

Detecting Large-Scale Collusion in Procurement Auctions

Kei Kawai Jun Nakabayashi
New York University Tohoku University

January 27, 2014

Abstract

This paper documents wide-spread collusion among construction firms in Japan using data that account for most of the construction projects procured by the national government of Japan during 2003 to 2006. By examining rebids that occur for auctions that failed to meet the reserve price, we use ideas similar to regression discontinuity for identifying collusion. We identify about 690 firms whose conduct is inconsistent with competitive behavior. The number of projects awarded to these bidders is about 7,000, or about 1/6 of the total number of construction projects awarded by the national government during the period. The value of these auctions sums to about \$8 billion, about 5-6% of which may have been saved absent collusion.

1 Introduction

One of the central themes of competition policy is to deter, detect, and punish collusion. While there is almost universal agreement among economists that collusion among firms is socially undesirable, firms often have private incentives to engage in collusive behavior absent regulatory sanctions. Therefore it is crucial to ensure that the antitrust agencies have the authority and the resources to go after collusion in order to promote competition among firms. To the extent that collusive activities remain undetected or unpunished, collusion may become the norm rather than the exception with possibly large detrimental effects on the economy.

In this paper, we document wide-spread collusion among construction firms using procurement data from Japan. Our data set covers April 2003 to December 2006 accounting for most of the construction projects procured by the national government of Japan during

this period, totaling more than 40,000 auctions and more than 42 billion US dollars. On an annual basis, the award amount sums to around 14 billion U.S. dollars, or about 3% of the national tax revenue. Using this large data set, we provide evidence of wide-spread collusion. We find patterns of collusion that persist across region, across types of construction projects and across time.

While there were only four collusion cases brought against the construction firms by the antitrust authorities (JFTC) in connection to procurement projects auctioned by the national government during the sample period, there is wide spread speculation that many construction firms were engaging in collusive behavior. For example, the Japanese Bar Association issued a study in 2001 in which they concluded that collusion is wide spread in Japan with extremely high probability based, in part, on the testimony of the defendants involved in five criminal collusion cases. In another widely publicized incident in 1997, Sakae Hirashima, a former corporate executive of Obayashi corporation – who was the ring leader in the Kansai area and often referred to as the “don” or the “emperor” of the construction industry according to news reports – filed a report to the JFTC which implicated more than 150 construction firms. He claimed that he was involved in allocating among the ring members, more than 50 billion dollars worth of construction projects in 1996 alone.¹ In fact, collusion among construction firms were deemed so pervasive that it became one of the sticking points during talks over the U.S. - Japan trade friction.²

There are numerous news reports and articles that document evidence of possible collusion even aside from the examples we discussed above. However, much of it are based on isolated and scattered anecdotal evidence. As far as we are aware, there has not been any concrete evidence regarding the pervasiveness of collusion. By examining the universe of construction projects procured by the Japanese national government, this paper seeks to provide a systematic account of collusion among construction firms in procurement auctions.

In principle, bidding rings can be organized in a variety of ways depending on whether members engage in side-payments or not, whether explicit communication between the members is feasible, etc. Whatever the exact arrangement, however, a very common feature of bidding rings is that ring members pick a predetermined winner beforehand and reduce competitive pressure in the actual auction. Hence, all the ring members except for the

¹The projects he claimed to allocate include those procured by the national government as well as local governments. The latter accounts for about 10 times the value of the former.

²See Japan Structural Impediments Initiative Joint Report, in particular, the Report by the Japanese delegation, Section IV (Exclusionary Business Practices) II (Measures to be Taken) -(7) (Effective Deterrence against Bidrigging).

predetermined winner submit non-serious high bids and the sole serious bid is submitted by the predetermined winner. Of course, even the serious bid is inflated relative to the competitive bid ensuring that the ring extracts surplus from the buyer. Almost all of the existing evidence – including all of the criminal collusion cases cited in the previously cited Bar Association study, the collusive scheme described by Sakae Hirashima, as well as the four bidding rings that were prosecuted by the JFTC – indicate that bidding rings in the construction industry in Japan are organized in this manner. That is, the bidding ring picks a designated winner in advance (who is typically the only firm that incurs the cost of estimating the project cost) and the rest of the ring members help the predetermined winner win. By examining rebids for auctions that failed to meet the reserve price, our paper exploits this feature of bidding rings for detecting collusion.

The auction mechanism for the Japanese public construction projects is a variant of the first-price sealed-bid (FPSB) mechanism with a secret reserve price. In fact, the auction mechanism is exactly the same as FPSB as long as the lowest bid is below the secret reserve price, in which case the lowest bidder becomes the winner and the auction ends. If none of the bids is below the reserve price, however, the buyer solicits a new round of bids from the same set of bidders immediately thereafter. The lowest bid (but not the identity of the bidder) from the initial round is revealed to all the bidders before the bidders submit second-round bids (supposedly to encourage second-round bids that meet the secret reserve price). No other bids are revealed. If the lowest bid in the second round is still higher than the reserve price (which is kept unchanged), there is a third (final) round of bidding. Approximately 20% of all auctions advance to the second round in our data.

In order to identify collusion, we use an idea that is similar to regression discontinuity design. In particular, we look for patterns in the data where the identity of the lowest bidder is very persistent – consistent with designating a predetermined winner among the ring members – beyond what can be explained by competitive behavior. To be more concrete, let $R1$ and $R2$ be the lowest and the second lowest bidders in the first round, respectively. We then examine the second-round bids of $R1$ and $R2$ for auctions where (1) the first-round bids of $R1$ and $R2$ are only ε apart; and (2) the auction proceeds to the second round. Note that conditional on the first-round being very close to each other, who turns out to be the lowest/second-lowest in the first-round is going to be as good as random under competition. Hence the two bidders can be thought of as symmetric, in terms of costs, risk aversion, beliefs over the distribution of the reserve price etc. Hence, absent informational advantage accrued to $R2$ given the way the first round bid is revealed to the participants, the likelihood of winning in the second round should be equally likely for $R1$ and $R2$ as long as

ε is small enough. Note as well that the informational advantage accrued to $R2$ would make it even less likely that $R1$ would be the lowest bidder in the second round under competitive behavior.³

Given this situation, we would expect $R2$ to win at least as often as $R1$ in the second round under competitive bidding.⁴ However, we find that $R2$ rarely outbids $R1$ in the second round in the actual data. For example, when we set ε to be 1% of the reserve price, $R2$ outbids $R1$ in only about 2.2% of the cases. The probability that $R1$ remains the lowest bidder in the second round is around 97.5%.⁵

Of course, it is possible that our results are still driven by inherent cost differences among the firms, i.e., the bandwidth we use (e.g., $\varepsilon = 1\%$ of the reserve price) is not small enough to adequately control for differences in costs, etc., among the bidders. In order to rule out this possibility, we apply the same test to the second round bids of $R2$ and $R3$ (the second and the third lowest bidders in the first round). More precisely, among the set of auctions where (1) the bid difference between $R2$ and $R3$ in the first round is less than 1% of the reserve price; and (2) the auction proceeds to the second round, we find that $R3$ outbids $R2$ in about 49.44% of the case. This gives assurance that that bandwidth we choose for ε is sufficiently small for purging any inherent differences among the bidders. Our results thus suggest that there is much more persistence – across multiple rounds within the same auction – in the identity of the lowest bidder than can be explained by competition.

In order to provide more evidence on collusion, we further examine the shape of the distribution of the difference between the normalized second round bids of $R1$ and $R2$ ($\equiv \Delta_{12}$) as well as the difference between the normalized second round bids of $R2$ and $R3$ ($\equiv \Delta_{23}$), for different values of ε .⁶ First, we find that the distribution of the latter (Δ_{23}) is symmetric around zero. The symmetry of the distribution around zero reconfirms that we can control well for underlying costs etc., and treat $R2$ and $R3$ as being symmetric by

³The buyer reveals the lowest bid, but none of the other bids. Hence $R1$ only knows that his bid was the lowest among the first-round bids while $R2$ gains knowledge of two bids, his own bid as well as the lowest bid. This means that conditional on the two lowest bids being very close to each other, $R2$ has an informational advantage in the second round: While $R1$ does not know that there is another bidder who bid just above his bid, $R2$ knows exactly how the lowest bidder bid.

⁴More precisely, $R2$ should do no worse than $R1$. This does not logically imply that $R2$ should win more often if profit margins are very thin. We address this concern by examining auctions that (1) go to the third round and (2) $R2$ submits substantially lower bids in the third round. For these set of auctions, we have an upper bound on $R2$'s costs, or equivalently, a lower bound on $R2$'s profit margin. Even for these auctions, we find that $R2$ almost never outbids $R1$ in the second round.

⁵The reason why this probability is lower than 97.8% ($= 100\% - 2.2\%$) is because occasionally third lowest bidder in the first round (or the fourth lowest, etc.) becomes the lowest bidder in the second round.

⁶We normalize the bids by dividing the bids by the reserve price. The precise definition of Δ_{12} and Δ_{23} are as follows: $\Delta_{12} = \frac{[\text{2nd round bid of } R2] - [\text{2nd round bid of } R1]}{\text{Reserve Price}}$ and $\Delta_{23} = \frac{[\text{2nd round bid of } R3] - [\text{2nd round bid of } R2]}{\text{Reserve Price}}$.

taking small values of ε . In contrast, however, we find that the histogram of Δ_{12} is highly skewed to the right of zero. The fact that the distribution is skewed to the right of zero (in fact, about 98% of the mass lie to the right of zero) confirms our previous finding that $R1$ almost never gets outbid by $R2$ in the second round.

The second finding – perhaps this is a more conclusive evidence of collusion – is that we find what appears to be a discontinuity at exactly zero for the histogram of Δ_{12} . That is, when we focus on a small band around zero, we find that there are hundreds of auctions that fall just to the right of zero ($\Delta_{12} \in [0, t]$ for some small positive t) whereas there are almost no auctions that fall to the left of zero ($\Delta_{12} \in [-t, 0]$). For example, when we take $t = 1\%$ and consider the ratio of auctions with $\Delta_{12} \in [0, -t]$ and $\Delta_{12} \in [0, t]$, the ratio is 1 : 50. This implies that there are many auctions where $R2$ loses to $R1$ in the second round by a tiny margin but almost no auctions where $R2$ wins by a tiny margin. The histogram for Δ_{23} , on the other hand, is continuous and symmetric around zero with a fair amount of variance.

The discontinuity exhibited in the histogram of Δ_{12} at zero strongly suggests that the bidders know how each other will bid in the second round, and moreover, that auction participants designate a predetermined winner in advance. To see this, suppose to the contrary. If $R1$ and $R2$ were uncertain as to how each other will bid, then one should observe a similar number of auctions where $R2$ wins by a tiny margin as auctions where $R2$ loses by a tiny margin. Hence the fact that Δ_{12} seems discontinuous at zero suggests that bidders are aware of how each other will bid. But if this is the case, why else would $R2$ *lose* by a small margin (rather than *win* by a small margin) other than to yield to the predetermined winner? We take the fact that the histogram of Δ_{12} is discontinuous at exactly zero as evidence of bidder collusion.

Lastly, we develop a formal test statistic for collusion based on the idea that discontinuity of Δ_{12} is evidence of collusive behavior and apply it to each firm in our data set. Our test is based on the variance of Δ_{23} as well as the frequency with which $\Delta_{12} \in [-t, 0]$ relative to $\Delta_{12} \in [0, t]$ among the set of auctions that the firm participates in. The variance of Δ_{23} puts a bound on how sharp the distribution of Δ_{12} can change around zero under the null of competitive bidding.

In our baseline result, we find 691 firms for whom we reject the null hypothesis of competitive behavior at 95%. The number of auctions won by these firms totals 6,900, or about 1/6 of the total number of auctions. The total award amount of these auctions sums to about \$7.9 billion. We estimate that absent collusion, taxpayers could have saved about \$425 million. Moreover, if we also consider the fact that the same firms that we identify

are also active in municipal and prefectural construction projects – the total value of these projects is close to 10 times the size for our data set – the overall impact of collusion can be staggering.

Overall, the bidding pattern that we identify as suggestive of collusion is prevalent across region, time, and type of project. We find that the shape of the distribution of Δ_{12} and Δ_{23} look remarkably consistent regardless of how we condition on observables, suggesting that collusion is wide-spread. Moreover, there is strong evidence that most bidders have participated in auctions that are highly suspicious.

1.1 Related Literature

This paper is most closely related to the empirical literature on the detection of collusion.⁷ Existing empirical studies tend to be related to court-documented cases of cartel activities, e.g., paving in highway construction in Nassau and Suffolk counties (Porter and Zona 1993), school milk in Ohio (Porter and Zona 1999), school milk in Florida and Texas (Pensendorfer 2000), and collectible stamps in North America (Asker 2010).⁸ While the analysis of this paper does not rely on data generated by known cartels, it is still useful to apply our analysis to the bids of known bidding rings for validation. We do this in Section 5 for the four known bidding cartels.

There is also a literature that tests for collusion in the absence of any prior knowledge of bidder conduct. Examples include bidding in seal coat contracts in three States in the U.S. Midwest (Bajari and Ye 1999), U.S. forest service timber sales (Baldwin, Marshall Richard 1997, Athey, Levin and Seira 2011), U.S. offshore oil and gas leases (Haile, Hendricks, Porter and Onuma 2013), roadwork contracts in Italy (Decaloris, 2013) and public-works consulting in Japan (Ishii, 2009). Ishii (2009) studies 175 auctions for consulting work in Naha, Okinawa and analyzes how the winner of the auctions can be explained by exchange of favors. While her identification is based on bid patterns across auctions, our identification strategy focuses on how bidders bid within a given auction. Our study also looks at the universe of construction projects procured by the national government whereas she studies a specific local market.⁹

⁷For a brief survey, see the entry Bidding Rings, by Asker in the Palgrave Dictionary of Economics.

⁸Bajari and Ye (2003) proposed a way to distinguish between competitive and collusive bidding with data on procurement auctions for seal coat contracts in the Midwest, where possible auction riggings were not documented by the antitrust authority. They structurally estimated the bidder's cost distribution from observed bid data and compare it with the recovered cost distribution based on the beliefs of the industry experts on the bidder's cost schedule.

⁹For a more general overview of bidding rings among procurement firms in Japan, see Mcmillan (1991).

2 Institutional Background

Auction Mechanism The auction mechanism used in our sample is a variant of the first-price sealed bid (FPSB) auction with a secret reserve price which is based on engineers' estimates. In fact, the auction mechanism is exactly the same as FPSB auction as long as the lowest bid is below the secret reserve price, in which case the lowest bidder becomes the winner and the auction ends. If no bid meets the reserve price, the buyer solicits rebids from the same set of bidders immediately thereafter.¹⁰ The lowest bid (but not the identity of the bidder) from the initial round is revealed to all the bidders before the bidders submit second-round bids (supposedly to encourage second-round bids that meet the secret reserve price). No other bids are revealed. This means that when bidding in the second round, the bidders know that the secret reserve price is lower than the lowest first round bid.

The second round proceeds in the same manner as the initial round; if the lowest bid is below the reserve price, the auction ends and the lowest bidder wins. Otherwise, the auction goes to the third round. As before, the lowest bid in the second round is revealed to all the bidders before the third round. The third round is the final round and there are no further rounds. If no bid meets the reserve price in the third round, bilateral negotiation takes place between the buyer and the the lowest third round bidder. The same secret reserve price is used in all three rounds.

Bidder Participation As is the case in many countries, participation in procurement auctions in Japan is not fully open. A contractor who wishes to participate must first go through screening to be pre-qualified. Because pre-qualification occurs at the regional level, a contractor needs to be pre-qualified for each region in which it wishes to bid on projects.

In addition to pre-qualification, there may be additional restrictions on participation: Depending on how restrictive they are, the auctions can be divided into one of four categories. The first and the second categories are the most restrictive, where only the invited bidders participate. In these two categories, the government invites typically 10 bidders from the pool of pre-qualified contractors. The difference between the two categories is that in the first category, the invited bidders are chosen randomly from the pool, while in the second category, the government takes into account the preferences of the contractors over the types of projects they wish to bid on.

The third and fourth categories are less restrictive. The set of potential bidders are

See also Ohashi (2009) who discusses how the change in auction design in Mie Prefecture may have affected collusion.

¹⁰Article 82 of Cabinet Order concerning the Budget,Auditing and Accounting.

still restricted to the pool of pre-qualified contractors, but participation is open to all pre-qualified contractors. The difference between the third and the fourth categories is that in the third category, the government reserves the right to exclude potential bidders under certain conditions.

Collusive Behavior In principle, bidding rings can be organized in a variety of ways depending on whether members engage in side-payments or not, whether explicit communication between the members is feasible, etc. Whatever the exact arrangement, however, a very common feature of bidding rings is that ring members pick a predetermined winner beforehand and reduce competitive pressure in the actual auction. Almost all of the existing evidence indicate that bidding rings in the construction industry in Japan are organized in this manner. That is, the bidding ring picks a designated winner in advance and the rest of the ring members help the predetermined winner win.

There is also some anecdotal evidence that describes the mechanism used by bidding rings to ensure that the designated bidder win the auction. For instance, according to the records of a criminal case in Nagoya,¹¹ the designated winner of a bidding ring would communicate to other members how it would bid in each of the three successive round so that the other members can bid higher. In addition, there is also anecdotal evidence that designating a predetermined winner also has the advantage of avoiding the cost of estimating the project cost for non-designated bidders. According to the same criminal case in Nagoya, only the designated winner would incur the substantial cost of estimating the cost of the project.¹²

[reserve price is sometimes predictable, but there are many cases in which it is unpredictable.]

3 Data

We use data on auctions for construction projects obtained from the Ministry of Land, Infrastructure and Transportation, who is the largest procurement buyer in Japan. The data set spans April 2003 to December 2006 and covers most of the construction works auctioned by the national government in Japan during this period. From this sample, we dropped scoring

¹¹see nichibenren-pdf p20 (or 22 of pdf page).

¹²Estimating the project cost involves understanding the specifications of the project, assessing the quantity and quality of materials required, negotiating prices for construction material and arranging for available subcontractors. These costs are often quite substantial.

auctions, unit-price auctions, and those with mistakenly recorded data for our analysis.¹³ We are left with approximately 42,000 auctions with a total award amount of nearly \$41 billion. On an annual basis, the award amount is around \$14 billion, accounting for close to 3 percent of the national government tax revenue.

The data includes information on all bids, bidder identity, the (secret) reserve price, auction date, auction category, location of the construction site as well as other characteristics of the project.¹⁴ In particular, we have data on whether the auction proceeded to the second round or the third round, as well as the bids in each round. Table 1 provides summary statistics of the data. In the table, we report the reserve price of the auction, the winning bid, the ratio of the winning bid to the reserve price, the fraction of auctions that advance to the second and the third round, and the number of bidders, by auction category (category 1 being the most restrictive and category 4 being the most open. See Section 2)

In the first column of the table, we find that the average reserve price of the auctions is about 96 million yen, although there is considerable heterogeneity among the categories. The most restrictive category has an average reserve price of 50 million yen, while the most open category has an average of about 545 million yen. In the third column, we report the ratio of the winning bid to the reserve price. We find that the winning bid ranges between 92% to 95% of the reserve price. In the next two columns, we report the fraction of auctions that advance to the 2nd and the 3rd round. We find that about 18% of the auctions go to the second round and a little less than 3% advance to the third round. In Table 1, we report the reserve price, the winning bid, and the ratio of the winning bid to the reserve price for the set of auctions that advance to the second round and the third round.

4 Analysis

In order to provide evidence of collusive behavior, we first discuss the persistence of the identity of the lowest bidder across rounds in a given auction. We show that while there is almost 50% probability that the third lowest bidder in the first round outbids the second lowest bidder in the second round, the probability that the lowest bidder is outbid by the second lowest bidder is about 2%. This result persists even after controlling for possible differences in bidder costs.

We next focus on what appears to be a discontinuity in the distribution of the difference

¹³The removed auction data which are missed or mistakenly recorded accounts for 5.6% or 1.3% of the entire dataset, respectively. The scoring auction data accounts for 15.3%.

¹⁴Ministry of Land Infrastructure, Transportation, procures 21 types of construction works including civil engineering (or heavy and general construction works), buildings, bridges, paving, dredging, and painting.

between the second round bids of the two lowest bidders from the first round. We document evidence that there are a large number of instances where the second lowest bidder of the first round loses *in the second round* to the lowest bidder of the first round by a very narrow margin. However, we find almost no cases in which the second lowest bidder of the first round narrowly outbids the first round lowest bidder in the second round. We argue that the pattern seen in the data is not consistent with competitive behavior, and moreover, it is suggestive of communication between the members.

Lastly, we discuss whether the second round bidding strategy of the second lowest bidder in the first round can be explained by competitive behavior. We find that the size of the profit margin implied by the bidding strategy is inconsistent with the margin implied by their third round bidding behavior.

4.1 Persistence of the Identity of the Lowest Bidder

Persistence in the Second Round We begin our analysis by studying the extent to which the lowest bidder in the first round is also the lowest bidder in later rounds. Recall that a key feature of bidding rings is that there is a designated winner. In order to let the designated winner win the auction, other ring members must submit bids in such a way as to ensure that the designated bidder is the lowest bidder. Because the reserve price is a random variable from the perspective of the bidding ring, the ring members must make sure that the designated bidder is the lowest bidder in each successive round in case the auction takes multiple rounds (which happens in about 20% of the auctions). This is especially important if the designated bidder is the only one who has incurred the cost of estimating the project cost. This implies that we should observe persistence in the identity of the lowest bidder across rounds. Then our first piece of evidence of collusion is patterns of persistence in the identity of the lowest bidder across rounds for a given auction.

In Table 2, we report how the rank of the bidders changes from the first round to the second round for all auctions that proceeded to the second round with five bidders or more. The (i, j) element of the matrix corresponds to the probability that a bidder submits the j -th lowest bid in the second round, conditional on submitting the i -th lowest bid in the first round, i.e., $\Pr(j\text{-th lowest} | i\text{-th lowest})$. Thus, the diagonal elements correspond to the probability that a given bidder remains in the same rank in both rounds and the horizontal sum of the probabilities is one.

What is striking about this table is the probability in the $(1, 1)$ cell. We find that approximately 97% of the case, the lowest bidder in the first round is still the lowest bidder in

Terminal Round	(R)eserve Yen M.	(W)inbid Yen M.	W/R	Lowest bid / Reserve			# Bidders	N
				Round 1	Round 2	Round 3		
1	103.449 (247.21)	96.581 (234.39)	0.922 (0.086)	0.922 (0.086)	-	-	9.86 (2.60)	33,886
2	80.032 (177.46)	77.433 (172.81)	0.962 (0.034)	1.057 (0.076)	0.962 (0.034)	-	9.94 (2.45)	6,606
3	60.269 (171.96)	56.936 (162.04)	0.942 (0.040)	1.128 (0.111)	1.054 (0.091)	0.942 (0.040)	9.41 (2.28)	1,156
All	98.537 (235.91)	92.443 (224.17)	0.929 (0.081)	0.949 (0.103)	0.976 (0.058)	0.942 (0.040)	9.86 (2.57)	41,648

Note: The numbers in parentheses are the standard deviations.

Table 1: Sample Statistics

	Round 2					
	1	2	3	4	5+	
Round 1	1	96.75%	1.55%	0.52%	0.28%	0.90%
	2	1.46%	30.72%	18.91%	13.04%	35.88%
	3	0.48%	22.69%	18.46%	14.25%	44.12%
	4	0.31%	17.26%	16.64%	16.12%	49.67%
	5+	0.18%	8.55%	10.18%	11.03%	70.06%

Note: The (i,j) element of the matrix denotes the probability that a bidder submits the j -th lowest bid in the second round conditional on submitting the i -th lowest bid in the first round. When there are ties, multiple bidders are assigned to the same rank.

Table 2: Rank of the Second-round Bid by Rank of the First-round Bid

the second round. The flip side of this is that if a bidder is not the lowest bidder in the first round, the bidder is almost never the lowest bidder in the second round. For example, the conditional probability that a second lowest bidder in round one becomes the lowest bidder in round two is only about 1.5%. Note also that the diagonal elements other than the (1, 1) element are much smaller: the probability that the second lowest bidder in the first round remains the second lowest bidder is just around 36%. There is very strong persistence in the identity of the lowest bidder, but not necessarily for other positions.

In order to illustrate this point further, we next examine more closely how the three lowest bidders in the first round behave in the second round. In what follows, we let $R_1(k)$, $R_2(k)$, and $R_3(k)$ denote the identity of the bidder who submits the k -th lowest bid in the 1st, 2nd, and 3rd rounds. We also denote the (normalized) bid of bidder i in round t by b_i^t . Because there is considerable variation in the project size, we work with the normalized bids by dividing the actual bids by the reserve price of the auction. Hence, $b_{R_1(1)}^2$, for example, denotes the second round bid of the first round lowest bidder as a percentage of the reserve price.

In the top left panel of Figure 1, we plot the histogram of $\Delta_{12}^2 \equiv b_{R_1(2)}^2 - b_{R_1(1)}^2$ for the set of auctions that go to the second round. That is, we plot the (normalized) difference in the second-round bid of the first round lowest bidder ($R_1(1)$) and the second round bid of the first round second lowest bidder ($R_1(2)$). Note that almost all of the mass lies to the right of zero, which confirms what we report in Table 2, i.e., a flip in the ordering between the lowest and the second lowest bidders almost never happens across rounds. In the top right panel of Figure 1, we plot the histogram of $\Delta_{23}^2 \equiv b_{R_1(3)}^2 - b_{R_1(2)}^2$, which is the normalized difference in the rebids of the second lowest bidder in the first round ($R_1(2)$) and the third lowest bidder in the first round ($R_1(3)$). In stark contrast to the left panel, the shape of the histogram for Δ_{23}^2 is quite symmetric around zero. This implies that the ranking between $R_1(2)$ and $R_1(3)$ flips with almost 50% probability. This also seems consistent with our previous finding that there is much less persistence in the ranking for the second and third ranks.

So far, the results that we have presented correspond to all of the auctions that proceeded to the second round. However, it is possible that our results are driven by inherent cost differences among firms. As long as there are significant cost differences between the lowest bidder and all of the other bidders, our results can be generated by competitive bidding. In order to rule out this possibility, we perform the same analysis by conditioning on the set of auctions where the first round bids are very close to each other. The idea is that if, for example, the bids of $R_1(1)$ and $R_1(2)$ are sufficiently close (i.e., $b_{R_1(2)}^1 - b_{R_1(1)}^1 < \varepsilon$

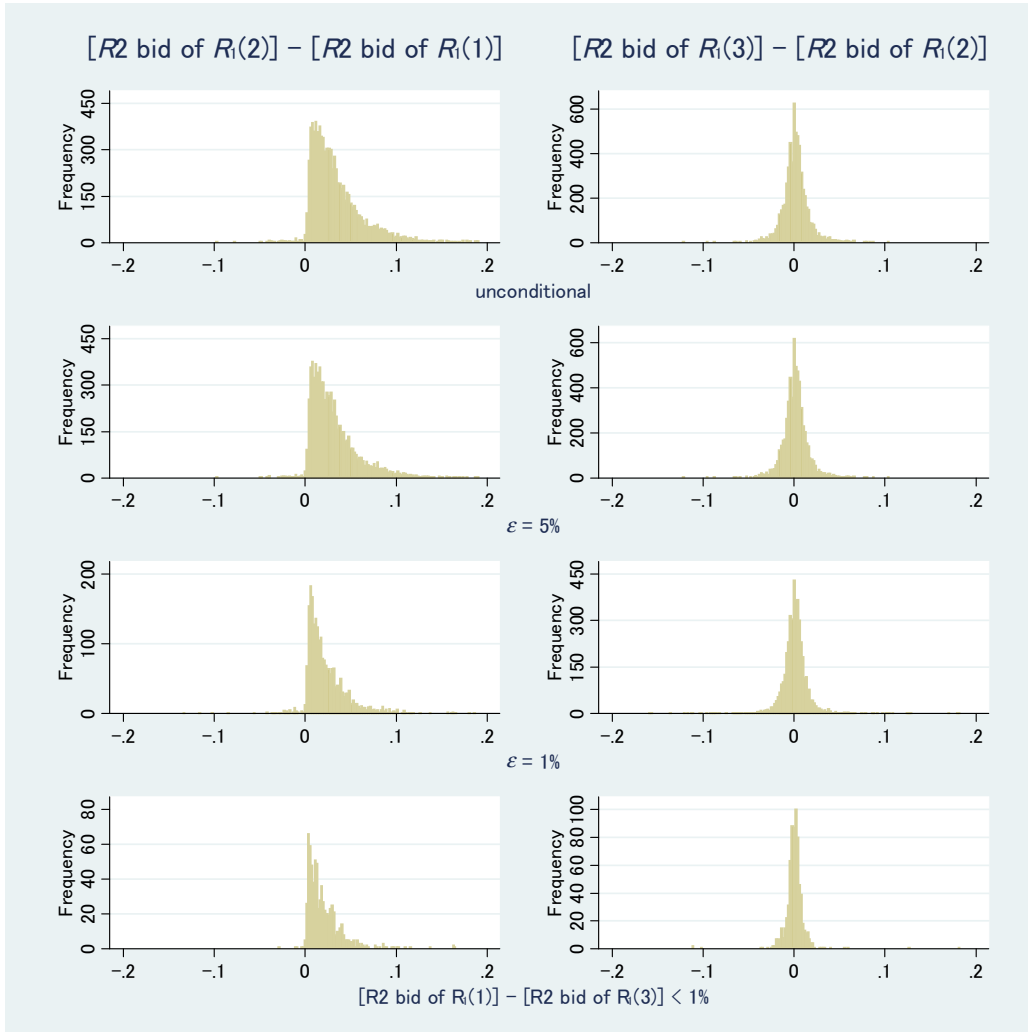


Figure 1: Difference in the Second Round Bids of $R_1(2)$ and $R_1(1)$ (Left Panels) and the Difference in the Second Round Bids of $R_1(3)$ and $R_1(2)$ (Right Panels). The first row is the histogram for the set of auctions that reach the second stage; and $R_1(1)$ and $R_1(2)$ (or $R_1(2)$ and $R_1(3)$) submit valid bids in the second round. The second to fourth rows plot the same histogram, but only for auctions where the differences in the first round bids are relatively small.

for some small ε), there should be little inherent cost differences among them. In fact, if ε is small enough, which bidder turns out to be the lowest/second lowest bidder in the first round is as good as random. Hence $R_1(1)$ and $R_1(2)$ should be symmetric ex-ante, including costs, but also on all other dimensions as well, such as risk attitude, beliefs over the reserve price, etc.

In the second row of Figure 1, we plot Δ_{12}^2 and Δ_{23}^2 for the subset of auctions for which the bids in the first round are within 5% of each other. In particular, we plot the histogram of Δ_{12}^2 for the set of auctions where $b_{R_1(2)}^1 - b_{R_1(1)}^1 < 0.05$ in the left panel and the histogram of Δ_{23}^2 for the set of auctions with $b_{R_1(3)}^1 - b_{R_1(2)}^1 < 0.05$ in the right panel. The shape of the distribution of Δ_{12}^2 in the left panel is still very skewed and asymmetric around zero, while the distribution of Δ_{23}^2 in the right panel remains symmetric around zero. The fact that the distribution of Δ_{23}^2 is symmetric around zero and very similar to the top panel suggests that cost differences between bidders do not seem to play a large role: If cost differences were driving the skewed bid pattern for Δ_{12}^2 in the left panel, we should also expect to see a distribution of Δ_{23}^2 that is skewed to the right of zero. The third row plots the distribution of Δ_{12}^2 and Δ_{23}^2 , but now conditioning on auctions with $b_{R_1(2)}^1 - b_{R_1(1)}^1 < 0.01$ and $b_{R_1(3)}^1 - b_{R_1(2)}^1 < 0.01$, respectively. Lastly, the bottom row shows the distribution of Δ_{12}^2 and Δ_{23}^2 conditional on the event that the three lowest bids in the first round are *all* within 1% of each other, $b_{R_1(3)}^1 - b_{R_1(1)}^1 < 0.01$.¹⁵ Taken together, the results of Figure 1 suggests that it is not differences in costs that is driving the persistence in the identity of the lowest bidder.

In the Online Appendix, we further explore whether the distribution of Δ_{12}^2 and Δ_{23}^2 exhibits similar patterns when we condition the sample by various auction characteristics, such as region, auction category, project type, and year. We find that the distribution of Δ_{12}^2 and Δ_{23}^2 often look very similar to those shown in Figure 1 – the distribution of Δ_{12}^2 is skewed to the right and displays what appears to be a discontinuity at zero while the distribution of Δ_{23}^2 is symmetric around zero.

Persistence in the Third Round For the subset of auctions that go to the third round, we can further examine whether a similar pattern continues to hold in the third round. In the top left panel of Figure 2, we plot the difference in the third round bids of $R_1(1)$ and $R_1(2)$, i.e., $\Delta_{12}^3 \equiv b_{R_1(2)}^3 - b_{R_1(1)}^3$. In the right panel we plot the differences in the third round bids of $R_1(3)$ and $R_1(2)$, i.e., $\Delta_{23}^3 \equiv b_{R_1(3)}^3 - b_{R_1(2)}^3$. In rows two to four of Figure

¹⁵Note that $b_{R_1(3)}^1 > b_{R_1(2)}^1 > b_{R_1(1)}^1$, by construction. Hence, $b_{R_1(3)}^1 - b_{R_1(1)}^1 < 0.01$ implies $b_{R_1(2)}^1 - b_{R_1(1)}^1 < 0.01$ and $b_{R_1(3)}^1 - b_{R_1(2)}^1 < 0.01$.

2, we plot the histogram conditioning on the set of auctions in which the first round bids were sufficiently close. The left panel of the second row plots Δ_{12}^3 for the set of auction where $b_{R_1(2)}^1 - b_{R_1(1)}^1 < 0.05$, the third row plots Δ_{12}^3 for which $b_{R_1(2)}^1 - b_{R_1(1)}^1 < 0.01$ and the last row plots Δ_{12}^3 for which $b_{R_1(3)}^1 - b_{R_1(1)}^1 < 0.03$. Similarly, the right panel plots Δ_{23}^3 for the set of auctions where $b_{R_1(3)}^1 - b_{R_1(2)}^1 < 0.05$, $b_{R_1(3)}^1 - b_{R_1(2)}^1 < 0.01$ and $b_{R_1(3)}^1 - b_{R_1(1)}^1 < 0.03$, respectively.

Informational Advantage of $R_1(2)$ Recall from Section 2 that the lowest bid is announced in each round, but none of the other bids are. This means that while $R_1(1)$ only gains knowledge that it was the lowest bidder in the first round, $R_1(2)$ learns exactly how the lowest bidder bid in the first round in addition to how it bid itself. This implies that conditional on the two lowest bids being very close to each other, $R_1(2)$ has an informational advantage over $R_1(1)$ in the second round. To see this, consider the case when $R_1(1)$ and $R_1(2)$ bid almost exactly the same amount, say $\$Z$. The information revealed to $R_1(1)$ at the end of the first round is that $\$Z$ is the lowest bid. The information revealed to $R_1(2)$, on the other hand, is that $\$Z$ is the lowest bid; and (at least) one other firm beside itself bid $\$Z$. Clearly, $R_1(2)$ has a bigger information set at the end of the first round.

So far, we have documented that the ordering between $R_1(1)$ and $R_1(2)$ is very persistent across rounds while the ordering between $R_1(2)$ and $R_1(3)$ is not. Given the informational advantage accrued to $R_1(2)$, however, the fact that the ordering between $R_1(1)$ and $R_1(2)$ does not change from round to round is even more surprising. Once we condition on auctions where $R_1(1)$ and $R_1(2)$ bid close to each other in round one, $R_1(2)$ should be aware that by bidding a little bit more aggressively, it can beat $R_1(1)$ in the next round. Hence, when we factor in the informational advantage of $R_1(2)$, we would normally expect the order of $R_1(1)$ and $R_1(2)$ to flip more frequently than 50%, and not less, under competitive behavior.

4.2 Discontinuity of Δ_{12}^2 at Zero

A striking feature of the distribution of Δ_{12}^2 is that there is what appears to be a discontinuous jump at exactly zero. This is in stark contrast to the distribution of Δ_{23}^2 , which is symmetric and continuous around zero. We argue that this pattern of bidding is inconsistent with competitive behavior.

Consider first the distribution of Δ_{23}^2 in the right panels of Figure X. Note that there is a certain amount of variance in Δ_{23}^2 in all of the four panels. This seems to indicate that for many auctions, there is a reasonable amount of idiosyncracies among the bidders with

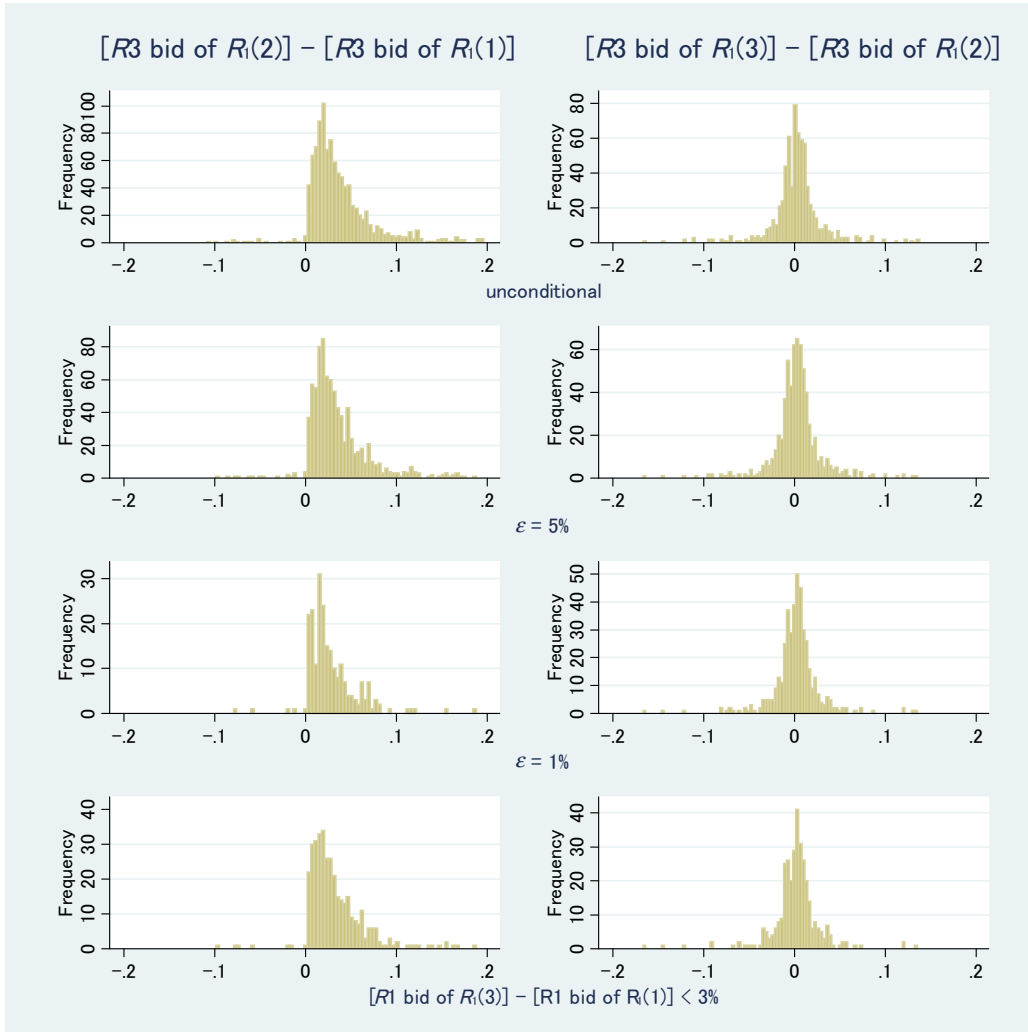


Figure 2: Difference in the Third Round Bids of $R_1(2)$ and $R_1(1)$ (Left Panels) and the Difference in the Third Round Bids of $R_1(3)$ and $R_1(2)$ (Right Panels). The first row corresponds to all auctions that reached the third round and $R_1(1)$ and $R_1(2)$ (in the case of the left panel) or $R_1(2)$ and $R_1(3)$ (in the case of the right panel) submitted valid bids in the third round. The second to fourth rows plot the same histogram, but only for auctions where the differences in the first round bids are relatively small.

regard to the beliefs over the distribution of the reserve rate, risk preference, etc., inducing variance in the second round bids. In other words, idiosyncratic reasons among the bidders seem to induce at least a certain amount of uncertainty in the second round bidding for many auctions.

Now consider the distribution of Δ_{12}^2 in the left panels of Figure X. To the extent that there exists a reasonable amount of idiosyncracies among the bidders, $R_1(2)$ should outbid $R_1(1)$ in the second round by a narrow margin just as often as $R_1(1)$ outbids $R_1(2)$ by a narrow margin. That is, there should be a similar number of observations in which $\Delta_{12}^2 \in [-t, 0]$ and $\Delta_{12}^2 \in [0, t]$ for small values of t – a feature which we clearly do not see in any of the histograms of the left panels of Figure 1. This is inconsistent with competitive behavior.

Moreover, the discreteness exhibited in the histogram of Δ_{12}^2 at zero suggests that the bidders know exactly how the other bidders will bid in the second round. If, to the contrary, $R_1(2)$ and $R_1(1)$ were both uncertain about each other’s bid, there should be just as many cases where $R_1(2)$ won by a tiny margin as cases where $R_1(2)$ lost by a tiny margin. Hence, the discontinuity of Δ_{12}^2 suggests that the bidders have prior knowledge about how each other will bid and that $R_1(2)$ is deliberately losing by submitting a slightly higher bid than $R_1(1)$ (rather than winning by slightly underbidding $R_1(1)$). This is evidence that $R_1(1)$ is the predetermined winner.

Regarding whether ring members can achieve collusion without communication, we feel that it seems unlikely. There is large heterogeneity in project size, specification, etc. between auctions. This makes it hard for bidders to predict a particular price which could serve as an obvious anchor of tacit (i.e., no communication) collusion, in general. Therefore, the observed bid pattern seems to indicate communication.

4.3 Optimality of $R_1(2)$ in the Second Round

We now explore the persistence of the lowest bidder and the discontinuity of Δ_{12}^2 from the perspective of the optimality of $R_1(2)$ ’s second round bid. Recall from the previous section that there are many cases in which $R_1(2)$ loses to $R_1(1)$ in the second round by a very small margin. In particular, focusing on the second row of the left panel of Figure 1, we find that about 4.48%, 16.37%, and 39.29% of the distribution lies within $[0, 0.005]$, $[0, 0.01]$, and $[0, 0.02]$, respectively. The probability that the distribution lies to the left of zero is only 1.73%. This suggests that $R_1(2)$ can increase the probability of outbidding $R_1(1)$ substantially by shading its bid only slightly.

Of course, outbidding $R_1(1)$ is not the same as winning the auction because one must outbid all the other bidders as well as the secret reserve price in order to win the auction. In order to take this into consideration, we shade the second-round bid of $R_1(2)$ in every auction by 0.5%, 1%, and 2% and count the number of instances in which the shaded bid is lower than the realized secret reserve price and all the other bids. We find that $R_1(2)$ would win the auction in 3.51%, 11.18% and 30.07% of the time, respectively. In contrast, the actual fraction of auctions in which $R_1(2)$ won (either in the second round or the third round) was a mere 1.23%. This means that $R_1(2)$ could have increased the probability of winning the auction by about three-fold from around 1.23% to 3.51% by lowering its bid by merely 0.5%. Unless profit margin of $R_1(2)$ is very thin – in fact, less than 0.52% of its second round bid – $R_1(2)$ can increase its profits by lowering its second-round bid by 0.5%, i.e., the observed bid of $R_1(2)$ is not optimal.¹⁶

While a profit margin of only around 0.52% of the second-round bid seems quite small to be reasonable, it is difficult to obtain direct cost measures that allows us to test this. What we do instead, is to consider a subset of auctions that (1) go to the third round; and (2) $R_1(2)$ bids substantially less in the third round than in the second round. For these auctions, the third round bid of $R_1(2)$ gives us a lower bound on $R_1(2)$'s profit margin: if $b_{R_1(2)}^3 < b_{R_1(2)}^2 \times x\%$, then we know that $R_1(2)$ was willing to win the auction at a price that is at least $x\%$ lower than its second round bid.

In the top-left panel of Figure 3, we plot $\Delta_{12}^2 \equiv b_{R_1(2)}^2 - b_{R_1(1)}^2$ for the set of auctions that proceeded to the third round; $b_{R_1(2)}^3$ was at least 10% lower than $b_{R_1(2)}^2$; and $\varepsilon = 0.05$. The first two conditions ensure that we are only examining the set of auctions where the profit margin of $R_1(2)$ was sufficiently high in the second round; and the last condition ($\varepsilon = 0.05$) ensures that the differences between $R_1(1)$ and $R_1(2)$ are relatively modest.

Note that the shape of Δ_{12}^2 remains more or less the same compared to the distribution of Δ_{12}^2 plotted in the left panels of Figure 1, i.e., there is a substantial mass just to the right of zero, but almost none to the left of zero. This suggests that low profit margins cannot explain the reluctance of $R_1(2)$ to outbid $R_1(1)$ in the second round. The top right panel of Figure 3 plots the distribution of Δ_{23}^2 for the same set of auctions for comparison. By and large, the distribution of Δ_{23}^2 is symmetric around zero, as before. The two panels in the second row plot the histogram when we further condition the sample to the set of auctions

¹⁶ $R_1(2)$ is better off lowering its second round bid by 0.5% unless its cost is higher than 99.48% of its second-round bid,

$$(b_{R_1(2)}^2 - c) * 1.23\% \leq (0.995b_{R_1(2)}^2 - c) * 30.07\%$$

Solving for c gives about $c \leq 99.48\%b_{R_1(2)}^2$.

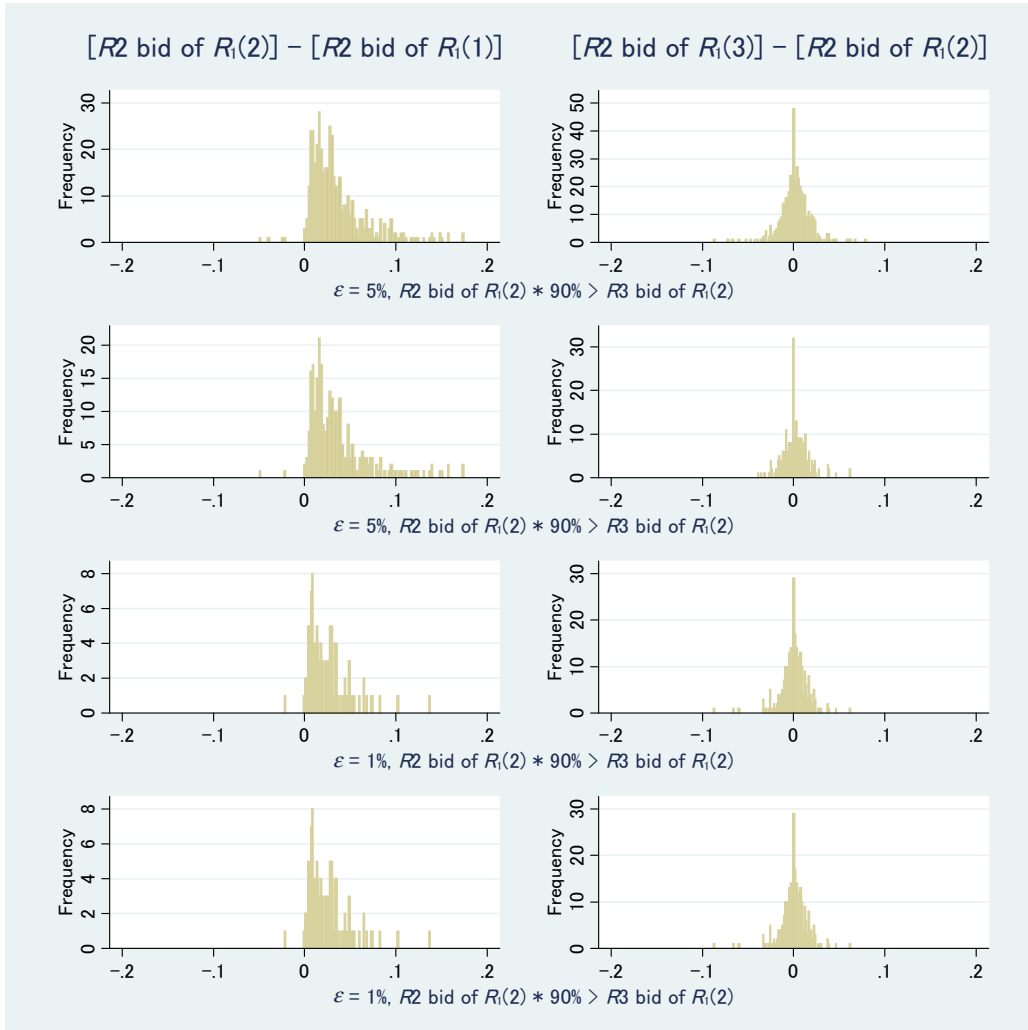


Figure 3: Difference in the Second Round Bids of $R_1(2)$ and $R_1(1)$ (Left Panels) and the Difference in the Second Round Bids of $R_1(3)$ and $R_1(2)$ (Right Panels) for Auctions with Large Profit Margin. The figure plots the histogram for the set of auctions that eventually reach the third round; and $R_1(2)$'s third round bid is less than 90% (first two rows) or 85% (last two rows) of its second round bid. See the main text for conditions on ϵ .

in which $R_1(2)$ bids at least 15% less in the third round ($b_{R_1(2)}^3 < 0.85 \cdot b_{R_1(2)}^2$). Again, we see a similar pattern as before. Lastly, the third and fourth panels plot Δ_{12}^2 and Δ_{23}^2 for auctions with $\varepsilon < 0.01$, but all the other conditions are kept the same as the top two panels. Our results seem quite robust to the choice of ε .

5 Case Study

In this section, we analyze four collusion cases that were implicated by the JFTC during our sample period. The four cases that we examine are the bidding ring of (A) prestressed concrete providers [JFTC Case H16(52e7)27-29], (B) firms installing traffic signs [JFTC Case H17(52e7)5-8], (C) builders of bridge upper structure [JFTC Case H17(52e7)27-29], and (D) floodgate builders [JFTC Case H19(63aa)2-5]). In all of these cases, the ring members were charged with deciding on a predetermined winner for each project (often based on factors such as whether a project is related to an existing project¹⁷ or the amount of auctions each firm has won in the past). The predetermined winner would then communicate to the rest of the ring members how each of them should bid. All of the implicated firms in cases (B), (C) and (D) admitted wrongdoing soon after? within a year? of the investigation, but all of the firms implicated in case (A) denied any wrongdoing and the case went to trial.¹⁸

Before we analyze these four cases, we point out one interesting feature of the bidding ring in case (A). According to the ruling in case (A), an internal rule existed among the ring members which prescribed that (1) the predetermined winner should aim to win the auction in the first round; (2) if the predetermined winner did not win in the first round, the predetermined winner should apply a prespecified discount to its first round bid; (3) the rest of the ring members should apply a smaller discount to the first round bid of the predetermined winner.¹⁹ The prespecified discount used in the ring was 3% for auctions with an expected value less than 100 million yen, 2% for auctions with an expected value between 100 million yen and 500 million yen, and 1.5% for auctions expected to worth more than 500 million yen. One consequence of this internal rule is that we would observe the same lowest bidder in round one and round two.

In Figure 4, we plot the winning bid (as a percentage of the reserve price) for auctions in which the winning bidder is a member of one of the implicated bidding rings. The winning bid is the bid submitted by the winner, i.e., the lowest bid of the concluding round. We

¹⁷p5 of [REDACTED]

¹⁸Out of 20 firms that were initially implicated, one firm was acquired by another firm, one firm was acquitted and the rest of the firms subsequently admitted wrongdoing.

¹⁹[REDACTED]

have also drawn in a vertical line that corresponds to the “end date” of collusion.²⁰ The “end date” corresponds to the day appearing in the initial JFTC ruling after which the ring members were deemed to stop colluding.

We see that for cases (2), (3), and (4), there is a general drop in the winning bid of about 5-6% after the collusion end-date. However, it is worth mentioning that there are still some auctions where the winning bid is extremely high even after the end-date – in fact, about 30.4% of auctions after the end-date have a winning bid higher than 95%. While the investigation and the ruling of the JFTC seemed to have made collusion harder, it is far from clear whether the price levels after the end-date are at truly competitive levels as we discuss more below. Hence, the 5-6% price drop may be a conservative estimate of the effect of collusion.

We now examine whether we observe persistence in the second round bid of $R_1(1)$ and $R_1(2)$ during the period in which the firms were colluding. Figure 5 plots the histogram of Δ_{12}^2 and Δ_{23}^2 before the collusion end-date for each of the four bidding rings. The sample used for the figure corresponds to the set of auctions where the first stage bids of $R_1(1)$ and $R_1(2)$ are less than 5%, i.e., $\varepsilon = 0.05$. We see that for all four bidding rings, Δ_{12}^1 is asymmetric around zero while Δ_{23}^2 is symmetric around zero, as before.

In Figure 6, we again plot the histogram of Δ_{12}^2 and Δ_{23}^2 with $\varepsilon = 0.05$ for each of the four bidding rings, but only for auctions occurring after the collusion “end-date”. Although the sample size is very small, the figure shows a similar pattern as before. This may seem to cast doubt on our analysis – to the extent that firms bid competitively after the collusion end date, why does the figure show a similar pattern in Δ_{12}^2 and Δ_{23}^2 as before?

We believe that asymmetry in the distribution of Δ_{12}^2 should be taken as evidence that firms may have been able to continue colluding at least on some auctions even after the “end date”, i.e., the date after which the firms were deemed to have stopped colluding by the JFTC. As the drop in the winning bid suggests in Figure 4, the investigation certainly seems to have made collusion harder among the members. But this does not necessarily mean that the firms completely ceased to collude. For example, in a subsequent ruling which determined the fines of each firm involved in case (A), the JFTC stated that there were still circumstances conducive to collusion even after the “end-date” that appeared in the initial ruling. With respect to case (A), the ring members may well have been able to continue using the prespecified discount (e.g., 3% for auctions with value less than 100

²⁰This date does not always coincide with the date the investigations started because some firms continued to collude even after the start of the investigation. We use the date deemed by the JFTC in its ruling to be the last day of collusion.




Figure 4: Winning Bid of Auctions of Which the Winner Was Involved in One of the Four Bidding Rings. The x-axis corresponds to the calendar date from the beginning of our sample (i.e. April 1 of 2003) and the y-axis corresponds to the winning bid as a percentage of the reserve rate. The vertical line in each of the four panels corresponds to the collusion “end date”.




Figure 5: Difference in the Second Round Bids of $R_1(2)$ and $R_1(1)$ (Left Panels) and the Difference in the Second Round Bids of $R_1(3)$ and $R_1(2)$ (Right Panels) Before the Collusion End Date. We use $\varepsilon = 0.05$, hence the differences in the first round bids are relatively small.

million yen, etc.) for auctions that proceeded to the second round. It is also worth pointing out that many firms that were implicated in these cases are repeat offenders. For example, one firm involved in case (A) had been found guilty in four previous collusion cases while a number of firms implicated in case (C) were subsequently charged and found guilty of collusion in a separate case by the JFTC. From this, we conclude that being implicated is no guarantee that a firm will behave competitively thereafter and that firms may have been able to continue colluding well beyond the “end date”, at least on some auctions.

6 Detection of Collusive Bidders

In this section, we develop a formal statistical test of collusive behavior based on the idea we discussed in section 4.2, namely, the distribution of Δ_{12}^2 should not be discontinuous at zero under competitive bidding. We then apply our test to each firm in order to examine whether its bidding behavior is consistent with competitive bidding.

Test Statistic Recall from section 4.2 that there is a reasonable amount of variance in Δ_{23}^2 among bidders who submit almost identical first-round bids. This suggests that there is a reasonable amount of idiosyncracies with regard to the beliefs over the distribution of the reserve rate, risk preference, etc., that induces variance in the second round bids. To the extent that this is the case, $R_1(2)$ should outbid $R_1(1)$ in the second round by a narrow margin just as often as $R_1(1)$ outbids $R_1(2)$ by a narrow margin, i.e., the amount of idiosyncracies measured by the variance of Δ_{23}^2 puts a bound on how sharp the distribution of Δ_{12}^2 can change around zero. The test statistic that we propose below formalizes this idea by looking for violations of this bound.

We begin by specifying the second round bids of $R_1(2)$ and $R_1(3)$ as follows,

$$\begin{aligned} b_{R_1(2)}^2 &= X + u_2 \\ b_{R_1(3)}^2 &= X + u_3, \end{aligned}$$

where X is a common component and u_2, u_3 are bidder specific shocks distributed independent and identically according to F_u . As long as we condition on auctions where the first round bids of $R_1(2)$ and $R_1(3)$ are close enough, this specification seems natural: Both $R_1(2)$ and $R_1(3)$ should have similar cost structures and similar information, which is captured in the common component, X . Note that X is a random variable whose distribution can arbitrarily depend on the object being auctioned, information revealed in the first round,




Figure 6: Difference in the Second Round Bids of $R_1(2)$ and $R_1(1)$ (Left Panels) and the Difference in the Second Round Bids of $R_1(3)$ and $R_1(2)$ (Right Panels) After the Collusion End Date. We use $\varepsilon = 0.05$, hence the differences in the first round bids are relatively small.

etc. The error terms, u_2 , and u_3 are independent bidder specific components that result from differences in the bidders' beliefs over the secret reserve price, heterogeneity in the bidders' risk preference, etc. We assume that u_i is independent of u_j ($j \neq i$) and X . Now, given that Δ_{23}^2 is just the difference between $b_{R_1(3)}^2$ and $b_{R_1(2)}^2$,

$$\begin{aligned}\Delta_{23}^2 &\equiv b_{R_1(3)}^2 - b_{R_1(2)}^2 \\ &= u_3 - u_2.\end{aligned}$$

Note that given our assumptions on (u_2, u_3) , we can recover F_u from realizations of Δ_{23}^2 .

We now consider putting bounds on the distribution of Δ_{12}^2 using F_u . Let us denote by Y , the second round bid of $R_1(1)$:

$$b_{R_1(1)}^2 = Y.$$

Given that $R_1(1)$ has a different information set than all of the other bidders (as well as perhaps having different cost), we do not impose any restrictions on the distribution of Y other than independence with respect to (u_2, u_3) , i.e., $Y \perp (u_2, u_3)$. In particular, Y can have arbitrary correlation with respect to X . Now, given that $\Delta_{12}^2 = b_{R_1(2)}^2 - b_{R_1(1)}^2$, we have $\Delta_{12}^2 = X + u_2 - Y$. Then, consider $d(t)$ ($t \in \mathbb{R}^+$) which we define as

$$d(t) = \Pr(\Delta_{12}^2 \in [0, t]) - \Pr(\Delta_{12}^2 \in [-t, 0]).$$

$\Pr(\Delta_{12}^2 \in [-t, 0])$ is just the probability that Δ_{12}^2 falls within $[-t, 0]$ and $\Pr(\Delta_{12}^2 \in [0, t])$ is the probability that Δ_{12}^2 falls within $[0, t]$. Hence, $d(t)$ is the difference between the probability that Δ_{12}^2 falls just to the right of zero and the probability that Δ_{12}^2 falls just to

the left of zero. We can derive a simple bound of $d(t)$ using F_u after some algebra,

$$\begin{aligned}
d(t) &= \Pr(\Delta_{12}^2 \in [0, t]) - \Pr(\Delta_{12}^2 \in [-t, 0]) \\
&= \int \mathbf{1}_{\{X+u_2-Y \in [0, t]\}} dF_{X,Y}(X, Y) dF_u(u_2) \\
&\quad - \int \mathbf{1}_{\{X+u_2-Y \in [-t, 0]\}} dF_{X,Y}(X, Y) dF_u(u_2) \\
&= \int F_u(Y - X + t) - F_u(Y - X) dF_{X,Y}(X, Y) \\
&\quad - \int F_u(Y - X) - F_u(Y - X - t) dF_{X,Y}(X, Y) \\
&= \int F_u(Y - X + t) + F_u(Y - X - t) - 2F_u(Y - X) dF_{X,Y}(X, Y) \\
&\leq \sup_x \|F_u(x + t) + F_u(x - t) - 2F_u(x)\|,
\end{aligned}$$

where the second line uses independence of u_2 with respect to Y and the fourth line replaces $Y - X$ with x .

Our test statistic simply compares $d(t)$ with the bound derived from F_u . Define $\tau(t)$ as

$$\tau(t) \equiv \sup_x \|F_u(x + t) + F_u(x - t) - 2F_u(x)\| - d(t).$$

Given that we can estimate F_u and $d(t)$, we can estimate $\tau(t)$. Under the null hypothesis of competitive behavior, $\tau(t)$ should be nonnegative.

Detecting Collusive Bidders We now apply this test to each firm that we observe in the data. In particular, for a given firm, we collect all auctions in which the firm participated. We then estimate $d(t)$ and F_u parametrically using realizations of Δ_{12}^2 and Δ_{23}^2 from a subset of these auctions where (1) the auction proceeded to the second round; and (2) the first round bids of $R_1(2)$ and $R_1(3)$ were sufficiently close to each other, $b_{R_1(3)}^1 - b_{R_1(2)}^1 < \varepsilon$.²¹ For $d(t)$, we use a frequency estimator, and for F_u , we use maximum likelihood by specifying F_u to be a mean-zero Normal distribution with parameter σ_u , $u \sim N(0, \sigma_u^2)$. While our test statistic can easily accommodate a nonparametric estimate of F_u , we imposed functional form assumptions on F_u because the number of auctions per firm is not very large. In practice, we estimate $\tau(t)$ for every firm who participated in at least five auctions that

²¹Note that we condition on the set of auctions where the second and third lowest bids in the first round are within ε , given the assumptions on u_2 and u_3 . Note also that we drop auctions if Δ_{23}^2 is bigger than 30% to make sure we exclude misrecordings etc. Note that this condition biases against us.

meet the two criteria mentioned above. We compute an estimate of $\tau(t)$ for each of 3,559 firms, in total. Given our parametric assumption on F_u , $\tau(t)$ has an asymptotically Normal distribution.

In the top left panel of Figure X, we plot the estimates of $\tau(t)$ for each firm for $t = 1\%$ and $\varepsilon = 5\%$. As shown in the panel, the estimated distribution of $\tau(t)$ lies somewhat to the right of zero, but there is also a substantial mass below zero as well. Under the null hypothesis of competitive bidding, the value of $\tau(t)$ should be positive, hence a negative estimate of $\tau(t)$ raises concerns about possible collusive behavior. In the top right panel, we plot the t -statistic for each firm.²² Again, we find that the estimated t -statistic is negative for a substantial fraction of firms. Moreover, there are 583 firms whose t -statistic is less than -1.65 , which is the one-sided critical value for rejecting the null hypothesis of competitive behavior at 95%.²³ Among these firms, 21 firms correspond to those that were implicated in one of the four bid rigging cases. In the second row of Figure X, we plot our estimate of $\tau(t)$ and the t -statistic for $t = 2\%$ and $\varepsilon = 5\%$. The results are qualitatively similar. For this case, we find that 508 firms have t -statistic less than -1.65 .

It should be clear from the construction of the test statistic that the value of $\tau(t)$ should be nonnegative for all values of t under competitive bidding. Hence we next conduct a joint hypothesis test for $t = 1\%$ and 2% , i.e., test whether $(\tau(1\%), \tau(2\%))$ is jointly nonnegative. Under the joint hypothesis test, we find that we can reject the null for 691 firms. The joint hypothesis test picks out 26 firms that were implicated in one of the four bid rigging cases.²⁴

To get a sense of the magnitude of our findings, note that the total number of all auctions that were won by the 691 “suspicious” firms is about 6,900 auctions, or about 1/6 of the total number of auctions. The total award amount of these auctions sums to about \$7.9 billion. Given that the four case studies show about a 5-6% average drop in the winning bid after bidding rings are implicated, our results suggest that taxpayers could have saved about \$425 million in the absence of collusion. Moreover, if we consider the fact that the total amount of municipal and prefectural construction projects awarded in Japan is close to 10 times larger than our data set, the total impact of collusion can even be bigger as a whole. There is also ample reason to believe that collusion is just as rampant among municipal and

²²We only include the estimated test statistic for firms with more than 5 observations, i.e., the firm has participated in at least 5 auctions that (1) proceeded to the second round; and (2) $b_{R_1(3)}^1 - b_{R_1(2)}^1 < \epsilon$. Same for the t -statistic.

²³Note that we conduct a one-sided test, so the critical value is -1.65 .

²⁴In practice, we estimated the joint (2dimensional) distribution of $(\tau(1\%), \tau(2\%))$. We then simulated 500 draws of (2×1) random vectors according to the simulated joint distribution. We test whether there are more than 25 ($= 5\%$ of 500) draws whose elements are both positive.

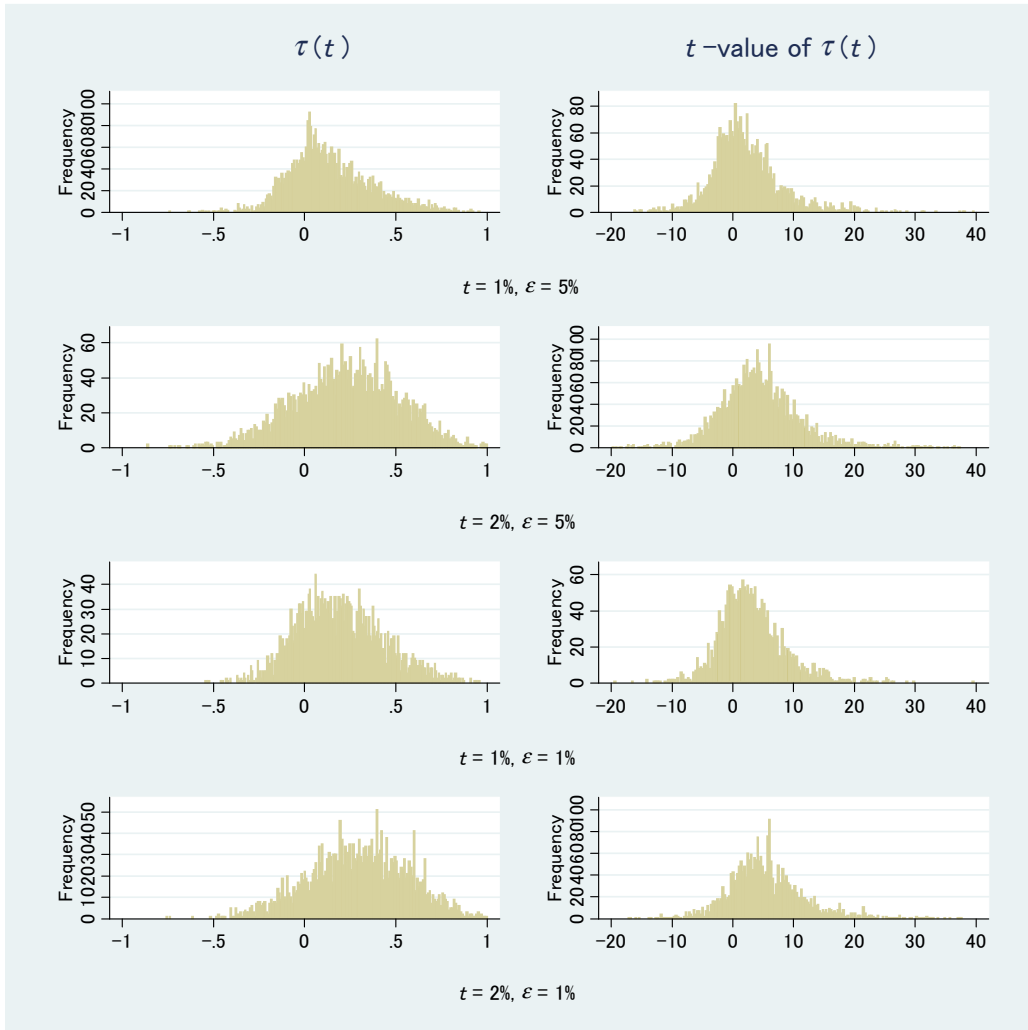


Figure 7: Estimate of $\tau(t)$ (Left Panel) and t -Statistic (Right Panel). We estimated $\tau(t)$ for each firm using only the subset of auctions in which it participated. Top two panels plot the histogram for $t = 1\%$ and $t = 2\%$ with $\epsilon = 5\%$. Bottom two panels plot the histogram for $t = 1\%$ and $t = 2\%$ with $\epsilon = 1\%$.

prefectural construction projects given that some of the same construction firms in our data set participate in these auctions as well.

7 Conclusion

In this paper, we document large-scale collusion among construction firms in Japan using bidding data of procurement auctions. We find evidence of collusion across region, type of construction projects and time. We then test, for each firm, whether its bidding behavior is consistent with competitive behavior. Our test identifies close to 700 “suspicious” firms, who won a total of about 6,900 auctions, or about 1/6 of the total number of auctions during our sample period.

whose bidding behavior is inconsistent find data collusion detection>for making arrests, need more evidence by JFTC. it’s only detection at 95% confidence. blah blah.Talk about institutional details that make collusion easy somewhere in the paper? E.g. restrictions on entry, low sanctions, not aggressive investigations etc. Include in Conclusion?

References

- [1] Asker, John, “Bidding Rings,” *New Palgrave Dictionary of Economics Online*, Eds. Steven Durlauf and Lawrence Blume, Palgrave Macmillan 2010.
- [2] Asker, John, “A Study of the Internal Organization of a Bidding Cartel,” *American Economic Review*, Vol. 100, No. 3, 2010, pp. 724-762
- [3] Athey, Susan, Jonathan Levin and Enrique Seira, “Comparing open and Sealed Bid Auctions: Evidence from Timber Auctions,” *Quarterly Journal of Economics* 126 (1), 2011, pp.207-257.
- [4] Bajari, Patrick and Lixin Ye, “Deciding Between Competition and Collusion,” *Review of Economics and Statistics*, Vol. 85, No. 4, 2003, pp. 971-989
- [5] Baldwin, Laura, Robert Marshall, and Jean-Francois Richard, “Bidder Collusion at Forest Service Timber Sales,” *Journal of Political Economy*, Vol. 105, No. 4, 1997, pp. 657-699
- [6] Conley, Timothy and Francesco Decarolis, “Detecting Bidders Groups in Collusive Auctions,” working paper, Boston University

- [7] Haile, Philip, Kenneth Hendricks, Robert Porter, and Toshi Onuma, "Testing Competition in U.S. Offshore Oil and Gas Lease," 2012, working paper, Yale University
- [8] Ishii, Rieko, "Favor Exchange in Collusion: Empirical Study of Repeated Procurement Auctions in Japan," *International Journal of Industrial Organization*, Vol. 27, Issue 2, 2009, pp. 137-144
- [9] McMillan, John, "Dango: Japan's Price-Fixing Conspiracies," *Economics and Politics*, vol. 3(3), 1991, pp.201-218.
- [10] Ohashi, Hiroshi, "Effects of Transparency in Procurement Practices on Bidding Behavior: The Experience of Municipal Public Works," *Review of Industrial Organization*, 34(3), 2009: 267-85
- [11] Pesendorfer, Martin, "A Study of Collusion in First-Price Auctions," *Review of Economic Studies* Vol. 67 (3), 2000, pp.381-411.
- [12] Porter, Robert and J. Douglas Zona, "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy*, Vol. 101, No. 3 1993, pp. 518-538
- [13] Porter, Robert and J. Douglas Zona, "Ohio School Milk Markets: An Analysis of Bidding," *RAND Journal of Economics*, Vol. 30, No. 2 1999, pp. 263-288

8 Appendix

Analysis on Collusive Behavior by Region, Auction Category, Project Type, and Time Many circumstances support the collusive bidding in the Japanese procurement market. As mentioned earlier, even in the fourth category of the auction, which imposes the weakest participation restriction, contractors must be pre-qualified at the regional level. This insulates bidders from outside competition and facilitates cooperation among them. In addition, almost all contractors are affiliated with some industrial and regional business associations, which are used as a vehicle for organizing collusion. Furthermore, although these have been strengthened year by year, penalties against implicated collusion members are still weak relative to other countries. According to news articles, almost all of the public construction projects in Kansai region in 1996 were allocated through bid riggings ().

Asymmetry of Δ_{12}^3 Another feature of the data that seems to be at odds with competitive bidding is that the distribution of Δ_{12}^3 continues to be very asymmetric even in situations

where the informational asymmetry among the bidders seems reasonably modest. Recall that the lowest bid in each round is announced before proceeding to the next stage. This introduces an informational asymmetry between the lowest bidder and all the other bidders. Even if we control for cost differences by conditioning on the set of auctions where the first round bids are sufficiently close to each other, there can be systematic differences in the bidding pattern between the lowest bidder and the other bidders in later rounds because of informational asymmetry. However, when the lowest bid in the first round ($b_{R_1(1)}^1$) and the lowest bid in the second round ($b_{R_2(1)}^2$) are sufficiently close, all bidders learn that the second round bid of everyone is in $[b_{R_2(1)}^2, b_{R_1(1)}^2]$ before they bid in the third round. That is, if we look at auctions that reach the third round and $[b_{R_2(1)}^2, b_{R_1(1)}^2]$ is sufficiently close, we would expect to see little informational asymmetry among the bidders going into the third round. As long as we look at auctions that reach the third round and (1) condition on auctions where $[b_{R_2(1)}^2, b_{R_1(1)}^2]$ is small, and (2) condition on the set of auctions where the first round bids are sufficiently close to each other so that the identities of $R_1(1)$, $R_1(2)$, and $R_1(3)$, etc. are as good as random, we should expect the third round bids to be roughly identical.

In the left panel of row of Figure ??, we plot the histogram of $\Delta_{12}^3 = b_{R_1(2)}^3 - b_{R_1(1)}^3$ for the subset of auctions where the lowest bid in the first and second rounds are within 3% of each other ($b_{R_1(1)}^2 - b_{R_2(1)}^2 < 0.03$) and the first round bids of $R_1(2)$ and $R_1(1)$ are within 3% ($b_{R_1(2)}^1 - b_{R_1(1)}^1 < 0.03$). In the right panel, we plot $\Delta_{23}^3 = b_{R_1(3)}^3 - b_{R_1(2)}^3$ conditioning on $b_{R_1(1)}^2 - b_{R_2(1)}^2 < 0.03$ and $b_{R_1(3)}^1 - b_{R_1(2)}^1 < 0.03$. Note that the shape of the histogram in the left panel is still asymmetric around zero while the distribution in the right panel is symmetric. Conditioning on events that seem

is announced the fact that have relatively symmetric information is almost invariant to the value of information revealed in the second round.