

Marginal Expected Revenue for Jointly Ranking Auction and Fixed-Price Listings in E-Commerce Sponsored Search

Greg Kocher*, Sanjana Arun

eBay, USA

Abstract

E-commerce search ranking must balance multiple objectives—relevance, user engagement, and platform revenue—when allocating impression slots to competing listings. Estimating the expected revenue component is well understood for fixed-price items, but becomes challenging when marketplace inventory includes mixed listing formats such as pure auctions and hybrid “Auction with Buy It Now” (ABIN) items, where prices evolve dynamically and the final transaction value is unknown at ranking time. Yet auction and ABIN listings account for a meaningful share of inventory and transaction volume on platforms such as eBay, and are a popular format for individual sellers and for unique items with unclear value. We extend the standard Expected Cost-per-Mille (eCPM) framework to auction and ABIN listings by deriving a **marginal eCPM (meCPM)** that captures the incremental value of showing one more impression of an item whose price is still evolving. The resulting formulation extends the familiar fixed-price eCPM—which is already inherently marginal—to auction dynamics, allowing unified ranking of fixed-price, auction, and ABIN listings under a single objective. We then describe a practical production implementation that approximates this objective, addressing cold-start challenges by bootstrapping from existing engagement models. Online A/B tests at a large e-commerce platform showed positive revenue gains and statistically significant improvements to user metrics, and the system was deployed to production.

Keywords

e-commerce, sponsored search, ad ranking, auction listings, expected revenue, eCPM

1. Introduction

E-commerce search ranking systems must allocate limited impression slots across competing listings to maximize user and platform objectives such as buyer engagement and revenue. When a buyer searches for a product, the platform aims to provide the user with the most relevant and high quality results for their search query. Meanwhile, sellers want to sell their item quickly and at a favorable price. To this end, some sellers may enroll their item in an advertising program in order to boost their item to gain extra visibility, while paying the platform a fee which depends on the type of ad program and certain details about the listing. Two common ad program types in sponsored search are Cost-per-Click (CPC) [1, 2], where the seller pays the platform a small fee each time a buyer clicks on the promoted item, and Cost-per-Acquisition (CPA) [3], where the seller pays the platform a predetermined fraction of the final selling price (the “ad rate”) only if the item sells. Finally, the platform aims to facilitate and monetize each transaction. This three-party interaction—buyer, seller, and platform—creates a multi-objective optimization problem where the search ranking system must balance relevance and user satisfaction against seller visibility and platform revenue, all within the tight latency and scale constraints of real-time search. This is typically implemented through multiple stages, including retrieval [4], followed by several rounds of filtering [5, 6], re-ranking, blending [7], de-duplicating, and layout optimization, which condense a large set of candidate items into a final ranked list on the Search Results Page (SRP).

In sponsored search systems, part of this overall allocation process is often driven by value-aware ranking models that estimate expected revenue per impression, also referred to as Expected Cost-per-Mille (eCPM) [8, 9, 10]. This scoring is well-understood in *fixed-monetization* settings, which include both CPC (where the advertiser’s bid per click is known in advance) and fixed-price CPA (where the

ECOM’26: SIGIR Workshop on eCommerce, Jul 24, 2026, Melbourne, Australia

*Corresponding author.

✉ gkocher@ebay.com (G. Kocher); sanjanaarun@ebay.com (S. Arun)

🆔 0009-0004-4832-1580 (G. Kocher); 0009-0008-3710-3958 (S. Arun)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

item’s selling price and ad rate are known), and where engagement—a click or a sale—is modeled as a binary outcome. However, at some marketplaces like eBay, a notable portion of inventory operates under auction dynamics, where prices evolve over time and final transaction value is unknown at ranking time. Auction-based listings account for a meaningful share of inventory and serve an important role in the marketplace: they provide a price discovery mechanism for unique, one-of-a-kind, or hard-to-value items where the seller cannot easily set a fixed price [11, 12], and are particularly popular among individual and consumer-to-consumer (C2C) sellers [13]. Yet these listings have historically been out of scope of CPA advertising programs due to the difficulty of estimating their value in a comparable way. For fixed-monetization programs—both CPC and fixed-price CPA—there is a single source of uncertainty (whether the buyer will engage), while the monetization value conditional on engagement is known in advance. For CPA auction items, however, there are *two* sources of uncertainty: not only whether a buyer will engage, but also what the final selling price (and therefore the platform’s monetization) will be. This creates the need for new approaches to estimate and compare expected revenue across all listing formats, even when that term is only one ingredient of a broader multi-objective ranker.

In summary, our main contributions are:

1. We extend the standard eCPM framework to auction listings by introducing a **marginal eCPM (meCPM)** that accounts for the stochastic, evolving price dynamics of proxy-bidding auctions.
2. We derive meCPM expressions for fixed-price, pure auction, and hybrid Auction with Buy It Now (ABIN) listings, enabling a unified comparison of all listing formats on a revenue basis.
3. We present a practical, simplified implementation that approximates these formulations and still leads to generally well-calibrated revenue prediction, allowing us to address cold-start challenges and launch a new auction CPA ad program.
4. We validate the approach in online A/B tests and achieve statistically significant user metrics and positive revenue gains, and have deployed a version of the system to production.

2. Background and Related Work

Sponsored search ranking in e-commerce platforms typically involves scoring each candidate ad item by its expected revenue per impression (eCPM), and then ranking items directly by revenue [14, 15, 16]; using multi-objective scoring functions [17, 18, 19]; introducing explicit constraints on relevance, quality, inventory, or performance [20, 21, 22]; or estimating the incremental (causal) value of showing each ad [23, 24, 25]. Formally, the expected revenue per impression has the general form [8, 9, 21, 26]:

$$\text{eCPM} = \mathbb{P}(\text{action} \mid \text{impression}) \times R_{\text{action}} \quad (1)$$

where throughout the paper, we omit the conventional factor of 1000 used in Cost-per-Mille notation, and where “action” is the monetizing event and R_{action} is the revenue earned by the platform when the action occurs. For CPC ads, the action is a click and the revenue per click is an advertiser’s bid. For fixed-price CPA, the action is a sale and the revenue per sale is $p_{\text{FP}} \cdot r$, where p_{FP} is the item’s fixed price and the ad rate r is the fraction of the final selling price the seller pays to the platform [27]. Thus,

$$\text{eCPM}_{\text{CPC}} = \mathbb{P}(\text{click} \mid \text{impression}) \times \text{bid} \quad (2)$$

$$\text{eCPM}_{\text{FP}} = \mathbb{P}(\text{sale} \mid \text{impression}) \times p_{\text{FP}} \times r \quad (3)$$

In both of these cases, the outcome is binary—a click happens or not, a sale happens or not—and the revenue gained from the impression, R_{action} , is fixed and known in advance and is realized immediately when the action is taken. Auction listings do not fit this pattern: the final price is unknown at ranking time, and although the transaction is not funded until auction close, a portion of the eventual revenue is already “locked in” by past bids and should not be attributed to future impressions. Therefore, a reasonable ranking objective is the *marginal* expected revenue, or *meCPM*: the incremental revenue that one more impression generates above the locked-in baseline. This distinction conceptually mirrors lift-based approaches to ad ranking [23, 24, 25] and the general economic definition of marginal revenue [28].

For CPC and fixed-price CPA, because revenue gain is “all or none” and there is no “locked-in” partial revenue accumulating over time, the standard definition of expected revenue per impression is the *marginal* expected revenue: the value of showing one more impression equals the full eCPM, so $\text{meCPM}_{\text{CPC}} = \text{eCPM}_{\text{CPC}}$ and $\text{meCPM}_{\text{FP}} = \text{eCPM}_{\text{FP}}$. Similarly, pure Cost-per-Impression ads (CPI, or CPM if scaled by 1000) are a trivial special case, where the impression itself is the action, and by definition, $\text{meCPM}_{\text{CPI}} = \text{eCPM}_{\text{CPI}}$. In the following section, we derive the marginal eCPM for CPA auction and ABIN listings under proxy bidding mechanics (such as at eBay), and show that expected revenue scoring of all listing formats fits into the marginal eCPM framework.

3. Methods

3.1. Pure Auction Format

Proxy Bidding Mechanism In *proxy bidding* (automatic bidding) systems, as used in eBay auctions, when a bidder places a bid, they submit the maximum value they are willing to pay, V , not a visible bid amount. The system automatically bids on their behalf, and on behalf of all previous bidders, raising the current visible price, p_{cur} , only as high as necessary to give the bidder with the highest valuation the leading position in the auction. How much p_{cur} increases depends on the values of several parameters (see (4)): the new bid amount, V ; the current leader’s hidden bid amount, M ; and the minimum bid increment, δ , which determines how much higher V must be relative to the current price p_{cur} (see Appendix A). Bids accumulate during the lifetime of the auction (typically 3, 7, or 10 days), and at auction close, the winner is the bidder with the highest submitted maximum, and they pay the second-highest submitted maximum (plus a small amount, which depends on V , M , and δ). This is similar in outcome to a second-price sealed-bid (Vickrey) auction [11, 29], and in this scenario, bidders are incentivized to submit their true maximum willingness to pay, and there is no incentive for strategic *bid shading* [30]. However, our derivation does not require truthfulness. We only require that V be the submitted maximum bid used by the proxy-bidding mechanism. This distinction is important because in practice, the visible ascending price and fixed close-time give rise to *auction sniping* behavior, where bids are strategically withheld until immediately before the auction ends [31]. Although not our main focus, the framework can in principle capture sniping behavior via the time-dependence of auction state s , which conditions both the bid-through-rate $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ and the conditional CDF F_s from which V is drawn.

Locked-in Revenue In ascending auctions, the price increases as new bids arrive. This leads to a key fact: if the current visible price is p_{cur} , then the platform has already *locked-in* revenue of $p_{\text{cur}} \times r$, which will be realized regardless of future impressions. An additional impression should therefore compete based on the *incremental* revenue it is expected to generate above this locked-in amount.

Greedy Next-Impression Approximation We compute the marginal value of only the single next impression, yielding a greedy, closed-form expression that depends on M , δ , the observable auction/context state s , and the conditional bid-value CDF F_s of submitted maximum bids given that a valid bid occurs. A richer treatment could model the full remaining impression sequence until auction close, e.g. via dynamic programming with Bellman recursions over the auction state.

3.1.1. Marginal eCPM Derivation

Let s denote the observable auction/context state, let p_0 denote the seller’s starting bid price, and let τ denote the valid-bid threshold: $\tau = p_0$ before any bid has been placed, and $\tau = p_{\text{cur}} + \delta$ once the auction has an active bid. The derivation in this subsection assumes the post-bid setting, in which $\tau = p_{\text{cur}} + \delta$. The bid-through-rate model estimates $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$. For a valid bid ($V \geq \tau$), i.e.,

a bid high enough to be successfully submitted to the platform, the price movement is (Appendix B.1):

$$\Delta p_{\text{cur}}(V; p_{\text{cur}}, M, \delta) = \begin{cases} (V + \delta) - p_{\text{cur}} & \text{if } p_{\text{cur}} + \delta \leq V \leq M - \delta & (p_{\text{cur}} \rightarrow V + \delta) \\ M - p_{\text{cur}} & \text{if } M - \delta < V \leq M & (p_{\text{cur}} \rightarrow M) \\ V - p_{\text{cur}} & \text{if } M < V < M + \delta & (p_{\text{cur}} \rightarrow V) \\ (M + \delta) - p_{\text{cur}} & \text{if } V \geq M + \delta & (p_{\text{cur}} \rightarrow M + \delta) \end{cases} \quad (4)$$

Submitted maxima below τ would not be valid bids and would produce no price change. Conditional on a valid bid occurring, the bidder submits a maximum V drawn from the distribution with conditional CDF F_s , where $F_s(v) = \mathbb{P}(V \leq v \mid \text{valid bid}, s)$ —i.e. F_s is supported on valid submitted maxima given the current state, s . Note that although δ is piecewise-constant in price, here we treat it as a single value at the current p_{cur} ; this is exact within a tier and an approximation at tier boundaries.

Taking the expectation of (4) over the bidder's maximum V , and multiplying by the ad rate r and the probability that the impression produces a valid bid, gives the marginal eCPM in factored form:

$$\text{meCPM}_{\text{auction}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot \mathbb{E}_{V \sim F_s}[\Delta p_{\text{cur}}(V; p_{\text{cur}}, M, \delta)], \quad (5)$$

where $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ is the *bid-through-rate*—the probability that a given impression results in a valid bid, conditional on features of the item, context, time, etc.—and the second factor is the expected price uplift given that a valid bid is placed. We expand the conditional expectation using (4) (full derivation in Appendix B), and define the conditional clipped survival integral $\mathcal{J}_s(x, y) \equiv \int_x^y S_s(v) dv$ when $x < y$ and $\mathcal{J}_s(x, y) = 0$ when $x \geq y$, where $S_s(v) = 1 - F_s(v)$ is the conditional survival function of submitted maximum bids given a valid bid in state s . In the post-bid setting (where $\tau = p_{\text{cur}} + \delta$, and M is defined since there is an existing leader):

$$\begin{aligned} \text{meCPM}_{\text{auction}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot & \left[(p_{\text{new}}(\tau) - p_{\text{cur}}) \right. \\ & \left. + \mathcal{J}_s(\tau, M - \delta) + \mathcal{J}_s(\max(\tau, M), M + \delta) \right] \end{aligned} \quad (6)$$

where $p_{\text{new}}(\tau)$ is the visible price after the minimum valid bid, given by $\min(M, \tau + \delta)$ when $\tau \leq M$ and $\min(\tau, M + \delta)$ when $\tau > M$ (see (11)). The three terms inside the bracket have intuitive interpretations. **(i)** $(p_{\text{new}}(\tau) - p_{\text{cur}})$ is the *baseline uplift* from the minimum valid bid, conditional on a valid bid occurring. **(ii)** $\mathcal{J}_s(\tau, M - \delta)$ captures additional upside when the old leader with hidden maximum M keeps the lead and the visible price is pushed up to $V + \delta$. **(iii)** $\mathcal{J}_s(\max(\tau, M), M + \delta)$ captures the case where the challenger barely takes the lead and the visible price moves just high enough to overtake M .

A few observations about this expression are worth noting (see Appendix B.3). Among the three bracket terms, the baseline uplift $(p_{\text{new}}(\tau) - p_{\text{cur}})$ and the third term are each $O(\delta)$, while the second term $\mathcal{J}_s(\tau, M - \delta)$ scales with $M - p_{\text{cur}}$ and can dominate when M is significantly above p_{cur} , reflecting the substantial expected upside from items with hidden headroom for the price to climb. Conditional on a valid bid, a higher current visible price p_{cur} reduces the expected price uplift from the next bid (the bracket term of (6)), although the full meCPM also depends on how $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ and the conditional bid-value distribution F_s vary with the auction state. In practice, items whose visible price has already risen close to their typical final price have lower marginal value per impression. Holding the valid-bid distribution and bid-through-rate fixed, the price-uplift expression is increasing in M : a larger hidden maximum from the current leader gives more room for the visible price to climb.

The zero-bid starting case (no bids yet, $p_{\text{cur}} = p_0$, no locked-in revenue) is simply $\text{meCPM}_{\text{zero-bid}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_0$, i.e. the full settlement price, since any valid first bid becomes the leader at visible price p_0 . See Appendix B.4.

3.1.2. Reserve Prices

Some auctions have a hidden reserve price p_R , a seller-specified threshold below which the item will not sell even if the auction closes with a winning bid. Bidders see whether the reserve has been met

but not its value [30, 32]. The reserve modifies the meCPM formulation. When at least one valid bid has been placed and the reserve has been met (so $p_{\text{cur}} \geq p_R$), the item will sell, locked-in revenue is $p_{\text{cur}} \times r$ as in the no-reserve case, and meCPM reduces to the post-bid pure-auction case in (6). If $p_{\text{cur}} < p_R$, the item is not yet guaranteed to sell, so locked-in revenue is *zero*, not $p_{\text{cur}} \times r$. Unlike the no-reserve derivation above where the zero-bid case needed a separate formula, here both the post-bid and zero-bid starting states are addressed in a single formula. A reserve-clearing bid ($V \geq p_R$) gives a settlement price of $\min(V, C)$, where the settlement cap C is piecewise: $C \equiv \max\{p_R, M + \delta\}$ post-bid, and $C \equiv p_R$ zero-bid (when M is undefined), and only valid bids with $V \geq L$, where $L \equiv \max\{p_R, \tau\}$, also clear the reserve. Taking the expectation over such bids (full derivation in Appendix C) we get:

$$\text{meCPM}_{\text{reserve not met}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot [\min(L, C) S_s(L) + \mathcal{J}_s(L, C)] \quad (7)$$

where $L \equiv \max\{p_R, \tau\}$ and C is given by the piecewise definition above, and $S_s(L) \equiv \mathbb{P}(V > L \mid \text{valid bid}, s)$ is the conditional survival probability that a valid bid also clears the reserve. In the common case $p_R \geq M + \delta$ (and trivially in the zero-bid case via the piecewise definition of C), we have $C = p_R \leq L$, so $\mathcal{J}_s(L, C) = 0$ and the formula reduces to $r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot S_s(L) \cdot p_R$, i.e. the reserve settlement price (scaled by the ad rate) weighted by the probability that the impression produces a reserve-clearing bid.

3.2. ABIN Format

An Auction with Buy It Now (ABIN) listing combines auction and fixed-price mechanics. The item runs as an auction with a starting bid but also has a Buy It Now price p_{BIN} at which a buyer may purchase immediately, ending the auction [33]. So, a single impression can trigger at most one of two mutually exclusive revenue generating events: a BIN purchase or a new bid. The marginal eCPM becomes:

$$\text{meCPM}_{\text{ABIN}} = r \cdot \mathbb{P}(\text{BIN} \mid \text{imp}, s) \cdot p_{\text{BIN}} + \text{meCPM}_{\text{auction}}(s) \quad (8)$$

where $\text{meCPM}_{\text{auction}}(s)$ is the auction-format meCPM appropriate to the listing’s current state: the reserve-not-met case (7) when a reserve is set and unmet; the post-bid pure-auction case (6) when at least one valid bid exists and any reserve is met; or the zero-bid no-reserve form $r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_0$ when no bid has been placed and no reserve applies (Appendix B.4). On eBay, the BIN option typically disappears once the first bid is placed, so $\mathbb{P}(\text{BIN} \mid \text{imp}, s) = 0$ post-bid and the ABIN formula collapses to the auction case. We note the corner case where for some categories the BIN option remains after getting bids, until p_{cur} reaches a fraction of p_{BIN} . In that case, if there is no reserve price, or if the reserve is already met, then $p_{\text{cur}} \cdot r$ is already locked in so the first term of (8) instead uses $(p_{\text{BIN}} - p_{\text{cur}}) \cdot r$.

3.3. Unified Ranking Across Listing Formats

All listing formats share the meCPM structure $r \times \mathbb{P}(\text{event} \mid \text{imp}) \times \text{price signal}$. For CPI, CPC, and fixed-price CPA, there is no locked-in partial revenue, so the standard eCPM is already marginal (meCPM = eCPM). The full meCPM equations (6), (7), and (8), extend this to incrementally evolving prices, and in principle, along with the standard CPC (2) and fixed-price CPA (3) equations, could be used to score all listing formats on a revenue basis and fit them into the overall ranking system with other terms and constraints. However, this requires multiple components not available at launch time due to the cold-start setting.

3.4. Simplified Implementation

We explain several practical simplifications to enable a deployable system which brings statistically significant gains (Section 4.1), discuss alternative modeling approaches, and leave the full meCPM implementation as future work. First, although a proxy-bidding platform may in principle have access to M or p_R if present, a ranking implementation may prefer not to require them, since a score which depends on M or p_R could, in theory, subtly reveal information about these values. Alternative

treatments of M and p_R are left for future work, and for now we omit these quantities. Second, δ : in many e-commerce auctions, the minimum bid increment is small relative to p_{cur} (e.g., \$0.50 for items under \$25, \$1.00 for \$25–\$100), so δ/p_{cur} is small by design. Some terms are $O(\delta)$ corrections that can be omitted to first order. Finally, we omit several mechanics such as buyer offers on zero-bid auction items, and Second Chance Offers, which are relatively uncommon and are downstream events so omitting them has negligible impact to ranking quality for this use case.

Even after these simplifications, the meCPM still requires components such as the valid-bid-through-rate $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ and the conditional bid-value CDF F_s . The bid-through-rate can be fit to data analogous to standard CTR or CVR modeling, with a binary label indicating whether the impression led to a valid bid being placed. Conditional on a valid bid, the bid-value CDF F_s can be estimated from the historical submitted bids in comparable auction states. This aligns with the production data-generating process: the platform can see the values of successfully submitted bids, while the bid-through-rate model captures the probability that an impression produces such a bid. A parametric approach could model f_s using a positive, right-skewed distribution such as a log-normal, gamma, Weibull, or mixture distribution, with parameters conditioned on item and auction state, and use tractable expressions for S_s . The integral $\mathcal{J}_s$ could be evaluated in several ways including, for $x \leq y$, via the standard stop-loss identity $\mathcal{J}_s(x, y) = \mathbb{E}_{V \sim F_s}[(V - x)_+] - \mathbb{E}_{V \sim F_s}[(V - y)_+]$ (see [34]). A non-parametric implementation could use the empirical plug-in estimator $\hat{\mathcal{J}}_s(x, y) = \frac{1}{n} \sum_{i=1}^n \max\{0, \min(V_i, y) - x\}$ operating directly on the discrete pool of valid bid samples $V_1, \dots, V_n$ from comparable auction states, without requiring a fitted continuous distribution. Hybrid and semi-parametric approaches are possible too, e.g. quantile regression to predict a small set of conditional quantiles and reconstruct S_s by interpolation, or k -nearest-neighbor estimators that pool over the most similar historical auction states. In all cases, individual auctions may have only a few discrete submitted bids over their lifetime, so estimation requires pooling across many items and other state features like price and time buckets. The granularity of s trades off bias against variance. Alternatively, one could bypass the full conditional bid-value distribution and summarize the price upside distribution with a point estimate, e.g. using $\hat{p}_f(s)$, the expected final selling price [35, 36] given the current state, and modeling the remaining price uplift as $\max(\hat{p}_f(s) - p_{\text{cur}}, 0)$; or conservatively assuming that any new bid only increases the price by the minimum bid increment, and using $\text{meCPM}_\delta = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot \delta$.

However, these estimators cannot be fit in the cold-start setting. Instead, one possibility is to leverage an existing conversion model but use a hybrid price variant which still captures auction dynamics while also allowing a unified treatment of the new auction and ABIN listing formats alongside fixed-price within a single meCPM score component:

$$\text{meCPM}_{\text{simplified}} = r \cdot p(\text{sale} \mid \text{imp}) \cdot p_{\text{cur}} \quad (9)$$

Comparing (9) to the full meCPM (6), this approximation involves two key substitutions. **(i)** $p(\text{sale} \mid \text{imp})$ replaces $p(\text{valid bid} \mid \text{imp}, s)$: the existing conversion model already captures auction-relevant signals such as item quality, relevance, time dynamics, and engagement activity. Though $p(\text{sale})$ and $p(\text{valid bid})$ predict different events, they are closely related, since high quality, relevant items likely to sell are also likely to attract bids (see Section 4.2). **(ii)** For fixed-price items, $p_{\text{cur}} = p_{\text{FP}}$ and (9) recovers (3) exactly, and for auction and ABIN items, p_{cur} serves as the price signal instead of the full expression. The resulting score preserves the meCPM’s structural form—probability of a revenue event, times a price signal, times the ad rate—while reusing an existing well-calibrated model. Importantly, the approximation is tightest in the regime that covers the most impressions: auctions typically run for 3–10 days and bidding often concentrates near close, so a large fraction of auction impressions are in the zero-bid starting state where $p_{\text{cur}} = p_0$ and the theoretical meCPM reduces to $r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_0$ (Appendix B.4). In this regime, the simplified form (9) has exactly the same structure, differing only in the substitution of $\mathbb{P}(\text{sale})$ for $\mathbb{P}(\text{valid bid})$. A partial additional justification is that the minimum bid increment δ is itself a monotonically increasing step function of p_{cur} (e.g., \$0.50 at \$5–\$25, \$1.00 at \$25–\$100, \$2.50 at \$100–\$250, and so on), so the floor on incremental revenue per bid also scales with p_{cur} (Pearson $r = 0.93$, Spearman $\rho = 0.93$, Kendall $\tau = 0.81$ between δ and p_{cur} for prices up to

Table 1

Online test results. Lift is relative to production control, with statistically significant lifts ($p < .05$) in bold.

Metric	Lift (%)
Ad Revenue Metric 1	+0.43%
Ad Revenue Metric 2	+1.24%
Platform Metric 1	+1.02%
Ad Quality Metric 1	+0.21%
Ad Quality Metric 2	+0.13%

\$10,000). This high rank-order correlation means using p_{cur} in $\text{meCPM}_{\text{simplified}}$ roughly preserves the relative ordering of items by their minimum guaranteed uplift, similar to meCPM_{δ} .

Finally, despite the simplicity of the form in (9), it serves primarily as a *bootstrapping* role: it generates on-policy, auction-specific data, enabling dedicated valid-bid-through-rate models $p(\text{valid bid} \mid \text{imp}, s)$, estimation of the conditional bid-value distribution F_s or survival function S_s , or final price predictor model $\hat{p}_f$, as needed for the full meCPM and its variants. As shown in Section 4.1, even this simplification delivers statistically significant improvements on live traffic.

4. Experiments

4.1. Online Tests

We evaluate the proposed system through experiments on live, session randomized traffic. As seen in Table 1, introducing auction listings into the ad ranking system delivered positive ad revenue gains and statistically significant ($p < .05$) improvements in engagement metrics, driven primarily by newly monetized auction-dominated queries that previously had little CPA coverage. Meanwhile, guardrail metrics were neutral or directionally positive, confirming no degradation in buyer experience. For confidentiality, we anonymize exact metric names but note that metrics related to Search Ad Revenue, Ad CTR, Ad Relevance, and transaction volume all increased.

4.2. Model Evaluation

We evaluate model performance metrics on $> 30\text{M}$ logged production impressions of auction and ABIN listings at rank 1. Discriminative performance is notable: the repurposed $p(\text{sale})$ model achieves an AUC-ROC on auction and ABIN listings that is 93% as high as its AUC-ROC on the existing fixed-price inventory, despite never having been trained specifically for auctions, indicating reasonable transfer of the sale prediction task from existing inventory to auction items. This is not entirely surprising since many of the signals that predict whether a fixed-price item will sell are also informative for predicting auction outcomes, such as buyer intent and item relevance and quality. Nonetheless, it demonstrates the effectiveness of this approach in addressing the cold-start problem to launch a new ad program with no prior data for these item listing types.

The discussion of calibration is more nuanced. We compute $p(\text{sale})$ calibration metrics using 10 bins of predicted probabilities and observe that MAPE on auction and ABIN listings is notably higher than on the existing fixed-price inventory. The model tended to over-predict $\mathbb{P}(\text{sale} \mid \text{imp})$ for auction listings, which is expected: per-impression conversion probability is inherently lower for auctions than for fixed-price items since they can have multiple impressions which collect bids but do not lead to an immediate purchase from that impression which removes it from the platform’s inventory. However, as in (9), in the component for expected revenue estimation, $\mathbb{P}(\text{sale} \mid \text{imp})$ is multiplied by other factors to get the $\text{meCPM}_{\text{simplified}}$, and for auction and ABIN items, since the item’s current bid can only increase in time and not decrease, the overall set of impressions of an ad auction item will necessarily have a current price term which is at or below the final selling price, i.e. the price component of the simplified meCPM score will systematically under-predict ground-truth transaction revenue. These

two points, and the fact that $p(\text{sale})$ models often have decreasing partial dependence on price, mean that various miscalibrating phenomena have opposing effects which partially offset each other. This led to well-calibrated revenue prediction which closely tracked the line $y = x$. Further, since the miscalibration was very nearly linear but with reduced slope, some of the miscalibration is absorbed by other system parameters which scale the meCPM score component and combine it with other terms, and hence also provide their own implicit form of calibration, and dilute the impact of miscalibration within the overall ranking process. Finally, the revenue calibration is further reinforced by the sizeable share of auction impressions in the zero-bid starting state, where the simplified formula closely matches the theoretical meCPM (Section 3.4).

5. Discussion

Overall, our approach demonstrates that auction listings can be successfully integrated into expected value ranking with positive business and user experience impact. We highlight several considerations that are not currently limiting, but may become important as the system scales.

Ad System Tradeoffs A natural question is impression redistribution: do sponsored auction items shift away impressions from fixed-price sponsored listings? Ultimately, this comes down to net platform value: if integrating a new ad format into the ranking system delivers overall positive user engagement and platform metrics, a small shift of impression share across ad programs may be a reasonable tradeoff as long as constraints on guardrail metrics are satisfied and as long as the net effects are positive, as observed in online tests (Section 4.1).

Segmenting queries by their auction-vs-fixed-price impression mix, we observe three groups: auction-heavy queries (a small share of Search Results Pages, SRP), mixed queries (a somewhat larger share), and fixed-price-heavy queries (the large majority of SRP). Revenue gains were concentrated in the mixed-query group, while the auction-heavy and fixed-price-heavy groups showed little net change. In aggregate across all queries, the net effect on fixed-price CPA was neutral, and auction listings exhibited higher click-through rates than fixed-price listings in comparable slots, suggesting that the newly monetized inventory is largely additive rather than just a reallocation of existing demand.

Toward the Full Marginal eCPM Formulation The system described in Section 3.4 is a practical and highly simplified approximation of the marginal eCPM objective derived in Section 3.1. A natural direction for future work is closing the gap between the two. The full formulation requires a valid-bid-through-rate estimator $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ and a conditional bid-value CDF $F_s(v) = \mathbb{P}(V \leq v \mid \text{valid bid}, s)$, which were not initially available due to the cold-start setting. As auction-specific data accumulates over time, however, the full meCPM-based approach and its variants become increasingly feasible to explore in future work, as does quantifying whether the resulting gain over the simplified score justifies the added modeling complexity.

6. Conclusion

We introduced marginal expected revenue per impression (meCPM) as a meaningful objective for ranking mixed listing formats within Cost-per-Acquisition e-commerce search advertising programs. Fixed-price, pure auction, and hybrid (ABIN) listings can all be unified under this generalized formulation of expected revenue in a way that naturally fits ad programs in sponsored search. Although ranking generally includes other terms and blended objectives, not solely revenue, the component for ad monetization for all listing formats can be understood as instances of a unified meCPM formulation: the probability of a revenue generating event, times the expected price uplift above locked-in revenue, times the ad rate. CPC and fixed-price CPA formats are inherently marginal, and the auction meCPM extends this to incrementally evolving prices.

The problem addressed here fits a unique space among major e-commerce platforms, as few marketplaces operate mixed fixed-price and auction listing formats at comparable scale, and even fewer integrate these formats into a unified ad ranking system. To address the cold-start problem of launching a new ad program without prior data, we adopt a simplified approximation of this framework. Live tests on production traffic showed statistically significant gains in user and business metrics. Future work could pursue dedicated auction ranking models for an exact marginal eCPM-based ranker, or intermediate options such as using a point estimate of final selling price to model expected remaining price uplift. Overall, our framework shows that mixed fixed-price and auction listings can be successfully integrated into expected value ranking with positive business and user experience impact.

Acknowledgments

We acknowledge the cross-team collaboration which made this effort possible, and we thank the passionate users of the eBay platform who make it a unique marketplace for buyers and sellers.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT 5.4/5.5 and Claude Opus 4.6/4.7 to check grammar and to suggest refinements for some sentences. After using these tools, the authors reviewed and edited the content as needed, and take full responsibility for the content in this work.

References

- [1] B. Edelman, M. Ostrovsky, M. Schwarz, Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords, *American Economic Review* 97 (2007) 242–259. doi:10.1257/aer.97.1.242.
- [2] H. R. Varian, Position auctions, *International Journal of Industrial Organization* 25 (2007) 1163–1178. doi:10.1016/j.ijindorg.2006.10.002.
- [3] M. Mahdian, K. Tomak, Pay-per-action model for online advertising, in: *Proceedings of the 1st International Workshop on Data Mining and Audience Intelligence for Advertising (ADKDD'07)*, ACM, 2007, pp. 1–6. doi:10.1145/1348599.1348600.
- [4] D. Saha, A. Zaman, H. Zou, N. Chen, X. Shu, N. Vase, A. Bagherjeiran, Improving ad matching via cluster-adaptive keyword expansion and relevance tuning, in: *Proceedings of the SIGIR Workshop on eCommerce (eCom'25)*, volume 4123 of *CEUR Workshop Proceedings*, 2025, pp. 1–9. https://ceur-ws.org/Vol-4123/paper_8.pdf.
- [5] D. Hillard, S. Schroedl, E. Manavoglu, H. Raghavan, C. Leggetter, Improving ad relevance in sponsored search, in: *Proceedings of the Third ACM International Conference on Web Search and Data Mining (WSDM '10)*, ACM, 2010, pp. 361–370. doi:10.1145/1718487.1718532.
- [6] M. S. Chaudhary, J. Nejati, M. Rahman, G. Adalinge, A. Bagherjeiran, Towards quality ad selection: A model-based approach to performance filtering, in: *Proceedings of the SIGIR Workshop on eCommerce (eCom'25)*, volume 4123 of *CEUR Workshop Proceedings*, 2025, pp. 1–9. https://ceur-ws.org/Vol-4123/paper_5.pdf.
- [7] C. Carrion, Z. Wang, H. Nair, X. Luo, Y. Lei, P. Gu, X. Lin, W. Chen, J. Jin, F. Zhu, C. Peng, Y. Bao, Z. Lin, W. Yan, J. Shao, Blending advertising with organic content in e-commerce via virtual bids, in: *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2023)*, volume 37, 2023, pp. 15476–15484. doi:10.1609/aaai.v37i13.26835.
- [8] R. Rosales, H. Cheng, E. Manavoglu, Post-click conversion modeling and analysis for non-guaranteed delivery display advertising, in: *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM '12)*, ACM, 2012, pp. 293–302. doi:10.1145/2124295.2124333.

- [9] B. Yuan, J.-Y. Hsia, M.-Y. Yang, H. Zhu, C.-Y. Chang, Z. Dong, C.-J. Lin, Improving ad click prediction by considering non-displayed events, in: Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM '19), ACM, 2019, pp. 329–338. doi:10.1145/3357384.3358058.
- [10] Z. Zhao, J. Gao, Y. Zhang, S. Han, S. Lou, X.-R. Sheng, Z. Wang, H. Zhu, Y. Jiang, J. Xu, B. Zheng, COPR: Consistency-oriented pre-ranking for online advertising, in: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23), ACM, 2023, pp. 4974–4980. doi:10.1145/3583780.3615465.
- [11] V. Krishna, Auction Theory, 2nd ed., Academic Press, 2009.
- [12] K.-P. Chen, H.-p. Lai, Y.-T. Yu, The seller's listing strategy in online auctions: Evidence from eBay, International Journal of Industrial Organization 56 (2018) 107–144. doi:10.1016/j.ijindorg.2017.11.003.
- [13] L. Einav, C. Farronato, J. Levin, N. Sundaresan, Auctions versus posted prices in online markets, Journal of Political Economy 126 (2018) 178–215. doi:10.1086/695529.
- [14] K. Aryafar, D. Guillory, L. Hong, An ensemble-based approach to click-through rate prediction for promoted listings at Etsy, in: Proceedings of the ADKDD'17, 2017, pp. 1–6. doi:10.1145/3124749.3124758.
- [15] F. Yang, W. Yuan, N. Anwar, F. Huang, B. Huffman, L. Chen, K. Shmakov, S. Goda, J. Jung, N. Moleyar, X. Peng, K.-c. Lee, M. Wen, Ad ranking for Walmart display ads system, in: Proceedings of the SIGIR Workshop on eCommerce (eCom'25), volume 4123 of *CEUR Workshop Proceedings*, 2025, pp. 1–12. https://ceur-ws.org/Vol-4123/paper_4.pdf.
- [16] W. Lu, S. Chen, K. Li, L. V. S. Lakshmanan, Show me the money: Dynamic recommendations for revenue maximization, Proceedings of the VLDB Endowment 7 (2014) 1785–1796. doi:10.14778/2733085.2733086.
- [17] D. Carmel, E. Haramaty, A. Lazerson, L. Lewin-Eytan, Multi-objective ranking optimization for product search using stochastic label aggregation, in: Proceedings of The Web Conference 2020 (WWW '20), 2020, pp. 373–383. doi:10.1145/3366423.3380122.
- [18] S. Kolbin, A. Jain, M. Malgiredy, K. Zimmer, M. Mosmer, P. Phelps, More ads, happier shoppers: Unified-valuation ad allocation at scale, in: A. Bagherjeiran, N. Djuric, K.-c. Lee, L. Pang, V. Radosavljevic, S. Rajan (Eds.), Computational Advertising, volume 2796 of *Communications in Computer and Information Science*, Springer Nature Switzerland, Cham, 2026, pp. 85–98. doi:10.1007/978-3-032-16358-5_7, AdKDD 2025.
- [19] R. Louca, M. Bhattacharya, D. Hu, L. Hong, Joint optimization of profit and relevance for recommendation systems in e-commerce, in: Proceedings of the RMSE Workshop at ACM RecSys 2019, Copenhagen, Denmark, 2019, pp. 1–4. <https://www.dsml.etsy.com/publication-summaries/revenue-optimization-recsys-2019>.
- [20] M. Malgiredy, D. Xie, S. Kolbin, K. Zimmer, M. Mosmer, P. Phelps, Profit aware ad ranking with relevance constraint, in: A. Bagherjeiran, N. Djuric, K.-c. Lee, L. Pang, V. Radosavljevic, S. Rajan (Eds.), Computational Advertising, volume 2796 of *Communications in Computer and Information Science*, Springer Nature Switzerland, Cham, 2026, pp. 71–84. doi:10.1007/978-3-032-16358-5_6, AdKDD 2025.
- [21] Y. Chen, P. Berkhin, B. Anderson, N. R. Devanur, Real-time bidding algorithms for performance-based display ad allocation, in: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '11), ACM, 2011, pp. 1307–1315. doi:10.1145/2020408.2020604.
- [22] H. Lu, L. Zhang, Y. Zhu, The power of linear programming in sponsored listings ranking: Evidence from a large-scale field experiment, arXiv preprint arXiv:2403.14862 (2024). <https://arxiv.org/abs/2403.14862>.
- [23] J. Xu, X. Shao, J. Ma, K.-c. Lee, H. Qi, Q. Lu, Lift-based bidding in ad selection, in: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2016), 2016, pp. 651–657. doi:10.1609/aaai.v30i1.10025.
- [24] B. He, Y. Weng, X. Tang, Z. Cui, Z. Sun, L. Chen, X. He, C. Ma, Rankability-enhanced revenue

- uplift modeling framework for online marketing, in: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '24, ACM, New York, NY, USA, 2024, pp. 5093–5104. doi:10.1145/3637528.3671516.
- [25] L. Bottou, J. Peters, J. Quiñonero-Candela, D. X. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, E. Snelson, Counterfactual reasoning and learning systems: The example of computational advertising, *Journal of Machine Learning Research* 14 (2013) 3207–3260.
 - [26] K.-c. Lee, B. Orten, A. Dasdan, W. Li, Estimating conversion rate in display advertising from past performance data, in: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '12), ACM, 2012, pp. 768–776. doi:10.1145/2339530.2339651.
 - [27] Z. Ge, W. Zhou, J. Lute, A. Ilardi, Relevance constrained re-ranking in sponsored listing recommendations, in: Proceedings of the 27th ACM SIGKDD Workshop on Data Mining for Online Advertising (AdKDD'21), ACM, 2021, pp. 1–7. <http://papers.adkdd.org/2021/papers/adkdd21-ge-relevance.pdf>.
 - [28] N. G. Mankiw, *Principles of Economics*, 10th ed., Cengage Learning, Boston, MA, 2024.
 - [29] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, *Journal of Finance* 16 (1961) 8–37. doi:10.1111/j.1540-6261.1961.tb02789.x.
 - [30] J. G. Riley, W. F. Samuelson, Optimal auctions, *American Economic Review* 71 (1981) 381–392. <https://www.jstor.org/stable/1802786>.
 - [31] A. E. Roth, A. Ockenfels, Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the internet, *American Economic Review* 92 (2002) 1093–1103. doi:10.1257/00028280260344632.
 - [32] P. Bajari, A. Hortaçsu, The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions, *RAND Journal of Economics* 34 (2003) 329–355. <https://www.jstor.org/stable/1593721>.
 - [33] X. Wang, A. L. Montgomery, K. Srinivasan, When auction meets fixed price: A theoretical and empirical examination of Buy-it-Now auctions, *Quantitative Marketing and Economics* 6 (2008) 339–370. doi:10.1007/s11129-008-9041-0.
 - [34] R. Kaas, M. Goovaerts, J. Dhaene, M. Denuit, *Modern Actuarial Risk Theory: Using R*, 2nd ed., Springer, 2008. doi:10.1007/978-3-540-70998-5.
 - [35] R. Ghani, Price prediction and insurance for online auctions, in: Proceedings of KDD 2005, 2005, pp. 411–418. doi:10.1145/1081870.1081918.
 - [36] D. van Heijst, R. Potharst, M. van Wezel, A support system for predicting eBay end prices, *Decision Support Systems* 44 (2008) 970–982. doi:10.1016/j.dss.2007.11.004.

A. Notation

Symbol	Meaning
r	Ad rate (fraction of final selling price paid by seller to platform, $r \in [0, 1]$)
p_{FP}	Price of a fixed-price (FP) listing
p_0	Starting bid price set by the seller on an auction listing
p_{cur}	Current visible price on an auction listing
p_R	Reserve price (if any) on an auction listing
p_{BIN}	Buy It Now price (for ABIN listings)
δ	Minimum bid increment, a piecewise-constant function of price (e.g. \$0.50 at \$5–\$25, \$1.00 at \$25–\$100). Treated as a single value at the current p_{cur} for the meCPM derivation; this is exact within a tier, and an approximation at tier boundaries
M	Current leader’s hidden submitted maximum bid (undefined when no bids have been placed yet)
V	A new potential bidder’s submitted maximum bid
s	Observable auction/context state: p_{cur} , δ , bid count, time remaining, bidder-visible flags (reserve presence, reserve-met, BIN availability), and item/query/context features. The reserve price p_R , BIN price p_{BIN} , and hidden leader maximum M are pulled out as explicit formula arguments
$f_s(v)$	Conditional density of submitted maximum bids given a valid bid in state s
$F_s(v)$	Conditional CDF: $F_s(v) = \mathbb{P}(V \leq v \mid \text{valid bid}, s)$
$S_s(v)$	Conditional survival function: $S_s(v) = 1 - F_s(v) = \mathbb{P}(V > v \mid \text{valid bid}, s)$
τ	Valid-bid threshold: $\tau = p_0$ before any bid; $\tau = p_{\text{cur}} + \delta$ once auction has an active bid
b	Lower limit of the new-leader integral: $b \equiv \max(\tau, M)$
L	Reserve-clearing bid threshold: $L \equiv \max\{p_R, \tau\}$
C	Settlement cap in reserve-not-met regime: $C \equiv \max\{p_R, M + \delta\}$ post-bid; $C \equiv p_R$ zero-bid
$\mathcal{J}_s(x, y)$	Conditional clipped survival integral: $\mathcal{J}_s(x, y) \equiv \int_x^y S_s(v) dv$ if $x < y$; 0 if $x \geq y$
$\hat{p}_f(s)$	Predicted (expected) final selling price for items of this type, in this state

B. meCPM Derivation: Auction Without Reserve Price

We derive the marginal eCPM for an auction with no reserve price from Section 3.1. Assume there is no reserve price in effect ($p_R = 0$, or equivalently $p_{\text{cur}} \geq p_R$). If at least one bid has been placed, the item will sell, so the current locked-in revenue is $p_{\text{cur}} \cdot r$ and it will be realized regardless of whether any additional impressions are shown.

B.1. Effect of One New Bid on the Final Price

Let s denote the observable auction/context state and let τ denote the valid-bid threshold; in the post-bid setting considered here, $\tau = p_{\text{cur}} + \delta$. The bid-through-rate model estimates $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$. Conditional on a valid bid occurring, the bidder submits a maximum V drawn from the conditional CDF F_s , where $F_s(v) = \mathbb{P}(V \leq v \mid \text{valid bid}, s)$. For completeness, a submitted maximum below τ would not constitute a valid bid and would not produce any price change. Since F_s is conditional on a valid bid, the expectation is taken only over $V \geq \tau$, for which the visible price changes as

$$\Delta p_{\text{cur}}(V; p_{\text{cur}}, M, \delta) = \begin{cases} (V + \delta) - p_{\text{cur}} & \text{if } p_{\text{cur}} + \delta \leq V \leq M - \delta & (p_{\text{cur}} \rightarrow V + \delta) \\ M - p_{\text{cur}} & \text{if } M - \delta < V \leq M & (p_{\text{cur}} \rightarrow M) \\ V - p_{\text{cur}} & \text{if } M < V < M + \delta & (p_{\text{cur}} \rightarrow V) \\ (M + \delta) - p_{\text{cur}} & \text{if } V \geq M + \delta & (p_{\text{cur}} \rightarrow M + \delta) \end{cases} \quad (10)$$

Each case of (10) is worth understanding:

1. The current leader's proxy has room to outbid the challenger: proxy bidding automatically raises the visible price from p_{cur} to $V + \delta$. The leader remains the same.
2. The challenger's maximum is within one δ increment of M , so the leader's proxy wants to bid $V + \delta$ but cannot exceed its own maximum M . Visible price rises to M ; leader stays the same.
3. The challenger outbids the leader ($V > M$) and takes the lead, but $V < M + \delta$, so the second-price settlement $M + \delta$ exceeds the new leader's own maximum. The visible price is capped at V .
4. The challenger outbids the leader by at least one full increment ($V \geq M + \delta$). The visible price rises to the old leader's maximum plus one increment, $M + \delta$.

The four case piecewise format above is intuitive, but for convenience in later steps, we rewrite it below in a different format split into two cases by whether or not the new bid V overtakes the current winning bid M . We write the new visible price as:

$$p_{\text{new}}(V) = \begin{cases} \min(M, V + \delta), & \tau \leq V \leq M, \quad (\text{old leader}) \\ \min(V, M + \delta), & V > M, \quad (\text{new leader}) \end{cases} \quad (11)$$

In the first case the current leader's proxy bids $V + \delta$ to stay ahead, capped at the leader's own maximum M . In the second the challenger becomes the new leader at one increment above M , capped at the challenger's own maximum V .

The price uplift is $\Delta p_{\text{cur}}(V) = p_{\text{new}}(V) - p_{\text{cur}}$ for valid bids and zero otherwise.

B.2. Expected Incremental Revenue

The marginal eCPM in factored form is:

$$\text{meCPM}_{\text{auction}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot \mathbb{E}_{V \sim F_s}[\Delta p_{\text{cur}}(V; p_{\text{cur}}, M, \delta)] \quad (12)$$

For notational simplicity, we now assume a continuous F_s ; discrete bid grids can be handled by the corresponding sum form. Since F_s is the conditional distribution of submitted maxima given a valid bid, and valid bids satisfy $V \geq \tau$, the conditional expectation is:

$$\mathbb{E}_{V \sim F_s}[\Delta p_{\text{cur}}(V)] = \int_{\tau}^{\infty} \Delta p_{\text{cur}}(V) f_s(V) dV \quad (13)$$

where f_s is the conditional density of submitted maximum bids given a valid bid in state s .

We now derive (13) in closed form by integrating by parts over each branch of (11). To handle regions that may be empty (when M is close to p_{cur}), we use a conditional clipped survival integral

$$\mathcal{J}_s(x, y) \equiv \begin{cases} \int_x^y S_s(v) dv, & x < y, \\ 0, & x \geq y, \end{cases} \quad (14)$$

which is zero whenever the integration interval is empty.

Using $\Delta p_{\text{cur}}(V) = p_{\text{new}}(V) - p_{\text{cur}}$, split the integral (13) into the two cases of (11):

$$\int_{\tau}^{\infty} \Delta p_{\text{cur}}(V) f_s(V) dV = \underbrace{I_1}_{\text{old-leader}} + \underbrace{I_2}_{\text{new-leader}}$$

where $I_1 = \int_{\tau}^M (\min(M, V + \delta) - p_{\text{cur}}) f_s(V) dV$ when $\tau < M$ (and $I_1 = 0$ when $\tau \geq M$), and $I_2 = \int_{\max(\tau, M)}^{\infty} (\min(V, M + \delta) - p_{\text{cur}}) f_s(V) dV$.

Old-leader region (I_1 , present when $\tau < M$). When $\tau < M$, integrate by parts with

$$u = \min(M, V + \delta) - p_{\text{cur}}, \quad dv = f_s(V) dV, \quad v = -S_s(V)$$

$$I_1 = \left[-(\min(M, V + \delta) - p_{\text{cur}}) S_s(V) \right]_{\tau}^M + \int_{\tau}^M \frac{d}{dV} [\min(M, V + \delta)] S_s(V) dV \quad (15)$$

The left term of (15) evaluates to

$$-\underbrace{(\min(M, M + \delta) - p_{\text{cur}})}_{= M - p_{\text{cur}}} S_s(M) + \underbrace{(\min(M, \tau + \delta) - p_{\text{cur}})}_{= p_{\text{new}}(\tau) - p_{\text{cur}}} S_s(\tau)$$

Under the continuous- F_s convention, there is no point mass at the valid-bid threshold, so $P(V = \tau \mid \text{valid bid}, s) = 0$. Since F_s is conditional on valid bids ($V \geq \tau$), this gives $F_s(\tau) = 0$ and hence $S_s(\tau) = 1 - F_s(\tau) = 1$, so the second term simplifies to $p_{\text{new}}(\tau) - p_{\text{cur}}$. For the right term of (15), the derivative of $\min(M, V + \delta)$ with respect to V is 1 when $V + \delta < M$ (i.e. $V < M - \delta$) and 0 when $V + \delta > M$ (i.e. $V > M - \delta$). Therefore

$$I_1 = (p_{\text{new}}(\tau) - p_{\text{cur}}) - (M - p_{\text{cur}}) S_s(M) + \mathcal{J}_s(\tau, M - \delta) \quad (16)$$

New-leader region (I_2). Use the same integration by parts, now with $u = \min(V, M + \delta) - p_{\text{cur}}$, $dv = f_s(V) dV$, $v = -S_s(V)$, and writing $b \equiv \max(\tau, M)$ for the lower limit:

$$I_2 = \left[-(\min(V, M + \delta) - p_{\text{cur}}) S_s(V) \right]_b^{\infty} + \int_b^{\infty} \frac{d}{dV} [\min(V, M + \delta)] S_s(V) dV \quad (17)$$

For the first term of (17), at $V \rightarrow \infty$ the boundary term vanishes (bounded integrand times $S_s(V) \rightarrow 0$) and at $V = b$, $\min(b, M + \delta) = b$, since $b = \max(\tau, M) \leq M + \delta$ (because $\tau = p_{\text{cur}} + \delta \leq M + \delta$). So the left term contributes $(b - p_{\text{cur}}) S_s(b)$. For the second term of (17), the derivative of $\min(V, M + \delta)$ is 1 when $V < M + \delta$ and 0 when $V > M + \delta$. Therefore

$$I_2 = (b - p_{\text{cur}}) S_s(b) + \mathcal{J}_s(b, M + \delta) \quad (18)$$

Combining I_1 and I_2 . Add (16) and (18). When $\tau < M$, $b = M$ and the terms $-(M - p_{\text{cur}}) S_s(M)$ from I_1 and $+(b - p_{\text{cur}}) S_s(b)$ from I_2 cancel, giving:

$$I_1 + I_2 = (p_{\text{new}}(\tau) - p_{\text{cur}}) + \mathcal{J}_s(\tau, M - \delta) + \mathcal{J}_s(M, M + \delta) \quad (19)$$

When $\tau \geq M$, $I_1 = 0$ (the integration interval $[\tau, M]$ is empty, per its piecewise definition above) and I_2 alone gives the result (with $b = \tau$, so $p_{\text{new}}(\tau) - p_{\text{cur}} = b - p_{\text{cur}}$):

$$\begin{aligned} I_1 + I_2 &= 0 + (\tau - p_{\text{cur}}) S_s(\tau) + \mathcal{J}_s(\tau, M + \delta) \\ &= (\tau - p_{\text{cur}}) + \mathcal{J}_s(\tau, M + \delta) \\ &= (p_{\text{new}}(\tau) - p_{\text{cur}}) + \mathcal{J}_s(\tau, M + \delta) \end{aligned} \quad (20)$$

where we use $S_s(\tau) = 1$ (since F_s is assumed continuous and is conditional on a valid bid, i.e. $V \geq \tau$), and $p_{\text{new}}(\tau) = \tau$ from the new-leader case of (11) (since $\tau \geq M$, and $\tau = p_{\text{cur}} + \delta \leq M + \delta$ gives $\min(\tau, M + \delta) = \tau$). So, in both cases (19) and (20), the sum of $I_1 + I_2$ is:

$$I_1 + I_2 = (p_{\text{new}}(\tau) - p_{\text{cur}}) + \mathcal{J}_s(\tau, M - \delta) + \mathcal{J}_s(\max(\tau, M), M + \delta) \quad (21)$$

where the 2nd term vanishes when $\tau \geq M$ (interval $[\tau, M - \delta]$ is empty, by definition of (14)), and the 3rd term unifies both cases via $b = \max(\tau, M)$. Therefore:

$$\boxed{\text{meCPM}_{\text{auction}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot \left[(p_{\text{new}}(\tau) - p_{\text{cur}}) + \mathcal{J}_s(\tau, M - \delta) + \mathcal{J}_s(\max(\tau, M), M + \delta) \right]} \quad (22)$$

in the post-bid setting (where $\tau = p_{\text{cur}} + \delta$).

B.3. Sanity Checks

Non-negativity. Each component of (22) is non-negative, as we would expect from an equation for marginal revenue: $\mathbb{P}(\cdot), r \in [0, 1]$; $p_{\text{new}}(\tau) \geq p_{\text{cur}}$ by definition; and $\mathcal{J}_s(x, y) \geq 0$. Therefore $\text{meCPM} \geq 0$, with units of price per impression (multiply by 1000 for Cost-per-Mille).

Monotonicity in p_{cur} . For a *fixed* conditional distribution F_s , raising p_{cur} increases $\tau = p_{\text{cur}} + \delta$, which shrinks both $\mathcal{J}_s$ intervals from the left, and the baseline term $(p_{\text{new}}(\tau) - p_{\text{cur}})$ also weakly decreases. However, the conditional distribution F_s itself depends on the auction state s , which includes p_{cur} : raising p_{cur} raises the valid-bid threshold, changing the distribution that S_s and $\mathcal{J}_s$ are computed from. Therefore the conditional expected uplift (the bracket term in (22)) is not guaranteed to be monotonically decreasing in p_{cur} . The full meCPM also depends on how $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ varies with the auction state, so strict monotonicity of the meCPM in p_{cur} requires conditions on both F_s and $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$. Qualitatively, items whose visible price has already climbed close to their expected final price may tend to have lower marginal value per impression, and we might expect the bid-through-rate to have a decreasing partial dependence on p_{cur} in many regimes. However, via possible *social-proof* effects, where prospective bidders interpret a higher p_{cur} or a larger accumulated bid count as a signal of item desirability and become more likely to bid themselves, may flip the dependence to increasing in some regions. In such cases, the full meCPM might not decrease monotonically in p_{cur} .

Monotonicity in M . Because M is hidden from bidders and is not included in the observable state s , the conditional bid-value distribution F_s and the valid-bid-through-rate $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ do not depend on M . Thus the monotonicity enters only through the integral bounds of the price-uplift term. In the bracket of (22), raising M extends the upper limit of $\mathcal{J}_s(\tau, M - \delta)$, adding $S_s(M - \delta) dM$ of area, while shifting $\mathcal{J}_s(\max(\tau, M), M + \delta)$ rightward, which changes it by $(S_s(M + \delta) - S_s(M)) dM$. The net rate of change is $S_s(M - \delta) - S_s(M) + S_s(M + \delta) \geq 0$, since S_s is decreasing and non-negative. The term $(p_{\text{new}}(\tau) - p_{\text{cur}})$ is also non-decreasing in M through $p_{\text{new}}(\tau)$. So overall, meCPM is monotonically increasing in M , which makes sense because impressions of listings with more hidden upside M have more marginal value to the platform.

Term magnitudes. The first and third bracket terms of (22) are each $O(\delta)$, while the second term $\mathcal{J}_s(\tau, M - \delta)$ scales with $M - p_{\text{cur}}$ and dominates when M is significantly above p_{cur} .

Extreme case $M = p_{\text{cur}}$. Both $\mathcal{J}_s$ terms vanish and $p_{\text{new}}(\tau) - p_{\text{cur}} = \delta$, so any valid bid would raise the visible price by exactly δ : $\text{meCPM} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot \delta$.

B.4. Zero-Bid Starting Case

The derivation of Appendix B.1 and Appendix B.2 above assumes at least one bid has been placed (so M is defined, with $M \geq p_{\text{cur}}$). In the zero-bid starting case, no leader exists and any valid bid $V \geq p_0$ becomes the first leader at a visible price of p_0 . The marginal revenue is the full settlement price (since locked-in revenue is zero), so $\text{meCPM}_{\text{zero-bid}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_0$.

C. meCPM Derivation: Auction With Reserve Price

We extend the pure-auction derivation of Appendix B to auctions with a reserve price p_R . The item sells only if the winning bid exceeds p_R . This creates the two regimes below.

C.1. Reserve Already Met ($p_{\text{cur}} \geq p_R$)

If the reserve is already met, i.e. the current visible price p_{cur} is greater than the reserve price p_R , the item will sell and locked-in revenue is $p_{\text{cur}} \cdot r$. The situation is identical to the pure-auction case of (22).

C.2. Reserve Not Yet Met ($p_{\text{cur}} < p_R$)

If the reserve is not met, locked-in revenue is *zero* since if the auction ended now there would be no sale. We start with the post-bid sub-case here, where M is defined and $M < p_R$. The zero-bid sub-case is recovered from the same formula. A new bid with $V < p_R$ does not clear the reserve, so under the greedy next-impression only assumption (i) in Section 3.1, the incremental revenue is zero. (We note that alternative modeling approaches, e.g. a dynamic-programming formulation, could assign non-zero value to such bids, since raising p_{cur} toward p_R , even without clearing it, brings the auction closer to a state where a future bid does clear the reserve and unlocks revenue.) On the other hand, a new bid with $V \geq p_R$ clears the reserve and takes the lead (since $V \geq p_R > M$ in this post-bid sub-case). The second-highest maximum is $M < p_R$, so the settlement price under the proxy-bidding rule (second case of (11)) would be $\min(V, M+\delta)$, but now in the reserve price setting is floored at p_R , i.e. the platform sets p_{cur} to the larger of the proxy-bidding price and p_R , so:

$$p_{\text{settle}}(V) = \max(p_R, \min(V, M+\delta))$$

Since locked-in revenue is zero in the reserve-not-met regime, under the greedy next-impression only assumption (i) in Section 3.1, the marginal revenue from a reserve-clearing bid is the full settlement price $p_{\text{settle}}(V) \cdot r$. The meCPM is therefore $r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ times the expected settlement price over valid bids, accounting for the reserve-clearing condition. To compute this expectation, we first simplify $p_{\text{settle}}(V)$.

Define $C \equiv \max\{p_R, M+\delta\}$. For reserve-clearing bids ($V \geq p_R$), this simplifies to $p_{\text{settle}}(V) = \min(V, C)$: if $M + \delta \geq p_R$ then $C = M + \delta$ and the reserve floor has no effect, so $p_{\text{settle}} = \min(V, M+\delta) = \min(V, C)$; if $M + \delta < p_R$ then $C = p_R$ and $p_{\text{settle}} = \max(p_R, \min(V, M+\delta)) = p_R = \min(V, C)$ since $V \geq p_R = C$. So we can write $p_{\text{settle}} = \min(V, C)$.

Next, let τ be the valid-bid threshold: $\tau = p_0$ if no valid bid has been placed yet, and $\tau = p_{\text{cur}} + \delta$ once the auction has an active bid. Conditional on a valid bid, $V \sim F_s$. Only valid bids with $V \geq L \equiv \max\{p_R, \tau\}$ are both valid and reserve clearing. The expected settlement price contribution over valid bids is:

$$\mathbb{E}_{V \sim F_s} [\min(V, C) \cdot \mathbf{1}\{V \geq L\}] = \int_L^\infty \min(V, C) f_s(V) dV$$

To evaluate this integral, split at C :

$$\int_L^\infty \min(V, C) f_s(V) dV = \int_L^C V f_s(V) dV + C \int_C^\infty f_s(V) dV \quad (23)$$

where the first integral on the right-hand side of (23) is present only when $L < C$ (and is zero when $L \geq C$). For the first integral, integrate by parts with $u = V$, $du = dV$, $dv = f_s(V) dV$, $v = -S_s(V)$:

$$\int_L^C V f_s(V) dV = [-V S_s(V)]_L^C + \int_L^C S_s(V) dV = -C S_s(C) + L S_s(L) + \mathcal{J}_s(L, C)$$

where $\mathcal{J}_s$ is the conditional clipped survival integral defined in (14). By the definition of survival function, the second integral of (23) evaluates to $C S_s(C)$. So, adding the two terms of (23), the $C S_s(C)$ terms cancel, and we get:

$$\int_L^\infty \min(V, C) f_s(V) dV = L S_s(L) + \mathcal{J}_s(L, C)$$

The above assumes $L < C$. When $L \geq C$, every reserve-clearing bidder has $V \geq L \geq C$, so $\min(V, C) = C$ throughout and the integral evaluates to $C S_s(L)$. Both cases ($L < C$ and $L \geq C$) are captured by the expression $\min(L, C) S_s(L) + \mathcal{J}_s(L, C)$, since when $L \geq C$ we have $\min(L, C) = C$ and $\mathcal{J}_s(L, C) = 0$. So, the meCPM in the reserve-not-met regime is r times $\mathbb{P}(\text{valid bid} \mid \text{imp}, s)$ times this expected settlement contribution:

$$\boxed{\text{meCPM}_{\text{reserve not met}} = r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot [\min(L, C) S_s(L) + \mathcal{J}_s(L, C)]} \quad (24)$$

where $L \equiv \max\{p_R, \tau\}$ and $C \equiv \max\{p_R, M+\delta\}$. Intuitively, the first term $\min(L, C) S_s(L)$ is the floor settlement value (either L itself or the cap C , whichever is smaller) earned conditional on the bid clearing L , and $\mathcal{J}_s(L, C)$ is the additional expected upside from V exceeding L on its way up to C .

The derivation above assumes M is defined (i.e. at least one bid has been placed). In the zero-bid sub-case (M undefined, no incumbent), the piecewise definition gives $C \equiv p_R$; combined with $L = p_R$ (since $\tau = p_0 < p_R$ in this regime), the formula collapses to $r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_R \cdot S_s(p_R)$, the marginal value of a reserve-clearing first bid, which pays exactly p_R in the absence of an existing bid.

C.2.1. Common Case

Post-bid, we are in the reserve-not-met regime with $p_R > M$. In all but the narrow edge case where M is within one bid increment of p_R (i.e. $M < p_R < M + \delta$), we have $p_R \geq M + \delta$ and thus $C = p_R \leq L$, so $\mathcal{J}_s(L, C) = 0$ and the formula (24) reduces to $r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_R \cdot S_s(L)$ – the reserve price weighted by the probability that a valid bid also clears the reserve. In the narrow edge case $M + \delta > p_R$, $\mathcal{J}_s(L, C)$ adds a small additional upside (bounded by δ). Thus, the marginal eCPM in the reserve-not-met regime is driven mainly by the value of just clearing the reserve at p_R (the first term of (24)), with the second term adding at most a δ -sized correction in the narrow edge case where M is within one increment of p_R .

D. Summary

The following table collects the marginal eCPM formulas for all listing types.

Listing Type	Marginal eCPM
Fixed Price	$r \cdot \mathbb{P}(\text{sale} \mid \text{imp}) \cdot p_{\text{FP}}$
Auction, post-bid, no reserve	$r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot [(p_{\text{new}}(\tau) - p_{\text{cur}}) + \mathcal{J}_s(\tau, M-\delta) + \mathcal{J}_s(\max(\tau, M), M+\delta)]$
Auction, zero-bid, no reserve	$r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_0$
Auction, post-bid, reserve already met	Same as Auction, post-bid, no reserve
Auction, post-bid, reserve not met	$r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot [\min(L, C) S_s(L) + \mathcal{J}_s(L, C)]$
Auction, zero-bid, reserve not met	$r \cdot \mathbb{P}(\text{valid bid} \mid \text{imp}, s) \cdot p_R \cdot S_s(p_R)$
ABIN	$r \cdot \mathbb{P}(\text{BIN} \mid \text{imp}, s) \cdot p_{\text{BIN}} + \text{meCPM}_{\text{auction}}(s)$, where $\text{meCPM}_{\text{auction}}(s)$ is whichever auction-format row above applies to the listing's current state