



AppNexus
A Xandr Company

Demystifying Auction Dynamics for Digital Buyers and Sellers

INTRODUCTION

At AppNexus, we frequently field questions about the impact of auction dynamics from both buyers and sellers. The great debate, “Which is better for my bottom line: a first- or second-price auction?” has raged at industry events and company meetings, and in the press. Our own CEO has [opined](#) on the subject.

According to auction theory, auction type *should not* impact seller revenue, and yet, we consistently see buyers using other DSPs bidding more than they need to win in first-price auctions. As a result, first-price auctions still offer publishers a significant upside.

AppNexus’ investments in bidding technology, however, have placed our buyers ahead of the curve. For buyers using our Bid Price Optimization algorithm, auction type *does not* impact seller revenue.

So the answer to the question is: it depends... on whether you are a buyer or a seller, as well as on the innovation and advancement of your technology provider.

We show you the evidence in the following pages.



TABLE OF CONTENTS

04	DEFINING AUCTION TYPES
05	WHICH AUCTION TYPE IS BEST?
06	WAIT, BUT WHY?
07	BUT, DO WE GET REVENUE EQUIVALENCE IN THE REAL WORLD?
07	BID PRICE OPTIMIZATION ALGORITHM
09	PRODUCTION RESULTS FROM THE APPNEXUS EXCHANGE
10	OBSERVATIONS
10	UNDERSTANDING DEVIATION FROM REVENUE EQUIVALENCE THEORY IN THE REAL WORLD
11	WHAT ABOUT BUYERS NOT USING BPO?
12	CONCLUSION AND RECOMMENDATIONS



DEFINING AUCTION TYPES

At the outset, we need to establish working definitions of first- and second-price auctions. As well, there are some auctions that do not match either definition. So, for the purposes of this white paper, we'll define a **second-price auction** as one in which

- 1. The winner pays a cent more than the second highest bid or the reserve price.
- 2. Bids below the reserve price cannot win.

By contrast, we define a **first-price auction** as one in which

- 1. The winner pays their submitted bid value.
- 2. Bids need to beat either the reserve price or the hard floor.

We'll use the term mixed-price auction¹ to refer to auctions with soft floors or auctions where the winning bid price is reduced, but not all the way down to the second-highest bid price or to the reserve price.

The table below illustrates the pricing dynamics of these three auction type with the bids held constant.

Example of Definitions

Auction Type	Buyer 1 Bid	Buyer 2 Bid	Winner	Winner Pays
1 st Price	\$10	\$6	Buyer 1	\$10
2 nd Price	\$10	\$6	Buyer 1	\$6.01
Mixed Price	\$10	\$6	Buyer 1	\$8.00



WHICH AUCTION TYPE IS BEST?

The table on the previous page suggests that first-price auctions are great for sellers and bad for buyers. Many in the industry agree. Publishers tend to like first-price while advertisers tend to like second-price.

But in practice, first-price auctions do not always favor the seller. After all, many publishers use mixed-price auctions or soft floors, so the evidence would show that first-price auctions are not one-sidedly advantageous to them. Clearly, buyers will reduce their bid in a first-price auction to avoid overpaying, but by how much and at what loss to the seller?

This seems like a hard question. Fortunately for us, two different Nobel Prize-winning economists have worked through the details. Their conclusion is startling, elegant, or absurd, depending on your point of view. The [Revenue Equivalence Theorem](#), proved originally by Thomas Vickery, holds that, in short, first- and second-price auctions *should* yield the same revenue for sellers in expectation. Mixed-price auctions should as well, at least under some technical assumptions that would seem to cover a wide range of use cases.



WAIT, BUT WHY?

Proving the Revenue Equivalence Theorem would require more Greek letters than anyone is interested in reading. So instead, let's apply some basic intuition to the example above to compare how each auction type would likely play out.

The only new piece of information we need is each buyer's valuation of the item being bid on. Think of the valuation as the amount of money the buyer gets if they win. In digital advertising, for example, it might be the amount of money the brand will pay the agency for an impression.

Below, we've written out each buyer's valuation for each auction scenario. In order to keep things simple, we've kept these valuations the same as the bid amounts in the previous example. Further, let's assume both buyers know each other's valuations. How would we expect each buyer to bid in our three auction types?

Valuations

Auction Type	Buyer 1 Valuation	Buyer 2 Valuation
1 st Price	\$10	\$6
2 nd Price	\$10	\$6
Mixed Price	\$10	\$6

In a first-price auction, Buyer 1 knows to bid one cent more than the valuation of Buyer 2. Buyer 2 has no incentive to bid any higher than his valuation. The same holds in a mixed-price auction because neither a soft floor nor price reduction impacts who wins and Buyer 1 always pays less when she bids less.

By contrast, in a second-price auction, Buyer 1 can bid \$10. She could bid lower, but there is no incentive to do so because she'll only end up paying a cent more than Buyer 2's bid regardless of how much more she bids. Similarly, Buyer 2 would have no incentive to bid higher than his valuation. Thus, we get revenue equivalence.

Full Information with Sophisticated Bidding

Auction Type	Buyer 1 Valuation	Buyer 2 Valuation	Buyer 1 Bid	Buyer 2 Bid	Buyer 1 pays
1 st Price	\$10	\$6	\$6.01	\$6	\$6.01
2 nd Price	\$10	\$6	\$10	\$6	\$6.01
Mixed Price	\$10	\$6	\$6.01	\$6	\$6.01



BUT, DO WE GET REVENUE EQUIVALENCE IN THE REAL WORLD?

In the hypothetical example above, we got revenue equivalence. What about in the more complicated real-world example of digital ad exchanges, in which there are billions of bid requests, thousands of buyers bidding, unknown valuations, repeated auctions, and auction logic that often isn't transparent or consistent? Will revenue equivalence still hold true?

The key issue appears to be technological investment rather than which assumptions hold. When buyers, or their technology platforms, try to bid intelligently across auction types, it is indeed the case that first- and second-price auctions produce very similar results. However, buyers need technology and expertise to optimize for all auction types. If they don't have that technology or expertise, then first-price auctions will yield more revenue for publishers.

Bid Price Optimization algorithm

We've seen the results first-hand. The AppNexus Programmable Platform (APP) recently added Bid Price Optimization (BPO), a bidding algorithm for first- and mixed-price auctions. BPO uses a form of machine learning called "reinforcement learning" to figure out how much to bid, so that the buyer wins without paying more than necessary. The key idea is that even if a buyer does not know ahead of time how others will bid or even how they have bid in the past, APP can learn, through trial and error, the correct amount to bid.

To get a feel for how this plays out, let's use the valuations above but concentrate only on first-price auctions. Let's also suppose that each buyer does not know how their opponent bids but uses the following highly simplified algorithm to decide its bid amount:

- The buyer starts by shading her bid by 20% of her valuation.
- If the buyer wins and has never lost, she reduces her bid by another 10% from her initial valuation.
- Once the buyer loses for the first time, she would increase her bid by 8% from her initial valuation.
- If the buyer wins a round but has also lost before, she reduces her bid by 4% from her initial valuation.
- If the buyer loses twice or more in a row, she increases her bid by 10%, up to 99% of her valuation.



The table below shows the results of seven rounds of auctions between two buyers following the formula above: Buyer 1 whose valuation is \$10, and Buyer 2 whose valuation is \$6.

1st Price Auction with Incomplete Information and Buyers Using Bid Shading Algorithms

Auction Round	Buyer 1 Valuation	Buyer 2 Valuation	Buyer 1 Bid	Buyer 2 Bid	Winner	Winner Pays
1	\$10	\$6	\$8.00	\$4.80	Buyer 1	\$8.00
2	\$10	\$6	\$7.20	\$5.10	Buyer 1	\$7.20
3	\$10	\$6	\$6.00	\$5.70	Buyer 1	\$6.00
4	\$10	\$6	\$5.00	\$5.94	Buyer 2	\$5.94
5	\$10	\$6	\$5.80	\$5.70	Buyer 1	\$5.80
6	\$10	\$6	\$5.40	\$5.94	Buyer 2	\$5.94
7	\$10	\$6	\$6.20	\$5.70	Buyer 1	\$6.20

By the third round, we get the expected result from auction theory. If we carried out the simulation for more auctions, we would get some oscillation but only because the algorithm is over-simplified. It's difficult but not impossible to build an algorithm that can "remember" how much it needs to bid. In addition, though the buyer may sometimes bid too little or too much, these mistakes will essentially cancel each other out. Over the course of many auctions, this yields the seller revenue that is similar to a second-price auction.




Production results from the AppNexus Exchange

The results above provide intuition, but not real evidence that algorithmically-powered bidding removes the difference between first- and second-price auctions for buyers. Fortunately, we have recently released BPO to the AppNexus platform and as an exchange, we have data from sellers that use our first-price auction, second-price auction, and our legacy auction, which includes mixed-price auctions. Thus, we can make comparisons across auction types, over time, at scale.

The main metric we will use is Average Bid Reduction (ABR). In first-price and mixed-price auctions, ABR is the percentage bid reduction from the buyer's original bid valuation to the bid that we submit in the auction. In second-price auctions, ABR is the percentage bid reduction from the valuation to the price paid. Revenue equivalence implies that the ABR should be about the same across auction types.

The following plot shows ABR for first- and second-price auctions on AppNexus platform for buyers using BPO. AppNexus has been rolling out BPO to the platform supply over the last nine months, and each dotted line represents a new portion of inventory that was enabled.



Observations

- We saw a convergence of ABR for first- and second-price auctions as BPO coverage increased to 100% of first-price inventory.
- The difference in ABR between first- and second-price auctions is less than ~5%.

Understanding deviation from revenue equivalence theory in the real world

In the real world, the Revenue Equivalence Theorem does not hold perfectly, as demonstrated by the difference in ABR for first- and second-price auction. Possible reasons include:

- **Difference in inventory in first- and second-price inventory.** Sellers may have different reserve prices, or buyers' bid valuations could vary across inventory.
- **Difference in buyer goals across inventory.** It's possible that not all buyers are optimizing towards revenue maximization.
- **BPO is still improving.** We're still in the early days for bid shading solutions, especially when it comes to estimating the second bid price. The buyer's estimate of the second bid will get better with each auction in which they use BPO.
- **Selection Bias.** Some buyers may choose not to buy first-price inventory at all. Or, they may lower their valuations even when BPO is in effect, which would affect the calculation of our valuation.

If we take all of these factors into account, there is very little evidence of a meaningful difference between first- and second-price auctions in the long run for buyers using BPO.




WHAT ABOUT BUYERS NOT USING BPO?

Today, buyers pay *almost* the same amount in first-, second-, and mixed-price auctions when using BPO. But what about external demand partners on AppNexus exchange who aren't using BPO? In addition to APP buyers, 100+ external demand partners (DSPs) buy media on the AppNexus Exchange. The available evidence suggests that many large buyers have yet to adjust their bidding behavior for first-price auctions. From a publisher's perspective, this means that first- and second-price auctions are not revenue equivalent.

Keep in mind that there's no way for us to calculate ABR for our external demand partners because we do not know their bid valuations. The best we can do is compare total seller revenue before and after our auction logic changes, while also controlling for seasonality. When we do that, the data shows that when publishers first switched over to first-price auctions, they tended to get a large increase in spend from buyers not using BPO. Over the following month or two, they lose about 40% of that spend bump.

For example, the chart below shows the year-over-year percent change in spend from non-BPO buyers for a sample of publishers who switched from second- to first-price auctions. After the switch to first-price, we see a massive jump in spend. Even after it falls over the subsequent months, there is still a substantial boost in seller revenue.

Non-BPO Buyer Spend on Publishers that Switched to 1st Price Auctions



We don't know how long this revenue bump will last, but as of the release of this white paper in July of 2018, it is still holding. We expect that as they see more and more first-price auctions, major buyers will invest in bidding algorithms that adjust bidding strategy based on auction type.



CONCLUSION AND RECOMMENDATIONS

Buyers

In the long run, auction type is irrelevant for buyers using a bid price reduction algorithm like BPO to adjust to first-price auctions; the buyer will ultimately pay the same price in first- and second-price auctions given enough iterations. However, external demand partners not using BPO tend to overpay in first-price auctions, even though we send those buyers the auction type (BidRequest.at) in the [OpenRTB request](#).

Thus, buyers should not overthink auction type. Instead, we recommend they focus on finding the right technology partners, supply, and users to meet their goals. Brands working with DSPs who don't optimize for auction type using BPO or a similar mechanism should consider alternative vendors, as they may be throwing money away.

Sellers

In contrast, auction type is currently quite relevant for sellers.

Publishers with header integrations should virtually always use first-price auctions. Why? Because the bid passed to the header is what the buyer pays and is what determines who wins, meaning that header bidding auctions functionally use first-price auction logic. So buyers and sellers alike have a clear reason to make their header bidding auctions first-price.

Publishers without header integrations may also use first-price auctions (and make more money) if they have a large share of spend coming from external buyers not using BPO or a similar bid reduction algorithm. On the other hand, publishers with a large share of spend from buyers using a bid reduction algorithm may do better employing second-price auctions with floor optimization. We welcome publishers to test and learn what works for them but caution them that the advantages of first-price auctions may not last forever.

END NOTES

1. The term mixed price auction comes from the book, [Auction Theory](#), by Vijay Krishna. There mixed price auction refers to auctions that are a weighted average of first and second price auctions. Auctions with soft floors work a bit differently, but "mixed" is still a good description of them.





AppNexus
A Xandr Company