) by assuming the center cannot control the bids ring members submit at the auction, but the center can enforce side payments between ring members. Lopomo, Marx, and Sun (2011) show in this case that no collusive mechanism exists that improves bidders' payoffs relative to noncooperative

⁴ McAfee and McMillan (1992) focus on private value auctions because the optimal ring mechanism in the pure common value case is trivial. Efficiency is attained regardless of which member gets the right to bid in the seller's auction. Thus, an all-inclusive ring can use some exogenous method to allocate the right to one of its members, such as a random allocation with equal probability weights, and ask each bidder to report her valuation.

bidding even if side payments that depend only on ring members' reported valuations are allowed.

Lopomo, Marshall, and Marx (2005) show that in ascending-bid auctions, ring members can have an incentive to bid in ways that can cause the ring to win the item even though a non-ring bidder has the highest valuation. The ring mechanism may cause a ring member's reported valuation to increase her payment if another member of the ring wins the auction. In this case, ring members have an incentive to report valuations in excess of their actual valuations. This causes the ring's highest valuation to increase, so that the ring may win the auction even though a non-ring bidder has the highest valuation. Similarly, if the ring uses a knockout auction, ring members may have an incentive to bid in excess of their valuations if their payment in the knockout auction depends on their bids.

Lopomo, Marshall, and Marx's inefficiency result depends on two assumptions. First, there is no preauction communication among the ring members regarding their valuations. Second, the ring must *ex post* balance its budget. If the first assumption does not hold, then the results of Mailath and Zemsky (1991) show that an efficient explicit collusion mechanism exists. If the second assumption does not hold, then a modified version of the collusive mechanism in Graham and Marshall (1987) yields efficient explicit collusion without preauction communication.

Finally, Hendricks, Porter, and Tan (2008) analyze whether efficient collusion is possible in FPSB common value auctions. Assuming a ring forms, efficiency is trivially obtained since the value of the item by definition is the same, that is, common, to all bidders. So which particular ring member acquires the item is irrelevant from the perspective of the efficiency of the collusion. However, Hendricks, Porter, and Tan show that the ring may be unable to form in the first place, even if the cartel is legal.

Hendricks, Porter, and Tan show that in a common value auction, buyers who receive high signals regarding the true value of the object may prefer not to join a bidding ring. In the absence of a bidding ring, a buyer's only expected payment is to the seller in the main auction. In the presence of a bidding ring, a buyer's expected payment is the sum of two components: (1) the payment to the seller in the main auction and (2) the payment to other buyers in the knockout round. The expected payment of a buyer with a high signal to the seller in the main auction is lower with the bidding ring. However, with the bidding ring, low-signal bidders, free from worry regarding the winner's curse, bid aggressively in the knockout round. Thus, the expected sum of (1) the payment to the seller in the main auction in the presence of a bidding ring and (2) payments to other buyers in the knockout round may be greater than the expected payment to the seller in the main auction in the absence of a bidding ring. In this case, a buyer who receives a high signal regarding the true value of the object prefers not to join a bidding ring.

In summary, the literature on collusion in one-shot auctions establishes that bidders can collude efficiently in an IPV environment if they can engage in preauction communication and make side transfers to each other. Communication is essential because members of the ring have to reveal their private values in order to allocate the object to

the member with the highest value. Side transfers are needed to ensure that the members tell the truth and bid accordingly.

20.2.2.2. *Bid Rigging in Repeated Auctions*

The theory of repeated games suggests that repeated play makes it easier for bidders to collude, since they can condition their behavior on bids and enforce collusive outcomes by threatening to respond to deviations with competitive bidding. This form of collusion is known as tacit collusion, and it is not illegal. The celebrated folk theorem establishes that sufficiently patient players can construct a self-enforcing scheme in which they act as would a single firm and, thereby, achieve first-best collusive profits. However, this result assumes that players do not have private information. Therefore, an important question addressed in the literature on collusion in repeated auctions is whether an all-inclusive ring in an IPV environment can earn first-best collusive profits when members do not communicate or make side transfers to each other.

Before proceeding, we define efficient collusion in repeated auctions. A collusive scheme is efficient if, in each auction, the ring (1) bids if and only if the highest valuation of its members exceeds the reserve price; (2) never pays the seller more than the reserve price; and (3) assigns the object to the member with the highest valuation. Efficient collusion is basically a (random) bid rotation scheme in which each member gets to win whenever her valuation is the highest and exceeds the reserve price. By definition, efficient collusion maximizes the ex post surplus of the ring and generates first-best collusive profits.

In a seminal contribution, Skrzypacz and Hopenhayn (2004) establish that tacit collusion cannot achieve first-best collusive profits when bids are private information. They consider a model in which bidders bid repeatedly for identical objects sold sequentially over time. The seller publicly announces whether the object is sold and the identity of the winner but does not provide any information on the bids or the identities of the losing bidders. The bidders' private valuations are distributed independently and identically across bidders and auctions. The auction format can be any one of the standard first-price or second-price auctions. Bidders cannot communicate or make side transfers to each other. In this environment, the authors derive an upper bound on the profits that a ring can obtain and show that this upper bound is substantially smaller than the profits in an efficient collusive scheme.

The authors also characterize the types of collusive schemes that can achieve the upper bound on profits. They first note that any scheme in which, following every public history, bidders adopt symmetric bidding strategies does no better than a simple bid rotation scheme. This result corresponds to the weak cartel result of McAfee and McMillan (1992). That is, symmetric continuation values imply that the bidders share equally in future punishments and rewards so there are no transfers and, in that case, the best the ring can do is bid rotation. The schemes that do better than bid rotation involve treating bidders differently based on their history of wins and losses. The idea is to punish bidders that have more wins with a lower probability of winning, and even possible exclusion, in future auctions. As a result, winners have lower continuation values than

losers, and these differences work like transfers. But, despite these transfers, the lack of communication reduces collusive profits below the efficient level.

Athey, Bagwell, and Sanchirico (2004) explore this trade-off between efficiency and profits in an infinitely repeated Bertrand game with inelastic demand. The firms' costs are independently and identically distributed across firms and periods. The stage game is equivalent to a procurement auction since the firm that sets the lowest price wins the entire market. Consequently, the model is essentially the same as Skrzypacz and Hopenhayn except that bids are publicly observed. The authors characterize the set of *symmetric* equilibria and obtain two striking results. First, they show that, in any collusive equilibrium in which bidders use monotone bid functions, the ring achieves efficiency (due to the sorting of types) but earns profits that are no better than those generated by the noncooperative equilibrium. Second, the symmetric scheme that yields the highest profit to the ring involves rigid pricing; that is, bidders bid the reserve price and the seller randomly chooses a winner. This result is similar to the result obtained by Skrzypacz and Hopenhayn and generalizes the weak cartel result of McAfee and McMillan (1992) to repeated auctions with observable bids.

If tacit collusion is inefficient, can communication among the members of the ring regarding their private information help them achieve first-best collusive profits? And, if so, what does the collusive scheme look like? These questions have been addressed in several papers.

Athey and Bagwell (2001) consider a simplified version of the infinitely repeated Bertrand game discussed above. There are only two firms and their costs in each period are identical and independent Bernoulli random variables: "high" with probability η and "low" with probability $1 - \eta$. As noted above, their model is equivalent to a repeated procurement auction. The firms explicitly, albeit illegally, collude. The authors assume that (1) the firms can communicate with each other in each period regarding their current costs, but (2) they cannot make side payments (so as to reduce the probability of detection by the antitrust authorities). Prices, that is, bids, are public information.

The main result of their paper is that a sufficiently patient ring can achieve the first-best collusive profits. The efficient collusive scheme is similar to one that Skrzypacz and Hopenhayn construct to show that a ring can do better than bid rotation. In each period, the high-cost bidder does not bid and the low-cost bidder bids the reserve price. However, the high-cost bidder is favored in future auctions with a higher probability of being selected in the event that both bidders have same cost. These future "market-share" favors have no impact on efficiency since they only occur in periods in which the firms' realized costs are the same. But the market-share favors mean that the continuation value of the low-cost bidder is lower than the continuation value of the high-cost bidder, and this difference acts like a transfer. The only issue is whether the magnitudes of the feasible market-share favors are sufficient to incentivize bidders to cooperate and report their costs truthfully. The authors show that, if the firms do not discount future profits too heavily, the answer is yes.

Aoyagi (2003) analyzes a model of infinitely repeated auctions consisting of two symmetric bidders who bid every period on a single indivisible good. The auction format can

be either first-price or second-price. The private signals of the bidders are real-valued and identically and independently distributed across bidders and auctions. The bidders report their private signals to a communication device, called a “center,” to coordinate their bids in each auction. After receiving the reported signals, the center subsequently instructs the bidders on how to bid in the auction. A “collusion scheme” consists of the center’s choice of instruction rule in every period as a function of the reports and the public histories. The public history is the sequence of instruction rules chosen by the center in past auctions and the bids in those auctions. A collusion scheme is an equilibrium if, for each bidder, telling the truth about her valuations is incentive compatible, and it is rational to adhere to the instructions. Aoyagi studies a class of dynamic bid rotation schemes with “grim trigger” punishments. In these schemes, the players begin in the collusion phase in which only one bidder is instructed to bid in a given stage auction. If a player deviates from the instructions, a punishment phase is triggered in which the one-shot Nash equilibrium results.

Aoyagi’s dynamic bid rotation scheme works as follows. In the first phase, the center uses the efficient instruction rule. That rule instructs the bidder with the higher valuation (based on the bidders’ reported valuations to the center) to bid the reserve price if and only if the valuation exceeds the reserve price. The center instructs the other bidder not to bid. The difficulty is that the efficient instruction rule is not incentive compatible. Each bidder has an incentive to overstate their reported valuations to the center in hopes of winning the auction at the reserve price. To solve this problem, the bid rotation scheme has a second phase in which the payoff to the bidder with the highest reported valuation is reduced relative to the other bidder with some positive probability. The instruction rule in this phase is incentive compatible. The bid rotation scheme proceeds in this second phase for a fixed number of periods before reverting to the first stage.

The collusive profits from Aoyagi’s scheme are not first-best but, as in Skrzypacz and Hopenhayn’s model, they are higher than the profits the ring would obtain from the scheme in which the bidders take turns winning at the reserve price independently of their values. The main difference between his model and that of Athey and Bagwell is the cardinality of the set of valuations. In the Athey and Bagwell model the set is finite, but in Aoyagi’s model the set is the unit interval. Thus, the kind of “market-share” favors that Athey and Bagwell use in their optimal collusion scheme is not possible in Aoyagi’s model, since the probability of both bidders having the same valuation is zero.

In a subsequent paper, Aoyagi (2007) uses the same model but with a finite type space to show that a ring can achieve the first-best collusive profits. The optimal collusion scheme is fully efficient, that is, each bidder’s equilibrium payoff is close to what she would get if the object were allocated at the reserve price to the highest-valuation bidder in every stage auction. Using a similar dynamic bid rotation scheme to that in his 2003 paper, which allows for bidder communication, Aoyagi identifies conditions under which an equilibrium collusion scheme is fully efficient.

One conclusion of the literature on repeated auctions is that bidders do not have to use side payments to implement the efficient bid rotation scheme. Strategies in which current losers are rewarded in future play can provide sufficient incentives for bidders to cooperate. A second conclusion is that sellers can make it difficult for bidders to achieve

first-best profits by not making bids public. In this case, bidders have to communicate to earn the full benefits from collusion, which makes a ring illegal and easier to detect. A third conclusion is that, when bidders are sufficiently patient, communication is sufficient to achieve first-best collusion profits.

20.3. BID RIGGING: EMPIRICAL STUDIES

20.3.1. Empirical Studies of Bid Rigging in Auctions

Porter and Zona (1993) examine bidding for state highway construction procurement auctions in Long Island, New York, from 1979 to 1985.⁵ The New York Department of Transportation (“DOT”) awarded approximately \$120 million in 186 separate highway contracts in this period. The DOT used the FPSB auction format in its highway construction procurement auctions. Porter and Zona first evaluate whether the characteristics of these procurement auctions would tend to facilitate collusion. The first characteristic of the auctions is the public nature of certain information. Prior to a given procurement auction, the DOT made public a “Plan Buyers List” that listed the firms that purchased the plans for that highway construction project. Thus, cartel members had knowledge of the set of potential bidders against whom they would bid before each procurement auction. On the day the winning (low) bidder was selected, the DOT publicly announced all bids and the identity of each bidder. As Porter and Zona note, this information allowed cartel members to detect deviations from an agreement. The second characteristic of the auctions is the DOT’s inelastic demand. Out of the 186 highway procurement auctions, 185 were ultimately funded and awarded to the low bidder. Thus, any increase in the winning (low) bid caused by successful collusion was captured as profits by the cartel. The third characteristic of the auctions is the structure of the market for highway construction services. Porter and Zona note that on the 25 largest construction jobs, the four largest firms accounted for 45% of the bids. The fourth characteristic of the market is the opportunity for the firms to communicate. Most of the bidders belonged to the same local trade associations, and joint bidding was allowed. The fifth characteristic was that the DOT tended to run its procurement auctions for larger jobs on a regular basis, at the beginning of each year. As Porter and Zona note, the regularity of the auctions would tend to make calculating the discounted present value of continued collusion more predictable. The final characteristic was that the highway construction firms were relatively homogeneous. They used the same technologies and purchased inputs from the same suppliers.

⁵ See also Doane et al. in this volume for a discussion of detecting explicit collusion.

Porter and Zona next identify five firms as possible members of a highway construction cartel. One of the firms was convicted in federal court of bid rigging on a highway construction job in Long Island in 1984, and the other four firms were unindicted coconspirators. The same five firms had been named as participants in bid-rigging conspiracies in other antitrust or racketeering suits in New York. Using this information, Porter and Zona divide their data into two sets: (1) bids from firms other than the five possible cartel firms (called bids from “competitive firms”) and (2) bids from the five possible cartel firms. They also restrict their data set to 75 road-paving jobs that had at least two bids from competitive firms.

Porter and Zona first estimate a regression model in which the dependent variable is the log of a firm’s bid. The independent variables consist of measures of a firm’s capacity and capacity utilization, and dummy variables indicating whether the firm’s headquarters was on Long Island, and for competitive firms, whether that firm previously had won a highway construction auction on Long Island. Using first the bids from competitive firms, Porter and Zona conclude that the estimated regression fits the data reasonably well and the coefficient estimates have the expected signs. In contrast, using bids from the cartel firms, many of the estimated regression coefficients do not have the expected signs. A Chow test for equality of the estimated coefficients in the two regressions rejects the null hypothesis that the coefficients in the two regressions are equal. On the basis of these results, Porter and Zona conclude that they can reject the null hypothesis of no bid rigging.

Finally, Porter and Zona estimate a second regression model that uses the same independent variables, but changes the dependent variable to the rank of a firm’s bid in a given auction, rather than the dollar value of its bid as used in their first regression model. They first estimate this model using bids from competitive firms. They estimate three regressions, one using all the bids, one using just the low bids, and another using all bids other than the low bid. Comparing the estimated coefficients in the three regressions, they cannot reject the null hypothesis of no model misspecification. On this basis, they cannot conclude that the bids from competitive firms are generated by different processes depending on whether the bids are low or not.

Porter and Zona then estimate the same regression model using bids from the cartel firms, again estimating three regressions, one using all the bids, one using just the low bids, and one using all bids other than the low bid. Comparing the estimated coefficients in the three regressions, they reject the null hypothesis of no model misspecification. On this basis, they find that they can conclude that the bids from cartel firms are generated by different processes depending on whether the bids are low or not. They reject the null hypothesis of no phantom bidding and conclude that these five firms engaged in bid rigging.

Baldwin, Marshall, and Richard (1997) examine the winning bids and characteristics for 108 oral auctions for timber sold by the Forest Service in the Pacific Northwest from 1975 to 1981. Baldwin, Marshall, and Richard described several salient features of the auctions. First, the Forest Service’s reserve prices in the auctions are very low. Second, old-growth timber is quite heterogeneous. Bidders invest substantial resources to

develop their own valuations of specific forest tracts. Third, second-growth or young timber is relatively homogenous and of lower quality than old-growth timber. Fourth, despite the fact that timber is extremely heavy per dollar of value, firms with mills within 100 miles of a given forest tract have approximately the same costs of cutting and transporting the logs. Fifth, the mills differ substantially in terms of their efficiency, which is measured by the quantity of merchantable board feet that can be extracted from a given log. Logging firms closely guard information regarding the efficiencies of their mills. Finally, prior to each oral auction, the Forest Service publicly releases the number of qualified bidders, and after each auction, the Forest Service makes public the quantity of timber purchased by each firm. In contrast, private sellers of timber do not make such information public.

Given the importance of private information held by logging firms regarding the efficiency of their mills, Baldwin, Marshall, and Richard use an IPV model. In order to reduce the heterogeneous nature of the product, they restrict their data set to auctions for second-growth timber. Baldwin, Marshall, and Richard attempt to determine whether, after controlling for demand conditions, variations in winning (low) bids are better explained by collusion or by variations in the supply of timber. They estimate several models: (1) the noncooperative (i.e., noncollusive) model without considering supply effects; (2) the collusive model without supply effects; (3) the noncooperative model with supply effects, and two nesting models that include both collusion and supply effects. They conclude that their results strongly suggest that variation in the winning bids is best explained by collusion and not by noncooperative behavior or changes in supply conditions. They estimate that collusion reduces the auction revenues received by the Forest Service by approximately 7.9% across all the auctions. However, for a subset of 13 auctions that had particularly low winning bids, the loss in revenues to the Forest Service was approximately 52.9%.

Athey, Levin, and Seira (2011) examine Forest Service auctions that occurred in Idaho-Montana and California between 1982 and 1990. (See also Athey and Levin, 2001, for analyses of Forest Service auctions.) Since the Forest Service used both open and sealed bids, Athey, Levin, and Seira are able to test how prices and revenues vary between the two auction formats. They show that sealed bid auctions attract more small bidders and tend to yield higher revenues to the Forest Service. Bidders may be loggers (i.e., small firms without manufacturing capacity) or mills (i.e., larger firms with manufacturing capacity). They estimate a structural IPV model that allows for entry on the part of bidders, and then they use that model to predict prices and revenues in Idaho-Montana and California with open or sealed bid auctions. They find that average open auction sales prices and average sealed bids in California are close to their predicted values, assuming firms behave competitively. This competitive result also holds for average sealed bids in Idaho-Montana. However, average open auction sales prices in Idaho-Montana are statistically different than both predicted competitive and collusive prices. Athey, Levin, and Seira (p. 251) conclude that “mildly cooperative behavior on the part of participating mills appears to provide a better match than either the competitive or fully collusive extremes.” To further test for the presence of collusion in

Idaho-Montana auctions, the authors compare average sales prices in both open and sealed bid auctions as a function of the number of mills. When the number of mills is zero or one (recall that bidders can be loggers or mills), their predicted competitive prices are close to the actual prices. However, when the number of mills is two or more, their predicted competitive prices are substantially above the actual prices. They conclude this finding is consistent with collusive bidding.

Given these results, Athey, Levin, and Seira investigate the welfare consequences of the Forest Service using either open or sealed bidding exclusively. They use two alternate specifications of mill bidding behavior: competitive behavior and collusive behavior 18% of the time, that is, the frequency of collusive behavior that best matches the observed open auction prices in Idaho-Montana. With the assumption of competitive behavior, they find that average revenues to the Forest Service are nearly the same with open or sealed bid auctions. However, with the more realistic assumption of collusive behavior (18% of the time), they find that the average revenues to the Forest Service would be substantially higher with sealed bid auctions than with open auctions. This result follows from their finding that sealed bid auctions encourage more participation from smaller bidders.

Porter and Zona (1999) study procurement auctions conducted by Ohio school districts for milk. The State of Ohio charged 13 dairies with bid rigging in the period 1980–1990. Porter and Zona prepared expert reports on behalf of the plaintiff in *State of Ohio v. Louis Trauth Dairies, Inc. et al.* Their article analyzes whether the behavior of three of the alleged conspirators located near Cincinnati is more consistent with competition or collusion. Using data collected in the litigation, they create a sample of bids for a control group of firms not accused of bid rigging and a sample of bids for the three firms accused of bid rigging.

They first estimate a probit regression to calculate the probability that a control firm will submit a bid as a function of the distance between a relevant school district and that firm's nearest milk processing plant. They find that the probability falls as distance increases. For example, the probability that a control firm with a plant very near a school district will bid in a milk procurement auction run by that district exceeds 50%, but falls essentially to zero when the distance exceeds 75 miles. They next estimate an OLS regression to calculate how a control firm's bid changes as distance increases. They find that, all else equal, a control firm's bid increases by approximately 0.5 cents per half pint when the distance increases by 50 miles. Given that the average price per half pint was approximately 13 cents in this period, this amounts to approximately a 4% increase in the control firms' bid.

Comparing these results for the group of control firms to the three alleged bid-rigging firms, Porter and Zona find that the alleged conspirators bid more frequently than the control group model predicts at longer distances, for example, in excess of 60 miles. They also find that two of the three alleged conspirators actually bid less the further away their plants are to the relevant school district. Porter and Zona interpret these results as showing an "inverted price umbrella," with higher bids in school milk auctions close to the conspirators' plants and lower bids at more distant locations. They conclude

the evidence is consistent with a local conspiracy in the Cincinnati area and inconsistent with competitive behavior. Finally, their empirical results show that the bid rigging increased the prices paid for milk by school districts by approximately 6.5%.

Bajari and Ye (2003) use auction theory and empirical methods to distinguish competitive from collusive bidding behavior. They develop a general procurement auction model with asymmetric bidders, for example, bidders with different costs. They then derive a series of necessary and sufficient conditions for a distribution of bids to be consistent with bids generated by a model with competitive bidding. For example, one condition is that, given publicly available information that affects all firms' bids, competitive bidding necessarily implies that the bids of any two firms must be independently distributed. Bajari and Ye use this condition to evaluate whether the bids made by construction firms in the Midwest are more consistent with competition or collusion. They regress firms' bids on several explanatory variables, for example, distance from a firm's location to the job site; the minimal distance of its rivals to the job site; the firm's capacity; and the maximum free capacity of its rivals. They test whether the regression residuals for any pair of firms are correlated. If they reject the null hypothesis of no correlation, then the two firms' bids are not independently distributed and so cannot be consistent with competition. They reject the null hypothesis for one pair of firms that bid against each other a large number of times.

Bajari and Ye derive a second condition stating that competitive bidding necessarily implies that each firm's estimated coefficient for any given explanatory variable in the bid regression should be the same. Their test results identify a specific pair of firms that fails this "exchangeability" test. Interestingly, one of the two firms in this pair is also one of the two firms in the pair that fails the "conditional independence" test. Bajari and Ye conclude that these three firms constitute a candidate cartel. However, they also emphasize that their test results could be explained by conduct other than collusion. They conclude that an economist finding such results should consult with industry experts to make sure the inconsistencies with the competitive bidding model do not result from ignorance of the industry cost structure.

Rather than attempting to distinguish firms engaged in bid rigging from firms acting competitively, Pesendorfer (2000) analyzes the characteristics of two documented school milk cartels in Florida and Texas from 1980 to 1991. All the firms in his study pled guilty to price-fixing allegations brought by the US Department of Justice. Pesendorfer first shows that the market shares of firms in the Florida cartel fluctuated more than the market shares of firms in the Texas cartel. He examines several noncollusive explanations for this fact: (1) less potential competition in Texas; (2) some cartel members in Texas may bid on only a small subset of the school milk contracts; (3) a decline in the number of contracts per school district in Florida; and (4) an individual cartel member in Florida may have had costs that were correlated over time (e.g., high cost realizations in one year followed by low realizations the next year) that would cause its market share to fluctuate substantially over time. Pesendorfer finds that the empirical evidence does not support any of these noncollusive explanations. He concludes that the best explanation for the larger market share fluctuations in Florida is that cartel members in that

state used side payments, rather than a market allocation agreement, to operate their cartel. That is, they operated a strong cartel. In contrast, the Texas cartel operated by dividing the state into different geographic regions and specifying which firms should win in which regions. That is, they operated a weak cartel.

Pesendorfer also analyzes theoretically the efficiency of weak cartels. Recall that an auction mechanism is efficient if the mechanism allocates the object to the bidder with the highest valuation. Or alternatively, a cartel mechanism is efficient if it designates the member with the lowest cost to submit the lowest cartel bid. Pesendorfer shows that weak cartels cannot be efficient as long as the number of procurement contracts offered simultaneously by the buyer is finite. However, he shows that a weak cartel can use a mechanism (the “Ranking Mechanism”) in which each cartel member announces a ranking of the procurement contracts according to her costs. The member who ranks a contract highest (assigns the highest preference) will be the sole cartel bidder for that contract. If several cartel members rank a contract at the same position, the sole bidder can be determined by a coin toss. Pesendorfer shows the Ranking Mechanism is incentive compatible (i.e., cartel members will truthfully reveal their costs). Moreover, the mechanism converges to the optimal, efficient outcome as the number of contracts increases. Finally, Pesendorfer notes that 136 contracts are awarded every year in Texas, which leads him to conclude that the weak Texas cartel is likely almost efficient.⁶

20.3.2. Experimental Studies of Bid Rigging in Auctions

Isaac and Plott (1981) is the first experimental study of explicit collusion in an auction format. Their study addresses two primary questions. First, antitrust laws presume that opportunities to conspire, attempts to conspire, and successful conspiracies are closely linked. Is this presumption correct when the organizational costs of conspiracy are low? Second, which economic model of the price and output effects of price-fixing conspiracies best explains the observed behavior in their experiments? Isaac and Plott perform seven experiments using a double-auction mechanism in which buyers and sellers both make price offers. They first run three auctions in which neither buyers nor sellers have the opportunity to collude explicitly. They then run two auctions in which sellers are allowed to collude explicitly, without the knowledge of buyers. Finally, they run two auctions in which buyers are allowed to collude explicitly, without the knowledge of sellers. The explicit collusion was restricted in that participants could not reveal their private valuations or discuss side payments.

Their basic finding is that, compared to the no-collusion auctions, both the seller- and buyer-collusion auctions are less competitive. Regarding their first question, Isaac and Plott reach several conclusions. First, in an environment with low costs of explicit collusion, buyers and sellers attempt to reach an agreement. Second, the attempts to reach an

⁶ See Abrantes-Metz et al. (2006) for a discussion of the use of econometric screening methods to detect collusion in the retail gasoline industry.

agreement result in a coordinated pricing strategy. Third, the pricing conspiracies have substantial, anticompetitive effects on equilibrium prices and quantities. With respect to price dynamics observed in the experiments, Isaac and Plott note that price changes in the three noncollusive experiments were always in the direction of the predicted equilibrium price. In contrast, prices in both the seller- and buyer-collusion auctions were more erratic, often moving away from the predicted cartel price. The average price change from period to period was also larger in the collusive experiments than in the noncollusive experiments.

Finally, regarding their second question, Isaac and Plott conclude that the textbook cartel model best describes the behavior observed in the collusive experiments. However, the prices in these experiments generally did not converge to the predicted cartel price, as participants had difficulty maintaining agreed-upon prices. In follow-on research, Isaac, Ramey, and Williams (1984) show that the double-auction mechanism used by Isaac and Plott (1981) tends to make the enforcement of price-fixing agreements more difficult than posted offer prices.

Isaac and Walker (1985) study sealed-bid auctions experimentally. They investigate two questions. First, in a series of FPSB auctions for identical items, are winning bids affected by revealing the losing bids from prior auctions? Second, does the ability of buyers to collude explicitly reduce the winning bids? They conducted a total of 30 experiments. In 10 of the experiments, explicit collusion was not allowed and buyers were not informed of losing bids in prior auctions. In 10 of the experiments, explicit collusion was not allowed and buyers were informed of losing bids in prior auctions. In the final 10 experiments, explicit collusion was allowed. In five of the collusive auctions buyers were not informed of losing bids in prior auctions, and in the other five collusive auctions buyers were informed of the losing bids. With respect to their two questions, Isaac and Walker find that the revelation of losing bids reduces the prices paid by buyers in subsequent auctions. They also find that the ability to collude explicitly reduces winning bids. However, the collusive prices were not affected by whether losing bids were revealed.

Kwasnica (2000) evaluates experimentally the strategies used by auction participants to collude. Using a series of simultaneous, first-price auctions, he addresses three empirical questions: (1) Do bidders form cooperative agreements in simultaneous first-price auctions? (2) If so, what types of strategies do they utilize? (3) What effect do these strategies have on the outcome of the auction? The auction format consisted of 10 separate experiments. In each experiment a total of five objects were sold to five buyers in five simultaneous single-unit first-price auctions. This format was chosen to facilitate collusion. The auctions ran for 17–22 periods. In periods 1–5 participants were not allowed to communicate, but they were allowed to communicate thereafter. The communication was restricted in that participants could not reveal their private valuations or discuss side payments. In six of the 10 experiments the bidders had symmetric valuations, and in the remaining four experiments they had asymmetric valuations. In the latter, one bidder's valuations were substantially higher than the symmetric valuations of the other four bidders. Finally, in six of the auctions, participants were informed of the winning bid and the identity of the winning bidder, while in the remaining four auctions participants were only informed of the winning bid.

Consistent with Issac and Walker (1985), Kwasnica finds that bidders formed cooperative agreements in all 10 experiments. More interestingly, he finds that in seven of the 10 experiments participants used a bid rotation mechanism. They were able to obtain approximately 90% of the total surplus in the auctions. In other words, they achieved approximately 90% of the maximum possible efficiency, which would occur if the winning bidder had the highest valuation in each auction. Kwasnica concludes that the Ranking Mechanism proposed by Pesendorfer (2000) best describes the behavioral strategy chosen by the bidders.

Finally, in three of the experiments, Kwasnica finds that bidders deviated from incentive compatible strategies. Instead, they selected a linear bid reduction strategy; for example, they agreed to bid 1% of their valuations. Participants used a linear bid reduction strategy only in experiments in which they had symmetric valuations and knew the identity of the winning bidder. Although a linear bid reduction strategy was not incentive compatible, under the experimental framework such a strategy led to higher profits than the incentive compatible, bid-rotation strategy. Kwasnica concludes that asymmetry of bidders' valuations and less information leads bidders to revert to incentive compatible strategies.⁷

20.4. DETERRING BID RIGGING

In theory, explicit collusion in auctions always can be prevented. The difficulty is that the information required to do so is essentially impossible to obtain. Che and Kim (2006, 2009) show how to convert any given auction mechanism into a collusion-proof auction, that is, an auction in which a bidding ring cannot earn profits. In such a collusion-proof auction, the seller's expected total revenue will be the same as in a revenue-maximizing auction in the absence of collusion. However, as discussed by Marshall and Marx (2012), the Che and Kim collusion-proof mechanism is impractical, as it requires the following:

1. The mechanism requires all bidders, including losing bidders, to make payments to the auctioneer. However, the auctioneer may cheat by demanding higher-than-required payments from the bidders. This means that all the bidders must be able to verify that the auctioneer has implemented the mechanism correctly.
2. To establish the collusion-proof mechanism, the auctioneer must know the number of bidders and the set of types of the bidders, that is, their valuations, as well as the identities and types of two of the bidders in the ring. In Che and Kim (2006), bidders must agree to participate in the mechanism prior to the ring-formation game. In addition, the identities of the bidders that will have the

⁷ Other experimental studies of explicit collusion in auctions include Davis and Wilson (2002) and Phillips, Menkhaus, and Coatney (2003).

opportunity to form a ring (only one ring is allowed to form) is fixed prior to the bidders agreeing or not to participate in the collusion-proof mechanism. As shown in Che and Kim (2009), under certain conditions, similar results hold when bidders first decide whether to collude and then decide whether to participate in the auctioneer's collusion-proof mechanism.

3. The bidding ring must submit a bid for every ring member, despite possible incentives for the ring to suppress some members' bids.

Given these practical drawbacks to the Che and Kim collusion-proof mechanism, economists have investigated other ways to deter explicit collusion in auctions. The literature offers several general suggestions (Klemperer, 2002; Marshall and Marx, 2009; Marshall and Marx, 2012). First, the use of FPSB auctions instead of English auctions should reduce collusion. Recall that bid-rigging agreements are stable in English auctions since no bidder has an incentive to cheat on the agreement since the cartel will bid up to the highest valuation of its members. In contrast, in FPSB auctions, the cartel must reduce its bid below the highest valuation of its members in order to earn positive expected profits. This reduction in the bid provides ring members the incentive to cheat on the agreement by outbidding the cartel. Hu, Offerman, and Onderstal (2011) provide experimental results that suggest using the Amsterdam second-price auction, which combines aspects of both FPSB and English auctions, to reduce collusion.⁸ Second, the auctioneer should limit the information provided to bidders regarding the number of bidders, their identities, their bids, including their losing bids if a similar auction will occur in the future. Finally, the auctioneer should not hold auctions at regular intervals and for relatively small items; instead the auctioneer should hold auctions at longer, irregular intervals and for relatively large items, for example, packages of separate items. This increases the costs of creating and maintaining a bidding ring by raising the gains to cheating.

20.5. CONCLUSIONS AND OPEN QUESTIONS

In his classic article on collusion, Stigler (1964) observed that “the system of sealed bids, publicly opened with full identification of each bidder's price and specifications,

⁸ The Amsterdam second-price auction has two phases (Goeree and Offerman, 2004). In phase 1, the auctioneer raises the price successively and bidders drop out when the price exceeds their individual valuation. This process continues until two bidders remain. The price at which the third-highest bidder drops out defines the endogenous reserve price or “bottom” price for the second phase. In phase 2, both remaining bidders independently submit sealed bids, which must be at least as high as the bottom price. The highest bidder wins and pays a price equal to the second-highest sealed bid. Both bidders in the phase two receive a premium, which equals a fraction $0 < \alpha < 0.5$ of the difference between the second-highest sealed bid and the bottom price.

is the ideal instrument for the detection of price-cutting” and argued that collusion is always more effective when the market is more transparent. Open auctions are often more transparent than sealed bid auctions because bidders learn about the participation and bids of their rivals in real time. But, even in sealed bid auctions, the seller can choose how much information to make public before and after the auction. The less information the seller reports, the harder it is for the ring to achieve the gains from collusion. Indeed, the conclusion of the literature surveyed in this paper is that a seller should use a first-price, sealed bid auction and reveal as little information as possible.

This conclusion, however, is based mainly on analyses of private value, single-item auctions. In common value auctions or multi-item auctions, the lack of transparency in the price-formation process can lead to inefficient allocations and lower revenues. For example, in common value auctions, bidders bid more aggressively in ascending bid auctions than in sealed bid auctions because they are able to learn about their rivals’ information, mitigating the effects of the winner’s curse. In auctions of spectrum licenses or oil and gas leases where spatial complementarities are important, bidders can more easily acquire their desired bundles and face less risk doing so in a multiround ascending bid auction than in a single round of sealed bids. Thus, in most situations, the choice of auction mechanism is likely to involve an important trade-off. On the one hand, more transparent auction mechanisms can yield greater efficiency and revenues when bidders behave competitively; on the other hand, they are also more vulnerable to collusion. A greater understanding of this trade-off is one of the big open questions in auction design.

The theory literature has focused almost exclusively on the case of an all-inclusive ring.⁹ However, in practice, most rings involve subsets of bidders—not all bidders. Extending the theory to include partial rings is an important, although difficult, topic. Partial rings introduce asymmetries into the auction, and characterizing equilibrium bidding when the bidder types are drawn from asymmetric distributions is complicated.

Finally, an important policy issue is to examine the rules under which the federal government in the United States operates as a buyer or seller. The Federal Acquisition Regulations, the Freedom of Information Act, and other regulations mandate transparency. The constraints that these rules impose on the choice of auction design may be quite costly.

⁹ Two exceptions are McAfee (1994) and McAfee and McMillan (1992), who show that, in first-price environments, the equilibrium ring in a cartel formation game generally does not include all of the bidders.

REFERENCES

- Abrantes-Metz, R., L. Froeb, J. Geweke, and C. Taylor. 2006. *A Variance Screen for Collusion*. *International Journal of Industrial Organization* 24: 467–86.
- Aoyagi, M. 2003. *Bid Rotation and Collusion in Repeated Auctions*. *Journal of Economic Theory* 112: 79–105.
- Aoyagi, M. 2007. *Efficient Collusion in Repeated Auctions with Communication*. *Journal of Economic Theory* 134: 61–92.
- Asker, J. 2010. *A Study of the Internal Organization of a Bidding Cartel*. *American Economic Review* 100: 724–62.
- Athey, S., and K. Bagwell. 2001. *Optimal Collusion with Private Information*. *Rand Journal of Economics* 32: 428–65.
- Athey, S., and J. Levin. 2001. *Information and Competition in U.S. Forest Service Timber Auctions*. *Journal of Political Economy* 109: 375–417.
- Athey, S., J. Levin, and E. Seira. 2011. *Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions*. *Quarterly Journal of Economics* 126: 207–57.
- Bajari, P., and L. Ye. 2003. *Deciding between Competition and Collusion*. *Review of Economics and Statistics* 85: 971–89.
- Baldwin, L., R. Marshall, and J-F. Richard. 1997. *Bidder Collusion at Forest Service Timber Sales*. *Journal of Political Economy* 105: 657–99.
- Blume, A., and P. Heidhues. 2006. *Private Monitoring in Auctions*. *Journal of Economic Theory* 131: 179–211.
- Boone J., R. Chen, J. Goeree, and A. Polydoro. 2009. *Risky Procurement with an Insider Bidder*. *Experimental Economics* 12: 417–36.
- Brown, A., C. Plott, and H. Sullivan. 2009. *Collusion Facilitating and Collusion Breaking Power of Simultaneous Ascending and Descending Price Auctions*. *Economic Inquiry* 47: 395–424.
- Brusco, S., and G. Lopomo. 2002. *Collusion via Signalling in Simultaneous Ascending Bid Auctions with Heterogeneous Objects, with and without Complementarities*. *Review of Economic Studies* 69: 407–36.
- Che, Y.-K., and J. Kim. 2006. *Robustly Collusion-Proof Implementation*. *Econometrica* 74: 1063–107.
- Che, Y.-K., and J. Kim. 2009. *Optimal Collusion-Proof Auctions*. *Journal of Economic Theory* 144: 565–603.
- Compte, O., A. Lambert-Mogiliansky, and T. Verdier. 2005. *Corruption and Competition in Procurement Auctions*. *Rand Journal of Economics* 36: 1–15.
- Cramton, P., and J. Schwartz. 2000. *Collusive Bidding: Lessons from the FCC Spectrum Auctions*. *Journal of Regulatory Economics* 17: 229–52.
- Dasgupta, P., and E. Maskin. 2000. *Efficient Auctions*. *Quarterly Journal of Economics* 115: 341–88.
- Davis, D., and B. Wilson. 2002. *Collusion in Procurement Auctions: An Experimental Examination*. *Economic Inquiry* 40: 213–30.
- Fabra, N. 2003. *Tacit Collusion in Repeated Auctions: Uniform versus Discriminatory*. *Journal of Industrial Economics* 51: 271–93.
- Goeree, J., and T. Offerman. 2003. *Competitive Bidding in Auctions with Private and Common Values*. *Economic Journal* 113: 598–613.
- Goeree, J., and T. Offerman. 2004. *The Amsterdam Auction*. *Econometrica* 72: 281–94.

- Graham, D., and R. Marshall. 1987. *Collusive Bidder Behavior at Single-Object Second-Price and English Auctions*. *Journal of Political Economy* 95: 1217–39.
- Graham, D., R. Marshall, and J.-F. Richard. 1990. *Differential Payments within a Bidder Coalition and the Shapley Value*. *American Economic Review* 80: 493–510.
- Haan, M., L. Schoonbeek, and B. Winkel. 2009. *Experimental Results on Collusion: The Role of Information and Communication*. In J. Hinloopen and H-T Normann, *Experiments and Competition Policy*, 9–33. Cambridge: Cambridge University Press.
- Hendricks, K., and R. Porter. 2007. *An Empirical Perspective on Auctions*. In M. Armstrong and R. Porter, *Handbook of Industrial Organization*, vol. 3: 2073–143. Amsterdam: Elsevier.
- Hendricks, K., R. Porter, and G. Tan. 2008. *Bidding Rings and the Winner's Curse*. *Rand Journal of Economics* 39: 1018–41.
- Hinloopen, J., and S. Onderstal. 2010. *Collusion and the Choice of Auction: An Experimental Study*. Tinbergen Institute Discussion Paper.
- Hu, A., T. Offerman, and S. Onderstal. 2011. *Fighting Collusion in Auctions: An Experimental Investigation*. *International Journal of Industrial Organization* 29: 84–96.
- Isaac, R., and C. Plott. 1981. *The Opportunity for Conspiracy in Restraint of Trade*. *Journal of Economic Behavior and Organization* 2: 1–30.
- Isaac, R., V. Ramey, and A. Williams. 1984. *The Effects of Market Organization on Conspiracies in Restraint of Trade*. *Journal of Economic Behavior and Organization* 5: 191–222.
- Isaac, R., and J. Walker. 1985. *Information and Conspiracy in Sealed Bid Auctions*. *Journal of Economic Behavior and Organization* 6: 139–59.
- Ishii, R. 2009. *Favor Exchange in Collusion: Empirical Study of Repeated Procurement Auctions in Japan*. *International Journal of Industrial Organization* 27: 137–44.
- Klemperer, P. 1999. *Auction Theory: A Guide to the Literature*. *Journal of Economic Surveys* 13: 227–86.
- Klemperer, P. 2002. *What Really Matters in Auction Design*. *Journal of Economic Perspectives* 16: 169–89.
- Klemperer, P. 2008. *Competition Policy in Auctions and "Bidding Markets."* In P. Buccirossi, *Handbook of Antitrust Economics*, 583–624. Cambridge, MA: MIT Press.
- Kwasnica, A. 2000. *The Choice of Cooperative Strategies in Sealed Bid Auctions*. *Journal of Economic Behavior and Organization* 42: 323–46.
- Kwasnica, A., and K. Sherstyuk. 2007. *Collusion and Equilibrium Selection in Auctions*. *Economic Journal* 117: 120–45.
- Li, J., and C. Plott. 2008. *Tacit Collusion in Auctions and Conditions for Its Facilitation and Prevention: Equilibrium Selection in Laboratory Experimental Markets*. *Economic Inquiry* 47: 425–48.
- Lopomo, G., R. Marshall, and L. Marx. 2005. *Inefficiency of Collusion at English Auctions*. *Contributions in Theoretical Economics* 5, article 4.
- Lopomo, G., L. Marx, and P. Sun. 2011. *Bidder Collusion at First-Price Auctions*. *Review of Economic Design* 15: 177–211.
- Mailath, G., and P. Zemsky. 1991. *Collusion in Second Price Auctions with Heterogeneous Bidders*. *Games and Economic Behavior* 3: 467–86.
- Marshall, R., and L. Marx. 2007. *Bidder Collusion*. *Journal of Economic Theory* 133: 374–402.
- Marshall, R., and L. Marx. 2009. *The Vulnerability of Auctions to Bidder Collusion*. *Quarterly Journal of Economics* 124: 883–910.
- Marshall, R., and L. Marx. 2012. *The Economics of Collusion: Cartels and Bidding Rings*. Cambridge, MA: MIT Press.

- McAfee, R. P., and J. McMillan. 1987. *Auctions and Bidding*. *Journal of Economic Literature* 25: 699–738.
- McAfee, R. P., and J. McMillan. 1992. *Bidding Rings*. *American Economic Review* 82: 579–99.
- McAfee, R.P., D. Vincent, M. Williams, and M. Havens. 1993. *Collusive Bidding in Hostile Takeovers*. *Journal of Economics and Management Strategy* 2: 449–82.
- Milgrom, P., and R. Weber. 1982. *A Theory of Auctions and Competitive Bidding*. *Econometrica* 50: 1089–122.
- Myerson, R. 1981. *Optimal Auction Design*. *Mathematics of Operations Research* 6: 58–73
- Pesendorfer, M. 2000. *A Study of Collusion in First-Price Auctions*. *Review of Economic Studies* 67: 381–411.
- Phillips, O., and D. Menkhaus. 2009. *Maintaining Tacit Collusion in Repeated Ascending Auctions*. *Journal of Law and Economics* 52: 91–109.
- Phillips, O., D. Menkhaus, and K. Coatney. 2003. *Collusive Practices in Repeated English Auctions: Experimental Evidence on Bidding Rings*. *American Economic Review* 93: 965–79.
- Porter, R. 1983. *A Study of Cartel Stability: The Joint Executive Committee 1880–1886*. *Bell Journal of Economics* 14: 301–14.
- Porter, R. 2005. *Detecting Collusion*. *Review of Industrial Organization* 26: 147–67.
- Porter, R., and J. Zona. 1993. *Detection of Bid Rigging in Procurement Auctions*. *Journal of Political Economy* 101: 518–53.
- Porter, R., and J. Zona. 1999. *Ohio School Milk Markets: An Analysis of Bidding*. *Rand Journal of Economics* 30: 263–88.
- Potters, J. 2009. *Transparency about Past, Present, and Future Conduct: Experimental Evidence on the Impact on Competitiveness*. In J. Hinloopen and H.-T. Normann, *Experiments and Competition Policy*, 81–104. Cambridge: Cambridge University Press.
- Riley, J., and W. Samuelson. 1981. *Optimal Auctions*. *American Economic Review* 71: 381–39.
- Robinson, M. 1985. *Collusion and the Choice of Auction*. *Rand Journal of Economics* 16: 141–45.
- Sherstyuk, K. 1999. *Collusion without Conspiracy: An Experimental Study of One-Sided Auctions*. *Experimental Economics* 2: 59–75.
- Sherstyuk, K. 2002. *Collusion in Private Value Ascending Price Auctions*. *Journal of Economic Behavior and Organization* 48: 177–95.
- Sherstyuk, K., and J. Dulatre. 2008. *Market Performance and Collusion in Sequential and Simultaneous Multi-object Auctions: Evidence from an Ascending Auctions Experiment*. *International Journal of Industrial Organization* 26: 557–72.
- Skrzypacz, A., and H. Hopenhayn. 2004. *Tacit Collusion in Repeated Auctions*. *Journal of Economic Theory* 114: 153–69.
- Stigler, G. 1964. *A Theory of Oligopoly*. *Journal of Political Economy* 72: 44–61.
- Vickrey, W. 1961. *Counterspeculation, Auctions, and Competitive Sealed Tenders*. *Journal of Finance* 16: 8–37.
- Waehrer, K. 1999. *Asymmetric Private Values Auctions with Application to Joint Bidding and Mergers*. *International Journal of Industrial Organization* 17: 437–52.
- Wilson, R. 1992. *Strategic Analysis of Auctions*. In R. Aumann and S. Hart, *Handbook of Game Theory*, vol. 1: 227–79. Amsterdam: Elsevier.