Cascaded Bid Floors for Heavy-Tailed Mobile Ad Auctions: A Deployed Production System

Rowan Swiers
Metica
London, United Kingdom
rowan@metica.com

Yash Jain
Metica
London, United Kingdom
yash@metica.com

Ajinkya Jawale
Metica
London, United Kingdom
ajinkya@metica.com

Andrew Maher
Metica
London, United Kingdom
andrew@metica.com

Abstract

Bid-floor optimization for publishers in mobile ad auctions is difficult: bids are heavy-tailed, user-level signals are sparse, and publishers observe only partial auction feedback. Mediation infrastructure also supports auction retry paths, where an auction that fails to fill at one floor can be attempted again at a lower floor. These retry paths complicate floor setting, but also create an opportunity: publishers can target high-value demand while preserving a low-floor safety path.

We study this problem through FloorGym, a publisher-side simulator calibrated on production data. We show that single-level bid floors often perform poorly under heavy-tailed first-price bids: floors high enough to extract price also lose enough fill to cancel out the gains. We show using simulations that a cascaded bid-floor policy can mitigate this trade-off by attempting an aggressive learned primary floor first and, if it does not fill, reissuing the auction opportunity at a near-zero safety floor.

We evaluate this cascaded bid-floor design in production through a seven-day A/B test on 44 mobile games from the portfolio of Game District, a large mobile-game publisher, covering approximately 5 million daily active users. Treatment increased average revenue per daily active user (ARPDAU) by +13.2% relative to incumbent floor strategies. The effect was primarily price-led: revenue per impression increased by +9.3%, while impressions per daily active user (ImpDAU) increased by +3.7%.

CCS Concepts

• **Information systems → Computational advertising**.

Keywords

bid-floor optimization, reserve prices, first-price auctions, mobile ad mediation, heavy-tailed distributions, online advertising

1 Introduction

In-game advertising has become a large advertising channel. Free-to-play titles dominate this market, monetizing a large share of their traffic through in-app advertising. An impression is a single opportunity to show an ad to a user. These impressions are sold through mediation platforms such as AppLovin MAX [2] and LevelPlay [24] which run per-impression auctions across many demand sources. For game publishers, the main price-control lever is the bid floor, hereafter also called the floor: a minimum acceptable price per impression, below which an auction attempt fails. Bid floors create a price/fill-rate trade-off. Raising the floor can increase revenue when bidders have high willingness-to-pay, but can also cause the auction to fail when no bidder clears the bid floor.

Mobile mediation auctions are typically first-price: the highest bidder wins and pays its own bid. In such auctions, bidders generally shade, meaning they bid below their private value; the amount of shading depends on demand, pacing, competition, and auction context [21, 25]. Publishers observe only partial auction telemetry: winning bids, no-bid events, clearing networks, and user/context features, but not the full vector of losing bids or bidder values. At the same time, bids are heavy-tailed: most impressions clear at low prices, while a small fraction command very high prices. A single learned floor is therefore fragile. Set too low, it leaves surplus in the upper tail; set too high, and fill loss in the body of the distribution can erase the price gains.

We study this problem using FloorGym [1], a publisher-side simulator for bid-floor optimization. FloorGym is calibrated to representative production games and models the problem faced by a publisher setting floors in mobile mediation auctions. We use it to compare single-level and high/low bid-floor policies under repeated auctions with bidder shading and heavy-tails.

In our simulator, we find that a *cascaded bid-floor policy* is effective at increasing revenue: an auction is first offered at a primary floor and, conditional on no fill, retried at a lower fallback floor. This strategy lets the publisher aim into the upper tail of the bid distribution without converting most failed primary attempts into lost impressions.

For the production experiment in this paper, we deploy the cascaded bid-floor system on 44 mobile games from a customer Game District covering approximately 5 million daily active users. Treatment increased ARPDAU by +13.2% relative to the games' incumbent bid-floor strategies. The effect is primarily price-led: CPM, defined as publisher revenue per thousand impressions, increased by +9.3%, while impressions per user increased by +3.7%. This pattern is consistent with the intended operating regime: higher prices with little fill loss.

Contributions.

[1] https://github.com/meticalabs/floor-gym

- We formulate the publisher-side bid-floor optimization problem for heavy-tailed first-price mobile ad auctions with bid shading and retry paths.
- We introduce FloorGym, a calibrated publisher-side simulator for evaluating bid-floor policies under bidder shading and heavy-tailed CPMs, and show empirically that single-level floors fail in this regime, while high/low redraw policies can extract upper-tail value while preserving fill.
- We report production evidence from an A/B test conducted on 44 mobile games from Game District's portfolio, covering approximately 5 million daily active users.

2 Related Work

Our work sits at the intersection of reserve-price learning, first-price bidder behavior, simulation and online experimentation.

Waterfall mediation. Mobile mediation has historically used waterfalls: an impression is offered to ad networks in sequence, with price cutoffs, until one fills. Prior work studies waterfall order and floor optimization [10], online waterfall learning [18], abort decisions that truncate unprofitable descents [23] and reserve-price failure risk [14]. Our cascade differs from this literature: it repeats the auction opportunity at different bid-floor levels, rather than routing the impression through a sequence of networks with network-specific price cutoffs.

First-price auctions and bidder response. Display and mediation auctions are now commonly first-price [4, 7]. In first-price auctions, bidders shade relative to value [9, 21, 25], and autobidding and pacing systems create an adaptive bid landscape for publishers setting floors [3, 6, 11, 13, 17]. FloorGym uses shading and non-adaptive bidders to examine the effect of publisher-side floor setting; the production system operates in a richer environment with heterogeneous adaptive bidders.

Reserve prices in online auctions. Revenue-optimal reserves are classical [19, 22]. Much subsequent reserve-learning work was developed for second-price advertising auctions, where the reserve problem is often framed as bid prediction, contextual pricing, or personalised reserve selection. In second price auctions the winning bid pays the second highest price. This includes repeated-auction reserve learning [5, 16], and contextual or personalised reserves [12, 20]. A key difference between first and second price auctions is that in first price auctions bidders shade their true value whereas in second price auction shading is not always necessary.

Reserve optimization in first-price auctions. First-price floor optimization has been studied directly. The shared difficulty is that floors censor losing bids, bidder strategies respond to floor changes, and the resulting revenue objective is nonconvex and only partially observed. Proposed approaches include stochastic gradient ascent on logged revenue with a demand/bid-response decomposition for variance reduction [8], parametric per-buyer-type bid CDFs plugged into a first-price revenue objective with explicit shading correction [1], and a multi-task neural network that uses both filled and unfilled impressions to predict a conservative lower bound on the top bid, which is then used as the reserve [15]. These approaches set a single floor per impression. Our deployed system is a cascaded bid-floor policy: a learned primary floor followed by a near-zero safety floor on a retry attempt so there can be up to two auctions rather than only one. The benefit is the cascade can target upper-tail value while preserving fallback fill, whereas a single-level floor often cannot.

3 FloorGym: Simulator and Mechanism

FloorGym is an open-source publisher-side simulator for bid-floor optimization in mobile ad mediation. It simulates auctions against a parametric bidder population calibrated to production mediation data. Below, we describe the bidder model, auction mechanics, policies, and metrics; results follow in §4.

3.1 Bidder model and auction mechanics

Each simulated impression is indexed by i. At the start of the impression we draw a context

$$x_i = (c_i, h_i, a_i),$$

where c_i is the country tier, h_i is the hour bucket, and a_i is the ad format. The context is held fixed across all auction attempts for the same impression.

For each auction attempt, the impression is exposed to a population of K contextual buyers. Let $\ell \in \{1, 2\}$ index the auction level or attempt. Buyer j participates in attempt ℓ independently with probability p. Conditional on participation, bidder j has private valuation

$$v_{i\ell j} = \exp\big(\mu_0 + \tau_{c_i} + \eta_{h_i} + \zeta_{a_i} + \sigma\varepsilon_{i\ell j}\big), \qquad \varepsilon_{i\ell j} \sim \mathcal{N}(0, 1). \quad (1)$$

Here μ_0 sets the global log-CPM scale, τ_{c_i} is the country-tier effect, η_{h_i} is the hour-of-day effect, and ζ_{a_i} is the ad-format effect. The parameter σ controls unobserved bidder variation: larger σ produces more dispersion across bidders within the same observed context, and hence a heavier upper tail of valuations.

We model bidder behavior in two steps. First, absent a binding floor, bidders use a common linear-shading rule: bidder j with private value $v_{i\ell j}$ submits the shaded bid $sv_{i\ell j}$, where $s \in (0, 1]$ is the shading factor. Linear shading is a common reduced-form approximation in first-price auctions: it is easy to interpret, captures the practical fact that bidders often bid below value, and gives a single parameter controlling the gap between values and submitted bids. It is also easier to reason about analytically than a fully heterogeneous or dynamically learned shading model.

Second, we specify how bidders respond when the seller imposes a floor f. We use a perfect-response model. To simplify notation, Eq. (2) suppresses the impression, attempt, and bidder indices. Under this response model, the bidder keeps its shaded bid when the floor does not bind; raises its bid to the floor when doing so is sufficient to clear the floor without exceeding its value; and drops out when the floor exceeds its value:

$$b(f) = \begin{cases} sv, & \text{if } f \le sv \quad \text{(floor does not bind)}, \\ f, & \text{if } sv < f \le v \quad \text{(bid lifts to floor)}, \\ 0, & \text{if } f > v \quad \text{(bidder drops out)}. \end{cases} \quad (2)$$

Both parts of this model are simplifications. The common linear-shading rule abstracts from bidder-specific shading. The perfect-response rule is also idealised: real demand sources may respond imperfectly to reserves, may update only after observing repeated floor changes, and may sometimes bid above the minimum required to clear the floor. We therefore use Eq. (2) as a mechanism model and report sensitivity to the shading factor and prediction noise rather than relying on a single calibrated value of s.

An auction attempt fills if at least one participating bidder submits a bid satisfying $b_{i\ell j}(f) \geq f$. Conditional on fill, the bidder with the highest submitted bid wins and pays that bid.

We evaluate three auction regimes.

- **Single-level floor.** The seller runs one auction attempt at floor f. If no bidder clears the floor, the impression is unfilled.
- **Repeated identical floor.** The seller runs up to two auction attempts at the same floor f. If the first attempt does not fill, the second attempt uses the same impression context x_i but a fresh draw of bidder participation and valuation noise.
- **High/low redraw cascade.** The seller chooses an ordered pair (f_1, f_2) with $f_1 > f_2$. The primary auction runs first at floor f_1. If it does not fill, the safety auction runs at floor f_2 using the same impression context x_i but a fresh draw of bidder participation and valuation noise.

The simulator parameters are fitted from 21.9M production impressions. Experiments use $K = 8$ buyers and bidder participation probability $p = 0.53$. The fitted valuation model has $\mu_0 = 0.63$, $\sigma = 1.19$, tier effects $(+2.49, -0.48, -0.18)$, and small hour and format effects $(|\bar{\eta}|, |\bar{\zeta}| \leq 0.21)$.

3.2 Policies and metrics

Baselines and learned policies. FloorGym includes baseline policies, learned policies, and simulator upper bounds. **No Floor** sets $f = 0$ and preserves fill without attempting price extraction. **Fixed** policies use constant floors independent of context.

The **full-information oracle** observes the actual drawn valuations $\{v_j\}$ for the impression and chooses the revenue-maximising floor for that realisation. It is a strict per-impression upper bound. The **CPM predictor** is a weaker, context-conditional oracle: it observes only $\mathbb{E}[\mathrm{CPM} \mid x]$ and sets the primary floor to

$$f_1 = m\,\mathbb{E}[\mathrm{CPM} \mid x],$$

where m is a multiplier tuned by grid search. To test robustness to prediction error, we also evaluate an **imperfect CPM predictor**

$$\hat{c}(x) = \mathbb{E}[\mathrm{CPM} \mid x]\,(1 + \epsilon), \qquad \epsilon \sim \mathcal{N}(0, \sigma^2_{\text{pred}}),$$

with floor $f_1 = m\hat{c}(x)$.

4 Simulator Results

4.1 Effect of bidder shading on floor strategies

Figure 1 compares three regimes across bidder shading levels: a single-level floor, a repeated same-floor policy, and a high/low cascade. Each is evaluated with a perfect CPM predictor and an imperfect CPM predictor with $\sigma_{\text{pred}} = 0.5$, corresponding to MAPE 0.39 against the true conditional mean CPM.

The single-level policy only improves on No Floor under very strong shading ($s \lesssim 0.5$), where there is a large gap between shaded bids and private values. The reason is that a floor only helps in the binding interval $(sv_j, v_j]$: below this interval the floor has no effect, and above it the bidder drops out. In our calibrated environment, floors high enough to bind on valuable impressions lose too much fill on ordinary impressions.

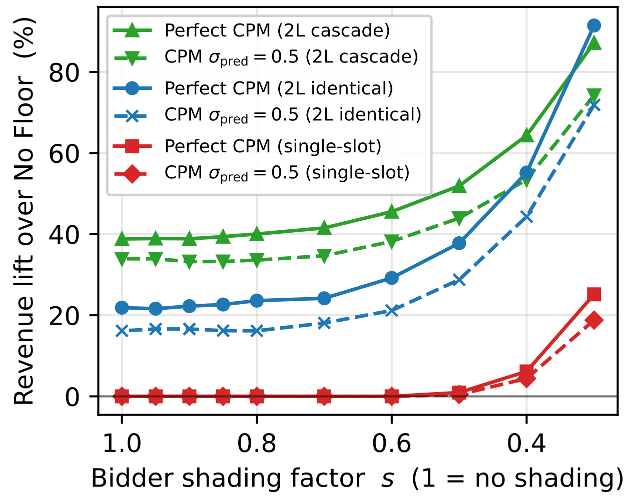

Figure 1: Revenue lift for different floor strategies as a function of bidder shading

A repeated identical floor policy gives failed impressions a second chance to fill, and therefore improves on a single attempt. However, because both attempts use the same floor, it still cannot target the upper tail aggressively without making both attempts miss. It adds redundancy, but it does not change the core price/fill trade-off.

The high/low cascade changes this trade-off by separating the two roles. The primary floor targets high-value demand, while the near-zero safety floor preserves fill when the primary attempt misses. This is why the cascade outperforms both single-level floors and repeated identical floors across a wide range of shading levels. Prediction noise reduces lift, but preserves the same ordering.

The cascade also delivers lift when there is little or no bid shading, because the two auction attempts expose the impression to different bidder draws. Thus, even when bidders do not strategically respond to the floor, the high/low structure can still create value by pairing an aggressive first attempt with a low-risk fallback.

4.2 Sensitivity to unobserved bidder variation

Figure 2 shows revenue lift as a function of the unobserved bidder variation σ in Eq. (1), with shading held fixed at $s = 0.8$.

As σ increases, the upper tail of the maximum valuation distribution becomes heavier, giving the cascade more opportunity to target high-value impressions so lift increases. The cascade consistently outperforms repeated identical floors.

Single-shot floor lift quickly shrinks towards zero as unobserved bidder variation increases. In this regime, the binding corridor becomes too narrow relative to the dispersion in bidder values, making a single floor increasingly fragile.

4.3 Predictors with high/low cascade

We report revenue per impression (RPI), high-slot fill rate (Hi-Fill), and revenue lift over the No Floor baseline. Hi-Fill is the fraction

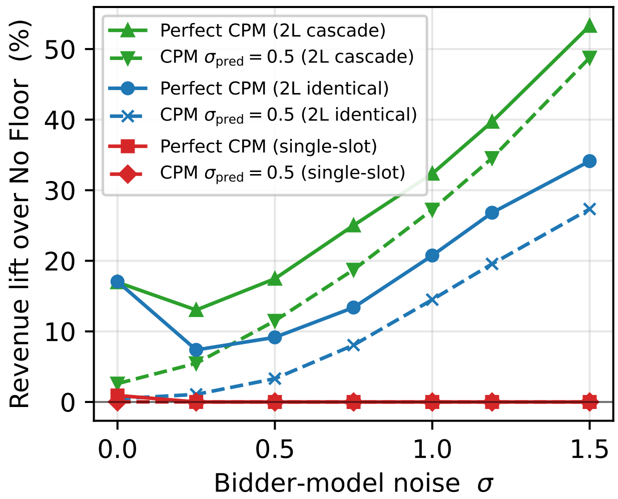

Figure 2: Revenue lift over No Floor as a function of unobserved bidder variation σ, with shading $s = 0.8$. Perfect CPM predictor ($\sigma_{\text{pred}} = 0$) and imperfect CPM predictor ($\sigma_{\text{pred}} = 0.5$) shown in three regimes: single-slot, 2L cascade with $f_2 = 0$ safety, and 2L identical (two shots at the same floor). Primary multiplier m tuned per cell.

Table 1: High/low redraw cascade benchmark with independent redraw on the calibrated first-price and shading model. Lift is over the No Floor baseline.

Strategy	RPI	Hi-Fill	Lift
Full-information oracle	48.35	0.41	**+53.19%**
Perfect CPM predictor ($m^* = 1.2$)	43.21	0.33	+36.91%
Perfect CPM predictor (misspecified $m = 1.5$)	42.96	0.25	+36.11%
Perfect CPM predictor (misspecified $m = 1.0$)	42.90	0.41	+35.92%
Fixed [$110, $0]	41.78	0.10	+32.38%
Fixed [$120, $0]	41.75	0.09	+32.28%
Fixed [$100, $0]	41.36	0.11	+31.06%
Imperfect CPM predictor, $\sigma = 0.5$ ($m^* = 1.5$)	41.34	0.32	+31.00%
No Floor	31.56	1.00	–

of impressions filled by the primary high-floor attempt; remaining filled impressions are served through the low-floor safety path.

Table 1 compares predictors for the primary floor in a high/low cascade with shading factor $s = 0.8$. m^* is the lift-optimal multiplier from a grid sweep; "misspecified m" marks a deliberately off-optimal setting. The perfect CPM predictor at m^* reaches +36.9% lift, about 0.69 of the full-information oracle's upper bound, and beats the best single fixed primary cascade floor ($110, +32.4%) by about 4.5 percentage points.

The cascade is forgiving of misspecified multipliers. Moving the predictor off its optimum to $m = 1.0$ or $m = 1.5$ costs only about 1 percentage point of lift, and the $\sigma = 0.5$ imperfect predictor still delivers +31% lift despite substantial multiplicative error in every floor. The safety floor absorbs the cost of bad primary picks.

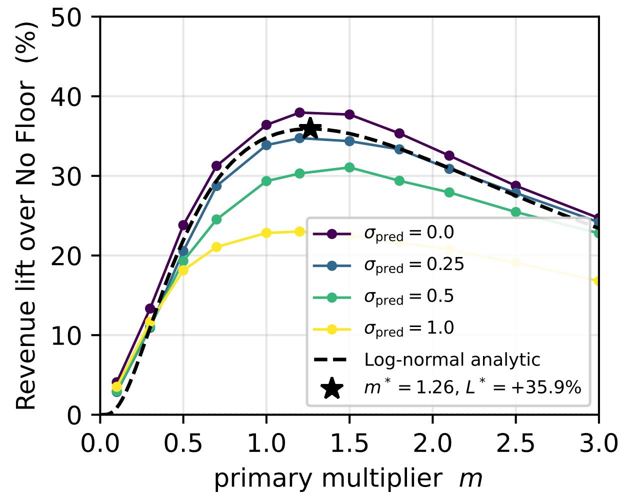

Figure 3: Redraw cascade (2L $[m, 0]$) revenue lift over No Floor on the calibrated bidder model, by primary multiplier m and added prediction error σ_{pred}. The dashed line is the log-normal closed form $L(f)$ (Eq. 3) with $s = 0.8$ and residual-variation fit $\mu = 1.66$, $\sigma = 0.90$.

4.4 Choosing the optimal multiplier for a CPM predictor

To further investigate sensitivity to prediction error, we start from a perfect predictor of CPM and plot the revenue uplift as the primary floor multiplier changes. Figure 3 reports the resulting revenue lift as a function of m and prediction-error level σ.

The cascade remains profitable even with substantial prediction error. As σ increases, the peak lift falls. The choice of multiplier matters: setting m too high is more costly than setting it too low.

4.5 Optimal primary floor multiplier

The simulated lift curve is well explained by a simple log-normal approximation. First approximate the max bidder valuation $V = \max_j v_j$ by a single log-normal with $\log V \sim \mathcal{N}(\mu, \sigma^2)$ and write $\mu_V := \mathbb{E}[V] = e^{\mu+\sigma^2/2}$, $z(x) := (\log x - \mu)/\sigma$. Under the bidder-response model of Eq. (2) with shading factor s, the cascade lift over the no-floor baseline is

$$L(f) = \Phi\big(z(f)\big) - \Phi\big(z(f/s) - \sigma\big) + \frac{f}{s\mu_V}\big[\,\Phi\big(z(f/s)\big) - \Phi\big(z(f)\big)\big], \quad (3)$$

and the optimal primary floor f^* is the unique root of the first-order condition

$$\phi(z^*)\,\frac{f^* - s\mu_V}{\sigma f^*} = \Phi(z^* + \tau) - \Phi(z^*), \quad z^* := z(f^*),\ \tau := -\log s/\sigma, \quad (4)$$

a one-dimensional root-find.

Fitting $\mu = 1.66$ and $\sigma = 0.90$ from residual bidder variation in $\log V$ in the simulator, and using the simulation shading setting $s = 0.8$, gives $f^* = 8.02$, or equivalently $m^* = f^*/(s\mu_V) = 1.26$ on the predictor-multiplier axis, with peak lift +36%. This matches the simulator's perfect-predictor cascade in Figure 3 to within ~1 percentage point; residual gap arises from the single-log-normal approximation of the max-of-K distribution.

This closed-form approximation is practically useful because it gives a principled default value for the production multiplier m, reducing the need for expensive per-game multiplier sweeps.

5 Production System

The simulator results motivate two production design principles: avoid single floor policies and separate price extraction from fill protection with a high/low cascade. We now describe how these principles become the deployed system.

5.1 Strategy portfolio

Each production strategy proposes the high floor in a two-level high/low policy. Strategy selection is performed per (game, segment), depending on the amount of available data, game type, and ad format. Depending on the mediation integration, this may be implemented as a sequential retry or through parallel auctions.

The strategy portfolio contains two families. The first is a *predictive CPM model*: a model trained on logged context and revenue tuples predicts the expected CPM $\hat{c}(x)$ for an auction opportunity, and the primary floor is set to $m\hat{c}(x)$ for some multiplier m. The second is an *off-policy reward model*: a model trained on logged context, bid floor, and revenue tuples predicts the expected revenue of candidate floors, and chooses the floor with the highest predicted revenue. This second approach is useful when sufficient exploration data is available; however typically the CPM predictor generates more lift.

Strategies are not a single per-game choice: they can be deployed at finer granularity — different segments of users within the same game (e.g., country, device tier, user-value cohort) can run different strategies. Selection is driven by A/B testing lift on a rolling window per (game, segment).

5.2 Architecture and Deployment

The service is deployed through a lightweight SDK integrated into each mobile game. The SDK observes request-time signals, including coarse geographic context, device type, ad format, and session context, and sends them to our servers.

The servers enrich these signals with aggregated user, game, and market features, including recent engagement, recent ad revenue, historical impression behavior, and country-level price signals. The resulting feature vector is routed to the deployed floor-setting model for the active game segment.

Models are trained offline on logged auction and outcome data and updated periodically. At serving time, the inference server returns a primary bid floor, which is passed to the mediation platform and applied to the relevant ad unit or auction request.

5.3 Simulation-to-Production Gap

In production, depending on the integration, the high- and low-floor opportunities may be issued sequentially or in parallel, and some systems include additional timeout or fallback logic. These choices affect realized latency, fill dynamics, and the exact magnitude of lift, but not the central mechanism studied here: pairing an aggressive high-floor opportunity with a low-floor safety path.

Table 2: Per-format revenue-weighted lift, 2026-04-23 to 2026-04-29

Format	ARPDAU	95% CI	CPM	Imp/DAU
Portfolio	+13.2%	[+10.5, +15.2]	+9.3%	+3.7%
Interstitial	+10.5%	[+8.3, +12.3]	+8.0%	+2.5%
Rewarded	+20.5%	[+14.3, +25.2]	+13.0%	+6.9%

Some integrations also restrict floors to a discrete set of preconfigured values, rather than allowing a continuous request-time floor. We treat this as an implementation constraint rather than a change to the underlying policy class: the learned floor is mapped to the nearest feasible configured floor.

6 Production Results

We evaluated the system using within-game A/B tests. Assignment is deterministic and persistent: a stable device-level identifier is hashed to a fixed arm, so no device crosses between treatment and holdout during the window. Traffic was split approximately 80/20 treatment/holdout, covering roughly 4 and 1 million daily active users per arm. Treatment traffic used the cascaded bid-floor system, while holdout traffic continued to use each game's incumbent floor strategy; these incumbent strategies were game-specific, but were typically either no floor or repeated fixed floors.

Over the seven-day window from 2026-04-23 to 2026-04-29, we ran a within-game A/B test on 44 mobile games from Game District's portfolio. These games covered approximately 5 million daily active users. Aggregating over the full window, treatment increased ARPDAU by +13.2% relative to holdout. Computed day by day, the revenue-weighted ARPDAU lift averaged +13.4%, with a standard deviation of 0.9 percentage points and a range of +11.9% to +14.4%.

Uncertainty quantification. To quantify uncertainty we use a non parametric game-level bootstrap: at each of 10,000 replications we resample the 44 game cells with replacement and recompute the revenue-weighted mean ARPDAU lift. This treats each game as the unit of analysis, the appropriate level given that users are randomised within game and lifts are aggregated per game and makes no distributional assumption on the (heavy-tailed) per-game lift. The resulting 95% percentile interval for the portfolio ARPDAU lift is [+10.5%, +15.2%] around the point estimate of +13.2%. Per-format bootstrap intervals are reported in Table 2.

Format split. Ads are served in different formats, including rewarded and interstitial. Splitting the same production window by ad format (Table 2), interstitial inventory had +10.5% revenue lift, while rewarded inventory had +20.5% revenue lift. Interstitial accounts for the larger share of the experiment's holdout traffic and thus most of the absolute lift. The difference between formats reflects differences in auction composition, bid shading, and ad-request frequency.

Decomposing ARPDAU into price and volume. Across the portfolio the CPM lift is +9.3% and the ImpDAU lift is +3.7%. The decomposition is informative: the bulk of the lift comes from *price* the floors successfully extract surplus revenue from shaded bidders

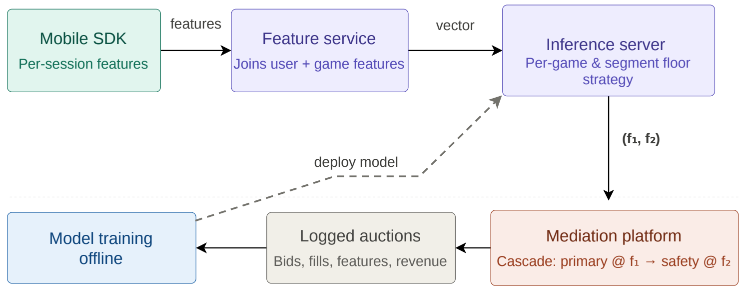

Figure 4: Production serving architecture for cascaded bid-floor selection.

while a smaller residual comes from *volume*, the safety floor prevents unfilled-impression rates from rising even with an aggressive primary floor.

7 Conclusion and Future Work

We studied publisher-side bid-floor optimization in heavy-tailed first-price mobile mediation auctions with partial feedback and retry paths. Using FloorGym, we showed that single-level floors are fragile: floors high enough to raise prices can lose enough fill to cancel the gains. A high/low cascade changes this trade-off by pairing an aggressive primary floor with a near-zero safety floor. In production, across 44 mobile games from Game District's portfolio covering approximately 5 million daily active users, treatment increased ARPDAU by +13.2% relative to incumbent floor strategies in a seven-day holdout window.

Future work could model production mediation more directly, including parallel calls, discrete preconfigured floors, latency cutoffs, overlapping bidder sets, and adaptive bidder shading.

References

[1] Miguel Alcobendas, Jonathan Ji, Hemakumar Gokulakannan, Dawit Wami, Boris Kapchits, Emilien Pouradier Duteil, Korby Satow, Maria Rosario Levy Roman, Oriol Diaz, Amado A. Diaz Jr., and Rabi Kavoori. 2023. Optimizing Floors in First Price Auctions: An Empirical Study of Yahoo Advertising. *arXiv preprint arXiv:2302.06018* (2023).

[2] AppLovin. 2026. AppLovin MAX: In-app bidding and mediation. https://www.applovin.com/max/. Accessed 2026-05-07.

[3] Santiago R. Balseiro, Kshipra Bhawalkar, Zhe Feng, Haihao Lu, Vahab Mirrokni, Balasubramanian Sivan, and Di Wang. 2024. A Field Guide for Pacing Budget and ROS Constraints. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research)*. PMLR. arXiv:2302.08530.

[4] Martin Bichler, Alok Gupta, and Matthias Oberlechner. 2024. Revenue in First- and Second-Price Display Advertising Auctions: Understanding Markets with Learning Agents. *arXiv preprint arXiv:2312.00243* (2024).

[5] Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. 2015. Regret Minimization for Reserve Prices in Second-Price Auctions. *IEEE Transactions on Information Theory* 61, 1 (2015), 549–564. doi:10.1109/TIT.2014.2365772

[6] Yuan Deng, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. 2024. Efficiency of the First-Price Auction in the Autobidding World. In *Proceedings of the 25th ACM Conference on Economics and Computation*. arXiv:2208.10650.

[7] Stylianos Despotakis, R. Ravi, and Amin Sayedi. 2021. First-Price Auctions in Online Display Advertising. *Journal of Marketing Research* 58, 5 (2021), 888–907. doi:10.1177/00222437211030201

[8] Zhe Feng, Sebastien Lahaie, Jon Schneider, and Jinchao Ye. 2021. Reserve Price Optimization for First Price Auctions in Display Advertising. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, 3230–3239.

[9] Djordje Gligorijevic, Tian Zhou, Bharatbhushan Shetty, Brendan Kitts, Shengjun Pan, Junwei Pan, and Aaron Flores. 2020. Bid Shading in The Brave New World of First-Price Auctions. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*. 2453–2460. doi:10.1145/3340531.3412689

[10] Dan Halbersberg, Matan Halevi, and Moshe Salhov. 2022. Search and Score-Based Waterfall Auction Optimization. In *Proceedings of the 16th International Conference on Learning and Intelligent Optimization*. arXiv:2201.06409.

[11] Yanjun Han, Zhengyuan Zhou, and Tsachy Weissman. 2024. Optimal No-regret Learning in Repeated First-price Auctions. *Operations Research* (2024). arXiv:2003.09795.

[12] Joey Huchette, Haihao Lu, Hossein Esfandiari, and Vahab Mirrokni. 2020. Contextual Reserve Price Optimization in Auctions via Mixed Integer Programming. In *Advances in Neural Information Processing Systems*, Vol. 33. 1287–1297.

[13] Olivier Jeunen, Sean Murphy, and Ben Allison. 2022. Learning to Bid with AuctionGym. In *Proceedings of the AdKDD 2022 Workshop, in conjunction with KDD 2022*. https://www.amazon.science/publications/learning-to-bid-with-auctiongym

[14] Achir Kalra, Chong Wang, Cristian Borcea, and Yi Chen. 2019. Reserve Price Failure Rate Prediction with Header Bidding in Display Advertising. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2819–2827. doi:10.1145/3292500.3330729

[15] Achir Kalra, Chong Wang, Cristian Borcea, and Yi Chen. 2023. Reserve Price Optimization in First-Price Auctions via Multi-Task Learning. In *Proceedings of the 23rd IEEE International Conference on Data Mining*. IEEE, 200–209.

[16] Yash Kanoria and Hamid Nazerzadeh. 2021. Incentive-Compatible Learning of Reserve Prices for Repeated Auctions. *Operations Research* 69, 2 (2021), 509–524. doi:10.1287/opre.2020.1398

[17] Rachitesh Kumar, Jon Schneider, and Balasubramanian Sivan. 2024. Strategically-Robust Learning Algorithms for Bidding in First-Price Auctions. In *Proceedings of the 25th ACM Conference on Economics and Computation*. arXiv:2402.07363.

[18] Branislav Kveton, Saied Mahdian, S. Muthukrishnan, Zheng Wen, and Yikun Xian. 2019. Waterfall Bandits: Learning to Sell Ads Online. *arXiv preprint arXiv:1904.09404* (2019).

[19] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73. doi:10.1287/moor.6.1.58

[20] Renato Paes Leme, Martin Pál, and Sergei Vassilvitskii. 2016. A Field Guide to Personalized Reserve Prices. In *Proceedings of the 25th International Conference on World Wide Web*. 1093–1102. doi:10.1145/2872427.2883071

[21] Shengjun Pan, Brendan Kitts, Tian Zhou, Hao He, Bharatbhushan Shetty, Aaron Flores, Djordje Gligorijevic, Junwei Pan, Tingyu Mao, San Gultekin, and Jianlong Zhang. 2020. Bid Shading by Win-Rate Estimation and Surplus Maximization. In *Proceedings of AdKDD 2020 (KDD Workshop on the Application of AI to Digital Advertising)*. arXiv:2009.09259.

[22] John G. Riley and William F. Samuelson. 1981. Optimal Auctions. *American Economic Review* 71, 3 (1981), 381–392.

[23] Michael Ting and Nicolas Grislain. 2018. Maximizing Net Income of the Auction Waterfall with an Abort Decision Tree. *arXiv preprint arXiv:1809.01245* (2018).

[24] Unity (ironSource). 2026. LevelPlay: Mobile ad mediation platform. https://www.is.com/levelplay/. Accessed 2026-05-07.

[25] Tian Zhou, Hao He, Shengjun Pan, Niklas Karlsson, Bharatbhushan Shetty, Brendan Kitts, Djordje Gligorijevic, San Gultekin, Tingyu Mao, Junwei Pan, Jianlong Zhang, and Aaron Flores. 2021. An Efficient Deep Distribution Network for Bid Shading in First-Price Auctions. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3996–4004. doi:10.1145/3447548.3467167