

Robust Multi-Item Demand Autobidding Auctions

Rigel Galgana, Negin Golrezaei, Yifeng Teng

MIT Operations Research, MIT Sloan School of Management, Google Research


Abstract. Over the last decade, advertising platforms have increasingly adopted autobidding interfaces in which advertisers only submit high level campaign objectives rather than individual, granular bids within each auction. In this paper, we study sponsored shopping auctions, in which advertisers compete for limited advertising slots that the platform displays to a prospective shopper. Our model has two primary features: 1) each advertiser crucially can possess multiple distinct products and can win several slots, and 2) advertisers are a mixture of quasi-linear (QL) utility maximizers and value-maximizers (VM). Under uniform bidding—where each advertiser’s vector of bids is proportional to her valuations—we compare the worst case welfare guarantees of the generalized first-price (GFP), generalized second-price (GSP), and Vickrey–Clarke–Groves (VCG) mechanisms.

Our main finding is that GFP guarantees a 0.4569 fraction of the optimal welfare. In contrast, GSP and VCG can obtain nearly 0 welfare. We further obtain equilibrium-support dependent welfare guarantees, which recover and explain several discrepancies between the welfare bounds achieved in prior literature. Lastly, we surprisingly show that a wide class of natural learning dynamics recovers our welfare guarantees even when equilibria do *not* exist and when valuations are correlated or even adversarial. We verify our results via simulation and show not only superior GFP welfare compared to GSP and VCG, but also revenue.

Key words: autobidding; multi-item demand; position auctions; price of anarchy

1. Introduction

Consider a prospective shopper (user) searching for items on a digital platform, e.g., Google Shopping. The user inputs a query and is shown an assortment of both organic results and sponsored ads (see Figure 1). While the exact way in which these sponsored ads are selected is quite complicated and has evolved over time, most platforms have typically employed auctions to determine which sponsored ads are shown to users. In the 2000s and early 2010s, advertisers would carefully choose the keywords they would bid for within a user’s query, maintained separate bids across those keywords, and estimated click rates and conversions (Edelman et al. 2007). This gave advertisers fine-grained control but required substantial operational effort, and a predefined keyword list could miss context or be incompatible with modern conversational or natural-language requests such as through LLM-assisted product search. Understandably, tailoring one’s desired keywords and corresponding bids to specific demographics, contexts, or conversations would be logistically infeasible.

Figure 1 Sponsored shopping example.

Note. A user looking for dress shoes types the query “men’s dress shoes”. Google Shopping displays both sponsored shopping results (top row) and organic results (bottom row). In the brief moment between the user sending their query and seeing their ads, Google quickly runs an auction between advertisers who wish to display possibly multiple products. For example, the advertisers ‘Cole Haan’ and ‘Amberjack’ have multiple sponsored ads.

In response to this increasing advertiser overhead, Google introduced Smart Bidding in 2016—the first modern autobidding system. At a high level, autobidding is a system in which advertisers only specify high level campaign objectives (e.g., a return-on-investment (ROI) target, a budget, duration), while the platform performs both the analytics and the exact bidding strategy. Specifically, the platform would supply query context, estimate click and conversion rates, and adjust for trends in recent campaign performance in order to implement agents that bid on behalf of the advertiser (Aggarwal et al. 2024, 2019, Balseiro et al. 2021). While this removes the necessity of advertisers having to perform much of the challenging analytics, the simpler interface still requires strategic care as a small number of parameters determine the behavior of the platform’s representative autobidding agent and consequently, the success of the advertising campaign. Indeed, prior literature has already shown how truthful specification of campaign objectives to an automated agent or platform need not be optimal (Alimohammadi et al. 2023, Feng et al. 2024, Kolumbus and Nisan 2022a,b). Several other works, as we discuss later, also show how such strategic behavior can drastically reduce welfare—i.e., worsen the quality or relevance of the ads shown to users (Athey and Ellison 2011, Chen and He 2011, Google 2026). This leads us to our main research question.

Which autobidding mechanisms are robust to strategic advertiser behavior? Specifically, how should ad allocations and corresponding payments be assigned in order to ensure high welfare?

In this paper, we study specifically *sponsored shopping* auctions where advertisers compete for limited ad slots on the page shown in response to a user’s query. In contrast to classic sponsored-search auctions in which each advertiser only possesses a single ad or there is only a single slot, an advertiser may wish to promote several products from her catalog and may win several slots. The main goal of our paper is to characterize the consequences of strategic behavior in these autobidding auctions.

Model overview. There are N advertisers and M slots with *click-through rates* (CTRs)—the probabilities of engagement with the respective slots—satisfying $1 = c_1 \geq \dots \geq c_M \geq 0$. While this is the typical setting of standard position auctions (Varian 2007), our model has two main differentiating features: *multi-item demand* and *mixed advertiser objectives*.

- (i) *Multi-item demand.* Advertiser n owns *multiple* ads with valuations $v_{n,1} \geq \dots \geq v_{n,M} \geq 0$ representing varying quality or user fit. Following the uniform bidding model of Dütting et al. (2026), each advertiser submits one multiplier $\alpha_n \geq 0$ which sets the bid corresponding to advertiser n ’s j th ad to be $b_{n,j} \doteq \alpha_n v_{n,j}$. That is, this multiplier simultaneously governs all of an advertiser’s bids. Allocation is done in decreasing order of bids, and importantly, advertisers can win *multiple* slots. Note that this is different from the aforementioned position auction setting with *unit-demand*—where each advertiser instead has one product.
- (ii) *Mixed advertiser objectives.* We consider auctions in which two types of advertisers with different objectives may be present. A quasi-linear (QL) advertiser maximizes value minus payment, while a value-maximizing (VM) advertiser maximizes the CTR-weighted value of the ads she receives, subject to an ROI constraint, which requires her value to be at least her payment. We consider this generalization as advertisers may optimize for different objectives (Booking Holdings Inc. 2025, Wayfair Inc. 2025): an established firm may optimize for profit margins (QL) whereas newer firms may optimize for total visibility (VM).

In this paper, we consider three different mechanisms: generalized first-price (GFP), generalized second-price (GSP), and the Vickrey–Clarke–Groves mechanism (VCG). Each of these mechanisms allocates slots in descending order of bids. Under GFP, the advertiser pays the bid corresponding to each assigned ad, weighted by CTR. Under GSP—the most popular of the three mechanisms—the payment per assigned ad is the largest smaller bid submitted by another advertiser, again weighted by CTR. Under VCG, the advertiser pays her externality cost. With this, we are now ready to summarize our main results.

1.1. Contributions

We compare the performance of GFP, GSP, and VCG when advertisers strategically select their multipliers. In particular, we are interested in characterizing their worst case welfare guarantees when the advertisers are

behaving according to some notion of a game theoretic equilibrium. Most famously, a *pure Nash equilibrium* (PNE) requires that advertisers are independently best responding to one another. We also consider the more general notion of *coarse correlated equilibrium* (CCE), to which independent learning dynamics necessarily converge. For each mechanism and equilibrium notion, we measure the welfare lost to strategic behavior by the *price of anarchy* (PoA) defined as the worst-case ratio of equilibrium welfare to optimal welfare. Our main result is that GFP achieves a CCE PoA of 0.4569, whereas GSP and VCG achieve a PNE PoA of 0 (see Table 1).

Table 1 Price of anarchy

Mechanism	Price of anarchy	Result	Remark
GFP	≥ 0.4569 (CCE)	Thm. 1	We obtain a tight $1 - 1/e$ PNE PoA
GSP, VCG	$= 0$ (PNE)	Thm. 2	We construct a 0 welfare PNE in the simplest setting of all-QL or all-VM advertisers with 1 slot

GFP achieves a positive welfare guarantee. Our main result for GFP is a PoA of 0.4569 for CCEs (see Table 1 and Theorem 1). This is quite surprising: the constant 0.4569 is consistent with the bound for simultaneous first-price auctions in Colini-Baldeschi et al. (2026), Deng et al. (2024b) despite our extension to multi-item demand and uniform bidding. We construct a matching upper bound CCE, showing that this 0.4569 lower bound is tight. In fact, this upper bound construction belongs to the narrower class of correlated equilibria (CE, Definition EC.1) sitting between CCE and PNE (EC.3.2). Lastly, we complement our CCE PoA results with a tight PNE PoA of $1 - \frac{1}{e} \approx 0.6321$.

GSP and VCG do not. As illustrated in Table 1, we show that GSP and VCG admit PNE with 0 welfare, consistent with the findings of Deng et al. (2024a). In fact, this low welfare can occur even without either of the two main features of our model, multi-item demand and mixed objectives. We construct a 0 welfare PNE in an $M = 1$ slot auction when the advertisers are either all-QL or all-VM. To further show the weakness of GSP specifically, we show that 3 strengthened variants of it still achieve 0 welfare. First, we construct a 0 welfare *strict* PNE in which every advertiser's multiplier is her unique best response (Proposition 1). Second, we show that GSP also admits 0 welfare PNE under *self-pricing*, in which an advertiser's own ads also set her prices (Section 4.2.2). Thirdly, we show that our PNE construction holds even when we do not consider randomized VM deviations (Proposition 2).

1.2. Relationship to Prior Work and Technical Contributions

Our GFP guarantee matches the constant of Deng et al. (2024b) and Colini-Baldeschi et al. (2026). However, their arguments do not carry over to our setting as they consider simultaneous first-price auctions where each item is sold in its own auction and advertisers bid on each item separately. Directly adapting their methods requires meticulously constructing their thresholding argument for each individual slot and ad pair, being careful not to accidentally double count the contributions towards welfare or revenue. We avoid this route

altogether with the key observation that under GFP, the difference between equilibrium revenue and optimal welfare can be upper bounded as a function of the advertisers' deviation values (Lemma EC.1). The CCE constraints trivially upper bound these deviation values which yields our 0.4569 bound after some algebra. However, these CCE constraints are actually stronger than necessary and our approach allows us to establish the 0.4569 welfare bound for certain learning dynamics *even when CCE do not exist* (Theorem 3).

As for GSP, 0 welfare was already obtained in the single-unit setting with all VM advertisers (Deng et al. 2024a). In seeming contradiction to this, Dütting et al. (2026) show a $1/2$ guarantee for both GSP and VCG for the more general sponsored shopping setting. We reconcile these seemingly conflicting bounds by deriving welfare guarantees as a function of the smallest and largest multipliers at equilibrium. Specifically, Dütting et al. (2026) assume that multipliers are at least 1—which guarantees their $1/2$ welfare guarantee almost as an immediate corollary—whereas Deng et al. (2024a) allow for multipliers as low as 0. Our support-dependent welfare bounds (Proposition 5) recover both of these guarantees for both GSP and VCG, and even show a similar welfare improvement in GFP as the equilibrium support shrinks.

1.3. Extensions


We extend our main results in three directions. We first consider special cases of our sponsored shopping auction when advertisers do not have multi-item demand or mixed objectives. Second, we obtain equilibrium-support dependent welfare guarantees. Lastly, we consider the setting when valuations are stochastic or time varying.

1.3.1. Unit-demand and homogeneous objectives. The two main features of our sponsored shopping model were that advertisers possess multiple ads and could be a mixture of QL and VM. We already know that GSP/VCG achieve 0 welfare in the unit-demand setting with homogeneous advertiser objectives (i.e., when advertisers are all-VM or all-QL). Instead, we show that the guarantees of GFP strictly improve. In particular, with unit-demand, GFP achieves a tight PNE PoA of $\frac{4}{5}$, compared with $1 - \frac{1}{e}$ in general (Proposition 3). Similarly, when advertisers are all-QL or all-VM, GFP achieves a tight CCE PoA of $1 - \frac{1}{e}$ and $\frac{1}{2}$, respectively (Proposition 4), compared with 0.4569 in general.

1.3.2. Equilibrium-support dependent PoA. We characterize the PoA when every multiplier at equilibrium lies between a *floor* $A_{\min}$ and a *ceiling* $A_{\max}$, with $0 < A_{\min} \leq 1 \leq A_{\max}$. We show that across GFP, GSP, and VCG, the welfare guarantees improve beyond their worst case guarantee at rate $O(A_{\max}^{-1} + A_{\min})$ (Proposition 5). This implies that low welfare requires both very small and very large equilibrium multipliers, which we verify with our simulations (Figure 2). We also consider the setting when advertisers are prohibited from using multipliers that are *weakly dominated*—multipliers that perform weakly worse than another multiplier across all possible competing multiplier distributions. Depending on the mechanism and the advertiser's objective, this *undominated* restriction confines her equilibrium multipliers to a range with $A_{\min} = 1$, $A_{\max} = 1$, or both, or leaves them unrestricted (Lemma 1). Under this undominated restriction, we

show that the GFP and GSP welfare guarantees remain the same, whereas VCG improves to a $\frac{1}{2}$ welfare guarantee (Proposition 6).

Figure 2 Simulated welfare as the range of multipliers widens.



Note. Learning dynamics with one QL advertiser, one VM advertiser, and two slots (Section 6), with multipliers in $[A_{\max}^{-1}, A_{\max}]$. Lines mark the 1st percentile of welfare divided by optimal welfare across 100 trials, and shading extends to the 5th percentile.

1.3.3. Stochastic and time-varying valuations. Thus far, we have assumed that valuations are fixed. Interestingly, we show that GFP still provides a 0.4569 welfare guarantee even when allowing for stochastic valuations, possibly correlated across advertisers (Corollary 1). Moreover, this 0.4569 welfare guarantee extends even to the repeated setting (Theorem 3) with time-varying valuations (Corollary 2), *even when no equilibrium exists* (Proposition EC.6). This enables us to run simulations such as in Section 6 and illustrate our welfare guarantees numerically. Our guarantees for correlated and adversarial valuations are surprising, since full-information welfare guarantees typically extend to stochastic valuations only when they are independent (Roughgarden et al. 2017). To our knowledge, the closest related works on autobidding (Table 2) consider fixed valuations only.

1.4. Organization

The remainder of the paper is organized as follows: in Section 2, we give an overview of related work. In Section 3, we formally define the auction, the advertisers' objectives, and the equilibrium notions. In Section 4, we bound the price of anarchy of GFP, GSP, and VCG. In Section 5, we extend these results to the special cases of unit-demand or homogeneous objectives, obtain support-dependent welfare bounds, and extend them to stochastic and time-varying valuations. In Section 6, we compare welfare and revenue in simulations, and Section 7 concludes. The e-companion contains a guide to notation, complete result tables, technical lemmas, all proofs, and experimental details.

2. Related Work

We first compare our results with the closest welfare analyses of autobidding auctions. We then discuss the broader literatures on multi-unit and position auctions and on autobidding.

Table 2 Closest welfare analyses of autobidding auctions.

Paper	GFP	GSP	VCG	Multi-item	QL and VM	Price of anarchy
Deng et al. (2024b)	✓				✓	GFP: 1/2 (all-VM), 0.4569 (mixed); MNE.
Deng et al. (2024a)		✓				GSP: 0; all-VM, MNE.
Colini-Baldeschi et al. (2026)	✓				✓	GFP: 0.4569 (mixed); CCE.
Dütting et al. (2026)		✓	✓	✓		GSP, VCG: 1/2 (all-VM, multipliers at least one).
This paper	✓	✓	✓	✓	✓	GFP: ≥ 0.4569 (CCE); GSP, VCG: 0 (PNE).

“Multi-item” means that an advertiser may win several slots in one auction, and “QL and VM” means that both types compete in the same auction. In the last column, “mixed” means that both types may be present, and MNE and PNE denote mixed and pure Nash equilibria.

2.1. Closest Related Work

Table 2 compares our results with the closest autobidding literature. Dütting et al. (2026) study the sponsored-shopping model that we adopt, in which one multiplier sets all of an advertiser’s bids (uniform bidding). Under the assumption that advertisers are all-VM and use multipliers of at least one, they show that autobidding equilibria exist and keep at least half of the optimal welfare under both GSP and VCG. In contrast, our results require no such restrictions, and like Deng et al. (2024a), we show that GSP and VCG can then lose almost all welfare even with one slot.

As for GFP, Deng et al. (2024b) show that when each item is sold in its own first-price auction and advertisers bid separately on each item, equilibria keep half of the optimal welfare when advertisers are all-VM, and a 0.4569 fraction when QL and VM advertisers compete. Colini-Baldeschi et al. (2026) obtain the same constant for CCE. We show that this constant still holds under the uniform bidding setting of Dütting et al. (2026).

Technical contributions. Our main welfare guarantee is the 0.4569 GFP lower bound which is consistent with Deng et al. (2024b) and Colini-Baldeschi et al. (2026). However, in our sponsored shopping setting, slots are sold through a single auction with uniform-bidding rather than through a simultaneous first-price auction. This prohibits a straightforward extension of their argument which optimizes over the probability of allocation, and instead requires a more careful characterization of the expected acquired values, utility, and revenue. Our generalized analysis actually reveals the crucial insight that the 0.4569 welfare lower bound does *not* require equilibrium existence (Theorem 3)—which may not exist in constrained games such as our sponsored shopping auction (Proposition EC.6). Using this insight, we show that a wide class of natural learning dynamics maintain a 0.4569 welfare guarantee. This is the first welfare guarantee for learning dynamics without requiring equilibrium existence in the autobidding literature. In particular, our dynamic welfare bound is unique in two ways: it holds when both QL and VM advertisers are present, whereas earlier guarantees for learning dynamics assume all-VM. Moreover, our welfare guarantee holds even when equilibria do not exist, as opposed to asserting only that the learning dynamics need not converge (Fikioris and Tardos 2025, Gaitonde et al. 2023, Lucier et al. 2024).

2.2. Multi-Unit and Position Auctions

Position auctions, in which each advertiser has one ad and slots differ in CTR, are the classic model of sponsored search (Edelman et al. 2007, Varian 2007), and the welfare of their equilibria under GSP is well studied (Caragiannis et al. 2011, Elzayn et al. 2022, Paes Leme and Tardos 2010). We generalize these models by considering multi-item demand and mixed advertiser objectives in our sponsored shopping auction model. We note that neither of these generalizations are new on their own: Dütting et al. (2026) consider the special case of multi-item demand but homogeneous advertiser objectives for GSP and VCG. Similarly, Colini-Baldeschi et al. (2026) analyzed single-slot auctions for mixed advertiser objectives.

It is also worth mentioning the related literature regarding multi-item demand. When all slots have the same CTR, our model is a *multi-unit auction* where each advertiser bids on several *identical* units, where GFP is referred to as the *pay-as-bid* auction. Several foundational works study the equilibria and efficiency of such pay-as-bid auctions (Ausubel et al. 2014, Birmpas et al. 2019, de Keijzer et al. 2013, Engelbrecht-Wiggans and Kahn 1998a,b, Fabra et al. 2006, Hortaçsu and McAdams 2010, Markakis et al. 2022, Markakis and Telelis 2015, Wilson 1979). More recent work studies equilibria of uniform-price auctions (Feng et al. 2026) (a multi-unit variant of second-price auctions), including with value-maximizing buyers (Golrezaei and Sahoo 2026), and learning in repeated multi-unit auctions (Brânzei et al. 2023, Galgana and Golrezaei 2025, 2026, Golrezaei and Sahoo 2025). The key difference between our model and these existing works—beyond the different weights in each slots—is that we consider uniform-bidding. In contrast, these aforementioned papers assume that buyers choose separate bids for each unit.

2.3. Autobidding

Objectives and constraints. Aggarwal et al. (2019) formalize value-maximization under ROI and budget constraints and show that uniform bidding is optimal in truthful auctions, and Balseiro et al. (2021) compare value- and utility-maximizing advertisers. Other work designs auctions for ROI-constrained buyers (Golrezaei et al. 2021b), studies bidding and pricing under budget and ROI constraints (Golrezaei et al. 2021a) and autobidding across several advertising channels (Deng et al. 2023), and examines advertisers’ incentives to misreport their goals to the autobidder (Alimohammadi et al. 2023, Feng et al. 2024). Aggarwal et al. (2024) survey the area.

Welfare. Truthful auctions keep half of the optimal welfare when advertisers are all-VM (Aggarwal et al. 2019). Later work proves welfare guarantees in non-truthful auctions and with budget constraints (Balseiro et al. 2023, Liaw et al. 2023, 2024, Liu and Shen 2023).

Pacing equilibria. Another line of work studies the existence and computation of *pacing equilibria* where bidders have explicit budget constraints. Here, each advertiser’s multiplier is pinned down by her budget and ROI constraint (Chen et al. 2024, Conitzer et al. 2022a,b, Li and Tang 2024). Our notion of equilibrium is separate as we consider primarily the static setting as well as the dynamic but non-budgeted setting in Section 5.3.2.

Learning. Bidding algorithms are well studied for repeated auctions with budget or ROI constraints (Balseiro and Gur 2019, Castiglioni et al. 2022, 2024, Susan et al. 2023) and for first-price auctions (Balseiro et al. 2019). When budget- or ROI-constrained advertisers use such algorithms, they keep a constant fraction of the optimal liquid welfare—a separate but related notion of welfare—regardless of whether or not the learning dynamics converge (Fikioris and Tardos 2025, Gaitonde et al. 2023, Lucier et al. 2024). We go one step beyond this and show that our welfare guarantees hold even when equilibria need not exist (Section 5.3).

3. Model

In this section, we define the sponsored-shopping environment, advertiser objectives, mechanisms, randomized deviations, and the price of anarchy. We first analyze fixed valuations. Section 5.3.1 considers privately observed stochastic valuations, and Section 5.3.2 extends the learning guarantees to valuations that vary across auctions.

3.1. Sponsored Shopping Auction Model

We adopt the sponsored shopping auction model of Dütting et al. (2026) which is a multi-item demand version of the standard position auction model (Edelman et al. 2007, Varian 2007).¹ Here, there are N advertisers and M slots with click-through rates (CTRs) given by

$$1 = c_1 \geq c_2 \geq \dots \geq c_M \geq 0,$$

with $c_m \doteq 0$ for all $m > M$ by convention. Each advertiser has M ads with valuations

$$v_{n,1} \geq \dots \geq v_{n,M} \geq 0, \quad n = 1, \dots, N, \quad (1)$$

where ad (n, j) is the j th ad of advertiser n , and its *valuation* $v_{n,j}$ is the expected value of a user clicking on it. Each ad corresponds to one of the advertiser’s products, where advertisers with fewer than M products are padded with valuations of 0. If ad (n, j) is assigned to slot m , it contributes $c_m v_{n,j}$ to advertiser n ’s value and to welfare.

3.2. Multipliers and Allocation

Each advertiser n may submit any multiplier $\alpha_n \in [0, \infty)$. The bid corresponding to ad (n, j) is $b_{n,j} \doteq \alpha_n v_{n,j}$. We write $\alpha \doteq (\alpha_1, \dots, \alpha_N)$ for the profile of multipliers. All ads are ranked by bid and assigned to slots in descending order with ties resolved by one fixed, value-independent priority order over all ads (e.g., ordered by advertiser index and then by ad index). At α , the *rank* $r_{n,j}(\alpha)$ of ad (n, j) is its position in this ranking; i.e., the ad of rank m receives slot m , and ads of rank above M receive no slot. We write $r_{n,j}$ when α is clear

¹ While our model assumes that the CTRs are independent of the configuration of ads among them, there exist other models of user engagement that account for possible dependencies (Athey and Ellison 2011, Craswell et al. 2008, Kempe and Mahdian 2008).

from context. Since allocations are by bid order (and ad index as per the tie-breaking rule above), advertisers are allocated their most valuable slots first. With this, we define advertiser n 's *acquired value* and *welfare* as

$$V_n(\alpha) \doteq \sum_{j=1}^M c_{r_{n,j}} v_{n,j}, \quad \text{Wel}(\alpha) \doteq \sum_{n=1}^N V_n(\alpha). \quad (2)$$

Recall that $c_m = 0$ for all $m > M$. Letting $v_{(m)}$ be the m th largest valuation among all ads, the optimal welfare is obtained when the top valuations are matched with the top slots:

$$\text{OPT} \doteq \sum_{m=1}^M c_m v_{(m)}.$$

3.3. Advertiser Objectives

Let $P_n(\alpha) \geq 0$ be advertiser n 's payment at profile α , which depends on the mechanism that will be defined in Section 3.4. Letting σ denote a distribution over profiles α , each advertiser has one of the following two objectives.

DEFINITION 1 (QUASI-LINEAR ADVERTISER). A quasi-linear (QL) advertiser maximizes expected utility $\mathbb{E}_{\alpha \sim \sigma} [V_n(\alpha) - P_n(\alpha)]$. For convenience, we write the utility as $u_n(\alpha) \doteq V_n(\alpha) - P_n(\alpha)$.

DEFINITION 2 (VALUE-MAXIMIZING ADVERTISER). A value-maximizing (VM) advertiser maximizes expected value $\mathbb{E}_{\alpha \sim \sigma} [V_n(\alpha)]$ subject to the return-on-investment (ROI) constraint $\mathbb{E}_{\alpha \sim \sigma} [P_n(\alpha)] \leq \mathbb{E}_{\alpha \sim \sigma} [V_n(\alpha)]$. We say that distribution σ is *feasible* for VM advertiser n if their ROI constraint is satisfied.

We note that we adopt the standard convention from previous literature by normalizing the ROI target to one (Deng et al. 2024b).² For simplicity, we omit the expectation subscript $\alpha \sim \sigma$ whenever it is clear from context. We write $\mathcal{Q}$ and $\mathcal{V}$ for the sets of QL and VM advertisers, where $\mathcal{Q} \subseteq \{1, \dots, N\}$ and $\mathcal{V} \doteq \{1, \dots, N\} \setminus \mathcal{Q}$. Unless otherwise specified, both types can be present.

3.4. Mechanisms

Our main goal is to compare the GFP, GSP, and VCG mechanisms, three of the most popular auction formats used in practice (Deng et al. 2024b, Dütting et al. 2026). All three allocate slots in descending bid order as per Section 3.2 but differ in payment. To define these, we first define $\hat{b}_{n,i}(\alpha)$ to be the i th largest bid among the ads of the other advertisers besides n , and as with ranks, we write $\hat{b}_{n,i}$ when α is clear from context. In particular, $\hat{b}_{n,r_{n,j}-j+1}$ is the largest bid of another advertiser's ad ranked below ad (n, j) and thus the payment for ad (n, j) is:

$$\begin{aligned} \text{GFP:} \quad & c_{r_{n,j}} b_{n,j}, \\ \text{GSP:} \quad & c_{r_{n,j}} \hat{b}_{n,r_{n,j}-j+1}, \\ \text{VCG:} \quad & \sum_{m=r_{n,j}}^M (c_m - c_{m+1}) \hat{b}_{n,m-j+1}, \end{aligned} \quad (3)$$

² A VM advertiser n with target ROI $\tau_n > 0$ requires $\mathbb{E}[V_n(\alpha)] \geq \tau_n \mathbb{E}[P_n(\alpha)]$, and dividing her valuations by τ_n turns this into $\mathbb{E}[P_n(\alpha)] \leq \mathbb{E}[V_n(\alpha)]$. Every instance with general τ_n thus has the same equilibria as an instance with target one, whether or not QL advertisers are present (Colini-Baldeschi et al. 2026). Our welfare guarantees are thus in terms of liquid welfare.

and advertiser n 's payment $P_n(\alpha)$ is the sum of the payments for all of her ads. To interpret these, under GFP, each ad pays its own bid per click, so $P_n(\alpha) = \alpha_n V_n(\alpha)$. Under GSP, it pays per click the largest bid of another advertiser's ad ranked below it. Under VCG, the payment per click of an ad (n, j) with CTR $c_{r_{n,j}} > 0$ is its payment divided by the probability of a click $c_{r_{n,j}}$. Thus, in this payment, the weights $(c_m - c_{m+1})/c_{r_{n,j}}$ on the bids $\hat{b}_{n,m-j+1}$ for $m = r_{n,j}, \dots, M$ sum to one. Each ad thus pays per click a weighted average of the bids of the other advertisers' ads ranked below it. Since this is at most its corresponding GSP payment, which is at most its own bid, the payment $P_n(\alpha)$ is lowest at VCG and highest at GFP. As shown in EC.2.3, the VCG payment corresponds precisely to the familiar interpretation as the advertiser's *externality cost*—the CTR-weighted bids of the other advertisers' ads without her, minus their CTR-weighted bids when she is present.

EXAMPLE 1 (PAYMENTS UNDER GFP, GSP, AND VCG). Consider $M = 2$ slots with CTRs $c_1 = 1$ and $c_2 = 1/2$. Let $N = 2$ with advertiser 1's valuations being 3 and 2, and advertiser 2's being 1 and 0. With multipliers of 1, advertiser 1 receives both slots and her competing bids are $\hat{b}_{1,1} = 1$ and $\hat{b}_{1,2} = 0$.

- Payment under GFP: $1 \cdot 3 + \frac{1}{2} \cdot 2 = 4$, her own bids weighted by CTR.
- Payment under GSP: $1 \cdot 1 + \frac{1}{2} \cdot 1 = \frac{3}{2}$; i.e., advertiser 1 pays $\hat{b}_{1,1}$ for both ads, weighted by CTR.
- Payment under VCG: $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 1 = 1$, her externality cost.

3.5. Equilibria and the Price of Anarchy

Advertisers choose their multipliers strategically and behave at equilibrium, i.e., a state in which no advertiser wishes to deviate. However, it is important to consider *randomized deviations* in our setting, as a VM advertiser may improve her outcome as her ROI constraints need only hold in expectation. In particular, she can submit a low multiplier that generates *ROI slack*—a strictly positive difference between value and payment—to subsidize a high multiplier that incurs an *ROI deficit*—a strictly positive difference between payment and value. The following example illustrates this.

EXAMPLE 2 (NECESSITY OF RANDOMIZED MULTIPLIERS). Consider GFP with a single slot with CTR of 1, and a VM advertiser n with one ad with valuation 2. Suppose that their highest competing bid is 1 or $2 + \varepsilon$, with probability $1/2$ each, for some $\varepsilon \in (0, 1/4]$ with ties broken against advertiser n . Whenever she wins the slot, she obtains value 2 and pays her own bid.

Fixed multipliers. If advertiser n reports multiplier α_n , her bid is $2\alpha_n$. We consider her possible bids:

- Consider selecting some multiplier that yields bid $2\alpha_n > 2 + \varepsilon$. Then, she always wins and pays $2\alpha_n > 2$, which exceeds her value and thus violates ROI.
- She considers selecting some bid less than $2 + \varepsilon$. She wins when the other bid is 1, which happens with probability $1/2$, yielding expected value $\frac{1}{2} \cdot 2 = 1$.

Thus, the optimal ROI feasible fixed multiplier is any α_n that wins whenever her competitor bids 1 but not at $2 + \varepsilon$. This multiplier yields an expected value of 1. Now suppose that she uses a randomized strategy: she submits multiplier $\frac{1}{2} + \varepsilon$ with probability 4ε and multiplier $1 + \varepsilon$ otherwise.

- Multiplier $\frac{1}{2} + \varepsilon$ yields bids $1 + 2\varepsilon$, so it wins only against the bid 1. Its expected value is $\frac{1}{2} \cdot 2 = 1$ and its expected payment is $\frac{1}{2}(1 + 2\varepsilon)$, so it generates ROI slack.
- Multiplier $1 + \varepsilon$ yields bids $2 + 2\varepsilon$, so it wins against both bids. Its value is 2 and its payment is $2 + 2\varepsilon$, so it incurs an ROI deficit.

Her expected value exceeds her payment by

$$4\varepsilon\left(\frac{1}{2} - \varepsilon\right) - (1 - 4\varepsilon)2\varepsilon = 4\varepsilon^2 \geq 0.$$

Thus, the randomized strategy satisfies her expected ROI constraint with expected value $4\varepsilon \cdot 1 + (1 - 4\varepsilon) \cdot 2 = 2 - 4\varepsilon$, nearly twice the value 1 of the best fixed multiplier. We can obtain a similar example with GSP and VCG (EC.2.1).

Randomized deviations. We therefore allow VM advertisers to deviate to randomized strategies. We write α_{-n} for the multipliers of the advertisers other than n , and (α', α_{-n}) for the profile in which advertiser n uses α' instead of α_n . For a *randomized strategy* $\tilde{\sigma}_n$, a distribution over $[0, \infty)$, draw $\tilde{\alpha}_n \sim \tilde{\sigma}_n$ independently of $\alpha \sim \sigma$. The strategy is *feasible* for a VM advertiser if

$$\mathbb{E}[P_n(\tilde{\alpha}_n, \alpha_{-n})] \leq \mathbb{E}[V_n(\tilde{\alpha}_n, \alpha_{-n})],$$

with expectations over both $\tilde{\alpha}_n$ and α . Randomizing cannot help a QL advertiser as her expected utility is linear in the deviating distribution. Randomized deviations, however, can help VM advertisers by improving their value whilst satisfying their ROI constraints as per (EC.2.1).

DEFINITION 3 (EQUILIBRIUM). Let σ be a feasible, possibly correlated distribution over profiles, that is, $\mathbb{E}_{\alpha \sim \sigma}[P_n(\alpha)] \leq \mathbb{E}_{\alpha \sim \sigma}[V_n(\alpha)]$ for every $n \in \mathcal{V}$.

- (i) σ is a *coarse correlated equilibrium* (CCE) if every QL advertiser $n \in \mathcal{Q}$ satisfies

$$\mathbb{E}_{\alpha \sim \sigma}[u_n(\alpha)] \geq \mathbb{E}_{\alpha \sim \sigma}[u_n(\alpha', \alpha_{-n})] \quad \text{for every multiplier } \alpha' \geq 0,$$

and every VM advertiser $n \in \mathcal{V}$ satisfies

$$\mathbb{E}_{\alpha \sim \sigma}[V_n(\alpha)] \geq \mathbb{E}_{\tilde{\alpha}_n \sim \tilde{\sigma}_n, \alpha_{-n} \sim \sigma_{-n}}[V_n(\tilde{\alpha}_n, \alpha_{-n})] \quad \text{for every feasible randomized strategy } \tilde{\sigma}_n.$$

- (ii) A profile α is a *pure Nash equilibrium* (PNE) if the distribution is a point mass CCE.

Price of anarchy. Before we formally define the price of anarchy, we first must formally define what it represents. We define an *instance* I to consist of 1) $N, M \geq 1$ the number of advertisers and slots, 2) CTRs $\mathbf{c}$, 3) valuations $\mathbf{v}_1, \dots, \mathbf{v}_N$, and 4) the sets $\mathcal{Q}, \mathcal{V} \subseteq \{1, \dots, N\}$ of QL and VM advertisers (and technically, a tie-breaking rule). Its optimal welfare is $\text{OPT}(I) \doteq \sum_{m=1}^M c_m v_{(m)}$, as in Section 3.2. Now, we are ready to define the welfare price of anarchy.

DEFINITION 4 (PRICE OF ANARCHY). Fix an equilibrium class $\mathcal{E} \in \{\text{CCE}, \text{PNE}\}$ and a mechanism $\mathcal{M} \in \{\text{GFP}, \text{GSP}, \text{VCG}\}$. The *price of anarchy* (PoA) is

$$\text{PoA}_{\mathcal{E}}(\mathcal{M}) \doteq \inf_{I, \sigma} \frac{\mathbb{E}_{\alpha \sim \sigma}[\text{Wel}(\alpha)]}{\text{OPT}(I)}, \quad (4)$$

where the infimum is over all instances I with $\text{OPT}(I) > 0$ and all equilibria σ of I under $\mathcal{M}$ in the class $\mathcal{E}$.

The denominator in (4) is the optimal welfare $\text{OPT}(I)$ (Section 3.2), while its numerator is the equilibrium welfare. Our PoA convention implies that every PoA lies in $[0, 1]$, and larger values indicate stronger welfare guarantees. Importantly, as $\text{PNE} \subseteq \text{CCE}$, a PNE upper-bound construction applies to every CCE and conversely, a CCE lower bound applies to every PNE. For simplicity, we omit the equilibrium subscript when a CCE lower bound and a PNE upper bound match.

4. Results

We present the welfare guarantees for GFP (Section 4.1) and for GSP and VCG (Section 4.2), which we prove in EC.3.

4.1. GFP

Our main result is that GFP guarantees a constant fraction of optimal welfare even in very general settings.

THEOREM 1. $\text{PoA}_{\text{CCE}}(\text{GFP}) = 0.4569$ and $\text{PoA}_{\text{PNE}}(\text{GFP}) = 1 - 1/e$.

Interestingly, this 0.4569 constant is also obtained by Deng et al. (2024b) and Colini-Baldeschi et al. (2026) in their analyses of simultaneous first-price auctions. Thus the common multiplier and multi-item demand do not weaken this guarantee. As apparent, Theorem 1 gives different PoA's for CCE and PNE, and both are tight. More specifically, the CCE bound remains tight for correlated equilibria (CE), a narrower equilibrium class which implies that 0.4569 is also the PoA over CE (EC.3.2).³ This tightness is important. Unlike a CCE, a CE lets advertisers swap their recommended multiplier for another, trivially ruling out recommending dominated multipliers (see Section 5.2.2). Thus the gap between the CCE and PNE constants does not come from CCE placing weight on dominated multipliers that always perform poorly.

Proof Overview.

4.1.1. CCE Lower Bound. Fix a CCE σ , and let $U_n \doteq \mathbb{E}[u_n(\alpha)]$, $V_n \doteq \mathbb{E}[V_n(\alpha)]$, $P_n \doteq \mathbb{E}[P_n(\alpha)]$ denote expected utilities, values, and payments. Let V_n^* denote advertiser n 's value in the optimal allocation, so that $\text{OPT} = \sum_n V_n^*$. We subscript these quantities QL or VM to sum over their respective populations, e.g., $U_{\text{QL}} \doteq \sum_{n \in \text{QL}} U_n$ is the QL advertisers' total expected utility. Let V_{QL}^* and V_{VM}^* be the two populations' contributions to OPT. Lastly, let $\text{Rev}(\alpha) \doteq \sum_n P_n(\alpha)$ denote the revenue at profile α . The proof has four steps.

³ We formally define CE in Definition EC.1 as a solution class between CCE and Nash equilibria that represent the outcome of no-swap-regret dynamics.

Lower Bound on Equilibrium Welfare. A QL advertiser's value is utility plus payment. Decomposing this, we have $\mathbb{E}[\text{Wel}] = U_{\text{QL}} + \mathbb{E}[\text{Rev}] - P_{\text{VM}} + V_{\text{VM}}$. Payments are nonnegative, so $P_{\text{VM}} \leq \mathbb{E}[\text{Rev}]$, and the VM advertisers' expected ROI constraints give $P_{\text{VM}} \leq V_{\text{VM}}$. Hence

$$\mathbb{E}[\text{Wel}] \geq U_{\text{QL}} + \max\{\mathbb{E}[\text{Rev}], V_{\text{VM}}\}. \quad (5)$$

The remaining steps show that revenue or VM value is large whenever QL utility is small.

Upper Bound on Optimal Value Using Recoverable Value. Advertiser n 's *weighted bid* $B_n(\alpha) \doteq \alpha_n V_n(\alpha) = \sum_{j=1}^M c_{r_{n,j}} b_{n,j}$ is the sum of her bids weighted by the CTRs of their slots, and the *total weighted bid* is $B(\alpha) \doteq \sum_n B_n(\alpha)$. Under GFP, each advertiser pays her weighted bid, so revenue equals the total weighted bid. The auction's allocation maximizes the total weighted bid over all assignments of ads to slots (EC.2.3). Hence, as advertiser n raises her multiplier with α_{-n} fixed, revenue grows at the rate of her acquired value V_n . For each advertiser n , we fix a *target* $\bar{\alpha}_n \in [0, 1]$ —the largest multiplier to which we let her deviate. Her *recoverable value* $\int_0^{\bar{\alpha}_n} V_n(\alpha', \alpha_{-n}) d\alpha'$ is the integral of the value that she could obtain by deviating to each multiplier between 0 and $\bar{\alpha}_n$. Now, consider raising the advertisers' multipliers to at least their targets one at a time. Each raise increases revenue by at most the advertiser's recoverable value. Note that this is true only because raising the other multipliers can only *lower* her acquired value. At this raised profile where every advertiser reports at least their target, every advertiser n bids at least $\bar{\alpha}_n$ times her valuations. Thus, this raised profile—which does at least as well as the optimal allocation in terms of total weighted bid—has a total weighted bid of at least $\sum_n \bar{\alpha}_n V_n^*$. This yields Lemma EC.1(i):

$$\sum_n \bar{\alpha}_n V_n^* \leq \text{Rev}(\alpha) + \sum_n \int_0^{\bar{\alpha}_n} V_n(\alpha', \alpha_{-n}) d\alpha'. \quad (6)$$

We call this the *recovery inequality* and it shows that revenue must be large whenever recoverable value is small and vice versa.

Equilibrium Bounds on Recoverable Value. The deviations in the recovery inequality (6) use multipliers $\alpha' \leq \bar{\alpha}_n \leq 1$, so they never bid above valuations, even when the equilibrium multipliers do. Under GFP, such a deviation pays $\alpha' V_n \leq V_n$, so it satisfies ROI, and the CCE conditions bound its value for both types of advertisers.

- A QL advertiser n who deviates to $\alpha' < 1$ obtains utility $(1 - \alpha')V_n(\alpha', \alpha_{-n})$, which is at most U_n in expectation. Hence $\mathbb{E}[V_n(\alpha', \alpha_{-n})] \leq U_n/(1 - \alpha')$, and integrating bounds her expected recoverable value by $-U_n \ln(1 - \bar{\alpha}_n)$.
- A VM advertiser n who deviates to $\alpha' \leq 1$ acquires at most her equilibrium value in expectation, that is, $\mathbb{E}[V_n(\alpha', \alpha_{-n})] \leq \mathbb{E}[V_n(\alpha)]$. With target one, her expected recoverable value is therefore at most $\mathbb{E}[V_n(\alpha)]$.

Target Selection. Take expectations in the recovery inequality (6), with a common target $\bar{\alpha} < 1$ for the QL advertisers and with target 0 or one for every VM advertiser. The two bounds above then give (Lemma EC.2)

$$\mathbb{E}[\text{Rev}] \geq \bar{\alpha} V_{\text{QL}}^* + U_{\text{QL}} \ln(1 - \bar{\alpha}), \quad \mathbb{E}[\text{Rev}] + V_{\text{VM}} \geq V_{\text{VM}}^* + \bar{\alpha} V_{\text{QL}}^* + U_{\text{QL}} \ln(1 - \bar{\alpha}).$$

Letting $\bar{u} \doteq U_{\text{QL}}/V_{\text{QL}}^* \in (0, 1)$ be the QL advertisers' utility relative to their optimal value, the best target is $\bar{\alpha} = 1 - \bar{u}$, which turns the common QL term into $V_{\text{QL}}^* \phi(\bar{u})$, where $\phi(\bar{u}) \doteq 1 - \bar{u} + \bar{u} \ln \bar{u}$. Substituting into Equation (5), and using that a maximum is at least either of its arguments and at least their average,

$$\mathbb{E}[\text{Wel}] \geq U_{\text{QL}} + \max \left\{ V_{\text{QL}}^* \phi(\bar{u}), \frac{V_{\text{VM}}^* + V_{\text{QL}}^* \phi(\bar{u})}{2} \right\}. \quad (7)$$

Normalizing $\text{OPT} = V_{\text{QL}}^* + V_{\text{VM}}^* = 1$ and substituting $U_{\text{QL}} = \bar{u} V_{\text{QL}}^*$, both lower bounds become affine in V_{QL}^* . For $\bar{u} \leq 1/e$, the first term increases and the second does not, so the worst case occurs when they are equal. As for $\bar{u} > 1/e$, the second term is at least $1/2$. Minimizing over $\bar{u}$ (which yields $\bar{u} \approx 0.1586$) gives

$$\frac{\mathbb{E}[\text{Wel}]}{\text{OPT}} \geq \min_{0 \leq \bar{u} \leq 1} \frac{1 + \bar{u} \ln \bar{u}}{2 - \bar{u} + \bar{u} \ln \bar{u}} \approx 0.4569.$$

4.1.2. PNE Lower Bound. At a PNE α , the ROI constraints and utility non-negativity hold at α . Since $P_n = \alpha_n V_n$ under GFP, every advertiser with positive value uses a multiplier $\alpha_n \leq 1$, so her bids never exceed her valuations. We can then compare each advertiser's part of the bound with her own value rather than with revenue, using the finer Lemma EC.1(ii). With a common target $\bar{\alpha}$, it implies

$$\bar{\alpha} \text{OPT} \leq \sum_n \left[\max\{\alpha_n, \bar{\alpha}\} V_n(\alpha) + \int_0^{\bar{\alpha}} (V_n(\max\{\alpha_n, \alpha'\}, \alpha_{-n}) - V_n(\alpha)) d\alpha' \right].$$

Unlike the recovery inequality (6), the deviations here only raise multipliers, and the integrand counts only the value gained beyond the current value. Each bracket is at most $V_n(\alpha)$. If $\alpha_n \geq \bar{\alpha}$, the integral is 0 and the bracket is $\alpha_n V_n(\alpha) \leq V_n(\alpha)$. If $\alpha_n < \bar{\alpha}$, we use the advertiser's equilibrium condition.

- A VM advertiser can raise her multiplier to any $\alpha' \leq 1$ without violating ROI, so at a PNE she gains no value from it. Her integral is 0, and her bracket is $\bar{\alpha} V_n(\alpha) \leq V_n(\alpha)$.
- A QL advertiser who raises her multiplier to $\alpha' \in [\alpha_n, \bar{\alpha}]$ obtains at most her current utility, so $(1 - \alpha') V_n(\alpha', \alpha_{-n}) \leq (1 - \alpha_n) V_n(\alpha)$. Her bracket is therefore at most

$$V_n(\alpha) \left[\bar{\alpha} + \int_{\alpha_n}^{\bar{\alpha}} \frac{\alpha' - \alpha_n}{1 - \alpha'} d\alpha' \right] = V_n(\alpha) \left[\alpha_n + (1 - \alpha_n) \ln \frac{1 - \alpha_n}{1 - \bar{\alpha}} \right],$$

which is at most $V_n(\alpha)$ for every $\alpha_n \in [0, \bar{\alpha})$ exactly when $\bar{\alpha} \leq 1 - 1/e$.

The target $\bar{\alpha} = 1 - 1/e$ therefore gives $(1 - 1/e)\text{OPT} \leq \sum_n V_n(\alpha) = \text{Wel}(\alpha)$.

4.1.3. Upper Bounds. Both bounds are tight. In the worst-case PNE, a QL advertiser owns many ads with valuation one but receives only about a $1/e$ fraction of the unit-CTR slots, and a VM advertiser receives the remaining slots with lower-valuation ads. In the worst-case CE, one of many symmetric VM advertisers, chosen uniformly at random, uses a very large multiplier and wins many low-valuation ads in place of higher-valuation ads, while the others use a low multiplier, and adding a QL advertiser to this construction gives the minimum above. We prove Theorem 1 in EC.3.1 and EC.3.2. $\square$

4.2. GSP and VCG

Unlike GFP, both GSP and VCG achieve arbitrarily low PoA even at a PNE with a single slot.

THEOREM 2. $\text{PoA}_{\text{PNE}}(\text{GSP}) = \text{PoA}_{\text{PNE}}(\text{VCG}) = 0$, even with one slot and two QL or two VM advertisers.

Theorem 2 illustrates the severe fragility of GSP and VCG. Their welfare can approach 0 even without multiple slots, multi-item demand, mixed objectives, or broader equilibrium classes such as CCE.

Proof Overview. A single slot and two advertisers with valuations 1 and $\varepsilon < 1$ already suffice for poor welfare. The higher-valuation advertiser uses multiplier 0, while the lower-valuation advertiser uses a multiplier above $1/\varepsilon$ and wins with a bid above one. The winner pays 0, satisfies ROI, and acquires all her value, so she cannot gain whether she is QL or VM. The other advertiser cannot win the slot without paying more than her value and thus remains indifferent among all losing bids. Consequently, the resulting profile is a PNE when the advertisers are all-QL or all-VM, with welfare ratio ε . With one slot and one ad per advertiser, GSP and VCG have identical allocations and payments, since the highest bid wins and pays the other advertiser's bid. This remains true after every unilateral deviation, so the construction applies unchanged to VCG. We prove Theorem 2 in EC.3.3. $\square$

Theorem 2 is disconcerting as GSP is the most commonly used in practice (Dütting et al. 2026, Edelman et al. 2007, Varian 2007). Even more troubling is that GSP in particular is especially fragile. In fact, we give three natural refinements of GSP and its equilibria and prove that even these are not enough to rescue its welfare.

4.2.1. Strict Equilibria. The PNE in the proof of Theorem 2 relies on indifference. The losing advertiser is indifferent among all losing multipliers, and the winning advertiser is indifferent among all winning multipliers. A *strict* PNE rules this out. It is a PNE in which every advertiser's multiplier is her unique best response, so every other multiplier gives a QL advertiser strictly less utility, and every other feasible randomized strategy gives a VM advertiser strictly less expected value. Loosely speaking, strict PNE are attracting whereas non-strict PNE are *not* under many natural learning dynamics in finite games (Giannou et al. 2021, Vlatakis-Gkaragkounis et al. 2020). As such, the familiar reader may discount the PNE in the proof of Theorem 2. To address this, we construct strict PNE with arbitrarily low welfare (full construction in EC.3.4):

PROPOSITION 1. $\text{PoA}_{\text{PNE}}(\text{GSP}) = 0$, even over strict PNE.

4.2.2. Self-Pricing. One may wonder whether the failure of GSP arises from the mechanism preventing an advertiser’s own bids from setting her prices (Section 3.4). However, *self-pricing*—allowing an advertiser’s own ads further down the ranking to determine her GSP payments (Dütting et al. 2026)—does not prevent poor welfare. Under self-pricing, ad (n, j) pays $c_{r_{n,j}} b_{n',j'}$, where (n', j') is the ad ranked directly below it and may be advertiser n ’s own ad $(n, j + 1)$, whereas under GSP it pays $c_{r_{n,j}} \hat{b}_{n,r_{n,j}-j+1}$, which ignores her own ads (Equation (3)). With one ad per advertiser, self-pricing and GSP have identical allocations and payments. As such, the single-slot PNE construction in the proof of Theorem 2 applies unchanged when the advertisers are all-QL or all-VM. Its welfare ratio approaches 0, so the PNE price of anarchy of GSP with self-pricing is 0 as well.

4.2.3. Deviations to Fixed Multipliers. Recall from Example 2 that a VM advertiser can gain value by randomizing her multiplier, since her ROI constraint needs to hold only in expectation. This is why our equilibria allow VM advertisers to deviate to randomized strategies. We now ask whether the PoA of GSP would still be 0 if we prohibited deviations to this larger class of strategies and considered only *fixed deviations*—deviations to a fixed multiplier. We call a profile a *PNE with fixed deviations* if no QL advertiser improves their utility with any fixed deviation, and every VM advertiser satisfies ROI and gains no value from any ROI feasible fixed multiplier. Our next result shows that mixed GSP can still have arbitrarily low welfare.

PROPOSITION 2. $\text{PoA}_{\text{PNE}}(\text{GSP}) = 0$, even over PNE with fixed deviations.

Proof Overview. In the PNE in the proof of Theorem 2, let the high-valuation advertiser be QL and the low-valuation advertiser be VM. The VM advertiser’s multiplier exceeds one, and since no randomized strategy helps either advertiser, no fixed multiplier does. We prove Proposition 2 in EC.3.5. $\square$

5. Extensions

We extend our main results to unit-demand and homogeneous objectives (Section 5.1), to a restricted equilibrium support (Section 5.2), and to stochastic and time-varying valuations (Section 5.3).

5.1. Unit-Demand and Homogeneous Objectives

Unit-demand. Multi-item demand allows each advertiser to submit several ads and receive several slots. Under *unit-demand*, each advertiser has a single ad, but the auction may still have several slots with different CTRs. In our model, an instance has unit-demand if $v_{n,j} = 0$ for every advertiser n and every $j \geq 2$. We first compare PNE welfare under unit-demand with the multi-item constructions of Section 4.

PROPOSITION 3. With unit-demand, $\text{PoA}_{\text{PNE}}(\text{GFP}) = 4/5$.

This stronger welfare guarantee is possible because the unit-demand setting—which is essentially just the standard position auction—is the special case of sponsored shopping with $v_{n,2} = \dots = v_{n,M} = 0$ for every advertiser n . In particular, Proposition 3 gives a $4/5$ GFP PNE PoA for the unit-demand setting, whereas there exists a $1 - 1/e$ PNE upper-bound construction in the multi-item setting (Theorem 1). In contrast, for GSP and VCG, poor welfare already arises with one ad per advertiser and one slot (Theorem 2).

Proof Overview. With unit-demand, welfare can be lost at PNE through pairs of advertisers in which a VM advertiser is allocated a slot with a high CTR where a higher-valuation QL advertiser is allocated a worse slot. The QL advertiser could take this slot by bidding just above the VM advertiser, but she does not improve her utility from doing so as per the PNE constraint. Hence the VM advertiser’s bid (and thus her valuation) is at least the QL advertiser’s valuation times 1 minus the CTR of the worse slot. If that CTR is high, the QL advertiser loses little in the worse slot, and if it is low, the VM advertiser’s valuation is close to hers. Each such pair therefore keeps at least $4/5$ of optimal welfare. For the matching upper bound, take two slots with CTRs 1 and $1/2$. A VM advertiser with valuation about $1/2$ wins the top slot by bidding her valuation, while a QL advertiser with valuation 1 takes the second slot with a bid near 0 that still beats a third advertiser with a negligible valuation. Outbidding the VM advertiser would cost the QL advertiser about $1/2$ and not raise her utility, so welfare is about $(1/2 + 1/2)/(1 + 1/4) = 4/5$ of the optimum. We prove Proposition 3 in EC.4.1. $\square$

Homogeneous objectives. GFP keeps a larger fraction of the optimal welfare when advertisers are all-QL or all-VM.

PROPOSITION 4. $\text{PoA}_{\text{CCE}}(\text{GFP}) = 1/2$ if advertisers are all-VM, and $\text{PoA}_{\text{CCE}}(\text{GFP}) = 1 - 1/e$ if all-QL.

Since the bound $0.4569 < 1/2 < 1 - 1/e$ of Theorem 1 is tight, mixing the two objectives is what lowers the guarantee of GFP. Together with the correlated equilibria of EC.3.2, the constants $1/2$ and 0.4569 are therefore tight for both CCE and CE. The all-QL bound $1 - 1/e$ is also tight for CCE. In our construction (Feldman et al. 2016b), two QL advertisers compete for one slot. With probability close to $1/e$, both bid little and the advertiser with the much lower valuation wins. Otherwise, she bids just below the other advertiser and loses. These losing bids make it costly for the other advertiser to win with a lower bid.

Proof Overview. When advertisers are all-VM, the bounds in the proof of Theorem 1 give one half. Specifically, in Equation (7), we have $U_{\text{QL}} = V_{\text{QL}}^* = 0$ in the all-VM setting, and this gives $\mathbb{E}[\text{Wel}] \geq \text{OPT}/2$. When advertisers are all-QL, the same deviation bound, with every target equal to $1 - 1/e$, gives a CCE guarantee of $1 - 1/e$. We prove Proposition 4 in EC.4.3. $\square$

5.2. Equilibrium Support

Our main results allow every multiplier in $[0, \infty)$, and the low-welfare equilibria of GSP and VCG in Theorem 2 use both very small and very large multipliers. To better understand the effect of the equilibrium

support, we first restrict equilibrium multipliers to an interval and further restrict by excluding weakly dominated multipliers as prior work does.

5.2.1. Floor and Ceiling. Fix a *floor* $A_{\min}$ and a *ceiling* $A_{\max}$ with $0 < A_{\min} \leq 1 \leq A_{\max}$. We emphasize that these parameters bound the equilibrium multipliers, not the underlying action space $[0, \infty)$. Here, we consider only equilibria in which every multiplier lies in the *equilibrium support* $[A_{\min}, A_{\max}]$ but deviations are unrestricted (i.e. to $[0, \infty)$). Defining the PoA with an equilibrium floor and ceiling $[A_{\min}, A_{\max}]^N$, we give the following support-dependent PoA guarantee:

PROPOSITION 5. *With equilibrium support $[A_{\min}, A_{\max}]$, $\text{PoA}_{\text{PNE}}(\text{GFP}) = 1 - (1 - A_{\min})/e$, and as $A_{\max}^{-1}, A_{\min} \rightarrow 0$,*

$$\text{PoA}_{\text{CCE}}(\text{GFP}) = 0.4569 + O(A_{\max}^{-1} + A_{\min}), \quad \text{PoA}(\text{GSP}), \text{PoA}(\text{VCG}) = \Theta(A_{\max}^{-1} + A_{\min}).$$

Proof Overview. Proposition 5 states that the welfare guarantees worsen as the equilibrium support $[A_{\min}, A_{\max}]$ grows. The intuition behind this is exemplified in the PNE construction from the proof overview of Theorem 2. Here, a high-valuation advertiser uses the floor $A_{\min}$ and loses to a low-valuation advertiser bidding larger than the former's valuation.

- The larger the value of $A_{\max}$, the higher the low-valuation advertiser can bid. This makes outcompeting her ROI infeasible or unprofitable for high-valuation advertisers.
- The smaller the value of $A_{\min}$, the lower the price the low-valuation advertiser needs to pay to win the item when the high-valuation advertiser reports multiplier $A_{\min}$. This enables increasingly lower-valuation advertisers to outcompete high-valuation ones without violating ROI feasibility/individual rationality.

For GFP, a similar floor dependence arises because a QL advertiser who bids far below her valuations keeps a large margin on every slot she wins, so she may forgo slots for which her valuations are high. A floor caps this margin at $1 - A_{\min}$, which raises the PNE guarantee to $1 - (1 - A_{\min})/e$. We give the full proof as well as the order-wise matching lower bound (Lemma EC.4) in EC.4.4. $\square$

5.2.2. Undominated Multipliers. In a seemingly conflicting set of results in prior literature and some of our own, Aggarwal et al. (2019) and Dütting et al. (2026) show a GSP and VCG welfare guarantee of $1/2$, whereas Deng et al. (2024a) and Theorem 2 show that GSP and VCG can yield 0 welfare. What explains this discrepancy is that the aforementioned two papers assume that underbidding (reporting $\alpha < 1$) is weakly dominated and thus restrict multipliers to be at least 1. This changes the welfare guarantees as per our support-dependent bounds (Proposition 5 and Lemma EC.4). To better understand why they make this assumption regarding underbidding domination, we first formalize the notion of dominance. We say that

advertiser n 's multiplier α is weakly dominated if there exists some α' that does at least as well as α against every profile α_{-n} of their competitors' multipliers. If advertiser n is QL, this means that

$$u_n(\alpha', \alpha_{-n}) \geq u_n(\alpha, \alpha_{-n}) \quad \text{for every } \alpha_{-n}.$$

If advertiser n is VM, her ROI constraint requires $\mathbb{E}[u_n(\alpha)] \geq 0$, where $u_n(\alpha) = V_n(\alpha) - P_n(\alpha)$ is her ROI slack or deficit (Section 3.5). Then α' weakly dominates α if it weakly increases both her value and u_n ,

$$V_n(\alpha', \alpha_{-n}) \geq V_n(\alpha, \alpha_{-n}) \quad \text{and} \quad u_n(\alpha', \alpha_{-n}) \geq u_n(\alpha, \alpha_{-n}) \quad \text{for every } \alpha_{-n}.$$

The next lemma identifies multipliers that are weakly dominated in every instance.

LEMMA 1. *In every instance, with any mix of advertiser objectives, multiplier one weakly dominates every QL multiplier above one under GFP and GSP. Under VCG, it weakly dominates every other QL multiplier and every VM multiplier below one.*

We prove Lemma 1 in EC.4.6. Under GFP and GSP, we cannot exclude VM multipliers below one in the same way, because their ROI slack may support larger multipliers within a randomized strategy, as Example 2 shows. This implies that the no-underbidding restriction on GSP in Dütting et al. (2026) is *not* without loss of generality: it fundamentally disallows deviations to multipliers that can be strictly profitable. To that end, we reconcile the results by also considering welfare guarantees when the weakly dominated multipliers of Lemma 1 are prohibited. We say that advertisers use *undominated multipliers*—the multipliers that Lemma 1 does not exclude.

PROPOSITION 6. *With undominated multipliers, the following hold.*

- $\text{PoA}_{\text{CCE}}(\text{VCG}) = \text{PoA}_{\text{PNE}}(\text{VCG}) = 1/2$ with tightness at an all-VM instance.
- The PoA guarantees for CCE and PNE of both GFP (Theorem 1 and Proposition 4) and GSP (Theorem 2 and Proposition 1) are unchanged.

Unlike GFP and GSP, removing dominated multipliers improves VCG's guarantee. With undominated multipliers, its PoA is $1/2$, strictly above GFP's guarantee of 0.4569. We prove Proposition 6 in EC.4.7.

5.3. Stochastic and Time-Varying Valuations

Our main model assumes that advertisers' valuations are fixed. We now let valuations be random, and then extend to the repeated setting where valuations may change from round to round.

5.3.1. Stochastic Valuations. We show in this section that our welfare guarantees extend to random, possibly correlated valuations. Let each advertiser's *valuation vector* $\mathbf{v}_n \doteq (v_{n,1}, \dots, v_{n,M})$ be random, and let the *valuation profile* $\mathbf{v} \doteq (\mathbf{v}_1, \dots, \mathbf{v}_N)$ have an arbitrary joint distribution with finite expected optimal welfare $\mathbb{E}[\text{OPT}(\mathbf{v})]$, where $\text{OPT}(\mathbf{v})$ is the optimal welfare at valuations $\mathbf{v}$ (Section 3.2). Advertiser n

privately observes $\mathbf{v}_n$ and chooses her multiplier by a *strategy* π_n , which maps each of her valuation vectors to a multiplier. A profile of strategies $\pi \doteq (\pi_1, \dots, \pi_N)$ then gives the multiplier profile $\alpha \doteq \pi(\mathbf{v})$, with $\alpha_n = \pi_n(\mathbf{v}_n)$. We write $V_n(\mathbf{v}, \alpha)$, $P_n(\mathbf{v}, \alpha)$, and $u_n(\mathbf{v}, \alpha)$ when the dependence on valuations matters.

We consider Bayesian CCE, in which each QL advertiser maximizes expected utility over unilateral changes to her strategy, and each VM advertiser maximizes expected acquired value over changes that satisfy her *ex-ante* ROI constraint which necessitates that ROI is satisfied only *after* averaging over their random valuations.

DEFINITION 5 (BAYESIAN CCE). Let $\pi \doteq (\pi_1, \dots, \pi_N)$ be a random profile of strategies, independent of $\mathbf{v}$, and let $\alpha \doteq (\pi_1(\mathbf{v}_1), \dots, \pi_N(\mathbf{v}_N))$. Advertiser n satisfies *ex-ante* ROI if $\mathbb{E}_{\mathbf{v}, \pi}[P_n(\mathbf{v}, \alpha)] \leq \mathbb{E}_{\mathbf{v}, \pi}[V_n(\mathbf{v}, \alpha)]$, where the expectations average over both valuations and strategies. Then the distribution of π is a *Bayesian CCE* if every VM advertiser satisfies ex-ante ROI, and

- (i) every QL advertiser n satisfies $\mathbb{E}_{\mathbf{v}, \pi}[u_n(\mathbf{v}, \alpha)] \geq \mathbb{E}_{\mathbf{v}, \pi}[u_n(\mathbf{v}, [\pi'_n(\mathbf{v}_n), \alpha_{-n}])]$ for every strategy π'_n , and
- (ii) every VM advertiser n satisfies $\mathbb{E}_{\mathbf{v}, \pi}[V_n(\mathbf{v}, \alpha)] \geq \mathbb{E}_{\mathbf{v}, \pi}[V_n(\mathbf{v}, [\tilde{\pi}_n(\mathbf{v}_n), \alpha_{-n}])]$, with the right side also averaging over $\tilde{\pi}_n$, for every strategy $\tilde{\pi}_n$ drawn at random, independently of $\mathbf{v}$ and π , with which she satisfies ex-ante ROI.

Now, we importantly note that the advertisers' multipliers may be correlated, but each advertiser's multiplier and her deviations depend only on her own valuations. This ensures that a Bayesian CCE with fixed valuations corresponds to our standard notion of a CCE.

Bayesian PoA. Replacing the valuations of an instance by a joint distribution of valuations gives a *Bayesian instance*. The *Bayesian PoA* of a mechanism $\mathcal{M}$ is

$$\text{PoA}_{\text{BCCE}}(\mathcal{M}) \doteq \inf_{I, \pi} \frac{\mathbb{E}_{\mathbf{v}, \pi}[\text{Wel}(\mathbf{v}, \alpha)]}{\mathbb{E}_{\mathbf{v}, \pi}[\text{OPT}(\mathbf{v})]}, \quad (8)$$

where the infimum is over all Bayesian instances I with $\mathbb{E}[\text{OPT}(\mathbf{v})] > 0$ and all random strategy profiles π whose distribution is a Bayesian CCE of I under $\mathcal{M}$, and the expectations average over valuations and strategies, with $\alpha = \pi(\mathbf{v})$. This is the PoA of Definition 4 with expected welfare in both the numerator and the denominator. With fixed valuations, Bayesian CCE and CCE coincide, so $\text{PoA}_{\text{BCCE}}(\mathcal{M}) \leq \text{PoA}_{\text{CCE}}(\mathcal{M})$. Despite this, we surprisingly can show that our CCE welfare guarantees still hold at BCCE.

COROLLARY 1. *The CCE lower bounds of Theorem 1 and Propositions 4 and 6 hold for PoA_{BCCE} . In particular, $\text{PoA}_{\text{BCCE}}(\text{GFP}) = 0.4569$. Moreover, $\text{PoA}_{\text{BCCE}}(\text{VCG}) = 1/2$ with undominated multipliers.*

For example, GFP keeps at least $0.4569 \mathbb{E}[\text{OPT}(\mathbf{v})]$, and VCG with undominated multipliers keeps at least $\mathbb{E}[\text{OPT}(\mathbf{v})]/2$, even though i) valuations may be correlated and ii) the benchmark is ex-post. This robustness is surprising. The closest related works in Table 2 consider fixed valuations only, and full-information welfare guarantees typically extend to Bayesian settings only when valuations are independent (Roughgarden et al. 2017). Correlation is natural in our setting, since a query that signals strong purchase intent should generally

raise many of the advertisers' valuations. In contrast, the vanishing-welfare instances of GSP and VCG are attained already with fixed valuations, and thus trivially extend to this more general, stochastic setting.

Proof Overview. We note that always using the same multiplier is a valid deviation, and bidding valuations satisfies ROI in every realization, so the equilibrium conditions give the same equilibrium welfare lower bounds and optimal welfare upper bounds (in expectation) as in the CCE proof of Theorem 1. Averaging the inequalities of Lemmas EC.1 and EC.2 over valuations and play, and choosing the targets after averaging, gives the GFP bounds. Under VCG with undominated multipliers, the argument for Proposition 6 needs only the restriction on multipliers and expected ROI. We prove Corollary 1 in EC.4.8. $\square$

5.3.2. Repeated Auctions. Our results so far support a simple message. GFP guarantees a constant fraction of optimal welfare even when advertisers' multipliers become very small or large, whereas GSP and VCG can lose almost all of it. These guarantees concern equilibria and thus are only non-vacuous if such equilibria *exist* and in the repeated setting, can converge to them. In the absence of constraints (such as ROI constraints), these CCEs necessarily exist and no-regret algorithms approach them (Cesa-Bianchi and Lugosi 2006). Thus, these CCE bounds describe the welfare of learning in unconstrained settings. However, ROI constraints break this link (Mannor et al. 2009) and some instances have no CCE (Proposition EC.6). In such instances, vanilla no-regret learning has no equilibrium to approach, and its average play need not converge. The main result of this section is that we can *still* preserve our welfare guarantees with a suitable learning dynamic.

Suppose the auction is repeated for T rounds. In each round, each advertiser draws her multiplier from a finite grid of multipliers in $[A_{\min}, A_{\max}]$ that contains multiplier one, and observes counterfactual payoffs (full information setting).⁴ Advertisers learn with *primal-dual algorithms*, as formally defined in EC.4.10.3. At a high level, these primal-dual algorithms guarantee that on every history of play, the distance between average welfare and our CCE welfare guarantees is at most a function of the advertisers' primal and dual regret, which these algorithms drive to 0 as T grows.

Dynamic PoA. Learning need not converge to an equilibrium, so we evaluate the history of play by its average welfare. Let α_t be the multiplier profile in round t . For a mechanism $\mathcal{M}$ and a given floor and ceiling, the *dynamic PoA* is

$$\text{PoA}_{\text{dyn}}(\mathcal{M}) \doteq \inf_I \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[T^{-1} \sum_{t \leq T} \text{Wel}(\alpha_t)]}{\text{OPT}(I)}, \quad (9)$$

where the infimum is over all instances I with $\text{OPT}(I) > 0$, the profiles α_t are generated by the primal-dual algorithms, and the expectation is over their random draws. This is the PoA of Definition 4 with the equilibrium replaced by the history of play, so unlike the PoA, it does not require an equilibrium to exist. Definition EC.2 in EC.4.10.1 states it formally, including the effect of grid discretization error. Now, we are ready to state our main dynamic welfare guarantee.

⁴ Unlike in Section 5.2.1, the interval $[A_{\min}, A_{\max}]$ also restricts the multiplier deviations; not just the equilibrium multipliers.

THEOREM 3. *Assume all learners adopt primal–dual algorithms. Then the following hold.*

- $\text{PoA}_{\text{dyn}}(\text{GFP}) \geq 0.4569$. For all-VM: $\text{PoA}_{\text{dyn}}(\text{GFP}) \geq 1/2$ and for all-QL: $\text{PoA}_{\text{dyn}}(\text{GFP}) \geq 1 - 1/e$ (as in their static counterparts Theorem 1 and Proposition 4).
- Assuming undominated multipliers, $\text{PoA}_{\text{dyn}}(\text{VCG}) \geq 1/2$ (as in its static counterpart Proposition 6).

Moreover, on every instance, the ratio of expected average welfare over T rounds to OPT satisfies these bounds up to an additive error of $O(T^{-1/4})$.

Consistent with our static setting guarantees, both GFP and undominated VCG maintain their 0.4569 and $1/2$ welfare guarantees respectively. For example, under GFP, the expected average welfare is at least $0.4569 \text{OPT} - O(T^{-1/4})$. The bound holds on every instance, whether or not an equilibrium exists or play converges. In contrast, GSP and VCG can still perform badly as one may repeat in each round the low-welfare PNE from Proposition 5 (see Lemma EC.6).

Proof Overview. We recall the proof of the CCE lower bound in Section 4.1.1. Throughout, we fix a history of play and write $\mathbb{E}[X] \doteq T^{-1} \sum_{t \leq T} X_t$ for the time average of a quantity X_t over the T rounds. As in Section 4.1.1, $U_n \doteq \mathbb{E}[u_n(\alpha)]$, $V_n \doteq \mathbb{E}[V_n(\alpha)]$, and $P_n \doteq \mathbb{E}[P_n(\alpha)]$, now with this time average. This average need not be a CCE or approach one (as per equilibrium non-existence Proposition EC.6), so we first identify where the proof of Theorem 1 uses equilibrium conditions. We show that we can relax these conditions and that these relaxed requirements are satisfied whenever primal and dual regret are small.

Comparison with the Proof of Theorem 1. Lemma EC.1(i) holds at every profile, so the recovery inequality (6) holds in every round and hence on average, and the choice of targets is pure algebra. The equilibrium enters only through three inequalities, in the welfare decomposition and in the bounds on recoverable value.

- (i) A QL advertiser gains no utility from a fixed multiplier, $\mathbb{E}[u_n(\alpha', \alpha_{-n})] \leq U_n$.
- (ii) A VM advertiser gains no value by always bidding her valuations, $\mathbb{E}[V_n(1, \alpha_{-n})] \leq \mathbb{E}[V_n(\alpha)]$. Since acquired value is nondecreasing in her own multiplier, this also bounds her value under every deviation to a multiplier $\alpha' \leq 1$.
- (iii) A VM advertiser satisfies ROI on average, $\mathbb{E}[P_n(\alpha)] \leq \mathbb{E}[V_n(\alpha)]$.

In particular, this proof never uses the full CCE condition for a VM advertiser. Condition (ii) instead compares her with a single fixed strategy of deviating to 1 which satisfies her ROI constraints in every round regardless of what the other advertisers do.

Welfare Bounds with Learning Errors. We first introduce three key quantities: the QL regret R_n , the VM value shortfall S_n , and the VM net ROI deficit v_n , defined by

$$R_n \doteq \left[\max_{\alpha'} \mathbb{E}[u_n(\alpha', \alpha_{-n})] - U_n \right]_+, \quad S_n \doteq [\mathbb{E}[V_n(1, \alpha_{-n})] - V_n]_+, \quad v_n \doteq [P_n - V_n]_+, \quad (10)$$

where the maximum is over her grid multipliers. These measure how much (i), (ii), and (iii) are violated over time. We show that each step of Theorem 1's proof holds up to these errors. Since payments are nonnegative and $P_{VM} \leq V_{VM} + \nu_{VM}$, we have that Equation (5) becomes

$$\mathbb{E}[\text{Wel}] + \nu_{VM} \geq U_{QL} + \max\{\mathbb{E}[\text{Rev}], V_{VM}\}.$$

A QL deviation to $\alpha' < 1$ gives $(1 - \alpha')\mathbb{E}[V_n(\alpha', \alpha_{-n})] \leq U_n + R_n$, so her recoverable value is at most $-(U_n + R_n) \ln(1 - \bar{\alpha}_n)$, and a VM advertiser's recoverable value with target $\bar{\alpha}_n = 1$ is at most $\mathbb{E}[V_n(\alpha)] + S_n$. The revenue bounds of Lemma EC.2 therefore hold with $U_{QL} + R_{QL}$ and $V_{VM} + S_{VM}$ in place of U_{QL} and V_{VM} , and we obtain:

$$\mathbb{E}[\text{Wel}] + R_{QL} + S_{VM} + \nu_{VM} \geq (U_{QL} + R_{QL}) + \max\{\mathbb{E}[\text{Rev}], V_{VM} + S_{VM}\}.$$

Now, we finish off the proof by plugging this in instead of Equation (5) and performing the same optimization over $\bar{u}$ as in Theorem 1's proof yields

$$\mathbb{E}[\text{Wel}] \geq 0.4569 \text{OPT} - R_{QL} - S_{VM} - \nu_{VM}.$$

Note here that we used the fact that reporting multiplier 1 gives a QL advertiser utility 0, so $U_n + R_n \geq 0$ and the minimization over $\bar{u}$ is unchanged. The other CCE bounds of Sections 4 and 5 use the same three inequalities, so they transfer in the same way (Lemma EC.5). Since Lemma EC.1(i) holds at every valuation profile, the argument also applies to time-varying valuations. Lastly, the primal–dual algorithms make these errors vanish. Exponential weights guarantees each QL advertiser's expected regret $R_n = O(T^{-1/2})$. For a VM advertiser, the multiplier update has small regret on the penalized rewards (primal regret), and the penalty update has small regret against every fixed penalty (dual regret). Lemma EC.7 shows that these two regrets bound S_n and, through the cap $T^{1/4}$ on the penalty, ν_n , which gives the rate $T^{-1/4}$. We give the details and full proof in EC.4.10. $\square$

Time-varying valuations. We now the formal extension of our learning guarantees to the setting where valuations change over time, possibly even in an adversarial manner. In particular, these each round's valuations may depend on each other and respond to past play. In round t , advertiser n has valuation vector $\mathbf{v}_{n,t}$, the valuation profile is $\mathbf{v}_t \doteq (\mathbf{v}_{1,t}, \dots, \mathbf{v}_{N,t})$, and $\text{OPT}_t \doteq \text{OPT}(\mathbf{v}_t)$ is that round's optimal welfare.

COROLLARY 2. *Theorem 3 extends to time-varying valuations, even adversarial ones, with OPT replaced by $\mathbb{E}[T^{-1} \sum_{t \leq T} \text{OPT}_t]$. For example, under GFP,*

$$\mathbb{E}\left[\frac{1}{T} \sum_{t \leq T} \text{Wel}(\mathbf{v}_t, \alpha_t)\right] \geq 0.4569 \mathbb{E}\left[\frac{1}{T} \sum_{t \leq T} \text{OPT}_t\right] - O(T^{-1/4}).$$

As in Theorem 3, these bounds need neither equilibrium existence nor convergence, and Theorem 3 is the special case of fixed valuations. Furthermore, to our knowledge, the closest related works on autobidding (Table 2) consider fixed valuations only. Our guarantees not only allow for correlated valuations, but even adversarially generated ones. We prove Corollary 2 in EC.4.10.4.

6. Experiments

Our results thus far have been purely theoretical and concern the worst case: GSP and VCG can lose almost all welfare, at equilibrium and under learning dynamics (Theorem 2, Proposition 5, and Lemma EC.6), whereas GFP keeps a constant 0.4569 fraction of optimal welfare (Theorems 1 and 3). The theory does *not* show whether learning dynamics actually reach the low-welfare outcomes. We thus simulate the primal–dual algorithms (Theorem 3) and compare the welfare and revenue of GFP, GSP, and VCG. Across all tables and figures in this section, GFP’s first-percentile welfare consistently outperforms GSP and VCG with the gap largest with $M = 1$ slot. With multipliers in $[0.1, 10]$, GSP and VCG’s first-percentile welfare lies near 0, whereas GFP’s is at least .64—considerably larger than its 0.4569 worst case guarantee. Moreover, as a curious side observation, we also see that GFP consistently obtains more revenue than GSP and VCG throughout.

6.1. Experimental Design

We run several experiments in order to show robustness of our theoretical findings across different settings and parameterizations. All experiments run the primal–dual algorithms of Theorem 3 for $T = 10^6$ rounds, with the implementation settings in EC.5.3, and normalize by optimal welfare. Since our theoretical results are with respect to worst case welfare, we report the 1st percentile welfare across trials.

- **Baseline comparison** (Section 6.2). We compare both the welfare and revenue of our mechanisms for mixed, all-VM, and all-QL advertiser populations, with varying numbers of slots.
- **Effect of equilibrium support** (Section 6.3). Proposition 5 shows that the welfare guarantees worsen as the equilibrium support grows larger. We test this directly by varying $A_{\min}$ and $A_{\max}$.
- **More competition** (Section 6.4). We double the number of advertisers to test whether more competition removes the low-welfare outcomes of GSP and VCG as it is known that inefficiency tends to fade as markets grow (Feldman et al. 2016a).
- **Robustness to stochastic valuations** (Section 6.5). Corollary 2 extends our learning guarantees to time-varying valuations. We test whether such randomness pushes GSP and VCG out of their worst-case outcomes.
- **Additional experiments** In Section 6.6, we briefly describe two further experiments with stronger learning dynamics (no-swap regret learning) and with undominated multipliers.

6.2. Baseline Comparison

Our main suite of experiments fixes $A_{\min} = 0.1$ and $A_{\max} = 10$. The baseline experiment has two advertisers, $M \in \{1, 3, 5, 10\}$ identical slots, and independently drawn lognormal valuations. Populations are mixed, all-VM, or all-QL, and a mixed instance has one advertiser of each type. Advertisers use the primal–dual algorithms, whose exponential-weights updates minimize external regret. We run 200 trials per configuration.


Table 3 reports first-percentile welfare and revenue for $M \in \{1, 3, 10\}$. The complete tables in EC.5 also include $M = 5$ and the 5th and 25th percentiles for every experiment, and EC.5.2 and EC.5.3 give the full experimental breakdown.

Welfare					Revenue				
Population	Mechanism	$M = 1$	$M = 3$	$M = 10$	Population	Mechanism	$M = 1$	$M = 3$	$M = 10$
Mixed	GFP	.65	.80	.85	Mixed	GFP	.10	.14	.33
	GSP	.14	.32	.60		GSP	.01	.02	.07
	VCG	.14	.32	.64		VCG	.01	.02	.03
All-VM	GFP	1	.96	.98	All-VM	GFP	.10	.62	.91
	GSP	.14	.47	.67		GSP	.01	.02	.08
	VCG	.14	.44	.67		VCG	.01	.02	.04
All-QL	GFP	.93	.92	.97	All-QL	GFP	.10	.10	.10
	GSP	.14	.44	.71		GSP	.01	.02	.06
	VCG	.14	.32	.67		VCG	.01	.02	.03

Table 3 Welfare and revenue (1st percentile) under primal–dual learning dynamics.

In Table 3, we see that GFP consistently obtains high welfare—at least .65 which is even stronger than the PNE worst case guarantee of $1 - 1/e \approx 0.63$. In contrast, GSP and VCG have low welfare. With one slot, both have first-percentile welfare and revenue close to 0, though we note that they improve with the number of slots. Interestingly, GFP consistently also earns more revenue than GSP or VCG.

6.3. Effect of Equilibrium Support



Note. Simulations of learning dynamics with multipliers in $[A_{\max}^{-1}, A_{\max}]$, showing percentiles of revenue divided by OPT on a logarithmic scale across 100 trials. Lines mark the 1st percentile, and shading extends to the 25th. GFP revenue is consistently far higher than GSP and VCG revenue at every value of $A_{\max}$. With one slot, GSP and VCG coincide. Details are in EC.5.1.

We vary the range of multipliers $[A_{\max}^{-1}, A_{\max}]$ from $A_{\max} = 1$, where every advertiser bids her valuations, to $A_{\max} = 64$. Each market has one QL and one VM advertiser and identical slots, each with CTR one. With one slot, we compare GFP with GSP, which then coincides with VCG. With two slots, we compare GFP, GSP, and VCG. We run 100 trials per configuration (EC.5.1).

As Proposition 5 predicts, the welfare of GSP and VCG falls as the range widens (Figure 2). Their first-percentile welfare falls roughly like $A_{\max}^{-1}$, to .04 at $A_{\max} = 64$. In contrast, GFP’s first-percentile welfare stays above .8. GFP’s revenue is also far higher at every value of $A_{\max}$ (Figure 3).

6.4. More Competition

To model more competition, we increase the number of advertisers from our baseline experiments (Section 6.2) from two to four, so a mixed population has two advertisers of each type.

Welfare					Revenue				
Population	Mechanism	$M = 1$	$M = 3$	$M = 10$	Population	Mechanism	$M = 1$	$M = 3$	$M = 10$
Mixed	GFP	.64	.84	.86	Mixed	GFP	.10	.23	.41
	GSP	.12	.37	.78		GSP	.03	.09	.38
	VCG	.12	.37	.64		VCG	.03	.07	.23
All-VM	GFP	1	.96	.98	All-VM	GFP	.18	.56	.93
	GSP	.13	.64	.87		GSP	.03	.12	.84
	VCG	.13	.37	.61		VCG	.03	.08	.21
All-QL	GFP	.90	.87	.92	All-QL	GFP	.10	.13	.16
	GSP	.12	.33	.72		GSP	.02	.09	.15
	VCG	.12	.28	.72		VCG	.02	.07	.15

Table 4 Welfare and revenue (1st percentile) with four advertisers.

More competition does not eliminate the low welfare of GSP and VCG (Table 4). With one slot, GSP and VCG’s 1st percentile welfare still remain near 0 whereas GFP’s. Just as in our baseline experiments, the gap between GFP versus GSP or VCG decreases with the number of slots. The revenue gap is also somewhat smaller than with two advertisers, demonstrating the effect of more competition increasing the corresponding payments in GSP and VCG.

6.5. Robustness to Stochastic Valuations

Corollary 2 extends our learning guarantees to time-varying valuations. We test whether such randomness pushes GSP and VCG out of their worst-case outcomes, using the same baseline experiments in Section 6.2. The only difference is that in every round, we multiply each ad’s original valuation by an independent uniform draw from $[0.75, 1.25]$. Valuations therefore fluctuate around the same underlying valuations. We normalize welfare and revenue by the realized optimum over the same rounds (see EC.5.3 for details).

We see that randomness fails to save GSP and VCG from their worst-case outcomes (Table 5). With one slot, their first-percentile welfare is still near 0, compared with at least the massively improved .85 under GFP—which we recall was around 0.65 with deterministic valuations. Additionally, similar to our previous experiments, GFP still earns more revenue consistently.

Welfare					Revenue				
Population	Mechanism	$M = 1$	$M = 3$	$M = 10$	Population	Mechanism	$M = 1$	$M = 3$	$M = 10$
Mixed	GFP	.85	.81	.87	Mixed	GFP	.10	.15	.22
	GSP	.17	.50	.71		GSP	.00	.02	.11
	VCG	.17	.46	.68		VCG	.00	.01	.05
All-VM	GFP	1	.97	.99	All-VM	GFP	.18	.67	.79
	GSP	.17	.59	.90		GSP	.00	.02	.65
	VCG	.17	.51	.73		VCG	.00	.01	.14
All-QL	GFP	.91	.93	.98	All-QL	GFP	.10	.10	.10
	GSP	.17	.50	.74		GSP	.00	.02	.06
	VCG	.17	.44	.74		VCG	.00	.01	.05

Table 5 Welfare and revenue (1st percentile) with stochastic valuations.

6.6. Additional Experiments

The e-companion includes two further experiments. First, we test whether stronger learning dynamics avoid low welfare by replacing each learner with a swap-regret version (Blum and Mansour 2007), described in EC.5.3. We find that it does not, and the results nearly match Table 3 (Table EC.8). Second, with undominated multipliers (Section 5.2.2), VCG’s welfare improves while GSP’s barely does, as Proposition 6 predicts. That being said, GFP still seems to outperform it in most settings, even without restricting GFP to undominated multipliers (Table EC.9).

7. Conclusion

In this paper, we studied sponsored-shopping auctions where advertisers have multi-item demand and have varying objectives. We establish the following results.

- **GFP retains constant welfare guarantees.** In Theorem 1, we prove a CCE PoA of at least 0.4569 and a PNE PoA of exactly $1 - 1/e$. When advertisers are all-VM or all-QL, the CCE PoA improves to exactly $1/2$ or $1 - 1/e$, respectively (Proposition 4).
- **GSP and VCG do not.** In Theorem 2, we prove that their PNE PoA is 0, already in single-slot auctions with advertisers who share the same objective.
- **The main welfare guarantees do not need equilibria.** Despite potentially not possessing any equilibria for learning dynamics to converge towards (Proposition EC.6), we are able to show that in repeated auctions, a natural family of primal–dual learning algorithms achieves our 0.4569 welfare guarantee for GFP with vanishing error. These guarantees also extend to stochastic, possibly correlated valuations and to valuations that vary (possibly adversarially) across auctions (Corollaries 1 and 2).
- **Simulations.** In the experiments in Section 6, we observe high welfare and revenue under GFP, while GSP and VCG consistently demonstrate low welfare and revenue across various settings.

7.1. Managerial Insights

With the interpretation of welfare as the relevance and quality of sponsored results, it is of utmost importance to both the platform’s advertisers and users that the welfare remains robust. Our results have several immediate implications for how platforms should design and manage their autobidding systems.

Choice of mechanism. GFP is the most robust of the three mechanisms to very small or large multipliers. It retains a constant welfare guarantee even when some multipliers are very small and others very large, whereas GSP and VCG can yield poor welfare even with one slot and advertisers who share the same objective (Theorem 2). We confirm this in our simulations, with both GSP and VCG producing low welfare. This GFP welfare guarantee superiority over GSP and VCG complement the many practical reasons for first-price over second-price type mechanisms, including simplicity and price transparency (Cox 2019, Federal Trade Commission 2026, Wong 2021).

Steady spending. An autobidder spreads a campaign’s budget over its duration. The multiplier submitted in one auction can therefore reflect both the advertiser’s ROI requirement and her remaining budget. If the campaign spends more than planned early on, the autobidder may need to conserve the depleted budget by raising its ROI target and shrinking its later multipliers. If it spends less than planned, the autobidder may instead lower that target and increase its multipliers to deplete the remaining budget. Such corrections can lead to very small or large multipliers, which we have demonstrated can lead to worsened welfare. Thus, the platform must ensure that expenditure over time stays close to the advertiser’s intended expenditure throughout the campaign.

Open questions. Two directions seem especially relevant for practice. First, bounding revenue in equilibrium could explain why GFP earns more revenue than GSP and VCG in our simulations. Second, modeling budget constraints directly alongside ROI constraints is a natural next step.

References

- Aggarwal G, Badanidiyuru A, Balseiro SR, Bhawalkar K, Deng Y, Feng Z, Goel G, Liaw C, Lu H, Mahdian M, Mao J, Mehta A, Mirrokni V, Paes Leme R, Perloff A, Piliouras G, Schneider J, Schwartzman A, Sivan B, Spendlove K, Teng Y, Wang D, Zhang H, Zhao M, Zhu W, Zuo S (2024) Auto-bidding and auctions in online advertising: A survey. *ACM SIGecom Exchanges* 22(1):159–183, URL <http://dx.doi.org/10.1145/3699824.3699838>.
- Aggarwal G, Badanidiyuru A, Mehta A (2019) Autobidding with constraints. *Web and Internet Economics (WINE 2019)*, volume 11920 of *Lecture Notes in Computer Science*, 17–30 (Springer), URL http://dx.doi.org/10.1007/978-3-030-35389-6_2.
- Alimohammadi Y, Mehta A, Perloff A (2023) Incentive compatibility in the auto-bidding world. *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 63 (Association for Computing Machinery), URL <http://dx.doi.org/10.1145/3580507.3597725>.
- Athey S, Ellison G (2011) Position auctions with consumer search. *The Quarterly Journal of Economics* 126(3):1213–1270.
- Ausubel LM, Cramton P, Pycia M, Rostek M, Weretka M (2014) Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies* 81(4):1366–1400, URL <http://dx.doi.org/10.1093/restud/rdu023>.
- Balseiro S, Golrezaei N, Mahdian M, Mirrokni V, Schneider J (2019) Contextual bandits with cross-learning. *Advances in Neural Information Processing Systems*, volume 32.

- Balseiro S, Kroer C, Kumar R (2023) Contextual standard auctions with budgets: Revenue equivalence and efficiency guarantees. *Management Science* 69(11):6837–6854, URL <http://dx.doi.org/10.1287/mnsc.2023.4719>.
- Balseiro SR, Deng Y, Mao J, Mirrokni V, Zuo S (2021) The landscape of auto-bidding auctions: Value versus utility maximization. *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 132–133 (Association for Computing Machinery), URL <http://dx.doi.org/10.1145/3465456.3467607>.
- Balseiro SR, Gur Y (2019) Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65(9):3952–3968, URL <http://dx.doi.org/10.1287/mnsc.2018.3174>.
- Birmpas G, Markakis E, Telelis O, Tsikiris A (2019) Tight welfare guarantees for pure Nash equilibria of the uniform price auction. *Theory of Computing Systems* 63(7):1451–1469, URL <http://dx.doi.org/10.1007/s00224-018-9889-7>.
- Blum A, Mansour Y (2007) From external to internal regret. *Journal of Machine Learning Research* 8:1307–1324, URL <https://www.jmlr.org/papers/v8/blum07a.html>.
- Booking Holdings Inc (2025) Annual report on Form 10-K for the year ended December 31, 2024. U.S. Securities and Exchange Commission, URL <https://www.sec.gov/Archives/edgar/data/1075531/000107553125000010/bkng-20241231.htm>.
- Brânzei S, Derakhshan M, Golrezaei N, Han Y (2023) Learning and collusion in multi-unit auctions. *Advances in Neural Information Processing Systems*, volume 36, 22191–22225.
- Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2011) On the efficiency of equilibria in generalized second price auctions. *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, 81–90 (Association for Computing Machinery), URL <http://dx.doi.org/10.1145/1993574.1993588>.
- Castiglioni M, Celli A, Kroer C (2022) Online learning with knapsacks: the best of both worlds. *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 2767–2783.
- Castiglioni M, Celli A, Kroer C (2024) Online learning under budget and ROI constraints via weak adaptivity. *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 5792–5816.
- Cesa-Bianchi N, Lugosi G (2006) *Prediction, Learning, and Games* (Cambridge: Cambridge University Press).
- Chen X, Kroer C, Kumar R (2024) The complexity of pacing for second-price auctions. *Mathematics of Operations Research* 49(4):2109–2135.
- Chen Y, He C (2011) Paid placement: Advertising and search on the internet. *The Economic Journal* 121(556):F309–F328.
- Colini-Baldeschi R, Klumper S, Kroll T, Leonardi S, Schäfer G, Tsikiris A (2026) Optimal type-dependent liquid welfare guarantees for autobidding agents with budgets. *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1795–1823, URL <http://dx.doi.org/10.1137/1.9781611978971.64>.
- Conitzer V, Kroer C, Panigrahi D, Schrijvers O, Stier-Moses NE, Sodomka E, Wilkens CA (2022a) Pacing equilibrium in first price auction markets. *Management Science* 68(12):8515–8535, URL <http://dx.doi.org/10.1287/mnsc.2022.4310>.
- Conitzer V, Kroer C, Sodomka E, Stier-Moses NE (2022b) Multiplicative pacing equilibria in auction markets. *Operations Research* 70(2):963–989, URL <http://dx.doi.org/10.1287/opre.2021.2167>.
- Cox S (2019) Simplifying programmatic: first price auctions for Google Ad Manager. Google Ad Manager, URL <https://blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager/>.
- Craswell N, Zoeter O, Taylor M, Ramsey B (2008) An experimental comparison of click position-bias models. *Proceedings of the 2008 International Conference on Web Search and Data Mining (WSDM)*, 87–94 (ACM), URL <http://dx.doi.org/10.1145/1341531.1341545>.
- de Keijzer B, Markakis E, Schäfer G, Telelis O (2013) Inefficiency of standard multi-unit auctions. *Algorithms – ESA 2013*, volume 8125 of *Lecture Notes in Computer Science*, 385–396 (Springer), URL http://dx.doi.org/10.1007/978-3-642-40450-4_33.
- Deng Y, Golrezaei N, Jaillet P, Liang JCN, Mirrokni V (2023) Multi-channel autobidding with budget and ROI constraints. *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 7617–7644.

- Deng Y, Mahdian M, Mao J, Mirrokni V, Zhang H, Zuo S (2024a) Efficiency of the generalized second-price auction for value maximizers. *Proceedings of the ACM Web Conference 2024 (WWW)*, 46–56, URL <http://dx.doi.org/10.1145/3589334.3645360>.
- Deng Y, Mao J, Mirrokni V, Zhang H, Zuo S (2024b) Efficiency of the first-price auction in the autobidding world. *Advances in Neural Information Processing Systems*, volume 37, 139270–139293, URL <http://dx.doi.org/10.52202/079017-4420>.
- Dütting P, Li Y, Paes Leme R, Spendlove K, Teng Y (2026) Autobidding equilibria in sponsored shopping. arXiv preprint arXiv:2602.21966, URL <https://arxiv.org/abs/2602.21966v2>.
- Edelman B, Ostrovsky M, 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, URL <http://dx.doi.org/10.1257/aer.97.1.242>.
- Elzayn H, Colini Baldeschi R, Lan B, Schrijvers O (2022) Equilibria in auctions with ad types. *Proceedings of the ACM Web Conference 2022*, 68–78, URL <http://dx.doi.org/10.1145/3485447.3512052>.
- Engelbrecht-Wiggans R, Kahn CM (1998a) Multi-unit auctions with uniform prices. *Economic Theory* 12(2):227–258, URL <http://dx.doi.org/10.1007/s001990050220>.
- Engelbrecht-Wiggans R, Kahn CM (1998b) Multi-unit pay-your-bid auctions with variable awards. *Games and Economic Behavior* 23(1):25–42, URL <http://dx.doi.org/10.1006/game.1997.0599>.
- Fabra N, von der Fehr NH, Harbord D (2006) Designing electricity auctions. *The RAND Journal of Economics* 37(1):23–46, URL <http://dx.doi.org/10.1111/j.1756-2171.2006.tb00002.x>.
- Federal Trade Commission (2026) FTC, states sue Amazon over secret ad surcharge scheme. Press release, URL <https://www.ftc.gov/news-events/news/press-releases/2026/08/ftc-states-sue-amazon-over-secret-ad-surcharge-scheme>.
- Feldman M, Immorlica N, Lucier B, Roughgarden T, Syrgkanis V (2016a) The price of anarchy in large games. *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 963–976, URL <http://dx.doi.org/10.1145/2897518.2897580>.
- Feldman M, Lucier B, Nisan N (2016b) Correlated and coarse equilibria of single-item auctions. Cai Y, Vetta A, eds., *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*, volume 10123 of *Lecture Notes in Computer Science*, 131–144 (Springer), URL http://dx.doi.org/10.1007/978-3-662-54110-4_10.
- Feng Y, Golrezaei N, Li Q (2026) Nash equilibria in uniform price auctions: Theory, computation, and market insights. ACM SIGMETRICS 2026; SSRN Working Paper 6175080, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6175080.
- Feng Y, Lucier B, Slivkins A (2024) Strategic budget selection in a competitive autobidding world. *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, 213–224 (Association for Computing Machinery), URL <http://dx.doi.org/10.1145/3618260.3649688>.
- Fikioris G, Tardos É (2025) Liquid welfare guarantees for no-regret learning in sequential budgeted auctions. *Mathematics of Operations Research* 50(2):1233–1249.
- Gaitonde J, Li Y, Light B, Lucier B, Slivkins A (2023) Budget pacing in repeated auctions: Regret and efficiency without convergence. *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics*, 52:1–52:1 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik), URL <http://dx.doi.org/10.4230/LIPIcs.ITCS.2023.52>.
- Galgana R, Golrezaei N (2025) Learning in repeated multiunit pay-as-bid auctions. *Manufacturing & Service Operations Management* 27(1):200–229, URL <http://dx.doi.org/10.1287/msom.2023.0403>.
- Galgana R, Golrezaei N (2026) Gap based learning in markets. SSRN working paper 6222981, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6222981.
- Giannou A, Vlatakis-Gkaragkounis EV, Mertikopoulos P (2021) Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information. Belkin M, Kpotufe S, eds., *Proceedings of the 34th Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, 2147–2148 (PMLR).

- Golrezaei N, Jaillet P, Liang JCN, Mirrokni V (2021a) Bidding and pricing in budget and ROI constrained markets. arXiv preprint arXiv:2107.07725, URL <https://arxiv.org/abs/2107.07725v1>.
- Golrezaei N, Lobel I, Paes Leme R (2021b) Auction design for ROI-constrained buyers. *Proceedings of the Web Conference 2021 (WWW)*, 3941–3952, URL <http://dx.doi.org/10.1145/3442381.3449841>.
- Golrezaei N, Sahoo S (2025) Learning safe strategies for value maximizing buyers in uniform price auctions. *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, volume 267 of *Proceedings of Machine Learning Research*, 19841–19877 (PMLR), URL <https://proceedings.mlr.press/v267/golrezaei25a.html>.
- Golrezaei N, Sahoo S (2026) Safe strategies in multi-unit auctions: Equilibria, liquid welfare, and revenue. SSRN working paper 6971218, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6971218.
- Google (2026) About Ad Rank. Google Ads Help, URL <https://support.google.com/google-ads/answer/1722122>.
- Hortaçsu A, McAdams D (2010) Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the Turkish treasury auction market. *Journal of Political Economy* 118(5):833–865, URL <http://dx.doi.org/10.1086/657948>.
- Kempe D, Mahdian M (2008) A cascade model for externalities in sponsored search. *Internet and Network Economics (WINE 2008)*, volume 5385 of *Lecture Notes in Computer Science*, 585–596 (Springer).
- Kolumbus Y, Nisan N (2022a) Auctions between regret-minimizing agents. *Proceedings of the ACM Web Conference 2022 (WWW)*, 100–111, URL <http://dx.doi.org/10.1145/3485447.3512055>.
- Kolumbus Y, Nisan N (2022b) How and why to manipulate your own agent: On the incentives of users of learning agents. *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, 28080–28094, URL <http://dx.doi.org/10.52202/068431-2036>.
- Li J, Tang P (2024) Vulnerabilities of single-round incentive compatibility in auto-bidding: Theory and evidence from ROI-constrained online advertising markets. *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, 2886–2894, URL <http://dx.doi.org/10.24963/ijcai.2024/320>.
- Liaw C, Mehta A, Perloth A (2023) Efficiency of non-truthful auctions in auto-bidding: The power of randomization. *Proceedings of the ACM Web Conference 2023 (WWW)*, 3561–3571, URL <http://dx.doi.org/10.1145/3543507.3583492>.
- Liaw C, Mehta A, Zhu W (2024) Efficiency of non-truthful auctions in auto-bidding with budget constraints. *Proceedings of the ACM Web Conference 2024 (WWW)*, 223–234, URL <http://dx.doi.org/10.1145/3589334.3645636>.
- Liu X, Shen W (2023) Auto-bidding with budget and ROI constrained buyers. *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, 2817–2825, URL <http://dx.doi.org/10.24963/ijcai.2023/314>.
- Lucier B, Pattathil S, Slivkins A, Zhang M (2024) Autobidders with budget and ROI constraints: Efficiency, regret, and pacing dynamics. *Proceedings of the Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, 3642–3643.
- Mannor S, Tsitsiklis JN, Yu JY (2009) Online learning with sample path constraints. *Journal of Machine Learning Research* 10:569–590.
- Markakis E, Sgouritsa A, Tsikiris A (2022) Towards a characterization of worst case equilibria in the discriminatory price auction. *Web and Internet Economics (WINE 2021)*, volume 13112 of *Lecture Notes in Computer Science*, 186–204 (Springer), URL http://dx.doi.org/10.1007/978-3-030-94676-0_11.
- Markakis E, Telelis O (2015) Uniform price auctions: Equilibria and efficiency. *Theory of Computing Systems* 57(3):549–575, URL <http://dx.doi.org/10.1007/s00224-014-9537-9>.
- Paes Leme R, Tardos É (2010) Pure and Bayes–Nash price of anarchy for generalized second price auction. *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 735–744 (IEEE), URL <http://dx.doi.org/10.1109/FOCS.2010.75>.
- Roughgarden T, Syrgkanis V, Tardos É (2017) The price of anarchy in auctions. *Journal of Artificial Intelligence Research* 59:59–101, URL <http://dx.doi.org/10.1613/JAIR.5272>.
- Sinha A, Vaze R (2024) Optimal algorithms for online convex optimization with adversarial constraints. *Advances in Neural Information Processing Systems*, volume 37, URL <http://dx.doi.org/10.52202/079017-1306>.

- Susan F, Golrezaei N, Schrijvers O (2023) Multi-platform budget management in ad markets with non-ic auctions. URL <https://arxiv.org/abs/2306.07352>.
- Varian HR (2007) Position auctions. *International Journal of Industrial Organization* 25(6):1163–1178, URL <http://dx.doi.org/10.1016/j.ijindorg.2006.10.002>.
- Vlatakis-Gkaragkounis EV, Flokas L, Lianas T, Mertikopoulos P, Piliouras G (2020) No-regret learning and mixed Nash equilibria: They do not mix. *Advances in Neural Information Processing Systems*, volume 33, 1380–1391.
- Wayfair Inc (2025) Annual report on Form 10-K for the year ended December 31, 2024. U.S. Securities and Exchange Commission, URL <https://www.sec.gov/Archives/edgar/data/1616707/000161670725000022/w-20241231.htm>.
- Wilson R (1979) Auctions of shares. *The Quarterly Journal of Economics* 93(4):675–689, URL <http://dx.doi.org/10.2307/1884475>.
- Wong M (2021) Moving AdSense to a first-price auction. Google AdSense, URL <https://blog.google/products/adsense/our-move-to-a-first-price-auction/>.
- Zinkevich M (2003) Online convex programming and generalized infinitesimal gradient ascent. *Proceedings of the 20th International Conference on Machine Learning*, 928–936.

Electronic Companion

Robust Multi-Item Demand Autobidding Auctions

EC.1–EC.5 contain the notation guide, complete result tables, technical lemmas, proofs, and detailed experimental results. The proofs follow the order of the main text. References are listed in the main manuscript.

EC.1. Notation and Complete Results

This section collects our notation and summarizes the welfare bounds proved in the paper.

EC.1.1. Guide to Notation

The table below lists our notation and terms for the model (Section 3), the results and extensions (Sections 4 and 5), the proofs (EC.2–EC.4), and repeated auctions (Section 5.3.2). The last column gives the place where each is defined.

Notation or term	Meaning	Definition
<i>Model</i>		
N, M	Numbers of advertisers and slots.	Section 3.1
$(n, j), v_{n,j}$	Ad j owned by advertiser n and its valuation.	Equation (1)
c_m	CTR of slot m .	Section 3.1
$v_{(m)}$	The m th largest valuation among all ads, counting repeated valuations separately.	Section 3.2
$\alpha_n, \alpha, \alpha_{-n}$	Advertiser n 's multiplier, the multiplier profile, and the opponents' multipliers.	Sections 3.2 and 3.5
$b_{n,j} = \alpha_n v_{n,j}$	Bid corresponding to ad (n, j) .	Section 3.2
$r_{n,j}(\alpha)$	Rank of ad (n, j) , written $r_{n,j}$ when α is clear from context.	Section 3.2
$\hat{b}_{n,i}(\alpha)$	The i th largest bid among the ads of advertisers other than n , written $\hat{b}_{n,i}$ when α is clear from context.	Section 3.4
Allocation, optimal allocation	The assignment of ads to slots by rank, and the allocation when every advertiser uses multiplier one.	Section 3.2
V_n, P_n, u_n	Acquired value, payment, and QL utility.	Sections 3.2 to 3.4
Wel, OPT	Welfare $\sum_n V_n$ and optimal welfare.	Section 3.2
QL, VM	Advertiser objectives.	Definitions 1 and 2
$\mathcal{Q}, \mathcal{V}$	QL and VM advertiser sets.	Section 3.3
Expected ROI, feasible	A VM advertiser's expected payment is at most her expected value, and a distribution of the profile is feasible for her if it satisfies expected ROI.	Definition 2
ROI slack	Strictly positive difference $V_n - P_n$ between a VM advertiser's value and payment.	Section 3.5
ROI deficit	Strictly positive difference $P_n - V_n$ between a VM advertiser's payment and value.	Section 3.5
σ	Distribution of the multiplier profile.	Section 3.5
$\tilde{\alpha}_n, \tilde{\sigma}_n$	Deviating multiplier and randomized strategy.	Section 3.5
CCE, CE, PNE	Equilibrium notions with randomized VM deviations.	Definitions 3 and EC.1
ψ_n	Generalized swap, which maps each recommendation to a replacement distribution.	Definition EC.1
I	An instance.	Section 3.5
$\text{PoA}_{\mathcal{E}}(\mathcal{M})$	Worst equilibrium welfare divided by optimal welfare, for equilibrium notion $\mathcal{E}$ and mechanism $\mathcal{M}$.	Definitions 4 and EC.1

Notation or term	Meaning	Definition
<i>Results and extensions</i>		
V_n^* , Rev	Advertiser n 's value in the optimal allocation, and revenue $\sum_n P_n$.	Section 4.1
Population sums	In the CCE analysis, U_n , V_n and P_n are expected utility, value and payment. A subscript QL or VM sums a quantity over that population, and omitting the advertiser subscript sums it over all advertisers.	Section 4.1.1 and EC.2.4
$B_n(\alpha)$, $B(\alpha)$	Weighted bid $\alpha_n V_n(\alpha)$, the sum of advertiser n 's bids weighted by the CTRs of their slots, and the total weighted bid $\sum_n B_n(\alpha)$, which equals GFP revenue.	Section 4.1.1
$\bar{\alpha}_n$, recoverable value	Target, the largest multiplier to which advertiser n deviates in the welfare bounds, and her value under these deviations, integrated over multipliers from 0 to her target.	Section 4.1.1
$\bar{u}$, ϕ	QL advertisers' expected utility relative to their optimal value, and $\phi(\bar{u}) = 1 - \bar{u} + \bar{u} \ln \bar{u}$, extended by 0 for $\bar{u} \geq 1$.	Section 4.1.1 and Lemma EC.2
Strict PNE	A PNE in which every advertiser's multiplier is her unique best response.	Section 4.2.1
PNE with fixed deviations	A profile at which no QL advertiser gains from a fixed multiplier, and every VM advertiser satisfies ROI and gains no value from a fixed multiplier that satisfies ROI.	Section 4.2.3
Unit-demand	An instance with $v_{n,j} = 0$ for every advertiser n and every $j \geq 2$.	Section 5.1
Homogeneous objectives	Advertisers are all-QL or all-VM.	Section 1.3
$\text{PoA}_{\mathcal{E}}(\mathcal{M}, \text{All-VM})$, $\text{PoA}_{\mathcal{E}}(\mathcal{M}, \text{All-QL})$	PoA over the instances in which advertisers are all-VM, or all-QL. We omit $\mathcal{E}$ for a common value or order.	EC.1.2
$A_{\max}$, $A_{\min}$, equilibrium support	Ceiling and floor of equilibrium multipliers, and the interval $[A_{\min}, A_{\max}]$ between them.	Section 5.2.1
Undominated multipliers	QL multipliers of at most one under GFP and GSP and of one under VCG, and VM multipliers of at least one under VCG.	Section 5.2.2
$\mathcal{A}_n$	Interval of advertiser n 's equilibrium multipliers, which in repeated play contains her multiplier grid.	Table EC.4
ν_n , ν , π_n	Advertiser n 's valuation vector, the random valuation profile, and a strategy, which maps advertiser n 's valuation vector to a multiplier.	Section 5.3.1
Bayesian CCE, ex-ante ROI	Equilibrium notion with stochastic valuations, and the ROI constraint averaged over valuations.	Definition 5
<i>Proofs</i>		
$\mathcal{X}$, $\mathbf{x}(\alpha)$, $V_n(\mathbf{x})$	Allocations, the auction's allocation at profile α , and advertiser n 's value under allocation $\mathbf{x}$.	EC.2.3
$P_n^{\mathcal{M}}$, $\text{Rev}^{\mathcal{M}}$	Payment and revenue under mechanism $\mathcal{M}$, which we omit when it is clear.	EC.2.3
Weighted bid under $\mathbf{x}$	Advertiser n 's weighted bid $\alpha_n V_n(\mathbf{x})$ under allocation $\mathbf{x}$, and its sum over advertisers.	EC.2.3
$\mathcal{X}_{-n}$, B_{-n}	Allocations that assign no ad of advertiser n , and the other advertisers' total weighted bid.	EC.2.3
V^{both} , V^{inef} , V^{omit}	Value from common CTR, additional CTR, and missing CTR relative to the optimal allocation, the last of which is the omitted value.	Equation (EC.1)
B^{inef}	Weighted bid on additional CTR, $\sum_n \alpha_n V_n^{\text{inef}}$.	Equation (EC.2)
$[x]_+$	Positive part, $\max\{x, 0\}$.	EC.2.4
<i>Repeated auctions</i>		
T , t	Number of rounds and a round. In EC.4.10, $\mathbb{E}$ is expectation under the empirical distribution of the realized profiles.	EC.4.10.1

Notation or term	Meaning	Definition
$\mathbf{v}_{n,t}, \text{OPT}_t$	Advertiser n 's valuation vector in round t , and that round's optimal welfare.	Section 5.3.2
$\mathcal{A}_n^L$	Advertiser n 's multiplier grid in repeated play, with at most L points.	EC.4.10.1
R_n, S_n, ν_n	QL regret, VM shortfall relative to bidding valuations, and VM net ROI deficit under learning dynamics.	Equation (10)
$g_{n,t}, \mu_{n,t}, \bar{\mu}$	Reward, penalty coefficient on the ROI deficit, and its cap in the primal–dual algorithm.	Equation (EC.15)
γ, γ_μ	Step sizes of the multiplier and penalty updates.	EC.4.10.3
K_n	Known bound on advertiser n 's value and payment in every round.	EC.4.10.3

EC.1.2. Complete Results

For convenience, Tables EC.1 to EC.3 summarize the welfare bounds proved in this paper, first without a floor or ceiling and then with a floor $A_{\min}$ and a ceiling $A_{\max}$ (Section 5.2.1). The tables and the precise statements in EC.4 use the following notation. For an equilibrium notion $\mathcal{E}$ and a mechanism $\mathcal{M}$, we write $\text{PoA}_{\mathcal{E}}(\mathcal{M}, \text{All-VM})$ and $\text{PoA}_{\mathcal{E}}(\mathcal{M}, \text{All-QL})$ for the PoA of Definition 4 with the infimum restricted to the instances in which advertisers are all-VM or all-QL, respectively, while $\text{PoA}_{\mathcal{E}}(\mathcal{M})$, over all instances, allows any mix of the two objectives. With a floor and ceiling, these are taken only over the equilibria whose multipliers lie in $[A_{\min}, A_{\max}]$ (Section 5.2.1), and with undominated multipliers, only over the equilibria in which advertisers use them (Section 5.2.2). The equilibrium notion $\mathcal{E}$ is CCE, CE (Definition EC.1), or PNE. As in Section 3.5, we omit $\mathcal{E}$ when a CCE lower bound and a PNE upper bound match. Entries with O , Θ , or Ω are asymptotic, as $A_{\max}^{-1}, A_{\min} \rightarrow 0$, and 0.4569 is rounded. The cited results give the exact bounds.

Table EC.1 PoA bounds without a floor or ceiling.

Mechanism	Mixed	All-VM	All-QL	Reference
GFP	$\text{PoA}_{\text{CCE}} = \text{PoA}_{\text{CE}} \approx 0.4569$; $\text{PoA}_{\text{PNE}} = 1 - 1/e$	$\text{PoA}_{\text{CCE}} = \text{PoA}_{\text{CE}} = 1/2$	$\text{PoA}_{\text{CCE}} = 1 - 1/e$	Thm. 1; Props. 4, EC.1; Def. EC.1
GFP, unit-demand	$\text{PoA}_{\text{PNE}} = 4/5$	$\text{PoA}_{\text{PNE}} \geq 4/5$	$\text{PoA}_{\text{PNE}} \geq 4/5$	Prop. 3
GSP	$\text{PoA} = 0$, also for strict PNE	$\text{PoA} = 0$	$\text{PoA} = 0$	Thm. 2; Prop. 1
VCG	$\text{PoA} = 0$	$\text{PoA} = 0$	$\text{PoA} = 0$	Thm. 2
VCG, undominated multipliers	$\text{PoA} = 1/2$	$\text{PoA} = 1/2$	$\text{PoA} = 1$	EC.4.7

With undominated multipliers, the entries for GFP, and for GSP with mixed or all-VM populations, are unchanged (Proposition 6), except that for all-QL advertisers, GFP is only known to satisfy $\text{PoA}_{\text{CCE}} \geq 1 - 1/e$ (Proposition EC.2).

Table EC.2 Complete PoA bounds with a floor and ceiling for mixed objectives.

Population	Mechanism	All multipliers	Undominated multipliers	Reference
Mixed	GFP	$\text{PoA}_{\text{ACCE}}, \text{PoA}_{\text{CE}} = 0.4569 + O(A_{\max}^{-1} + A_{\min});$ $\text{PoA}_{\text{PNE}} = 1 - 1/e + A_{\min}/e$	Same bounds	Props. EC.1, EC.2
Mixed	GSP	$\text{PoA} = \Theta(A_{\max}^{-1} + A_{\min})$	Same bound	Prop. EC.3
Mixed	VCG	$\text{PoA} = \Theta(A_{\max}^{-1} + A_{\min})$	$\text{PoA} = 1/2 + O(A_{\max}^{-1})$	Props. EC.3, EC.5

We write PoA without a subscript when the bound holds for every equilibrium class from PNE to CCE. A common asymptotic order does not assert equality of the exact PoAs.

Table EC.3 Complete PoA bounds with a floor and ceiling for homogeneous objectives.

Population	Mechanism	All multipliers	Undominated multipliers	Reference
All-VM	GFP	$\text{PoA}_{\text{ACCE}}, \text{PoA}_{\text{CE}} = 1/2 + O(A_{\max}^{-1} + A_{\min})$	Same bounds	Props. 4, EC.2
All-VM	GSP	$\text{PoA} = \Theta(A_{\max}^{-1} + A_{\min})$	Same bound	Prop. EC.3
All-VM	VCG	$\text{PoA} = \Theta(A_{\max}^{-1} + A_{\min})$	$\text{PoA} = 1/2 + O(A_{\max}^{-1})$	Props. EC.3, EC.5
All-QL	GFP	$\text{PoA}_{\text{ACCE}} = 1 - 1/e + O(A_{\max}^{-1} + A_{\min})$	$\text{PoA}_{\text{ACCE}} \geq 1 - 1/e$	Props. 4, EC.2
All-QL	GSP	$\text{PoA} = \Theta(A_{\max}^{-1} + A_{\min})$	$\text{PoA}_{\text{ACCE}} = \Omega(A_{\max}^{-1} + A_{\min})$	Prop. EC.3; Lem. EC.4
All-QL	VCG	$\text{PoA} = \Theta(A_{\max}^{-1} + A_{\min})$	$\text{PoA} = 1$	Props. EC.3, EC.5

We write PoA without a subscript when the bound holds for every equilibrium class from PNE to CCE. A common asymptotic order does not assert equality of the exact PoAs.

EC.2. Preliminaries and Technical Lemmas

We first prove the GSP and VCG analogues of Example 2 and define the correlated equilibrium. We then define the allocations and weighted bids used in the proofs and prove the allocation inequalities used in the welfare proofs and the resulting nonlinear GFP bound on utility and revenue.

EC.2.1. GSP and VCG Analogues of Example 2

In addition to GFP as discussed in Example 2, randomization can improve on every feasible fixed choice of α_n under both GSP and VCG. We note, however, that under VCG the low multiplier of the randomized strategy need not lie below one. With one slot, GSP and VCG coincide, so one construction proves both analogues.

Verification. Take two VM advertisers whose single ads have valuation one. Advertiser 1 uses $\alpha_1 = 1$. Advertiser 2 uses $1/2$ or $3/2 + \delta$ with equal probability, where $0 < \delta < 1/2$, and has priority at equality. Her choice $1/2$ loses, while $3/2 + \delta$ wins and pays advertiser 1's bid of one. Thus her strategy is feasible, with expected value and payment both equal to $1/2$. Against this strategy, advertiser 1 loses in both states whenever she fixes $\alpha_1 \leq 1/2$. A choice $\alpha_1 \in (1/2, 3/2 + \delta]$ wins only against $\alpha_2 = 1/2$ and gives expected value–payment pair $(1/2, 1/4)$. A choice $\alpha_1 > 3/2 + \delta$ wins in both states and gives pair $(1, 1 + \delta/2)$, which violates ROI. Hence every feasible fixed α_1 obtains value at most $1/2$.

Fixing $0 < \varepsilon < 1/2 - \delta$, advertiser 1 can use $1/2 + \varepsilon$ with probability $2\delta/(1 + 2\delta)$ and $3/2 + \delta + \varepsilon$ otherwise. The former lies below one and gives pair $(1/2, 1/4)$, and the latter lies above one and gives pair $(1, 1 + \delta/2)$. Therefore

$$\mathbb{E}[V_1] = \mathbb{E}[P_1] = \frac{1 + \delta}{1 + 2\delta} \rightarrow 1 \quad \text{as } \delta \rightarrow 0.$$

The choice $1/2 + \varepsilon$ creates exactly the ROI slack needed to support $3/2 + \delta + \varepsilon$. Thus, as in Example 2, restricting α_1 to a fixed value would miss the profitable randomized deviation. As a clarification, in this single-slot example, replacing $1/2 + \varepsilon$ by one gives the same allocation and payment under both GSP and VCG. Thus the example establishes the benefit of randomization, not a need for multipliers below one. For VCG, Lemma 1 justifies this replacement more generally, including when an advertiser may receive several slots. $\square$

EC.2.2. Correlated Equilibrium

A further equilibrium notion, the correlated equilibrium, appears in some of our results and experiments. We use the notation of Section 3.5.

DEFINITION EC.1 (CORRELATED EQUILIBRIUM). Let σ be a distribution over profiles that is feasible for every VM advertiser. Draw $\alpha \sim \sigma$, and view each multiplier α_n as a *recommendation* to advertiser n . A *generalized swap* ψ_n for advertiser n maps each recommendation α_n to a distribution $\psi_n(\alpha_n)$ over $[0, \infty)$. Following it, advertiser n uses $\tilde{\alpha}_n \sim \psi_n(\alpha_n)$ instead of α_n . The swap is *feasible* for a VM advertiser n if

$$\mathbb{E}[P_n(\tilde{\alpha}_n, \alpha_{-n})] \leq \mathbb{E}[V_n(\tilde{\alpha}_n, \alpha_{-n})].$$

Then σ is a *correlated equilibrium* (CE) if every QL advertiser $n \in Q$ satisfies

$$\mathbb{E}_{\alpha \sim \sigma}[u_n(\alpha)] \geq \mathbb{E}[u_n(\tilde{\alpha}_n, \alpha_{-n})] \quad \text{for every generalized swap } \psi_n,$$

and every VM advertiser $n \in \mathcal{V}$ satisfies

$$\mathbb{E}_{\alpha \sim \sigma}[V_n(\alpha)] \geq \mathbb{E}[V_n(\tilde{\alpha}_n, \alpha_{-n})] \quad \text{for every feasible generalized swap } \psi_n.$$

In each expectation without a subscript, $\alpha \sim \sigma$ is drawn first and then $\tilde{\alpha}_n \sim \psi_n(\alpha_n)$. We write $\text{PoA}_{\text{CE}}(\mathcal{M})$ for the price of anarchy of Definition 4 with this equilibrium notion.

A generalized swap that ignores the recommendation is a randomized strategy, so every CE is a CCE. Every PNE is a CE (Roughgarden et al. 2017). In games without constraints, the empirical distribution of play approaches the set of CCE when every player's external-regret vanishes, and the set of CE when every player's swap-regret vanishes (Blum and Mansour 2007, Roughgarden et al. 2017). Our swap-regret experiments (Section 6.6) use the second kind of learning.

EC.2.3. Allocations and Weighted Bids

The proofs compare the auction's allocation with other ways of assigning ads to slots. An *allocation* $\mathbf{x}$ assigns ads to slots so that each slot receives at most one ad and each ad at most one slot. Let $\mathcal{X}$ be the set of allocations, and let $V_n(\mathbf{x})$ be advertiser n 's value under $\mathbf{x}$, the sum of her ads' valuations weighted by the CTRs of their slots. Write $\mathbf{x}(\alpha)$ for the auction's allocation at profile α , which assigns slot m to the ad of rank m for $m = 1, \dots, M$ (Section 3.2), so that $V_n(\alpha) = V_n(\mathbf{x}(\alpha))$. We write $P_n^{\mathcal{M}}(\alpha)$ and $\text{Rev}^{\mathcal{M}}(\alpha)$ for advertiser n 's payment and the revenue (Section 4.1) under mechanism $\mathcal{M}$, and we omit $\mathcal{M}$ when the mechanism is clear from context. The weighted bids of Section 4.1.1 extend to every allocation, with advertiser n 's weighted bid under $\mathbf{x}$ being $\alpha_n V_n(\mathbf{x})$, so that $B_n(\alpha) = \alpha_n V_n(\mathbf{x}(\alpha))$. The auction's allocation assigns the highest bids to the slots in decreasing order of CTR, so it maximizes the total weighted bid over $\mathcal{X}$, a property that we call *allocation optimality*:

$$B(\alpha) = \sum_{n=1}^N B_n(\alpha) = \max_{\mathbf{x} \in \mathcal{X}} \sum_{n=1}^N \alpha_n V_n(\mathbf{x}).$$

VCG payments as externality costs. Fix a profile α and an advertiser n , and recall the competing bids $\hat{b}_{n,i}$ of Section 3.4. Without advertiser n 's ads, the other advertisers' ads fill the slots in decreasing order of bid. We show that the VCG payment in (3) is her externality cost, that is,

$$P_n^{\text{VCG}}(\alpha) = \sum_{i=1}^M c_i \hat{b}_{n,i} - \sum_{\substack{n'=1 \\ n' \neq n}}^N \sum_{j=1}^M c_{r_{n',j}} b_{n',j},$$

the CTR-weighted bids of the other advertisers' ads without her, minus their CTR-weighted bids when she is present. Number the other advertisers' ads $i = 1, 2, \dots$ in rank order, so that the i th of them has bid $\hat{b}_{n,i}$, and let k_i be the number of advertiser n 's ads ranked above it. The i th of them then has rank $i + k_i$, so the right side equals

$$\sum_{i \geq 1} \hat{b}_{n,i} (c_i - c_{i+k_i}) = \sum_{i \geq 1} \hat{b}_{n,i} \sum_{j=1}^{k_i} (c_{i+j-1} - c_{i+j}),$$

with $c_m = 0$ for $m > M$ and $\hat{b}_{n,i} = 0$ when the other advertisers have fewer than i ads. Advertiser n 's ads ranked above the i th of the other ads are $(n, 1), \dots, (n, k_i)$, and ad (n, j) has $r_{n,j} - j$ ads of other advertisers ranked above it. Hence $j \leq k_i$ exactly when $i \geq r_{n,j} - j + 1$. Substituting $m = i + j - 1$ and exchanging the order of summation gives

$$\sum_{j=1}^M \sum_{m \geq r_{n,j}} (c_m - c_{m+1}) \hat{b}_{n,m-j+1},$$

which is the VCG payment in (3), because the terms with $m > M$ vanish.

Similarly, let $\mathcal{X}_{-n} \subseteq \mathcal{X}$ be the allocations that assign no ad of advertiser n , and let $B_{-n}(\mathbf{x}) \doteq \sum_{n' \neq n} \alpha_{n'} V_{n'}(\mathbf{x})$ be the other advertisers' total weighted bid. Ranking the other advertisers' ads by bid, without advertiser n 's ads, maximizes B_{-n} over $\mathcal{X}_{-n}$, so the VCG payment equals

$$P_n^{\text{VCG}}(\alpha) = \max_{\mathbf{x} \in \mathcal{X}_{-n}} B_{-n}(\mathbf{x}) - B_{-n}(\mathbf{x}(\alpha)).$$

No advertiser pays more than her bids under GSP or VCG. More specifically, $P_n^{\text{GSP}}(\alpha) \leq B_n(\alpha)$, because each GSP price is a bid ranked below the ad's own bid, and $P_n^{\text{VCG}}(\alpha) \leq B_n(\alpha)$, because $X_{-n} \subseteq \mathcal{X}$ gives $\max_{x \in \mathcal{X}_{-n}} B_{-n}(x) \leq B(\alpha) = B_n(\alpha) + B_{-n}(x(\alpha))$. In particular, multiplier one satisfies ROI under every mechanism, as payment then is at most acquired value.

EC.2.4. Allocation Values and Deviations

Throughout the proofs, α denotes the profile before any deviation, possibly drawn from σ . A unilateral deviation replaces α_n by α' with α_{-n} fixed, and a superscript $\star$ refers to the welfare-maximizing allocation. We first define the positions used in the welfare identity (EC.1) below. As in Section 3, ad (n, j) has valuation $v_{n,j}$, bid $b_{n,j} = \alpha_n v_{n,j}$, and rank $r_{n,j}(\alpha)$ at profile α . Let $\mathcal{I}_n \doteq \{(n, 1), \dots, (n, M)\}$ be the ads owned by advertiser n , and let $r_{n,j}^\star$ be the rank of ad (n, j) when ads are ranked by valuation, i.e., at profile $(1, \dots, 1)$. Ads ranked below M receive no slot, and $c_m = 0$ for $m > M$. Advertiser n 's acquired value and her value $V_n^\star$ in the optimal allocation (Section 4.1) are therefore

$$V_n(\alpha) = \sum_{j=1}^M c_{r_{n,j}} v_{n,j}, \quad V_n^\star = \sum_{j=1}^M c_{r_{n,j}^\star} v_{n,j}.$$

To distinguish the CTR common to the two allocations from the CTR that an ad loses or gains relative to the optimum:

$$\begin{aligned} V_n^{\text{both}}(\alpha) &\doteq \sum_{j=1}^M v_{n,j} \min\{c_{r_{n,j}}, c_{r_{n,j}^\star}\}, \\ V_n^{\text{omit}}(\alpha) &\doteq \sum_{j=1}^M v_{n,j} [c_{r_{n,j}^\star} - c_{r_{n,j}}]_+, \\ V_n^{\text{inef}}(\alpha) &\doteq \sum_{j=1}^M v_{n,j} [c_{r_{n,j}} - c_{r_{n,j}^\star}]_+, \end{aligned}$$

where $[x]_+ = \max\{x, 0\}$. Thus V_n^{omit} is advertiser n 's part of the *omitted value*, the value that the optimal allocation obtains because some ads receive more CTR there than at α , and V_n^{inef} is the value obtained from assigning other ads more CTR than in the optimum. For clarity, these are portions of value, rather than sets of ads. As in Section 4.1.1, a subscript QL or VM sums a quantity over that population, and omitting the advertiser subscript sums it over all advertisers. We suppress the profile argument when it is clear. These definitions, along with the fact that $\text{Wel}(1, \dots, 1) = \text{OPT}$, give the welfare identities

$$\text{Wel} = V^{\text{both}} + V^{\text{inef}}, \quad \text{OPT} = V^{\text{both}} + V^{\text{omit}}. \quad (\text{EC.1})$$

For completeness, we also write $B^{\text{inef}} \doteq \sum_n \alpha_n V_n^{\text{inef}}$ for the weighted bid on the additional CTR. Allocation optimality implies

$$\sum_n \alpha_n V_n^{\text{omit}} \leq B^{\text{inef}} = \sum_n \alpha_n V_n^{\text{inef}} \leq B(\alpha). \quad (\text{EC.2})$$

The first two parts of the next lemma bound optimal or omitted value using the current total weighted bid and acquired values under deviations, and the third bounds the omitted weighted bid by payments. The targets $\bar{\alpha}_n$ of Section 4.1.1, the upper limits of the integrals over deviations, specify the range of α' in each part.

- **Optimal value and fixed deviations.** In part (i), acquired value is the rate at which the total weighted bid increases with α_n . Increasing each α_n in turn gives a bound on $\sum_n \bar{\alpha}_n V_n^\star$ using deviations with the original α_{-n} . The current total weighted bid, which equals GFP revenue, appears only once.
- **Omitted value and raises.** In part (ii), we bound omitted optimal value by B^{inef} and the additional value acquired under α' -raises, $\alpha_n \mapsto \max\{\alpha_n, \alpha'\}$. We use this finer inequality in the GFP PNE proof.
- **Omitted bids and second-price payments.** In part (iii), the weighted bid $\sum_n \alpha_n V_n^{\text{omit}}$ is bounded by revenue under GSP or VCG.

LEMMA EC.1. Fix α and any targets $\bar{\alpha}_1, \dots, \bar{\alpha}_N \geq 0$.

(i) The optimal allocation satisfies the recovery inequality (Section 4.1.1)

$$\sum_n \bar{\alpha}_n V_n^\star \leq B(\alpha) + \sum_n \int_0^{\bar{\alpha}_n} V_n(\alpha', \alpha_{-n}) d\alpha'. \quad (\text{EC.3})$$

(ii) The omitted value satisfies

$$\sum_n \bar{\alpha}_n V_n^{\text{omit}}(\alpha) \leq B^{\text{inef}}(\alpha) + \sum_n \int_0^{\bar{\alpha}_n} [V_n(\max\{\alpha_n, \alpha'\}, \alpha_{-n}) - V_n(\alpha)] d\alpha'. \quad (\text{EC.4})$$

(iii) For $\mathcal{M} \in \{\text{GSP}, \text{VCG}\}$,

$$\sum_n \alpha_n V_n^{\text{omit}} \leq \text{Rev}^{\mathcal{M}}(\alpha). \quad (\text{EC.5})$$

Proof. Part (i): optimal value and fixed deviations. We first express changes in the total weighted bid through acquired value. By allocation optimality, $B(\alpha) = \max_{x \in \mathcal{X}} \sum_n \alpha_n V_n(x)$ at every profile (EC.2.3). With α_{-n} fixed, each of the finitely many allocations gives an affine function of α_n . Their maximum is continuous and piecewise linear. Wherever it is differentiable, its slope is the actual acquired value $V_n(\alpha', \alpha_{-n})$. Thus, for any α_{-n} ,

$$B(\max\{\alpha_n, \bar{\alpha}_n\}, \alpha_{-n}) - B(\alpha_n, \alpha_{-n}) = \int_{\alpha_n}^{\max\{\alpha_n, \bar{\alpha}_n\}} V_n(\alpha', \alpha_{-n}) d\alpha'. \quad (\text{EC.6})$$

Next, increasing the components of α_{-n} cannot improve the position of any owned ad, because bids determine the ranking with a fixed priority at ties. It therefore cannot increase acquired value at a fixed α_n . Starting from $\alpha^{(0)} = \alpha$, apply a $\bar{\alpha}_n$ -raise to each advertiser in turn by setting

$$\alpha^{(n)} = (\max\{\alpha_1, \bar{\alpha}_1\}, \dots, \max\{\alpha_n, \bar{\alpha}_n\}, \alpha_{n+1}, \dots, \alpha_N), \quad n = 1, \dots, N.$$

These profiles only compare allocation objectives. They are *not* coordinated equilibrium deviations and need not satisfy ROI. At step n , the components of $\alpha_{-n}^{(n-1)}$ are at least their original values, and combining this with the integral identity (Equation EC.6) guarantees that $B(\alpha^{(N)}) - B(\alpha)$ is

$$= \sum_n \int_{\alpha_n}^{\max\{\alpha_n, \bar{\alpha}_n\}} V_n(\alpha', \alpha_{-n}^{(n-1)}) d\alpha' \leq \sum_n \int_{\alpha_n}^{\max\{\alpha_n, \bar{\alpha}_n\}} V_n(\alpha', \alpha_{-n}) d\alpha' \leq \sum_n \int_0^{\bar{\alpha}_n} V_n(\alpha', \alpha_{-n}) d\alpha'.$$

For the last inequality, the increment is 0 if $\bar{\alpha}_n \leq \alpha_n$. Otherwise, its integration interval lies within $[0, \bar{\alpha}_n]$, and acquired value is nonnegative. Finally, the profile $\alpha^{(N)}$ submits at least $\bar{\alpha}_n$ times every valuation of advertiser n . Evaluating its allocation objective at the welfare-optimal allocation gives $B(\alpha^{(N)}) \geq \sum_n \bar{\alpha}_n V_n^*$. Combining the preceding integral bound with $B(\alpha^{(N)}) \geq \sum_n \bar{\alpha}_n V_n^*$ gives the recovery inequality (EC.3), with each $V_n(\alpha', \alpha_{-n})$ evaluated at the original α_{-n} .

Part (ii): omitted value and raises. Unlike part (i), we bound omitted value using only $V_n(\max\{\alpha_n, \alpha'\}, \alpha_{-n}) - V_n(\alpha)$, so the deviation retains the current $V_n(\alpha)$. We identify the missing owned ads within the first m positions, bound their recovery using selected nonoptimal bids, and then combine the position qualities with weights $c_m - c_{m+1}$. Fix $m \leq M$ and write

$$\mathcal{S}_m = \{(n', j) : r_{n', j} \leq m\}, \quad \mathcal{O}_m = \{(n', j) : r_{n', j}^* \leq m\}, \quad s_n = |\mathcal{S}_m \cap \mathcal{I}_n|, \quad o_n = |\mathcal{O}_m \cap \mathcal{I}_n|.$$

Here $\mathcal{S}_m$ and $\mathcal{O}_m$ contain the ads in the first m positions of the current and welfare-optimal allocations, respectively, and s_n and o_n count advertiser n 's ads in these sets. All counts below refer to this m . List advertiser n 's ads $(n, 1), \dots, (n, M)$ in index order, that is, in decreasing valuation. Both selections are therefore prefixes of this ordered list:

$$\mathcal{S}_m \cap \mathcal{I}_n = \{(n, 1), \dots, (n, s_n)\}, \quad \mathcal{O}_m \cap \mathcal{I}_n = \{(n, 1), \dots, (n, o_n)\}.$$

Within this prefix, an advertiser missing an optimal ad has no nonoptimal ad, and her omitted ads follow her selected ads:

$$o_n > s_n \implies (\mathcal{S}_m \setminus \mathcal{O}_m) \cap \mathcal{I}_n = \emptyset, \quad (\mathcal{O}_m \setminus \mathcal{S}_m) \cap \mathcal{I}_n = \{(n, s_n + 1), \dots, (n, o_n)\}.$$

Let $q_n = [o_n - s_n]_+$ count advertiser n 's ads that appear among the first m positions in the optimal allocation but not in the current allocation, and let $q = \sum_n q_n$ be their total number. Since both allocations assign m ads to these positions, $q = |\mathcal{O}_m \setminus \mathcal{S}_m| = |\mathcal{S}_m \setminus \mathcal{O}_m|$ is also the number of selected nonoptimal ads. If $q = 0$, there is no omitted value to bound at these positions. Otherwise, list the bids in $\mathcal{S}_m \setminus \mathcal{O}_m$ increasingly as $b'_1 \leq \dots \leq b'_q$. For advertiser n 's j th omitted ad $(n, s_n + j)$, use b'_j as a *comparison bid*—a selected nonoptimal bid used to give a sufficient condition for recovering that ad. We note that these bids need not identify the smallest bids that do so.

Now, suppose $\alpha' v_{n, s_n + j} > b'_j$. Since every selected bid is at least every unselected bid, this implies $\alpha' > \alpha_n$. To count the opposing ads that could remain ahead of the first $s_n + j$ owned ads, we may exclude all originally

unselected ads, advertiser n 's s_n selected ads, and the j nonoptimal ads with bids at most b'_j . These last two groups are disjoint because an advertiser missing an optimal ad has no selected nonoptimal ad among the first m positions. Thus

$$|\{(n', j') \notin \mathcal{I}_n : b_{n', j'} > b'_j\}| = |\{(n', j') \in \mathcal{S}_m \setminus \mathcal{I}_n : b_{n', j'} > b'_j\}| \leq m - s_n - j.$$

Each of the first $s_n + j$ owned ads now has bid strictly above b'_j . At most $m - s_n - j$ opposing ads can precede them, so all $s_n + j$ are selected among the first m positions. The conditions $\alpha' v_{n, s_n + j} > b'_j$ form an initial segment in j , because valuations decrease and comparison bids increase. The α' -raise also retains the previously selected ads. Thus

$$\sum_{\substack{i=1, \dots, M: \\ r_{n,i}(\max\{\alpha_n, \alpha'\}, \alpha_{-n}) \leq m}} v_{n,i} - \sum_{i=1}^{s_n} v_{n,i} \geq \sum_{j=1}^{q_n} v_{n, s_n + j} \mathbf{1}\{\alpha' v_{n, s_n + j} > b'_j\}.$$

For $\alpha' \leq \alpha_n$, the right side is 0 and the same inequality holds. Integrating recovers the portion of each omitted ad's valuation beyond its comparison bid. For $v \geq 0$,

$$\bar{\alpha}_n v \leq b'_j + \int_0^{\bar{\alpha}_n} v \mathbf{1}\{\alpha' v > b'_j\} d\alpha' = b'_j + [\bar{\alpha}_n v - b'_j]_+.$$

The sum of the assigned comparison bids is at most the sum of the bids in $\mathcal{S}_m \setminus \mathcal{O}_m$. Indeed, split the increasing list $b'_1 \leq \dots \leq b'_q$ into consecutive blocks, one for each advertiser n , with q_n bids in her block. Her q_n comparison bids are the smallest q_n bids in the list, so they are bounded term by term by the bids in her block, and the blocks exhaust the list because $\sum_n q_n = q$. Hence

$$\sum_n \sum_{j=1}^{q_n} b'_j \leq \sum_{j=1}^q b'_j = \sum_{(n', j') \in \mathcal{S}_m \setminus \mathcal{O}_m} b_{n', j'}.$$

Combining the lower bound on acquired value in the first m positions, the integral for each omitted ad, and the upper bound on the sum of assigned bids yields

$$\sum_n \bar{\alpha}_n \sum_{j=s_n+1}^{o_n} v_{n,j} \leq \sum_{(n', j') \in \mathcal{S}_m \setminus \mathcal{O}_m} b_{n', j'} + \sum_n \int_0^{\bar{\alpha}_n} \left[\sum_{\substack{i=1, \dots, M: \\ r_{n,i}(\max\{\alpha_n, \alpha'\}, \alpha_{-n}) \leq m}} v_{n,i} - \sum_{i=1}^{s_n} v_{n,i} \right] d\alpha'.$$

To pass to the complete allocation, multiply by $c_m - c_{m+1}$ and sum over m . For each ad (n, i) , telescoping gives

$$\sum_{m=1}^M (c_m - c_{m+1}) \mathbf{1}\{r_{n,i}^* \leq m < r_{n,i}\} = [c_{r_{n,i}^*} - c_{r_{n,i}}]_+.$$

The reverse position inequality similarly gives $[c_{r_{n,i}} - c_{r_{n,i}^*}]_+$, while $\sum_{m=1}^M (c_m - c_{m+1}) \mathbf{1}\{r \leq m\} = c_r$ for every rank r . The omitted-value sum, the sum of selected nonoptimal bids, and the integrated value gain therefore become exactly $\sum_n \bar{\alpha}_n V_n^{\text{omit}}$, B_n^{inef} , and the integrated additional value in Equation (EC.4).

Part (iii): weighted bids and second-price payments. We use the sets $\mathcal{S}_m$ and $\mathcal{O}_m$ from part (ii) to bound omitted bids by opposing selected ads' prices under GSP and, by adding feasible replacements, by payments under VCG. Fix any bijection $h_m : \mathcal{O}_m \setminus \mathcal{S}_m \rightarrow \mathcal{S}_m \setminus \mathcal{O}_m$. Since an advertiser with an omitted ad has no selected nonoptimal ad among the first m positions, the two ads in every pair have different owners. Each omitted ad (n, j) also ranks below its paired ad $h_m(n, j)$, so its bid can determine the GSP price of the paired ad. Write $p(n', j')$ for the GSP price per click of ad (n', j') . Then $p(h_m(n, j)) \geq b_{n,j}$ for every $(n, j) \in \mathcal{O}_m \setminus \mathcal{S}_m$, so

$$\sum_{(n', j') \in \mathcal{S}_m} p(n', j') \geq \sum_{(n, j) \in \mathcal{O}_m \setminus \mathcal{S}_m} b_{n,j}.$$

Multiplying by $c_m - c_{m+1}$ and summing gives

$$\text{Rev}^{\text{GSP}}(\alpha) = \sum_{m=1}^M (c_m - c_{m+1}) \sum_{(n', j') \in \mathcal{S}_m} p(n', j') \geq \sum_{m=1}^M (c_m - c_{m+1}) \sum_{(n', j') \in \mathcal{O}_m \setminus \mathcal{S}_m} b_{n', j'} = \sum_n \alpha_n V_n^{\text{omit}}.$$

For VCG, let $\mathcal{S}_m^{-n}$ contain the largest m opposing bids, or all opposing ads if there are fewer than m . When advertiser n is removed, the other advertisers can retain the ads in $\mathcal{S}_m \setminus \mathcal{I}_n$ and add the omitted ads paired with her selected nonoptimal ads. This is a feasible set of at most m opposing ads. Consequently,

$$\sum_{(n', j') \in \mathcal{S}_m^{-n}} b_{n', j'} - \sum_{(n', j') \in \mathcal{S}_m \setminus \mathcal{I}_n} b_{n', j'} \geq \sum_{\substack{(n', j') \in \mathcal{O}_m \setminus \mathcal{S}_m: \\ h_m(n', j') \in \mathcal{I}_n}} b_{n', j'}.$$

The left side is advertiser n 's VCG payment for m identical positions, that is, the loss in other advertisers' total weighted bid caused by her presence. Summing over advertisers counts each paired omitted bid once. Her VCG payment for the original CTR vector is the weighted sum of these differences, because the same ranking by bid maximizes every prefix in the counterfactual allocation. Hence

$$\begin{aligned} \text{Rev}^{\text{VCG}}(\alpha) &= \sum_{m=1}^M (c_m - c_{m+1}) \sum_n \left(\sum_{(n', j') \in \mathcal{S}_m^{-n}} b_{n', j'} - \sum_{(n', j') \in \mathcal{S}_m \setminus \mathcal{I}_n} b_{n', j'} \right) \\ &\geq \sum_{m=1}^M (c_m - c_{m+1}) \sum_{(n', j') \in \mathcal{O}_m \setminus \mathcal{S}_m} b_{n', j'} = \sum_n \alpha_n V_n^{\text{omit}}. \end{aligned}$$

□

EC.2.5. GFP Utility and Revenue

Low QL utility restricts the value obtainable from fixed deviations. We combine this restriction with the feasible VM deviation to one in the joint allocation inequality from Lemma EC.1(i). The lemma evaluates ϕ at $\mathbb{E}[u_n]/V_n^*$, a QL advertiser's expected utility divided by her contribution to optimal welfare. Choosing a target for each QL advertiser gives the following two revenue bounds. In the proof of Theorem 1, a common target will suffice.

LEMMA EC.2. Fix a GFP CCE, with U_n , V_{VM} and V_{VM}^* as in Section 4.1.1. Extend the function ϕ of Section 4.1.1 by $\phi(\bar{u}) = 0$ for $\bar{u} \geq 1$, with $\bar{u} \ln \bar{u} = 0$ at 0. Then

$$\mathbb{E}[\text{Rev}] \geq \sum_{n \in Q} V_n^* \phi(U_n/V_n^*), \quad \mathbb{E}[\text{Rev}] + V_{\text{VM}} \geq V_{\text{VM}}^* + \sum_{n \in Q} V_n^* \phi(U_n/V_n^*),$$

where $V_n^* \phi(U_n/V_n^*) = 0$ by convention whenever $V_n^* = 0$.

Proof. We first bound the integrals of QL and VM acquired value under fixed deviations, then choose the targets. A QL advertiser can obtain utility 0 by choosing $\alpha_n = 1$, so $U_n \geq 0$. Under GFP, fixing $\alpha_n = \alpha' \in [0, 1)$ gives utility $(1 - \alpha')V_n(\alpha', \alpha_{-n})$. The CCE inequalities imply

$$\mathbb{E}[V_n(\alpha', \alpha_{-n})] \leq \frac{U_n}{1 - \alpha'}, \quad \int_0^{\bar{\alpha}_n} \mathbb{E}[V_n(\alpha', \alpha_{-n})] d\alpha' \leq -U_n \ln(1 - \bar{\alpha}_n) \quad (0 \leq \bar{\alpha}_n < 1).$$

For a VM advertiser, acquired value is nondecreasing in α_n , and choosing $\alpha_n = 1$ is ROI feasible because payment then equals value. Applying CCE to this fixed deviation gives

$$\sum_{n \in \mathcal{V}} \int_0^1 \mathbb{E}[V_n(\alpha', \alpha_{-n})] d\alpha' \leq \sum_{n \in \mathcal{V}} \mathbb{E}[V_n(1, \alpha_{-n})] \leq V_{\text{VM}}.$$

With the QL targets fixed, apply the recovery inequality (EC.3) first with VM targets 0 and then with VM targets one. Substituting the two integral bounds and rearranging yields

$$\mathbb{E}[\text{Rev}] \geq \sum_{n \in Q} [\bar{\alpha}_n V_n^* + U_n \ln(1 - \bar{\alpha}_n)], \quad \mathbb{E}[\text{Rev}] + V_{\text{VM}} \geq V_{\text{VM}}^* + \sum_{n \in Q} [\bar{\alpha}_n V_n^* + U_n \ln(1 - \bar{\alpha}_n)].$$

For $0 < U_n < V_n^*$, each bracket is maximized at $\bar{\alpha}_n = 1 - U_n/V_n^*$, where it equals $V_n^* \phi(U_n/V_n^*)$. If $U_n \geq V_n^*$ or $V_n^* = 0$, choose target 0. If $U_n = 0 < V_n^*$, let the target tend to one from below. $\square$

EC.3. Proofs for Section 4

We prove the GFP CCE and PNE bounds and the CE upper bounds, followed by the GSP and VCG theorem, a strict GSP equilibrium with vanishing welfare, and GSP with fixed deviations. The shared lemmas are proved in EC.2. We write the proofs of Theorems 1 and 2 for a general floor and ceiling (Section 5.2.1), so that they also give most of Proposition 5.

EC.3.1. Proof of Theorem 1

We begin with the main guarantee for GFP, which holds for every CCE and sharpens to $1 - 1/e$ for PNE.

THEOREM 1. $\text{PoA}_{\text{CCE}}(\text{GFP}) = 0.4569$ and $\text{PoA}_{\text{PNE}}(\text{GFP}) = 1 - 1/e$.

Proof. We prove Theorem 1 for any floor $A_{\min} \in [0, 1]$ and ceiling $A_{\max} \in [1, \infty]$ on the equilibrium multipliers, as in Section 5.2.1, where $A_{\min} = 0$ and $A_{\max} = \infty$ allow every multiplier. We first prove the 0.4569 guarantee for CCE, which uses only deviations and therefore holds for every floor and ceiling. We then show

that every CCE also keeps an $A_{\min}$ fraction of the optimal welfare, and that $\text{PoA}_{\text{PNE}}(\text{GFP}) = 1 - (1 - A_{\min})/e$. The case $A_{\min} = 0$ and $A_{\max} = \infty$ gives Theorem 1.

CCE Lower Bound

We use the notation U_n , U_{QL} , V_{VM} , P_{VM} , V_{QL}^* and V_{VM}^* of Section 4.1.1, and the bounds on integrals of acquired value under deviations from the proof of Lemma EC.2, first when the QL advertisers contribute nothing to the optimal welfare and then with a common QL target. The lemma also bounds the sum of revenue and VM value using a separate target for each QL advertiser:

$$V_{\text{VM}} + \mathbb{E}[\text{Rev}] \geq V_{\text{VM}}^* + \sum_{n \in Q} V_n^* \phi(U_n/V_n^*). \quad (\text{EC.7})$$

To convert revenue bounds into welfare bounds, use nonnegative payments and expected ROI to obtain $P_{\text{VM}} \leq \min\{V_{\text{VM}}, \mathbb{E}[\text{Rev}]\}$. Since QL value is utility plus payment,

$$\mathbb{E}[\text{Wel}] = U_{\text{QL}} + \mathbb{E}[\text{Rev}] + V_{\text{VM}} - P_{\text{VM}} \geq U_{\text{QL}} + \max\{\mathbb{E}[\text{Rev}], V_{\text{VM}}\}.$$

If $V_{\text{QL}}^* = 0$, the sum in Equation (EC.7) is 0 and $V_{\text{VM}}^* = \text{OPT}$. The maximum in the welfare bound is at least half the sum of its arguments, so $\mathbb{E}[\text{Wel}] \geq U_{\text{QL}} + \text{OPT}/2 \geq \text{OPT}/2$. The last inequality uses $U_{\text{QL}} \geq 0$, since choosing $\alpha_n = 1$ gives QL utility 0.

Now assume $V_{\text{QL}}^* > 0$. Give every QL advertiser the same target $0 \leq \bar{\alpha} < 1$ in the proof of Lemma EC.2. Its QL integral bound sums to $-U_{\text{QL}} \ln(1 - \bar{\alpha})$. Using VM targets 0 and one, respectively, in the recovery inequality (EC.3):

$$\mathbb{E}[\text{Rev}] \geq \bar{\alpha} V_{\text{QL}}^* + U_{\text{QL}} \ln(1 - \bar{\alpha}), \quad \mathbb{E}[\text{Rev}] + V_{\text{VM}} \geq V_{\text{VM}}^* + \bar{\alpha} V_{\text{QL}}^* + U_{\text{QL}} \ln(1 - \bar{\alpha}).$$

Optimizing the common QL term in these revenue bounds, as at the end of the proof of Lemma EC.2, gives the maximizing target $\bar{\alpha} = 1 - U_{\text{QL}}/V_{\text{QL}}^*$ when $0 < U_{\text{QL}} < V_{\text{QL}}^*$. Choose 0 when $U_{\text{QL}} \geq V_{\text{QL}}^*$, and let $\bar{\alpha}$ tend to one from below when $U_{\text{QL}} = 0$. In each case its supremum is $V_{\text{QL}}^* \phi(U_{\text{QL}}/V_{\text{QL}}^*)$. Substituting these revenue bounds into $\mathbb{E}[\text{Wel}] \geq U_{\text{QL}} + \max\{\mathbb{E}[\text{Rev}], V_{\text{VM}}\}$ gives

$$\mathbb{E}[\text{Wel}] \geq U_{\text{QL}} + \frac{V_{\text{VM}}^* + V_{\text{QL}}^* \phi(U_{\text{QL}}/V_{\text{QL}}^*)}{2}, \quad \mathbb{E}[\text{Wel}] \geq U_{\text{QL}} + V_{\text{QL}}^* \phi(U_{\text{QL}}/V_{\text{QL}}^*).$$

The first inequality again uses that the maximum of revenue and VM value is at least half their sum, and the second uses the revenue bound alone.

Normalize $\text{OPT} = 1$, and recall $\bar{u} = U_{\text{QL}}/V_{\text{QL}}^*$ from Section 4.1.1. The two preceding bounds are affine in the QL contribution V_{QL}^* . For $0 \leq \bar{u} \leq e^{-1}$, the first is nonincreasing and the second is nondecreasing in that contribution, and they meet at $V_{\text{QL}}^* = 1/(1 + \phi(\bar{u}))$. For $\bar{u} \geq e^{-1}$, the first bound is already at least $1/2$. Minimizing over $\bar{u}$ therefore gives

$$\mathbb{E}[\text{Wel}] \geq \min_{0 \leq \bar{u} \leq 1} \frac{1 + \bar{u} \ln \bar{u}}{2 - \bar{u} + \bar{u} \ln \bar{u}} > 0.4569, \quad (\text{EC.8})$$

where the minimum, about 0.4569, is attained at $\bar{u} \approx 0.1586$. In the case $V_{\text{QL}}^* = 0$, we showed that $\mathbb{E}[\text{Wel}] \geq \text{OPT}/2$. Hence every CCE satisfies $\mathbb{E}[\text{Wel}] \geq 0.4569 \text{ OPT}$, and $\text{PoA}_{\text{CCE}}(\text{GFP}) \geq 0.4569$.

CCE Bound from the Floor

At a realized equilibrium profile, the selected allocation maximizes the total weighted bid $B(\alpha)$, which equals GFP revenue (EC.2.3). Evaluating that objective at the optimal allocation and using $\alpha_n \geq A_{\min}$ gives

$$B(\alpha) = \max_{x \in \mathcal{X}} \sum_n \alpha_n V_n(x) \geq \sum_n \alpha_n V_n^* \geq A_{\min} \text{OPT}. \quad (\text{EC.9})$$

A QL advertiser can obtain utility 0 by choosing $\alpha_n = 1$, and each VM advertiser satisfies expected ROI. Thus

$$\mathbb{E}[\text{Wel}] - \mathbb{E}[\text{Rev}] = \sum_{n \in Q} \mathbb{E}[u_n] + \sum_{n \in \mathcal{V}} \mathbb{E}[V_n - P_n] \geq 0.$$

Taking expectations in the revenue lower bound in Equation (EC.9) gives $\mathbb{E}[\text{Wel}] \geq A_{\min} \text{OPT}$.

PNE Lower Bound

Fix a PNE α and suppress its argument in current values and payments. Under GFP, $P_n = \alpha_n V_n$, so every advertiser with $V_n > 0$ has $\alpha_n \leq 1$ by nonnegative QL utility or VM feasibility. If $A_{\min} = 1$, Equation (EC.9) gives efficiency. Otherwise, set $\bar{\alpha} = 1 - (1 - A_{\min})/e$. Applying the bound on omitted value in Equation (EC.4) with this common target and adding $\bar{\alpha} V^{\text{both}}$ gives

$$\bar{\alpha} \text{OPT} \leq \sum_n \left[\bar{\alpha} V_n^{\text{both}} + \alpha_n V_n^{\text{inef}} + \int_0^{\bar{\alpha}} [V_n(\max\{\alpha_n, \alpha'\}, \alpha_{-n}) - V_n] d\alpha' \right].$$

We now bound each advertiser's contribution by V_n .

VM advertisers. For a VM advertiser with $\alpha_n < \bar{\alpha}$, we use the feasible deviation to multiplier one and monotonicity of acquired value to obtain

$$V_n \leq V_n(\max\{\alpha_n, \alpha'\}, \alpha_{-n}) \leq V_n(1, \alpha_{-n}) \leq V_n \quad (0 \leq \alpha' \leq \bar{\alpha}).$$

Her integral is therefore 0. For any advertiser with $\alpha_n \geq \bar{\alpha}$, the integral is also 0 because every α' -raise with $\alpha' \leq \bar{\alpha}$ leaves α_n unchanged. In either case, the remaining terms are at most V_n , since $\bar{\alpha} \leq 1$, $V_n^{\text{both}} + V_n^{\text{inef}} = V_n$, and $\alpha_n \leq 1$ whenever $V_n > 0$.

QL advertisers. For a QL advertiser with $\alpha_n < \bar{\alpha}$, her incentive inequality implies

$$(1 - \alpha') V_n(\alpha', \alpha_{-n}) \leq (1 - \alpha_n) V_n \implies V_n(\alpha', \alpha_{-n}) - V_n \leq \frac{\alpha' - \alpha_n}{1 - \alpha'} V_n \quad (\alpha_n \leq \alpha' \leq \bar{\alpha}).$$

Thus her contribution to the upper bound on $\bar{\alpha} \text{OPT}$ is at most

$$V_n \left[\bar{\alpha} + \int_{\alpha_n}^{\bar{\alpha}} \frac{\alpha' - \alpha_n}{1 - \alpha'} d\alpha' \right] = V_n \left[\alpha_n + (1 - \alpha_n) \ln \left(\frac{1 - \alpha_n}{1 - \bar{\alpha}} \right) \right] \leq V_n.$$

The final inequality uses $\alpha_n \geq A_{\min}$, which gives $(1 - \alpha_n)/(1 - \bar{\alpha}) \leq e$. The same bounds apply when $V_n = 0$. Summing over advertisers gives $\text{Wel} \geq [1 - (1 - A_{\min})/e]\text{OPT}$.

PNE Upper-Bound Construction

If $A_{\min} = 1$, the claimed upper bound equals one and is immediate. Assume $A_{\min} < 1$. Take M unit-CTR slots and an integer $m' < M$. Advertiser n is QL, owns M ads with valuation one, and uses $A_{\min}$. Advertiser n' is VM, uses $\alpha_{n'} = 1$, and owns $M - m'$ ads whose i th smallest valuation is

$$\frac{i + m' A_{\min}}{m' + i}, \quad i = 1, \dots, M - m'.$$

Her other ads have valuation 0. Advertiser n'' is VM, uses $\alpha_{n''} = 1$, and owns m' ads with valuation $A_{\min}$. At ties at bid $A_{\min}$, advertiser n has priority. Advertiser n' receives all of her ads, advertiser n receives m' slots, and advertiser n'' receives none. Advertiser n has equilibrium utility $m'(1 - A_{\min})$. If advertiser n lowers α_n below $A_{\min}$, advertiser n'' receives the remaining m' slots and advertiser n receives no value. If α_n overtakes the i lowest bids of advertiser n' , then advertiser n receives $m' + i$ slots and must choose $\alpha_n \geq (i + m' A_{\min})/(m' + i)$. Her utility is therefore at most

$$(m' + i) \left(1 - \frac{i + m' A_{\min}}{m' + i} \right) = m' (1 - A_{\min}).$$

The utility bound holds for every fixed α_n and therefore, by linearity, for every randomized QL deviation. Advertiser n' already receives every ad she owns, so she cannot increase her value. Advertiser n'' receives no value when $\alpha_{n''} \leq 1$. For every $\alpha' \geq 0$, her ROI deficit is $P_{n''}(\alpha', \alpha_{-n''}) - V_{n''}(\alpha', \alpha_{-n''}) = (\alpha' - 1)V_{n''}(\alpha', \alpha_{-n''}) \geq 0$, with strict inequality whenever she receives positive value. A randomized deviation with positive expected value therefore violates expected ROI. Hence neither VM advertiser has a profitable randomized deviation satisfying expected ROI, and the profile is a PNE.

Since the optimal welfare is M , the equilibrium welfare ratio is

$$1 - \frac{m'}{M} (1 - A_{\min}) \sum_{i=1}^{M-m'} \frac{1}{m' + i}.$$

Choose $M \rightarrow \infty$ with $m'/M \rightarrow 1/e$. The harmonic sum satisfies $\sum_{i=1}^{M-m'} (m' + i)^{-1} \rightarrow \ln e = 1$, so the preceding welfare ratio tends to $1 - 1/e + A_{\min}/e$.

With $A_{\min} = 0$ and $A_{\max} = \infty$, these bounds prove Theorem 1. We discuss the relationship between our results and those of prior literature in Sections 1.2 and 2.1. $\square$

EC.3.2. Worst-Case Correlated Equilibria

The next proposition shows that the CCE bound of Theorem 1 and the bound for VM advertisers in Proposition 4 are tight, even for correlated equilibria (EC.2.2). Its constructions can also keep every multiplier in a fixed interval $[A_{\min}, A_{\max}]$, which we use for the equilibrium support in Section 5.2.

PROPOSITION EC.1. $\text{PoA}_{\text{CE}}(\text{GFP}) = 0.4569$, and $\text{PoA}_{\text{CE}}(\text{GFP}, \text{All-VM}) = 1/2$. With equilibrium support $[A_{\min}, A_{\max}]$ and undominated multipliers, both incur an additional error of $O(A_{\max}^{-1} + A_{\min})$ as $A_{\max}^{-1}, A_{\min} \rightarrow 0$.

Proof Overview. In both constructions, VM advertisers are recommended either an aggressive multiplier $A_{\max}$ or a conservative multiplier just above $A_{\min}$. The aggressive advertiser wins many low-valuation ads at the expense of higher-valuation ads, and the conservative advertisers keep only their high-valuation ads. The recommendation is random, so expected ROI holds on average, and no swap helps any advertiser. For the mixed bound, we add a QL advertiser whose valuations make every multiplier up to a threshold give her the same utility. This reproduces the worst case of the minimization in the proof of Theorem 1. $\square$

In both CE constructions below, each VM advertiser is recommended either a low multiplier $\alpha^{\text{lo}} < 1$, with which she receives one ad, or $A_{\max}$, with which she also receives additional ads. We verify her incentives by choosing a penalty $\mu \geq 0$ on payment in excess of value for which both recommendations attain the same reward $V_n - \mu(P_n - V_n)$ and no other α_n attains more. Equation (EC.10) in the lemma below ensures that her expected payment equals her expected value.

LEMMA EC.3. Consider GFP with $c_1 = \dots = c_M = 1$. Fix $A_{\max} > 1$, a low multiplier $\alpha^{\text{lo}} \in (0, 1)$, a common bid level $\bar{b} > 0$, and the numbers $N' \in \mathbb{N}$ of advertisers and $m' \in \mathbb{N}$ of additional ads per advertiser, satisfying

$$(N' - 1)(1 - \alpha^{\text{lo}}) = (A_{\max} - 1) \left(1 + \frac{\alpha^{\text{lo}} m'}{A_{\max}} \right). \quad (\text{EC.10})$$

Suppose N' VM advertisers each own one ad with valuation $\bar{b}/\alpha^{\text{lo}}$ and m' additional ads with valuation $\bar{b}/A_{\max}$, with all their other ads having valuation 0. Exactly one of these advertisers chosen uniformly at random is recommended $A_{\max}$, and every other is recommended α^{lo} . Assume the remaining bids and the tie-breaking rule ensure the following for each of these advertisers and every induced opposing profile.

- (i) With a multiplier below α^{lo} , she receives no ad.
- (ii) With a multiplier in $[\alpha^{\text{lo}}, A_{\max})$, she receives only the ad with valuation $\bar{b}/\alpha^{\text{lo}}$.
- (iii) With a multiplier of at least $A_{\max}$, she receives at most m' additional ads, and following recommendation $A_{\max}$ gives her all of them.

Then each of these advertisers satisfies expected ROI and has no profitable feasible generalized swap.

Proof. The high-valuation ad remains selected as α_n increases. Conditions (ii)–(iii) and unit CTRs therefore give the following value–payment pairs when advertiser n follows her recommendation:

$$(V_n, P_n) = \begin{cases} (\bar{b}/\alpha^{\text{lo}}, \bar{b}), & \alpha_n = \alpha^{\text{lo}}, \\ (\bar{b}/\alpha^{\text{lo}} + m' \bar{b}/A_{\max}, A_{\max} \bar{b}/\alpha^{\text{lo}} + m' \bar{b}), & \alpha_n = A_{\max}. \end{cases}$$

The high recommendation has probability $1/N'$. The deficit at this recommendation and the slack at the low recommendation cancel by Equation (EC.10), so expected ROI is binding:

$$\mathbb{E}_{\alpha \sim \sigma}[P_n - V_n] = \frac{\bar{b}}{N' \alpha^{\text{lo}}} \left[(A_{\max} - 1) \left(1 + \frac{\alpha^{\text{lo}} m'}{A_{\max}} \right) - (N' - 1)(1 - \alpha^{\text{lo}}) \right] = 0.$$

We choose μ to equalize $V_n - \mu(P_n - V_n)$ at the two recommendation pairs above. The bound $0 < \mu(A_{\max} - 1) < 1$ also ensures that acquired value has a positive coefficient at $\alpha_n = A_{\max}$:

$$\begin{aligned} \mu &= \frac{\alpha^{\text{lo}} m'}{A_{\max}(A_{\max} - \alpha^{\text{lo}}) + (A_{\max} - 1)\alpha^{\text{lo}} m'}, \quad 0 < \mu(A_{\max} - 1) < 1, \\ (1 + \mu(1 - \alpha^{\text{lo}})) \frac{\bar{b}}{\alpha^{\text{lo}}} &= (1 + \mu(1 - A_{\max})) \left(\frac{\bar{b}}{\alpha^{\text{lo}}} + \frac{m' \bar{b}}{A_{\max}} \right) > 0. \end{aligned}$$

Fix any induced α_{-n} and any deviation $\alpha' \geq 0$. Under GFP, $P_n(\alpha', \alpha_{-n}) = \alpha' V_n(\alpha', \alpha_{-n})$, so the reward at α' is $(1 + \mu(1 - \alpha')) V_n(\alpha', \alpha_{-n})$. We bound it in each of the three ranges of α' .

For $\alpha' < \alpha^{\text{lo}}$, condition (i) gives a reward of 0. For $\alpha^{\text{lo}} \leq \alpha' < A_{\max}$, condition (ii) gives the reward

$$(1 + \mu(1 - \alpha')) \frac{\bar{b}}{\alpha^{\text{lo}}} \leq (1 + \mu(1 - \alpha^{\text{lo}})) \frac{\bar{b}}{\alpha^{\text{lo}}}.$$

For $\alpha' \geq A_{\max}$, the acquired value is at most $\bar{b}/\alpha^{\text{lo}} + m' \bar{b}/A_{\max}$, the sum of the valuations of all her ads. We first use $\alpha' \geq A_{\max}$ and $V_n \geq 0$ to replace $1 + \mu(1 - \alpha')$ by $1 + \mu(1 - A_{\max})$, then apply this value bound with the positive coefficient $1 + \mu(1 - A_{\max})$:

$$\begin{aligned} (1 + \mu(1 - \alpha')) V_n(\alpha', \alpha_{-n}) &\leq (1 + \mu(1 - A_{\max})) V_n(\alpha', \alpha_{-n}) \\ &\leq (1 + \mu(1 - A_{\max})) \left(\frac{\bar{b}}{\alpha^{\text{lo}}} + \frac{m' \bar{b}}{A_{\max}} \right) = (1 + \mu(1 - \alpha^{\text{lo}})) \frac{\bar{b}}{\alpha^{\text{lo}}}. \end{aligned}$$

We note that this last case includes every deviation strictly above $A_{\max}$. Following either recommendation attains $(1 + \mu(1 - \alpha^{\text{lo}})) \bar{b}/\alpha^{\text{lo}}$. Averaging the bounds for all three ranges of α' over a feasible generalized swap, and using the equality of expected payment and value at equilibrium from Equation (EC.10), gives

$$\begin{aligned} \mathbb{E}[V_n(\tilde{\alpha}_n, \alpha_{-n})] &\leq \mathbb{E}[V_n(\tilde{\alpha}_n, \alpha_{-n}) - \mu(P_n(\tilde{\alpha}_n, \alpha_{-n}) - V_n(\tilde{\alpha}_n, \alpha_{-n}))] \\ &\leq (1 + \mu(1 - \alpha^{\text{lo}})) \frac{\bar{b}}{\alpha^{\text{lo}}} = \mathbb{E}_{\alpha \sim \sigma}[V_n - \mu(P_n - V_n)] = \mathbb{E}_{\alpha \sim \sigma}[V_n]. \end{aligned}$$

The first inequality uses feasibility of the swap, and the last equality uses the equality of expected payment and value at equilibrium. Hence the swap cannot increase expected value. $\square$

Proof of Proposition EC.1. Since every CE is a CCE, the lower bounds follow from the proofs of Theorem 1 and Proposition 4. For the upper bounds, we construct CE whose welfare ratios approach $1/2$ and the minimum in Equation (EC.8).

All-VM Advertisers

We construct a CE in which one of N' symmetric VM advertisers uses $A_{\max}$, while the others use the same $\alpha^{\text{lo}} < 1$. Each symmetric advertiser owns one high-valuation ad and m' additional lower-valuation ads. Assume $A_{\min} < 1$ and choose parameters satisfying expected ROI (Equation (EC.10)). For each sufficiently large integer m' , let N' be the smallest integer strictly larger than $1 + (A_{\max} - 1)(1 + m'A_{\min}/A_{\max})/(1 - A_{\min})$ and set the conservative recommendation

$$\alpha^{\text{lo}} \doteq \frac{N' - A_{\max}}{N' - 1 + (A_{\max} - 1)m'/A_{\max}}.$$

Then $N' \rightarrow \infty$ and $\alpha^{\text{lo}} \rightarrow A_{\min}$ from above as $m' \rightarrow \infty$. For all sufficiently large m' , $A_{\min} < \alpha^{\text{lo}} < 1$, and the balance identity in Equation (EC.10) holds with $\bar{b} = 1$.

There are N' symmetric VM advertisers. Each owns one ad with valuation $1/\alpha^{\text{lo}}$ and m' additional ads with valuation $A_{\max}^{-1}$. A final VM advertiser owns $m' + 1$ ads with valuation one and chooses $\alpha_n = 1$. There are $M = N' + m'$ unit-CTR slots. At ties at bid one, all high-valuation ads have first priority, followed by the additional ads of the symmetric advertisers and then the final advertiser's ads. Choose one symmetric advertiser uniformly, recommend $A_{\max}$ to her, and recommend α^{lo} to every other symmetric advertiser. At this recommended profile, the allocation contains all N' high-valuation ads and the additional ads of the advertiser recommended $A_{\max}$. Fix any symmetric advertiser. Below α^{lo} , all of her bids are below one, whereas the other high-valuation ads and the final advertiser provide $(N' - 1) + (m' + 1) = M$ opposing bids at least one. She therefore receives no ad. In $[\alpha^{\text{lo}}, A_{\max})$, she receives only her high-valuation ad. At or above $A_{\max}$, she receives at most all m' additional ads, and following recommendation $A_{\max}$ attains that allocation. These verify conditions (i)–(iii) of Lemma EC.3. Since the CTRs equal one and the balance identity in Equation (EC.10) holds, the lemma establishes expected ROI and optimality against all generalized swaps for the symmetric advertisers. The final advertiser receives value 0 whenever $\alpha_n \leq 1$. For $\alpha_n > 1$, her GFP payment exceeds acquired value whenever that value is positive. Since losing choices give both value and payment 0, randomizing over α_n cannot make a deviation that acquires value feasible. The proposed distribution is consequently a CE, and all of its multipliers lie in $[A_{\min}, A_{\max}]$.

Equilibrium welfare is $N'/\alpha^{\text{lo}} + m'/A_{\max}$, optimal welfare is $N'/\alpha^{\text{lo}} + m'$, and the balance identity in Equation (EC.10) implies $N'/(\alpha^{\text{lo}}m') \rightarrow (A_{\max} - 1)/[A_{\max}(1 - A_{\min})]$. Hence

$$\frac{\mathbb{E}[\text{Wel}]}{\text{OPT}} \rightarrow \frac{A_{\max} - A_{\min}}{2A_{\max} - 1 - A_{\max}A_{\min}} \quad \text{as } m' \rightarrow \infty. \quad (\text{EC.11})$$

The right side exceeds $1/2$ by $O(A_{\max}^{-1} + A_{\min})$, which proves the statement for all-VM advertisers.

QL and VM Advertisers

We add a QL advertiser to the preceding construction. Her ads have groups of equal valuations, chosen so that no choice of α_n gives utility above her equilibrium utility. Fix $\bar{u} \in (0, 1)$ as the limiting ratio $U_{\text{QL}}/V_{\text{QL}}^*$ used in the minimization in Equation (EC.8). For sufficiently small $A_{\min}$, set the QL recommendation to

$a_0 = A_{\min} < 1 - \bar{u}$ and choose a fine grid $a_0 < a_1 < \dots < a_L = 1 - \bar{u}$. For each k , group k contains m'_k QL ads with valuation $1/a_k$. Start with a large integer m'_0 and set $U_{\text{QL}} = (1 - a_0)m'_0/a_0$. We choose each subsequent group as large as possible while keeping the utility from a_k at most U_{QL} :

$$m'_k \doteq \left\lfloor a_k \left(\frac{U_{\text{QL}}}{1 - a_k} - \sum_{s=0}^{k-1} \frac{m'_s}{a_s} \right) \right\rfloor.$$

Thus $\sum_{s=0}^k m'_s/a_s \leq U_{\text{QL}}/(1 - a_k)$ at every grid point. Between adjacent grid points, the set of ads won is unchanged and GFP payment increases with α_n , so every QL choice $\alpha_n \in [0, 1]$ gives utility at most U_{QL} , while a_0 attains it. Choices $\alpha_n > 1$ give nonpositive GFP utility. Let $V_{\text{QL}} = m'_0/a_0$ be the QL advertiser's equilibrium value, let $V_{\text{QL}}^* = \sum_{k=0}^L m'_k/a_k$ be her contribution to optimal welfare, and let $m' = \sum_{k=1}^L m'_k$ be the number of additional QL ads. For a fixed grid, the rounding error in each m'_k is less than one, so it vanishes after normalization as $m'_0 \rightarrow \infty$. Refining the grid afterward gives $V_{\text{QL}}^*/U_{\text{QL}} \rightarrow 1/\bar{u}$ and

$$\frac{m'}{U_{\text{QL}}} \rightarrow \int_{a_0}^{1-\bar{u}} \frac{\alpha'}{(1-\alpha')^2} d\alpha' = \frac{1}{\bar{u}} + \ln \bar{u} - \frac{1}{1-a_0} - \ln(1-a_0).$$

Since $V_{\text{QL}} = U_{\text{QL}}/(1 - a_0)$, these limits imply

$$\frac{V_{\text{QL}}}{V_{\text{QL}}^*} = \frac{\bar{u}}{1-a_0} + o(1), \quad \frac{m'}{V_{\text{QL}}^*} = 1 + \bar{u} \ln \bar{u} - \frac{\bar{u}}{1-a_0} - \bar{u} \ln(1-a_0) + o(1). \quad (\text{EC.12})$$

The first ratio equals $\bar{u} + O(A_{\min})$, while the second equals $1 - \bar{u} + \bar{u} \ln \bar{u} + O(A_{\min}^2)$.

Use the all-VM CE construction leading to Equation (EC.11) with this m' , keeping its choices of N' and α^{lo} so that Equation (EC.10) holds. Add the QL advertiser, m'_0 slots, and m'_0 ads with valuation one for the final VM advertiser. Recommend a_0 to the QL advertiser. At bid one, place her ads after the symmetric advertisers' high-valuation ads but before their additional ads. All other valuations, recommendations, and priorities remain unchanged, and $M = N' + m'_0 + m'$.

We check only the allocation conditions affected by these additions. A symmetric VM advertiser choosing $\alpha_n < \alpha^{\text{lo}}$ still faces $(N' - 1) + (m'_0 + m' + 1) = M$ opposing bids at least one. She therefore receives no ad. The m'_0 new slots accommodate the initial QL group, so choosing $\alpha_n \in [\alpha^{\text{lo}}, A_{\max})$ still gives only the high-valuation ad, and following recommendation $A_{\max}$ still receives all m' additional ads. Thus Lemma EC.3 still gives expected ROI feasibility and rules out profitable generalized swaps. The final VM advertiser's incentive check is also unchanged, since choosing $\alpha_n \leq 1$ gives value 0, and any choice that acquires positive value pays more than that value.

For the QL advertiser, every ad whose bid reaches one fits alongside the N' high-valuation VM ads, because she owns $m'_0 + m'$ ads. The final advertiser excludes every lower QL bid. Hence, for $\alpha_n \leq 1$, the inequalities $\sum_{s=0}^k m'_s/a_s \leq U_{\text{QL}}/(1 - a_k)$ for the QL ad groups bound utility by U_{QL} against every opposing recommendation profile. Choices $\alpha_n > 1$ give nonpositive GFP utility. Linearity then rules out profitable generalized swaps, and the distribution is a CE, and its QL multiplier $a_0 = A_{\min}$ is at most one.

The allocation contains all high-valuation VM ads, the initial QL group, and the aggressive VM advertiser's additional ads. The optimal allocation instead assigns all high-valuation VM ads and all QL ads, since their valuations exceed one. Therefore

$$\frac{\mathbb{E}[\text{Wel}]}{\text{OPT}} = \frac{N'/\alpha^{\text{lo}} + V_{\text{QL}} + m'/A_{\text{max}}}{N'/\alpha^{\text{lo}} + V_{\text{QL}}^*}. \quad (\text{EC.13})$$

Substitute $N'/(\alpha^{\text{lo}} m') \rightarrow (A_{\text{max}} - 1)/[A_{\text{max}}(1 - A_{\text{min}})]$ from the all-VM construction leading to Equation (EC.11) and the QL value ratios in Equation (EC.12) into Equation (EC.13). Choosing $\bar{u} \approx 0.1586$, the minimizer in Equation (EC.8), makes the welfare ratio tend to that minimum plus $O(A_{\text{max}}^{-1} + A_{\text{min}})$ as $m'_0 \rightarrow \infty$ and the grid is refined. This proves the upper bound with both types of advertisers and equilibrium support $[A_{\text{min}}, A_{\text{max}}]$. Letting $A_{\text{max}}^{-1}, A_{\text{min}} \rightarrow 0$ gives it without a floor or ceiling. $\square$

EC.3.3. Proof of Theorem 2

We now turn to GSP and VCG, whose equilibria can lose almost all of the optimal welfare.

THEOREM 2. $\text{PoA}_{\text{PNE}}(\text{GSP}) = \text{PoA}_{\text{PNE}}(\text{VCG}) = 0$, even with one slot and two QL or two VM advertisers.

Proof. As in the proof of Theorem 1, we allow a floor $A_{\text{min}} \in [0, 1]$ and a ceiling $A_{\text{max}} \in [1, \infty)$ and construct single-slot PNE whose welfare ratios tend to $\max\{A_{\text{max}}^{-1}, A_{\text{min}}\}$. Letting $A_{\text{min}} = 0$ and $A_{\text{max}} \rightarrow \infty$ proves the theorem, with the advertisers all-QL or all-VM.

GSP – Single-Slot PNE Construction

When $\max\{A_{\text{max}}^{-1}, A_{\text{min}}\} = 1$, the upper bound is immediate. Otherwise, take one unit-CTR slot and one ad per advertiser. Advertiser n has valuation $v_{n,1} = 1$, uses A_{min} , and may be QL or VM. Advertiser n' is VM, uses A_{max} , and has valuation $v_{n',1} = (1 + \delta) \max\{A_{\text{max}}^{-1}, A_{\text{min}}\} < 1$, where $\delta > 0$ will tend to 0. Advertiser n' wins, pays $A_{\text{min}} < v_{n',1}$, and already acquires her maximum value. For every $\tilde{\alpha}_n$,

$$P_n(\tilde{\alpha}_n, \alpha_{n'}) = A_{\text{max}} v_{n',1} V_n(\tilde{\alpha}_n, \alpha_{n'}), \quad A_{\text{max}} v_{n',1} > 1.$$

When advertiser n wins, she pays $A_{\text{max}} v_{n',1}$ for value one, and otherwise both value and payment are 0. Since $A_{\text{max}} v_{n',1} > 1$, advertiser n pays more than her value whenever she wins, which rules out positive QL utility and positive expected value under any feasible randomized VM deviation. Thus the profile is a PNE for either type of advertiser n , and these multipliers are undominated under GSP. Its welfare ratio tends to $\max\{A_{\text{max}}^{-1}, A_{\text{min}}\}$.

GSP – All-QL

For the all-QL upper bound, take advertiser n to be QL in the single-slot construction above and change advertiser n' to QL. Whenever advertiser n' wins, she pays $A_{\text{min}} < v_{n',1}$, so her current $\alpha_{n'}$ maximizes utility. The single-slot incentive check for advertiser n is unchanged. Thus $\text{PoA}_{\text{PNE}}(\text{GSP}, \text{All-QL}) \leq$

$\max\{A_{\max}^{-1}, A_{\min}\}$. For clarity, this construction does not use undominated multipliers, because $\alpha_{n'} = A_{\max} > 1$.

VCG

With one slot and one ad per advertiser, GSP and VCG have identical allocations and payments, also after every unilateral deviation, so their incentive checks apply unchanged to VCG, whether the advertisers are all-QL, all-VM, or of both types. In every case, the welfare ratio is $(1 + \delta) \max\{A_{\max}^{-1}, A_{\min}\}$, and letting $\delta \rightarrow 0$ proves the upper bounds.

With one ad per advertiser, GSP with self-pricing coincides with GSP, so these PNE are also PNE under self-pricing. $\square$

EC.3.4. Proof of Proposition 1

Requiring every advertiser's multiplier to be her unique best response does not remove these low-welfare equilibria of GSP.

PROPOSITION 1. $\text{PoA}_{\text{PNE}}(\text{GSP}) = 0$, even over strict PNE.

In the single-slot construction in the proof of Theorem 2, the losing advertiser is indifferent among all losing multipliers. Strictness does not rescue GSP. Exact repricing at tied bids can make both advertisers' multipliers uniquely optimal while welfare still vanishes. We show strictness only and do not study whether learning dynamics reach this equilibrium.

Proof. Let $a \geq 4$ and $\eta \in (0, 1/4]$. *Instance.* Take six slots with CTRs

$$(c_1, \dots, c_6) = \left(1, \eta, \eta\left(1 - \frac{1}{2a}\right), \eta\left(1 - \frac{1}{a}\right), \eta\left(1 - \frac{1}{a}\right), \eta\left(1 - \frac{2}{a}\right)\right).$$

The QL advertiser has ads with valuations a and $a/2$ and uses multiplier $\alpha_{\text{QL}} = 1/a$, so her bids are 1 and $1/2$. The VM advertiser has ads with valuations $v_{\text{VM},1}$, $1/a$, $1/(2a)$, and $1/(2a)$, where $v_{\text{VM},1} \doteq 1 + \eta(1 - 2/a + 3/(2a^2))$, and uses multiplier $\alpha_{\text{VM}} = a$, so her bids are $av_{\text{VM},1}$, 1, $1/2$, and $1/2$. All other ads have valuation 0. We write (QL, j) and (VM, j) for the j th ads of the two advertisers. The priority order ranks $(\text{VM}, 1)$ and $(\text{VM}, 2)$ first, then $(\text{QL}, 1)$ and $(\text{QL}, 2)$, then $(\text{VM}, 3)$ and $(\text{VM}, 4)$, and then the ads with valuation 0.

Outcome. The ranking is $(\text{VM}, 1)$, $(\text{VM}, 2)$, $(\text{QL}, 1)$, $(\text{QL}, 2)$, $(\text{VM}, 3)$, $(\text{VM}, 4)$, because ad $(\text{VM}, 2)$ wins the tie at bid one and ad $(\text{QL}, 2)$ wins the tie at bid $1/2$. Under GSP, ads $(\text{VM}, 1)$ and $(\text{VM}, 2)$ pay the bid of $(\text{QL}, 1)$, which is one. Ads $(\text{QL}, 1)$ and $(\text{QL}, 2)$ pay the bid of $(\text{VM}, 3)$, which is $1/2$, and ads $(\text{VM}, 3)$ and $(\text{VM}, 4)$ pay 0. The VM advertiser's payment $P_{\text{VM}} = 1 + c_2$ equals her value $V_{\text{VM}} = v_{\text{VM},1} + c_2/a + (c_5 + c_6)/(2a)$, so her ROI constraint binds. Welfare equals $1 + 3a\eta/2$, while $\text{OPT} \geq a$ because ad $(\text{QL}, 1)$ could occupy the first slot. With $\eta = 1/a$, the welfare ratio is at most $5/(2a)$.

The QL advertiser. Her utility at $1/a$ is $u_{QL} = c_3(a - \frac{1}{2}) + c_4 \frac{a-1}{2}$. Lowering her multiplier to $[1/(2a), 1/a)$ moves (QL, 2) below (VM, 3) and (VM, 4) into slot six, giving $c_3(a - \frac{1}{2}) + c_6 \frac{a}{2}$. Below $1/(2a)$, both of her ads fall to slots five and six at price 0, giving $c_5a + c_6 \frac{a}{2}$. Raising it to $(1/a, 2/a]$ moves (QL, 1) above (VM, 2) into slot two, where it pays the bid of (VM, 2), which is one, giving $c_2(a - 1) + c_4 \frac{a-1}{2}$. On $(2/a, v_{VM,1}]$, ad (QL, 2) also passes (VM, 2) and pays one, giving $c_2(a - 1) + c_3(\frac{a}{2} - 1)$. Above $v_{VM,1}$, ad (QL, 1) wins the first slot at price $av_{VM,1} > a$, so her utility is below $c_3(\frac{a}{2} - 1)$. Each of these is strictly below u_{QL} , because

$$\begin{aligned} c_4(a - 1) - c_6a &= \frac{\eta}{a}, & u_{QL} - c_5a - c_6 \frac{a}{2} &= \frac{3\eta}{4a}, \\ c_3(a - \frac{1}{2}) - c_2(a - 1) &= \frac{\eta}{4a}, & u_{QL} - c_2(a - 1) - c_3(\frac{a}{2} - 1) &= \frac{\eta}{4} \left(1 + \frac{1}{a}\right). \end{aligned}$$

Thus $1/a$ is her unique best response.

The VM advertiser. Let $\mu \doteq 1/(2a^2 - 3a - 1)$. We show that multiplier a uniquely maximizes the reward $V_{VM} - \mu(P_{VM} - V_{VM})$ against the QL advertiser's bids. Her outcome changes only when one of her bids crosses one of the QL advertiser's bids. Relative to multiplier a , her deviations have the following effects.

- On $(a, 2a]$, ads (VM, 3) and (VM, 4) pass (QL, 2) into slots four and five at price $1/2$, changing the reward by $(1 + \mu)(c_4 - c_6)/(2a) - \mu(c_4 + c_5)/2$.
- Above $2a$, they also pass (QL, 1) into slots three and four at price one, changing it by $(1 + \mu)(c_3 + c_4 - c_5 - c_6)/(2a) - \mu(c_3 + c_4)$.
- On $[a/2, a)$, ad (VM, 2) falls below (QL, 1) into slot three at price $1/2$, changing it by $(1 + \mu)(c_3 - c_2)/a + \mu(c_2 - c_3/2)$.
- On $[1/v_{VM,1}, a/2)$, ad (VM, 2) falls into slot four at price 0, changing it by $(1 + \mu)(c_4 - c_2)/a + \mu c_2$.

Substituting the CTRs, all four changes are negative exactly when

$$\frac{1}{2a^2 - 2a - 1} < \mu < \frac{2}{2a^2 + a - 2},$$

and $\mu = 1/(2a^2 - 3a - 1)$ lies in this interval for $a \geq 4$. Below $1/v_{VM,1}$, ad (VM, 1) loses the first slot, so $V_{VM} \leq \eta(v_{VM,1} + 1) < 1/(1 + \mu)$ and the reward is at most $(1 + \mu)V_{VM} < 1$, whereas at multiplier a , where ROI binds, the reward equals $V_{VM} \geq 1$. Hence every feasible randomized deviation $\tilde{\alpha}_{VM}$ other than a satisfies

$$\begin{aligned} \mathbb{E}[V_{VM}(\tilde{\alpha}_{VM}, \alpha_{QL})] &\leq \mathbb{E}[(V_{VM} - \mu(P_{VM} - V_{VM}))(\tilde{\alpha}_{VM}, \alpha_{QL})] \\ &< (V_{VM} - \mu(P_{VM} - V_{VM}))(a, \alpha_{QL}) = V_{VM}(a, \alpha_{QL}), \end{aligned}$$

so it gives strictly less expected value. $\square$

With a floor and ceiling. Let $a = \min\{A_{\max}, 1/A_{\min}\}$, and suppose that $a \geq 4$. The strict PNE above uses the multipliers $1/a \in [A_{\min}, 1]$ for the QL advertiser and $a \in [1, A_{\max}]$ for the VM advertiser, which are undominated under GSP (Section 5.2.2). Its welfare ratio is at most $5/(2a) \leq \frac{5}{2}(A_{\max}^{-1} + A_{\min})$.

EC.3.5. Proof of Proposition 2

The low-welfare equilibria of GSP also remain when VM advertisers deviate only to fixed multipliers.

PROPOSITION 2. $\text{PoA}_{\text{PNE}}(\text{GSP}) = 0$, even over PNE with fixed deviations.

Proof. We simply use the single-slot PNE from the proof of Theorem 2 with advertiser n QL and advertiser n' VM. The QL advertiser uses $A_{\min}$ and the VM advertiser uses $A_{\max}$. Since both advertisers were already optimal against randomized deviations, they are also optimal against fixed ones, so this is still a PNE with fixed deviations. Its welfare ratio tends to $\max\{A_{\max}^{-1}, A_{\min}\}$, and letting $A_{\min} = 0$ and $A_{\max} \rightarrow \infty$ proves the proposition. $\square$

EC.4. Proofs for Section 5

We prove the results of Section 5 in the order of the main text. Along the way, we state the precise bounds with a floor and ceiling, in the notation of EC.1.2, and show that with a floor and ceiling, some instances have no CCE.

EC.4.1. Proof of Unit-Demand PNE Lower Bound

We begin with unit-demand instances, in which every PNE of GFP keeps at least $4/5$ of the optimal welfare.

PROPOSITION 3. *With unit-demand, $\text{PoA}_{\text{PNE}}(\text{GFP}) = 4/5$.*

We prove Proposition 3 in this subsection and the next.

Proof of the lower bound. With unit-demand, write $v_{n,1}$ for the valuation of advertiser n 's ad and $b_{n,1}$ for her bid. Her other ads have valuation 0, so they bid 0 at every multiplier. We group positions with equal CTRs and treat an ad in a slot with CTR 0 as unallocated. An allocated VM advertiser with a positive valuation below the top CTR must have $P_n = V_n$. Otherwise, mixing her current multiplier with a raise that wins a top position improves expected value while preserving ROI. The same argument rules out lowering her payment while retaining her CTR, as do QL incentives. In any inefficient equilibrium, all allocated bids below the top CTR therefore equal a common bid $\bar{b}$. Unallocated advertisers have valuations at most $\bar{b}$, and valuations above $\bar{b}$ below the top occur only among QL advertisers at the highest lower CTR. Whenever the allocation is inefficient, every top advertiser is VM. Thus an optimal allocation can be obtained by disjoint pairwise reassignments between these two groups, leaving all other welfare unchanged. We give the boundary cases and the pairing argument next.

We now justify these disjoint reassignments. Order advertisers by bid, and, with some abuse of notation, write c_n for advertiser n 's assigned CTR, with $c_n = 0$ when she is unallocated. Every allocated advertiser has $b_{n,1} \leq v_{n,1}$ by VM feasibility or nonnegative QL utility:

$$c_n b_{n,1} = P_n \leq V_n = c_n v_{n,1} \quad (\text{VM}), \quad u_n = c_n (v_{n,1} - b_{n,1}) \geq 0 \quad (\text{QL}).$$

If no advertiser with a positive valuation is allocated below the top CTR, any inefficiency would allow an excluded advertiser to outbid a lower-valuation allocated advertiser at a bid below her own valuation. This would improve QL utility or VM value while satisfying ROI. Thus we may assume that some advertiser with a positive valuation is allocated below the top CTR, and we let

$$\bar{c} \doteq \max\{c_n : v_{n,1} > 0, 0 < c_n < 1\}$$

be the highest such CTR.

A VM advertiser with a positive valuation allocated below the top CTR must be ROI-tight:

$$c_n b_{n,1} = c_n v_{n,1}, \text{ so } b_{n,1} = v_{n,1}, \quad \text{for every VM advertiser with } v_{n,1} > 0 \text{ and } 0 < c_n < 1.$$

Otherwise, a sufficiently small mixture with a finite bid above every opposing bid would preserve expected ROI and increase her value. No advertiser with a positive valuation allocated below the top CTR can therefore reduce her bid while retaining her CTR. A QL advertiser would gain utility, while a VM advertiser would create slack and then use the same mixture to improve on her original value.

Let $\bar{b} \doteq \max(\{0\} \cup \{b_{n,1} : c_n = 0\})$ be the largest bid among unallocated advertisers, or 0 if there is none. If a positive bid below the top CTR strictly exceeded the next bid, or had no bid below it, its owner could lower it while retaining her CTR. Hence each positive lower bid equals the bid just below it. Applying this down the bid order, every positive lower bid equals the largest unallocated bid, $\bar{b}$, and an allocated lower bid can be 0 only if all of them are, in which case $\bar{b} = 0$ as well. Thus every allocated lower position has bid $\bar{b} \geq 0$:

$$b_{n,1} = \bar{b} \quad \text{for every advertiser with } 0 < c_n < 1.$$

Every lower VM consequently has valuation $\bar{b}$, and every lower QL has valuation at least $\bar{b}$. An unallocated advertiser with valuation above $\bar{b}$ could bid slightly above $\bar{b}$, win a slot, and improve her objective while satisfying ROI. Thus excluded valuations are at most $\bar{b}$. A lower QL with valuation above $\bar{b}$ must receive CTR $\bar{c}$, because a sufficiently small increase above $\bar{b}$ would otherwise increase her CTR to at least $\bar{c}$ and strictly improve utility. All allocated positions below $\bar{c}$ therefore have valuation $\bar{b}$. In summary,

$$v_{n,1} \begin{cases} = \bar{b} & \text{for a VM advertiser with } 0 < c_n < 1, \\ \geq \bar{b}, \text{ with } v_{n,1} > \bar{b} \text{ only if } c_n = \bar{c} & \text{for a QL advertiser with } 0 < c_n < 1, \\ \leq \bar{b} & \text{for an unallocated advertiser } (c_n = 0). \end{cases}$$

If a top advertiser bids $\bar{b}$, any lower QL with valuation above $\bar{b}$ could similarly gain top CTR at a bid just above $\bar{b}$. Hence in this case all lower valuations equal $\bar{b}$ and the allocation is efficient. Otherwise every top bid exceeds $\bar{b}$. No top advertiser can then be QL, because she could lower her bid toward $\bar{b}$ while remaining in a top position. Thus in every inefficient equilibrium all top advertisers are VM, with valuations strictly above $\bar{b}$:

$$v_{n,1} \geq b_{n,1} > \bar{b} \quad \text{for every advertiser with } c_n = 1.$$

To obtain an optimal allocation, select the largest valuations among the top and CTR- $\bar{c}$ groups for the top positions. Pair each entrant with a distinct displaced top advertiser of smaller valuation, leaving ties in place. Every entrant is QL because lower VM valuations equal $\bar{b}$. Moving each displaced VM to her paired entrant's slot gives disjoint reassignments and leaves all other welfare unchanged. Let $\mathcal{P}$ be the set of pairs (n, n') of a displaced VM advertiser n and her entrant n' . All other advertisers contribute the same to welfare in both allocations, so

$$\text{Wel} - \sum_{(n, n') \in \mathcal{P}} (v_{n,1} + \bar{c} v_{n',1}) = \text{OPT} - \sum_{(n, n') \in \mathcal{P}} (v_{n',1} + \bar{c} v_{n,1}).$$

Consider a pair $(n, n') \in \mathcal{P}$ of a displaced top VM advertiser n and a higher-valuation QL advertiser n' at the highest lower CTR $\bar{c} < 1$. The QL incentive inequality and the VM advertiser's ROI constraint give

$$\bar{c}(v_{n',1} - \bar{b}) \geq v_{n',1} - b_{n,1}, \quad v_{n,1} \geq b_{n,1} \geq (1 - \bar{c})v_{n',1} + \bar{c}\bar{b} \geq (1 - \bar{c})v_{n',1}.$$

Hence each reassignment satisfies

$$\frac{v_{n,1} + \bar{c}v_{n',1}}{v_{n',1} + \bar{c}v_{n,1}} \geq \frac{1}{1 + \bar{c}(1 - \bar{c})} \geq \frac{4}{5}.$$

Summing over these disjoint reassignments and adding the unchanged welfare of the other advertisers gives $\text{Wel} \geq (4/5)\text{OPT}$. $\square$

EC.4.2. Matching Unit-Demand PNE Upper Bound Construction

We now show that the bound $4/5$ in Proposition 3 is tight.

Proof of the upper bound. To obtain the matching upper bound, fix $0 < \delta < 1$. Take two slots with CTRs $(1, 1/2)$ and three advertisers, each with one ad. Their types are (VM, QL, QL), their valuations are $((1 + \delta)/2, 1, \delta)$, and $\alpha = (1, \delta, 1)$. Give advertiser 2 priority over advertiser 3 at ties. Advertiser 1 wins the first slot with ROI equality and already acquires her maximum possible value. Advertiser 2 wins the second slot at bid δ , obtaining utility $(1 - \delta)/2$. A smaller bid loses, a bid between δ and $(1 + \delta)/2$ cannot improve her utility, and a bid that wins the first slot gives utility at most $(1 - \delta)/2$. Advertiser 3 cannot win at a bid below her valuation, so she cannot obtain positive utility. Thus the profile is a PNE. The welfare ratio is

$$\frac{\text{Wel}}{\text{OPT}} = \frac{(1 + \delta)/2 + 1/2}{1 + (1 + \delta)/4} = \frac{4 + 2\delta}{5 + \delta} \rightarrow \frac{4}{5} \quad \text{as } \delta \rightarrow 0.$$

$\square$

EC.4.3. Proof of Proposition 4

We next prove the exact guarantees of GFP when advertisers share the same objective.

PROPOSITION 4. $\text{PoA}_{\text{CCE}}(\text{GFP}) = 1/2$ if advertisers are all-VM, and $\text{PoA}_{\text{CCE}}(\text{GFP}) = 1 - 1/e$ if all-QL.

Proof. We first prove the lower bounds and then show that both are tight.

Lower Bounds

When advertisers are all-VM, $V_{\text{QL}}^* = 0$, and the CCE lower bound in EC.3.1 gives $\mathbb{E}[\text{Wel}] \geq \text{OPT}/2$. When advertisers are all-QL, we apply the recovery inequality (EC.3) with every target equal to $1 - 1/e$. The QL CCE inequality gives $\mathbb{E}[V_n(\alpha', \alpha_{-n})] \leq \mathbb{E}[u_n]/(1 - \alpha')$ for $\alpha' < 1$. Since $\int_0^{1-1/e} (1 - \alpha')^{-1} d\alpha' = 1$, we obtain $(1 - 1/e)\text{OPT} \leq \mathbb{E}[\text{Rev}] + \sum_n \mathbb{E}[u_n] = \mathbb{E}[\text{Wel}]$.

Upper Bounds

For all-VM advertisers, Proposition EC.1 gives $\text{PoA}_{\text{CCE}}(\text{GFP}, \text{All-VM}) = 1/2$. Every CE is a CCE, so $\text{PoA}_{\text{CCE}}(\text{GFP}, \text{All-VM}) \leq 1/2$. For all-QL advertisers, the bound follows from the CCE construction in the single-slot example of Feldman et al. (2016b) achieving welfare $1 - 1/e$. $\square$

EC.4.4. Proof of Proposition 5

We now bound the PoA when every equilibrium multiplier lies between a floor and a ceiling.

PROPOSITION 5. *With equilibrium support $[A_{\min}, A_{\max}]$, $\text{PoA}_{\text{PNE}}(\text{GFP}) = 1 - (1 - A_{\min})/e$, and as $A_{\max}^{-1}, A_{\min} \rightarrow 0$,*

$$\text{PoA}_{\text{CCE}}(\text{GFP}) = 0.4569 + O(A_{\max}^{-1} + A_{\min}), \quad \text{PoA}(\text{GSP}), \text{PoA}(\text{VCG}) = \Theta(A_{\max}^{-1} + A_{\min}).$$

Except for the lower bound for GSP and VCG, the bounds of Proposition 5 come directly from the proofs of Theorems 1 and 2, which allow a floor and ceiling (EC.3.1 and EC.3.3). We state the precise bounds for GFP and for GSP and VCG, prove the remaining lower bound, and then combine them to prove Proposition 5.

EC.4.4.1. GFP. The next proposition gives the precise bounds for GFP, including those for homogeneous objectives.

PROPOSITION EC.2. *With equilibrium support $[A_{\min}, A_{\max}]$, with or without undominated multipliers, $\text{PoA}_{\text{PNE}}(\text{GFP}) = 1 - (1 - A_{\min})/e$ and*

$$\text{PoA}_{\text{CCE}}(\text{GFP}) \geq \max\{0.4569, A_{\min}\}, \quad \text{PoA}_{\text{CCE}}(\text{GFP}, \text{All-VM}) \geq \max\{1/2, A_{\min}\},$$

$$\text{PoA}_{\text{CCE}}(\text{GFP}, \text{All-QL}) \geq \max\{1 - 1/e, A_{\min}\}.$$

Without undominated multipliers, also $\text{PoA}_{\text{CCE}}(\text{GFP}, \text{All-QL}) \leq 1 - 1/e + O(A_{\max}^{-1} + A_{\min})$ as $A_{\max}^{-1}, A_{\min} \rightarrow 0$.

Proof. The proofs of Theorem 1 and Proposition 4 use only deviations to multipliers in $[0, 1]$, which remain available with or without undominated multipliers, and they do not require QL multipliers to be at most one in equilibrium. They therefore give the CCE lower bounds 0.4569, $1/2$, and $1 - 1/e$. For all-QL advertisers, the same straightforward equilibrium support analysis as in the all-VM and mixed constructions

of Proposition EC.1 carries over to the upper-bound construction of Feldman et al. (2016b). The proof of Theorem 1 also gives the bound $A_{\min}$ for CCE and the PNE statement (EC.3.1). Its PNE construction uses the QL multiplier $A_{\min} \leq 1$ and VM multipliers equal to one, so it remains a PNE with undominated multipliers (Lemma 1). $\square$

EC.4.4.2. GSP and VCG. The next proposition gives the precise bounds for GSP and VCG.

PROPOSITION EC.3. *With equilibrium support $[A_{\min}, A_{\max}]$, as $A_{\max}^{-1}, A_{\min} \rightarrow 0$,*

$$\text{PoA}(\mathcal{M}), \text{PoA}(\mathcal{M}, \text{All-VM}), \text{PoA}(\mathcal{M}, \text{All-QL}) = \Theta(A_{\max}^{-1} + A_{\min}), \quad \mathcal{M} \in \{\text{GSP}, \text{VCG}\},$$

and with undominated multipliers, the same holds for $\text{PoA}(\text{GSP})$ and $\text{PoA}(\text{GSP}, \text{All-VM})$. In each case, single-slot PNE have welfare ratios arbitrarily close to $\max\{A_{\max}^{-1}, A_{\min}\}$.

The next lemma gives the lower bound in Proposition EC.3. It uses only a fixed deviation of either the ceiling or the floor, along with Lemma EC.1(iii) and the fact that neither mechanism charges more than the bid.

LEMMA EC.4. *Let $\mathcal{M} \in \{\text{GSP}, \text{VCG}\}$, with equilibrium support $[A_{\min}, A_{\max}]$, with or without undominated multipliers, and for any mix of objectives. Every CCE satisfies*

$$\mathbb{E}[\text{Wel}] \geq \max \left\{ \frac{1}{2A_{\max} + 2}, \frac{A_{\min}}{1 + A_{\min}} \right\} \text{OPT} \geq \frac{1}{8} (A_{\max}^{-1} + A_{\min}) \text{OPT}.$$

Proof. We first use the ceiling. Fix a realized profile and apply the recovery inequality (EC.3) with every target equal to $1/2$. Acquired value is nondecreasing in the advertiser's own multiplier, so each integral is at most $\frac{1}{2} V_n(\frac{1}{2}, \alpha_{-n})$. Every multiplier is at most $A_{\max}$, so the total weighted bid satisfies $B(\alpha) = \sum_n \alpha_n V_n(\alpha) \leq A_{\max} \text{Wel}(\alpha)$. Multiplying by two,

$$\text{OPT} \leq 2A_{\max} \text{Wel}(\alpha) + \sum_n V_n(\frac{1}{2}, \alpha_{-n}).$$

Under both mechanisms, an advertiser pays at most her bid. Multiplier $1/2$ therefore gives a QL advertiser utility at least $\frac{1}{2} V_n(\frac{1}{2}, \alpha_{-n})$ and satisfies ROI for a VM advertiser. The CCE conditions give $\mathbb{E}[V_n(\frac{1}{2}, \alpha_{-n})] \leq 2\mathbb{E}[u_n] \leq 2\mathbb{E}[V_n]$ for QL advertisers and $\mathbb{E}[V_n(\frac{1}{2}, \alpha_{-n})] \leq \mathbb{E}[V_n]$ for VM advertisers. Hence $\text{OPT} \leq (2A_{\max} + 2)\mathbb{E}[\text{Wel}]$.

We next use the floor. Every multiplier is at least $A_{\min}$, so Lemma EC.1(iii) gives $A_{\min} V^{\text{omit}} \leq \text{Rev}$ at every profile, and Equation (EC.1) gives $\text{OPT} \leq \text{Wel} + V^{\text{omit}}$. A QL advertiser obtains nonnegative utility at multiplier one, and VM advertisers satisfy expected ROI, so $\mathbb{E}[\text{Rev}] \leq \mathbb{E}[\text{Wel}]$. Hence

$$\text{OPT} \leq \mathbb{E}[\text{Wel}] + \frac{\mathbb{E}[\text{Rev}]}{A_{\min}} \leq \left(1 + \frac{1}{A_{\min}}\right) \mathbb{E}[\text{Wel}].$$

The second inequality in the lemma follows from $1/(2A_{\max} + 2) \geq 1/(4A_{\max})$, $A_{\min}/(1 + A_{\min}) \geq A_{\min}/2$, and the fact that a maximum is at least the average of its arguments. $\square$

Proof of Proposition EC.3. Lemma EC.4 gives the lower bound over all instances, and hence over the all-VM and all-QL instances, with or without undominated multipliers. For the upper bound, the single-slot PNE in the proof of Theorem 2 (EC.3.3) have welfare ratios tending to $\max\{A_{\max}^{-1}, A_{\min}\} \leq A_{\max}^{-1} + A_{\min}$ for all-QL, all-VM, and mixed objectives, under both mechanisms. Except in the all-QL construction, their multipliers are undominated under GSP. Since every PNE is a CCE, the two bounds give the stated order for every equilibrium class. We end with a multi-slot construction used for Propositions 6 and EC.5.

Undominated Multipliers – Multi-Slot PNE Construction

Fix $\delta \in (0, A_{\max} - 1)$ and set $\bar{v} = (1 + \delta)/A_{\max} < 1$. Take M unit-CTR slots. Advertiser n' is VM, uses $A_{\max}$, and owns $M - 1$ ads with valuation $\bar{v}$ and one ad with valuation $v_{n',1} = (M - 1)(1 - \bar{v})$. Advertiser n uses multiplier one, may be QL or VM, and owns $M - 1$ ads with valuation one. For sufficiently large M , $v_{n',1} > 1$, so every bid of advertiser n' , at least $A_{\max}\bar{v} = 1 + \delta$, exceeds every bid of advertiser n , and advertiser n' fills all slots. Under VCG, removing her would allocate the $M - 1$ ads of advertiser n , so $P_{n'} = M - 1 = V_{n'}$. Advertiser n' therefore satisfies ROI and already acquires all of her value. Whenever advertiser n wins k slots, she displaces bids of at least $1 + \delta$, so she pays at least $(1 + \delta)k = (1 + \delta)V_n$. This rules out positive QL utility and positive value under any feasible randomized VM deviation, so the profile is a PNE. The optimal allocation keeps ad $(n', 1)$ and replaces her other ads with the ads of advertiser n , so

$$\frac{\text{Wel}}{\text{OPT}} = \frac{V_{n'}}{v_{n',1} + M - 1} \rightarrow \frac{1}{2 - \bar{v}} \quad \text{as } M \rightarrow \infty,$$

which tends to $1/(2 - A_{\max}^{-1}) = 1/2 + O(A_{\max}^{-1})$ as $\delta \rightarrow 0$. $\square$

EC.4.4.3. Completing the Proof of Proposition 5. The prices of anarchy in Proposition 5 are those over all instances in Propositions EC.1 to EC.3. For GFP, Proposition EC.2 gives the PNE price of anarchy $1 - (1 - A_{\min})/e$, and the proof of Theorem 1 shows that the CCE price of anarchy is at least the minimum in Equation (EC.8), about 0.4569. Since every CE is a CCE, Proposition EC.1 shows that it is at most this minimum plus $O(A_{\max}^{-1} + A_{\min})$ as $A_{\max}^{-1}, A_{\min} \rightarrow 0$. For GSP and VCG, Proposition EC.3 gives the order $A_{\max}^{-1} + A_{\min}$. Its lower bound holds for every CCE and its upper bound is attained by PNE, so the order applies to both equilibrium notions. $\square$

EC.4.5. Alternative Assumptions for GSP with a Floor and Ceiling

The three changes considered in Section 4.2—strict PNE, self-pricing, and deviations to fixed multipliers—do not restore a welfare guarantee for GSP when every equilibrium multiplier lies in $[A_{\min}, A_{\max}]$, with or without undominated multipliers. For self-pricing and strict PNE, this follows from the proofs of Theorem 2 and Proposition 1 (EC.3.3 and EC.3.4). For deviations to fixed multipliers, the next proposition is the analogue of Proposition 2.

PROPOSITION EC.4. *Suppose that VM advertisers deviate only to fixed multipliers. With equilibrium support $[A_{\min}, A_{\max}]$, $\text{PoAPNE}(\text{GSP}) \leq \max\{A_{\max}^{-1}, A_{\min}\}$.*

Proof. We simply use the single-slot PNE in EC.3.3, in which the QL advertiser uses $A_{\min}$ and the VM advertiser uses $A_{\max}$. Since both advertisers were already optimal against randomized deviations, they are also optimal against fixed ones, so this is still a PNE. Its welfare ratio tends to $\max\{A_{\max}^{-1}, A_{\min}\}$. $\square$

EC.4.6. Proof of Lemma 1

Table EC.4 Intervals $\mathcal{A}_n$ of equilibrium multipliers with a floor and ceiling.

Multipliers	QL advertiser n			VM advertiser n		
	GFP	GSP	VCG	GFP	GSP	VCG
All	$[A_{\min}, A_{\max}]$	$[A_{\min}, A_{\max}]$	$[A_{\min}, A_{\max}]$	$[A_{\min}, A_{\max}]$	$[A_{\min}, A_{\max}]$	$[A_{\min}, A_{\max}]$
Undominated	$[A_{\min}, 1]$	$[A_{\min}, 1]$	$\{1\}$	$[A_{\min}, A_{\max}]$	$[A_{\min}, A_{\max}]$	$[1, A_{\max}]$

These intervals restrict only the multipliers used in equilibrium, and every deviation may use any nonnegative multiplier.

Weak dominance allows us to exclude certain multipliers at equilibrium (Section 5.2.2). Table EC.4 lists the resulting intervals of equilibrium multipliers with a floor and ceiling.

LEMMA 1. *In every instance, with any mix of advertiser objectives, multiplier one weakly dominates every QL multiplier above one under GFP and GSP. Under VCG, it weakly dominates every other QL multiplier and every VM multiplier below one.*

Proof. Fix α_{-n} .

GFP. A QL advertiser with $\alpha_n > 1$ satisfies

$$u_n(\alpha_n, \alpha_{-n}) = (1 - \alpha_n)V_n(\alpha_n, \alpha_{-n}) \leq 0 = u_n(1, \alpha_{-n}).$$

Thus choosing $\alpha_n = 1$ weakly improves utility.

GSP. Lower $\alpha_n > 1$ to one for the QL advertiser. Her own ads retain their order. Consider an adjacent exchange in which an owned ad with valuation v moves from slot m to slot $m + 1$, below an opposing bid $b \geq v$. Its GSP price is b immediately before the exchange and some $p \leq b$ afterward. The utility change from this ad is

$$c_{m+1}(v - p) - c_m(v - b) = (c_m - c_{m+1})(b - v) + c_{m+1}(b - p) \geq 0.$$

For the boundary at the last allocated slot, we set $c_{M+1} = 0$. Every other owned ad keeps its rank and the same opposing ads below it. Its GSP price and utility contribution are therefore unchanged. Crossings below the allocated slots have no effect on utility. For completeness, at a simultaneous tie the same expression for the utility change is nonnegative at each successive adjacent exchange consistent with the fixed priority

rule. Between exchanges the allocation and GSP payments are constant. Summing the nonnegative changes proves $u_n(1, \alpha_{-n}) \geq u_n(\alpha', \alpha_{-n})$.

VCG (QL advertisers). Fix $\alpha' \geq 0$, and let $B_{-n}(\mathbf{x})$ be the other advertisers' total weighted bid under allocation $\mathbf{x}$, as defined in EC.2.3. By allocation optimality at the profiles (α', α_{-n}) and $(1, \alpha_{-n})$,

$$\begin{aligned} \alpha' V_n(\alpha', \alpha_{-n}) + B_{-n}(\mathbf{x}(\alpha', \alpha_{-n})) &\geq \alpha' V_n(1, \alpha_{-n}) + B_{-n}(\mathbf{x}(1, \alpha_{-n})), \\ V_n(1, \alpha_{-n}) + B_{-n}(\mathbf{x}(1, \alpha_{-n})) &\geq V_n(\alpha', \alpha_{-n}) + B_{-n}(\mathbf{x}(\alpha', \alpha_{-n})). \end{aligned}$$

The term $\max_{\mathbf{x} \in \mathcal{X}_{-n}} B_{-n}(\mathbf{x})$ in the VCG payment is independent of α' , so it cancels when we subtract utilities. Applying the second inequality gives

$$u_n(1, \alpha_{-n}) - u_n(\alpha', \alpha_{-n}) = V_n(1, \alpha_{-n}) + B_{-n}(\mathbf{x}(1, \alpha_{-n})) - V_n(\alpha', \alpha_{-n}) - B_{-n}(\mathbf{x}(\alpha', \alpha_{-n})) \geq 0.$$

VCG (VM advertisers). For $\alpha' < 1$, adding the two optimality inequalities cancels the opposing weighted bids and gives

$$(1 - \alpha')(V_n(1, \alpha_{-n}) - V_n(\alpha', \alpha_{-n})) \geq 0 \implies V_n(1, \alpha_{-n}) \geq V_n(\alpha', \alpha_{-n}).$$

The nonnegative utility difference just proven is also an increase in the VM advertiser's ROI slack $V_n - P_n$:

$$V_n(1, \alpha_{-n}) - P_n(1, \alpha_{-n}) \geq V_n(\alpha', \alpha_{-n}) - P_n(\alpha', \alpha_{-n}).$$

Both inequalities hold for every α_{-n} . For any randomized deviation $\tilde{\alpha}_n$, we replace each multiplier below one by one, giving $\max\{1, \tilde{\alpha}_n\}$. Taking expectations of these inequalities, we obtain

$$\mathbb{E}[V_n(\max\{1, \tilde{\alpha}_n\}, \alpha_{-n})] \geq \mathbb{E}[V_n(\tilde{\alpha}_n, \alpha_{-n})], \quad \mathbb{E}[(V_n - P_n)(\max\{1, \tilde{\alpha}_n\}, \alpha_{-n})] \geq \mathbb{E}[(V_n - P_n)(\tilde{\alpha}_n, \alpha_{-n})].$$

If the original deviation is feasible, the right side of the preceding inequality for expected ROI slack is nonnegative. The replacement therefore preserves expected ROI and weakly increases value. $\square$

EC.4.7. Proof of Proposition 6

We next show that excluding weakly dominated multipliers improves the guarantee of VCG, but not those of GFP and GSP.

PROPOSITION 6. *With undominated multipliers, the following hold.*

- $\text{PoA}_{\text{CCE}}(\text{VCG}) = \text{PoA}_{\text{PNE}}(\text{VCG}) = 1/2$ with tightness at an all-VM instance.
- The PoA guarantees for CCE and PNE of both GFP (Theorem 1 and Proposition 4) and GSP (Theorem 2 and Proposition 1) are unchanged.

Proof. We first prove the bounds for VCG and then show that those for GFP and GSP are unchanged.

VCG

With undominated multipliers, QL advertisers use $\alpha_n = 1$ and VM advertisers use $\alpha_n \geq 1$ under VCG. Every multiplier is then at least one, so Lemma EC.1(iii) gives $V^{\text{omit}} \leq \text{Rev}$, and Equation (EC.1) gives $\text{OPT} \leq \text{Wel} + \text{Rev}$ at every profile. QL advertisers obtain nonnegative utility and VM advertisers satisfy expected ROI, so $\mathbb{E}[\text{Rev}] \leq \mathbb{E}[\text{Wel}]$ and $\text{OPT} \leq 2\mathbb{E}[\text{Wel}]$ for every CCE. For the matching upper bound, use the multi-slot PNE construction in EC.4.4.2 with any $A > 1$ in place of $A_{\max}$. It uses multiplier A for the winning VM advertiser and one for the other advertiser, so these multipliers are undominated whether the latter advertiser is QL or VM. Its welfare ratio tends to $1/(2 - (1 + \delta)/A)$ as $M \rightarrow \infty$, and letting $\delta \rightarrow 0$ and $A \rightarrow \infty$ gives $\text{PoA}_{\text{PNE}}(\text{VCG}) \leq 1/2$. As for the full efficiency when advertisers are all-QL, every QL advertiser uses $\alpha_n = 1$, so the allocation maximizes welfare and $\text{Wel} = \text{OPT}$.

GFP and GSP

We note that excluding weakly dominated multipliers only removes equilibria, so the lower bounds of Theorems 1 and 2 still hold. For the upper bounds, the CE of Proposition EC.1 use QL multipliers of at most one, and the PNE in the proof of Theorem 1 uses QL multiplier 0, and under GFP, both are undominated. Under GSP, the strict PNE in the proof of Proposition 1 uses the multiplier $1/a < 1$ for the QL advertiser and a for the VM advertiser, which are undominated, and its welfare ratio vanishes as $a \rightarrow \infty$. Its multipliers are unique best responses, so they are not weakly dominated in that instance either. The prices of anarchy of GFP and GSP are therefore unchanged. $\square$

With a floor and ceiling, Propositions EC.1 to EC.3 show that the bounds of GFP and GSP are again unchanged by excluding weakly dominated multipliers, and the argument above gives a slightly larger guarantee for VCG.

PROPOSITION EC.5. *With equilibrium support $[A_{\min}, A_{\max}]$ and undominated multipliers, $\text{PoA}(\text{VCG})$ and $\text{PoA}(\text{VCG}, \text{All-VM})$ are $1/2 + O(A_{\max}^{-1})$, and $\text{PoA}(\text{VCG}, \text{All-QL}) = 1$.*

Proof. With undominated multipliers under VCG, QL advertisers use $\alpha_n = 1$ and VM advertisers use $\alpha_n \in [1, A_{\max}]$. Every multiplier then lies in $[1, A_{\max}]$, so Lemma EC.1(iii) gives $V^{\text{omit}} \leq \text{Rev}$, and Equation (EC.2) gives $V^{\text{omit}} \leq B^{\text{inef}} \leq A_{\max} V^{\text{inef}}$ at every profile. The welfare identity in Equation (EC.1) then gives $\text{OPT} \leq \text{Wel} + (1 - A_{\max}^{-1})\text{Rev}$ at every profile. QL advertisers obtain nonnegative utility and VM advertisers satisfy expected ROI, so $\mathbb{E}[\text{Rev}] \leq \mathbb{E}[\text{Wel}]$ and $\text{OPT} \leq (2 - A_{\max}^{-1})\mathbb{E}[\text{Wel}]$ for every CCE. For the matching upper bound, use the multi-slot PNE construction in EC.4.4.2. It uses multiplier $A_{\max}$ for the winning VM advertiser and one for the other advertiser, so these multipliers are undominated whether the latter advertiser is QL or VM. As for the full efficiency when advertisers are all-QL, every QL advertiser uses $\alpha_n = 1$, so the allocation maximizes welfare and $\text{Wel} = \text{OPT}$. $\square$

EC.4.8. Proof of Corollary 1

We now extend our CCE guarantees to stochastic, possibly correlated valuations.

COROLLARY 1. *The CCE lower bounds of Theorem 1 and Propositions 4 and 6 hold for PoA_{BCCE} . In particular, $\text{PoA}_{\text{BCCE}}(\text{GFP}) = 0.4569$. Moreover, $\text{PoA}_{\text{BCCE}}(\text{VCG}) = 1/2$ with undominated multipliers.*

Proof. All expectations are under the joint distribution of valuations and play in a Bayesian CCE (Definition 5). We keep the allocation and revenue inequalities of Lemmas EC.1 and EC.2 and the calculation in EC.3.1. Only the incentive comparisons change.

GFP

For a fixed $a \in [0, 1)$, the strategy that always uses multiplier a bids $av_{n,j}$ at the current valuation vector and uses no information about the other advertisers, so it is a valid deviation. QL optimality and the GFP mechanism give

$$(1 - a) \mathbb{E}[V_n(\mathbf{v}, [a, \alpha_{-n}])] \leq U_n \doteq \mathbb{E}[u_n(\mathbf{v}, \alpha)],$$

and comparison with multiplier one gives $U_n \geq 0$. For a VM advertiser, multiplier one pays exactly its acquired value in every realization, so it is a valid deviation, and

$$\mathbb{E}[V_n(\mathbf{v}, [1, \alpha_{-n}])] \leq \mathbb{E}[V_n], \quad \mathbb{E}[P_n] \leq \mathbb{E}[V_n].$$

In this proof only, let V_{QL}^* and V_{VM}^* be the two populations' expected contributions to the realized optimal allocation, so $V_{\text{QL}}^* + V_{\text{VM}}^* = \mathbb{E}[\text{OPT}(\mathbf{v})]$, and let U_{QL} , V_{VM} and P_{VM} be the population sums of Section 4.1.1 under these expectations. Fix a common QL target $\bar{\alpha} < 1$ and integrate the QL comparison up to it. For VM advertisers, acquired value is nondecreasing in the advertiser's own multiplier, so each integral up to one is at most the deviation value at one. Applying Lemma EC.1(i) at every pair of valuations and multipliers, averaging, and using VM targets 0 and then one gives

$$\mathbb{E}[\text{Rev}] \geq \bar{\alpha} V_{\text{QL}}^* + U_{\text{QL}} \ln(1 - \bar{\alpha}), \quad \mathbb{E}[\text{Rev}] + V_{\text{VM}} \geq V_{\text{VM}}^* + \bar{\alpha} V_{\text{QL}}^* + U_{\text{QL}} \ln(1 - \bar{\alpha}).$$

Payments are nonnegative and expected ROI holds, so $P_{\text{VM}} \leq \min\{V_{\text{VM}}, \mathbb{E}[\text{Rev}]\}$ and

$$\mathbb{E}[\text{Wel}] = U_{\text{QL}} + \mathbb{E}[\text{Rev}] + V_{\text{VM}} - P_{\text{VM}} \geq U_{\text{QL}} + \max\{\mathbb{E}[\text{Rev}], V_{\text{VM}}\}.$$

These are the inequalities optimized in Equation (EC.8). The same calculation with total optimum $\mathbb{E}[\text{OPT}(\mathbf{v})]$ gives 0.4569, and with no QL contribution, half the sum gives 1/2. When advertisers are all-QL, the same argument in the proof of Proposition 4 gives $(1 - 1/e)\mathbb{E}[\text{OPT}(\mathbf{v})]$. An expected optimum of 0 is trivial, and a QL optimal contribution of 0 is handled as in EC.3.1.

No expectation is factored across advertisers, because a deviation keeps the actual joint distribution of the deviator's valuations and the other advertisers' valuations and multipliers, and the targets are chosen from

the averaged inequalities rather than separately for each valuation profile. Correlated valuations are therefore harmless.

Undominated VCG

By Lemma EC.1(iii), undominated multipliers—which are all at least one under VCG—yield $V^{\text{omit}} \leq \text{Rev}$, so Equation (EC.1) gives $\text{OPT}(\mathbf{v}) \leq \text{Wel} + \text{Rev}$. QL advertisers bid their valuations and have nonnegative utility, and expected ROI gives $\mathbb{E}[\text{Rev}] \leq \mathbb{E}[\text{Wel}]$. Averaging and rearranging gives the bound without any VM incentive comparison. When advertisers are all-QL, every advertiser bids her valuations, so welfare equals $\text{OPT}(\mathbf{v})$ at every realization. $\square$

EC.4.9. Nonexistence of CCE with a Floor and Ceiling

Before turning to repeated auctions, we show that with a floor and a ceiling, GFP and GSP may have no CCE at all.

PROPOSITION EC.6. *With equilibrium support $[A_{\min}, A_{\max}]$, with or without undominated multipliers, there exist a GFP instance and a GSP instance with no CCE, and hence with no CE or PNE.*

Proof. Both examples rely on deviations in a narrow range or outside the equilibrium support. In the GFP example, a QL advertiser gains from a bid in a narrow interval just above $A_{\max}$, and in the GSP example, QL advertisers gain from bids below the floor $A_{\min}$.

GFP. For GFP, a feasible randomized VM deviation forces a profitable QL deviation. Take two unit-CTR slots. Advertiser 1 is QL with one ad with valuation $v_{1,1} = 2A_{\max}^2$, and advertiser 2 is VM and has two ads with valuation one. Give advertiser 2 priority at ties. Suppose that σ is a CCE and set $\theta = \Pr_{\alpha \sim \sigma}(v_{1,1}\alpha_1 \leq A_{\max})$. Here, advertiser 2 deviates to $A_{\max}$ with probability $1/(2A_{\max} - 1)$ and 0 otherwise. At 0 she receives at least one slot for free, while at $A_{\max}$ she receives at most two slots. The expected payment deficit therefore satisfies

$$\mathbb{E}[P_2(\tilde{\alpha}_2, \alpha_1) - V_2(\tilde{\alpha}_2, \alpha_1)] \leq \frac{2(A_{\max} - 1)}{2A_{\max} - 1} - \left(1 - \frac{1}{2A_{\max} - 1}\right) = 0.$$

It always acquires value at least one and acquires a second slot whenever it selects $A_{\max}$ and $v_{1,1}\alpha_1 \leq A_{\max}$. CCE optimality thus gives $\mathbb{E}_{\alpha \sim \sigma}[V_2] \geq 1 + \theta/(2A_{\max} - 1)$. At equilibrium, advertiser 1 wins exactly when advertiser 2 receives one slot, so her winning probability is $2 - \mathbb{E}_{\alpha \sim \sigma}[V_2]$. On $\{v_{1,1}\alpha_1 > A_{\max}\}$ she wins and pays more than $A_{\max}$. Substituting the lower bound on $\mathbb{E}_{\alpha \sim \sigma}[V_2]$ yields

$$\mathbb{E}_{\alpha \sim \sigma}[u_1] \leq v_{1,1}(2 - \mathbb{E}_{\alpha \sim \sigma}[V_2]) - A_{\max}(1 - \theta) \leq v_{1,1} - A_{\max} - \frac{A_{\max}\theta}{2A_{\max} - 1}.$$

If $\theta > 0$, this is strictly below $v_{1,1} - A_{\max}$. If $\theta = 0$, advertiser 1 pays more than $A_{\max}$ almost surely, so the first inequality is strict. In either case, choose a bid $\tilde{b}_1$ with $A_{\max} < \tilde{b}_1 < v_{1,1} - \mathbb{E}_{\alpha \sim \sigma}[u_1]$. Using $\tilde{\alpha}_1 = \tilde{b}_1/v_{1,1}$ wins against every opposing α_2 and gives utility $v_{1,1} - \tilde{b}_1 > \mathbb{E}_{\alpha \sim \sigma}[u_1]$, contradicting CCE.

GSP. For GSP, each QL advertiser can obtain utility equal to her full valuation, although their utilities cannot attain these valuations simultaneously. Take three unit-CTR slots and three advertisers, each with one ad. Advertisers 1 and 2 are QL with valuations $v_{1,1} = 1$ and $v_{2,1} = 2/3$, and advertiser 3 is VM with valuation $v_{3,1} = 1/2$. Suppose a CCE σ exists. Each advertiser's ad has an equilibrium bid of at least $A_{\min}/2$. Either QL advertiser can therefore deviate to bid $A_{\min}/4$, receive the third slot, and pay 0. CCE optimality implies $\mathbb{E}_{\alpha \sim \sigma}[u_1] = 1$ and $\mathbb{E}_{\alpha \sim \sigma}[u_2] = 2/3$. At every equilibrium profile, the higher-ranked QL ad has the other QL bid below it, so

$$u_1(\alpha) + u_2(\alpha) \leq \frac{5}{3} - \min\{\alpha_1, \frac{2}{3}\alpha_2\} \leq \frac{5}{3} - \frac{2A_{\min}}{3} < \frac{5}{3}.$$

Taking expectations contradicts the two utility equalities. Thus neither GFP nor GSP admit a CCE. $\square$

EC.4.10. Proof of Theorem 3 and Corollary 2

Finally, we show that learning dynamics recover our CCE guarantees, with fixed or time-varying valuations.

THEOREM 3. *Assume all learners adopt primal–dual algorithms. Then the following hold.*

- $\text{PoA}_{\text{dyn}}(\text{GFP}) \geq 0.4569$. For all-VM: $\text{PoA}_{\text{dyn}}(\text{GFP}) \geq 1/2$ and for all-QL: $\text{PoA}_{\text{dyn}}(\text{GFP}) \geq 1 - 1/e$ (as in their static counterparts Theorem 1 and Proposition 4).
- Assuming undominated multipliers, $\text{PoA}_{\text{dyn}}(\text{VCG}) \geq 1/2$ (as in its static counterpart Proposition 6).

Moreover, on every instance, the ratio of expected average welfare over T rounds to OPT satisfies these bounds up to an additive error of $O(T^{-1/4})$.

We prove the learning guarantee in two steps. First, we identify how violations of the three conditions used in our CCE proofs affect welfare. We then show that the primal–dual updates make these violations small. We allow valuations to vary throughout, so the same argument also proves Corollary 2. We give the concluding step in EC.4.10.4.

EC.4.10.1. Setup and Error Terms. We fix a mechanism and the multipliers allowed in equilibrium (a floor and ceiling, with or without undominated multipliers, as in Table EC.4) and repeat the auction for T rounds, keeping slots, CTRs, ads, objectives, and tie-breaking unchanged. In round t , advertiser n has valuations $\mathbf{v}_{n,t}$ and multiplier $\alpha_{n,t}$. Valuations may be correlated across advertisers and rounds and may respond to past play, but are fixed before the current multiplier draws. Each advertiser independently draws her multiplier, conditional on the past and current valuations, from a grid $\mathcal{A}_n^L \subseteq \mathcal{A}_n$ with at most $L \geq 2$ points. She then observes the value and payment that every grid multiplier would have obtained against the realized opposing multipliers. Every grid contains one and, when allowed, $A_{\min}$. As in Section 5.3.2, we suppress the separate discretization error, which vanishes as the grids are refined. To formalize this, we give the following formal definition of the dynamic PoA over a multiplier grid.

DEFINITION EC.2 (DYNAMIC POA). Fix a mechanism $\mathcal{M}$. For each $L \geq 2$, fix grids $\mathcal{A}_n^L$ as above, such that the largest gap between consecutive grid points vanishes as $L \rightarrow \infty$. For an instance I with $\text{OPT}(I) > 0$, let $W_{T,L}(I)$ be the expected average welfare over T rounds when every advertiser runs the primal–dual algorithm of EC.4.10.3 on her grid, with the step sizes in (EC.16) and $\bar{\mu} = T^{1/4}$. The *dynamic PoA* of $\mathcal{M}$ is

$$\text{PoA}_{\text{dyn}}(\mathcal{M}) \doteq \inf_I \liminf_{L \rightarrow \infty} \liminf_{T \rightarrow \infty} \frac{W_{T,L}(I)}{\text{OPT}(I)}.$$

Taking $T \rightarrow \infty$ first removes the learning error, and refining the grids then removes the discretization error. Hence the welfare bounds of Theorem 3, which hold on every instance and for every T , give its corresponding bounds on PoA_{dyn} .

Now that we have described our notion of dynamic PoA formally, we can relate it to several important notions of regret, shortfall, and deficit over time. We use $\mathbb{E}$ for expectation under the empirical distribution that puts probability $1/T$ on each realized pair $(\mathbf{v}_t, \boldsymbol{\alpha}_t)$. All values, payments, and counterfactuals use the valuations of the same pair, and we suppress this argument when clear. Thus $\mathbb{E}[\text{Wel}]$ is average welfare and $\mathbb{E}[\text{OPT}(\mathbf{v})] = T^{-1} \sum_{t \leq T} \text{OPT}_t$ is the average realized optimum. We measure the three departures from the CCE conditions by the QL regret R_n for $n \in \mathcal{Q}$, and the VM value shortfall S_n and net ROI deficit ν_n for $n \in \mathcal{V}$, as defined in (10) (Section 5.3.2), with $\mathbb{E}$ the empirical average above and the maximum over the grid $\mathcal{A}_n^L$. The QL *regret* R_n is the average utility the advertiser could have gained by using her best fixed grid multiplier throughout. The VM value *shortfall* S_n is the average value she lost relative to always bidding her valuations. The VM *net ROI deficit* ν_n is the amount by which her average payment exceeds her average value. Because we take positive parts after averaging, an advertiser’s surplus in one round can offset her own deficit in another, but not another advertiser’s deficit.

We write R_{QL} , S_{VM} , and ν_{VM} for the sums over the respective populations. At a CCE with fixed valuations, all three errors are 0. Our welfare proofs need only these three conditions, however, not the full CCE requirement that a VM advertiser be optimal against every deviation satisfying ROI on average. We can therefore use the same welfare arguments whenever these errors are small, even without an equilibrium.

EC.4.10.2. Welfare Bounds with Learning Errors. We write β for the applicable CCE welfare fraction from Theorem 1, Propositions EC.2, 6, and EC.5, and Lemma EC.4, using the exact bounds in their proofs. For all-VM and all-QL GFP, we use Proposition 4. The next lemma shows that learning does not change these leading constants. We absorb only the coefficients of the learning errors into $O(\cdot)$.

LEMMA EC.5. *Under the corresponding restrictions on multipliers, every history satisfies*

$$\mathbb{E}[\text{Wel}] \geq \beta \mathbb{E}[\text{OPT}(\mathbf{v})] - O(R_{\text{QL}} + S_{\text{VM}} + \nu_{\text{VM}}),$$

with an absolute implied constant. For VCG with undominated multipliers, the error is only $O(\nu_{\text{VM}})$.

Proof. We apply the earlier welfare proofs with the following relaxed incentive and ROI inequalities:

$$\begin{aligned}\mathbb{E}[u_n(\alpha', \alpha_{-n})] &\leq \mathbb{E}[u_n] + R_n & (n \in Q, \alpha' \in \mathcal{A}_n^L), \\ \mathbb{E}[V_n(1, \alpha_{-n})] &\leq \mathbb{E}[V_n] + S_n, \quad \mathbb{E}[P_n] \leq \mathbb{E}[V_n] + v_n & (n \in \mathcal{V}).\end{aligned}$$

Bidding one gives nonnegative QL utility, so $\mathbb{E}[u_n] \geq -R_n$. Adding payments across advertisers therefore gives

$$\mathbb{E}[\text{Rev}] \leq \mathbb{E}[\text{Wel}] + R_{\text{QL}} + v_{\text{VM}}. \quad (\text{EC.14})$$

This is the only change to the earlier step that bounded revenue by welfare. The allocation inequalities themselves hold at every valuation–multiplier pair, so we may average them over rounds without assuming independence of valuations.

GFP. We use the notation of Section 4.1.1, with V_{QL}^* and V_{VM}^* now denoting average contributions to the realized optimum. We replace QL utility U_n by $U_n + R_n \geq 0$. The QL incentive inequality becomes $(1 - \alpha')\mathbb{E}[V_n(\alpha', \alpha_{-n})] \leq U_n + R_n$, while monotonicity bounds each VM integral up to one by $\mathbb{E}[V_n] + S_n$. Thus the two revenue bounds in the CCE proof hold with U_{QL} replaced by $U_{\text{QL}} + R_{\text{QL}}$ and V_{VM} replaced by $V_{\text{VM}} + S_{\text{VM}}$.⁵

We can now finish exactly as in the CCE proof, because adding the three errors to welfare gives

$$\mathbb{E}[\text{Wel}] + R_{\text{QL}} + S_{\text{VM}} + v_{\text{VM}} \geq U_{\text{QL}} + R_{\text{QL}} + \max\{\mathbb{E}[\text{Rev}], V_{\text{VM}} + S_{\text{VM}}\}.$$

The right side has the same form as before. The minimization in Equation (EC.8) therefore gives 0.4569, and the all-VM specialization gives 1/2. For the floor guarantee, we combine the unchanged revenue lower bound in Equation (EC.9) with (EC.14). For all-QL GFP, we make the same utility replacement in the proof of Proposition 4.

GSP and VCG. We follow the ceiling and floor arguments of Lemma EC.4. When $A_{\min} \leq 1/2$, its ceiling argument uses multiplier 1/2. A QL advertiser’s utility at this multiplier is at least half her acquired value, while a VM advertiser acquires no more value than by bidding one. The relaxed incentive inequalities therefore give

$$\mathbb{E}[\text{OPT}(\mathbf{v})] \leq (2A_{\max} + 2)\mathbb{E}[\text{Wel}] + 2R_{\text{QL}} + S_{\text{VM}}.$$

For the floor argument, the inequality $\text{OPT}(\mathbf{v}) \leq \text{Wel} + \text{Rev}/A_{\min}$ still holds at every pair. We average it and use (EC.14) to obtain the same floor guarantee up to the stated errors. Taking the better of these two bounds proves the claim. When $A_{\min} > 1/2$, the floor guarantee is already stronger, so we do not need the half-multiplier deviation.

⁵ We use the version of Lemma EC.1(i) whose integrals start at the realized multipliers, so targets below $A_{\min}$ contribute no integral. For a finite grid, we round each remaining QL target down on that advertiser’s grid and use monotone right-endpoint sums, which introduce only the vanishing discretization error suppressed above.

Multipliers of at least one. In the settings of Propositions 6 and EC.5, the earlier proofs use only the restrictions on multipliers and ROI, rather than optimality against deviations. We therefore retain their pointwise inequality $\text{OPT}(\mathbf{v}) \leq \text{Wel} + (1 - A_{\max}^{-1})\text{Rev}$. QL advertisers bid one, so $\mathbb{E}[\text{Rev}] \leq \mathbb{E}[\text{Wel}] + \nu_{\text{VM}}$. Averaging and rearranging gives the same fraction $\beta = A_{\max}/(2A_{\max} - 1)$ with an error at most ν_{VM} . When advertisers are all-QL and use undominated multipliers under VCG, everyone bids her valuations, so welfare is exactly optimal. $\square$

The preceding lemma does not rule out the low welfare allowed by the GSP and VCG bounds. In fact, we can make all three learning errors 0 simply by repeating one of the earlier low-welfare equilibria.

LEMMA EC.6. *Let the mechanism be GSP or VCG, and let the population be mixed or all-VM. For every $\delta > 0$ with $(1 + \delta) \max\{A_{\max}^{-1}, A_{\min}\} < 1$, there is a single-slot instance with fixed valuations and a history with $R_{\text{QL}} = S_{\text{VM}} = \nu_{\text{VM}} = 0$ and*

$$\mathbb{E}[\text{Wel}] = (1 + \delta) \max\{A_{\max}^{-1}, A_{\min}\} \text{OPT}.$$

The same holds for GSP with undominated multipliers.

Proof. We repeat the single-slot PNE of EC.3.3 in every round, with grids containing its multipliers. QL optimality and ROI feasibility give $R_{\text{QL}} = \nu_{\text{VM}} = 0$. Bidding one is a feasible VM deviation, so VM optimality also gives $S_{\text{VM}} = 0$. Repetition leaves the welfare ratio unchanged. $\square$

EC.4.10.3. The Primal–Dual Algorithms. We now specify learning rules that make the three errors vanish. We write $V_{n,t}(\alpha')$, $P_{n,t}(\alpha')$, and $u_{n,t}(\alpha')$ for the value, payment, and utility that grid multiplier α' would have obtained in round t , at the current valuations against the realized opposing multipliers. We omit the multiplier argument for realized outcomes, for example $V_{n,t} = V_{n,t}(\alpha_{n,t})$.

A QL advertiser learns directly from utility. We use *exponential weights*, under which she draws α' with probability proportional to $\exp\{\gamma \sum_{t' < t} u_{n,t'}(\alpha')\}$, with a step size $\gamma > 0$ given in Equation (EC.16) below. A VM advertiser instead seeks value while keeping payment covered by value, so we apply the same rule to the rewards

$$g_{n,t}(\alpha') \doteq V_{n,t}(\alpha') - \mu_{n,t} [P_{n,t}(\alpha') - V_{n,t}(\alpha')]. \quad (\text{EC.15})$$

The coefficient $\mu_{n,t}$ penalizes the ROI deficit. We start it at $\mu_{n,1} = 0$ and update it by

$$\mu_{n,t+1} \doteq \min\{\bar{\mu}, \max\{0, \mu_{n,t} + \gamma_{\mu}(P_{n,t} - V_{n,t})\}\}.$$

with a step size $\gamma_{\mu} > 0$, also given in Equation (EC.16). Thus the penalty rises when payment exceeds acquired value and falls when value exceeds payment, staying between 0 and the cap $\bar{\mu}$. It is an internal learning weight, not an additional payment. A larger penalty makes multipliers that overspend less attractive to the advertiser.

We assume a known bound $K_n > 0$ on advertiser n 's value and payment at every grid multiplier, uniformly over rounds and valuation realizations. For example, $A_{\max} \sum_j c_j v_{n,j,t} \leq K_n$ for every t suffices. We start all advertisers from uniform weights and choose step sizes up to fixed positive constant factors:

$$\gamma = \begin{cases} \Theta\left(\frac{\sqrt{\log L/T}}{(1+\bar{\mu})K_n}\right) & \text{for VM advertisers,} \\ \Theta\left(\frac{\sqrt{\log L/T}}{K_n}\right) & \text{for QL advertisers,} \end{cases} \quad \gamma_\mu = \Theta\left(\frac{\bar{\mu}}{K_n\sqrt{T}}\right). \quad (\text{EC.16})$$

For a VM advertiser n , we track the cumulative regrets of her multiplier and penalty updates:

$$R^{\text{pr}} \doteq \left[\max_{\alpha' \in \mathcal{A}_n^L} \sum_{t \leq T} (g_{n,t}(\alpha') - g_{n,t}) \right]_+, \quad R^{\text{du}} \doteq \left[\max_{\mu \in [0, \bar{\mu}]} \sum_{t \leq T} (\mu - \mu_{n,t})(P_{n,t} - V_{n,t}) \right]_+.$$

The first measures the gain from using one fixed multiplier throughout, evaluated at the penalties actually used. The second measures the gain from using one fixed penalty throughout, evaluated at the realized ROI deficits. We next show that small regret in these two updates controls both value shortfall and net overspending.

LEMMA EC.7. *For every history, each VM advertiser n satisfies*

$$S_n \leq \frac{R^{\text{pr}} + R^{\text{du}}}{T}, \quad \bar{\mu} v_n \leq \mathbb{E}[V_n] + \frac{R^{\text{pr}} + R^{\text{du}}}{T},$$

where R^{pr} and R^{du} are her regrets. With (EC.16) and $\bar{\mu} = T^{1/4}$, in expectation,

$$R_n = O\left(K_n \sqrt{\frac{\log L}{T}}\right) \quad (n \in \mathcal{Q}), \quad S_n, v_n = O\left(K_n T^{-1/4} \sqrt{\log L}\right) \quad (n \in \mathcal{V}).$$

Proof. We fix a VM advertiser and omit n . Bidding one satisfies ROI under every mechanism, so its penalized reward satisfies $g_t(1) \geq V_t(1) \geq 0$. The multiplier update's regret limits how far its total reward can fall below that from bidding one. The penalty update's regret limits the gain from replacing its penalties by any fixed $\mu \in [0, \bar{\mu}]$. Adding these two regret inequalities and substituting the definition of g_t gives

$$\sum_{t \leq T} [V_t(1) - V_t] + \mu \sum_{t \leq T} (P_t - V_t) \leq R^{\text{pr}} + R^{\text{du}}. \quad (\text{EC.17})$$

With $\mu = 0$, we bound the cumulative value shortfall relative to bidding one. With $\mu = \bar{\mu}$, we bound the cumulative ROI deficit after dropping the nonnegative term $\sum_{t \leq T} V_t(1)$. We divide by T and take positive parts to obtain the two stated inequalities.

It remains to bound the regrets. The primal rewards have range $O((1+\bar{\mu})K)$, and the penalty update takes projected gradient steps on an interval of length $\bar{\mu}$, with gradients bounded by K . Applying the standard exponential-weights and projected-gradient bounds (Cesa-Bianchi and Lugosi 2006, Zinkevich 2003) with (EC.16), we obtain

$$R^{\text{pr}} = O\left((1+\bar{\mu})K\sqrt{T \log L}\right) \quad \text{in expectation,} \quad R^{\text{du}} = O(\bar{\mu}K\sqrt{T}) \quad \text{on every history.}$$

To include the effect of drawing multipliers, we condition on the past, current valuations, and opposing draws. The counterfactual rewards are then fixed before the advertiser’s own draw, so realized reward differs from its conditional expectation by a martingale increment. Over T rounds, the expected positive part of the accumulated difference is $O((1 + \bar{\mu})K\sqrt{T})$, already within the displayed bound. QL utility has range $O(K)$, so we obtain the stated QL regret by the same argument.

A larger penalty cap improves the ROI bound but increases the reward range and hence the learning error. In particular, the two pathwise inequalities above give expected shortfall $O((1 + \bar{\mu})K\sqrt{\log L/T})$ and expected deficit $O(K/\bar{\mu} + (1 + \bar{\mu})K\sqrt{\log L/T}/\bar{\mu})$. We choose $\bar{\mu} = T^{1/4}$ so both vanish at the stated rate. We also obtain high-probability versions from Azuma–Hoeffding and the same orders without knowing T by restarting at doubling horizons. $\square$

EC.4.10.4. Completing the Proof of Theorem 3 and Corollary 2. We now substitute the learning-error bounds from Lemma EC.7 into the welfare inequality of Lemma EC.5. Writing Wel_t for the welfare in round t and taking expectation over the entire history, including the learners’ draws and valuations, gives

$$\mathbb{E} \left[\frac{1}{T} \sum_{t \leq T} \text{Wel}_t \right] \geq \beta \mathbb{E} \left[\frac{1}{T} \sum_{t \leq T} \text{OPT}_t \right] - O \left(T^{-1/4} \sqrt{\log L} \sum_n K_n \right).$$

This proves Corollary 2. With fixed valuations, $\text{OPT}_t = \text{OPT}$ in every round, so we also obtain Theorem 3. The main-text rate suppresses the fixed payoff bounds and grid-size factor. Finally, Lemma EC.6 supplies the low-welfare histories satisfying all three learning conditions exactly. $\square$

Remark. Our goal is to recover the CCE welfare bounds up to a vanishing learning error, rather than to obtain tight regret rates. We can theoretically achieve $T^{-1/2}$ regret rates with a more careful analysis. In particular, the algorithm of Sinha and Vaze (2024) applies to distributions over finite grid multipliers, because negative expected value and expected ROI deficit are linear with bounded coefficients, and bidding one satisfies every round’s ROI constraint. Their $O(\sqrt{T})$ cumulative regret against bidding one and $O(\sqrt{T} \log T)$ cumulative constraint-violation bounds, together with the sampling argument above, give expected shortfall $O(T^{-1/2})$ and net ROI deficit $O(T^{-1/2} \log T)$, with constants depending on the grid and payoff bound.

EC.5. Experimental Implementation and Complete Results

For clarity, the experiments study learning outcomes rather than estimate worst-case PoA. We first describe the sweep over the range of multipliers behind Figures 2, 3, and EC.1, then the remaining experiments, and finally the results with undominated multipliers.


EC.5.1. Varying the Range of Multipliers

We set $A_{\max} = 2^{k/3}$ for $k = 0, 1, \dots, 18$ and $A_{\min} = A_{\max}^{-1}$, so at $A_{\max} = 1$, every advertiser bids her valuations. Each market has one QL and one VM advertiser. With $M = 1$, we run GFP and GSP, which then coincides

with VCG. With $M = 2$, we run GFP, GSP, and VCG, and also VCG with undominated multipliers. Every slot has CTR one. In each trial, we fix the QL advertiser’s top valuation at 1 and draw the VM advertiser’s top valuation log-uniformly from $[1/128, 128]$. With $M = 2$, each advertiser’s second valuation is her top valuation times an independent draw from the uniform distribution on $[1/2, 1]$.

Each grid is the lattice $\{2^{j/6} : j \in \mathbb{Z}\}$ intersected with the advertiser’s interval in Table EC.4, so a grid on $[A_{\max}^{-1}, A_{\max}]$ has $4k + 1$ points. We use the learning procedure in EC.5.3 for $T = 10^6$ rounds, except that we also multiply the QL step size by 100, which keeps QL learning speed comparable across ranges. We run 100 trials per configuration and range.

Results. GFP’s 1st-percentile welfare stays above .84 with one slot and .80 with two. Under GSP and VCG, it falls roughly like $A_{\max}^{-1}$, to .04 at $A_{\max} = 64$, within a factor of about 2.5 of the worst case in Proposition EC.3. These losses sit in the lower tail, as median welfare exceeds .97 for every mechanism and range. VCG with undominated multipliers, which the figures omit, keeps its 1st-percentile welfare above .86, consistent with Proposition 6. GFP’s 1st-percentile revenue equals the floor $A_{\min}$ and is 93 to 123 times that of GSP and VCG. Figure EC.1 shows welfare and revenue with one and two slots.



Note. Percentiles of Wel/OPT (top) and Rev/OPT (bottom, on a logarithmic scale) across 100 trials. Lines mark the 1st percentile, and shading extends to the 5th for welfare and to the 25th for revenue. Panel (b) is Figure 2, and panels (c) and (d) are Figure 3. With one slot, GSP and VCG coincide.

EC.5.2. Instances and Multiplier Grids

In the remaining experiments, every instance has uniform CTRs, $c_1 = \dots = c_M = 1$, with $M \in \{1, 3, 5, 10\}$ and either $N = 2$ or $N = 4$ advertisers. A two-advertiser mixed population has one VM and one QL advertiser, and

a four-advertiser mixed population has two of each, with VM advertisers indexed first. We draw independent lognormal valuations with log-mean 0 and log-variance one. We divide all NM valuations by their common maximum and sort within each advertiser so that $v_{n,1} \geq \dots \geq v_{n,M}$. Bid ties favor the lower advertiser index and then the lower ad index. We set $A_{\min} = 0.1$, $A_{\max} = 10$, and $L = 11$. The grid $\mathcal{A}_n^L$ has L logarithmically spaced points, including both endpoints, in each nonsingleton interval of Table EC.4, and a single-point interval is its own grid. Thus, by default, QL and VM advertisers use $[A_{\min}, A_{\max}]$. With undominated multipliers, QL GFP and GSP advertisers use $[A_{\min}, 1]$, QL VCG advertisers use multiplier one, VM GFP and GSP advertisers use $[A_{\min}, A_{\max}]$, and VM VCG advertisers use $[1, A_{\max}]$. Every grid includes multiplier one. We note that narrowing an interval while retaining L points changes its logarithmic resolution, rather than merely deleting actions from a common grid. The learning procedures use these grids in both settings.

For the stochastic-valuation experiment in Section 6.5, we draw the normalized two-advertiser baseline once per trial. In every round, we multiply each original valuation by an independent uniform draw from $[0.75, 1.25]$, without compounding the shocks or renormalizing. We recompute bids, allocations, payments, and OPT_t from the current valuations, and all other instance and grid settings remain unchanged.

EC.5.3. Learning and Reporting

Each advertiser independently chooses an endpoint and places almost all initial probability there, with a small full-support regularization to keep every multiplier learnable.⁶ Advertisers with only one available multiplier always use it.

In each round, every advertiser observes the *counterfactual value and payment*—what each grid multiplier would have obtained against the realized opposing multipliers. QL advertisers run exponential weights on counterfactual utility, and VM advertisers run the primal–dual algorithm of EC.4.10.3 with cap $\bar{\mu} = T^{1/4}$, for $T = 10^6$ rounds. Unless stated otherwise below, we multiply the baseline VM primal step size by 100 for this finite horizon and keep the QL and dual step sizes unchanged.⁷ We run 200 trials per configuration and report all of them, without selecting or discarding any trial based on regret or outcomes.

In the stochastic-valuation experiment, both primal learners update every round, with the baseline QL and VM step sizes multiplied by 125 and 1250, respectively, and the dual step size and cap are unchanged.

In the swap-regret experiment, each advertiser instead runs the swap-regret reduction of Blum and Mansour (2007), which plays one learner per grid multiplier through the stationary distribution of the resulting stochastic matrix, with one ROI penalty shared by a VM advertiser’s learners. Policies update every 100 rounds from aggregate feedback, and we multiply the QL step sizes by 100 and the VM primal step sizes by 1000.

⁶ The theoretical algorithm uses uniform weights, but any fixed full-support starting distribution gives the same horizon rates, with constants depending on its smallest probability.

⁷ For reproducibility, the baseline step sizes used in the experiments are $\gamma = \sqrt{8 \log L/T} / [(1 + 2\bar{\mu})K_n]$ for VM advertisers, $\gamma = \sqrt{8 \log L/T} / (2K_n)$ for QL advertisers, and $\gamma_\mu = \bar{\mu} / (K_n \sqrt{T})$. These are particular choices within the orders in (EC.16).

Tables EC.5–EC.8 report the four experiments in the default setting, and Tables EC.9–EC.12 report them with undominated multipliers (EC.5.4). Each table reports the 1st, 5th, and 25th percentiles across the 200 trials, in that order. Each trial contributes the ratios $\sum_{t \leq T} \text{Wel}_t / \sum_{t \leq T} \text{OPT}_t$ and $\sum_{t \leq T} \text{Rev}_t / \sum_{t \leq T} \text{OPT}_t$, where Wel_t and Rev_t are the welfare and revenue in round t . For fixed valuations, these equal average welfare and revenue divided by OPT. For stochastic valuations, the optimum is recomputed each round, as in Corollary 2. When advertisers are all-VM, undominated multipliers do not restrict GFP or GSP (Table EC.4), so these rows are the same in both sets of tables.

EC.5.4. Undominated Multipliers

Tables EC.9–EC.12 report every experiment when advertisers use only undominated multipliers (Section 5.2.2). As in Section 6, the values below are 1st percentiles. For mixed and all-VM populations, GFP keeps high welfare and GSP low welfare. Across all four experiments, the smallest mixed welfare value under GFP is .80. In contrast, GSP stays low while VCG gains the most. With one slot and fixed valuations, mixed welfare under GSP and VCG is .14 in the default setting (Table EC.5), but .22 under GSP and .98 under VCG with undominated multipliers (Table EC.9), and all-VM welfare under VCG rises from .14 to .97. Across all four experiments, the smallest mixed and all-VM welfare values under VCG are .64 and .62, consistent with Proposition 6. With two advertisers, GFP’s revenue still exceeds that of VCG.

Table EC.5 $N = 2$, fixed valuations.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.65/ .74/ .98	.10/ .10/ .28
		3	.80/ .83/ .88	.14/ .29/ .44
		5	.84/ .86/ .90	.22/ .33/ .50
		10	.85/ .88/ .91	.33/ .38/ .54
	GSP	1	.14/ .27/ .98	.01/ .02/ .09
		3	.32/ .51/ .77	.02/ .05/ .12
		5	.49/ .63/ .78	.05/ .08/ .19
		10	.60/ .73/ .80	.07/ .12/ .37
	VCG	1	.14/ .27/ .98	.01/ .02/ .09
		3	.32/ .49/ .74	.02/ .03/ .07
		5	.49/ .60/ .78	.02/ .04/ .09
		10	.64/ .72/ .79	.03/ .05/ .32
All-VM	GFP	1	1/ 1/ 1	.10/ .13/ .45
		3	.96/ .97/ .98	.62/ .78/ .94
		5	.97/ .98/ .99	.90/ .92/ .96
		10	.98/ .98/ .99	.91/ .94/ .97
	GSP	1	.14/ .36/ 1	.01/ .02/ .09
		3	.47/ .65/ .82	.02/ .05/ .14
		5	.59/ .71/ .89	.04/ .08/ .61
		10	.67/ .84/ .96	.08/ .14/ .93
	VCG	1	.14/ .36/ 1	.01/ .02/ .09
		3	.44/ .55/ .77	.02/ .03/ .08
		5	.49/ .59/ .76	.02/ .05/ .17
		10	.67/ .74/ .87	.04/ .05/ .53
All-QL	GFP	1	.93/ .94/ .98	.10/ .10/ .24
		3	.92/ .94/ .98	.10/ .10/ .12
		5	.95/ .97/ .98	.10/ .11/ .13
		10	.97/ .98/ .99	.10/ .10/ .12
	GSP	1	.14/ .27/ .99	.01/ .02/ .09
		3	.44/ .61/ .87	.02/ .05/ .09
		5	.50/ .65/ .92	.04/ .07/ .11
		10	.71/ .80/ .95	.06/ .08/ .11
	VCG	1	.14/ .27/ .99	.01/ .02/ .09
		3	.32/ .49/ .81	.02/ .03/ .06
		5	.49/ .56/ .90	.02/ .04/ .08
		10	.67/ .78/ .99	.03/ .05/ .15

Table EC.6 $N = 4$, fixed valuations.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.64/ .71/ .98	.10/ .25/ .41
		3	.84/ .88/ .93	.23/ .34/ .57
		5	.85/ .88/ .92	.32/ .43/ .63
		10	.86/ .88/ .92	.41/ .53/ .65
	GSP	1	.12/ .22/ .75	.03/ .10/ .17
		3	.37/ .71/ .86	.09/ .21/ .46
		5	.55/ .78/ .87	.08/ .28/ .56
		10	.78/ .83/ .89	.38/ .52/ .64
	VCG	1	.12/ .22/ .75	.03/ .10/ .17
		3	.37/ .57/ .76	.07/ .10/ .33
		5	.42/ .68/ .81	.07/ .18/ .47
		10	.64/ .71/ .80	.23/ .39/ .57
All-VM	GFP	1	1/ 1/ 1	.18/ .39/ .63
		3	.96/ .97/ .98	.56/ .73/ .92
		5	.96/ .97/ .98	.79/ .92/ .96
		10	.98/ .98/ .99	.93/ .96/ .98
	GSP	1	.13/ .23/ .93	.03/ .10/ .25
		3	.64/ .75/ .88	.12/ .37/ .72
		5	.66/ .83/ .92	.09/ .55/ .86
		10	.87/ .93/ .96	.84/ .90/ .95
	VCG	1	.13/ .23/ .93	.03/ .10/ .25
		3	.37/ .57/ .76	.08/ .12/ .45
		5	.41/ .64/ .79	.07/ .16/ .60
		10	.61/ .72/ .83	.21/ .40/ .71
All-QL	GFP	1	.90/ .94/ .98	.10/ .16/ .28
		3	.87/ .92/ .96	.13/ .14/ .23
		5	.92/ .93/ .95	.14/ .17/ .24
		10	.92/ .95/ .97	.16/ .17/ .23
	GSP	1	.12/ .22/ .97	.02/ .10/ .13
		3	.33/ .85/ .96	.09/ .13/ .30
		5	.49/ .86/ .96	.08/ .13/ .29
		10	.72/ .93/ .97	.15/ .20/ .28
	VCG	1	.12/ .22/ .97	.02/ .10/ .13
		3	.28/ .72/ .99	.07/ .10/ .26
		5	.41/ .80/ .99	.06/ .10/ .33
		10	.72/ .97/ .99	.15/ .27/ .36

Table EC.7 $N = 2$, stochastic valuations.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.85/ .86/ .96	.10/ .10/ .26
		3	.81/ .87/ .92	.15/ .22/ .42
		5	.84/ .89/ .92	.17/ .25/ .50
		10	.87/ .88/ .92	.22/ .29/ .47
	GSP	1	.17/ .35/ .96	.00/ .01/ .06
		3	.50/ .63/ .80	.02/ .05/ .15
		5	.67/ .72/ .81	.06/ .08/ .33
		10	.71/ .76/ .82	.11/ .26/ .41
	VCG	1	.17/ .35/ .96	.00/ .01/ .06
		3	.46/ .61/ .77	.01/ .03/ .09
		5	.61/ .68/ .79	.03/ .04/ .24
		10	.68/ .73/ .80	.05/ .21/ .36
All-VM	GFP	1	1/ 1/ 1	.18/ .26/ .61
		3	.97/ .98/ .98	.67/ .84/ .94
		5	.98/ .98/ .99	.80/ .88/ .94
		10	.99/ .99/ .99	.79/ .85/ .93
	GSP	1	.17/ .35/ .97	.00/ .01/ .07
		3	.59/ .71/ .89	.02/ .04/ .19
		5	.72/ .82/ .91	.07/ .08/ .65
		10	.90/ .93/ .97	.65/ .77/ .90
	VCG	1	.17/ .35/ .97	.00/ .01/ .07
		3	.51/ .62/ .77	.01/ .03/ .09
		5	.63/ .67/ .82	.03/ .05/ .20
		10	.73/ .77/ .90	.14/ .33/ .62
All-QL	GFP	1	.91/ .92/ .98	.10/ .10/ .16
		3	.93/ .95/ .98	.10/ .10/ .11
		5	.97/ .98/ .99	.10/ .10/ .10
		10	.98/ .98/ .99	.10/ .10/ .10
	GSP	1	.17/ .34/ 1	.00/ .01/ .04
		3	.50/ .63/ .96	.02/ .04/ .08
		5	.61/ .72/ .97	.05/ .06/ .09
		10	.74/ .94/ .98	.06/ .08/ .12
	VCG	1	.17/ .34/ 1	.00/ .01/ .04
		3	.44/ .59/ .93	.01/ .03/ .07
		5	.57/ .71/ .98	.03/ .04/ .11
		10	.74/ .94/ .99	.05/ .10/ .16

Table EC.8 $N = 2$, fixed valuations, swap-regret.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.65/ .74/ 1	.10/ .10/ .37
		3	.81/ .84/ .89	.14/ .28/ .43
		5	.84/ .87/ .91	.20/ .32/ .51
		10	.85/ .88/ .91	.34/ .39/ .54
	GSP	1	.14/ .27/ 1	.01/ .02/ .10
		3	.32/ .52/ .78	.02/ .06/ .18
		5	.50/ .64/ .78	.06/ .10/ .32
		10	.67/ .74/ .81	.08/ .15/ .42
	VCG	1	.14/ .27/ 1	.01/ .02/ .10
		3	.32/ .51/ .75	.02/ .03/ .09
		5	.50/ .64/ .77	.03/ .05/ .25
		10	.64/ .71/ .78	.04/ .06/ .36
All-VM	GFP	1	1/ 1/ 1	.10/ .10/ .40
		3	.96/ .97/ .98	.57/ .77/ .94
		5	.97/ .98/ .99	.91/ .95/ .97
		10	.98/ .98/ .99	.94/ .96/ .98
	GSP	1	.14/ .36/ 1	.01/ .02/ .09
		3	.47/ .70/ .87	.02/ .05/ .33
		5	.61/ .74/ .92	.04/ .08/ .66
		10	.75/ .86/ .96	.08/ .14/ .91
	VCG	1	.14/ .36/ 1	.01/ .02/ .09
		3	.47/ .60/ .78	.02/ .03/ .08
		5	.54/ .60/ .77	.02/ .05/ .14
		10	.67/ .74/ .86	.04/ .05/ .47
All-QL	GFP	1	.91/ .93/ 1	.10/ .10/ .25
		3	.92/ .96/ .98	.10/ .10/ .11
		5	.95/ .97/ .99	.10/ .10/ .12
		10	.96/ .98/ .99	.10/ .10/ .12
	GSP	1	.14/ .27/ 1	.01/ .02/ .09
		3	.47/ .62/ .91	.02/ .05/ .09
		5	.54/ .71/ .97	.05/ .07/ .10
		10	.74/ .85/ .98	.06/ .08/ .11
	VCG	1	.14/ .27/ 1	.01/ .02/ .09
		3	.32/ .49/ .81	.02/ .03/ .06
		5	.49/ .56/ .90	.02/ .04/ .08
		10	.67/ .78/ 1	.03/ .05/ .15

Table EC.9 $N = 2$, fixed valuations, undominated multipliers.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.85/ .96/ 1	.10/ .10/ .32
		3	.80/ .86/ .90	.13/ .28/ .44
		5	.86/ .87/ .92	.23/ .33/ .49
		10	.85/ .88/ .92	.32/ .38/ .54
	GSP	1	.22/ .41/ 1	.01/ .02/ .08
		3	.32/ .52/ .76	.04/ .06/ .18
		5	.50/ .63/ .77	.06/ .10/ .33
		10	.67/ .72/ .80	.08/ .15/ .44
	VCG	1	.98/ 1/ 1	.05/ .08/ .30
		3	.64/ .69/ .79	.08/ .15/ .33
		5	.64/ .67/ .78	.16/ .24/ .36
		10	.67/ .71/ .78	.20/ .28/ .40
All-VM	GFP	1	1/ 1/ 1	.10/ .13/ .45
		3	.96/ .97/ .98	.62/ .78/ .94
		5	.97/ .98/ .99	.90/ .92/ .96
		10	.98/ .98/ .99	.91/ .94/ .97
	GSP	1	.14/ .36/ 1	.01/ .02/ .09
		3	.47/ .65/ .82	.02/ .05/ .14
		5	.59/ .71/ .89	.04/ .08/ .61
		10	.67/ .84/ .96	.08/ .14/ .93
	VCG	1	.97/ .99/ 1	.05/ .14/ .41
		3	.64/ .71/ .80	.16/ .23/ .37
		5	.62/ .67/ .80	.18/ .29/ .50
		10	.67/ .74/ .86	.33/ .42/ .58
All-QL	GFP	1	1/ 1/ 1	.10/ .10/ .24
		3	.91/ .94/ .98	.10/ .10/ .11
		5	.96/ .97/ .99	.10/ .10/ .12
		10	.98/ .99/ .99	.10/ .10/ .12
	GSP	1	1/ 1/ 1	.01/ .02/ .08
		3	.94/ .97/ .98	.04/ .05/ .11
		5	.95/ .97/ .98	.05/ .08/ .12
		10	.97/ .98/ .99	.07/ .08/ .12
	VCG	1	1/ 1/ 1	.03/ .06/ .23
		3	1/ 1/ 1	.05/ .10/ .18
		5	1/ 1/ 1	.09/ .12/ .17
		10	1/ 1/ 1	.09/ .13/ .17

Table EC.10 $N = 4$, fixed valuations, undominated multipliers.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.81/ .94/ 1	.10/ .20/ .50
		3	.89/ .91/ .95	.22/ .33/ .56
		5	.88/ .90/ .94	.32/ .41/ .61
		10	.88/ .90/ .94	.43/ .51/ .65
	GSP	1	.13/ .23/ .99	.03/ .07/ .20
		3	.56/ .74/ .86	.15/ .30/ .51
		5	.66/ .78/ .88	.12/ .38/ .57
		10	.79/ .85/ .90	.41/ .52/ .65
	VCG	1	.98/ 1/ 1	.07/ .22/ .51
		3	.67/ .73/ .85	.22/ .30/ .50
		5	.72/ .75/ .82	.30/ .39/ .57
		10	.71/ .75/ .83	.38/ .48/ .62
All-VM	GFP	1	1/ 1/ 1	.18/ .39/ .63
		3	.96/ .97/ .98	.56/ .73/ .92
		5	.96/ .97/ .98	.79/ .92/ .96
		10	.98/ .98/ .99	.93/ .96/ .98
	GSP	1	.13/ .23/ .93	.03/ .10/ .25
		3	.64/ .75/ .88	.12/ .37/ .72
		5	.66/ .83/ .92	.09/ .55/ .86
		10	.87/ .93/ .96	.84/ .90/ .95
	VCG	1	.99/ 1/ 1	.20/ .39/ .63
		3	.67/ .73/ .81	.29/ .44/ .65
		5	.70/ .75/ .82	.37/ .50/ .74
		10	.73/ .77/ .84	.58/ .65/ .79
All-QL	GFP	1	.99/ 1/ 1	.10/ .16/ .32
		3	.88/ .93/ .97	.12/ .16/ .25
		5	.92/ .94/ .97	.13/ .16/ .25
		10	.94/ .96/ .98	.15/ .17/ .23
	GSP	1	1/ 1/ 1	.02/ .06/ .21
		3	.95/ .96/ .99	.13/ .17/ .31
		5	.96/ .97/ .98	.13/ .21/ .30
		10	.96/ .97/ .98	.18/ .21/ .28
	VCG	1	1/ 1/ 1	.06/ .17/ .35
		3	1/ 1/ 1	.14/ .20/ .35
		5	1/ 1/ 1	.17/ .23/ .39
		10	1/ 1/ 1	.26/ .29/ .38

Table EC.11 $N = 2$, stochastic valuations, undominated multipliers.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.88/ .91/ .96	.10/ .12/ .26
		3	.81/ .87/ .92	.15/ .22/ .42
		5	.86/ .89/ .92	.17/ .26/ .50
		10	.88/ .89/ .92	.22/ .29/ .48
	GSP	1	.27/ .49/ .99	.00/ .01/ .05
		3	.50/ .67/ .80	.02/ .05/ .18
		5	.67/ .72/ .81	.06/ .09/ .33
		10	.74/ .78/ .83	.11/ .26/ .39
	VCG	1	.93/ .96/ 1	.02/ .05/ .17
		3	.66/ .69/ .78	.09/ .14/ .30
		5	.66/ .69/ .79	.10/ .19/ .35
		10	.68/ .73/ .81	.17/ .26/ .38
All-VM	GFP	1	1/ 1/ 1	.18/ .26/ .61
		3	.97/ .98/ .98	.67/ .84/ .94
		5	.98/ .98/ .99	.80/ .88/ .94
		10	.99/ .99/ .99	.79/ .85/ .93
	GSP	1	.17/ .35/ .97	.00/ .01/ .07
		3	.59/ .71/ .89	.02/ .04/ .19
		5	.72/ .82/ .91	.07/ .08/ .65
		10	.90/ .93/ .97	.65/ .77/ .90
	VCG	1	.97/ .97/ 1	.03/ .05/ .27
		3	.68/ .72/ .81	.11/ .20/ .35
		5	.68/ .72/ .84	.16/ .29/ .43
		10	.75/ .80/ .90	.37/ .45/ .67
All-QL	GFP	1	.92/ .94/ .97	.10/ .10/ .20
		3	.92/ .94/ .98	.10/ .10/ .11
		5	.97/ .98/ .99	.10/ .10/ .11
		10	.98/ .99/ .99	.10/ .10/ .11
	GSP	1	1/ 1/ 1	.00/ .01/ .04
		3	.95/ .96/ .98	.02/ .05/ .10
		5	.95/ .97/ .98	.06/ .07/ .12
		10	.97/ .98/ .99	.07/ .08/ .12
	VCG	1	1/ 1/ 1	.02/ .05/ .15
		3	1/ 1/ 1	.06/ .09/ .17
		5	1/ 1/ 1	.07/ .10/ .16
		10	1/ 1/ 1	.09/ .12/ .17

Table EC.12 $N = 2$, fixed valuations, swap-regret, undominated multipliers.

Pop.	Mechanism	M	Welfare	Revenue
Mixed	GFP	1	.85/ .96/ 1	.10/ .10/ .32
		3	.81/ .86/ .89	.15/ .27/ .43
		5	.86/ .87/ .92	.22/ .32/ .50
		10	.85/ .89/ .92	.33/ .39/ .54
	GSP	1	.22/ .57/ 1	.01/ .02/ .09
		3	.32/ .52/ .78	.03/ .06/ .18
		5	.50/ .64/ .78	.06/ .10/ .31
		10	.67/ .75/ .82	.08/ .15/ .41
	VCG	1	.98/ 1/ 1	.05/ .08/ .30
		3	.64/ .71/ .80	.08/ .15/ .33
		5	.64/ .67/ .79	.16/ .24/ .36
		10	.67/ .71/ .78	.20/ .28/ .40
All-VM	GFP	1	1/ 1/ 1	.10/ .10/ .40
		3	.96/ .97/ .98	.57/ .77/ .94
		5	.97/ .98/ .99	.91/ .95/ .97
		10	.98/ .98/ .99	.94/ .96/ .98
	GSP	1	.14/ .36/ 1	.01/ .02/ .09
		3	.47/ .70/ .87	.02/ .05/ .33
		5	.61/ .74/ .92	.04/ .08/ .66
		10	.75/ .86/ .96	.08/ .14/ .91
	VCG	1	.97/ .99/ 1	.05/ .14/ .41
		3	.62/ .71/ .80	.16/ .24/ .37
		5	.64/ .69/ .80	.20/ .29/ .52
		10	.67/ .74/ .86	.33/ .42/ .58
All-QL	GFP	1	1/ 1/ 1	.10/ .10/ .25
		3	.91/ .94/ .98	.10/ .10/ .11
		5	.96/ .97/ .99	.10/ .10/ .12
		10	.98/ .99/ .99	.10/ .10/ .12
	GSP	1	1/ 1/ 1	.01/ .02/ .08
		3	.94/ .97/ .98	.04/ .06/ .10
		5	.95/ .97/ .98	.05/ .08/ .11
		10	.97/ .98/ .99	.07/ .08/ .12
	VCG	1	1/ 1/ 1	.03/ .06/ .23
		3	1/ 1/ 1	.05/ .10/ .18
		5	1/ 1/ 1	.09/ .12/ .17
		10	1/ 1/ 1	.09/ .13/ .17