

Modeling and Estimating the Welfare Consequences of Competitive Collusion in Average Price Auctions with Generalized Linear Models

Benjamin Thomas

June 10, 2024

Abstract

Limited research has been done on average price auctions (APAs) since they do not maximize profit for the auctioneer and appear to be, by design, less efficient than first-price auctions (FPAs). However, despite this, numerous countries prefer APAs to FPAs for procurement auctions because they may be effective at preventing harms deriving from the Winner's Curse and bidding rings. While there is some evidence that APAs are effective at stopping the Winner's Curse, there is very limited work on the effectiveness of APAs at reducing harms from collusion. I use a dataset of Italian APAs where legal proceedings have made the relevant bidding rings known and estimate the marginal harm from bidding rings. I find that there is a linear relationship between the number of bidders in an auction and the marginal welfare effect of an increased collusive bidder and, unexpectedly, that there is a marginal welfare gain in auctions with 60 or less bidders. However, in large auctions, I find marginal welfare losses as high as 5% of reserve value.

1 Introduction

1.1 Why You Should Care About Average Price Auctions

There is an unfortunate dearth of research on Average Price Auctions (henceforth APAs). This can be attributed, in large part, to the fact that APAs are both uncommon and unintuitive, as they directly decrease the profit of the auctioneer. This effect is further compounded by the complexity of APAs' strategy space. In fact, considering second-order strategies where players attempt to react to each other's strategies, each player i 's best response function to the other players $-i$ is exponential with the size of $||-i||$, and excluding the special case of "play the mode", finding Bayesian Nash equilibria is an NP-hard problem (see McLennan). This suggests that APAs are not just unattractive to research,

but unattractive in general. However, many countries, including Italy, Taiwan, and Malaysia, as well as some U.S. states, such as New York, still prefer the use of APAs to First Price Auctions (henceforth FPAs) for procurement auctions. In fact, while Italy was forced to switch to FPAs in 2008 by a regulatory change in the European Union, when this regulation was repealed in 2011, Italy voluntarily transitioned back to APAs (Conley and Decarolis).

Empirical and theoretical evidence both suggest that, when there is little uncertainty over the true cost of a project, APAs result in less efficient outcomes and higher procurement costs than FPAs (Decarolis 2018, Chang et al.). Given that Italy alone procures over 10 billion euros of public projects every year using APAs, it seems shocking that Italy would volunteer to keep these costs high, as would New York or Malaysia. However, if we zoom out and consider institutional behavior more broadly, there are two main benefits attributed to APAs.

The first set of benefits can be described by the intuition that "the cheapest bridge is probably not the best bridge." If an auction is sufficiently competitive, then the firm that can complete the project for the cheapest will win the contract. However, in real life, this gives ample reasons to suspect that the winning firm may be cutting corners or taking risks, and even assuming well-specified contracts and cost-less legal adjudication, it would have been better to pay a firm more to do it right the first time. Even worse, we can easily imagine that the firm is unable to pay the fee prescribed by the contract and declares bankruptcy with the project unfinished. Furthermore, and endemically, winning firms fall victim to the Winner's Curse. Despite the winning firm's best intentions, it becomes evident that the winning firm underestimated the cost of the contract and as a consequence, the firm is unable to complete the project for the auction price, leading to delays, extra costs, and legal proceedings. Unfortunately, even in auctions where we would think the price is pretty clear-cut and there should be little uncertainty over the true cost of a project, such behavior is common and frustrating (Chen et al.).

The second set of benefits relates to bidding rings. Unlike an FPA, it is difficult for a bidding ring to collude to win an APA, since the result is contingent on every bidder, not just the highest bidders, making the winner largely random. Further, collusive strategies usually result in lowering the procurement price rather than raising it. While unfair, this suggests bidding rings may actually be desirable.

Given the numerous practical difficulties which APAs may help to solve, more ink should be spilled assessing to what extent these benefits manifest, as well as the broader welfare implications of these benefits. It certainly is not the case that, *a priori*, FPAs are always better than APAs, especially when we expect collusive behavior or are uncertain about the true cost of a contract. I present an overview of previous research assessing whether APAs actualize these benefits below, and I contribute to this

discussion by estimating the marginal welfare impact of collusion in a data set of APAs where it is known which firms are colluding. A reader may object that even if APAs are sometimes more useful than FPAs, they surely are never a better mechanism than Vickrey auctions. However, note that, while the degree is lessened, Winner's Curse and collusion are still serious issues in Vickrey auctions.

1.2 Is The Cheapest Bridge the Best Bridge?

It is well known that the dominant strategy in an FPA for a firm i is to bid precisely its expected costs, plus some extra amount based on its expectations regarding the bids of other players, firms $-i$. By contrast, the analysis of the dominant strategy of an APA is more difficult. Unfortunately, while many Bayesian Nash equilibria are known to exist, these equilibria are arbitrary and calculating them is an NP-hard problem, which suggests that these equilibria never occur in practice (Thomas 2024). However, there is known to be a bidder-optimal Nash Equilibrium where all players bid precisely the reserve price. On inspection, there is clearly no profitable deviation, as the average will always be precisely the reserve price, so both higher bids and lower bids never win. However, this is a special case of a class of Bayesian Nash equilibria, which I call "play the mode."

Perhaps due to the incorrect application of intuition learned in FPAs, or perhaps due to the implicit threat that future auctions could be switched to FPAs, which are less profitable for bidders if they do not cooperate, the "bid the reserve" equilibrium does not seem to emerge in practice. Instead, a consensus about a "correct" discount emerges over time, and all players randomly select a value near this discount (for example, as in Conley and Decarolis). As a result, the winning bid in APAs is usually the mode. Clearly, no profitable deviation from this strategy exists without collusion, since any bidder who bids away from the publicly known "discount" will not have any chance of winning the auction.

This suggests that the real-world result of a competitive APA is to assign the procurement contract to a random bidder at a known, but somewhat random, discount to the reserve price. This analysis essentially holds even when there are bidding rings (argued in section 1.3), although collusion can steer the average to favor cartels by lowering the average bid. The "bid the mode" equilibrium has concerning implications for the efficacy of APAs. Namely, since the prices in APAs are set mostly without reference to the true cost of a service, this suggests that APAs pay an arbitrarily higher price than an FPA in hopes that this avoids risky cost-cutting and the Winner's Curse.

Unfortunately, as of the writing of this paper, there is no empirical evidence on the question of whether APAs are effective at preventing cost-cutting and the Winner's Curse, although the theory would clearly suggest it. However, an experiment conducted by Chang et al. found that APAs essentially eliminated the Winner's Curse, as expected. Unfortunately, we have strong reason to be

skeptical of these results, as the experimental subjects used indistinguishable bid strategies in both APAs and FPAs and did not converge to the "play-the-mode" Nash equilibrium observed in real APAs. Further, while failing to complete procurement contracts within budget is an endemic problem, empirical evidence suggests that firms are responsive to incomplete contracts and price in the implied risk, which gestures towards both the Winner's Curse and cost-cutting being symptoms of inexperienced firms winning contracts (Chen et al.). Substantial evidence, using licensing and auction-approval mechanisms as proxies for firm skill, suggests that inexperienced firms are substantially more likely to fall into these traps (Moretti and Valbonesi, Decarolis 2014). Taken together, current empirical evidence suggests that well-written contracts and effective firm licensing are the most direct solutions to this problem, but given the direct and implied costs of these mechanisms, and their middling results in practice (Hojin), the relative efficacy of APAs should be more thoroughly considered.

1.3 Collusion in APAs: How-to

We now consider the efficacy and research on bidding rings in APAs. For simplicity, we will only consider the real-world equilibrium where most bidders use the "play the mode" strategy, and we assume the mode discount is publicly known. This assumption is justified by the medium-term tendency of repeated APAs over similar goods to converge to this strategy profile (Decarolis 2018).

Unlike in FPAs, where bidding rings intentionally bid high to win a contract at a high price (so-called "demand reduction"), in APAs bidding rings have to "steer" the bid, usually through having a number of firms engage in demand exaggeration by bidding unreasonably low. Collusive firms hope that designated winning firms are more likely to win the bid as a result of the steering. Consider the following auction as an example:




Figure 1: Visualization of "bid steering" in an APA from Conley and Decarolis. Shapes correspond to cartel membership; the goal of bid steering is to make the dark rhombus win.

However, this strategy has a number of issues. For one, other bidding rings may be steering the bid as well. Above, we see that the "X" cartel and "Rhombus" cartel are both attempting to steer the median upwards, unbeknownst to each other. Even more interestingly, "norms" around steering appear to develop (as seen by the X cartel and the Rhombus cartel using a similar discount for steering) to reduce competition between cartels. In theory, this allows firms outside the bidding rings to free-ride by anticipating the level of steering, and it also allows cartels to defect by steering in the opposite direction. As a result, *a priori*, it is not clear what the net results on welfare are (Decarolis 2018). There are also logistical questions: how many firms should steer, and how many should try to win? Further, does a bidding ring make this behavior look legal? However, despite the complications, steering does appear to be effective. Consider the following:




Figure 2: Empirical Chance of a cartel winning based on how many firms from the cartel participate in the auction. Counterfactual from "honest play" in red.

As we can see, there is a clear linear relationship between the number of cartel members engaging in bid steering and its likelihood of success. However, even with many firms engaged in steering schemes, success is hardly guaranteed. In fact, even in the best case with 13 allied collusive firms, the success rate is only 40%! This is not because APAs appear to give firms that play fair a real chance. Roughly 69% of auctions in the dataset were won by firms participating in bidding rings, compared to only 31% by non-collusive firms. Even two firms colluding are collectively 6 times more likely to win than a single firm playing fair. By contrast, a blockbuster paper on detecting collusion in Japan's FPA procurement of construction projects (Kawai and Nakabayashi) found that only 37% of projects were won by bidding rings. While these situations are heterogeneous in many important regards, this does seem to strikingly contradict the notion that APAs deter cartel behavior by making bidding rings too difficult, as well as that non-collusive firms can defeat collusion by playing well.

Despite this, while it is well known that collusion in FPAs causes substantial inefficiency, it is unclear that this is the case in APAs. While collusion is, of course, unfair, it usually results in a lower price, and since APAs, by design, pay an arbitrary amount above the true cost of a contract, this lower price should be more efficient and may therefore be good. The overall welfare implications of this, however, are not well understood. Further, there are many more direct ways to combat collusion in FPAs. See, for example, Padhi et al. However, since the question of whether collusion is, in some sense, efficient, is empirical, I spend the rest of the paper attempting to answer it.

2 Estimating the Marginal Welfare Impact of Collusion

2.1 Data

I use a dataset collected by Conley and Decarolis, which comes from the legal office of the municipality of Turin. The dataset consists of 276 APAs conducted by Turin between 2000-2003 to procure roadwork jobs of homogeneous cost and scale. In 2008, Turin convicted many managers of these construction firms for colluding in the APAs. As a result, we know of 8 bidding rings as well as the firms present in each bidding ring. While there is substantial heterogeneity between the number of firms in each cartel, as well as each firm's high-level strategic behavior (for example, deciding when to participate in a particular auction and, if so, how many firms to send), these are analyzed to great effect by Conley and Decarolis. However, given participation in an auction, I find that the underlying data is not very sensitive to these higher-level strategic decisions. This data also includes the discount each firm bid in any given auction and whether the firm won the auction.

2.2 Auction Rules

The rules of the relevant APAs are as follows. First, firms must have appropriate licensing to enter. It is possible to enter with licensing that is only sufficient for some portion of the contract and to sell the rest as subcontracts. Second, all firms anonymously send bids. Third, the average of these bids is taken, μ_1 . Fourth, the bottom 10% and top 10% of bids around μ_1 are then eliminated, and μ_2 is calculated from the remaining bids. Fifth, and finally, the lowest bid which is strictly greater than μ_2 is selected.

2.3 Method

I provide a sketch of my method here. Technical details follow for those who are interested. My process is as follows: first, I find the expected revenue for firm i 's participation in an auction, conditioned on how many firms collude with i and how many firms collude against i . I clean the resulting data and then fit a generalized linear model (which is a higher dimension analog of linear regression). I then check each auction in the dataset, and for each cartel c in each auction j , I use the model to predict the change in welfare of every firm in auction j if a new firm joined the auction, colluding with cartel c . This data is disaggregated by the total number of firms in the auction and is my final result.

I now provide the technical details. I represent the data as a number of independent trials over a "firm position" plane in $\mathbb{R}^2$. For every firm, i , in any particular auction, j , some number of firms are known to be colluding with firm i in auction j , and some number of firms are known to be colluding against firm i in auction j . These two factors map to the X and Y coordinates of the plane. For example, consider an auction with $N = 10$ firms. Three are in cartel A , four are in cartel B , and three are in no cartel. We would then represent the firms in A as (3, 4), the firms in B as (4, 3), and the other firms as (1, 7). Henceforth, this plane will be referred to as the "firm position."

I next construct the expected revenue from participation in an auction based on firm position. This required two mappings of the data to firm position. First, I map firm position to the chance a firm in that position wins the auction (henceforth, WinRate). To calculate WinRate, I iterated over every firm i within every auction j and found firm i 's firm position. I then added 1 to "observations" of firm i 's firm position, and if a member of firm i 's cartel won, I also added 1 to the "victories" of firm i 's firm position. WinRate was then given by:

$$\text{WinRate}(\text{firm position}) = \frac{\text{Victory}(\text{firm position})}{\text{Observations}(\text{firm position})}$$

Next, I created an expected profit function, PerMemberWinningDiscount. As before, I iterated

over every firm, i , within every auction, j , and converted firm i to its firm position. Then, contingent on a member of firm i 's cartel winning, I appended the discount rate at which firm i 's ally won to "discounts" and added 1 to "victories." I then found the overall discount the whole cartel won at, WinningDiscount. WinningDiscount was given by:

$$\text{WinningDiscount}(\text{firm position}) = 1 - \frac{\text{Sum of Discounts}(\text{firm position})}{\text{Victories}(\text{firm position})}$$

However, cartels have multiple members and have to distribute funds among their members. For simplicity, I assume all funds are spread evenly and that all firms are autonomous actors. The truth of this assumption is dubious, as many firms appeared to be shells created just to vote, and many colluding firms were run by family members. Nonetheless, disaggregation by this fact is not possible, nor does it impact the analysis. Thus we have that:

$$\text{PerMemberWinningDiscount}(\text{firm position}) = \frac{\text{WinningDiscount}(\text{firm position})}{\text{Average \# of Cartel Members}(\text{firm position})}$$

For visualization purposes, I have included 2-dimension plots of these variables plotted against only the size of firm i 's allied cartel.




Figure 3: (Repeated) Empirical Chance of a cartel winning based on how many firms from the cartel participate in the auction. Counterfactual from "honest play" in red.




Figure 4: Empirical Percent of Reserve Price in Average Victories by how many firms from the cartel participated in the auction. Counterfactual from "honest play" in red.

We now create the expected revenue function. This follows clearly from PerMemberWinningDiscount and WinRate. In fact, we have that:

$$\text{ExpectedRevenue}(\text{firm position}) = \text{PerMemberWinningDiscount}(\text{firm position}) * \text{WinRate}(\text{firm position})$$

The two-dimension version of this data appears as follows:




Figure 5: Empirical Expected Revenue by Cartel size. Counterfactual "honest play" in red.

While I use a linear regression to draw a trend line through this data, it is quite messy. When we examine the full data in $\mathbb{R}^3$, the data only appears noisier:




Figure 6: Empirical Expected Revenue by Firm Position

However, conducting principal component analysis reveals that the primary two principal components are decreasing revenue as the number of firms colluding against firm i increases and expected revenue increasing as the number of firms colluding with firm i increases. These trends are what we would expect, so we are justified in cleaning the data to extract them.

Unfortunately, this data has a prodigious amount of noise. This is to be expected, since there were 800 discrete firm positions, which were then mapped to ExpectedRevenue, which is continuous. I take three steps to extract features. First is the removal of outliers. Any data point with a firm position that mapped to an expected revenue of over 0.02 was removed. This is justified since any expected revenue corresponding to 2% of the reserve price indicated that a particular firm position had very good luck, rather than indicating a real trend. Since the data for many firm positions was sparse, such outliers would be expected. Next, I implemented smoothing for the number of competing firms. Specifically, for the purposes of calculating ExpectedRevenue, I counted not just auctions from the same firm position as firm i , but also auctions where the number of firms colluding against i was in the range $[\text{of opposing colluding firms} \cdot (1 - q), \text{of opposing colluding firms} \cdot (1 + q)]$. This was justified, as noise prevented capturing the marginal impact of firms colluding against i clearly, which meant there was no substantial trade-off to smoothing, and smoothing would reduce variance. Third, I removed data points where, after smoothing, there were not at least k trials available. This was justified since this removed the data we were the least confident about, thereby reducing noise, without being likely to alter the underlying structure of the data since the data which we are the most confident about has not been removed.

I then fitted a generalized linear model (GLM) to the cleaned data. Since GLMs struggle with noisy data, I used hyperparameter tuning to extract the trends as clearly as possible. To do this, I searched for the best value of the GLM's r^2 , subject to different values of q and k . I tested q in a range

of $[0, 0.5]$ and k in a range of $[0, 30]$. I found the best result at $q = 0.15$, $k = 28$, with $r^2 = 0.66$. An image of this model is included in the results.

Finally, I used the GLM to test the marginal cost of adding another collusive firm to the original dataset of auctions. For each cartel, c , in each auction, j , I found the change in expected revenue for each firm, i , in auction j if a new firm was added to cartel c . Importantly, this change in expected revenue included the large positive expected revenue that the new firm received from participating in the auction, as opposed to not participating. I then disaggregated this data by the number of firms n in each auction j and took the average of the change in expected revenue. This yielded the average per-firm change in expected revenue by the disaggregated number of firms. I then multiplied each disaggregated data point by the number of firms n in its associated auctions to find the total marginal change in revenue caused by the introduction of a new collusive bidder. This yielded the image included in results.

2.4 Results

First, I fitted a generalized linear model to the $\mathbb{R}^3$ dataset of expected revenue by firms colluding with and against firm i . I report an r^2 of 0.662. Further, the generalized linear model is given by the following equation, where the sum 1 represents the reserve price:

$$\text{Expected Revenue} = 0.02 + 0.003(\text{Firms colluding with firm } i) - 0.0006(\text{Firms colluding against firm } i))$$




Figure 7: Generalized Linear Model of Expected Revenue by Collusion in APA

Further, I estimate the marginal change in welfare of bidders by the introduction of a new collusive bidder. Since this trend is linear, I fit a linear regression, although it trivially returns $r^2 = 0.99$ since

this trend results from the application of the generalized linear model. The equation of the linear regression is given by the following, where n is the number of firms in the auction:

$$\Delta \text{Welfare} = 0.034 - 0.00053n$$

Note that this implies that $\Delta \text{Welfare}$ is negative when $n \geq 65$.




Figure 8: Net change in Welfare by the addition of a collusive firm by the number of firms in auction.

2.5 Discussion

These results are both surprising and counterintuitive. In general, we would not expect there to be situations where collusion has positive welfare effects. Further, the marginal welfare effects of collusion, as auctions get larger, become quite significant. We should take a moment to consider what this represents to understand why this makes sense. When a new collusive firm is added, its chance of expected revenue comes at the cost of competing firms, decreasing their revenue. However, firms in a cartel with the new firm have a higher chance of winning, increasing expected revenue. In small auctions where the marginal increase in win rate caused by the entrant is high, and each individual cartel is quite large relative to the pool of bidders, the "helping of friends" effect is quite strong. However, as the pool gets larger, the number of other firms grows at a larger rate, causing the "harming of enemies" effect to outpace the "helping of friends" effect. Further, non-collusive bidders have a preference for larger auctions and are harmed dramatically by collusion since they have no network of collusive allies. As auctions get bigger, the number of defenseless "fair players" who are harmed increases.

A close reader may remember that collusion tends to result in contracts being procured for lower prices and may wonder what portion of this welfare is transferred to the government. Unfortunately,

developing measures that test this directly is complex and this paper is already quite long. However, based on the more general estimates of (Conley and Decarolis), I would speculate that at least half of the welfare loss is in fact a transfer from bidders to the government, although it seems unlikely that this is the entirety of the effect.

3 Conclusion

3.1 Parting Thoughts

In this paper, I have argued that average price auctions may be a more appropriate option than first-price auctions or Vickrey auctions in situations where we are concerned about knock-on effects from the Winner's Curse or about the harms caused by collusion. However, average price auctions are known to not prevent collusion but are only suspected of limiting the harms caused by collusion. I contribute to this conversation by providing the first empirical estimation of the welfare impacts of collusion in average price auctions. I find that while collusion is welfare-positive in small auctions, roughly $n < 65$, collusion has significant negative impacts on welfare in larger auctions. This suggests that average price auctions are an effective means of reducing the harms caused by collusion, but only in small auctions. Therefore, we should rely on traditional methods for large auctions.

3.2 References

- Conley, Timothy and Francesco Decarolis. 2016. "Detecting Bidders Groups in Collusive Auctions," *American Economic Journal: Microeconomics*, 8(2).
- Chang et al. 2015. "An Investigation of the Average Bid Mechanism for Procurement Auctions," *Management Science*, 61:6, pp. 1237-1254.
- Chen et al. 2009. "Managing Project Failure Risk Through Contingent Contracts in Procurement Auctions," *Decision Analysis*, 7(1).
- Decarolis, Francesco. 2018. "Comparing Public Procurement Auctions," *International Economic Review*, 59(2).
- Decarolis, Francesco. 2014. "Awarding Price, Contract Performance, and Bids Screening: Evidence from Procurement Auctions," *American Economic Journal: Applied Economics*, 6(1).
- Hojin, Jung. 2016. "Renegotiation on Incomplete Procurement Contracts," *Applied Economics*, 23.

-
- Kawai, Kei and Jun Nakabayashi. 2022. "Detecting Large-Scale Collusion in Procurement Auctions," *Journal of Political Economy*, 130(5).
 - McLennan, Andrew. 2011. "The Computational Complexity of Games and Markets: An Introduction for Economists," https://andymclennan.droppages.com/complexity_survey.pdf.
 - Moretti, Luigi, and Paola Valbonesi. 2015. "Firms' Qualifications and Subcontracting in Public Procurement: An Empirical Investigation," *Journal of Law Economics and Organization*, 31(3).
 - Padhi et al. 2016. "Design of Auction Parameters to Reduce the Effect of Collusion," *Decision Sciences*, 47(6).
 - Thomas, Benjamin. 2024. "Arrow-Debreu Markets are Intractable but Efficient when $P \neq PPAD$," Final for CS254.
 - Munoz-Garcia, Felix. 2017. https://felixmunozgarcia.com/wp-content/uploads/2017/08/average_bid_auc

References