TI 2026-075/VII
Tinbergen Institute Discussion Paper

Randomizing in Search Ad Auctions to Extract Revenue from Runaway Winners

Simon van Tartwijk [1]

1 University of Amsterdam

Randomizing in Search Ad Auctions to Extract Revenue from Runaway Winners*

Simon van Tartwijk†

September 22, 2026

Abstract

This paper studies a randomizing threshold auction of the kind online platforms use to sell search-advertising slots. The mechanism awards the slot outright only when the highest bid sufficiently exceeds the runner-up's, and otherwise assigns it by lottery among the close bidders. Such relative-bid randomization can raise revenue above the standard second-price auction when advertisers' values are bimodal: usually low but occasionally very high, reflecting whether an ad matches the user's query. It does so by threatening *runaway winners* with the risk of losing, and charging them a *certainty premium* to restore a sure win. Welfare effects cut both ways. A second-price auction having only its reserve with which to tax a runaway winner can, once the best-matched advertisers are far enough ahead, only refuse to sell to the weaker ones outright. The randomizing auction has a second instrument and keeps serving weaker bidders. Over the range, randomization raises both revenue and total welfare. However, the best-matched advertisers are nonetheless taxed and ad relevance for users is lowered. A learning exercise further shows that randomization prevents advertisers from separately inferring rival values and the pricing rule from observed outcomes. These results are then interpreted in the light of the U.S. antitrust case against Google, which ability to impose the randomized generalized second-price auction (rGSP) was held to be direct evidence of monopoly power and the resulting remedies have begun to require auction transparency.

Keywords: Randomizing auctions, sponsored search auctions, online advertising platforms, mechanism design, antitrust
JEL-codes: D44, D82, L41, L13

1 Introduction

Each time a user submits a query to a major search engine, an auction runs in milliseconds to decide which advertisements appear alongside the results, at the scale of billions of queries

*I thank Maarten Pieter Schinkel, Jan Tuinstra, Sander Onderstal, Michael D. Whinston, and Sindri Engilbertsson for helpful discussions and comments. All remaining errors are my own.

†University of Amsterdam, and Tinbergen Institute. Email: s.i.vantartwijk@uva.nl

per day. In 2019, Google modified the rules of its search-advertising auction through a launch internally codenamed *Polyjuice*, after the Harry Potter potion that lets one person assume the appearance of another.[1,2] The launch introduced the *Randomized Generalized Second-Price Auction* (rGSP): when the top two ad scores are close, the mechanism artificially inflates the runner-up's score so that the runner-up *appears* as a higher bidder and sometimes wins the slot.[3] The threshold for "close" can be wide: in some cases "the auction winner would have to bid 3.7 times higher than the runner-up to avoid swapping."[4,5]

rGSP featured prominently in the antitrust trial against Google. In August 2024, Judge Amit Mehta held that "Google is a monopolist" under Section 2 of the Sherman Act and cited Google's use of "pricing knobs," explicitly including rGSP, as direct evidence of monopoly power, applying the standard from *United States v. Microsoft Corp.* that certain conduct is itself evidence of monopoly power because "a firm without a monopoly would have been unable" to engage in it.[6] Plaintiffs' posed that a platform under competitive pressure could not impose mechanisms like rGSP because advertisers facing an opaque, stochastic allocation introduced to extract rent, would shift spend elsewhere. Google's economic expert Mark Israel read it instead as an investment in platform attractiveness, softening winner-take-all dynamics and broadening the variety of ads shown, and disputed that prices had risen at all.[7] The court resolved that dispute against Google's favour on the broader empirical record.[8]

The contribution of this paper is a tractable model that shows whether a mechanism of

[1]Unless explicitly stated otherwise, all citations to court filings, trial transcripts, trial exhibits, and the court's opinion in this paper refer to *United States v. Google LLC* (consolidated with the parallel state-plaintiff case *State of Colorado et al. v. Google LLC*), Case Nos. 20-cv-3010 and 1:20-cv-03715-APM, U.S. District Court for the District of Columbia. Subsequent footnotes use the following short forms for materials from this case: *Mem. Op.* for the court's Memorandum Opinion on liability (Mehta, J., filed Aug. 5, 2024, Doc. 1033); *Final Judgment* for the court's Final Judgment on remedies (Mehta, J., filed Dec. 5, 2025, Doc. 1462); *Pls.' Proposed FOF* for *Plaintiffs' Proposed Findings of Fact*; *Def.'s Proposed FOF* for *Defendant Google LLC's Proposed Findings of Fact and Conclusions of Law*; *Pls.' Closing Statement* for *Plaintiffs' Closing Statement on Search Advertising* (Doc. 421661); *Trial Tr.* for the trial transcript (cited with ECF Doc. number and page); and "Ex. UPX####" / "Ex. DX####" for trial exhibits (the UPX/DX prefix denotes Plaintiffs' or Defendant's, respectively). All case material is available at `https://www.justice.gov/atr/case/us-and-plaintiff-states-v-google-llc`.

[2]*Trial Tr.*, Doc. 1014, pp. 422–423 (plaintiffs' closing argument); see also *Trial Tr.*, Doc. 959, p. 4180 (witness confirming the codename on cross-examination).

[3]*Pls.' Closing Statement*, pp. 94–101.

[4]*Pls.' Proposed FOF*, ¶ 674.

[5]The exact randomizing mechanism remains obscure. As one internal Google document stated (Ex. UPX0059; see also *Pls.' Proposed FOF*, ¶ 673): "We don't want to have to say we randomize – that will have perception problems" to which Adam Juda, Vice President of Product Management at Google, responded: "If I have to say 'we randomly disable you if you don't bid high enough,' I'm going to have another bad year at GMN [Google Marketing Next]."

[6]*Mem. Op.*, pp. 189–190: Google's ability to set prices in its ad auctions "without considering [its] rivals' prices," the court found, qualifies.

[7]*Trial Tr.*, Doc. 993, pp. 8582–8584 (Israel direct examination) and p. 8567 (empirical defence on prices).

[8]*Mem. Op.*, p. 259 (price-effect finding).

this kind can raise revenue at all, under what conditions, and at whose expense. Central to the model is a prior-light randomized mechanism to auction off a single advertisement slot, referred to here as the *randomizing threshold auction* (rTA). In rTA, the highest bidder wins with certainty when the relative difference between her bid and that of the runner-up exceeds some threshold. When one or more rival bids are sufficiently close to the highest bidder's in relative terms, the mechanism allocates the object by fair lottery among them and the highest bidder. To obtain revenue-results in analytical form, the paper focuses on the unique payment rule that implements this allocation under *dominant-strategy incentive compatibility* (DSIC). The main question is whether rTA can improve revenue over the standard second-price auction (SPA). The SPA is the natural benchmark because Google's underlying auction is a *generalized* second-price auction, which for a single slot reduces to the SPA, the pre-randomization status quo that rGSP modifies.[9]

Answering this question requires taking a stand on the value environment, and online advertising supplies a natural one. Bidder valuations are shaped by a matching process between advertisers and users, and by the platform's *screening* of advertisers: the institutional choices (keyword-matching rules, quality-score filters) that determine which advertisers are eligible to bid for a given query. In setting the screening, a platform balances having thicker, more competitive auctions against the user-side cost of showing low-relevance ads. As advertisers are increasingly targeted to specific impressions and screening choices widen the eligible pool, markets for a given query come to span a wide range of relevance: most advertisers are poor matches while only a small number are highly relevant (Celis et al., 2014). This creates a bimodal value distribution: bidder valuations for the object are usually low but occasionally very high.

Under the second-price auction, this can generate *runaway winners*: advertisers who are far ahead of their rivals and therefore win with certainty while paying only what is needed to beat a much weaker competitor.[10] Although they are efficient, these outcomes leave substantial revenue on the table, which is why preventing runaway winners is a major objective for Google's auction design.[11] Randomization addresses runaway winners by weakening the allocation advantage of bidders who are only moderately ahead, while preserving certainty

[9]Under symmetric independent private values the SPA is moreover revenue-equivalent to the other standard formats that award the object to the highest bidder, so among non-randomized auctions the choice of benchmark is without loss.

[10]A vivid illustration comes from New Zealand's 1990 sealed-bid second-price spectrum auctions, run without reserve prices: in one case the winner had bid NZ$100,000 but paid the second bid of NZ$6, and in another a firm that bid NZ$7 million paid only NZ$5,000. Total receipts came in far below the government's forecasts. The episode became a cautionary tale for the role of reserves (McMillan, 1994); this paper shows that randomization can also resolve the runaway-winner problem.

[11]*Pls.' Proposed FOF*, ¶ 649.

for those whose lead is large. Once the mechanism introduces a sufficient probability that the winner can be swapped, it can charge strictly more for the certainty of winning. Plaintiffs' economic expert Michael D. Whinston characterizes the resulting effect as "introducing inefficiency into the auction ... to extract more out of ... the top advertiser," drawing an analogy with the sprinter Usain Bolt. A runner so far ahead of the field that he can coast to the finish line will run flat-out only if rivals are secretly placed at random head-start positions on the track that he cannot see.[12] The resulting payment increment is the *certainty premium* that drives the revenue gain.

The main results show that relative-bid randomization can improve revenue, but only under specific conditions on the value distribution and on the platform's own screening. In the two-bidder case, with both mechanisms posting the same low reserve, randomization pays only when the certainty premium extracted from well-matched advertisers outweighs the misallocation cost. That cost is incurred when the lottery falls between two advertisers of similar value and hands the slot to the wrong one. The first can only dominate when the eligible pool of bidder is sufficiently screened to prevent costly lotteries in which a poorly-matched advertiser wins.

Letting each mechanism post the reserve it prefers sharpens the comparison, and is where the paper's central point lies. An additional benefit of posting a reserve is that it removes the lotteries with poorly-matched advertisers, so it is worth more in rTA than in SPA. The deeper difference is that the second-price auction has one instrument, the reserve, for two objectives: tax runaway winners and serve lower-valued advertisers. Once the best-matched advertisers are far enough ahead it can only effectively do the first by giving up the second. The randomizing auction has two instruments and can do both jobs at once: it serves the lower-valued advertisers, but can still charge runaway winners a high price by threatening them with the lottery if they are unwilling to pay it.

This opens a range of environments in which the randomizing auction delivers both more revenue and more *welfare*, defined as total expected surplus,[13] than a second-price auction, each posting its own optimal reserve, because the second-price auction buys its revenue by excluding advertisers rather than by misallocating among them.Welfare in this context means total expected surplus: However, even if welfare increases, the redistribution rTA induces falls systematically on matched advertisers and on users, a result that cuts against the procompetitive defence of rGSP advanced by Google's experts.

Two extensions complement the analysis. First, extending the model to N bidders

[12] *Trial Tr.*, Doc. 963, pp. 4793–4794 (Whinston direct examination).

[13] This equals the expected valuation of the advertiser who is allocated the impression, unless the slot is not allocated at all, in which case it is zero. Surplus of users that get to see the ads are not explicitly modelled.

strengthens the revenue case: additional competitors raise the punishment the top bidder faces for failing to escape the randomization region. With more bids potentially falling close to the top bid, the lottery expands and the top bidder's winning probability inside the randomizing region drops further. This lets the auctioneer charge a higher certainty premium. Second, relaxing common knowledge of the mechanism and posing an *inflated second-price* payment rule shows that randomization distorts what advertisers can infer about their rivals from observed outcomes. The distortion goes in opposite directions depending on the advertiser's own value. Combined with inflated pricing, randomization also prevents advertisers from separately identifying rival values and the pricing rule, a transparency concern in its own right.

These technical results feed into an interpretive contribution: the analysis has implications for whether rGSP-style randomization is procompetitive or evidence of market-power abuse, both under the *US v. Microsoft* standard, where rGSP is itself not unlawful but is evidence for monopoly power, and under Article 102 of the Treaty on the Functioning of the European Union (TFEU), where the same conduct maps onto established categories of abuse and could be a violation in itself.

The remainder of the paper proceeds as follows. Section 2 places the paper within the related literature. Section 3 introduces the environment, the optimal mechanism used throughout as benchmark, and defines the randomizing threshold auction together with its DSIC payment rule. Section 4 presents the revenue comparison of rTA and SPA for two bidders. Section 5 turns to welfare. Section 6 extends the mechanism and its payment rule to N bidders. Section 7 presents a short extension that relaxes the common-knowledge assumption. Section 8 draws on the results to interpret the analysis through the lens of the antitrust concerns raised in the Google Search case. Section 9 concludes. All derivations and proofs can be found in the appendix.

2 Related literature

Early analyses of sponsored search auctions study the equilibrium and efficiency properties of the generalized second-price (GSP) position auction (Aggarwal et al., 2008; Edelman et al., 2007; Varian, 2007; Lahaie, 2006; Bae and Kagel, 2019; Athey and Ellison, 2011). GSP allocates ranked slots by bid and charges the position-k winner the bid of position $k+1$, an extension of the second-price logic (Vickrey, 1961) that is not incentive-compatible; for a single slot, GSP collapses to the standard second-price auction, the benchmark used throughout. Although the randomizing auction studied here is a single-unit auction, its design is motivated by rGSP, the randomizing variant of GSP now used by Google.

In Google's GSP-variant, allocations and payments are based on bids, predicted click-through rates and quality measures, summarized in a long-term value (LTV) score. The opaque determination of LTV was itself an antitrust concern: the DOJ describes how Google could artificially boost the runner-up's LTV relative to that of the winner to increase the winner's price, a tactic called *squashing*.[14] Thompson and Leyton-Brown (2013), study how squashing can improve revenue in GSP. In their design squashing entails that the auctioneer strategically throws away some quality information. This can be useful when the winner's and runner-up's LTVs are far apart because the winner has a high quality score. This creates a 'runaway winner', who then pays significantly below their willingness-to-pay. By ignoring quality information, the runner-up's LTV becomes closer to that of the winner, which increases their payment.[15] This is related to the randomization studied in this paper: both squashing and randomization introduce controlled inefficiencies that soften the allocation advantage of strong bidders and can thereby increase revenue when runaway winners occur frequently.

Another related line of work studies how targeting and information disclosure affect market thickness in advertising auctions. Even-Dar et al. (2007) introduce context-dependent auctions and show that incorporating additional information improves allocative efficiency, but may benefit advertisers and the platform asymmetrically. Building on this, Ghosh et al. (2007) show that while running separate auctions for each context maximizes advertiser welfare, it can severely reduce revenue due to a lack of competition within each context, motivating the use of bundling to thicken markets.

This trade-off between improved matching and reduced competition is emphasized more broadly by Levin and Milgrom (2010), who argue that excessive targeting can lead to overly thin markets and advocate "conflation," or the deliberate pooling of similar impressions, to restore competition. Bergemann and Bonatti (2011) show that increased targeting improves match quality but reduces the number of competing advertisers in each market, which can lower equilibrium prices. McAfee et al. (2013) highlight that fine targeting creates many small or "orphan" categories with little demand, creating additional challenges for pricing and allocation.

Rather than studying how much matching information the platform should provide, the present paper takes that heterogeneity as given. It focuses on environments in which most advertisers are poor matches but a small number have substantially higher valuations. The resulting gaps between the top bids are what, in a second-price auction, create runaway

[14] *Pls.' Closing Statement*, pp. 87–88.

[15] In the past, search engines have implemented "squashing" of quality scores. Yahoo! researchers explicitly discussed such practices (Metz, 2010).

winners.

This perspective is closest to Celis et al. (2014), who study a similar match-based environment and likewise take thin markets created by targeting as given. Their buy-it-now or take-a-chance (BIN-TAC) mechanism uses randomization as a screening device: high-value bidders are induced to accept a posted price by the threat of facing a lottery otherwise. Instead, the randomizing threshold auction studied here uses a dominant-strategy incentive-compatible (DSIC) payment rule and has a scale-free allocation rule based on relative bid closeness. This removes the need to calibrate a posted price for each environment and therefore requires less information to implement correctly.

This connect to the literature that studies revenue guarantees when the value distribution is unknown. Bulow and Klemperer (1996) show that a simple second-price auction without reserves achieves a $(n-1)/n$ fraction of optimal revenue, so that competition can substitute for detailed optimization. Subsequent work studies how robust mechanisms can approximate optimal revenue without precise distributional knowledge (Hartline and Roughgarden, 2009; Dhangwatnotai et al., 2015; Devanur et al., 2015; Allouah and Besbes, 2020). Fu et al. (2015) combine this with randomization, and show that randomized mechanisms can strictly outperform the second-price auction *without reserve* in a worst-case analysis. Their mechanism, termed *bid inflation*, randomly inflates the second-highest bid and withholds allocation if the highest bid does not exceed this inflated benchmark.

Finally, Mehta (2022) and Liaw et al. (2023) study how randomized mechanisms affect social welfare when bids are generated by automated bidding systems. In these environments, advertisers maximize value subject to spending constraints, so bids need not reflect true marginal values. In particular, when advertisers face limited opportunities, they may continue spending their budget even when the implied payment exceeds the underlying value of the ad slot. This breaks the usual link between bids and welfare, and can lead to inefficient allocations. Randomization can then improve welfare by weakening the allocation advantage of such aggressively bidding advertisers and reallocating probability towards competitors.

The randomizing threshold auction in this paper is a direct extension of Mehta's design to include a reserve value and, in the extension, multiple bidders. Yet, the economic focus of this paper is more general: it shows when such a randomizing mechanism outperforms the second-price auction even with optimally-behaving bidders. The two settings can be linked. In autobidding environments, occasional budget-driven or poorly targeted bids generate extreme outliers similar to the high-value matches in the bimodal environment studied here. The revenue-based rationale developed in this paper therefore extends naturally to the welfare-based rationale in the autobidding literature.

3 A model of randomizing threshold auctions

This section introduces the model used to study when randomization based on relative bids can outperform the second-price auction in revenue. Section 3.1 describes the environment, motivated by sponsored-search auctions and capturing the advertiser heterogeneity that targeting creates through a bimodal value distribution. This generates the configurations in which large bid gaps and "runaway winners" arise. Section 3.2 introduces the virtual value for this environment and shows that the resulting distribution is irregular, which warrants randomization as a solution. Section 3.3 shows how such randomization can be implemented optimally. Section 3.4 defines the randomizing threshold auction (rTA), a simple and prior-light mechanism that randomizes among bidders when their bids are sufficiently close in relative terms, together with its incentive-compatible payment rule. Throughout the section, the design is directly motivated by institutional evidence on Google's randomized generalized second-price auction (rGSP).

3.1 The sponsored-search environment

Consider a sponsored-search environment, in which online advertisers face a search result from a user for which they can either be a good *match* or not. When a user submits a query to a platform's search bar, the platform auctions off advertisement slots to a pool of potential bidders. To do so, it generates an *impression*: a targeting signal summarizing observable characteristics (e.g. keywords, demographics, location). Advertiser decide their bid profile beforehand: they indicate characteristics they value or explicitly want to exclude to determine their potential bid for the impression.

In choosing how informative the impression is, the platform makes a trade-off. Finer targeting raises the probability that an eligible advertiser infers that their ad is an excellent match for the query. Anticipating a potential click or sale, such a matched advertiser has a higher willingness to pay for the ad space than without the signal. Yet, fine targeting thins the pool of competitors. With less informative targeting signals, unmatched advertisers still value displaying their ads; they cast a wider net in hope of a sale or higher consumer awareness for their brand or product. With targeting too fine, these advertisers focus their marketing budget solely on matched impressions. Platforms therefore do not screen tightly, and the resulting bidder pool spans a wide range of relevance.

To reflect this, we model a platform that auctions off one advertisement slot and faces $N = 2$ bidders. Bidder i's value v_i is private and the two values are independent draws from a common distribution F on $[\ell, H]$. This distribution has two pieces. With probability μ

a bidder is *unmatched* and her value is uniform on the low segment $[\ell, 1]$; with probability $1-\mu$ she is *matched* and her value is uniform on the high segment $(1, H]$. The density is therefore a step function,

$$f(v) = \begin{cases} f_L \equiv \dfrac{\mu}{1-\ell}, & v \in [\ell, 1], \\ f_H \equiv \dfrac{1-\mu}{H-1}, & v \in (1, H], \end{cases} \qquad F(v) = \begin{cases} f_L(v-\ell), & v \in [\ell, 1], \\ \mu + f_H(v-1), & v \in (1, H]. \end{cases} \tag{1}$$

Importantly, we can write the match probabilities as $\mu = f_L(1-\ell), 1-\mu = f_H(H-1)$.

The *floor* parameter $\ell \in [0, 1)$ controls the width of the low segment. The normalization at $v = 1$ is without loss of generality and simply separates low and high types.[16] The parameter $H > 1$ governs the length of the high segment tail. We maintain throughout

$$\tfrac{1}{2} < \mu < 1, \tag{2}$$

which implies most bidders are unmatched.

The interesting configuration is the one in which exactly one advertiser is matched, which happens with probability $2\mu(1-\mu)$. The winner's value then lies above 1 and the runner-up's below it, so the gap between them can be as wide as $H-\ell$, and under a second-price auction the winner pays the runner-up's value however large that gap is. We refer to such a winner that wins by a large margin and pays a low price as *runaway winner*; H governs how extreme they can be.

The record in *US v. Google* speaks to how Google reduced the informativeness of keywords. The court found that, in 2014, Google removed an opt-out from "broader matching" that around 25% of advertisers had been using, while acknowledging that this would "[r]emove[] control from advertisers."[17] As a consequence, the only way for advertisers to prevent a particular query from triggering an ad is to specify a negative keyword, which the court describes as "a far more cumbersome way for advertisers to avoid undesirable auctions".[18] This is in line with the bimodal distribution: a platform that screened tightly would admit only likely matches, so μ would be near zero and F would have a single segment; a large μ is what a broad match envelope looks like in the model.

[16] Revenue equivalence requires a strictly increasing, atomless type distribution without breaks. Otherwise, the marginal-surplus argument that pins down expected payments breaks down and revenue can depend on the tie-breaking rule (Klemperer, 1999, fn. 117, p. 274). The restriction $\ell < 1$ keeps the analysis in the atomless regime throughout, while the common end point at $v = 1$ of the low and high segment ensures that the distribution has no breaks.

[17] *Mem. Op.*, ¶ 277.

[18] *Mem. Op.*, ¶ 278. On the basic keyword-targeting mechanism see also ¶ 184; on the complexity of the matching problem from the advertiser's side, see Jerry Dischler's testimony quoted at ¶ 275.

The skewed distribution with large runaway winners is also in line with the trial record. Adam Juda contrasts "lawyers looking for additional people to sign on to a mesothelioma lawsuit or people offering auto insurance or credit cards [who] are willing to pay hundreds of dollars" with "that small retailer [who] may be willing to pay only a few dollars."[19] At the aggregate revenue level, Michael Whinston testified that "the top 2.5 percent of Google's advertisers make up 90 percent of the revenue. . . really, there's a long tail here."[20] A wide gap between the segments is a large H; a small share of well-matched advertisers is a large μ.

Loosening match quality signals is nonetheless bounded, because showing users too many low-relevance ads damages the value of the ad space itself. Google's internal documents call this "ad blindness," the phenomenon by which users "become blind to these low-quality ads" and "issue fewer searches in the future."[21] The parameter ℓ records this bound. It is the valuation of the worst advertiser the platform is willing to admit, so $\ell > 0$ says that loosening stops somewhere short of admitting everyone. A larger ℓ means even unmatched advertisers retain some relevance, and the cost of randomization landing on one of them is small. The platform therefore admits neither only matched advertisers nor everyone, and (ℓ, μ) move together: tighter screening means larger ℓ, smaller μ (and fewer bidders), and looser screening pushes them the other way. The analysis below treats these parameters as primitives of any single auction. Across auctions, however, they are to some extent platform-design choices.

3.2 Virtual values

Expected revenue of an auction in the environment above is best described in terms of the bidders' *virtual values* (Myerson, 1981). These are most clearly introduced through a monopoly-pricing analogy (Bulow and Roberts, 1989; Bulow and Klemperer, 1996).

Consider a seller facing a single buyer with valuation $v \sim F$ who posts a price p and sells with probability $q(p) = 1 - F(p)$, so expected revenue as a function of quantity is $R(q) \equiv p(q)\, q$ with $p(q) \equiv F^{-1}(1 - q)$ the inverse demand. Differentiating in q gives the marginal revenue and expressing the result in prices gives the virtual value

$$\phi = v - \frac{1 - F(v)}{f(v)}.$$

We apply this to the environment from Section 3.1. Because each segment is uniform

[19] *Trial Tr.*, Doc. 927, p. 1330.

[20] *Trial Tr.*, Doc. 993, p. 8584.

[21] Ex. UPX0010, p. 7, fn. 15 (Google internal long-term-value document); see also *Def.'s Proposed FOF* ¶ 1156 (recounting Juda testimony on the same phenomenon).

and the mass above the break at $v = 1$ is $1 - F(1) = 1 - \mu$, virtual value collapses to a single expression,

$$\phi(v) \; = \; 2v - 1 - \frac{1-\mu}{f(v)} = \begin{cases} 2v - 1 - \kappa, & v \in [\ell, 1], \\ 2v - H, & v \in (1, H], \end{cases} \qquad \kappa = \frac{(1-\mu)(1-\ell)}{\mu}. \tag{3}$$

At $v = 1$ the density drops from f_L to f_H and (3) drops with it, by

$$J \; \equiv \; \phi(1^-) - \phi(1^+) \; = \; (1-\mu)\Big(\frac{1}{f_H} - \frac{1}{f_L}\Big) \; = \; (H-1) - \kappa. \tag{4}$$

So F is *irregular*, meaning that $\phi(v)$ is non-monotonic in v, if and only if $f_L > f_H$: the virtual value falls at $v = 1$ exactly when the low segment is the denser one.

Intuitively, the low marginal revenue of a matched advertiser with $v = 1^+$ reflects what serving her costs elsewhere; the auctioneer can then no longer charge types with value close to H a high price. Matches are relatively rare, so a matched advertiser beats the unmatched ones whatever her own value, and her winning probability is nearly flat across the high segment. A type near H would then happily mimic one just above the discontinuity: she gives up almost no chance of winning and pays much less. Serving those intermediate types therefore caps what can be extracted from the highest ones, and the larger is H the more it caps, which is why $\phi(1^+) = 2 - H$ falls with it.

3.3 The optimal benchmark

When the virtual value is non-decreasing, the revenue-maximizing mechanism allocates to the bidder with the highest marginal revenue whenever that is positive: a second-price auction with a reserve at the zero of the virtual value. Under (3) the optimum departs from that rule in two ways at once: it pools an interval of types into a lottery, and it prices that interval by its two ends.

To see why the optimal mechanism pools types, take an unmatched type $\underline{v} \in [\ell, 1]$ and a matched type $\overline{v} \in (1, H]$ whose marginal revenues are ranked the wrong way, $\phi(\underline{v}) > \phi(\overline{v})$, which (3) allows because ϕ drops at the break. Awarding the object to $\underline{v}$ raises virtual surplus but makes the allocation non-monotone in v, so no incentive-compatible mechanism can do it. The revenue-maximizing mechanism restores incentive compatibility by treating the two types identically: it replaces the revenue curve by its concave envelope, whose derivative is an *ironed* marginal revenue $\bar{\phi}$ that is flat on one interval $[v_L, v_H]$ including the break,

$$\bar{\phi}(v) = \begin{cases} \phi(v), & v < v_L, \\ m, & v \in [v_L, v_H], \\ \phi(v), & v > v_H, \end{cases} \qquad v_L = \tfrac{1}{2}\big(1+\kappa+m\big), \qquad v_H = \tfrac{1}{2}\big(H+m\big). \tag{5}$$

The ironed virtual value is derived in Appendix A.1 and drawn in the second panel of Figure 1. The endpoints follow from $\phi(v_L) = \phi(v_H) = m$ whenever the pool $[v_L, v_H]$ fits inside the support $[\ell, H]$, and the height is then

$$m \;=\; 1 - \frac{1-\mu}{\sqrt{f_L f_H}} \;=\; 1 - \sqrt{\kappa(H-1)}. \tag{6}$$

When it does not fit, $[v_L, v_H] \not\subset [\ell, H]$, which requires $\ell > \frac{1}{2}(1+\kappa)$, the chord runs from the bottom of the support instead: the pool is $[\ell, v_H]$ and the plateau is the average of ϕ over the pool rather than (6).

Figure 1: Marginal revenue and ironing

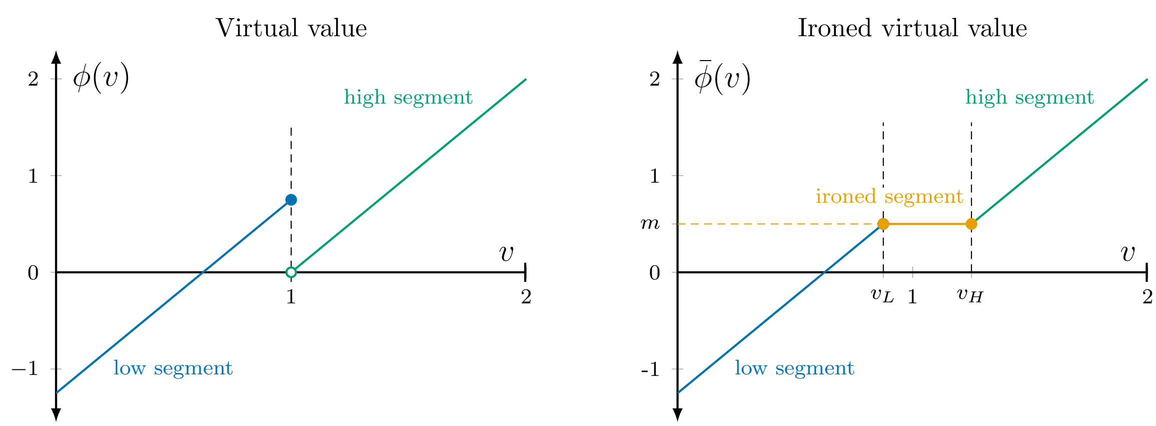

Notes. Marginal revenue $\phi(v)$ for a two-piece uniform mixture $F = \mu\, U[0,1] + (1-\mu)\, U(1,2]$ with $\mu = 0.8$. Marginal revenue rises within each segment but drops discontinuously at $v = 1$, where the low segment ends and the high segment begins, so it is not monotone. The revenue-maximizing mechanism responds by pooling every type on the ironed segment: these types are treated identically and the object is allocated among them by lottery. The right panel shows the resulting curve $\bar{\phi}(v)$, which is flat across the pooled range.

Because $\bar{\phi}$ is constant across the pool, the sign of m settles the reserve and the randomization together. If $m > 0$ the pooled types carry positive marginal revenue and are served: the optimum randomizes among them and posts its reserve in the low segment, at the zero of ϕ there whenever that zero lies inside the segment,

$$\rho_L \;\equiv\; \max\big\{\ell,\; \tfrac{1}{2}(1+\kappa)\big\}, \qquad \text{against} \qquad \rho_H \;\equiv\; \tfrac{1}{2}H$$

in the high segment.[22] If instead $m \le 0$ the pool is worthless, every pooled type is withheld, and the optimum is a second-price auction at ρ_H.

Both conditions on m have a reading that does not mention ironing. A randomizing pool exists at all only if the virtual value really drops at the break, (4). It is worth serving only while a seller facing a *single* buyer would still rather price at the low segment than at the high one. Together they show when the optimal mechanism randomizes.

Lemma 1 (The randomizing window)**.** *The revenue-maximizing mechanism randomizes with positive probability if and only if*

$$1+\kappa \;<\; H \;<\; H_m \;\equiv\; \begin{cases} 1+\frac{1}{\kappa}, & \ell(1+\mu)\le 1, \\ \frac{2\ell}{1-\mu}\left(1+\sqrt{1-\frac{1-\mu}{\ell}}\right), & \ell(1+\mu)>1. \end{cases} \tag{7}$$

Outside this range the optimum is a second-price auction with a reserve, at ρ_L when $H < 1+\kappa$ and at ρ_H otherwise.

Proof. See Appendix A.1. □

For configurations in the randomizing window, there is a unique dominant-strategy incentive compatible (DSIC) payment rule that implements it, and it shows why the mechanism works (Appendix A.2.1). Against a rival outside the pool the winner pays the ordinary second price, or the reserve when the rival falls below it. Against a rival *inside* the pool, a bidder who is also in the pool wins with probability one half and, conditional on winning, pays the bottom of the pool, v_L, which lies above the reserve because $\phi(v_L) = m > 0 = \phi(\rho_L)$; and a bidder who clears the pool outright, $v > v_H$, pays its midpoint, $\frac{1}{2}(v_L + v_H)$.

3.4 Randomized threshold auction

As Section 3.3 showed, the revenue-maximizing mechanism in an irregular environment of this kind does not always award the object to the highest bidder: over a range of types it allocates by lottery instead. Implementing it requires detailed knowledge of the value distribution, including the location and extent of that range. In practice, such information is difficult to estimate reliably in environments like online advertising, where value distributions vary across queries, and small misspecifications in the ironing region can be costly (Roughgarden and Schrijvers, 2016). This motivates the study of simpler, prior-light mechanisms that approximate the effects of ironing.

[22]The unconstrained zero $\frac{1}{2}(1+\kappa)$ exceeds ℓ if and only if $\ell(1+\mu) < 1$. With a higher floor it does not and the reserve is censored at ℓ, since a reserve below the bottom of the support cannot bind.

The *randomizing threshold auction* (rTA) is one such mechanism. It allocates deterministically to the highest bidder whenever her bid exceeds the runner-up's by a factor $\alpha > 1$, and otherwise allocates by a fair lottery between them. With two bidders this is the $\textsc{Rand}(\alpha, 1/2)$ auction of Mehta (2022) adapted to include a reserve value, so that the platform has two instruments rather than one: the factor α and the reserve ρ.

3.4.1 The allocation rule

Write $\mathbf{b} = (b_1, b_2)$ for the vector of bids, where we write b_i for the bid of $i = 1, 2$ and b_{-i} for her rival's. The randomizing threshold auction rTA(α, ρ) has two parameters: a randomization factor $\alpha \geq 1$ and a reserve $\rho \geq \ell$.[23]

Call a profile *close* when both bids clear the reserve and neither exceeds the other by more than a factor α. These profiles form the *randomizing region*

$$\mathcal{R}(\alpha, \rho) \equiv \Big\{ \mathbf{b} : \ \min\{b_1, b_2\} \geq \rho \ \text{ and } \ \max\{b_1, b_2\} \leq \alpha \min\{b_1, b_2\} \Big\}, \tag{8}$$

drawn in Figure 2. The allocation rule is

$$x_i(\mathbf{b}) = \begin{cases} \frac{1}{2}, & \mathbf{b} \in \mathcal{R}(\alpha, \rho), \\ 1, & \mathbf{b} \notin \mathcal{R}(\alpha, \rho), \ \text{ and } \ b_i > \max\{b_{-i}, \rho\}, \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

In words: the slot is withheld when no bid clears the reserve, split by a fair coin when the profile is close, and otherwise awarded to the higher bidder. The rule is scale free, since only the reserve and the ratio of the two bids matter, and it uses no feature of F. At $\alpha = 1$, it is simply the second-price auction with reserve ρ, for which we write SPA$(\rho) =$ rTA$(1, \rho)$.

This one-slot advertisement auction captures important elements of the description of rGSP from the exhibits of the Google Search case, which suggest two main features. First, allocation appears to depend on whether ads are close in ranking score: internal notes explain that rGSP "introduces a probability of swapping between two ads when LTV scores are close."[24] Second, there appears to be a threshold above which the leading bidder wins

[23] A reserve below ℓ binds no one, so $\rho \geq \ell$ restricts notation rather than the platform. Raising ℓ and raising ρ both truncate the bottom of the pool, but they differ in what they leave behind. A higher ℓ replaces the excluded advertisers with better ones, and the auction still has two competing bidders. A higher ρ can leave the slot empty. At $\rho = \ell$ the competitive set is therefore the same as at $\rho = 0$; what changes is payments, since the lowest type is charged her own value when she wins rather than a fraction of the leader's bid, which is what makes her rent vanish.

[24] Ex. UPX1045.

Figure 2: Randomizing region

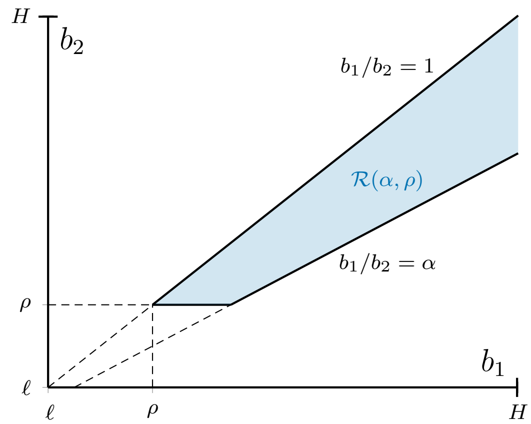

Notes. Graphical depiction of the randomizing region for parameters $\ell = 0.2, H = 2, \alpha = 1.5$. The reserve is set at the lower segment reserve candidate $\rho = \rho_L = 0.6$.

with certainty: as Google's expert witness Dr. Mark Israel put it, "if you are bidding high enough, you don't have to worry about the swap."[25]

Moreover, Google's own materials concede that "[w]e still don't know value distributions,"[26]. A rule based on the absolute difference $b_1 - b_2$ would need a separate cut-off for every query; a rule based on the ratio needs one α for all of them and is *detail-free* in the sense of the *Wilson doctrine* (Wilson, 1987). What the platform can still do is experiment: by running variants across many auctions it tunes α to maximize aggregated revenue. Within any single auction α is fixed and known, so it behaves like the screening choices ℓ and μ: a primitive of that auction, a design choice across auctions.

3.4.2 The payment rule

We look for the payments that make truthful bidding a dominant strategy, so that bids can be read as values. Fix the rival's bid and vary b_i. By (9), raising b_i can only move bidder i into the randomizing region and then past it, so her winning probability is non-decreasing and rises at two bids: she enters the randomizing region at

$$\tau_i \equiv \max\left\{\frac{b_{-i}}{\alpha},\ \rho\right\}, \tag{10}$$

the lowest bid that clears the reserve and keeps her within a factor α of her rival, and she leaves it at αb_{-i}.

[25] *Pls.' Closing Statement*, p. 100.
[26] *Pls.' Closing Statement*, p. 105.

Lemma 2 (Truthful payment rule)**.** *The allocation rule* (9) *is monotone, hence implementable. The unique expected payment rule p such that the randomizing threshold auction* (x, p) *is dominant-strategy incentive compatible is*

$$p_i(\mathbf{b}) = \begin{cases} 0, & x_i(\mathbf{b}) = 0, \\ \frac{1}{2}\,\tau_i, & x_i(\mathbf{b}) = \frac{1}{2}, \\ \frac{1}{2}\big(\tau_i + \alpha\, b_{-i}\big), & x_i(\mathbf{b}) = 1 \quad \textit{and} \quad b_{-i} \geq \rho, \\ \rho, & x_i(\mathbf{b}) = 1 \quad \textit{and} \quad b_{-i} < \rho. \end{cases}$$

Proof. Appendix A.2.2. □

The truthful price jumps exactly when the win probability jumps. Both bidders pay the same entry price, half the critical bid τ_i needed to be in the lottery, and a bidder who pulls clear of her rival pays a surcharge for avoiding the lottery, priced at αb_{-i}. Since bidders are risk neutral the mechanism can be run, without changing expected revenue, by charging the winner of the lottery p_i/x_i, which equals τ_i inside the lottery, and the loser nothing.[27]

Because incentive compatibility holds against every fixed bid of the rival, truthful bidding is weakly dominant, and hence a Bayes–Nash equilibrium under incomplete information. A bidder's best response depends only on her own valuation and the mechanism, never on what she believes about her rival, which is what lets the analysis below evaluate the allocation rule directly at values.

Figure 3 shows the allocation of rTA together with its payment rule. Below αb_{-i} the price on winning sits under the bid of the losing bidder, and that gap is the discount for being in the lottery. After the jump at αb_{-i} the winner pays a *certainty premium*, which raises the price above b_{-i}. It is paid precisely by the bidders who would have been runaway winners.

4 Analysis: when randomizing threshold auctions can beat second-price auctions

This section asks whether distorting the allocation, by randomizing on the basis of relative bids, can raise revenue above the second-price auction in the irregular environment of Sec-

[27]The pricing described in the rGSP record looks closer to an inflated second-price rule. Under any non-DSIC rule bidders shade, so the allocation, and with it revenue, need not coincide with the figures below; those are exact for the DSIC mechanism and a benchmark for its institutional counterpart.

Figure 3: Allocation and payments in the randomizing threshold auction

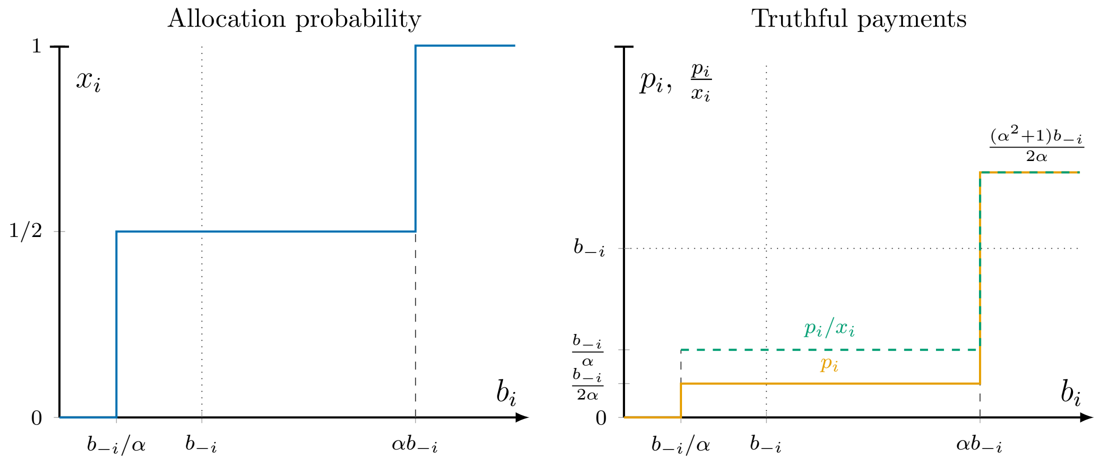

Notes. Allocation probability and truthful payments as a function of bidder i's own bid b_i, holding the rival's bid fixed at $b_{-i} = 2$ with $\alpha = 2.5$ and a reserve that does not bind. Bidder i enters the lottery at b_{-i}/α, where her winning probability jumps from 0 to 1/2, and excludes her rival at αb_{-i}, where it jumps to 1. The right panel plots the expected payment p_i and the price paid on winning p_i/x_i; the latter is undefined where $x_i = 0$ and so begins at b_{-i}/α. The dotted vertical marks b_{-i}, a reference point and not a jump threshold.

tion 3.1. The case is one about runaway winners: the wider the expected gap between the top two valuations, the more a second-price auction leaves on the table. Section 4.1 derives the revenue difference at a low common reserve, Section 4.2 reads off the condition any profitable randomization must meet, and asks what it takes for that condition to hold inside the window of Lemma 1. Section 4.3 then lets both mechanisms set their reserves optimally.

4.1 Revenue difference at the screening reserve

We first consider the revenue difference between the randomizing threshold auction (rTA) and the second-price auction (SPA) at the common reserve $\rho = \ell$. We denote the vector of valuations by $\mathbf{v} = (v_1, v_2)$, and without loss of generality let $v_1 > v_2$. Valuations are i.i.d. draws from the mixture F introduced in Section 3.1.

When we fix a common reserve $\rho \geq \ell$, the lowest type earns zero expected utility under both mechanisms, so both expected revenues equal their expected virtual surplus,[28]

$$R^{\text{rTA}}(\alpha, \rho) = \mathbb{E}\big[\phi(v_1)\, x_1(\mathbf{v}) + \phi(v_2)\, x_2(\mathbf{v})\big], \qquad R^{\text{SPA}}(\rho) = \mathbb{E}\big[\phi(v_1)\mathbf{1}\{v_1 \geq \rho\}\big],$$

where x_i is the allocation rule (9) evaluated at $\mathbf{v}$. The two coincide except when valuations

[28]In general expected revenue equals expected virtual surplus net of the lowest types' rents, $-\sum_i U_i(\underline{v})$ (Myerson, 1981). That term vanishes in both mechanisms studied here: under the SPA the lowest served type wins only if her rival draws below the reserve, an event of probability zero at $\rho = \ell$, and under the rTA she is charged her own value whenever she wins.

are in the randomizing region $\mathbf{v} \in \mathcal{R}(\alpha, \ell)$ in which case rTA allocates to bidder 2 with probability 1/2. Therefore, the expected revenue difference between rTA and SPA is

$$\Delta^{\text{rev}}(\alpha, \rho) \equiv R^{\text{rTA}}(\alpha, \rho) - R^{\text{SPA}}(\rho) = -\tfrac{1}{2}\,\mathbb{E}\left[\big(\phi(v_1) - \phi(v_2)\big) \cdot \mathbf{1}\{\mathcal{R}(\alpha, \rho)\}\right]. \tag{11}$$

We have three possible partitions,

$$\phi(v_1) - \phi(v_2) \;=\; \begin{cases} 2(v_1 - v_2), & LL \text{ and } HH, \\ 2(v_1 - v_2) - J, & LH, \end{cases}$$

where LL, HH indicate the events that both top bidders are in the same segment of the distribution (low or high), i.e. $v_1, v_2 \in [\ell, 1]$ or $v_1, v_2 \in [1, H]$, respectively. Similarly, LH indicates the event that the runner-up bidder is from the low distribution, $v_2 \in [\ell, 1]$, and the highest bidder from the high distribution, $v_1 \in [1, H]$.

Writing for each state $S \in \{LL, LH, HH\}$,

$$P_S(\alpha) \equiv \Pr(S \cap \mathcal{R}(\alpha, \ell)), \qquad M_S(\alpha) \equiv \mathbb{E}\big[(v_1 - v_2)\big|S \cap \mathcal{R}(\alpha, \ell)\big] \cdot P_S(\alpha)$$

for the probability of a randomizing gain and the expected misallocation loss in state S, the lemma below decomposes the expected revenue difference.

Lemma 3 (Revenue difference)**.** *Set $\rho = \ell$ and let $\underline{v} \equiv \max\{\ell, 1/\alpha\}$, then*

$$\Delta^{\text{rev}}(\alpha, \ell) = \tfrac{1}{2} J P_{LH}(\alpha) - \big(M_{LL}(\alpha) + M_{LH}(\alpha) + M_{HH}(\alpha)\big),$$

where

$$P_{LH} = f_L f_H\,(1 - \underline{v})\big(\alpha(1 + \underline{v}) - 2\big), \qquad M_{LL} = \tfrac{1}{3} f_L^2 \Big[(1 - \underline{v})^3 + (\alpha - 1)^2\big(\underline{v}^3 - \ell^3\big)\Big],$$

$$M_{LH} = \tfrac{1}{3} f_L f_H \Big[(\alpha - 1)^2\big(1 - \underline{v}^3\big) - (1 - \underline{v})^3\Big], \qquad M_{HH} = \tfrac{1}{3} f_H^2\, \frac{(\alpha - 1)^2\big(H^3 - \alpha^2\big)}{\alpha^2}.$$

Proof. See Appendix B.1. □

The decomposition shows that the sign of the expected revenue difference depends on whether the virtual-value drop J is large enough and sufficiently likely to be covered by the randomizing region. The formulation also clarifies the role of the floor ℓ on the low segment. Raising ℓ compresses the support of the low values and therefore reduces the valuation gaps in LL states, which directly lowers the cost of randomizing among low types.

The floor proves central in signing the revenue difference, so it is important to understand what it captures. Raising ℓ at a fixed μ is not the screening choice mentioned in Section 3.1, where tighter screening raises ℓ and lowers μ together and, in a setting with more bidders, would shrink the eligible pool. In this model, with higher ℓ, two advertisers always enter the auction, and the only thing that changes is that each drawn low type is worth more in expectation.

4.2 When randomization pays without optimal reserves

In this section, we show when randomization, as applied in rTA, can raise revenue at the suboptimal reserve $\rho = \ell$. Whether the revenue difference can be positive turns on $H - 1$. Widening it raises the drop $J = (H-1) - \kappa$ one for one, and with it the premium a runaway winner can be charged, but it also spreads the high segment out, which makes randomizing among mixed pairs correspondingly rarer if we keep μ and the number of bidders fixed.

For randomizing to make sense, H cannot be too high: by Lemma 1, once it grows past H_m even the revenue-maximizing mechanism gives up on the low segment and posts ρ_H, so H_m caps the values of H at which randomization is relevant. Write

$$\bar{H}(\ell, \mu) \;\equiv\; \inf\big\{H > 1 :\; \Delta^{\text{rev}}(\alpha, \ell) > 0 \;\text{ for some } \alpha\big\} \tag{12}$$

for the smallest tail size at which randomization pays at all. It is worth considering exactly when $\bar{H} < H_m$, and the two limits of the floor ℓ settle whether it is possible at all.

First, setting $\ell = 0$ leaves $\underline{v} = 1/\alpha$, so the advertisers between 0 and $1/\alpha$ are pooled only with each other, and the decomposition from Lemma 3 reads

$$\Delta^{\text{rev}}(\alpha, 0) = \frac{(\alpha-1)^2}{3\alpha^2}\Bigg[\underbrace{\frac{3\alpha}{2}\cdot\frac{\mu(1-\mu)}{H-1}\left(H - \frac{1}{\mu}\right)}_{\text{gain}} - \Bigg(\underbrace{\mu^2}_{LL} + \underbrace{\frac{\mu(1-\mu)}{H-1}\big(\alpha^2 - 1\big)}_{LH} + \underbrace{\frac{(1-\mu)^2}{(H-1)^2}\big(H^3 - \alpha^2\big)}_{HH}\Bigg)\Bigg]. \tag{13}$$

Randomization pays exactly when the bracket is positive. The gain carries the drop $H - 1/\mu$, positive precisely when the environment is irregular, but it is divided by $H - 1$: a longer high segment makes each runaway winner more valuable and makes mixed pairs rarer at the same rate. Hence, the LH loss also decreases in H. The HH loss, on the other hand, keeps growing with H because a lottery between two high types becomes costlier when they are worth more in expectation. The LL loss is the constant μ^2, independent of H.

Yet, with a larger H, there is also more room to widen the scope of the lottery by raising

α, which almost linearly increases the randomizing gain. The platform does not get that room for free; higher α also increasing misallocation losses, and the losses are quadratic. In the bracket, the coefficient on α^2 is $(1-\mu)\big[\mu(H-1)-(1-\mu)\big]/(H-1)^2$, positive exactly when $H > 1/\mu$. At $\ell = 0$ that is both the lower bound of the randomizing window and what makes the drop $H - 1/\mu$ positive in the first place, so wherever randomizing is worth doing at all the losses eventually outrun the gain as the lottery widens. The bracket is a concave quadratic in α and it peaks at

$$\alpha^* \;=\; \tfrac{3}{4}\,(H-1). \tag{14}$$

All that is left is whether the bracket clears zero there.

Proposition 1. *Set $\rho = \ell = 0$. Then $\Delta^{\text{rev}}(\alpha, 0) < 0$ for every $\alpha > 1$ and every H inside the randomizing window* (7)*: without floor, randomization never pays at any H at which the optimum would still randomize, that is $\bar{H}(0,\mu) > H_m$ for every μ.*

Proof. See Appendix B.2.2. □

The reason that randomization never pays in the randomizing window is as follows. Without floor the randomizing band $\{v_1 \le \alpha v_2\}$ narrows to nothing at the bottom of the support, where it pools pairs the optimum would never pool because they sit far below v_L.

Now, consider the other limit of the floor $\ell \to 1$. The admitted pool of unmatched advertiser compresses to a point, and therefore any LL draw leads to a lottery that brings zero misallocation loss. Taking the limit, such that $\underline{v} = \ell \to 1$, in Lemma 3,

$$\lim_{\ell\to 1} \Delta^{\text{rev}}(\alpha,\ell) = \underbrace{\mu(1-\mu)(\alpha-1)}_{\text{gain}} - (\alpha-1)^2\left(\underbrace{\frac{\mu(1-\mu)}{H-1}}_{LH} + \underbrace{\frac{(1-\mu)^2\big(H^3-\alpha^2\big)}{3\alpha^2(H-1)^2}}_{HH}\right). \tag{15}$$

With the admitted unmatched segment collapsed to a point, the relative randomizing band and the interval $[\max\{v_L, \ell\}, v_H]$ agree on the lower bound, and the only mismatch left is on the upper bound. The gain is now first order in the size of the lottery and no longer involves H: a longer high segment increases the virtual value drop and decreases the likelihood of randomizing at exactly offsetting rates. The losses are second order in α; a small lottery therefore always pays, and since the losses eventually dominate, a large one never does.

Proposition 2. *Set $\rho = \ell$ and fix any μ and any gap $H > 1$. In the limit $\ell \to 1$,*

(i) there is $\delta > 0$ such that $\Delta^{\text{rev}}(\alpha,\ell) > 0$ for every $\alpha \in (1, 1+\delta)$, so an arbitrarily small amount of randomization is profitable;

(ii) there is $\varepsilon > 0$ such that $\Delta^{\text{rev}}(\alpha,\ell) < 0$ for every $\alpha \in (H-\varepsilon, H]$, so sufficiently aggressive randomization is unprofitable.

Hence $\bar{H}(\ell,\mu) \to 1$ *as the floor rises to one.*

Proof. See Appendix B.3.2. □

The two limits of ℓ pull in opposite directions and both thresholds move monotonically between them. As the floor rises $\bar{H}$ falls, from above H_m at $\ell = 0$ to one in the limit, while H_m rises with the floor. Therefore, well short of one, there are floors at which $1 < \bar{H} < H_m$: randomization pays, and it pays at gaps the optimum is still willing to randomize at. Figure 4 draws the two ends and one such floor in between, all at $\mu = 0.9$.

Figure 4: Revenue effects of randomization as the floor rises

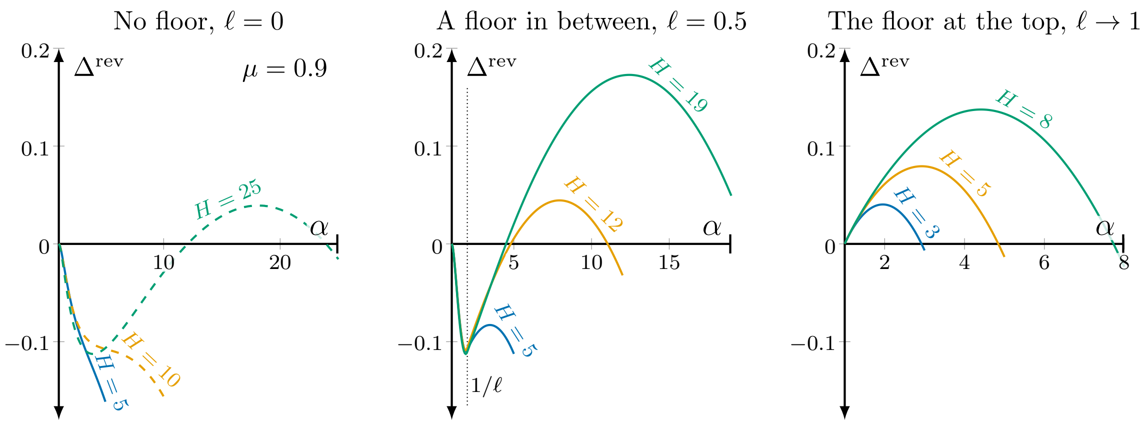

Notes. Expected revenue difference between the randomizing threshold auction and the second-price auction at the screening reserve $\rho = \ell$, as a function of the randomization factor α, with $\mu = 0.9$ fixed and one curve per H. Left panel: no floor, $\ell = 0$, so $\kappa = 1/9$, $H_m = 10$ and $\bar{H} = 21.7$. Middle panel: $\ell = 0.5$, so $\kappa = 1/18$, $H_m = 19$ and $\bar{H} = 9.6$. Right panel: the limit $\ell \to 1$ of (15), where H_m is unbounded and $\bar{H} = 1$. Solid curves lie inside the randomizing window (7); dashed curves do not. The dotted vertical in the middle panel marks $\alpha = 1/\ell$, below which some admitted advertisers can never be pooled with a matched rival.

4.3 Randomization as complement to optimal reserves

So far both mechanisms have posted the screening reserve. This subsection lets each of them choose its own instead, and shows what this does to the revenue difference between the mechanisms.

4.3.1 The second-price auction's optimal reserve

Consider first the optimal reserve price in a second-price auction. The auctioneer does not have randomization in its toolbox and therefore tries to extract revenue from runaway winners solely by setting a smart reserve price. Expected revenue and its derivative with respect to ρ are

$$R^{\text{SPA}}(\rho) = \int_{\rho}^{H} \phi(x)\, 2F(x)f(x)\,\mathrm{d}x, \qquad \frac{\partial R^{\text{SPA}}}{\partial \rho} = -\phi(\rho)\, 2F(\rho)f(\rho), \tag{16}$$

where $2F(x)f(x)$ is the density of the higher of two draws. The derivative has the sign of $-\phi(\rho)$, and on each segment ϕ increases through a single root. So, the second-price auction faces the same two reserve candidates as the optimal benchmark, ρ_L and ρ_H, and which of them is chosen depends on H.

When ρ^L is chosen, bidders from the low segment are sometimes allowed to win, which earns an amount $C(\mu, \ell)$ independent of H.[29] Instead, pricing the unmatched advertisers out by setting ρ^H raises what every matched advertiser pays in expectation, and that is worth more the larger is H. The two balance at

$$\underbrace{C(\mu, \ell) \;+\; \frac{(1-\mu)^2}{3}\, H}_{R^{\text{SPA}}(\rho_L)} \;=\; \underbrace{qH\Big(1 - \frac{q}{3}\Big)}_{R^{\text{SPA}}(\rho_H)}, \qquad q \equiv 1 - F(\rho_H) = \frac{1-\mu}{2}\left(1 + \frac{1}{H-1}\right). \tag{17}$$

Both sides cross once, at H^{SPA}, and SPA posts ρ_L below it and ρ_H above. The crossing is unique because the exclusion branch gains on the low branch as the gap widens, and it comes strictly before the optimal benchmark gives up on the low segment, $H^{\text{SPA}} < H_m$. Appendix B.4 verifies both.

The economics of (17) is that the second-price auction has one instrument for two jobs. The only way it can charge a runaway winner more is to post a very high reserve, which prices out the low segment. An H large enough to make the first worth doing forces it to stop doing the second. The optimal benchmark holds on longer because the lottery lets it collect a certainty premium from matched advertisers, and that premium is useful only because unmatched advertisers are still in the market.

4.3.2 The randomizing threshold auction's optimal reserve

The randomizing threshold auction carries a second instrument, α, and the first thing to say about it is where it does *not* combine with the first, ρ. A reserve above the break $\rho > 1$ leaves no unmatched advertiser in the market, so the randomizing region can hold only two matched advertisers, and pooling those is a pure loss by Lemma 3. Switching the lottery off is then strictly better, and the best the mechanism can do while pricing the low segment out

[29] Explicitly, $C(\mu, \ell) = \rho_L\, G(\rho_L) + (1-\mu) f_L\big(1 - \rho_L^2\big) + 2f_L^2\big[\frac{1}{6} - \frac{1}{2}\rho_L^2 + \frac{1}{3}\rho_L^3\big] + \frac{2}{3}(1-\mu)^2$, where $G(x) = 2F(x)\big(1 - F(x)\big)$ is the probability that exactly one bidder clears x.

is the second-price auction at ρ_H,

$$\max_{\alpha\geq 1,\ \rho>1} R^{\mathrm{rTA}}(\alpha,\rho) \;=\; R^{\mathrm{rTA}}(1,\rho_H) \;=\; R^{\mathrm{SPA}}(\rho_H). \tag{18}$$

Below the discontinuity, the instruments do combine, and they reinforce each other. Compare the two mechanisms at the same reserve $\rho > \ell$: by (16) their marginal returns to raising it differ by

$$I(\rho,\alpha) \equiv \frac{\partial R^{\mathrm{rTA}}(\alpha,\rho)}{\partial\rho} - \frac{\partial R^{\mathrm{SPA}}(\rho)}{\partial\rho} = \frac{\partial\Delta^{\mathrm{rev}}(\alpha,\rho)}{\partial\rho} = f(\rho)\int_{\rho}^{\alpha\rho} \big(\phi(x)-\phi(\rho)\big) f(x)\,\mathrm{d}x. \tag{19}$$

Raising the reserve removes the marginal type ρ from the randomizing region. This benefits rTA because that type is no longer pooled, inefficiently, with the types $[\rho,\alpha\rho]$. The benefit scales in the virtual value difference.

Evaluate this at SPA's own reserve, assuming $\ell(1+\mu)<1$, so that $\rho_L>\ell$ and $\phi(\rho_L)=0$. A factor too narrow to reach across the discontinuity, $\alpha < 1/\rho_L$, keeps the whole region of (19) inside the low segment, where $\phi(x)-\phi(\rho_L) = 2(x-\rho_L) > 0$: the reserve removes pure misallocation and $I>0$ outright. Otherwise, the integral is split at $v=1$. Below the discontinuity the reserve still removes misallocation; above it the rival is matched and the integrand is negative until $\alpha\rho_L = \rho_H$, so the reserve also removes a pool the lottery was profiting from. In this case,

$$I(\rho_L,\alpha) \;=\; f_L f_H\Bigg[\; \underbrace{\frac{H-1}{\kappa}\left(\frac{1-\kappa}{2}\right)^2}_{LL:\ \text{misallocation removed}} \;-\; \underbrace{(\alpha\rho_L-1)\big(H-1-\alpha\rho_L\big)}_{LH:\ \text{pooling gain forgone}} \;\Bigg]. \tag{20}$$

The first term does not involve α: the reserve removes the LL misallocation at the same rate because every LL pairing already falls in the randomizing region. The second is a downward parabola in $\alpha\rho_L$ with roots 1 and $H-1$, so the randomizing factor that costs the most is the one such that $\alpha\rho_L = \rho_H = \frac{1}{2}H$ exactly, where the term equals $\frac{1}{4}(H-2)^2$. Comparing the two, $I(\rho_L,\alpha)>0$ at every factor if and only if

$$(H-1)(1-\kappa)^2 \;>\; \kappa\,(H-2)^2,$$

and that inequality holds on exactly the randomizing window (7): its two roots in H are $1+\kappa$, which is $f_L = f_H$, and $1+1/\kappa$, which is H_m when the reserve is interior. The two instruments are complements precisely where the optimal benchmark randomizes as well.

Write $H^{\mathrm{rTA}}(\alpha)$ for the value of H at which the factor-α rTA stops serving the low segment, and $H^{\mathrm{rTA}} = \sup_\alpha H^{\mathrm{rTA}}(\alpha)$ for the largest at which any factor still serves it.

Lemma 4 (The randomizing auction's reserve). *Let H lie in the randomizing window* (7) *and let SPA's reserve lie strictly above the floor, $\rho_L > \ell$. Then $I(\rho, \alpha) > 0$ for every $\rho \in [\ell, \rho_L]$ and every $\alpha \in (1, H]$, so that*

(i) rTA prices the low segment higher than SPA does,

$$\rho^{\mathrm{rTA}}(\alpha) \;>\; \rho_L \qquad \textit{for every } \alpha > 1;$$

(ii) rTA abandons the low segment no sooner than SPA and no later than the optimum,

$$H^{\mathrm{SPA}} \;\leq\; H^{\mathrm{rTA}} \;\leq\; H_m,$$

the first inequality being strict if and only if some factor strictly beats SPA at H^{SPA}, for which $\bar{H}(\ell, \mu) \leq H^{\mathrm{SPA}}$ is sufficient.

Proof. See Appendix B.5. □

The intuition of the second part is as follows. At H^{SPA} the second-price auction is exactly indifferent between serving the low segment and pricing it out, so any factor that earns a strictly positive return there tips the balance towards serving it. Where the second-price auction has one instrument for two jobs, the randomizing auction has two, and it therefore both prices the low segment higher and holds on to it longer. The exception is where randomization has no return to earn at all: at low μ and a low floor, where by Proposition 1 no factor pays anywhere inside the window, the two mechanisms give up on the low segment at the same H and $H^{\mathrm{rTA}} = H^{\mathrm{SPA}}$. Wherever randomization pays, it also holds on longer.

4.3.3 The revenue difference at optimal reserves

Write $\Delta^*(\alpha, H)$ for the revenue difference when each mechanism posts its own optimal reserve. Below H^{SPA} the two both serve the low segment and differ only in how they allocate among them, so qualitatively Δ^* gives the same results as $\Delta(\alpha, \rho)$, though slightly more favourable to rTA. Above H^{SPA} that is no longer true. The second-price auction has switched to ρ_H, so what rTA must beat is $R^{\mathrm{SPA}}(\rho_H)$, which by (18) is the most either mechanism can earn once the low segment is priced out. Its advantage is then what randomization is worth net of what exclusion of the low segment is worth,

$$\Delta^*(\alpha, H) = \underbrace{\max_{\rho \leq 1} R^{\mathrm{rTA}}(\alpha, \rho) - R^{\mathrm{SPA}}(\rho_L)}_{\text{what randomization is worth}} \;-\; \underbrace{\Big[R^{\mathrm{SPA}}(\rho_H) - R^{\mathrm{SPA}}(\rho_L)\Big]}_{\text{what exclusion is worth}}, \qquad H \geq H^{\mathrm{SPA}}, \quad (21)$$

the last term being zero at H^{SPA} and absent below it.

Proposition 3 (The revenue difference at optimal reserves)**.** *Let H lie in the randomizing window* (7) *and let $\rho_L > \ell$. Then, at every factor $\alpha > 1$,*

(i) letting each mechanism optimize its reserve widens rTA's advantage over what it was at the screening reserve,

$$\Delta^*(\alpha, H) \; > \; \Delta^{\mathrm{rev}}(\alpha, \rho_L) \; > \; \Delta^{\mathrm{rev}}(\alpha, \ell) \qquad \textit{for every } H \leq H^{\mathrm{SPA}};$$

(ii) rTA strictly beats SPA over the whole range on which SPA has abandoned the low segment and rTA has not, and nowhere beyond it,

$$\Delta^*(\alpha, H) > 0 \;\; \textit{for } H \in \big(H^{\mathrm{SPA}}, H^{\mathrm{rTA}}(\alpha)\big), \qquad \Delta^*(\alpha, H) = 0 \;\; \textit{for } H \geq H^{\mathrm{rTA}}(\alpha).$$

Proof. See Appendix B.6. □

Figure 5 draws the logic behind the result. The right panel shows one more thing the two parts of the proposition imply together: Δ^* bends downward at H^{SPA}, because the revenue rTA has to beat switches there from the low branch of (17), which barely improves as H grows, to the high branch, which improves fast. Randomization is therefore worth most to the platform just where SPA is on the verge of abandoning the low segment.

Figure 5: Optimal reserves, the switch, and the revenue difference

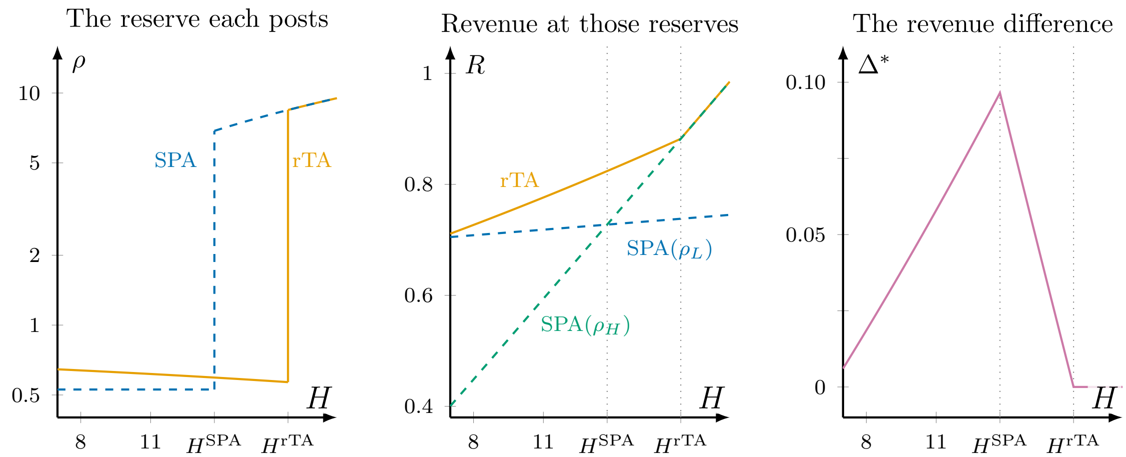

Notes. All panels use $\mu = 0.9$ and $\ell = 0.5$, with the randomization factor optimized at each H; SPA is dashed and rTA solid. The left panel gives the reserve each mechanism posts, on a log scale: SPA's jumps from ρ_L to $\rho_H = \frac{1}{2}H$ at H^{SPA}, while rTA's decreases continuously but stays above ρ_L until H^{rTA}. The middle panel gives the revenue those reserves earn. The two dashed curves are SPA's branches of (17) and SPA chooses the highest. The high-reserve branch rises much faster and overtakes at H^{SPA}. The solid curve is rTA's revenue, which lies above both until H^{rTA} and coincides with SPA(ρ_H) beyond it, where rTA prices the low segment out as well. The right panel is the difference between the two mechanisms at those reserves. It rises up to H^{SPA}, then bends downward as SPA switches branches, and reaches zero at H^{rTA}.

5 Welfare effects of randomizing

This section examines welfare from two complementary angles. Section 5.1 compares aggregate welfare between rTA and SPA, asking what a given revenue gain costs in total surplus when each mechanism posts the reserve that maximizes its own revenue. Section 5.2 then turns to how welfare is distributed across bidder types under rTA, showing that the certainty premium documented in Section 4 corresponds to a systematic redistribution of allocation probability and payments between matched and unmatched advertisers.

Throughout, *welfare* means total expected surplus, which in this private-values setting equals the expected value of the advertiser who is allocated the impression, and zero when a reserve leaves it unsold. This is allocative efficiency; payments are pure transfers and drop out. Section 5.2 adds a distinct margin: because an advertiser's value need not track the relevance of her ad to the user, reallocating impressions can move *user* welfare separately from advertiser surplus.

5.1 Welfare when the second-price auction sets a high reserve

At a fixed reserve, increasing randomization by setting $\alpha > 1$ can only reduce welfare: whenever the mechanism hands the impression to the lower-value bidder, it creates a distortion and expected surplus falls. But a reserve is a distortion too, and one that destroys trade altogether; this lowers welfare by the whole of the winner's value. Randomization misdirects trade and lowers welfare only by the gap between two advertisers' values, which are within a factor α of one another by definition. The first loss is extensive, the second intensive. Therefore, when an auctioneer uses rTA to extract revenue from runaway winners while also choosing a lower reserve value that she would under SPA, this can benefit welfare.

When this can happen is settled by Section 4.3: on the interval between the two thresholds H^{SPA} and H^{rTA}, the second-price auction withholds the impression from every unmatched advertiser while the randomizing threshold auction still serves them, and the welfare comparison for $H \in \left(H^{\mathrm{SPA}}, H^{\mathrm{rTA}}\right]$ is

$$W^{\mathrm{rTA}} - W^{\mathrm{SPA}} = \underbrace{\mathbb{E}\Big[v_{(1)}\,\mathbf{1}\big\{\rho^{\mathrm{rTA}} \le v_{(1)} < \rho_H\big\}\Big]}_{\text{trade the reserve destroys}} - \underbrace{\tfrac{1}{2}\,\mathbb{E}\Big[\big(v_{(1)} - v_{(2)}\big)\,\mathbf{1}\big\{\mathcal{R}(\alpha, \rho^{\mathrm{rTA}})\big\}\Big]}_{\text{trade the lottery misdirects}}. \qquad (22)$$

By Proposition 3 the randomizing auction also earns weakly more revenue on that interval, so whenever (22) is positive the welfare-improvement is not bought with revenue.

Figure 6: Welfare at optimal reserves

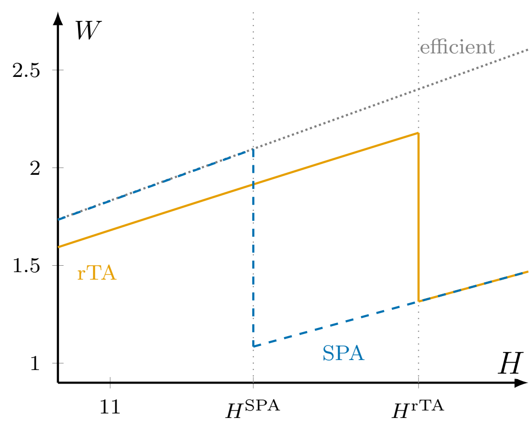

Notes. Expected welfare against H at $\mu = 0.9$ and $\ell = 0.5$, with the reserve, and for the randomizing threshold auction the factor, chosen to maximize each mechanism's own revenue; SPA is dashed and rTA solid. The dotted line is the efficient benchmark $\mathbb{E}[v_{(1)}]$. At H^{SPA} the second-price auction's reserve jumps from ρ_L to $\rho_H = \frac{1}{2}H$ and its welfare falls discontinuously, while its revenue is unchanged. The randomizing threshold auction keeps serving the low segment until H^{rTA}, where its revenue advantage is exhausted and it excludes too. On the interval between the two thresholds it delivers both more revenue and more welfare.

Figure 6 draws the welfare under both mechanisms. At H^{SPA} the second-price auction's welfare falls discontinuously although its revenue does not move at all. From that point onwards, the randomizing threshold auction delivers both more revenue and more welfare, until at H^{rTA} the two mechanism collapse onto one another.

Consider this final threshold H^{rTA}: here the platform is exactly indifferent revenue-wise between randomizing with a lower reserve and using SPA with a high reserve. One of the two options leaves most of the market trading while the other leaves many advertisement slots empty. A platform that cares about returning users and advertisers, will therefore likely choose the first.

5.2 Distributional effects across bidder types

Section 5.1 compared aggregate welfare across mechanisms. We now turn to how welfare is distributed across bidder types within rTA, which sharpens the interpretation of the revenue gains.

Figure 7 gives a more detailed view of the welfare effects of randomization by tracing, for each bidder type, the interim winning probability $X(v_i) \equiv \mathbb{E}_{v_j}\left[x_i(v_i, v_j)\right]$, the expected payment conditional on winning $P(v_i)/X(v_i)$ with $P(v_i) \equiv \mathbb{E}_{v_j}\left[p_i(v_i, v_j)\right]$, and the resulting interim utility $U(v_i) = v_i X(v_i) - P(v_i)$, under SPA and under the randomizing threshold auction at the screening reserve. Closed-form derivations are given in Appendix C.1.

Figure 7: Redistribution effects on advertisers

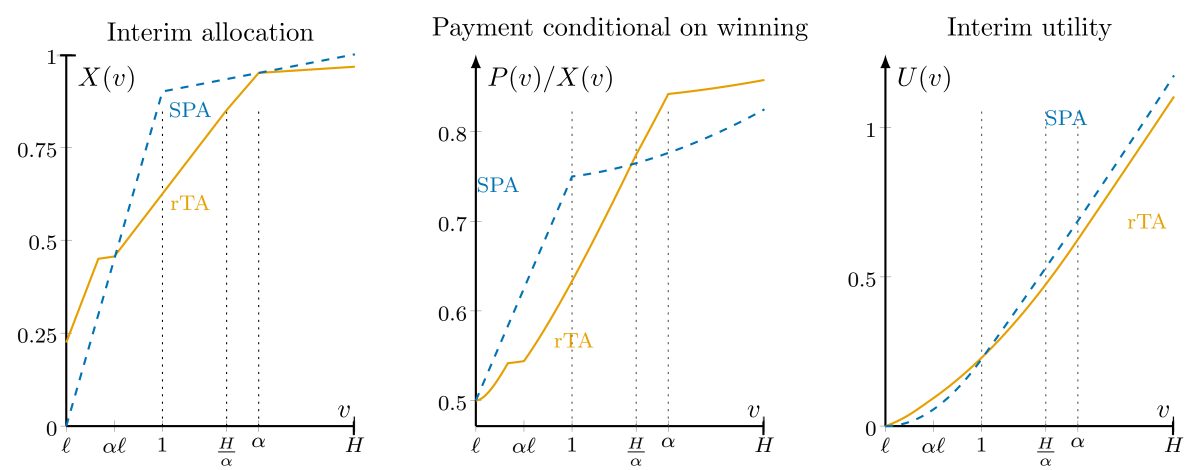

Notes. The figure plots interim outcomes for a bidder of type v in the two-bidder example with $(\alpha, \mu, \ell, H) = \left(\frac{3}{2}, \frac{9}{10}, \frac{1}{2}, 2\right)$ and the screening reserve $\rho = \ell$. The left panel shows the interim winning probability $X(v)$, the middle panel the expected payment conditional on winning $P(v)/X(v)$, and the right panel the interim utility $U(v)$. Dashed lines correspond to the second-price auction and solid lines to the randomizing threshold auction. Vertical dotted lines mark the break at $v = 1$ and the two kinks H/α and α at which a bidder's band reaches the top of the support.

The left panel shows that randomization does not affect all advertisers symmetrically. Very low types can gain winning probability relative to SPA because they now have a chance of obtaining the object when facing a slightly stronger rival. By contrast, strong types lose winning probability for the same exact reason. In this sense, the mechanism weakens the position of *runaway winners*: it compresses the allocation advantage of bidders who are ahead, but do not bid enough to escape randomization altogether.

A key feature of the redistribution is that it are primarily the poorly-matched advertisers that gain interim allocation, while the advertisers with a strong match to the user win less frequently. As a result, even if the allocative inefficiency of randomization is quantitatively small (it occurs among relatively close bids), it is systematically biased towards lower-quality matches. This introduces a qualitatively different welfare margin: randomization can reduce the relevance of displayed ads and therefore lower user welfare.

The middle panel shows the corresponding change in payments. For high types, the expected payment conditional on winning is higher under rTA than under SPA due to the certainty premium. Once a bidder is exposed to randomization when bids are close, securing the object with probability one becomes valuable, and the DSIC payment rule can extract part of that value. By contrast, for lower types, conditional payments are strictly lower under rTA, reflecting that these bidders now win primarily through lottery states where they receive a discount. The figure thus makes clear that the additional revenue generated by rTA is not coming from all advertisers uniformly, but primarily from stronger bidders whose allocation advantage is partially taxed through higher payments: the mechanism extracts revenue from would-be runaway winners.

The right panel combines these two effects into interim utility. The strongest types in the high segment lose utility under rTA: they win less often and, when they do win, they typically pay more. The lower types, by contrast, may gain utility either because their winning probability increases or because their expected payment in case of wins decreases. Randomization therefore has a clear redistributive effect. It transfers surplus away from dominant advertisers whose ads match the searcher's query and towards competitors with poorly-matching ads who would otherwise lose with certainty under SPA.

6 Extension to N bidders

The two-bidder case carries the economics, but the mechanism can be defined for any number of bidders. This section states the general mechanism, derives its truthful payment rule, and asks how the revenue comparison changes as competition thickens.

6.1 The mechanism with N bidders

Formally, write b_i for the bid of bidder $i \in \{1, \dots, N\}$, $b^{(k)}$ for the kth highest bid, and $b_{-i}^{(k)}$ for the kth highest bid among bidder i's rivals. Every bid is compared to a single cut-off,

$$\tau(\mathbf{b}) \equiv \max\left\{\frac{b^{(1)}}{\alpha},\ \rho\right\},$$

where ρ is the reserve value set by the auctioneer, which we take throughout as $\rho = \ell$.

Define the number of competing bidders as

$$K(\mathbf{b}) \equiv \begin{cases} \big|\{\, j :\ b_j \geq \tau(\mathbf{b}) \,\}\big|, & \text{if } b^{(1)} \geq \rho, \\ 0, & \text{if } b^{(1)} < \rho. \end{cases}$$

Thus $K(\mathbf{b})$ is the realized number of competitive bidders, whose bids lie within a factor α of the highest bid and clear the reserve. If $b^{(1)} < \rho$, no bidder meets the reserve and the object is not allocated.

The payment rule will turn on the same cut-off evaluated at bidder i's rivals alone,

$$\tau_{-i} \equiv \max\left\{\frac{b_{-i}^{(1)}}{\alpha},\ \rho\right\}, \qquad K_{-i} \equiv \big|\{\, j \neq i :\ b_j \geq \tau_{-i} \,\}\big|.$$

Unless bidder i is the highest bidder these coincide with the objects above, in the sense that $\tau_{-i} = \tau(\mathbf{b})$ and K_{-i} counts her competitive rivals. When she does lead, $\tau_{-i} \leq \tau(\mathbf{b})$, and it is precisely that gap which lets her exclude rivals by bidding higher.

The randomizing threshold *allocation rule* assigns

$$x_i(\mathbf{b}) = \begin{cases} \dfrac{1}{K(\mathbf{b})}, & b_i \geq \tau(\mathbf{b}), \\ 0, & b_i < \tau(\mathbf{b}). \end{cases} \tag{23}$$

Equivalently, when $b_i > \alpha\, b_{-i}^{(1)}$ bidder i is the only competitive bidder, so $K(\mathbf{b}) = 1$ and she wins with certainty; when $K(\mathbf{b}) \geq 2$, the object is assigned by a fair lottery over the competitive bidders.

Lemma 5 (Truthful payment rule with N bidders)**.** *The allocation rule* (23) *is monotone, hence implementable. The unique expected payment rule p such that the randomizing thresh-*

old auction (x,p) is dominant-strategy incentive compatible is

$$p_i(\mathbf{b}) = \begin{cases} 0, & \text{if } b_i < \tau_{-i}, \\ \\ \underbrace{\dfrac{\tau_{-i}}{K_{-i}+1}}_{\text{entering the lottery}} + \underbrace{\sum_{m=K(\mathbf{b})}^{K_{-i}} \alpha\, b_{-i}^{(m)} \left(\dfrac{1}{m} - \dfrac{1}{m+1}\right)}_{\text{excluding rivals}}, & \text{if } b_i \geq \tau_{-i}, \end{cases}$$

where an empty sum is zero. The second term vanishes unless $b_i > b_{-i}^{(1)}$.

Proof. See Appendix D.1. □

6.2 Revenue with N bidders

We now extend the revenue analysis to $N > 2$ bidders. Throughout, valuations are i.i.d. draws from the mixture distribution F, such that $\mathbf{v} = (v_1, \ldots, v_N)$ is the vector of order statistics with $v_1 > \cdots > v_N$.

Since truthful bidding is dominant, bids equal values and $K(\mathbf{v})$ is the competitive set of Section 6.1 evaluated at $\mathbf{v}$. Let

$$\bar{v}_K(\mathbf{v}) \equiv \frac{1}{K(\mathbf{v})} \sum_{j=1}^{K(\mathbf{v})} v_j$$

denote the average value in the competitive set, and let

$$L(\mathbf{v}) \equiv \sum_{j=1}^{K(\mathbf{v})} \mathbf{1}\{v_j \leq 1\}$$

denote the number of low-segment bidders in the competitive set. By construction, $0 \leq L(\mathbf{v}) \leq K(\mathbf{v})$. When $L(\mathbf{v}) = K(\mathbf{v})$, the mechanism randomizes only among low types; when $L(\mathbf{v}) = 0$, it randomizes only among high types.

Furthermore, define the ironing event as

$$\mathcal{I} \equiv \{1 \leq L(\mathbf{v}) < K(\mathbf{v})\},$$

that is, the event that the competitive set spans both segments of the mixture.

Lemma 6 (Revenue difference decomposition with N bidders)**.** *The expected revenue dif-*

ference can be written as

$$\Delta_N^{\text{rev}}(\alpha,\rho) = J \cdot \Pr\big(\mathcal{I}\big)\,\mathbb{E}\left[\frac{L(\mathbf{v})}{K(\mathbf{v})} \mid \mathcal{I}\right] - 2\,\mathbb{E}\Big[\big(v_1 - \bar{v}_K(\mathbf{v})\big)\,\mathbf{1}\{v_1 \geq \rho\}\Big]. \tag{24}$$

Proof. Under SPA, the object is allocated to the highest value, yielding virtual surplus $\phi(v_1)$. Under the randomized mechanism, expected virtual surplus equals the average among the top $K(\mathbf{v})$ bidders. The revenue difference conditional on the realised profile $\mathbf{v}$ is therefore

$$\Delta_N^{\text{rev}}(\alpha,\rho;\mathbf{v}) = \tfrac{1}{K(\mathbf{v})}\textstyle\sum_{i=1}^{K(\mathbf{v})}\phi(v_i) - \phi(v_1). \tag{25}$$

As for the two-bidder case, we can decompose the revenue difference by using that

$$\phi(v_1) - \phi(v_i) = 2(v_1 - v_i) - \mathbf{1}\{v_i \leq 1 < v_1\} \cdot \Big(H - \tfrac{1-(1-\mu)\ell}{\mu}\Big).$$

Since $\Delta_N^{\text{rev}}(\alpha,\rho) = \mathbb{E}\big[\Delta_N^{\text{rev}}(\alpha,\rho;\mathbf{v})\big]$, taking expectations over (25) permits the decomposition from the lemma. □

This decomposition separates four margins. First, the size of the virtual-value jump increases the revenue difference, but this is unaffected by N.

Second, if the probability that randomizing across that virtual-value jump occurs, increases with N, i.e. $\Pr(\mathcal{I}) \uparrow$, the expected revenue difference also increases. Moving from $N = 2$ to larger N changes the likelihood of ironing in the following ways. The probability that all bidders draw from the low segment, $\Pr(L = K)$, decreases rapidly in N. As a result, pure low-segment type configurations, where randomization only induces misallocation, become less frequent. Moreover, larger N raises the highest low-type order statistic. Conditional on $1 \leq v_1 < \alpha$, this makes it more likely that a low-type bidder lies falls in the competitive set, $\Pr(1 \leq L < K)$, so that the mechanism randomizes across the virtual-value discontinuity. Finally, the probability that the mechanism randomizes only among high types, $\Pr(L = 0)$, increases. When there are more high-type draws, increasing N is likely to generate configurations in which the leading bidder has $v_1 > \alpha$, so that no ironing occurs.

Third, if the intensity of ironing conditional on occurring, $\mathbb{E}[L(\mathbf{v})/K(\mathbf{v}) \mid \mathcal{I}]$, increases in N, so does the expected revenue difference. The low-type share L/K in allocation probability can move in either direction as N rises: additional low-type draws increase it, while additional high-type draws decrease it. Which force dominates depends on μ, which governs the composition of additional draws, and on (ℓ, H, α), which govern how likely it is that extra low-type or high-type draws fall inside the randomizing region. Generally, as μ is large, the ironing intensity tends to increase with N.

Fourth, increasing N tends to increase misallocation losses, measured by the expected average gap between the top value and the competitive set, $\mathbb{E}[(v_1 - \bar{v}_K(\mathbf{v}))]$. As N rises, the size of the competitive set, $K(\mathbf{v})$, increases with high probability, so the allocation is spread over more bidders. This lowers $\bar{v}_K(\mathbf{v})$ and therefore raises the gap $v_1 - \bar{v}_K(\mathbf{v})$.

Figure 8 presents the overall revenue effects of randomization as the number of bidders increases. To connect these patterns to the underlying mechanism, Table 1 decomposes, for selected α, the revenue difference into the three components from (24) that depend on N: the probability of ironing, its conditional intensity, and the misallocation loss.

Figure 8: Revenue effects with multiple bidders

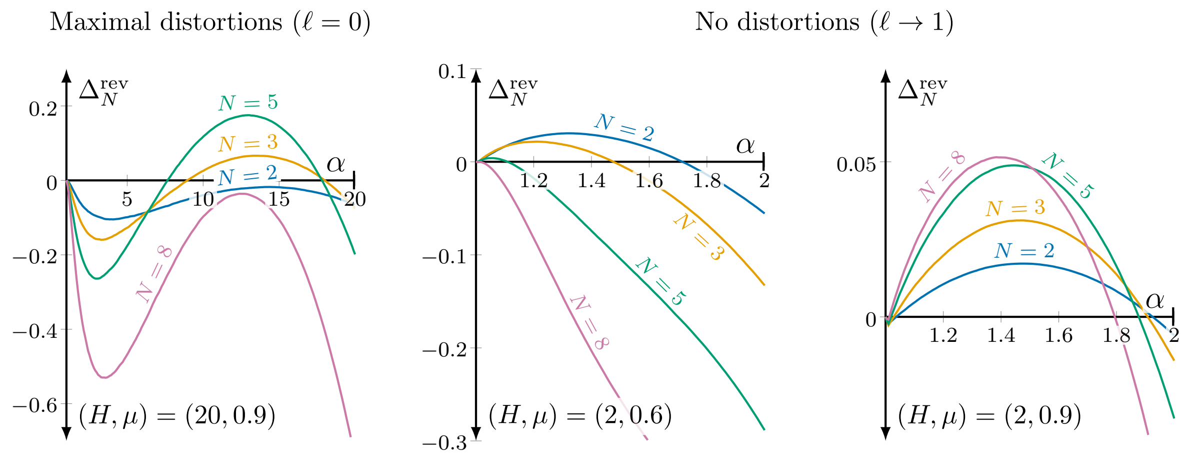

Notes. Expected change in revenue for N bidders.

We begin at $(\ell, H, \mu) = (0, 20, 0.9)$, the parameter point at which the distortion is largest. As N increases from 2 to 3 and 5, the revenue gains from randomization become larger for intermediate values of α. The table shows that this is driven primarily by increases in both the probability of ironing, $\Pr(\mathcal{I})$, and the conditional intensity, $\mathbb{E}[L/K \mid \mathcal{I}]$. In this range, the expansion and strengthening of ironing dominates the rise in misallocation losses . When N increases further to 8, however, the revenue curve shifts downward. The table indicates that while ironing remains frequent, the misallocation term $\mathbb{E}[v_1 - \bar{v}_K]$ grows sufficiently large to outweigh the gains.

The two no-distortion panels illustrate a different mechanism. When $\mu = 0.6$, increasing N reduces revenue for most values of α. The table shows that although ironing probability initially rises, the composition of the competitive set shifts towards high types, which increases misallocation losses and weakens the relative gain from ironing. By contrast, when $\mu = 0.9$, increasing N raises revenue over an intermediate range of α. In this case, additional bidders are mostly low types, which increases both the likelihood and intensity of ironing while keeping misallocation costs limited.

Table 1: Decomposition of revenue effects with N bidders

N	$\mathbb{E}[K(\mathbf{v})]$	$\Pr(L=K)$	$\Pr(L=0)$	$\Pr(\mathcal{I})$	$\mathbb{E}[L/K \mid \mathcal{I}]$	$\mathbb{E}[v_1 - \bar{v}_K]$	Δ_N^{rev}
$\ell = 0,\ H = 20,\ \mu = 0.9,\ \alpha = 12$							
2	1.80	0.81	0.14	0.05	0.50	0.24	-0.03
3	2.50	0.73	0.18	0.09	0.57	0.46	0.06
5	3.62	0.59	0.25	0.16	0.64	0.91	0.16
8	4.80	0.43	0.34	0.23	0.70	1.55	-0.05
$\ell \to 1,\ H = 2,\ \mu = 0.6,\ \alpha = 1.5$							
2	1.73	0.36	0.40	0.24	0.50	0.05	0.02
3	2.27	0.22	0.50	0.28	0.58	0.08	-0.00
5	2.99	0.08	0.68	0.25	0.67	0.13	-0.10
8	3.67	0.02	0.83	0.15	0.71	0.18	-0.25
$\ell \to 1,\ H = 2,\ \mu = 0.9,\ \alpha = 1.5$							
2	1.91	0.81	0.10	0.09	0.50	0.01	0.02
3	2.73	0.73	0.14	0.13	0.65	0.03	0.03
5	4.14	0.59	0.23	0.18	0.78	0.05	0.05
8	5.76	0.43	0.34	0.23	0.85	0.07	0.05

Notes. For each panel and each $N \in \{2,3,5,8\}$, the table reports the average size of the competitive set $\mathbb{E}[K(\mathbf{v})]$, the probability that all randomizing bidders are low types $\Pr(L = K)$, the probability that all randomizing bidders are high types $\Pr(L = 0)$, the probability of ironing, the conditional ironing intensity, the expected misallocation gap $\mathbb{E}[v_1 - \bar{v}_K]$, and the total revenue difference Δ_N^{rev}.

Taken together, the figure and table highlight a common pattern: increasing N tends to expand the scope for ironing, but also enlarges the competitive set, thereby increasing misallocation losses. Whether randomization becomes more profitable therefore depends on which force dominates. In environments with many low types, the expansion of ironing is the key margin; in environments with more high-type competition, the growth in misallocation costs becomes decisive.

7 Extension on pricing and inference

This section relaxes the assumption, maintained in the previous sections, that the mechanism is common knowledge, and asks what happens to bidder inference when the actual auction set-up differs from what advertisers assume. Section 7.1 reconstructs from the trial record a pricing rule that matches what is known of the pricing rule that underlies rGSP. Section 7.2 runs a simple learning exercise showing that an SPA-style inference rule, applied under randomization, produces specific belief distortions in opposite directions across the bidder's own value support. Section 7.3 draws the key implication for policy: under randomization, advertisers can no longer separately identify rival values and the pricing rule from observed outcomes alone.

7.1 The pricing rule

The trial record suggests that rGSP does not simply randomize positions, but also modifies the pricing rule. In particular, testimony by Dr. Kinshuk Jerath describes a mechanism in which the runner-up's score is artificially inflated, so that when the highest-ranked ad is not swapped out, the winning advertiser pays the amount needed to beat this inflated benchmark.[30] The court similarly finds that rGSP both changes the allocation (the runner-up sometimes wins) and "driv[es] up final prices,"[31] and internal Google documents describe its purpose as raising prices "in small increments over time" through what is called "inflation."[32] Mechanically, this resembles an *inflated second-price* rule: conditional on winning, the advertiser pays a scaled version of the runner-up's bid rather than the standard second price. This is closely related to earlier practices such as "squashing," in which the platform manipulates quality scores to bring competitors closer together, thereby increasing payments, as well as to the "bid inflation" mechanisms studied in the theoretical literature (e.g. Fu et al., 2015).

We formalize this institutional rule in a two-bidder setting designed to isolate the inference issue. Writing $b^{(1)}$ and $b^{(2)}$ for the higher and lower of the two bids, we consider the following auction:

- if $b^{(1)} > \alpha b^{(2)}$, the highest bidder wins for sure and pays $\alpha b^{(2)}$;
- otherwise the bids are close and the object is allocated by a fair coin flip: if the highest bidder wins, she pays $b^{(2)}$; if the runner-up wins, she pays $b^{(1)}/\alpha$.

The winner's payment in each case follows a single principle: it scales the loser's bid by one of the factors α, 1, or $1/\alpha$, and the auctioneer applies the largest of these scalings consistent with the winner's participation: the largest that does not exceed the winner's own bid. When the highest bidder wins deterministically ($b^{(1)} > \alpha b^{(2)}$), the inflated runner-up bid $\alpha b^{(2)}$ lies below $b^{(1)}$ and is feasible; the price is set at this ceiling. When the highest bidder wins through the lottery ($\alpha b^{(2)} > b^{(1)} > b^{(2)}$), $\alpha b^{(2)}$ now exceeds her own bid and the rule reverts to the next-largest scaling $b^{(2)}$, which is the standard second price. When the runner-up wins through the lottery ($b^{(2)} < b^{(1)} \leq \alpha b^{(2)}$), both $b^{(1)}$ and $\alpha b^{(1)}$ exceed her own bid $b^{(2)}$, leaving

[30] *Trial Tr.*, Doc. 968, pp. 5491–5492 (Jerath direct examination); see also *Pls.' Proposed FOF*, ¶ 672, which block-quotes the same testimony.

[31] *Mem. Op.*, ¶ 246: "rGSP occasionally randomly switches the LTV scores of the two top auction entrants, thereby allowing the runner-up to win the auction despite its originally lower LTV score. . . Much like squashing, rGSP artificially enhances the runner-up's score, creating more competitive auctions and driving up final prices."

[32] Ex. UPX0059.

the deflated $b^{(1)}/\alpha$ as the largest feasible scaling. The rule therefore applies inflation by factor α whenever the winner's bid permits it, and softens to a smaller scaling otherwise. It produces a symmetric inflation–deflation pattern around the lottery region: the loser's bid is scaled by α when the highest bidder wins deterministically, and by $1/\alpha$ when the runner-up wins the lottery.

7.2 Misspecified inference under randomization

With the rule fixed, we return to Jerath's argument about how rGSP shapes advertiser beliefs. Jerath argues that rGSP can affect advertisers' behaviour through two channels. First, when a high-ranked advertiser is occasionally swapped out, she may infer that her bid is insufficient to secure the position. As he puts it, "a winner had a high enough ad rank but they were flipped randomly to be the loser . . . and they would think that, oh, my bid is not allowing me to win and so they might perceive that I need to bid higher."[33] Second, even when a high-ranked advertiser is *not* swapped out, the inflation of the runner-up's score raises the price she pays, and over repeated auctions she comes to read that elevated price as the going rate for the position, again perceiving that she must bid higher. As Jerath puts it, "even if the winner stays the winner, they're paying a higher price because of the artificial inflation of ad rank . . . over time, they would think that this is sort of the price that gets me the position and would also perceive a higher bid requirement."[34]

This characterization isolates the side of each channel that pushes inferred competitiveness, and therefore bids, up, but the same mechanism produces a counterpart in the opposite direction. First, a runner-up who is randomly swapped to be the winner observes a victory she did not deserve under the underlying ranking, and reads it as evidence that her bid was already high enough to win, a less competitive perception. Second, a runner-up who wins through the lottery pays a deflated price[35] and reads this through inference that assumes second-price pricing as evidence that the rival was lower-valued than under unmodified second-price pricing, again a less competitive perception. The exercise that follows traces both directions through repeated play and characterizes the perceived rival distribution that emerges. Whether the net effect of these distortions pushes beliefs in a more or less competitive direction depends on the bidder's value, not on the mechanism alone.

To trace the joint effect of these channels through repeated play, consider a bidder who observes outcomes from the actual mechanism but incorrectly believes that the mechanism is

[33] *Trial Tr.*, Doc. 968, p. 5493 (Jerath direct examination).

[34] *Trial Tr.*, Doc. 968, p. 5493 (Jerath direct examination).

[35] The runner-up pays $b^{(1)}/\alpha$ in the rule developed above, but even if this is not the actual pricing rule used in rGSP, the runner-up's payment is weakly lower than what the winner would have paid in a SPA.

a truthful second-price auction. In every round, the bidder's own value v and the opponent's value u are independently drawn from the same uniform distribution $U[0,1]$, but they do not know that the other bidder's value follows the same distribution. Therefore, her inferences are based solely on what she observes: her own draw v, whether she wins or loses, and the payment p if she wins. She uses these observations to "learn" the opponent's value distribution. When updating beliefs, she treats every loss as evidence that $u > v$, and every winning payment p as evidence that $u = p$, exactly as in a truthful SPA.

Figure 9 shows the bidder's belief on the opponent's distribution inferred from the repeated interactions. Derivations are provided in Appendix E. This simple learning exercise shows how misinformation on the auction mechanism translates into distorted beliefs about the competitive environment. First, in case of a win, the bidder observes payment information from which she thinks she can infer the value of her competitor. She finds that payments are very often just below $1/\alpha$, so that she assumes probability mass accumulates just below that threshold. Second, when the bidder loses, she perceives to obtain no information on her competitor's value other than that $u > v$. As she faces many hard-to-explain losses from randomizing, which do not disappear if her own value is very high, $v \to 1$, she infers that there must be a mass of "unbeatable" high-value competitors with $u \geq 1$.

Figure 9: Misspecified learning under randomization

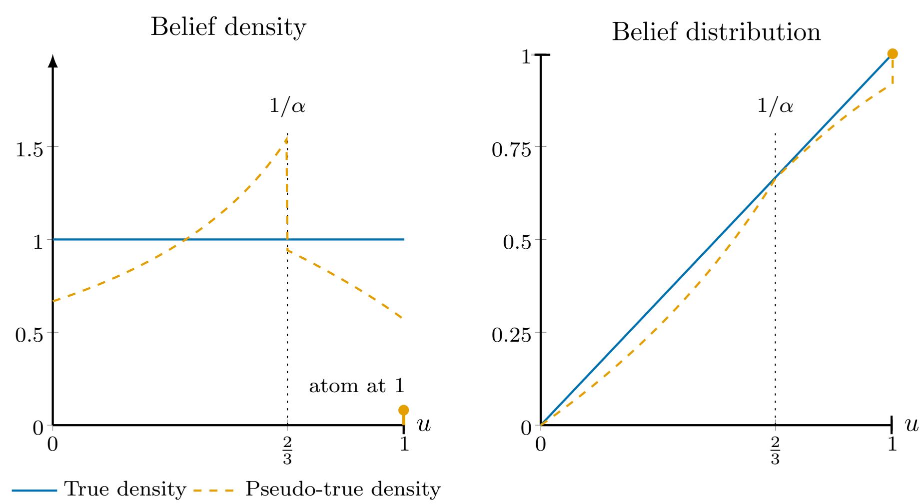

Notes. The true opponent distribution is uniform on $[0,1]$ (blue). When a bidder incorrectly interprets outcomes from the randomizing auction with $\alpha = \frac{3}{2}$ through the lens of a truthful second-price auction, learning converges to a pseudo-true distribution (dashed orange) that overweights values near the threshold $u = 1/\alpha = \frac{2}{3}$ and places a positive atom at the maximal type $u = 1$.

This pattern confirms the flip-side framing developed earlier: the perceived rival distribution is distorted in two qualitatively different ways that show up differently across the bidder's own value support. For intermediate values just below $1/\alpha$, the concentration of perceived mass near $1/\alpha$ shifts rivals upward relative to the true uniform distribution, so the bidder reads rivals as sitting just above her own threshold, the competitive perception Jerath emphasizes. For values in $[1/\alpha, 1]$, by contrast, the dominant distortion is the perceived atom at $u \geq 1$: rivals appear to include a positive mass of "unbeatable" types that no admissible bid would have beaten. This is qualitatively different from a tighter perception of competition over the relevant value range: it produces a phantom upper segment of unreachable types rather than a more competitive read of the rivals the bidder actually faces. The two distortions therefore reshape the perceived rival distribution in different directions across the support. The fact that Google obscures the description of the rGSP mechanism and payment rule does not, in general, imply that the resulting perception of rivals is uniformly more competitive than the truth.

7.3 Identification and transparency

The exercise keeps the bidder's behaviour fixed: she infers, but does not change her bid because also a SPA induces truthful bidding. Tying inference to bidding would require a behavioural rule mapping perceived rival strength into bids. Whatever its form, such a rule would translate inference into action only to the extent that the bidders can infer their rival distribution. Randomization makes this requirement much harder to satisfy. Under a deterministic second-price auction with an unknown squashing factor, which is plausible for the pre-rGSP era of search advertising, the bidder can identify the rival distribution from win/loss frequencies alone, since the allocation rule does not depend on the squashing factor. The factor itself is then identifiable as a residual from observed payments. Randomization breaks this separation: the same observed payment can be rationalized by different combinations of rival value and pricing parameters, and unexplained losses can reflect either randomization or genuinely stronger rivals. Pricing rule and rival distribution become entangled in observed outcomes. This is a transparency concern in its own right: randomization obscures not only the allocation rule but the advertiser's ability to verify either rival strength or pricing from her own auction data.

8 Some implications for antitrust

This section asks what the analysis above implies for antitrust enforcement, and the answer depends on the jurisdiction. Section 8.1 takes the U.S. case, where rGSP is not itself the violation but evidence bearing on whether Google holds monopoly power, and organizes the trial record under the procompetitive and anticompetitive readings. Section 8.2 turns to EU competition law, which can ask a question U.S. law does not: whether the mechanism is itself an abuse.

8.1 rGSP under US antitrust law

Under U.S. antitrust law, monopolization is a violation of Section 2 of the Sherman Act. This means that rGSP itself, as a means of rent extraction, is not the antitrust violation. Instead Google's ability to use rGSP enters the case as evidence on whether Google possesses monopoly power. The *Microsoft* rubric treats the ability to impose conduct that "a firm without a monopoly would have been unable to do" as direct evidence of monopoly power. Google's use of rGSP was accepted by the court in *US v. Google* as evidence of monopoly power. Once established that Google has monopoly power, the antitrust violations concern practices by which Google would have obtained or maintained the established monopoly position, against which, if any, remedies can be pursued. The mechanism itself was not directly outlawed: rGSP was only considered as part of the proof of market power. The two readings below organize the evidence as it was presented at trial.

8.1.1 A procompetitive reading

The procompetitive reading of Google's rGSP, advanced by Google's economic expert witness Mark Israel, is that rGSP served the platform's competitive interest in attracting advertisers and users. If advertisers are viewed as producers and users as the ultimate consumers, randomization can broaden the set of ads shown to users by giving smaller advertisers a positive probability of winning premium positions and stimulating participation by firms that would otherwise rarely appear. In the record, defendants emphasized that rGSP helped to "avoid a winner-takes-all phenomenon" and helped the identification of "new companies" that could offer a good user experience.[36] More broadly, the platform relies heavily on experimentation and learning at scale, using variation in ad positions to refine predictions and better identify advertiser performance:[37] repeated randomization generates information about closely competing advertisers and can improve future matching and ranking decisions.

Israel also disputed the empirical premise of the rent-extraction story. He cited Google's internal revenue data showing that "price goes down for several months after rGSP," arguing that the mechanism did not in fact raise advertiser payments and is therefore inconsistent with an exploitative interpretation.[38]

[36] *Def.'s Proposed FOF*, ¶¶ 1190–1191.
[37] *Pls.' Proposed FOF* ¶¶ 1012–1014.
[38] *Trial Tr.*, Doc. 993, p. 8567 (Israel direct examination).

The welfare analysis of Section 5.1 supports an additional related procompetitive angle. A platform that wants more revenue out of runaway winners can either raise the reserve or it can randomize. The two approaches differ in cost. Excluding poorly-matched advertisers by raising the reserve leaves slots unfilled and pushes advertisers out of the auction altogether. Randomization keeps every advertiser bidding and every slot filled, and moves impressions only between advertisers who were near-tied to begin with. A platform whose product is to connect advertisers to users, and that cares about the long-run health of its ad market, will therefore put some weight on advertiser participation and on what users see. So, the platform prefers to extract revenue using randomization, which aligns with higher surplus for advertisers.

8.1.2 An anticompetitive reading

The anticompetitive reading, adopted by the court in its monopoly-power finding, treats rGSP as a rent-extraction device made possible by monopoly power. The argument requires two parts. First, the mechanism must extract revenue, which it does by pooling the top bidders in a lottery when their bids are close. In this way, randomizing creates a willingness to pay for certain wins and lets the platform charge a "certainty premium." In the language of the Google case, rGSP was one of several pricing knobs used to "extract value more directly" from advertisers and to increase prices.[39] Second, the ability to impose such a mechanism must itself be suggestive of market power. The DOJ filings treat Google's use of pricing knobs to raise prices while paying little attention to rivals as evidence of monopoly power:[40] under genuine competitive pressure, advertisers confronted with opaque, stochastic allocation would spend their advertisement budget elsewhere. Hence, the ability to introduce such rules broadly was taken as direct evidence of monopoly power.

The revenue analysis of Section 4 provides the auction-theoretic foundation for the court's reading: in the irregular environments characteristic of targeted advertising, a simple relative-bid randomization *can* raise revenue over the standard second-price auction. This rules out the claim that no rent-extraction mechanism is available within rGSP-style designs. The court reached the same conclusion, finding that Google has "profitably raised prices substantially above the competitive level."[41]

[39] *Pls.' Proposed FOF* ¶¶ 648–653, 677–680.

[40] *Pls.' Proposed FOF* ¶¶ 677–681.

[41] *Mem. Op.*, p. 259. This was opposed by the defendants, but plaintiffs correctly replied that the relevant comparison is to the counterfactual world without rGSP, in which prices would have fallen further; an internal Google document characterizes rGSP as "a better pricing knob than format pricing" with an "initial impact of 10+% RPM with the current tuning knob" (Ex. DX0153, at slide deck p. 102). The court resolved the empirical dispute against Google.

The welfare analysis of Section 5.2 adds another concern that cuts directly against the procompetitive reading: the redistribution induced by rTA is not random. The bidders who lose surplus are the matched advertisers, whose ads are most relevant to the user; the bidders who gain are the unmatched advertisers, whose ads are less relevant. This cuts against the exploration rationale on which the procompetitive case rests: exploration that systematically reallocates impressions towards inferior matches is directed at inferior alternatives rather than at uncovering high-quality ones, which limits its informational benefit while directly lowering user welfare.

A potentially more serious concern is that under the guise of exploration, randomization may also conceal anticompetitive practices. The Google Search case repeatedly stresses that the platform could use pricing knobs to influence both ranking and prices.[42] This allows randomization to be used with other opaque interventions, including self-preferencing, in ways that are hard to detect. A related concern emerges from the inference analysis of Section 7. Randomization not only obscures the allocation and pricing rules; it also breaks the bidder's ability to identify them separately from observed outcomes. Even an advertiser who participates in many auctions and applies a careful inference rule cannot easily back out either the rival distribution or the pricing parameters when both are filtered through a randomization device.

Mandated transparency requirements on the allocation rule, the pricing rule, or both would directly address this concern. The court's remedies orders have begun to impose such requirements: Google must periodically report all material changes to its Search Text Ads auction, which is the first regulatory acknowledgement that auction opacity is itself a harm that can be remedied rather than an acceptable feature of platform design.[43]

8.2 Implications for EU antitrust law

Under European competition law the mechanism itself can be the violation. Article 102 of the Treaty on the Functioning of the European Union (TFEU) prohibits the abuse of a dominant position, and lists categories of conduct that are themselves unlawful when engaged in by a firm established to be dominant. Where under U.S. law only the "does rGSP show monopoly power?" has relevance, under EU law it is also relevant to ask, given dominance is established: "is the use of rGSP an abuse of dominance?"

The main types of abuse of dominance identified in European case law are exclusionary and exploitative abuses (Akman, 2009). The conduct analysed here is exploitative, and so falls under Article 102(a), which prohibits "directly or indirectly imposing unfair purchase

[42] *Pls.' Proposed FOF* ¶¶ 646, 651–653.

[43] *Final Judgment*, § VI.A, p. 11, implementing the September 2, 2025 Remedies Opinion (Doc. 1435).

or selling prices or other unfair trading conditions." The advertisers who pay more as a direct consequence of randomizing are Google's customers, not its competitors. Whether Google also excludes competitors is a separate matter, and one the Commission has already pursued: in September 2025 it fined Google €2.95 billion for favouring its own exchange in the ad auctions it runs for publishers.[44]

Pursuing rGSP itself as an exploitative abuse would be breaking new ground, and we are not aware of any decision treating the design of an auction mechanism as an exploitative abuse. That gap is less surprising than it looks. Every auction is designed to make money for the seller, and a reserve price already does this in the same way randomizing does: by raising what bidders pay at the cost of an inefficient allocation. In fact, much of the work on auction theory has taken this as its goal, and presumed that the bidders would otherwise bid uncompetitively. However, this literature has always assumed that the seller keeps to the rules it announces (Akbarpour and Li, 2020). Indeed, what is problematic about Google's use of rGSP is not that it raises revenue, but that it does so through a rule advertisers cannot observe or verify, which is the concern of Section 7 and the reason transparency is a more natural remedy than prohibition.

9 Concluding remarks

This paper has studied a simple class of randomizing auctions motivated by Google's rGSP auction for sponsored-search advertising, in which the highest bidder wins with certainty only when her bid sufficiently exceeds the runner-up's, and otherwise faces a lottery among nearby bidders. The mechanism can outperform a second-price auction in revenue by extracting a certainty premium from advertisers with high valuations for the object, *runaway winners*. Because the second-price auction has only its reserve with which to do that job, it sometimes sets this at such a high level that a large share of advertisers can no longer partake in the auction. Instead, the randomizing auction may optimally set a lower reserve in combination with the aforementioned lottery because the threat of that lottery still motivates runaway winners to bid high. Over that range randomization raises revenue and total welfare together.

However, the gains systematically redistribute away from the well-matched advertisers with the highest valuations and, thereby, randomizing also hurts users who prefer advertisements that are good matches to their queries. A learning exercise in which the bidder applies the wrong model further shows that randomization breaks the ability of advertisers to learn separately rival valuations and the pricing rule, which raises a transparency concern for auctions that use randomization without explicitly stating the rule used to determine the

[44]Case AT.40670, *Google – Adtech and Data-related practices.*

winner's payment.

Placing this in the context of antitrust enforcement, the paper theoretically supports the rent-extraction reading of randomizing auctions. In the *US v. Google* case on search advertising, the court indeed found rGSP to be evidence of market power, concluding that Google used it to profitably raise prices. They further imposed remedies to force Google to be more transparent about the mechanics of their search ad auctions. Under EU competition law the mechanism could instead be attacked directly, as an abuse of dominance, though no decision has yet treated an auction design that way.

Several limitations qualify these results, and point to natural extensions. The randomization rule analysed lets all bidders within the randomization region participate in the lottery on equal terms and switches between deterministic allocation and uniform lottery at the threshold. Alternatives in which only the top two bidders randomize, or in which the threshold is smooth or stepwise rather than hard, are equally consistent with the institutional record. This would generate revenue in between SPA and rTA as currently defined. Therefore, the revenue gains identified here can be read as upper bounds on what relative-bid randomization can achieve in this environment.

Furthermore, the match-based environment of Section 3.1 is the cleanest tractable irregular distribution but is only one example, and it abstracts from advertisers' use of automated bidding tools. The N-bidder analysis of Section 6 hints at the direction of the effect of multiple bidders, but does not yet produce sharp analogues of the two-bidder propositions. It could also be extended to a setting with multiple ad slots. The analysis of bidder inference of Section 7 is a stylized learning exercise rather than a comprehensive treatment. The pricing rule introduced there is one of several plausible reconstructions of rGSP from the trial record. Alternatives consistent with the same institutional evidence would yield different inference distortions. A deeper study of the bidder's learning problem under randomizing auctions with non-DSIC pricing, could be a useful next step.

The conclusions from this paper do not only apply to sponsored-search auctions. One setting where the same logic may apply is procurement. A specialized task, say a construction job suited to a few expert firms, can attract a pool of mostly ill-suited bidders alongside occasional strong matches, giving the cost distribution the same irregularity, so relative-bid randomization could discipline a firm whose advantage is large. Whether it should turns on whether the buyer knows in advance which firms are specialized. If it does, the setting is one of *asymmetry*, and directly handicapping the strong bidder dominates randomization. Randomization earns its place only when it is genuinely not known ex ante which bidder will prove strong.

These extensions and potential application of the randomizing threshold auction format suggest a fruitful agenda for further research.

References

Aggarwal, G., Feldman, J., Muthukrishnan, S., and Pál, M. (2008). Sponsored search auctions with Markovian users. In Papadimitriou, C. and Zhang, S., editors, *Internet and Network Economics*, pages 621–628, Berlin, Heidelberg. Springer.

Akbarpour, M. and Li, S. (2020). Credible auctions: A trilemma. *Econometrica*, 88(2):425–467.

Akman, P. (2009). Exploitative abuse in article 82EC: Back to basics? *Cambridge Yearbook of European Legal Studies*, 11.

Allouah, A. and Besbes, O. (2020). Prior-independent optimal auctions. *Management Science*, 66(10):4417–4432.

Athey, S. and Ellison, G. (2011). Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3):1213–1270.

Bae, J. and Kagel, J. H. (2019). An experimental study of the generalized second price auction. *International Journal of Industrial Organization*, 63:44–68.

Bergemann, D. and Bonatti, A. (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3):417–443.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *The American Economic Review*, 86(1):180–194.

Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–1090.

Celis, L. E., Lewis, G., Mobius, M., and Nazerzadeh, H. (2014). Buy-it-now or take-a-chance: Price discrimination through randomized auctions. *Management Science*, 60(12):2927–2948.

Devanur, N. R., Hartline, J. D., and Yan, Q. (2015). Envy freedom and prior-free mechanism design. *Journal of Economic Theory*, 156:103–143.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2015). Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333.

Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259.

Even-Dar, E., Kearns, M., and Wortman, J. (2007). Sponsored search with contexts. In Deng, X. and Graham, F. C., editors, *Internet and Network Economics*, pages 312–317, Berlin, Heidelberg. Springer.

Fu, H., Immorlica, N., Lucier, B., and Strack, P. (2015). Randomization beats second price as a prior-independent auction. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 323–323, Portland Oregon USA. ACM.

Ghosh, A., Nazerzadeh, H., and Sundararajan, M. (2007). Computing optimal bundles for sponsored search. In Deng, X. and Graham, F. C., editors, *Internet and Network Economics*, pages 576–583, Berlin, Heidelberg. Springer.

Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, EC '09, pages 225–234, New York, NY, USA. Association for Computing Machinery.

Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286.

Lahaie, S. (2006). An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the 7th ACM conference on Electronic commerce*, EC '06, pages 218–227, New York, NY, USA. Association for Computing Machinery.

Levin, J. and Milgrom, P. (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review*, 100(2):603–607.

Liaw, C., Mehta, A., and Perlroth, A. (2023). Efficiency of non-truthful auctions in autobidding: The power of randomization. In *Proceedings of the ACM Web Conference 2023*, pages 3561–3571, Austin TX USA. ACM.

McAfee, R. P., Papineni, K., and Vassilvitskii, S. (2013). Maximally representative allocations for guaranteed delivery advertising campaigns. *Review of Economic Design*, 17(2):83–94.

McMillan, J. (1994). Selling spectrum rights. *Journal of Economic Perspectives*, 8(3):145–162.

Mehta, A. (2022). Auction design in an auto-bidding setting: Randomization improves efficiency beyond VCG. In *Proceedings of the ACM Web Conference 2022*, WWW '22, pages 173–181, New York, NY, USA. Association for Computing Machinery.

Metz, C. (2010). Yahoo! 'handicaps' its search ad auctions. Does Google 'squash' too? The Register, `https://www.theregister.com/2010/09/16/yahoo_does_squashing/`.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Roughgarden, T. and Schrijvers, O. (2016). Ironing in the dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 1–18, Maastricht The Netherlands. ACM.

Thompson, D. R. M. and Leyton-Brown, K. (2013). Revenue optimization in the generalized second-price auction. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, EC '13, pages 837–852, New York, NY, USA. Association for Computing Machinery.

Varian, H. R. (2007). Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.

Wilson, R. (1987). Game-theoretic analyses of trading processes. In Bewley, T. F., editor, *Advances in Economic Theory: Fifth World Congress*, pages 33–70. Cambridge University Press.

A Derivations for Section 3

A.1 Proof of Lemma 1

First, it is only worth pooling some types if the virtual value really drops at the break, so $f_L > f_H$, or equivalently, $H > 1 + \kappa$. Second, that pool is worth serving only while $\bar{\phi} > 0$.

Ironing pools a single interval of types $[a, b]$ with $a \leq 1 < b$, assigns them their average virtual value, the *ironed* virtual value,

$$\bar{\phi} = \frac{1}{F(b) - F(a)} \int_a^b \phi(x) f(x) \, \mathrm{d}x,$$

and serves them only if $\bar{\phi} > 0$. Since $\phi(x) f(x) = \frac{\mathrm{d}}{\mathrm{d}x}\big[x\big(F(x) - 1\big)\big]$,

$$\int_a^b \phi(x) f(x) \, \mathrm{d}x = b\big(F(b) - 1\big) - a\big(F(a) - 1\big), \tag{A1}$$

and the ends are pinned by $\phi(a) = \phi(b) = \bar{\phi}$ whenever they are interior, which by (3) gives

$$a = \tfrac{1}{2}\big(\bar{\phi} + 1 + \kappa\big), \qquad b = \tfrac{1}{2}\big(\bar{\phi} + H\big). \tag{A2}$$

The interior case. Suppose $a \geq \ell$, which since $\bar{\phi} > 0$ holds whenever $\ell \leq \frac{1}{2}(1 + \kappa)$, that is $\ell(1 + \mu) \leq 1$. Substituting (A2) into (A1) and solving for $\bar{\phi}$ gives the plateau (6),

$$\bar{\phi} = m = 1 - \sqrt{\kappa(H - 1)},$$

taking the root below one, since the other would put a above the break. The pool is served, $\bar{\phi} > 0$, exactly when $\kappa(H - 1) < 1$, that is $H < 1 + 1/\kappa$.

The censored case. Suppose instead $\ell > \frac{1}{2}(1 + \kappa)$, so that for H large enough (A2) would put $a < \ell$. The pool is then $[\ell, b]$ with b still interior, and $F(\ell) = 0$, so (A1) gives

$$\bar{\phi} = \frac{b\big(F(b) - 1\big) + \ell}{F(b)},$$

which has no closed form in H, but its sign does. Setting $\bar{\phi} = 0$ puts b at $\rho_H = \frac{1}{2}H$ and leaves

$$\ell = \rho_H\big(1 - F(\rho_H)\big) = \frac{(1 - \mu)H^2}{4(H - 1)}.$$

The right side rises in H above $H = 2$, starting from $1 - \mu$, and the censored case requires $\ell > 1/(1+\mu) \geq 1 - \mu$, so it equals ℓ once. Solving gives

$$(1-\mu)H^2 - 4\ell H + 4\ell = 0,$$

and taking the root above two, gives

$$H = \tfrac{2\ell}{1-\mu}\left(1 + \sqrt{1 - \tfrac{1-\mu}{\ell}}\right)$$

and since $\bar{\phi}$ is continuous in H and becomes zero only there, the pool is served exactly below it. □

A.2 Truthful payment rules

Every mechanism in the paper allocates by a step function of a bidder's own bid, so all the payment rules follow from one calculation.

Fix bidder i and hold her rivals' bids fixed, writing $X(b)$ for her winning probability and $P(b)$ for her expected payment as functions of her own bid. Incentive compatibility between two values $v < w$ requires $vX(v) - P(v) \geq vX(w) - P(w)$ and $wX(w) - P(w) \geq wX(v) - P(v)$. Hence,

$$v\big(X(w) - X(v)\big) \;\leq\; P(w) - P(v) \;\leq\; w\big(X(w) - X(v)\big).$$

Now consider two very close bids $w = \tau_k$ and $v = \tau_k - \epsilon$, with $\epsilon > 0$ very small, just on both sides of a jump in winning probability, such that $\Delta X(\tau_k) = X(\tau_k) - X(\tau_k - \epsilon) > 0$. Letting $\epsilon \to 0$, the jump in payments is

$$\Delta P(\tau_k) = \tau_k \, \Delta X(\tau_k)$$

by the squeeze theorem. Furthermore, where X is flat, the equation gives $0 \leq P(w) - P(v) \leq 0$, so P is flat between jumps. And, finally, a bidder who never wins with $X = 0$ pays nothing, so $P = 0$.

Then, repeating this over all points where the win probability jumps yields

$$P(v_i) = \sum_{k \,:\, \tau_k \leq v_i} \tau_k \, \Delta X_k, \tag{A3}$$

which is the payment that guarantees truthfull bidding as a dominant strategy.

A.2.1 The optimal mechanism

Write $[\underline{p}, v_H]$ for the randomizing pool, where $\underline{p} = \max\{\ell, v_L\}$ allows for a floor high enough to censor it. Note $\underline{p} \geq \rho_L$: if $v_L \geq \ell$ then $\phi(v_L) = \bar{\phi} > 0$ puts v_L above the zero of ϕ on the low segment, and otherwise $\frac{1}{2}(1+\kappa) < v_L < \ell$, so $\underline{p} = \ell = \rho_L$. The optimal mechanism awards the object to the higher ironed marginal revenue when that is positive and splits it evenly when the two are equal. Fixing the rival at v_j and raising v_i gives

$$
\begin{array}{lll}
\text{(a) } v_j < \rho_L: & \text{one jump } 0 \to 1 \text{ at } \rho_L, & P = \rho_L; \\
\text{(b) } \rho_L \leq v_j < \underline{p}: & \text{one jump } 0 \to 1 \text{ at } v_j, & P = v_j; \\
\text{(c) } v_j > v_H: & \text{one jump } 0 \to 1 \text{ at } v_j, & P = v_j; \\
\text{(d) } v_j \in [\underline{p}, v_H]: & \text{jumps of } \frac{1}{2} \text{ at } \underline{p} \text{ and at } v_H, & P = \frac{1}{2}\underline{p} \text{ or } \frac{1}{2}\big(\underline{p} + v_H\big),
\end{array}
$$

the payments being (A3) read off the jumps.

A.2.2 Proof of Lemma 2

Fix the rival at b_{-i} and recall the cut-off $\tau_i = \max\{b_{-i}/\alpha, \rho\}$ from (10). The two bids lie within a factor α of one another exactly on $[\tau_i, \alpha b_{-i}]$, so raising b_i from zero gives

$$
\begin{array}{lll}
\text{(a) } b_{-i} < \rho: & \text{one jump } 0 \to 1 \text{ at } \rho, & p_i = \rho; \\
\text{(b) } b_{-i} \geq \rho: & \text{jumps of } \frac{1}{2} \text{ at } \tau_i \text{ and at } \alpha b_{-i}, & p_i = \frac{1}{2}\tau_i \text{ or } \frac{1}{2}\big(\tau_i + \alpha b_{-i}\big),
\end{array}
$$

the payments being (A3) read off the jumps. □

B Derivations and proofs for Section 4

B.1 Derivation revenue difference for two bidders

Throughout this derivation the two mechanisms post the same reserve $\rho \in [\ell, 1]$, and $\alpha \leq H$. Write $t \equiv \max\{\ell, \rho\}$ for the lowest served type; the screening reserve is the case $t = \ell$. Starting from the pointwise difference (11), the expected revenue difference is

$$\Delta^{\text{rev}}(\alpha, \ell) = \tfrac{1}{2} J \cdot P_{LH}(\alpha) - \big(M_{LL}(\alpha) + M_{LH}(\alpha) + M_{HH}(\alpha)\big),$$

where $P_{LH}(\alpha) = \Pr\big(LH \cap \mathcal{R}(\alpha, \rho)\big)$ and, for each state S,

$$M_S(\alpha) := \mathbb{E}\Big[(v_1 - v_2) \cdot \mathbf{1}\big\{S \cap \mathcal{R}(\alpha, \rho)\big\}\Big]$$

is the expected misallocation loss from randomizing in that state.

Because the two valuations are independent, the joint density of the pair is $f(v_1)f(v_2)$, which by (1) takes one of three constant values,

$$f(v_1)f(v_2) = \begin{cases} f_L^2, & LL \text{ states}, \\ f_L f_H, & LH \text{ states}, \\ f_H^2, & HH \text{ states}. \end{cases} \tag{B4}$$

Each of these already carries the probability of its state, so state probabilities never appear separately below. Two identities are used repeatedly,

$$f_L(1-\ell) = \mu, \qquad f_H(H-1) = 1-\mu, \tag{B5}$$

and each double integral below runs over the ordered pair, carrying the factor 2 for the two orderings.

Probability of randomizing gain Since $\alpha \le H$ and $v_1 \ge 1$, the event $\mathcal{R}$ requires $v_2 \ge 1/\alpha$: a low type can only be pooled with a matched rival if the band reaches across $v = 1$. Since it also requires $v_2 \ge t$, only

$$v_2 \in [\underline{v}, 1], \qquad \underline{v} = \max\left\{t, \frac{1}{\alpha}\right\},$$

can lie in the band, and for each such v_2 we must have $v_1 \in [1, \alpha v_2]$. Hence, by (B4),

$$\begin{aligned} P_{LH}(\alpha) &= 2 f_L f_H \int_{\underline{v}}^{1} \int_{1}^{\alpha v_2} \mathrm{d}v_1 \, \mathrm{d}v_2 \\ &= 2 f_L f_H \int_{\underline{v}}^{1} (\alpha v_2 - 1) \, \mathrm{d}v_2 \\ &= f_L f_H \big(\alpha(1-\underline{v}^2) - 2(1-\underline{v})\big) \\ &= f_L f_H (1-\underline{v})\big(\alpha(1+\underline{v}) - 2\big). \end{aligned}$$

When $\alpha t < 1$, so that $\underline{v} = 1/\alpha$, this simplifies to

$$P_{LH}(\alpha) = \frac{f_L f_H (\alpha - 1)^2}{\alpha}.$$

Misallocation losses *(i) Low-low states.* Conditional on LL both values lie in $[\ell, 1]$. A type below $1/\alpha$ reaches only up to αv_2, one above it reaches the break, so splitting at $\underline{v}$ gives

$$\begin{aligned}
M_{LL}(\alpha) &= 2f_L^2\left(\int_t^{\underline{v}}\int_{v_2}^{\alpha v_2}(v_1 - v_2)\,\mathrm{d}v_1\,\mathrm{d}v_2 + \int_{\underline{v}}^1\int_{v_2}^1 (v_1 - v_2)\,\mathrm{d}v_1\,\mathrm{d}v_2\right)\\
&= f_L^2\left(\int_t^{\underline{v}}(\alpha-1)^2 v_2^2\,\mathrm{d}v_2 + \int_{\underline{v}}^1 (1-v_2)^2\,\mathrm{d}v_2\right) = \tfrac{1}{3}f_L^2\Big[(\alpha-1)^2\big(\underline{v}^3 - t^3\big) + (1-\underline{v})^3\Big].
\end{aligned}$$

When $\alpha t \geq 1$ every admitted low type is within reach of the break, $\underline{v} = t$, the first term drops out and this is $\frac{1}{3}f_L^2(1-t)^3$. When $\alpha t < 1$ the types in $[t, 1/\alpha)$ are out of reach, $\underline{v} = 1/\alpha$, and the two terms collect into $f_L^2(\alpha-1)^2\big(1-\alpha^2 t^3\big)/(3\alpha^2)$.

(ii) Low-high states. Over the same region as the gain,

$$\begin{aligned}
M_{LH}(\alpha) &= 2f_L f_H \int_{\underline{v}}^1 \int_1^{\alpha v_2}(v_1 - v_2)\,\mathrm{d}v_1\,\mathrm{d}v_2\\
&= f_L f_H \int_{\underline{v}}^1 \Big((\alpha-1)^2 v_2^2 - (1-v_2)^2\Big)\,\mathrm{d}v_2\\
&= \frac{f_L f_H}{3}\Big((\alpha-1)^2\big(1-\underline{v}^3\big) - \big(1-\underline{v}\big)^3\Big).
\end{aligned}$$

When $\underline{v} = 1/\alpha$ this simplifies to

$$M_{LH}\left(\alpha \,\middle|\, \alpha t < 1\right) = \frac{f_L f_H(\alpha-1)^2(\alpha^2-1)}{3\alpha^2}.$$

(iii) High-high states. Conditional on HH both values lie in $(1, H]$. Since $\alpha \leq H$, randomization occurs when $v_1 \leq \alpha v_2$, and by (B4),

$$\begin{aligned}
M_{HH}(\alpha) &= 2f_H^2\left(\int_1^{H/\alpha}\int_{v_2}^{\alpha v_2}(v_1 - v_2)\,\mathrm{d}v_1\,\mathrm{d}v_2 + \int_{H/\alpha}^H\int_{v_2}^H (v_1 - v_2)\,\mathrm{d}v_1\,\mathrm{d}v_2\right)\\
&= f_H^2\left(\int_1^{H/\alpha}(\alpha-1)^2 v_2^2\,\mathrm{d}v_2 + \int_{H/\alpha}^H (H-v_2)^2\,\mathrm{d}v_2\right)\\
&= \frac{f_H^2}{3}\left[(\alpha-1)^2\left(\frac{H^3}{\alpha^3} - 1\right) + \left(H - \frac{H}{\alpha}\right)^3\right]\\
&= \frac{f_H^2(\alpha-1)^2}{3}\left(\frac{H^3}{\alpha^2} - 1\right).
\end{aligned}$$

(iv) Total. Adding the three losses, $M \equiv M_{LL} + M_{LH} + M_{HH}$ takes two forms. If $\alpha t < 1$,

$$M\left(\alpha \,\middle|\, \alpha t < 1\right) = \frac{(\alpha-1)^2}{3\alpha^2}\Big(f_L^2\big(1-\alpha^2 t^3\big) + f_L f_H\big(\alpha^2-1\big) + f_H^2\big(H^3-\alpha^2\big)\Big), \tag{B6}$$

and if $\alpha t \geq 1$,

$$\begin{aligned} M\left(\alpha \,\middle|\, \alpha t \geq 1\right) = \frac{1}{3}\bigg(& f_L^2(1-t)^3 + f_L f_H\Big((\alpha-1)^2\big(1-t^3\big) - (1-t)^3\Big) \\ & + f_H^2 \frac{(\alpha-1)^2}{\alpha^2}\big(H^3-\alpha^2\big)\bigg). \end{aligned} \tag{B7}$$

Total revenue difference Summing the gain and the losses gives the same two cases. If $\alpha t < 1$,

$$\begin{aligned} \Delta^{\text{rev}}(\alpha,\ell) = \frac{(\alpha-1)^2}{\alpha}\bigg[& \tfrac{1}{2} f_L f_H J - \frac{1}{3\alpha}\Big(f_L^2\big(1-\alpha^2 t^3\big) \\ & + f_L f_H\big(\alpha^2-1\big) + f_H^2\big(H^3-\alpha^2\big)\Big)\bigg]. \end{aligned} \tag{B8}$$

and if $\alpha t \geq 1$,

$$\begin{aligned} \Delta^{\text{rev}}(\alpha,\ell) = & \tfrac{1}{2} f_L f_H (1-t)\big(\alpha(1+t)-2\big) J \\ & - \frac{1}{3}\bigg(f_L^2(1-t)^3 + f_L f_H\Big((\alpha-1)^2\big(1-t^3\big) - (1-t)^3\Big) \\ & + f_H^2 \frac{(\alpha-1)^2}{\alpha^2}\big(H^3-\alpha^2\big)\bigg). \end{aligned} \tag{B9}$$

B.2 No floor ($\ell = 0$)

B.2.1 Revenue difference

At $\ell = \rho = 0$ every factor has $\alpha t < 1$, so $\underline{v} = 1/\alpha$ and (B8) applies. Substituting

$$f_L = \mu, \qquad f_H = \frac{1-\mu}{H-1}, \qquad \kappa = \frac{1-\mu}{\mu}, \qquad J = (H-1) - \kappa = H - \frac{1}{\mu},$$

and using $f_L f_H J = \dfrac{(1-\mu)(\mu H - 1)}{H-1}$, gives (13). Multiplying the bracket by $2(H-1)^2$ and collecting terms,

$$\Delta^{\text{rev}}(\alpha,0) = \frac{(\alpha-1)^2}{6\alpha^2(H-1)^2}\, K(\alpha,H), \tag{B10}$$

with

$$K(\alpha,H) \equiv 3(1-\mu)(\mu H-1)(H-1)\,\alpha - 2\mu^2(H-1)^2 - 2\mu(1-\mu)(H-1)\big(\alpha^2-1\big)$$
$$- 2(1-\mu)^2\big(H^3-\alpha^2\big). \tag{B11}$$

B.2.2 Proof of Proposition 1

Fix $\mu \in (1/2, 1)$, suppose $\ell = \rho = 0$, and let H lie inside the randomizing window (7), which at $\ell = 0$ reads $1/\mu < H < 1/(1-\mu)$. We show that $\Delta^{\text{rev}}(\alpha, 0) < 0$ at every $\alpha \in (1, H]$. Since the first term in (B10) is strictly positive, it suffices to sign K. For fixed H this is a strictly concave quadratic in α, because

$$\frac{\partial^2 K}{\partial \alpha^2} = -4(1-\mu)(\mu H - 1) < 0,$$

the lower bound of the window being exactly $\mu H > 1$. Hence it is enough to evaluate K at its maximizer. The first-order condition yields

$$\alpha^*(H) = \frac{3(H-1)}{4},$$

which is (14) and does not depend on μ. Substituting this value gives, writing $h = H - 1$,

$$K\big(\alpha^*(H), H\big) = \frac{Q(h,\mu)}{8},$$

$$Q(h,\mu) \equiv (1-\mu)(25\mu-16)h^3 - \big(16\mu^2+57(1-\mu)^2\big)h^2 + 16(1-\mu)(4\mu-3)h - 16(1-\mu)^2. \tag{B12}$$

It remains to show that $Q < 0$ throughout the window. The upper bound of the window, $h < 1/\kappa$, says exactly that $s := (1-\mu)h/\mu < 1$, so substituting $h = \mu s/(1-\mu)$ turns (B12) into a cubic in the single variable s on $(0, 1)$. Its constant and quadratic terms are negative for every μ, which already settles the sign whenever the linear term is non-positive, that is whenever $\mu \le \frac{3}{4}$. For $\mu > \frac{3}{4}$ the linear term turns positive, and completing the square reduces the claim to $16\mu^2 + 2\mu - 21 < 0$, which holds on $(0, 1)$ because the left-hand side rises with μ and is still -3 at $\mu = 1$.[45]

Hence $K\big(\alpha^*(H), H\big) = Q/8 < 0$, and by concavity $K(\alpha, H) < 0$ for every $\alpha \in (1, H]$. The prefactor in (B10) is positive, so $\Delta^{\text{rev}}(\alpha, 0) < 0$ at every factor and every gap in the

[45] Writing $(1-\mu)^2 Q = a_3 s^3 + a_2 s^2 + a_1 s + a_0$ with $a_3 = \mu^3(25\mu - 16)$, $a_2 = -\mu^2\big(73\mu^2 - 114\mu + 57\big)$, $a_1 = 16\mu(1-\mu)^2(4\mu - 3)$ and $a_0 = -16(1-\mu)^4$: here $a_0 < 0$, and $a_2 < 0$ because $73\mu^2 - 114\mu + 57 = 73\big(\mu - \frac{57}{73}\big)^2 + \frac{912}{73} > 0$. Since $0 < s < 1$, dropping the cubic term when $a_3 \le 0$ and replacing s^3 by s^2 when $a_3 > 0$ weakly raises the expression and leaves a quadratic whose leading coefficient is negative in either case, after which the two cases in the text follow.

window, as claimed.[46] Since $\bar{H}(0,\mu)$ is by (12) the smallest gap at which some factor pays, it exceeds H_m. $\square$

B.3 Maximal floor ($\ell \to 1$)

B.3.1 Revenue difference

Fix μ, a gap $H > 1$ and a factor $\alpha \in (1, H]$. As $\ell \to 1$ we have $\kappa = (1-\mu)(1-\ell)/\mu \to 0$, so the lower bound of the randomizing window (7) is eventually met, and $\underline{v} = \ell$ once $\ell > 1/\alpha$, so that (B9) applies. Using the identities (B5), the four blocks of the revenue difference are

$$\begin{aligned}
\tfrac{1}{2} JP_{LH} &= \tfrac{1}{2}\big(H - 1 - \kappa\big)\,\mu f_H \big(\alpha(1+\ell) - 2\big) \;\longrightarrow\; \mu(1-\mu)(\alpha - 1),\\
M_{LL} &= \tfrac{1}{3}\mu^2(1-\ell) \;\longrightarrow\; 0,\\
M_{LH} &= \tfrac{1}{3}\mu f_H \Big[(\alpha-1)^2\big(1+\ell+\ell^2\big) - (1-\ell)^2\Big] \;\longrightarrow\; \frac{\mu(1-\mu)(\alpha-1)^2}{H-1},
\end{aligned}$$

while M_{HH} does not involve ℓ. Total misallocation loss is thus

$$M(\alpha) = (\alpha - 1)^2 \left(\frac{\mu(1-\mu)}{H-1} + \frac{(1-\mu)^2(H^3 - \alpha^2)}{3\alpha^2(H-1)^2} \right) \tag{B13}$$

Subtracting this from the gain gives (15).

B.3.2 Proof of Proposition 2

For (i), M from (B13) carries a factor $(\alpha-1)^2$, while the other factor is finite at $\alpha = 1$, so $M(1) = M'(1) = 0$ and

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} \lim_{\ell \to 1} \Delta^{\mathrm{rev}}(\alpha, \ell)\bigg|_{\alpha=1} = \mu(1-\mu) \;>\; 0.$$

The revenue difference vanishes at $\alpha = 1$, so it is strictly positive on some interval $\alpha \in (1, 1+\delta)$.

For (ii), at $\alpha = H$ the second term of M collapses, since $H^3 - H^2 = H^2(H-1)$, and

$$M(H) = (H-1)^2 \left(\frac{\mu(1-\mu)}{H-1} + \frac{(1-\mu)^2}{3(H-1)} \right) = \mu(1-\mu)(H-1) + \frac{(1-\mu)^2(H-1)}{3},$$

[46] At the upper edge of the window, where $H = 1/(1-\mu)$, the same computation gives $K(\alpha^*, H) = -\big(7H^3 - 12H^2 + 35H - 14\big)/(8H)$, and $7H^3 - 12H^2 + 35H - 14$ has derivative $21H^2 - 24H + 35 > 0$ and value 64 at $H = 2$, so it is positive for every $H > 2$.

so the limit equals $-\frac{1}{3}(1-\mu)^2(H-1) < 0$ and, being continuous, stays negative on some $(H-\varepsilon, H]$.

Finally, $\Delta^{\text{rev}}(\alpha, \ell)$ is continuous in ℓ, so at any $\alpha \in (1, 1+\delta)$ it is strictly positive for ℓ close enough to one, and then $\bar{H}(\ell, \mu) \leq H$ by (12). Since $H > 1$ was arbitrary, $\bar{H}(\ell, \mu) \to 1$ as $\ell \to 1$. $\square$

B.4 The second-price auction's reserve

By (16) the derivative of R^{SPA} carries the sign of $-\phi(\rho)$, and by (3) ϕ rises through one root on each segment and drops by $J > 0$ at the break. So R^{SPA} has two local maxima, at ρ_L and at ρ_H, and the second-price auction posts the low one exactly when

$$D(H) \;\equiv\; R^{\text{SPA}}(\rho_L) - R^{\text{SPA}}(\rho_H) \;=\; \int_{\rho_L}^{\rho_H} \phi(x)\, 2F(x) f(x)\, \mathrm{d}x \;\geq\; 0. \tag{B14}$$

Each bidder clears ρ_H with probability q, and a value conditioned to exceed $\frac{1}{2}H$ is uniform on $[\frac{1}{2}H, H]$, so the smaller of two of them has mean $\frac{2}{3}H$:

$$R^{\text{SPA}}(\rho_H) = 2q(1-q)\,\tfrac{1}{2}H + q^2\,\tfrac{2}{3}H = qH\Big(1 - \frac{q}{3}\Big).$$

At the low reserve, split by pair type, writing $G(x) = 2F(x)\big(1 - F(x)\big)$ for the chance that exactly one bidder clears x:

$$R^{\text{SPA}}(\rho_L) = \underbrace{\rho_L\, G(\rho_L)}_{\text{sold at the reserve}} + \underbrace{(1-\mu) f_L\big(1-\rho_L^2\big)}_{LH} + \underbrace{2f_L^2\Big[\tfrac{1}{6} - \tfrac{1}{2}\rho_L^2 + \tfrac{1}{3}\rho_L^3\Big]}_{LL} + \underbrace{(1-\mu)^2\Big[\tfrac{2}{3} + \tfrac{1}{3}H\Big]}_{HH},$$

which is the left side of (17): only the HH block contains the gap, so the low branch is affine in it with slope $\frac{1}{3}(1-\mu)^2$, whether ρ_L is interior or censored at the floor. Differentiating both branches,

$$\begin{aligned} D'(H) &= \frac{(1-\mu)^2}{3} - \frac{\mathrm{d}}{\mathrm{d}H}\Big[qH\Big(1 - \frac{q}{3}\Big)\Big] \\ &= -\,\frac{(1-\mu)\,(H-2)\Big[(1+5\mu)H(H-1) - 2(1-\mu)\Big]}{12\,(H-1)^3} \;<\; 0 \end{aligned} \tag{B15}$$

for every $H > 2$, since $H(H-1) \geq 2$ there. The restriction is empty: at $H \leq 2$ the root $\rho_H = \frac{1}{2}H$ does not lie in the high segment and there is no exclusion candidate. So D has at most one zero, which is the H^{SPA} of (17).

It remains to sign D at the ceiling. By Lemma 1 and $\frac{\mathrm{d}}{\mathrm{d}p}\big[p(1-F(p))\big] = -\phi(p)f(p)$,

$$\int_{\rho_L}^{\rho_H} \phi(x)f(x)\,\mathrm{d}x = 0 \qquad \text{at } H = H_m,$$

which is (B14) without the weight $2F$. That weight is increasing and equals 2μ at the break, while ϕf is positive below the break and negative above it, so it shrinks the positive part of the integral and magnifies the negative part:

$$D(H_m) \;<\; 2\mu\int_{\rho_L}^{\rho_H} \phi(x)f(x)\,\mathrm{d}x \;=\; 0.$$

So the second-price auction has already switched to the exclusion reserve by the time the ceiling is reached: $H^{\mathrm{SPA}} < H_m$.

B.5 Proof of Lemma 4

A reserve $\rho > 1$ leaves only HH pairs in the randomizing region, where $\Delta^{\mathrm{rev}}(\alpha,\rho) = -M_{HH}(\alpha) \leq 0$ by Lemma 3, so setting $\alpha = 1$ is better and (18) holds. Below the break, (11) gives

$$\Delta^{\mathrm{rev}}(\alpha,\rho) = -\iint_{\mathcal{R}(\alpha,\rho)} \big(\phi(v_1)-\phi(v_2)\big)f(v_1)f(v_2)\,\mathrm{d}v_1\,\mathrm{d}v_2$$

over the ordered pair $v_1 > v_2$. Raising the reserve drops the pairs whose lower value is ρ, which is (19), so

$$\Delta^{\mathrm{rev}}(\alpha,\rho_L) - \Delta^{\mathrm{rev}}(\alpha,\ell) = \int_{\ell}^{\rho_L} I(\rho,\alpha)\,\mathrm{d}\rho. \tag{B16}$$

Fix $\rho \leq \rho_L \leq 1$. If $u \leq 1$ then $\phi(x)-\phi(\rho) = 2(x-\rho)$ throughout and $I = f_L^2(u-\rho)^2 > 0$. If $u > 1$, splitting (19) at the break, where the integrand loses J,

$$I(\rho,\alpha) = f_L\Big[f_L(1-\rho)^2 + f_H(u-1)\big(u+1-2\rho-J\big)\Big], \tag{B17}$$

a quadratic in u whose bracket is smallest at $u = \rho+\frac{1}{2}J$, where it equals $f_L(1-\rho)^2 - f_H\big(\rho+\frac{1}{2}J-1\big)^2$. Minimizing over every u rather than only the attainable ones weakens the bound, so $I \geq 0$ at every factor whenever

$$\sqrt{f_L}\,(1-\rho) \;\geq\; \sqrt{f_H}\,\Big|\rho+\tfrac{1}{2}J-1\Big|.$$

The two sides have derivatives $-\sqrt{f_L}$ and $\pm\sqrt{f_H}$, so their difference falls in ρ whenever $f_L > f_H$ and the condition binds at $\rho = \rho_L$. There $1-\rho_L = \frac{1}{2}(1-\kappa)$ and $\rho_L+\frac{1}{2}J-1 = \frac{1}{2}(H-2)$, so with $f_L = (1-\mu)/\kappa$ and $f_H = (1-\mu)/(H-1)$ it becomes

$$\kappa\,(H-2)^2 \;\le\; (1-\kappa)^2(H-1) \qquad\Longleftrightarrow\qquad \kappa H^2 - (1+\kappa)^2 H + (1+\kappa)^2 \;\le\; 0,$$

the quadratic of Appendix A.1 again, with roots $1+\kappa$ and $1+1/\kappa$. It therefore holds strictly on the interior of the window (7), where $I(\rho,\alpha) > 0$ on all of $[\ell, \rho_L]$ and at every factor.

Part (i) follows, since $\phi(\rho_L) = 0$ sets the second-price auction's own derivative in (16) at zero and leaves

$$\left.\frac{\partial R^{\mathrm{rTA}}(\alpha,\rho)}{\partial \rho}\right|_{\rho=\rho_L} = I(\rho_L,\alpha) \;>\; 0,$$

so ρ_L is not a local maximum and $\rho^{\mathrm{rTA}}(\alpha) > \rho_L$.

For part (ii), (18) makes $R^{\mathrm{SPA}}(\rho_H)$ the best either mechanism can do while excluding the low segment, so each serves that segment exactly while it can beat $R^{\mathrm{SPA}}(\rho_H)$, and the second-price auction does so exactly while $D \ge 0$. At $\alpha = 1$ the randomizing auction *is* the second-price auction, so whatever SPA earns while serving the low segment the randomizing auction can match. It therefore keeps serving at least as long, $H^{\mathrm{rTA}} \ge H^{\mathrm{SPA}}$, and because D has the single zero H^{SPA} the inequality is strict exactly when the randomizing auction does strictly better there. By part (i) and (B16) the difference at optimal reserves exceeds the difference at the screening reserve, so $\bar{H}(\ell,\mu) \le H^{\mathrm{SPA}}$ is sufficient for that. The upper bound is Lemma 1: at $H \ge H_m$ no mechanism beats the second-price auction at ρ_H, so $H^{\mathrm{rTA}} \le H_m$. □

B.6 Proof of Proposition 3

While the randomizing auction still serves the low segment, that is for $H < H^{\mathrm{rTA}}(\alpha)$,

$$\Delta^*(\alpha,H) = \max_{\rho\le 1} R^{\mathrm{rTA}}(\alpha,\rho) - \max\left\{R^{\mathrm{SPA}}(\rho_L),\; R^{\mathrm{SPA}}(\rho_H)\right\}.$$

For $H \le H^{\mathrm{SPA}}$ the second term is $R^{\mathrm{SPA}}(\rho_L)$, and by part (i) of Lemma 4 the first exceeds $R^{\mathrm{rTA}}(\alpha,\rho_L)$, so

$$\Delta^*(\alpha,H) \;>\; \Delta^{\mathrm{rev}}(\alpha,\rho_L) \;>\; \Delta^{\mathrm{rev}}(\alpha,\ell),$$

the second inequality being (B16). For $H \in \big(H^{\mathrm{SPA}}, H^{\mathrm{rTA}}(\alpha)\big)$ the second term is $R^{\mathrm{SPA}}(\rho_H)$, which the factor-α auction beats by the definition of $H^{\mathrm{rTA}}(\alpha)$, so $\Delta^* > 0$; at $H \ge H^{\mathrm{rTA}}(\alpha)$ it cannot beat it and both mechanisms earn $R^{\mathrm{SPA}}(\rho_H)$, so $\Delta^* = 0$. □

The kink discussed after the proposition is the switch in that second term. Its right derivative at H^{SPA} exceeds its left derivative by $-D'(H^{\mathrm{SPA}}) > 0$, so Δ^* bends downward there by the same amount.

C Derivations for Section 5

C.1 Interim allocations, payments, and utilities

Fix a mechanism and write $X(v)$, $P(v)$ and $U(v)$ for a bidder's winning probability, expected payment and expected surplus when her own value is v and the rival's is integrated out, so that $U(v) = vX(v) - P(v)$. Both mechanisms are truthful, so by the envelope theorem (see Myerson, 1981) surplus rises with value at the rate at which she wins, $U'(v) = X(v)$, and a bidder at the screening reserve gets nothing. Hence

$$U(v) = \int_{\ell}^{v} X(t)\,\mathrm{d}t,$$

and every interim object follows from the winning probability alone. It is convenient to write $G(v) := \int_{\ell}^{v} F(t)\,\mathrm{d}t$ for the primitive of the value distribution, extended linearly above the support, $G(v) = G(H) + (v - H)$ for $v > H$, since αv can exceed the top.[47]

Under SPA a bidder wins whenever the rival is below her, so $X^{\mathrm{SPA}} = F$ and

$$U^{\mathrm{SPA}}(v) = G(v), \qquad P^{\mathrm{SPA}}(v) = vF(v) - G(v).$$

Under rTA she wins outright when the rival falls below v/α and wins the lottery half the time when the rival lies between v/α and αv, so

$$X^{\mathrm{rTA}}(v) = \tfrac{1}{2}\Big(F(v/\alpha) + F(\alpha v)\Big).$$

Integrating this and changing variables in each term,

$$U^{\mathrm{rTA}}(v) = \frac{1}{2}\left[\alpha G(v/\alpha) + \frac{G(\alpha v) - G(\alpha\ell)}{\alpha}\right], \qquad P^{\mathrm{rTA}}(v) = vX^{\mathrm{rTA}}(v) - U^{\mathrm{rTA}}(v).$$

[47] Integrating (1) twice, $G(v) = \frac{\mu(v-\ell)^2}{2(1-\ell)}$ below the break and $G(v) = \frac{\mu(1-\ell)}{2} + \mu(v-1) + \frac{(1-\mu)(v-1)^2}{2(H-1)}$ above it.

D Derivations for Section 6

D.1 Proof of Lemma 5

Fix bidder i and hold the rivals' bids $\mathbf{b}_{-i}$ fixed, so that the cut-off τ_{-i} and the count K_{-i} are fixed with them. Raising her own bid she first joins the competitive set and then pushes rivals out of it one at a time (Figure 3, left panel), so her winning probability only rises and (A3) applies. The jumps are

$$\begin{aligned}
&\text{(a) } b_i = \tau_{-i}: \qquad 0 \to \tfrac{1}{K_{-i}+1}, \quad \text{contributing } \frac{\tau_{-i}}{K_{-i}+1};\\
&\text{(b) } b_i = \alpha b_{-i}^{(m)}: \quad \tfrac{1}{m+1} \to \tfrac{1}{m}, \quad \text{contributing } \alpha b_{-i}^{(m)}\Big(\frac{1}{m} - \frac{1}{m+1}\Big),
\end{aligned}$$

the second arising once she leads: raising her bid then raises the cut-off $\tau(\mathbf{b}) = \max\{b_i/\alpha, \rho\}$, and the rival ranked m among her own rivals drops out exactly when b_i crosses $\alpha b_{-i}^{(m)}$. Summing the jumps she has crossed, with $K(\mathbf{b})$ the realized size of the competitive set,

$$p_i(\mathbf{b}) = \frac{\tau_{-i}}{K_{-i}+1} + \sum_{m=K(\mathbf{b})}^{K_{-i}} \alpha\, b_{-i}^{(m)} \left(\frac{1}{m} - \frac{1}{m+1}\right),$$

an empty sum being zero. Whenever i does not lead, $K(\mathbf{b}) = K_{-i} + 1$ and the sum is empty, so the display covers both cases and recovers the payment rule of Lemma 5. □

E Derivations for Section 7

E.1 Misspecified learning under a randomizing auction

This appendix derives the belief plotted in Figure 9. The bidder plays the rule of Section 7.1 with factor $\alpha > 1$ against an opponent whose value u, like her own v, is drawn from $U[0,1]$, but reads her outcomes as if the auction were a truthful second-price auction: a loss at v as evidence that $u > v$, a winning payment p as evidence that $u = p$.

The pseudo-true distribution Let H denote her belief about the opponent's value distribution, with density h on $[0,1)$ and survival function $S := 1 - H$. Since her model is wrong, learning does not converge to $U[0,1]$ but to the *pseudo-true* distribution $H^\star$, the one that maximizes the expected log-likelihood of her data under her own second-price model.

Under that likelihood, a losing observation at value v contributes $\log S(v)$ and a winning observation with payment p contributes $\log h(p)$, so the pseudo-true distribution solves

$$H^\star \in \arg\max_H Q(H), \qquad Q(H) := \int_0^1 m_\alpha(p) \log h(p)\, dp + \int_0^1 \ell_\alpha(v) \log S(v)\, dv, \tag{E18}$$

where m_α is the true density of observed winning payments and $\ell_\alpha(v)$ is the true probability of losing when the bidder's value equals v.

A direct calculation gives

$$\ell_\alpha(v) = \begin{cases} 1 - \dfrac{\alpha+\alpha^{-1}}{2}\, v, & 0 \le v \le \alpha^{-1}, \\ \dfrac{1}{2}\Big(1 - \dfrac{v}{\alpha}\Big), & \alpha^{-1} < v \le 1, \end{cases}$$

and

$$m_\alpha(p) = \begin{cases} \dfrac{1-p}{\alpha} + \dfrac{\alpha^2-1}{2}\, p, & 0 \le p \le \alpha^{-1}, \\ (1-p)\Big(\dfrac{1}{\alpha} + \dfrac{1}{2}\Big), & \alpha^{-1} < p \le 1. \end{cases}$$

Hazard-rate representation Write $\lambda := h/S$ for the hazard rate of her belief. Substituting $h = \lambda S$ and $S(v) = \exp\big(-\int_0^v \lambda\big)$ turns (E18) into

$$Q(H) = \int_0^1 m_\alpha(v) \log \lambda(v)\, dv - \int_0^1 C_\alpha(v)\lambda(v)\, dv, \qquad C_\alpha(v) := \int_v^1 \big(m_\alpha(t) + \ell_\alpha(t)\big)\, dt,$$

which is pointwise concave in λ, so the first-order condition gives $\lambda^\star = m_\alpha/C_\alpha$. Evaluating C_α gives

$$C_\alpha(v) = \begin{cases} 1 - \dfrac{\alpha+1}{\alpha}\, v + \dfrac{3+\alpha+\alpha^2-\alpha^3}{4\alpha}\, v^2, & 0 \le v \le 1/\alpha, \\ \dfrac{(1-v)\big(3\alpha+1-(\alpha+3)v\big)}{4\alpha}, & 1/\alpha < v \le 1. \end{cases}$$

and hence the pseudo-true hazard

$$\lambda^\star(v) = \begin{cases} \dfrac{\frac{1}{\alpha} + \frac{\alpha^3-\alpha-2}{2\alpha}\, v}{1 - \frac{\alpha+1}{\alpha}\, v + \frac{3+\alpha+\alpha^2-\alpha^3}{4\alpha}\, v^2}, & 0 \le v \le 1/\alpha, \\ \dfrac{2(\alpha+2)}{3\alpha+1-(\alpha+3)v}, & 1/\alpha < v \le 1. \end{cases} \tag{E19}$$

The atom at the top What matters for the economics is that this hazard stays bounded up to the top of the support: as $v \to 1$ the payment density m_α and the weight C_α vanish at the same linear rate, so their ratio tends to $(\alpha + 2)/(\alpha - 1)$ rather than exploding. A bounded hazard integrates to something finite, so the fitted belief does not use up all its mass on $[0, 1)$,

$$\gamma^\star \;:=\; S^\star(1^-) \;=\; \exp\left(-\int_0^1 \lambda^\star(t)\,\mathrm{d}t\right) \;>\; 0,$$

and $\gamma^\star$ sits as an atom at $v = 1$.[48] Winning payments pin down an interior shape for the opponent's distribution, but the bidder loses more often than that shape can explain under a truthful second-price auction. The atom is how the fitted belief reconciles the two.

[48]Both pieces of (E19) integrate in closed form. Above $1/\alpha$ the hazard is the reciprocal of a linear function, and below it the denominator factors into two linear terms, so a partial-fraction step writes $S^\star$ as a product of powers on each piece.