Product Vision for Auction Floors

Author: Rahul Kooverjee
Contributors: Ali RadhaEric CheungChristopher Tignor ssg-aq@google.com

Self Link: go/auction-floor-vision
Last updated: 10 March 2021

Goal of this doc
This doc aims to tie together all of the ideas and work we've done on auction floors into a cohesive narrative on what we envision as the future of auction floors.

Vision
Today, publishers use auction floors for 3 reasons:

1. Maximizing auction revenue (equivalent to minimizing bid shading)
2. Representing business constraints (e.g. channel conflict, ad quality, $ opportunity cost)
3. Cross platform yield management (e.g. mediation)

In the future, we envision that:

- We can completely automate use case #1, as we can likely we can do a better job than publishers do today (though some pubs may still choose to maintain control)
- Use case #3 becomes redundant as mApp monetization shifts from mediation to bidding and we enforce against non-compliant waterfalls on web
- Publishers continue to use floors for a subset of use case #2, but we help them do so in a more efficient and optimal way

How we get there
Automated floors

- Build a performant model for automatically setting revenue maximizing first-price floors
- Transition publishers away from using floors for yield management by offering OPR as a way to evaluate our optimization

Business constraints

- Build more focussed solutions to specific constraints
- Drive targetCPM adoption

How do publishers use floors?

Use Cases

Broadly speaking, our publishers use floors for three use cases:

1. Maximizing auction revenue
    - Metric of interest: ad request eCPM (revenue / ad requests)
    - "Maximize revenue ⇐minimize bid shading"
2. Business constraints
    - Metric of interest: eCPM (revenue / matched requests)

[redacted]

PTX1021
1:23-cv-00108

- "The opportunity cost of serving an ad is $X"
- Examples: channel conflict, ad quality, ad density, non-ad monetization, value of house ads

3. Cross platform yield management (via Waterfall yield management
   - Metric of interest: total revenue across platforms
   - "I'm confident another platform will pay $Y, so I don't want to sell in your auction for less than $Y"
   - Examples: mediation on app, non-compliant waterfalls on web

[redacted]

Note that use cases #1 and #3 are mutually exclusive. However, a publisher can be using floors for both use cases #1 and #2 or #2 and #3 simultaneously.

#1 - Maximizing auction revenue

In a 1st price auction, a bidder is incentivized to shade their bids as much as possible since they pay what they bid. To counteract this, a publisher can set auction floors to force bidders to bid more if they'd like to win the impression. Publishers using floors for this use case are worried that without a floor bidders may aggressively shade their bids over time, especially in sparse auctions, thereby devaluing their inventory.

Note that there is an equivalence here between maximizing auction revenue and minimizing bid shading.

The primary metric of interest for this use case is ad request eCPM (revenue / ad requests) as a publisher aims to maximize the revenue they gain from every query that enters our auction.

#2 - Business constraints

Publishers often have business constraints or objectives that can (perhaps imprecisely) be translated into a price below which the net value of serving an ad via the auction is negative. These include:

- Channel conflict - where publishers do not want to sell inventory to buyers cheaper than they can buy it through a direct deal.
- Ad quality - publishers believe (perhaps incorrectly?) that low CPM ads are low quality and will hurt their brand.
- User experience / ad density - fewer ads results in a better user experience and so it might be worth sacrificing ads below a certain CPM if that means users are more engaged on the site. This is especially true for publishers who monetize through both ads and "transactions" on their sites (e.g. autotrader.ca who also makes money via listings) and are concerned that ads might degrade the UX such that the likelihood of a "transaction" declines.
- Value of a house ad - some house ads can drive revenue for publishers (e.g. subscriptions) and so it may be more valuable to serve a house ad vs. a low CPM auction ad.

[redacted]

The primary metric of interest for this use case is eCPM (revenue / matched requests) as a publisher has some minimum eCPM that represents the opportunity cost of serving an ad via the auction.

#3 - Cross platform yield management (via waterfalls)

It is common for publishers, especially those who are more sophisticated, to work with multiple SSPs and/or ad networks. These publishers aim to maximize their yield across platforms and use floors as one mechanism to do so when calling the different SSPs / networks in a waterfall.

Mediation on apps is the predominant case. As a simple example, a publisher might know, based on historical data, that network X generally pays a $1 CPM and so it doesn't make sense to accept less than a $1 CPM from our auction because they're confident they can achieve a better outcome from network X. So they place us in a mediation chain above network X with a $1 floor.

Other waterfall structures where additional ad calls are made after calling our auction are against policy. But we know that publishers have tried to find workarounds such as trafficking AdX or 3P demand as a house line item, which is partly why we instituted our policy on line item use. If the 3P demand was trafficked via a remnant line (as it should be per policy), then there would not be a need to use floors for cross-platform yield management as our unified auction allows remnant line items to compete with OA / OB demand on price and in real-time.

The primary metric of interest for this use case is total publisher revenue across platforms.

Empirical Observations

We've done extensive analysis on how publishers use UPRs, which can be found at go/upr-analysis. Here are a few key takeaways:

- Networks who use UPRs constitute >96% of AdX revenue and impressions
- The vast majority of networks using pricing rules have very few (< 25), indicating that they don't segment their inventory very granularly.
- The vast majority of publishers (~70%) using auction floors have them apply to all of their traffic.
- Publishers are not updating their floor prices very often and there is no real discernable pattern to the updates they perform except for a spike during COVID. So they likely aren't adjusting floor prices to capture the changes auction CPMs as market conditions change(e.g. Q4 holiday spike)
- Most floor prices (~80%) are between $0 and $2.00, with a small amount of higher floors

[redacted]

Vision

Overview

We envision a future where:

- We completely automate revenue maximization / bid shading protection, meaning publishers do not need to manage floors for yield maximization (though they still can if they choose to).
- Floors are no longer relevant for cross platform yield management as mobile app monetization shifts away from mediation to bidding and we launch enforcement against non-compliant waterfalls.
- Publishers continue to use floors to represent some (but not all) business constraints but in a more efficient and optimal way.

Put another way, publishers will no longer need to use floors for use cases #1 and #3 and while they will continue to do so for some subset of use case #2, we will help them do so more efficiently and optimally.

Motivation

The motivation behind this vision is threefold:

1. We can do a better job setting revenue maximizing floor prices for publishers
    - Setting revenue maximizing floor prices requires a ton of effort - segmenting inventory into similarly priced slices, understanding complex auction dynamics, modeling bidder behavior, and updating floors frequently as market conditions change.
    - We know from our empirical observations that most publishers are not granularly segmenting their inventory and are not updating floor prices very often, meaning they're likely leaving money on the table.
    - We're able to segment inventory more granularly than publishers do today, do the modeling, experimentation, and analysis necessary to determine a good floor price to minimize bid-shading, and automatically update floor prices as market conditions change.
    - We also have access to much more data than a single publisher does and therefore can develop more holistic pricing models.
2. Cross-platform yield management will no longer be a relevant use case
    - With the app ecosystem shifting away from mediation to bidding, using floors to represent an opportunity cost of filling via a given demand source vs. another becomes redundant as all demand sources compete on price in real time.
    - With our line item use policy and pending enforcement in Q3, publishers will no longer be able to set up non-compliant waterfalls.
3. Publishers will continue to have business constraints, but right now we offer a very blunt and at times hard to use tool to manage them
    - Our publishers are running complex businesses and will continue to have constraints on how they monetize via the auction.
    - But what they care about are those *constraints* not the floor price itself. Floor prices are just a very blunt means to an end which opens opportunities for sub-optimal use (e.g. using floors for blocking).
    - We can't help publishers correct the sub-optimal use because we don't have insight into the outcome they're trying to achieve; all we see is a single price.

The vision laid out here is similar to that in go/indirect-optimization-platform.

How do we get there?

The changes needed that remove the need for floors for cross-platform yield management are already underway (i.e. industry shift to bidding + scaling OB on our platform, line item policy enforcement).

So the rest of the document will focus on (1) how we get to a world with completely automated floors for revenue maximization and (2) how we could help publishers represent their business constraints in a more efficient and optimal way.

Automated floors

To transition publishers away from using floors for revenue maximization, we need to:

1. Build model that determines revenue maximizing floors
2. Get publishers to adopt our automated solution

Modeling

There are two questions a model needs to answer: (1) how should inventory be segmented and (2) how to set floors on those segments. RPO and AFLAC are attempts at building a model to set revenue maximizing floors based on different approaches to both segmentation and floor setting.

We will need to reconcile the different assumptions and approaches of each model and land on a single, performant model. Our proposed approach is iterative and consists of running buyer-reaction experiments to compare different approaches to setting floors. go/setfloor and go/opr-rpo-expt discusses this work further.

The requirements for our floor setting model are that it:

- Outperforms how publisher set floors today
- Does not make assumptions based on a publishers current pricing set up (segmenting or existing floors)

Adoption of our automation

With "optimized pricing", the external name for RPO as a product, we have the flexibility to set higher floors than what a publisher set in order to maximize revenue / minimize bid shading. As was the case with 2nd price RPO, 1st price RPO will be always on.

But that doesn't fully solve the problem. Publishers currently use floors for yield maximization today and will need to be convinced that our optimization solves the problem just as well. Plus, we know from our AFLAC experiments that most of the publishers are setting their floors too high, likely because they fear aggressive bid shading. Just telling publishers "hey, we have RPO now so you don't have to worry about using floors for yield optimization" is unlikely to convince them to transition to the ideal end state.

We need to give publishers a way to experiment with our optimization, get accurate reporting on the uplift, and evaluate whether we can indeed deliver on our promises.

That's where OPR comes in. It would essentially be a wrapper around RPO that removes the constraint of respecting existing floors. That means that OPR needs to:

1. Allow publishers to designate a portion of their inventory to try our optimization on
2. Give publishers reporting on the optimization performance and show uplift relative to existing pricing

3. Let publishers continue to be able to set floors that represent business constraints, which we won't go below

Requirement #1 is easily satisfied by using our existing targeting picker on pricing rules.

Requirement #2 is challenging as performance depends on buyer-reactions, which happen over time (~9 days according to our experiments) and aren't able to be captured in a traditional A/B experiment. That said, we do give publishers the ability to run traditional A/B experiments with different floor prices today. So perhaps that might be sufficient? We are also thinking about whether buyer reaction experiments on a per-publisher basis are possible, as they'd be useful for this and auto-refresh (although it is a *very* hard problem).

Requirement #3 is also challenging. Today, publishers use floors for both yield optimization and business constraints. We need to be able to "override" the former type of floor while respecting the latter. The problem is, we don't which use case a given UPR falls under. Some ideas on how we might satisfy it:

- [not recommended] ask publishers to turn off their non-business constraint floors.
    - *The entire reason we need OPR is that we believe this won't work. It also doesn't allow publishers to experiment, as we'd lose the counterfactual (i.e. their own floors)*
- [not recommended] add an additional input to OPR rules that lets publishers (optionally) set a price the rule would not go below (see go/constrained-opr for more).
    - *This could get very confusing. We'd essentially be asking publishers to recreate their entire UPR set-up minus the rules used for yield management. Put another way, UPRs are already the mechanism by which publishers specify business constraints so why build another?*
- [preferred] let publishers demarcate which existing (or newly created) UPRs represent business constraints. TBD on the UX, but perhaps something like a checkbox on existing UPRs or something targeting picker-esque that lives in settings and lets pubs select which of their UPRs to exempt from optimization
    - *This is the preferred option as leverages the fact publishers do already have rules set up for their business constraints. The key will be nailing the UX. Over 75% of publishers have less than 25 sub-rules, so this would not be too onerous an option.*

Business constraints

Our approach to helping publishers express their business constraints in a more efficient and optimal way has two components

1. Offering more focussed solutions to specific constraints
2. Driving tCPM adoption

Focussed solutions to specific constraints

Rather than think of "business constraints" in abstract, we should focus on the specific goals a publisher has and consider how we could best help them achieve those goals. Some of those goals are indeed best addressed via an auction floor, while others might be better addressed with a different tool (e.g. protections).

Channel conflict

Today, we do offer a solution to channel conflict in the form of advertiser specific targeting for UPRs. Some publishers use it. But many publishers don't, electing to set a blanket floor price across all advertisers insteads. Some reasons we've heard for this are that:

- Advertiser targeting doesn't apply to remnant (as we don't know the advertiser for 3P demand), so it doesn't offer complete protection against channel conflict across *all* demand sources.
- Setting up and maintaining rules for every one of the advertisers they have deals with is too onerous.

Advertiser blocks in protections are also a solution to preventing channel conflict. But they're a very extreme solution as they block the advertiser entirely when really what the publisher cares about is ensuring the advertiser pays at least as much as they would in a direct deal.

The net result is that publishers end up leaving money on the table as they filter non-channel conflict demand with their overly conservative setups. It's also bad for Google as it limits our inventory access.

Ideas (some taken from go/channel-conflict-prd by cgaonr@) to helping publishers better manage channel conflict include:

- Automatically detecting advertisers that publishers have (or have had) deals with and either (a) automatically setting floors or (b) making it easier to target these advertisers
- Making it easier to set up advertiser specific floors by offering the ability to do it by category or vertical (e.g. "all automotive advertisers")
- Creating more granular targeting capabilities to allow publishers target inventory / formats they don't sell directly (e.g. skippable formats on video)
- Offering discount rules more broadly, to give publishers to ability to give specific advertisers (e.g. smaller advertisers) lower floors as those smaller advertisers aren't candidates for direct deals

Ad quality

Our stance has been that protections, not floors, should be used to control ad quality. Indeed, we literally tell publishers that UPRs don't provide blocking functionality. But protections today aren't robust enough (e.g. they don't apply to open bidding) and so publishers return to the only tool they have - floors. So in order to move publishers away from floors to protections, we need to enhance protections to cover all demand sources, which would require changes to our creative scanning and review process.

- TODO: link to karenscruggs@ doc with her proposal to widen the applicability of protections once it's ready

Moreover, we're not actually confident that there's a strong correlation between ad quality and CPM. This is an area we should study further to understand if floors are even a good solution to ad quality. If they're not, then perhaps we can help better educate our publishers - e.g. by publishing a whitepaper.

Ad density

If publishers care about ad density, then a better solution than floors would allow them to specify a target ad density and we'd do the translation on our end from that target into a floor value that gives them that density. This could be similar to the recently deprecated Ad Balance feature that AdSense offered. That said,

we should research how prevalent this constraint is and whether ad density is the objective or whether it's a proxy for some other, more fundamental goal (i.e. impact to other monetization sources).

Ads impact to other monetization / value of house ads

Floors are likely the right tool to represent the monetary opportunity cost of filing an ad via the auction. But that doesn't mean we can't make it easier for publishers. For example, we might help publishers better understand the value of serving a house ad and automatically setting a floor equal to that value by converting some CPA value to a CPM floor.

Behind all of this is the idea that we need to develop a better understanding of the business goals of our publishers. All we know is they want an $X CPM, we don't know *why*. A more speculative idea is to have publishers tell us directly in the product what they're hoping to achieve when setting pricing floors. Armed with this information, we'd be able to better optimize while respecting their business goals.

Driving TargetCPM adoption

TargetCPM is an alternative way for publishers to represent their business constraints. Rather than an absolute constraint ("all impressions must pay $X"), the constraint becomes an average ("on average, queries should pay $X") allowing us to more optimally set floors to extract incremental revenue while still respecting a publishers business goals. Both Ad Manager and GBO have an OKR to increase the % of queries eligible for tCPM from ~47% today to 55% by EOY 2021.

Opportunities at the sub-rule level (ariane/4051553) and in-line opportunities (ariane/4068900) both launching in Q1 will help drive targetCPM adoption.

In addition, we only satisfy the average CPM constraint ~85% of the time which has led to escalations and publishers disabling targetCPM. It also means that it's hard to pitch targetCPM as an alternative way to represent their business constraints. So we're investigating how we might be able to improve the hit rate. go/tcpm-h1-2021 and go/tcpm-report discuss ideas on how we might do this, but there's a clear tradeoff between the uplift from tCPM and it's hit rate. So we'll need to decide whether we're comfortable trading off significant uplift in exchange for a higher hit rate.

Next steps

The work on automated floors is a higher priority than most of the business constraint work, except for perhaps channel conflict (as solving channel conflict could unlock significant inventory in our auction).. So in the near term, we will focus our efforts on automated floors (RPO + OPR) and addressing channel conflict.

Appendix

Open questions on OPR

- If OPR is a stepping stone to "full fledged RPO", do we need to wait for RPO to launch before launching OPR?

- *I think the answer to this is no presuming we have a way of implementing OPR that satisfies the model requirements listed above. There are benefits to not waiting for RPO - e.g. getting publishers comfortable with our optimizations earlier. Plus, if we expect significant revenue upside, there's high opportunity cost of waiting for RPO*

- How do we let publishers run experiments / evaluate the performance of our floor setting models? Is it even a solvable problem?
- What should the UX be for exempting existing rules from OPR?

[redacted]

Relevant links

- Indirect optimization platform by mhop@ and ctignor@
  - Lots of the same ideas as those espoused here
- Optimizing Floors for Yield | Designing for an Ideal State by aradha@
- Strat review notes
  -

Edit history

Date	Edits
9-23 February 2021	Document created
4 March 2021	Minor edits based on feedback
10 March 2021	Minor edits / polishing in preparation for review