PRICE OF ANARCHY OF FIRST-PRICE AUCTIONS IN AGENTIC MARKETS

Anonymous authors
Paper under double-blind review

ABSTRACT

An agentic marketplace is a decentralized, serving-time market where AI agents bid to claim each user query. An agent's bid is based on its expected value for the query, which is the expected reward for producing a high-quality response, less the agent's expected inference costs. First-price auctions (FPAs), in which the highest bidder wins and pays its bid, are a natural mechanism for this, since they are simple and transparent to run and require no central knowledge of private capabilities and costs. But a potential weakness is that FPAs are not truthful. Participants may bid less than their true value in the hope of a larger profit, so the query can be allocated to someone other than the bidder who values it most highly. This welfare loss is quantified by the *price-of-anarchy*. In this paper, however, we argue that agentic markets are likely to possess special structural features that reduce it. The first is *value concentration*. If the gap between the top two values is usually small (e.g., because the best models are often distilled into cheaper alternatives which perform similarly), then the welfare gap can go to zero. The second is *weak dependence* between values after conditioning on the prompt, in which case the efficiency guarantee approaches $1 - 1/e^2$. In both cases, we derive bounds on the price of anarchy that are significantly better than the well-known factor $1 - 1/e$ from the general case, and validate them empirically on prompt routing in the EmbedLLM dataset.

1 INTRODUCTION

Consider a user with a task that could be serviced by any one of many AI models, where each model has a different cost and response quality. She could choose among them directly, but that requires knowing something about every option on the market. Alternatively, aggregators and routers can make this choice on her behalf—e.g., by establishing a central view of which model is good at what tasks, and directing user queries accordingly (Chen et al., 2024; Ong et al., 2025; Aubakirova et al., 2025). That view has to be measured and maintained, however, which becomes harder as the candidate pool grows beyond well-known providers. A third, and emerging, approach is the "agentic marketplace," which is a serving-time market where possible responders bid to claim each query, and are rewarded according to the quality of the response they provide (Qi et al., 2026; Chang et al., 2020; Liu et al., 2026). Each responder is assumed to privately know (or estimate) its expected quality and cost on the current query, which the market must elicit efficiently and transparently, without assuming that responders are honest. Being decentralized, the market also scales with the number of participants rather than with the platform's ability to track them.

The main design choice in such a marketplace is how queries are allocated (based on the set of bids), and what price the chosen responder ends up paying to claim the query. There are many possible mechanisms, and recent work on agentic markets has largely favored those that are dominant-strategy incentive compatible, such as second-price and VCG variants, under which bidding one's true value is optimal regardless of what other bidders do (Chang et al., 2020; Liu et al., 2026; Cao et al., 2026).

In this paper, we analyze the viability of the first-price auction (FPA), which is the simplest of these mechanisms. Here, each responder bids for the right to service the current query. The highest bidder wins, pays its bid to the platform, submits its response, and collects a reward based on quality (Qi et al., 2026; Chang et al., 2020). FPAs have several advantages. The payment rule is transparent, since the winner pays its own bid, with no need to trust the platform to truthfully report a second price (Rothkopf et al., 1990). Bidders can be strategic, but do not need to be: a nonnegative bid

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

below expected quality minus cost leaves a nonnegative expected surplus when it wins, regardless of what others bid. FPAs are also simple to implement, which is part of why some of the largest markets, such as display advertising auctions, use them (Cox, 2019).

Nevertheless, FPAs are traditionally less preferred in theory than VCG-style alternatives since they are not truthful. Strategic agents shade their bids below their true values to increase profit, and if higher-valued agents shade too aggressively, they can lose to lower-valued agents, which decreases *social welfare* (the total value generated by the allocation). The ratio of the welfare at a strategic equilibrium to the best achievable welfare is called the *price of anarchy* (Roughgarden et al., 2017), and for FPAs it can be as low as (but not worse!) than $1 - 1/e \approx 63\%$ (Syrgkanis & Tardos, 2013).

A potential $1/e \approx 37\%$ loss in efficiency may seem severe, but it is a worst case bound over arbitrary value distributions. Here, we empirically show that in representative agentic markets, simulated on the large-scale EmbedLLM routing benchmark (Zhuang et al., 2025), realized welfare losses are often much smaller. To understand this gap, we derive new bounds showing that the worst case improves substantially under structural conditions natural in agentic markets. The primary condition is *value concentration*: when the gap between the top values is usually small, so is the welfare loss, as intuition suggests. Value concentration arises when agents have comparable substitutes, e.g., due to distillation or similar training data (Cao et al., 2026; Aubakirova et al., 2025). We derive bounds that formalize this relationship, depending on how many bidders are competitive, how close they are to the top, and how reliably this holds. As a secondary result, we show that the bound can also improve through *near-independence*, even when values are not concentrated. When bidder values are independent, Jin & Lu (2023) show that the price of anarchy improves to $1 - 1/e^2 \approx 86.5\%$. We prove that this guarantee degrades smoothly under small perturbations from independence in terms of total-variation distance. This is practically relevant because shared query features (such as prompt difficulty or type) account for much of the co-movement across models; the idea being that once this public information is conditioned out, the residual dependence may be much weaker.

Taken together, these results explicitly show that when models compete closely on capability or share task difficulty that can be conditioned out, the equilibrium allocation preserves substantially more welfare than worst-case theory predicts. The upshot for market design is that much of the motivation for more complex mechanisms can then fall away—the simple and transparent FPA can already capture most of the available welfare under conditions that agentic markets tend to satisfy. A platform also does not have to treat these conditions as purely exogenous. For example, in retail and platform economics, operators often stock "private labels" alongside third-party sellers to discipline markups (Shopova, 2023; Ding et al., 2025). In the same spirit, we examine how a platform can actively steer market conditions, for instance by introducing its own in-house model to increase competition (and act like a reserve price), or by filtering out low-performing responders before bidding begins.

Contributions of this work. This paper provides theoretical and empirical tools to analyze first-price auctions under what practitioners may expect from typical agentic market structure:

1. **Value concentration (§3):** We derive price-of-anarchy bounds that improve when multiple bidders have valuations close to the top. In the limit of close competition, the welfare loss vanishes entirely.
2. **Near-independence (§4):** We prove that the $1 - 1/e^2 \approx 0.865$ bound by Jin & Lu (2023) for independent values is stable under small perturbations, which also allows us to derive stronger welfare guarantees if values are only weakly dependent, e.g., after conditioning on the query.
3. **Platform design (§5):** We show that platforms can steer these conditions by combining routing with auctions to filter weak responders, or by injecting competition through in-house models.
4. **Empirical analysis (§6):** We run thousands of first-price auction simulations on the EmbedLLM dataset (Zhuang et al., 2025), and empirically show that value concentration and conditional independence naturally emerge, and drive high realized efficiency, even in adversarial settings.

2 Preliminaries

We begin by briefly reviewing the first-price auction (FPA). Consider a single user query that is to be serviced by one of n bidders (also called "providers" or "agents") for some user-provided reward.[1]

[1]Where do these rewards come from? The theory that follows is agnostic to this question, but a few concrete instantiations can be helpful to keep in mind. The simplest is, e.g., thumbs up / down reported back by the platform

Each bidder i has a value v_i for the request, defined as its expected reward minus its cost of service:

$$v_i \;=\; \mathbb{E}[\text{reward}(\text{response}_i) - \text{cost}_i \mid \text{query}]. \tag{1}$$

Note that in a standard auction, the item for sale is the only source of randomness (each bidder knows its own value, which is random only from the perspective of others), whereas here, the reward and cost of a response is typically not known until it is produced. Simply letting bidder i take the expectation of the response reward minus cost reduces the setting to the standard one. If bidder i wins the request and pays a price p_i, its expected payoff (or *surplus*) is then simply $v_i - p_i$.

In a sealed-bid FPA, every bidder i submits a bid b_i. The bidder with the highest bid wins, and pays its own bid. We write $i^*(\mathbf{b})$ for the index of the winning bidder where $\mathbf{b} = (b_1, \ldots, b_n)$, with ties broken by an announced measurable bid-only rule, possibly randomized. A rational bidder, however, may prefer not to bid its true value, as doing so guarantees zero surplus. Bidders may therefore strategically *shade* their bids below their values to gain positive surplus. We model this strategic interaction via Bayesian games of incomplete information (Roughgarden et al., 2017). Specifically, we model bidder values $\mathbf{v} = (v_1, \ldots, v_n)$ as draws from a joint value prior μ, and a strategy for bidder i is a possibly randomized function s_i mapping a value v_i to a bid $b_i = s_i(v_i)$. A strategy profile $\mathbf{s} = (s_1, \ldots, s_n)$ is a *Bayes-Nash equilibrium* (BNE) if, for every bidder i and almost every value v_i, the bid $s_i(v_i)$ maximizes bidder i's expected surplus given the strategies of the other bidders and the prior μ.

On a technical level, we assume throughout that bidders never bid more than their value at equilibrium, i.e., $0 \le s_i(v_i) \le v_i$ almost surely for each $i \in [n]$ and almost every $v_i \ge 0$. This is standard in the literature, since overbidding risks negative surplus (Syrgkanis & Tardos, 2013; Jin & Lu, 2023).

2.1 Social welfare and the price of anarchy

It is natural to measure the performance of an auction via the total *social welfare* that it generates, which is the sum of the total utilities of all the bidders plus the revenue of the auctioneer.[2] In our single-item setting, only one response is selected, so the social welfare simply reduces to the value of the winning bidder, $W(\mathbf{b}, \mathbf{v}) := v_{i^*(\mathbf{b})}$, whereas the *optimal welfare* is defined as $\mathrm{OPT}(\mathbf{v}) := \max_i \; v_i$.[3]

Since FPAs allocate to the highest *bidder* rather than the highest-value provider, shading can cause the auction to select a suboptimal provider. The *price of anarchy* (PoA) quantifies the worst-case welfare loss due to strategic behavior. Formally, the PoA is the infimum, over all settings (number of bidders, joint value distribution μ) and BNE, of the ratio the expected welfare to the optimal welfare, i.e.,

$$\text{PoA} := \inf_{\mu} \inf_{\mathbf{s} \in \mathrm{BNE}(\mu)} \frac{\mathbb{E}_\mu[W(\mathbf{s}(\mathbf{V}), \mathbf{V})]}{\mathbb{E}_\mu[\mathrm{OPT}(\mathbf{V})]}. \tag{2}$$

The infimum is restricted to priors with $0 < \mathbb{E}_\mu[\mathrm{OPT}(\mathbf{V})] < \infty$ and a nonempty set of equilibria. In symmetric settings (where all bidders draw values from the same distribution), the BNE of a first-price auction is unique and fully efficient (Chawla & Hartline, 2013), and the PoA is 1. In asymmetric settings (which we *do* expect in agentic markets), however, equilibria are generally inefficient.

A foundational result establishes that every BNE of a first-price single-item auction achieves $\ge (1 - 1/e) \approx 63.2\%$ of optimal welfare (Syrgkanis & Tardos, 2013; Roughgarden et al., 2017). Under the stronger assumption that bidders' values are drawn independently (but not identically), Jin & Lu (2023) obtain the tight bound of $(1 - 1/e^2) \approx 86.5\%$. These bounds are important, though they do paint a pessimistic picture of first-price auctions. A potential welfare loss of 36% (or even 13.5%) is large enough to make a platform hesitant to deploy first-price orchestration at all. These bounds, however, are worst-case lower bounds over all possible value distributions. As we argue in this paper, agentic marketplaces often possess structural features that push their realized welfare higher.

3 First-price auctions with concentrated values

The price of anarchy fundamentally depends on what is at stake. If all bidders value an item identically, any allocation achieves full efficiency. More generally, the welfare of any allocation is at least the min-

or user (assumed, honestly). More generally, the platform could use any agreed-upon function $\phi(x, y)$ that takes a query x and response y and outputs a reward funded by the user (e.g., a reward in escrow is released contingent upon the agent's response y passing a collection of unit tests or clearing a third-party autorater / verifier).

[2] We model a user's value as being equal to the price they pay for the response, so that their utility is always 0.

[3] We do not consider externalities, such as i^* getting additional private utility by placing an advertisement.

imum bidder value. Write $V_{(1)} \geq \cdots \geq V_{(n)}$ for the bidder values in decreasing order. Then for any μ,

$$\inf_{\mathbf{s}\in \mathrm{BNE}(\mu)} \frac{\mathbb{E}_\mu[W(\mathbf{s}(\mathbf{V}), \mathbf{V})]}{\mathbb{E}_\mu[\mathrm{OPT}(\mathbf{V})]} \geq \frac{\mathbb{E}_\mu[V_{(n)}]}{\mathbb{E}_\mu[V_{(1)}]}, \tag{3}$$

which requires no strategic analysis at all. When values are concentrated—e.g., if all participating models perform similarly on user queries—this ratio is already close to 1. The ratio depends on the lowest value $V_{(n)}$, however, so a single weak bidder can pull it far below 1. The rest of this section uses the strategic structure of the auction so that only the strongest few bidders need to be close to the top.

3.1 Deterministic competitive floor

Intuitively, markets only require competition at the top—i.e., between the bidders who matter—to be efficient. As a warmup, suppose at least two bidders *always* have values exceeding some fixed floor. For convenience, we write this floor as a fraction $\alpha \in [0, 1]$ of the expected optimal value, $\alpha\mathbb{E}_\mu[V_{(1)}]$. This floor forces competition that drives welfare (and thus the price of anarchy) linearly upward in α.

Proposition 1. *Suppose* $0 < \mathbb{E}_\mu[V_{(1)}] < \infty$ *and* $V_{(2)} \geq \alpha\mathbb{E}_\mu[V_{(1)}] \geq 0$ *almost surely. Then*

$$\inf_{\mathbf{s}\in \mathrm{BNE}(\mu)} \frac{\mathbb{E}_\mu[W(\mathbf{s}(\mathbf{V}), \mathbf{V})]}{\mathbb{E}_\mu[\mathrm{OPT}(\mathbf{V})]} \geq 1 - \frac{1-\alpha}{e}. \tag{4}$$

All proofs are in the appendix. The main insight is that when at least two bidders value every item above a known floor v_{floor}, competition forces the winning bid in any equilibrium up to at least v_{floor} (Lemma D.1); then, on top of this guaranteed revenue, at most a $1/e$-fraction of the top value above the floor, $V_{(1)} - v_{\text{floor}}$, is lost (Lemma D.2). As a concrete example, if a rival always values the query at least half that of the expected top value ($\alpha = 0.50$), the bound rises from $1 - \frac{1}{e} \approx 63\%$ to $\approx 82\%$.

3.2 Probabilistic competitive floors

The deterministic floor above requires at least two bidders to clear v_{floor} on every query. This is, of course, not realistic: unusual queries may occasionally leave only one qualified model, even if multiple frontier models compete closely on the vast majority of requests. We now relax to probabilistic guarantees and show that *market depth* (the number of competitors) plays an important role.

The intuition behind Proposition 1 is that when the runner-up always values the query at least some threshold amount v_{floor}, competition pushes the winning bid up, and with it, welfare. When a close rival is only present sometimes, however, bidders can shade more aggressively and the winning bid may fall. Let $R = \max_i s_i(v_i)$ be the winning bid. We call the expected gap $\mathbb{E}_\mu[(v_{\text{floor}} - R)_+]$ for some fixed v_{floor} and R the *bid shortfall* with respect to v_{floor}. Our first result bounds this shortfall.

Lemma 2. *Fix* $v_{\text{floor}} \geq 0$ *and* $m \in \{2, \ldots, n\}$. *Suppose* $\mathbb{E}_\mu[V_{(1)}] < \infty$ *and* $\mathbb{P}_\mu(V_{(m)} \leq v_{\text{floor}}) \leq \epsilon$, *for some* $\epsilon \in [0, 1]$. *Then every* $\mathbf{s} \in \mathrm{BNE}(\mu)$ *satisfies* $\mathbb{E}_\mu[(v_{\text{floor}} - R)_+] \leq v_{\text{floor}}\mathcal{R}_m(\epsilon)$ *where*

$$\mathcal{R}_m(\epsilon) := \inf_{p > -\frac{m-2}{m-1}} \left[\epsilon \frac{m-1}{m - \frac{p+2}{p+1}}\right]^{\frac{1}{p+2}}. \tag{5}$$

Remark 3. *The somewhat involved form of* $\mathcal{R}_m(\epsilon)$ *is primarily an artifact of the proof, in which we maximize welfare over a family of polynomially-weighted unilateral deviation bids; analyzing the welfare of these deviation bids allows us to derive bounds for the true equilibria bids. For some additional intuition,* $\mathcal{R}_m(0) = 0$ *and* $\mathcal{R}_m(1) = 1$. *Choosing* $p = 0$ *for* $m > 2$ *gives* $\mathcal{R}_m(\epsilon) \leq \sqrt{\frac{m-1}{m-2}}\,\epsilon$, *while choosing* $p = \frac{2}{\log(1/\epsilon)}$ *for* $m = 2$ *and* $0 < \epsilon < 1$ *gives* $\mathcal{R}_2(\epsilon) \leq \sqrt{e\epsilon(1 + \frac{1}{2}\log(1/\epsilon))}$.

Combining the bound on the bid shortfall with the results from Proposition 1 gives the following bound. For convenience, we again write the v_{floor} as a fraction $\alpha \in [0, 1]$ of the expected optimal value.

Theorem 4. *Suppose* $0 < \mathbb{E}_\mu[V_{(1)}] < \infty$. *Then for any* $\alpha \in [0, 1]$, $m \in \{2, \ldots, n\}$, *and* $\epsilon \in [0, 1]$ *such that* $\mathbb{P}_\mu(V_{(m)} \leq \alpha\mathbb{E}_\mu[V_{(1)}]) \leq \epsilon$, *we have*

$$\inf_{\mathbf{s}\in \mathrm{BNE}(\mu)} \frac{\mathbb{E}_\mu[W(\mathbf{s}(\mathbf{V}), \mathbf{V})]}{\mathbb{E}_\mu[\mathrm{OPT}(\mathbf{V})]} \geq 1 - \frac{1 - \alpha[1 - e\mathcal{R}_m(\epsilon)]_+}{e}. \tag{6}$$

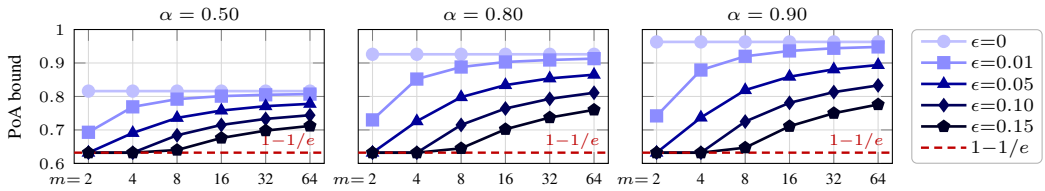


Figure 1: Price-of-anarchy bounds under value concentration, for various concentration settings. Each entry represents the bound for any distribution with $\mathbb{P}_\mu(V_{(m)} \le \alpha\mathbb{E}[V_{(1)}]) \le \epsilon$. At $\epsilon = 0$, we recover the deterministic floor bound from Equation (4); higher ϵ start to degrade the efficiency. Larger numbers of credible competitors ($m \ge 4$) stabilize the market, and preserve efficiency even under high ϵ.

As Theorem 4 holds for all $\alpha \in [0,1]$ and $m \in \{2,\dots,n\}$, taking the supremum over all valid α and m, and setting $\delta(\alpha,m) := \alpha\left[1 - e\mathcal{R}_m(F_m(\alpha\mathbb{E}_\mu[V_{(1)}]))\right]_+$ where $F_m(v) := \mathbb{P}(V_{(m)} \le v)$ gives

$$\inf_{\mathbf{s}\in \mathrm{BNE}(\mu)} \frac{\mathbb{E}_\mu[W(\mathbf{s}(\mathbf{V}),\mathbf{V})]}{\mathbb{E}_\mu[\mathrm{OPT}(\mathbf{V})]} \;\ge\; 1 - \frac{1 - \sup_{\alpha,m}\delta(\alpha,m)}{e}. \tag{7}$$

Theorem 4 takes on a similar form to the simple bound in Equation (4), but incurs a penalty (based on the *bid shortfall* from Lemma 2) that is function of the failure rate of the other bidders to achieve the minimum (normalized) value of $\alpha\mathbb{E}_\mu[V_{(1)}]$. The welfare bonus term $\delta(\alpha,m)$ envelops this penalty.

To illustrate how structure affects the the PoA bound from Theorem 4, in Figure 1 we plot the bound while varying m, α, and a hypothetical failure rate $F_m(\alpha\mathbb{E}_\mu[V_{(1)}]) = \epsilon$. We see that even if the failure rate ϵ is high (e.g., ≈ 0.10), we can still guarantee a stronger PoA if ϵ applies to multiple competitors. For example, if 16 competitors can achieve 90% of the mean best value at least 90% of the time, then the price-of-anarchy rises from $1 - \frac{1}{e} \approx 63\%$ to almost 80%. On the other hand, if only 4 competitors can match this threshold 90% of the time, then our new bound does not improve over the standard one.

Additional theoretical results related to value concentration are included in Section A. Specifically, in Section A.1 we establish a connection between value *correlation* and the notion of value concentration defined here, while Section A.2 gives a version of Theorem 4 that is conditioned on any public information X associated with the query (such as the prompt text or its general topic). The threshold $\alpha\,\mathbb{E}[V_{(1)} \mid X = x]$ can then adapt to how hard each query is, and often leads to a much stronger bound.

4 First-Price Auctions with Weakly-Dependent Values

Value concentration controls the allocation loss when the wrong bidder wins. A separate question is how the *dependence structure* among values affects strategic behavior. More specifically, under independent private values (IPV; where $V_1 \perp \cdots \perp V_n$), Jin & Lu (2023) show that the price of anarchy improves from $1 - 1/e \approx 0.632$ to $1 - 1/e^2 \approx 0.865$. In practice, bidder values are rarely strictly independent—yet they may exhibit only weak dependence. In agentic marketplaces, valuations co-vary across queries due to shared task difficulty, but we conjecture that much of this common variation can be explained by observable query features. If so, once query metadata X is observed (e.g., the prompt text, its topic, or its difficulty as estimated and communicated by the platform or user), agent values may be approximately independent: $\mu(v_1,\dots,v_n \mid X) \;\approx\; \prod_{i=1}^n \mu_i(v_i \mid X)$.

Whether this holds in practice is an empirical question that we examine in Section B.1. Regardless, it motivates an intriguing robustness question: does the $1 - 1/e^2$ guarantee degrade smoothly as values deviate from independence? We prove that it does. Specifically, for each realization of the public state $X = x$, let $\mu_x := \mu(\cdot \mid X = x)$ and $\nu_x := \prod_{i=1}^n \mu_i(\cdot \mid X = x)$, and write $\tau(x) := \|\mu_x - \nu_x\|_{\mathrm{TV}}$ for the total variation distance between the conditional joint distribution and the product of its marginals. We use $\|\mu - \nu\|_{\mathrm{TV}} := \sup_A |\mu(A) - \nu(A)|$. When $\tau(x) = 0$, values are conditionally independent given $X = x$. We then can show the following result.

Theorem 5. *For every fixed $n \ge 2$, there is a nondecreasing function $\eta_n^{\mathrm{cont}} : [0,1] \to [0,\infty)$ with $\eta_n^{\mathrm{cont}}(0) = 0$ and $\eta_n^{\mathrm{cont}}(t) \to 0$ as $t \downarrow 0$ such that the following holds. Suppose $0 \le V_i \le \overline{v}(X) < \infty$ almost surely. Then every no-overbidding Bayes–Nash equilibrium satisfies, for almost every x,*

$$\mathbb{E}[W \mid X = x] \;\ge\; \left(1 - \frac{1}{e^2}\right)\mathbb{E}[V_{(1)} \mid X = x] \;-\; \overline{v}(x)\,\eta_n^{\mathrm{cont}}(\tau(x)). \tag{8}$$

When $\tau(x) = 0$, the error term vanishes and we recover the $1 - 1/e^2$ bound, even if values are arbitrarily correlated *across* queries. Small departures from conditional independence incur an additive penalty scaled by the conditional value bound $\overline{v}(x)$. Converting to a ratio bound requires that $\overline{v}(x)$ is not too large relative to $\mathbb{E}[V_{(1)} \mid X = x]$; integrating over X gives the unconditional guarantee. The proof accounts for changes in both welfare expectations and bidding incentives when the joint law is replaced by the product law. Appendix E gives the full argument, with additional refinements.

Note that value concentration and weak dependence are not mutually exclusive, despite perhaps seeming so. For instance, suppose $V_i = a(X) + \theta Z_i$, where $a(X) \geq \theta$ and, conditional on X, Z_i are zero-mean, independent random variables in $[-1, 1]$. Then V_i are conditionally independent and lie in an interval of width 2θ. When θ is small relative to $a(X)$, they also share a high competitive floor.

5 On implications for agentic market platform design

The structure behind our bounds is partly under the platform's control, leading to some very simple design recommendations to protect market efficiency. In this section we explore what happens if the platform restricts which models bid on a query, or if it enters its own in-house model (e.g., if it has privately developed one, or is able to host an open-weight model) into the auction.

Pre-auction filtering. Many platforms already run a lightweight router that picks a subset $\mathcal{S}(X) \subseteq [n]$ of models based on the public query information X (Hu et al., 2024; Ong et al., 2025, *inter alia*). Let $W_{\mathcal{S}}$ be the welfare at an equilibrium of the auction in which only models in $\mathcal{S}(X)$ bid (note that the randomness in $W_{\mathcal{S}}$ also depends on $\mathcal{S}(X)$). The winner of this auction is always in $\mathcal{S}(X)$ by definition, so simply applying the naive bound of Equation (3) within it and marginalizing over X shows that every such equilibrium satisfies

$$\frac{\mathbb{E}[\min_{i \in \mathcal{S}(X)} V_i]}{\mathbb{E}[V_{(1)}]} \leq \frac{\mathbb{E}[W_{\mathcal{S}}]}{\mathbb{E}[V_{(1)}]} \leq \frac{\mathbb{E}[\max_{i \in \mathcal{S}(X)} V_i]}{\mathbb{E}[V_{(1)}]}. \tag{9}$$

While this bound might be weak for a market with open participation (where a single poor model can drag the welfare bound down, as discussed at the start of Section 3), a platform that *controls* who competes can make it strong by removing weak models, which raises the lower bound. Removing the best model (i.e., if $V_{(1)} \notin \mathcal{S}(X)$), however, also lowers the upper bound; thus filtering does incur some risk. Still, that risk can be mitigated by being more conservative in filtering, and as long as the router keeps the strongest models, the bounds of Section A.2 still apply to the filtered market $\mathcal{S}(X)$.

Platform participation. A platform with its own baseline model can also enter it into the auction, announcing its value V_0 for the query and then bidding it (effectively serving as a reserve price). We include V_0 in the public state X, so every model can condition on it, and the platform's model acts as a query-specific reserve price. This mirrors private-label strategies from retail economics (see, e.g., Shopova, 2023; Ding et al., 2025). Given $V_0 = v_0$, no bid below v_0 can win, so $v_{\text{floor}} = v_0$. Applying Proposition 1 conditionally on V_0 and averaging shows that every equilibrium satisfies (Section D.2)

$$\frac{\mathbb{E}[W]}{\mathbb{E}[V_{(1)}]} \geq 1 - \frac{1 - \mathbb{E}[V_0]/\mathbb{E}[V_{(1)}]}{e}, \tag{10}$$

where $V_{(1)}$ now includes the platform value V_0. As an example, a platform model worth half the expected top value ($\mathbb{E}[V_0] = 0.5\,\mathbb{E}[V_{(1)}]$) raises the guarantee from $\approx 63\%$ to $\approx 82\%$, like the deterministic floor with $\alpha = 0.5$ in Proposition 1. But the platform's model only needs to be good on average.

6 Experiments

We now test how well our bounds describe welfare in a realistic LLM market. We run one auction with all models on the full EmbedLLM dataset (Zhuang et al., 2025), compare the value concentration bound to realized welfare over many random sub-markets, and study hand-picked cohorts that fix optimal welfare while varying market structure. The appendix contains additional results that show how close, empirically, values are to conditionally independent given the query (Section B.1), as well as experimental results with platform-owned bidders and pre-auction filtering (Section B.2-B.3).

		Oracle	Random	Best single	Best per cluster	FPA	FPA (adv)
Full dataset	Empirical welfare	0.768	0.416	0.600	0.686	0.761	0.758
	Empirical PoA	1.000	0.542	0.781	0.894	0.992	0.988
Worst cluster	Empirical PoA	1.000	0.083	0.060	0.640	0.961	0.932

Table 1: Realized welfare on the full EmbedLLM market under different allocation rules, reported both raw and as a ratio to the optimal welfare. Full dataset welfare is averaged over all prompts, while worst-cluster welfare is averaged on each rule's lowest-welfare cluster (out of 20 embedding-based clusters).

6.1 Setting

Dataset. We simulate auctions on EmbedLLM (Zhuang et al., 2025), which aggregates over $30k$ prompts across math, code, question answering, and reasoning benchmarks, along with the scored responses from 112 models of various sizes. To align with our convention that a model's value is its expected reward minus its cost, we set $v_i(q) = \max(0, \hat{\mathbb{E}}[\text{acc}_i \mid q] - c_i)$, where $\hat{\mathbb{E}}[\text{acc}_i \mid q]$ estimates the probability that model i answers query q correctly, computed by k-nearest neighbors ($k = 16$) over Gemini prompt embeddings (Shanbhogue et al., 2026), excluding q. Following Jitkrittum et al. (2025), we set c_i to be linear in parameter count, at 1 per trillion parameters, so that, e.g., Qwen-1.5-7B-Chat has a cost of .00772 and Qwen-1.5-32B-Chat has a cost of .0325. Two models were excluded because their parameter count could not be determined, leaving 110 total.

Strategic bidding. Each model submits a sealed bid $b_i(q) \leq v_i(q)$ for each query q, and the highest bidder wins and pays its bid. We partition the prompts into 20 K-means clusters over the embeddings, and bidders condition their strategies on the cluster and their own value $v_i(q)$. A bidder therefore knows the approximate type of prompt it is bidding on and what it is worth to it, but knows its rivals only through prior distribution of values, conditioned on the cluster. Within each cluster we solve for (approximate) Bayes–Nash equilibria by fictitious play using a discretized grid over both values and possible bids (see Section C for details). Equilibria are not necessarily unique, and certain equilibria might have substantially worse welfare than others. We therefore solve every market twice: once for a *typical* equilibrium, where at each iteration bidders update their strategy to the best (surplus maximizing) response nearest its current bid until convergence, and once for an (also approximate) *adversarial* equilibrium, where at each iteration each bidder takes the welfare-*minimizing* response that still retains at most 5% relative regret of its best possible individual surplus until convergence.

Metrics. Welfare within a market is pooled over all prompts and reported as a fraction of the first-best oracle, $\text{OPT} = \mathbb{E}[\max_i v_i(q)]$, which we call the "realized PoA" to distinguish it from the worst-case bounds of Section 3. By a "market" we mean a subset of the 110 models bidding against each other, and most of our experiments draw many such subsets at random, with the sampling scheme stated in each case. Every number we report is therefore computed over prompts and over sampled markets. Since the bounds we are testing are worst-case statements, we report both the mean and the 0.1-percentile (`p0.1`) of realized PoA across markets,[4] and we compare both against the value concentration bound of Theorem 4 evaluated on the same market (where like the equilibrium strategies, the bound is computed using the empirical distribution of the bidder values). We also track how much the winner shades, $1 - \mathbb{E}[B_{\text{winner}}/V_{\text{winner}}]$, along with the platform's revenue, $\mathbb{E}[B_{\text{winner}}]$. Note, however, that shading or low revenue (i.e., low winning bids) on their own do not cost welfare, since queries are misallocated only when bidders shade by (problematically) different amounts.

6.2 Results

Results on the full EmbedLLM market. Table 1 reports a single auction run on the full dataset with all 110 models bidding against each other. The auction at equilibrium is nearly perfectly efficient, reaching 99.2% of oracle welfare at the typical equilibrium and 98.8% at the adversarial one. Both exceed the baseline of routing each cluster to its own best model (89.4%), and are well above the baseline of committing to one model everywhere (78.1%). The gap widens on each rule's worst cluster, where the single best model retains only 6.0% of first-best and per-cluster routing 64.0%, while the auction holds 96.1% and 93.2%. Bidders see their own value $v_i(q)$ before bidding, so a

[4] We use 99.9% coverage instead of the minimum to avoid outlier markets where the solver has not converged.

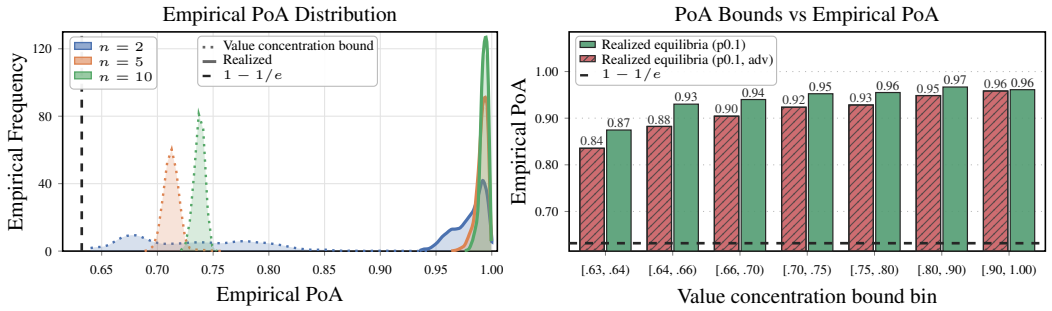


Figure 2: **Left:** Empirical distributions of the Theorem 4 bound (dotted) and the empirical PoA (solid) across $100k$ random n-model subsets of EmbedLLM, for $n \in \{2, 5, 10\}$. Both distributions tighten and shift upward as n grows, and are well above the distribution-free baseline (dashed). **Right:** Bottom 0.1% of empirical PoA across sub-markets grouped by bound bin ($15k$ samples per bin), for both typical equilibria (solid green) and approximate, adversarial equilibria (hatched red).

specialist can outbid the cluster leader on exactly the prompts where the leader is weak, which is what a fixed rule cannot do. On the full market and prompt distribution, Theorem 4 guarantees a welfare ratio $\geq 78.9\%$ (and a minimum bound of 65.9% on the worst cluster). Both are substantially above the naive $1 - 1/e$ bound, though still below what the market achieves. The $1 - 1/e$ bound holds for every market and so cannot indicate anything more specific about this one, whereas 78.9% is computed from the empirical value prior of this market, and correctly signals that welfare will be high. Note that we do not expect it to be tight, since the bound only uses value concentration and treats everything else as worst case, while near-independence (Section 4) is a second and separate reason this market can be efficient. We only report the concentration bound here, and leave combining the two to future work.

Results on randomly sampled EmbedLLM sub-markets. Moving beyond the full EmbedLLM market setting, we next look at the welfare for thousands of auctions held just between subsets of EmbedLLM models. Figure 2 (left) plots the empirical distributions of both (a) the value concentration bound of Theorem 4, and (b) the realized welfare ratio at a typical equilibrium over $100k$ randomly sampled model subsets of market size $n = 2$, 5, and 10. Both distributions increase in mean and concentrate as n grows. Realized welfare is much higher than $1 - 1/e$, as are our bounds (though while the value concentration bounds are tighter, they are still conservative, as in the result for the full market). In practice, the realized welfare ratio of the vast majority of sampled auctions falls above 95%. That said, unlike the worst-case $1 - 1/e$ bound, which is the same for every market, the value concentration bounds are *sensitive* to the market they are computed for. Figure 2 (right) groups randomly sampled model subsets for market sizes $n \in \{2, \ldots, 25\}$ into bins according to their value concentration bound, where the bins span $[1 - 1/e, 1.00)$. We sample $15k$ subsets per bin, and report the bottom 0.1% of realized welfare ratios within each bin, for both the typical and the adversarial strategy solver. The realized welfare at an adversarial equilibrium rises with the bins, from 0.84 in the $[0.63, 0.64)$ bin to 0.96 in the $[0.90, 1.00)$ bin. At a typical equilibrium it rises from 0.87 to 0.97 in the $[0.80, 0.90)$ bin and stays at 0.96 in the top bin (while the standard $1 - 1/e$ bound stays flat, of course).

Results on illustrative EmbedLLM cohorts. Random sub-markets vary in both how good their best model is and how much competition it faces. To separate the two, Table 2 shows results for four hand-picked cohorts of $n = 4$ models with the same first-best welfare ($\mathbb{E}[\max_i V_i] \approx 0.55$) but different market structure. In the *similar* cohort, the four models are comparable 7B chat models (Starling-LM-7B-beta, OpenChat-3.5-0106, Zephyr-7B-beta, and Mistral-7B-Instruct-v0.1) with average values between 0.43 and 0.51. In the *diverse* cohort (Qwen-1.5-7B-Chat, BioMistral-7B-DARE, ANIMA-Nectar-v2, and Zephyr-beta-math), the models have an even tighter range of average values (0.42 to 0.46), but they are good at different things, so on a given prompt the runner-up tends to trail the winner by more (17.1% versus 7.7%). For both cohorts the value concentration bound is well above $1 - 1/e$, and it is higher for the similar cohort (0.748 versus 0.719). Realized welfare is high in both, but the winner shades much more in the diverse cohort (19.8% versus 12.4%), so the similar cohort ends slightly higher (0.996 versus 0.988). Nevertheless, this diversity is mild, as every model is still a credible rival on many prompts. When a "generalist" (Yi-34B-Chat) is paired with three weaker, "niche" models (Qwen-1.5-0.5B-Chat, DeepSeek-Coder-6.7B-Instruct, and Gemma-2B-it), the runner-up

Cohort	$\mathbb{E}[\max_i V_i]$	$1 - \mathbb{E}[V_{(2)}]/\mathbb{E}[V_{(1)}]$	Bound	Shading	FPA (adv.)
Generalists (similar)	0.557	7.7%	0.748	12.4%	0.996
Generalists (diverse)	0.553	17.1%	0.719	19.8%	0.988
Generalists (1) + niche (3)	0.552	35.6%	0.687	29.2%	0.949
Generalists (2) + niche (2)	0.555	15.9%	0.707	22.3%	0.992

Table 2: Four cohorts constructed to have approximately equal $\mathbb{E}[\max_i V_i]$, but different value structure. Shading is $1 - \mathbb{E}[B_{\text{winner}}/V_{\text{winner}}]$, and empirical PoA is evaluated under the adversarial solver.

trails by 35.6% on average, the winner shades by 29.2%, and both the bound (0.687) and realized welfare (0.949) drop. Swapping one niche model for a second generalist is enough to bring welfare back up to 0.992, even though half the cohort is still weak (here the generalists are SOLAR-10.7B-Instruct and Neural-Chat-7B-v3-3, paired with Qwen-1.5-0.5B-Chat and Gemma-2B-it). Welfare stays high as long as the leader faces at least one close rival, and the remaining models can be much weaker.

7 Related work

Prior work on agentic markets has focused on dominant strategy incentive-compatible (DSIC) mechanisms, where bidding one's true value (or cost, in a procurement auction) is optimal. A sequential auction approach to solving Markov decision problems through sequential Vickrey (second-price) auctions requires subsidies from the mechanism, prompting the use of "clones" of each participant to make the first- and second-price bids coincide (Chang et al., 2020). In concurrent work, procurement auctions have been studied under a Vickrey–Clarke–Groves (VCG) mechanism, which is also DSIC (Liu et al., 2026). Both mechanisms are susceptible to collusion (Robinson, 1985) and require bidders to trust the auctioneer (Rothkopf et al., 1990). We argue that the FPA's lack of truthfulness is less of a concern under value structures we expect agentic markets to have, such as concentrated and/or only weakly dependent values.

The price of anarchy (PoA) of first-price auctions is well studied. A universal $1 - 1/e \approx 0.632$ bound has been established under arbitrary correlation via smoothness arguments (Syrgkanis & Tardos, 2013; Christodoulou et al., 2016; Roughgarden et al., 2017). For independent private values, the lower bound was raised to 0.743 by reasoning beyond smoothness (Hoy et al., 2018), and subsequently closed by proving that $1 - 1/e^2 \approx 0.865$ is tight (Jin & Lu, 2023). All of these bounds are worst case over the value distribution. We show how they *can* improve when additional structure on the values is assumed. More generally, our experiments connect to the broader agenda of studying the PoA beyond worst-case distributions, drawing on data-driven and empirical mechanism design (Hoy et al., 2017; Hartline, 2026). A similar view of equilibria weights each outcome by how likely learning dynamics are to reach it (Panageas & Piliouras, 2016).

8 Conclusion

This paper analyzes the welfare of first-price auctions (FPAs) among AI agents, where the main objection to non-truthful mechanisms like the FPA is that they are vulnerable to welfare loss from bid shading. Worst-case price of anarchy bounds are pessimistic, but agentic markets often have structure that narrows the gap. We prove that the welfare guarantee rises toward full efficiency when the strongest few bidders have similar values, and that the $1-1/e^2$ bound for independent values degrades smoothly when values are only weakly dependent (e.g., given the query). Finally, we show that platforms themselves can take simple steps to protect the efficiency of their markets, such as filtering out weak bidders with routing rules, or entering their own models to set competitive reserve prices.

Limitations and future work. We assume that models and users are aligned in the sense that positive utility derives only from providing value to the user; external sources of utility, such as advertising placement, would obviously require additional analysis. Second, while this paper focuses only on welfare, follow-up work should consider whether the special features of agentic markets may also help to tighten price-of-anarchy bounds on the auctioneer's revenue (Hartline et al., 2014).

AI USE

We used generative AI to refine mathematical claims and develop proof strategies. All proofs are written and checked by the authors. AI assistance was also used to produce the code for the experiments and the figures, and to help draft or edit parts of the paper for readability. All code and text was rigorously checked, revised, refined, and validated by the authors. We used an AI automated reviewer to evaluate an early draft of the manuscript and to suggest improvements. We take complete responsibility for all aspects of the final content of this work.

REPRODUCIBILITY

The empirical part of the paper is based on public data (Zhuang et al., 2025). The main implementation is the BNE approximate equilibrium solver, which is described in detail in Section C. Other aspects of the experiments are described in Section 6.

REFERENCES

Olivier Armantier, Jean-Pierre Florens, and Jean-Francois Richard. Approximation of nash equilibria in bayesian games. *Journal of Applied Econometrics*, 23(7):965–981, 2008. doi: https://doi.org/10.1002/jae.1040. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.1040.

Malika Aubakirova, Alex Atallah, Chris Clark, Justin Summerville, and Anjney Midha. State of AI: An empirical 100 trillion token study with OpenRouter. https://openrouter.ai/state-of-ai, December 2025. a16z and OpenRouter Inc.

Vitor Bosshard, Benedikt Bünz, Benjamin Lubin, and Sven Seuken. Computing bayes-nash equilibria in combinatorial auctions with verification. *Journal of Artificial Intelligence Research*, 69:531–570, 2020.

George W. Brown. Iterative solution of games by fictitious play. In T. C. Koopmans (ed.), *Activity Analysis of Production and Allocation*, pp. 374–376. Wiley, New York, 1951.

Yuhan Cao, Yu Wang, Sitong Liu, Miao Li, Yixin Tao, and Tianxing He. Pay for the second-best service: A game-theoretic approach against dishonest llm providers. In *Proceedings of the ACM Web Conference 2026*, WWW '26, pp. 249–260, New York, NY, USA, 2026. Association for Computing Machinery. ISBN 9798400723070. doi: 10.1145/3774904.3792446. URL https://doi.org/10.1145/3774904.3792446.

Michael Chang, Sid Kaushik, S Matthew Weinberg, Tom Griffiths, and Sergey Levine. Decentralized reinforcement learning: Global decision-making via local economic transactions. In *International Conference on Machine Learning*, pp. 1437–1447. PMLR, 2020.

Shuchi Chawla and Jason D. Hartline. Auctions with unique equilibria. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pp. 181–196, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450319621. doi: 10.1145/2492002.2483188. URL https://doi.org/10.1145/2492002.2483188.

Lingjiao Chen, Matei Zaharia, and James Zou. FrugalGPT: How to use large language models while reducing cost and improving performance. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=cSimKw5p6R. Featured Certification.

George Christodoulou, Annamária Kovács, Alkmini Sgouritsa, and Bo Tang. Tight bounds for the price of anarchy of simultaneous first-price auctions. *ACM Trans. Econ. Comput.*, 4(2), January 2016. ISSN 2167-8375. doi: 10.1145/2847520. URL https://doi.org/10.1145/2847520.

Sam Cox. Simplifying programmatic: First price auctions for Google Ad Manager. https://blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager/, March 2019. Google Ad Manager Blog.

Tiffany Ding, Dominique Perrault-Joncas, Orit Ronen, Michael I. Jordan, Dirk Bergemann, Dean Foster, and Omer Gottesman. Marketplace operators can induce competitive pricing, 2025. URL https://arxiv.org/abs/2503.06582.

Jason Hartline. The economics of no-regret learning algorithms, 2026. URL https://arxiv.org/abs/2601.22079.

Jason Hartline, Darrell Hoy, and Sam Taggart. Price of anarchy for auction revenue. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pp. 693–710, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450325653. doi: 10.1145/2600057.2602878. URL https://doi.org/10.1145/2600057.2602878.

Benjamin Heymann and Panayotis Mertikopoulos. An empirical study of fictitious play for estimating nash equilibria in first-price auctions with correlated values. *arXiv preprint arXiv:2606.16389*, 2026. URL https://arxiv.org/abs/2606.16389.

Darrell Hoy, Denis Nekipelov, and Vasilis Syrgkanis. Welfare guarantees from data. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/d3a7f48c12e697d50c8a7ae7684644ef-Paper.pdf.

Darrell Hoy, Samuel Taggart, and Zihe Wang. A tighter welfare guarantee for first-price auctions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pp. 132–137, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355599. doi: 10.1145/3188745.3188944. URL https://doi.org/10.1145/3188745.3188944.

Qitian Jason Hu, Jacob Bieker, Xiuyu Li, Nan Jiang, Benjamin Keigwin, Gaurav Ranganath, Kurt Keutzer, and Shriyash Kaustubh Upadhyay. Routerbench: A benchmark for multi-llm routing system. *arXiv preprint arXiv: 2403.12031*, 2024.

Yaonan Jin and Pinyan Lu. First price auction is $1 - 1/e^2$ efficient. *J. ACM*, 70(5), October 2023. ISSN 0004-5411. doi: 10.1145/3617902. URL https://doi.org/10.1145/3617902.

Wittawat Jitkrittum, Harikrishna Narasimhan, Ankit Singh Rawat, Jeevesh Juneja, Congchao Wang, Zifeng Wang, Alec Go, Chen-Yu Lee, Pradeep Shenoy, Rina Panigrahy, et al. Universal model routing for efficient LLM inference. *arXiv preprint arXiv:2502.08773*, 2025.

Xiao Liu, Haoyang Li, Songwei Li, Hongbo Fang, Fengli Xu, Feng Shi, and James Evans. Markets, not planners: Decentralized orchestration of llm agents with private information. *arXiv preprint arXiv:2608.23867*, 2026.

Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M Waleed Kadous, and Ion Stoica. RouteLLM: Learning to route LLMs from preference data. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=8sSqNntaMr.

Ioannis Panageas and Georgios Piliouras. Average case performance of replicator dynamics in potential games via computing regions of attraction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pp. 703–720, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450339360. doi: 10.1145/2940716.2940784. URL https://doi.org/10.1145/2940716.2940784.

Zhenting Qi, Huangyuan Su, Ao Qu, Chenyu Wang, Yu Yao, Han Zheng, Kushal Chattopadhyay, Guowei Xu, Zihan Wang, Weirui Ye, Vijay Janapa Reddi, Ju Li, Paul Pu Liang, Himabindu Lakkaraju, Sham Kakade, and Yilun Du. Economy of minds: Emerging multi-agent intelligence with economic interactions. *ArXiv preprint: 2606.02859*, 2026. URL https://arxiv.org/abs/2606.02859.

Zinovi Rabinovich, Victor Naroditskiy, Enrico H. Gerding, and Nicholas R. Jennings. Computing pure bayesian-nash equilibria in games with finite actions and continuous types. *Artificial Intelligence*, 195:106–139, 2013. ISSN 0004-3702. doi: https://doi.org/10.1016/j.artint.

2012.09.007. URL https://www.sciencedirect.com/science/article/pii/S0004370212001191.

Marc S Robinson. Collusion and the choice of auction. *The RAND Journal of Economics*, pp. 141–145, 1985.

Michael H Rothkopf, Thomas J Teisberg, and Edward P Kahn. Why are vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.

Tim Roughgarden, Vasilis Syrgkanis, and Éva Tardos. The price of anarchy in auctions. *J. Artif. Int. Res.*, 59(1):59–101, May 2017. ISSN 1076-9757.

Madhuri Shanbhogue, Zhe Li, Shanfeng Zhang, Gustavo Hernández Ábrego, Shih-Cheng Huang, Aashi Jain, Daniel Salz, Sonam Goenka, Chaitra Hegde, Ji Ma, Feiyang Chen, Jiaxing Wu, Tanmaya Dabral, Babak Samari, Kevin Poulet, Daniel Cer, Kaifeng Chen, Paul Suganathan, Hui Hui, Jovan Andonov, Philippe Schlattner, Jay Han, Iftekhar Naim, Wing Lowe, Vladimir Pchelin, Albert Yang, Yi-Ting Chen, Zhongli Ding, Grace Zhang, Georg Heigold, Yichang Chen, Antoine Reveillon, Brendan Mccloskey, Wenlei Zhou, Dahun Kim, Rui Meng, Emma Wang, Jack Zheng, Halley Fede, Zhen Yang, Keegan Mosley, Brian Potetz, Sahil Dua, Henrique Schechter Vera, Shen Gao, Hesen Zhang, Andreas Hess, Hengxuan Ying, Alberto Montes, Karan Gill, Min Choi, Sebastian Russo, Anja Hauth, Jinhyuk Lee, Michael Boratko, Megan Barnes, Vikram Rao, Claudiu Musat, Cyril Allauzen, Ehsan Variani, Shankar Kumar, Tom Bagby, Junyi Jiao, Yang Gu, Tengxin Li, Ayush Agrawal, Roberto Santana, Dev Nath, Stephen Karukas, Shuoxuan Han, Lucia Loher, Alice Twu, Nidhi Vyas, Siddharth Bhai, Frank Palma Gomez, Wangyuan Zhang, Chaoren Liu, Jizheng Yang, Steve Qiu, Shijie Zhang, Sujay Kulkarni, Sascha Rothe, Sean Nakamoto, Raphael Hoffmann, Zach Gleicher, Yunhsuan Sung, Qin Yin, Tom Duerig, and Mojtaba Seyedhosseini. Gemini embedding 2: A native multimodal embedding model from gemini. *arXiv preprint: arXiv 2605.27295*, 2026. URL https://arxiv.org/abs/2605.27295.

Radostina Shopova. Private labels in marketplaces. *International Journal of Industrial Organization*, 89:102949, 2023. ISSN 0167-7187. doi: https://doi.org/10.1016/j.ijindorg.2023.102949. URL https://www.sciencedirect.com/science/article/pii/S0167718723000309.

Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pp. 211–220, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450320290. doi: 10.1145/2488608.2488635. URL https://doi.org/10.1145/2488608.2488635.

Richard Zhuang, Tianhao Wu, Zhaojin Wen, Andrew Li, Jiantao Jiao, and Kannan Ramchandran. Embedllm: Learning compact representations of large language models. In *International Conference on Learning Representations*, volume 2025, pp. 76913–76926, 2025.

CONTENTS

A Additional theoretical results **14**

A.1 Correlation and value concentration 14

A.2 Query-conditional value concentration 14

B Additional experimental results **15**

B.1 Weak dependence given the prompt 15

B.2 Platform-owned bidder 15

B.3 Pre-auction filtering 16

C Equilibrium solver **17**

D Proofs for value concentration **17**

D.1 Convexity of the floor bonus 20

D.2 Platform reserve bound 21

D.3 Robustness of the $\left(1 - \frac{1}{e}\right)$ bound 21

E Proofs and additional results for weak dependence **21**

E.1 The welfare modulus 22

E.2 Proof of the main qualitative theorem (Theorem 5). 26

A Additional theoretical results

A.1 Correlation and value concentration

Positive correlation between the best and second-best values can sharpen the bound from competitive floors, even when their means and variances are held fixed. The following is a corollary of Theorem 4.

Corollary A.1. *Under the assumptions of Theorem 4, denote $U = V_{(2)}$ and $Y = V_{(1)}$. Suppose U, Y have finite positive standard deviations σ_U, σ_Y and $\mathrm{Corr}(U, Y) \geq \theta$ for some $\theta \in [0, 1)$. For $0 \leq \alpha < \mathbb{E}[U]/\mathbb{E}[Y]$, define $a_\alpha := \frac{\mathbb{E}[U] - \alpha\mathbb{E}[Y]}{\sigma_U}$, $b_\alpha := \frac{(1-\alpha)\mathbb{E}[Y]}{\sigma_Y}$, and $\epsilon_\theta(\alpha) := \left(1 + a_\alpha^2 + \frac{[\theta a_\alpha - b_\alpha]_+^2}{1-\theta^2}\right)^{-1}$. Then $\mathbb{P}(U \leq \alpha\mathbb{E}[Y]) \leq \epsilon_\theta(\alpha)$, and every no-overbidding Bayes–Nash equilibrium satisfies $\frac{\mathbb{E}[W]}{\mathbb{E}[Y]} \geq 1 - \frac{1}{e} + \alpha\left[\frac{1}{e} - \mathcal{R}_2\big(\epsilon_\theta(\alpha)\big)\right]_+$.*

The term $[\theta a_\alpha - b_\alpha]_+^2/(1-\theta^2)$ is the improvement over the one-sided Chebyshev bound, which uses only the runner-up's mean and variance. It is positive when $\theta a_\alpha > b_\alpha$. The two values need not have equal means or variances.

A.2 Query-conditional value concentration

The results in the main text apply to a fixed prior μ. In practice, each query comes with public information X, such as the prompt text, that all bidders see before they bid. Let μ_x be the law of $\mathbf{V}$ given $X = x$, and define

$$M_x := \mathbb{E}[V_{(1)} \mid X = x], \qquad F_m^x(c) := \mathbb{P}(V_{(m)} \leq c \mid X = x).$$

Write $g_m(\epsilon) := [1 - e\mathcal{R}_m(\epsilon)]_+$. The conditional welfare bonus is $\delta(\alpha, m \mid x) := \alpha\, g_m\big(F_m^x(\alpha M_x)\big)$ when $M_x > 0$, and $\delta(\alpha, m \mid x) := 0$ when $M_x = 0$.

Corollary A.2. *Suppose $0 < \mathbb{E}[V_{(1)}] < \infty$, and that for almost every x the assumptions of Theorem 4 hold under μ_x and bidding is a no-overbidding Bayes–Nash equilibrium under μ_x. Then for any $\alpha \geq 0$ and $m \in \{2, \ldots, n\}$,*

$$\mathbb{E}[W] \geq \mathbb{E}_X\left[\left(1 - \frac{1 - \delta(\alpha, m \mid X)}{e}\right) M_X\right]. \tag{11}$$

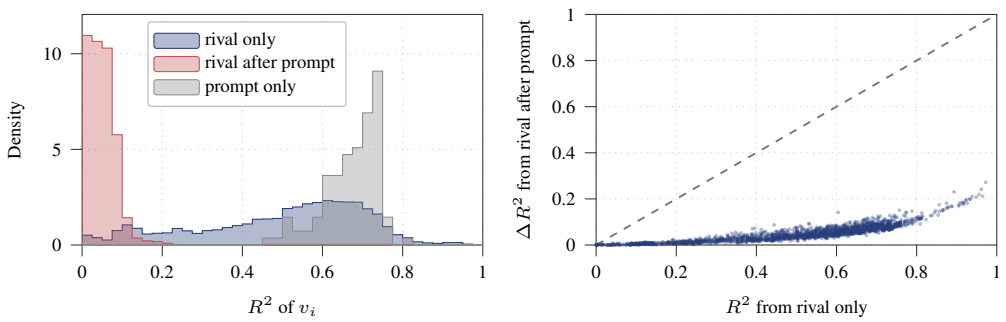
Because Theorem 4 is applied separately for each x, α and m can also be chosen per query. With that freedom, the conditional bound is never worse than applying Theorem 4 directly to μ, and in practice it can be much better, for two reasons.

The first is that the threshold αM_x adapts to each query, while the unconditional bound uses one threshold $\alpha\,\mathbb{E}[V_{(1)}]$ for all queries. On an easy query where every model does well, the threshold is high but most models still clear it. On a hard query the threshold drops, so again few models fall short. A single fixed threshold is too high for hard queries, which inflates the failure rate, and too low for easy ones, where α could have been larger.

The second is that conditioning helps even when the threshold is held fixed. Fix an absolute threshold $c \geq 0$, and let $F_m(c) := \mathbb{P}(V_{(m)} \leq c) = \mathbb{E}_X[F_m^X(c)]$ be the failure rate pooled over all queries. Since g_m is convex (Theorem D.3), Jensen's inequality gives

$$\mathbb{E}_X\big[g_m(F_m^X(c))\big] \geq g_m\big(\mathbb{E}_X[F_m^X(c)]\big) = g_m(F_m(c)). \tag{12}$$

At threshold c, Theorem 4 gives $\mathbb{E}[W \mid X = x] \geq (1 - 1/e)M_x + (c/e)\, g_m(F_m^x(c))$ for each query. Averaging over X and applying (12), the result is at least $(1-1/e)\mathbb{E}[V_{(1)}] + (c/e)\, g_m(F_m(c))$, which is the unconditional bound at the same threshold. So averaging the per-query bonuses never does worse than computing a single bonus from the pooled failure rate, and it can do strictly better when failure rates vary across queries.



(a) Variance of v_i explained

(b) Rival information left after prompt

Figure 3: How well the prompt and a rival predict each model's value v_i. **(a)** Distribution of R^2 from the held-out ridge prediction $g_i(q)$ alone, from a rival's value v_j alone, and the gain from adding the rival once the prompt is known. **(b)** Each point is one of 2,000 randomly sampled model pairs (i, j), plotting the rival's R^2 on its own against the gain it adds once the prompt is known. The dashed line is $y = x$. An (i, j) pair on this line would gain as much from the rival j's value v_j after the prompt as before, while points far below it are pairs whose dependence the prompt mostly explains.

B Additional Experimental Results

B.1 Weak Dependence Given the Prompt

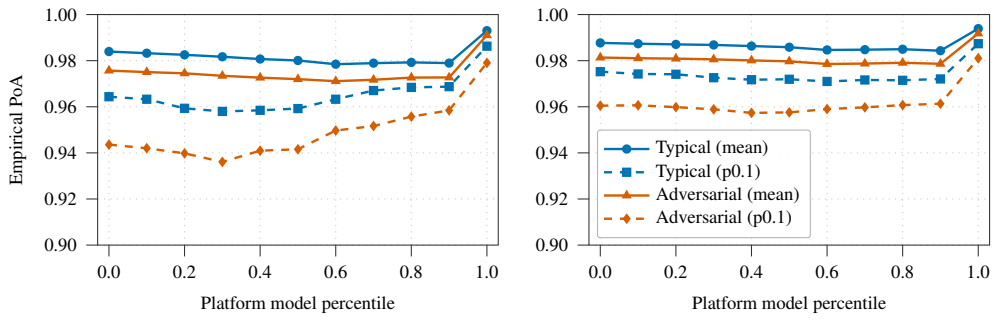
Section 4 conjectures that most of the dependence between bidder values comes from the query, so that values are close to independent once the query is known. Recall from Section 6 that $v_i(q) = \max(0, \hat{\mathbb{E}}[\text{acc}_i \mid q] - c_i)$, where the accuracy estimate averages over the $k = 16$ nearest neighbors of q. Models' values on the same prompt are thus built from the same neighbors, which makes them move together. To test how much of this the prompt alone can predict, we fit a ridge regression g_i from the prompt embedding to v_i for each model, splitting the prompts into five parts and predicting each part with a regression trained on the other four. For every ordered pair (i, j) of the 110 models, we then compare the R^2 of v_i from $g_i(q)$ alone, from the rival's value v_j alone, and from both. Each $R^2 = 1 - \sum_q (v_i(q) - \hat{v}_i(q))^2 / \sum_q (v_i(q) - \bar{v}_i)^2$ is the fraction of the variance of v_i across prompts that a linear fit $\hat{v}_i$ on the given predictors explains, where $\bar{v}_i$ is the mean of v_i. The gain from adding v_j to $g_i(q)$ is what the rival still tells us about v_i once the prompt is known.

Figure 3 shows the results. On average, the prompt alone explains $R^2 = 0.67$ of the variance of v_i, and a rival alone explains 0.50. Once the prompt is known, adding the rival raises R^2 by only 0.05 on average (0.09 at the 90th percentile). The prompt therefore accounts for 90% of what a rival's value explains, and it is the better predictor in 84% of pairs. The pairs with the most dependence left over are near-duplicate models whose values are almost perfectly correlated to begin with.

B.2 Platform-Owned Bidder

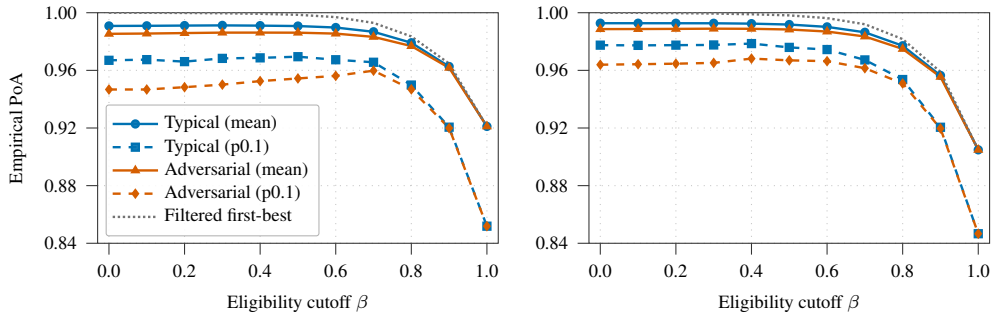
We test the platform participation strategy of Section 5 on randomly sampled markets from EmbedLLM. Each market contains the strongest EmbedLLM model, n rivals drawn at random, and a platform model that bids its own value on every prompt. The platform model is the EmbedLLM model at a given percentile of the ranking by average value, where 0 is the weakest model and 1 is the strongest (so at 1 the platform runs a copy of the best model). Rivals are never drawn from the platform candidates. For each percentile and $n \in \{5, 10\}$ we draw 10,000 random markets and solve for typical and adversarial equilibria as in Section 6.

Results are shown in Figure 4. These markets always include the strongest model and are already efficient, with a mean PoA of about 0.98, so a weak or average platform model changes welfare very little. The adversarial lower tail, however, improves noticeably as the platform model moves into the top half, from 0.94 to 0.96 at $n = 5$, and a platform model that matches the best model lifts every curve to about 0.99 (note that the best *single* model is not the same as the optimal composite model



(a) $n = 5$ (b) $n = 10$

Figure 4: Empirical PoA with a platform-owned bidder, as a function of where the platform model sits in the ranking of EmbedLLM models by average value. Each point averages 10,000 markets made of the strongest model, n random rivals, and the platform model. Solid lines are means and dashed lines are the 0.1 percentile across markets, for typical (blue) and adversarial (orange) equilibria.



(a) $n = 5$ (b) $n = 10$

Figure 5: Empirical PoA under pre-auction filtering, as a function of the eligibility cutoff β. Each point averages 10,000 random markets of n models. Solid lines are means and dashed lines are the 0.1 percentile across markets, for typical (blue) and adversarial (orange) equilibria. The dotted line is the filtered first-best, the welfare of always allocating to the best admitted model.

that yields $\mathbb{E}[\max_i V_i]$ welfare). As expected, the larger markets at $n = 10$ are already inherently more efficient, and depend less on the platform model to keep bids competitive.

B.3 Pre-auction filtering

We next test pre-auction filtering with a router that admits a model to a prompt's auction only if the model is close to the strongest model on the prompt's cluster. Specifically, model i may bid on cluster c if $\mathbb{E}[V_i \mid c] \geq \beta \max_j \mathbb{E}[V_j \mid c]$, where the maximum runs over the models in the market. Setting $\beta = 0$ admits everyone, and $\beta = 1$ admits only the cluster leader. Because the rule works at the cluster level, it can drop the model with the highest value on a particular prompt. We therefore measure welfare against the first-best of the unfiltered market, and also report the filtered first-best, which is the upper bound of Equation (9). For each β and $n \in \{5, 10\}$ we draw 10,000 random markets.

Results are shown in Figure 5. For $\beta \leq 0.6$ the router removes only weak models, and the filtered first-best stays near 1 while mean welfare is unchanged at about 0.99. The adversarial lower tail, however, *does* improve slightly over this range, from 0.947 to 0.960 at $n = 5$ and from 0.964 to 0.968 at $n = 10$. Past $\beta = 0.7$ the router starts dropping models that would have been best on some prompts. Consequently, the filtered first-best falls, and realized welfare falls with it, down to about 0.91 when only the cluster leader bids. As expected, filtering with a router helps only while it keeps the strongest models.

C EQUILIBRIUM SOLVER

We compute approximate Bayes–Nash equilibria with a damped form of fictitious play (Rabinovich et al., 2013; Heymann & Mertikopoulos, 2026), implemented directly on the empirical distribution of values. Fictitious play goes back to Brown (1951), and approximating Bayes–Nash equilibria numerically by iterating best responses over a restricted family of strategies is standard for auctions without closed-form solutions (Armantier et al., 2008; Bosshard et al., 2020). Each cluster of prompts is solved on its own. Within a cluster, every prompt is one draw from the joint distribution of the bidders' values, so the solver keeps whatever correlation the data has and never replaces it with a product of marginals or a smoothed prior. Fictitious play has no general convergence guarantee in first-price auctions, but Heymann & Mertikopoulos (2026) find that it converges numerically to approximate equilibria across a wide range of first-price auctions with correlated values, which is our setting.

Strategies. Within a cluster, we split each bidder's values into 5 bins of equal mass. A bidder's strategy is a shading factor for each bin, so a bidder in bin k of cluster c bids $b_i(q) = \alpha_{i,c,k}\, v_i(q)$. A bid of this form never exceeds the bidder's value. The bidder knows the prompt's cluster and its own value, and it knows its rivals only through the empirical distribution of their values and bids within the cluster.

Best responses. Fix bidder i, cluster c, and bin k, and let Q be the prompts of cluster c on which v_i falls in bin k. For every prompt $q \in Q$ we record the highest rival bid $r(q)$ and that rival's value. When the platform enters its own model (Section B.2), its bid counts as one more rival bid. For a candidate shading factor α, bidder i's average utility over Q is

$$U(\alpha) = \frac{1}{|Q|} \sum_{q \in Q} (1-\alpha)\, v_i(q) \left(\mathbf{1}[\alpha v_i(q) > r(q)] + \tfrac{1}{2}\, \mathbf{1}[\alpha v_i(q) = r(q)] \right), \tag{13}$$

where ties are split evenly. The welfare $W(\alpha)$ is the average over Q of the winner's value, counting bidder i's value when it wins and the rival's value when it loses. We search over the grid $\alpha \in \{0, 0.01, \ldots, 1\}$ together with the factors $\alpha = (r(q) + \delta)/v_i(q)$ that just outbid the rival on some prompt $q \in Q$, where $\delta = 10^{-6}$. For the *typical* equilibrium, the best response is the factor with the highest utility, and ties go to the factor closest to the current one. For the *adversarial* equilibrium, the best response is the factor with the lowest welfare among those whose utility is positive and at least 95% of the highest utility. If no factor meets both conditions, the bidder keeps its current bids.

Updates. All bidders start by bidding their values. In each round, we visit the bidders in turn, and for each of its bins a bidder moves its bids toward its best response to the current bids of the others,

$$b_i(q) \leftarrow (1-\gamma)\, b_i(q) + \gamma\, \hat{\alpha}\, v_i(q), \qquad q \in Q, \tag{14}$$

where $\hat{\alpha}$ is the best response. We use $\gamma = 1$ in the first round and $\gamma = 0.25$ afterwards; bids are then an exponentially weighted average of the bidder's past best responses. This averaging is what makes the procedure a form of fictitious play, and it helps avoid the cycling that undamped best responses can produce. The solver stops once no bid has moved by more than 10^{-3} for 10 consecutive rounds, or after 100 rounds. Note that starting from truthful bids means the dynamics mostly shade down, so the solver finds equilibria reachable by lowering bids from truthful bidding and may miss others.

Checks. The typical and adversarial runs share everything except the choice among best responses, so an adversarial equilibrium should never have higher welfare than a typical one on the same market. When it does, the solver has not settled, and we drop the market (for example, this removes at most 9% of markets in the filtering experiment of Section B.3). We also report the 0.1 percentile of welfare across markets in place of the minimum, so that a few unsettled markets do not drive the results.

D PROOFS FOR VALUE CONCENTRATION

Before presenting the proofs, we first isolate two core deviation arguments that were used throughout the concentration section. They are stated for mixed equilibria: every occurrence of B_i may include bidder i's private random seed, and the deviation may reproduce that seed before transforming the resulting bid. Throughout this appendix, we may write $M := \mathbb{E}_\mu[V_{(1)}]$ and $\kappa := 1/e$.

Lemma D.1. *If $V_{(2)} \geq c$ almost surely for a constant $c \geq 0$, then every Bayes–Nash equilibrium satisfies $R \geq c$ almost surely.*

Proof. Suppose first that $R \geq r$ almost surely for some $r < c$, and choose

$$r < b < \frac{c+r}{2}.$$

Bidder i considers the deviation that leaves its equilibrium bid unchanged unless $V_i \geq c$ and $B_i < b$, in which case it raises the bid to b. This uses only bidder i's type and private randomization.

Fix a realized value and bid profile with $R < b$. Let $A := \{i : V_i \geq c\}$, put $k := |A| \geq 2$, and let x_i be bidder i's equilibrium allocation probability at that profile. A deviating bidder in A wins for certain at bid b. Because $x_i > 0$ only when $B_i = R$, the sum of the high bidders' deviation gains is

$$\begin{aligned}\sum_{i \in A} [(V_i - b) - x_i(V_i - R)] &= \sum_{i \in A} [(1 - x_i)V_i + x_i R - b] \\ &\geq k(c - b) - x_A(c - R),\end{aligned}$$

where $x_A := \sum_{i \in A} x_i \leq 1$. Since $R < c$, this is at least

$$k(c - b) - (c - R) \geq c + R - 2b \geq c + r - 2b > 0.$$

On a profile with $R > b$, the transformed bid still loses whenever it was not already at least b, so the gain is zero. At $R = b$ the deviation can only add a high-value bidder to the top tie and has nonnegative gain. Therefore, if $\mathbb{P}(R < b) > 0$, the sum of the bidders' expected deviation gains is strictly positive. At least one bidder would have a profitable deviation, contradicting equilibrium. Hence $R \geq b$ almost surely for every $b \in (r, (c + r)/2)$. Taking a countable sequence of such bids increasing to $(c + r)/2$ gives $R \geq (c + r)/2$ almost surely.

Begin with $r_0 = 0$ and repeat the argument with $r_{t+1} = (c + r_t)/2$. Then $r_t = c(1 - 2^{-t}) \uparrow c$, so $R \geq c$ almost surely. □

Lemma D.2. *For every constant $c \geq 0$, every Bayes–Nash equilibrium satisfies*

$$\mathbb{E}[W] \geq (1 - \kappa)\mathbb{E}[(V_{(1)} - c)_+] + \mathbb{E}[\min\{R, c\}]. \tag{15}$$

Proof. Fix bidder i with realized value $v_i > c$. Let it deviate to $c + T_i$, where T_i has density

$$f_i(t) = \frac{1}{v_i - c - t}, \qquad 0 \leq t \leq (1 - \kappa)(v_i - c).$$

Because $\kappa = 1/e$, this density integrates to $\log(1/\kappa) = 1$. The deviation never exceeds v_i. If p_i is the highest rival bid, direct integration gives deviation utility at least

$$(1 - \kappa)(v_i - c)_+ - (p_i - c)_+. \tag{16}$$

When the expression on the right is negative, the inequality follows simply because the deviation itself has nonnegative utility.

Use this deviation for every bidder, and use the zero bid when $v_i \leq c$. For each realized profile, all deviation utilities are nonnegative. Select a measurable highest-value bidder i^*. The sum of deviation utilities is at least the lower bound in (16) for i^*, and $p_{i^*} \leq R$. Summing the equilibrium best-response inequalities therefore yields

$$\mathbb{E}\left[\sum_i u_i\right] \geq (1 - \kappa)\mathbb{E}[(V_{(1)} - c)_+] - \mathbb{E}[(R - c)_+].$$

Adding expected revenue and using $R - (R - c)_+ = \min\{R, c\}$ proves the result. Notice that every bidder's deviation depends only on its own value; no independence assumption is used. □

Proof of Proposition 1. Set $c := \alpha M$. Lemma D.1 gives $R \geq c$ almost surely, and the hypothesis also gives $V_{(1)} \geq c$ almost surely. Therefore Lemma D.2 yields

$$\mathbb{E}[W] \geq (1 - \kappa)(M - c) + c = (1 - \kappa)M + \kappa c.$$

Divide by M. □

Proof of Lemma 2. Set $c := v_{\text{floor}}$. The claim is immediate if $c = 0$. Otherwise fix $p > (2 - m)/(m - 1) > -1$, and draw a target bid $b \in [0, c]$ with density $f_p(b) = \frac{p+1}{c^{p+1}}(c - b)^p$. Each bidder reproduces its equilibrium bid and raises it to b exactly when $V_i \geq c$ and $B_i < b$. These are feasible unilateral deviations. Linearity permits coupling their target draws when summing expected gains.

Let $A = \{i : V_i \geq c\}$, $k = |A|$, and $x_A = \sum_{i \in A} x_i$, where x_i is the on-path allocation probability. On $R < b$, the sum of deviation gains is at least $k(c - b) - x_A(c - R)$. On $G = \{V_{(m)} \geq c\}$, this is at least $m(c - b) - (c - R)$. On G^c with $k \geq 1$, it is at least $(c - b) - (c - R)$. If $k = 0$, the gain is zero and the same latter lower bound still holds. On $R \geq b$, gains are nonnegative.

Put $D = (c - R)_+$ and $a = m(p + 1) - (p + 2) > 0$. When $R < c$, integrate the preceding bounds over $b \in [R, c]$. The resulting lower bounds are

$$\frac{a}{p+2}\frac{D^{p+2}}{c^{p+1}} \quad \text{on } G, \qquad -\frac{1}{p+2}\frac{D^{p+2}}{c^{p+1}} \quad \text{on } G^c.$$

When $R \geq c$, the integrated contribution is zero. Equilibrium makes the sum of expected deviation gains nonpositive, so $a\,\mathbb{E}[D^{p+2}\mathbf{1}_G] \leq \mathbb{E}[D^{p+2}\mathbf{1}_{G^c}]$. Since $D \leq c$ and $\mathbb{P}(G^c) \leq \epsilon$,

$$\mathbb{E}[D^{p+2}] \leq \frac{a+1}{a}\,\epsilon c^{p+2} = \epsilon c^{p+2}\frac{(m-1)(p+1)}{m(p+1) - (p+2)}.$$

The exponent $p + 2 > 1$, so the moment inequality $\mathbb{E}[D] \leq (\mathbb{E}[D^{p+2}])^{1/(p+2)}$ proves $\mathbb{E}[(c - R)_+] \leq c\,\mathcal{R}_m(\epsilon)$ after taking the infimum over p. □

Proof of Theorem 4. Fix an admissible triple (m, α, ϵ) and set $c := \alpha M$. Lemma 2 and Lemma D.2 give

$$\begin{aligned}\mathbb{E}[W] &\geq (1 - \kappa)\mathbb{E}[(V_{(1)} - c)_+] + c - c\mathcal{R}_m(\epsilon)\\ &\geq (1 - \kappa)(M - c) + c - c\mathcal{R}_m(\epsilon)\\ &= (1 - \kappa)M + c[\kappa - \mathcal{R}_m(\epsilon)].\end{aligned}$$

The anchor $c = 0$ independently gives $(1 - \kappa)M$, so the bracket may be replaced by its positive part. Divide by M to obtain (6). Since $\mathcal{R}_m$ is nondecreasing, the smallest admissible failure probability is optimal; namely,

$$\epsilon = F_m(\alpha M).$$

This gives the second line of Equation (6).

More generally, fix m and $c \geq 0$ and apply Lemma 2 with $\epsilon = F_m(c)$ in Lemma D.2:

$$\mathbb{E}[W] \geq (1 - \kappa)\mathbb{E}[(V_{(1)} - c)_+] + c[1 - \mathcal{R}_m(F_m(c))].$$

This auxiliary inequality holds for every $c \geq 0$ and $m \in \{2, \ldots, n\}$. □

Properties of $\mathcal{R}_m$. For each admissible p, the objective in Equation (5) is nondecreasing in ϵ. Taking an infimum preserves this order, so $\mathcal{R}_m$ is nondecreasing in its argument. For admissible p, the multiplier inside (5) can be written as

$$\frac{(m-1)(p+1)}{m(p+1) - (p+2)} = 1 + \frac{1}{m(p+1) - (p+2)}.$$

Observe that this is decreasing in m, and the admissible domain expands as m rises. When $\epsilon \in (0, 1)$, the objective tends to infinity at the lower endpoint of its domain and to one as $p \to \infty$, while it is below one for all sufficiently large finite p. Its minimum is therefore attained at an interior finite point. Evaluating the $(m + 1)$ objective at an optimizer for m makes the pointwise inequality strict and proves that $\mathcal{R}_{m+1}(\epsilon) < \mathcal{R}_m(\epsilon)$. Directly from the definition, $\mathcal{R}_m(0) = 0$ and $\mathcal{R}_m(1) = 1$.

For $0 < \epsilon < 1$, write $L = \log(1/\epsilon)$ and set $t = (m - 1)(p + 1) - 1 > 0$. Then

$$\mathcal{R}_m(\epsilon) = \inf_{t>0} \left[\epsilon(1 + 1/t)\right]^{(m-1)/(t+m)}. \tag{17}$$

The logarithm of the objective is $g(t) = \frac{m-1}{t+m}\left[-L + \log(1 + 1/t)\right]$, and its derivative has the sign of $L - H_m(t)$, where $H_m(t) = \log(1 + 1/t) + \frac{t+m}{t(t+1)}$. The function H_m is strictly decreasing

from infinity to zero; indeed $H'_m(t) = -\frac{2t^2+(2m+1)t+m}{t^2(t+1)^2} < 0$. Thus the minimum is unique and satisfies $H_m(t_*) = L$. As $L \to \infty$ with m fixed, $H_m(t) = \frac{m}{t} + \log(1/t) + 1 - m + O(t)$, so $t_* = \frac{m}{L - \log L + \log m + m - 1 + o(1)}$. Substituting in $g(t_*)$ gives $\log \mathcal{R}_m(\epsilon) = \frac{m-1}{m}\big[-L + \log L - \log m + 1\big] + o(1)$, which gives the small-$\epsilon$ asymptotics for fixed m in Remark 3.

There is also an explicit bound at every m. Choose $t = m/L$, and put $A = \epsilon(1 + L/m)$. Since $0 < A < 1$,

$$\begin{aligned}\mathcal{R}_m(\epsilon) &\le A^{(m-1)L/[m(L+1)]} \\ &= A^{(m-1)/m} \exp\left\{\frac{m-1}{m}\frac{L - \log(1+L/m)}{L+1}\right\} \\ &\le (eA)^{(m-1)/m}.\end{aligned}$$

Consequently

$$\mathcal{R}_m(\epsilon) \le \min\left\{1, \left[e\epsilon\left(1 + \frac{\log(1/\epsilon)}{m}\right)\right]^{(m-1)/m}\right\}. \tag{18}$$

The $m = 2$ case gives the shortfall expression given by $\mathcal{R}_2(\epsilon)$ in Remark 3, whereas the $m > 2$ bound follows directly from $p = 0$. The bound follows by taking $p \to \infty$ in (5).

Proof of Corollary A.1. Fix α and write $c = \alpha\mathbb{E}[Y]$, $a = a_\alpha$, and $b = b_\alpha$. Both a and b are positive. For $0 \le \lambda < \frac{\mathbb{E}[U]-c}{\mathbb{E}[Y]-c} \le 1$, the ordering $U \le Y$ gives $\{U \le c\} \subseteq \{U - \lambda Y \le (1-\lambda)c\}$. The mean of $U - \lambda Y$ exceeds $(1-\lambda)c$ by $\mathbb{E}[U] - c - \lambda(\mathbb{E}[Y] - c) > 0$, while $\mathrm{Var}(U - \lambda Y) \le \sigma_U^2 + \lambda^2\sigma_Y^2 - 2\lambda\theta\sigma_U\sigma_Y$. Cantelli's inequality therefore implies, with $r = \lambda\sigma_Y/\sigma_U$, $\mathbb{P}(U \le c) \le \left(1 + \frac{(a-rb)^2}{1+r^2-2\theta r}\right)^{-1}$. Choose

$$r = \begin{cases} 0, & \theta a \le b, \\ \dfrac{\theta a - b}{a - \theta b}, & \theta a > b. \end{cases}$$

In the second case $a - \theta b > 0$ and $0 < r < a/b$; thus in both cases the corresponding λ is admissible. Substitution gives

$$\frac{(a-rb)^2}{1+r^2-2\theta r} = a^2 + \frac{[\theta a - b]_+^2}{1-\theta^2}.$$

This proves the probability bound. Applying Theorem 4 with $m = 2$ and $\epsilon = \epsilon_\theta(\alpha)$ proves the welfare guarantee. □

Proof of Corollary A.2. For $M_x > 0$, apply Theorem 4 under μ_x with threshold αM_x and failure probability $F_m^x(\alpha M_x)$, for the fixed pair (α, m). This gives

$$\mathbb{E}[W \mid X = x] \ge \left(1 - \frac{1 - \delta(\alpha, m \mid x)}{e}\right) M_x.$$

The inequality is also valid when $M_x = 0$. Taking expectations over X proves the claim. □

D.1 Convexity of the floor bonus

Lemma D.3. *For an integer $m \ge 2$, define*

$$\mathcal{R}_m(\epsilon) := \inf_{p > \max\{-1, (2-m)/(m-1)\}} \left[\epsilon \frac{(m-1)(p+1)}{m(p+1) - (p+2)}\right]^{1/(p+2)}, \qquad \epsilon \in [0,1].$$

Then $\mathcal{R}_m$ is concave on $[0,1]$. Consequently, writing $[z]_+ := \max\{z, 0\}$, the function

$$g_m(\epsilon) := [1 - e\mathcal{R}_m(\epsilon)]_+$$

is convex on $[0,1]$, as is $\epsilon \mapsto \alpha g_m(\epsilon)$ for every fixed $\alpha \ge 0$.

Proof. For each admissible p, put

$$A_{m,p} := \frac{(m-1)(p+1)}{m(p+1)-(p+2)}, \qquad q_p := \frac{1}{p+2}.$$

The restriction on p ensures that $A_{m,p} > 0$ and $0 < q_p < 1$. Thus $f_{m,p}(\epsilon) := A_{m,p}^{q_p} \epsilon^{q_p}$ is concave on $[0,1]$. For $\epsilon_0, \epsilon_1 \in [0,1]$ and $t \in [0,1]$, every admissible p therefore satisfies

$$\begin{aligned} f_{m,p}\big(t\epsilon_0 + (1-t)\epsilon_1\big) &\geq t f_{m,p}(\epsilon_0) + (1-t) f_{m,p}(\epsilon_1) \\ &\geq t\mathcal{R}_m(\epsilon_0) + (1-t)\mathcal{R}_m(\epsilon_1). \end{aligned}$$

Taking the infimum over p on the left proves concavity of $\mathcal{R}_m$. Hence $1 - e\mathcal{R}_m$ is convex. Its pointwise maximum with the constant zero is convex, proving the assertion for g_m. Multiplication by $\alpha \geq 0$ preserves convexity. □

D.2 Platform reserve bound

We prove Equation (10). The platform announces $V_0 = v_0(X)$ before bidding and always bids it, and the other bidders play any Bayes–Nash equilibrium of the induced auction. Write $V_{(1)} := \max\{V_0, V_1, \ldots, V_n\}$.

Proof. Fix $X = x$ and write $v_0 = v_0(x)$. Since the platform bids v_0, the winning bid satisfies $R \geq v_0$. Run the argument of Lemma D.2 under the conditional law given $X = x$ with $c = v_0$, using deviations only for the strategic bidders; these deviations may depend on x because x is public. On profiles where a strategic bidder has the top value, the argument is unchanged, since the highest rival bid is still at most R. On profiles where the platform has the top value, $(V_{(1)} - c)_+ = 0$, and the required inequality holds because deviation utilities are nonnegative. The platform's utility is always zero, so welfare still equals total bidder utility plus revenue. Hence, using $V_{(1)} \geq v_0$ and $R \geq v_0$,

$$\mathbb{E}[W \mid X = x] \geq (1-\kappa)\left(\mathbb{E}[V_{(1)} \mid X = x] - v_0\right) + v_0 = (1-\kappa)\,\mathbb{E}[V_{(1)} \mid X = x] + \kappa v_0.$$

Taking expectations over X and dividing by $\mathbb{E}[V_{(1)}]$ gives Equation (10). □

D.3 Robustness of the $\left(1 - \frac{1}{e}\right)$ bound

Before moving to the proof of the weak dependence result (Theorem 5), we first show that corresponding robustness results are immediate for the smoothness-based $\left(1 - \frac{1}{e}\right)$ guarantee for arbitrarily correlated values (Roughgarden et al., 2017). Fix $X = x$, and suppose $0 \leq V_i \leq \bar{v}(x)$. If, under ν_x, no bidder can improve its ex ante expected utility by more than ε through a unilateral change of strategy, the proof of Lemma D.2 with $c = 0$ gives

$$\mathbb{E}_{\nu_x}[W] \geq (1 - 1/e)\mathbb{E}_{\nu_x}[V_{(1)}] - n\varepsilon.$$

Indeed, each of the n best-response inequalities incurs at most ε, and adding expected revenue gives the displayed bound. This is the finite-regret version of the private-deviation smoothness argument (Roughgarden et al., 2017, Definition 4.5, Eq. (18)). Now keep a no-overbidding equilibrium under μ_x fixed and evaluate it under ν_x. Both its utility and the utility from any no-overbidding deviation lie in $[0, \bar{v}(x)]$, so changing the law by total-variation distance $\tau(x)$ gives $\varepsilon \leq 2\bar{v}(x)\tau(x)$. The resulting welfare error under ν_x is therefore at most $2n\bar{v}(x)\tau(x)$; transferring welfare and optimal-welfare expectations back to μ_x adds at most $(2 - 1/e)\bar{v}(x)\tau(x)$ to this error.

In the next section, we establish the corresponding stability at the sharper coefficient $1 - 1/e^2$, for which this smoothness argument does not supply a welfare bound.

E Proofs and additional results for weak dependence

In order to prove the qualitative Theorem 5, we would like to compare the actual value law μ_x with the independent law ν_x having the same marginals. The core difficulty is of course that equilibrium strategies for μ_x need not remain optimal under ν_x, so we cannot simply apply the independent-values welfare bound after changing the distribution. Indeed, bids that were distinct can become tied in a

limit, and this can change who wins. The task, thus, is to show that an independent-values auction still has nearly the Jin–Lu welfare guarantee when no bidder can gain much by changing its strategy. We build the argument in the following sequence of steps: Lemma E.4 accounts for deviations that pass to the limit. Lemma E.5 then provides an exact equilibrium as regret vanishes. Using Lemmas E.4 and Lemma E.5, we show Proposition E.2 that establishes a uniform stability result i.e. that small regret gives a small welfare shortfall. Using these we are able to establish Theorem 5 that gives the guarantee under the dependent law.

E.1 The welfare modulus

We first define basic terms, and a measure of the welfare defect when there is a deviation from the independent values case.

Write $c_{\mathrm{JL}} := 1 - 1/e^2$ and fix the public information $X = x$. For an independent prior ν, let ν_i be its marginal law of V_i. For a mixed strategy profile s, let $\widehat{s}_i$ denote an alternative strategy for bidder i, while s_{-i} denotes the fixed profile of rival strategies. Define bidder i's ex ante regret by

$$\mathrm{Reg}_i^{\nu}(s) := \sup_{\widehat{s}_i} \mathbb{E}_{\nu}\left[u_i\big(V_i, \widehat{s}_i(V_i), s_{-i}(V_{-i})\big) - u_i\big(V_i, s_i(V_i), s_{-i}(V_{-i})\big)\right], \tag{19}$$

Here u_i is first-price utility: $v_i - b_i$ on winning and zero otherwise. Deviations may be measurable and mixed and may use every bid in $[0, V_i]$. Write $W(s)$ for the welfare generated by s and $x_+ := \max\{x, 0\}$.

The regret asks how much a bidder could improve its expected utility while its rivals keep their strategies. The alternative $\widehat{s}_i$ is an entire bidding strategy: it may choose a different bid for each value, rather than use one constant bid for every value. The expectation is taken over the bidder's own value as well as its rivals' values and all private randomization. Thus small ex ante regret means that the average available gain is small; it does not require that gain to be small at every individual value.

For $r \geq 0$, define the worst additive welfare defect among normalized independent-values auctions by

$$\delta_n(r) := \sup\Big\{ \big[c_{\mathrm{JL}}\mathbb{E}_{\nu}[V_{(1)}] - \mathbb{E}_{\nu}[W(s)]\big]_+ : \\ \nu = \textstyle\bigotimes_i \nu_i \text{ on } [0,1]^n, \quad s \text{ is mixed and no-overbidding}, \quad \max_i \mathrm{Reg}_i^{\nu}(s) \leq r \Big\}. \tag{20}$$

The supremum also ranges over publicly announced measurable randomized bid-only tie rules. In each regret comparison the tie rule is held fixed.

Remark E.1. $\big[c_{\mathrm{JL}}\mathbb{E}_{\nu}[V_{(1)}] - \mathbb{E}_{\nu}[W(s)]\big]_+$ *is the amount by which welfare falls below the Jin–Lu guarantee. The supremum measures how large that shortfall could be in any normalized independent-values auction in which every bidder's regret is at most* r*. Values lie in* $[0,1]$*, so both* r *and* $\delta_n(r)$ *are measured on a common scale.* $\delta_n(r) \to 0$ *means that, for any desired accuracy, a sufficiently small regret guarantees that accuracy in* every *such auction with the given number of bidders.*

We first state the core proposition that is key to the proof of Theorem 5.

Proposition E.2 (Uniform stability under vanishing regret)**.** *For every fixed* $n \geq 2$*, the function* δ_n *in* (20) *is nondecreasing, takes values in* $[0,1]$*, satisfies* $\delta_n(0) = 0$*, and obeys*

$$\delta_n(r) \longrightarrow 0 \qquad \text{as } r \downarrow 0. \tag{21}$$

This conclusion is uniform over independent value laws on $[0,1]^n$*, mixed no-overbidding strategies, bid atoms, and publicly announced randomized bid-only tie-breaking rules.*

Define

$$\eta_n^{\mathrm{cont}}(t) := (1 + c_{\mathrm{JL}})t + \delta_n(2t) = \left(2 - \frac{1}{e^2}\right) t + \delta_n(2t). \tag{22}$$

For $t \in [0,1]$, this is a nondecreasing modulus with $\eta_n^{\mathrm{cont}}(0) = 0$ and $\eta_n^{\mathrm{cont}}(t) \to 0$ as $t \downarrow 0$.

Remark E.3. $\delta_n(2t)$ *accounts for strategies that acquire regret at most* $2t$ *when the value law is changed. It is thus a statement about incentives. The second term i.e.* $(1 + c_{\mathrm{JL}})t$ *accounts for moving the welfare expectation and the expected optimum between the two laws. Neither term requires the strategies themselves to vary continuously with the prior.*

E.1.1 Proof of Proposition E.2

A proof for Proposition E.2 has two core obstacles, which we motivate and prove one after the other. The first problem is that even small changes in bids can create or remove ties, so expected utility might not converge even when the underlying distributions converge. Ties require special care. For example, bids t and $t + 1/k$ have a unique winner for every k, but both converge to t. The limiting tie rule may select a different bidder, so convergence of bids alone does not guarantee convergence of utilities or welfare. Lemma E.4 avoids this problem for deviations: it chooses bids at which the limiting rival distributions put no positive probability. At those bids, whether equality counts as winning is immaterial in the limit. We later handle ties in the strategies actually being played by recording the winner as well as the bids.

For a distribution function H, write $H(t-) := \lim_{z \uparrow t} H(z)$.

Lemma E.4 (Approximation by safe deviations)**.** *Here $\Rightarrow$ denotes weak convergence of probability measures. Let $F^k \Rightarrow F$ be value laws on $[0,1]$, independent of the rivals' independent bid draws. For each rival $j \neq i$, let $H_j^k \Rightarrow H_j$ be bid laws on $[0,1]$. Put*

$$G_i(t) := \prod_{j \neq i} H_j(t), \qquad M_i(v) := \sup_{0 \leq t \leq v} (v - t) G_i(t).$$

In the kth first-price auction, let $y_i^k(t)$ be bidder i's interim winning probability from the constant bid t, under an arbitrary bid-only tie rule. Then, for every $\varepsilon > 0$, there is a measurable no-overbidding deviation $d : [0,1] \to [0,1]$ such that

$$\liminf_{k \to \infty} \int (v - d(v)) y_i^k(d(v))\, dF^k(v) \geq \int M_i(v)\, dF(v) - \varepsilon. \tag{23}$$

Intuition behind deviation lemma and proof. $G_i(t)$ is the probability that every rival bids at most t. Consequently, $(v-t)G_i(t)$ is the utility from bidding t under the most favorable possible treatment of ties, and $M_i(v)$ is the best such utility. This is a useful upper benchmark even when the actual tie rule is less favorable. A bid just above t beats every rival bid equal to t, while increasing the payment by arbitrarily little. Provided $t < v$, we can make this increase without overbidding; when $t = v$, the benchmark payoff is zero. We also need one measurable deviation that works across values. The proof therefore divides $[0,1]$ into short intervals and selects a nearly optimal bid using the lower endpoint of each interval. That bid is feasible for every value in the interval. The 1-Lipschitz property of M_i says that using the lower endpoint loses at most the interval's width in utility. Each selected bid can be moved slightly to avoid the countably many atoms of the limiting rival bid laws. Winning probabilities then converge at these selected bids, and choosing interval boundaries without value atoms allows the expected utilities to converge as well. The partition is chosen for a fixed accuracy and kept fixed as k grows; only afterwards is the accuracy sent to zero.

Proof. The function M_i is 1-Lipschitz and $M_i(0) = 0$: it is a supremum of functions of v with slopes in $[0,1]$, and allowing additional bids when v increases changes the supremum by at most the increase in v. Choose a partition

$$0 = a_0 < a_1 < \cdots < a_m = 1$$

with mesh at most ζ, where every interior endpoint is a continuity point of F. Such a partition exists because F has only countably many atoms.

Fix $\ell \geq 2$. The union of the atoms of the finitely many limiting rival bid laws H_j is countable. Right continuity permits us to choose $t_\ell < a_{\ell-1}$ outside that union such that

$$(a_{\ell-1} - t_\ell) G_i(t_\ell) \geq M_i(a_{\ell-1}) - \zeta. \tag{24}$$

To see this, start from an approximately optimal bid strictly below $a_{\ell-1}$ and approach it from above through common continuity points while remaining below $a_{\ell-1}$. If the only candidate is $t = a_{\ell-1}$, its payoff is zero and the claim is immediate.

Define $d(v) = 0$ on $[0, a_1]$ and $d(v) = t_\ell$ on $(a_{\ell-1}, a_\ell]$. Values at partition endpoints may be assigned to either adjacent interval. This deviation never overbids. A bid t_ℓ wins whenever every rival bid is strictly below it, regardless of the tie rule, so

$$y_i^k(t_\ell) \geq \prod_{j \neq i} H_j^k(t_\ell -).$$

Because t_ℓ is not an atom of any H_j, weak convergence gives $H_j^k(t_\ell-) \to H_j(t_\ell)$. For completeness, the upper bound follows from $H_j^k(t_\ell-) \le H_j^k(t_\ell) \to H_j(t_\ell)$; for any continuity point $z < t_\ell$, the lower bound follows from $H_j^k(t_\ell-) \ge H_j^k(z) \to H_j(z)$, followed by $z \uparrow t_\ell$. The interior endpoints are continuity points of F. The intervals are continuity sets relative to $[0,1]$, including the last interval ending at 1, so the integrals converge even if F has an atom at 1.

For $v \in (a_{\ell-1}, a_\ell]$, Equation (24) and the Lipschitz property imply

$$(v - t_\ell)G_i(t_\ell) \ge M_i(v) - 2\zeta.$$

On the first interval, $0 \le M_i(v) \le \zeta$, while the zero bid has nonnegative utility. The left side of (23) is therefore at least $\int M_i\, dF - 3\zeta$. Let $\zeta < \varepsilon/3$. □

Lemma E.5. *Fix n. Consider a sequence of independent-values first-price auctions on $[0,1]^n$. In auction k, let s^k be a mixed no-overbidding profile, let π^k be a publicly announced randomized bid-only tie rule, and suppose*

$$\max_i \mathrm{Reg}_i^{\nu^k}(s^k) \longrightarrow 0.$$

There are a subsequence, an independent prior ν, a mixed no-overbidding profile s, and a publicly announced randomized bid-only tie rule π such that $s \in \mathrm{BNE}(\nu)$ and, along the subsequence,

$$\mathbb{E}_{\nu^k}[V_{(1)}] \longrightarrow \mathbb{E}_\nu[V_{(1)}], \qquad \mathbb{E}_{\nu^k}[W(s^k)] \longrightarrow \mathbb{E}_\nu[W(s)]. \tag{25}$$

Since the argument involved is delicate, we provide some intuition behind the lemma and sketch the proof idea first.

Intuition and proof idea A sequence of bidding functions need not have a useful pointwise limit, especially when bidders mix. Instead, the proof keeps each bidder's joint distribution of its value and bid. Since $0 \le b \le v \le 1$, these pairs remain in a compact set, so a subsequence of their distributions converges. The limit again describes a value distribution and, conditional on each value, a distribution of bids. Independence across bidders is preserved by taking the product of these limiting distributions. The winner is recorded separately because the limit of the bids does not tell us how a newly created tie was resolved. Once the winner I is part of the outcome, welfare is simply v_I and bidder utility is $(v_i - b_i)\mathbf{1}_{\{I=i\}}$; these quantities pass to the limit. At this stage the limiting winner is known to have a highest bid, but its selection has not yet been shown to depend only on bids. The rest of the proof establishes that the limiting utilities and welfare can indeed be obtained with an announced bid-only tie rule.

Proof. Let $D := \{(v,b) \in [0,1]^2 : 0 \le b \le v\}$. For bidder i in auction k, let q_i^k be the law of (V_i^k, B_i^k) on D. Independence of values and private randomizations implies that the law of all value–bid pairs is $q^k := \bigotimes_i q_i^k$. Probability measures on the compact metric space D are weakly compact. After taking a common subsequence,

$$q_i^k \Rightarrow q_i, \qquad q^k \Rightarrow q := \bigotimes_i q_i. \tag{26}$$

Write F_i and H_i for the value and bid marginals of q_i. Disintegrate $q_i(dv, db) = F_i(dv)s_i(db \mid v)$. The product of the value marginals is the limiting independent prior $\nu := \bigotimes_i F_i$, and closedness of D preserves no overbidding.

Let $I^k \in [n]$ be the winner, including the randomization in π^k, and let λ^k be the law of all value–bid pairs together with I^k. Pass to a further subsequence for which $\lambda^k \Rightarrow \lambda$. Its marginal on D^n is q, and

$$\lambda\{b_I = \max_j b_j\} = 1, \tag{27}$$

because the displayed event is closed. With $[n]$ carrying the discrete topology, the functions

$$(v,b,I) \longmapsto (v_i - b_i)\mathbf{1}_{\{I=i\}}, \qquad (v,b,I) \longmapsto v_I$$

are continuous. Hence, writing

$$U_i^* := \int (v_i - b_i)\mathbf{1}_{\{I=i\}}\, d\lambda, \qquad W^* := \int v_I\, d\lambda,$$

the on-path utilities converge to U_i^* and auction welfare converges to W^*. Also,

$$\mathbb{E}_{\nu^k}[V_{(1)}] \longrightarrow \int \max_i v_i \, dq = \mathbb{E}_\nu[V_{(1)}], \tag{28}$$

because the maximum is continuous.

For the limiting bid marginals, put

$$G_i(t) := \prod_{j \neq i} H_j(t), \qquad M_i(v) := \sup_{0 \le t \le v} (v - t) G_i(t).$$

Apply Lemma E.4 to bidder i. The regret inequality in auction k, followed by the lemma and vanishing regret, gives

$$U_i^* \ge \int M_i(v) \, dF_i(v). \tag{29}$$

There is a matching upper bound that uses only feasibility of the limiting winner. On the event $I = i$, every rival bid is at most b_i. Since $v_i - b_i \ge 0$ and λ has marginal q,

$$\begin{aligned} U_i^* &\le \int (v_i - b_i) \mathbf{1}_{\{b_j \le b_i \ \forall j \neq i\}} \, dq \\ &= \int (v - b) G_i(b) \, dq_i(v, b) \le \int M_i(v) \, dF_i(v). \end{aligned} \tag{30}$$

Every inequality in (29)– (30) is therefore an equality. In particular,

$$(v - b) G_i(b) = M_i(v) \quad \text{for } q_i\text{-almost every } (v, b), \tag{31}$$

and

$$(v_i - b_i) \left(\mathbf{1}_{\{b_j \le b_i \ \forall j \neq i\}} - \mathbf{1}_{\{I = i\}} \right) = 0 \quad \lambda\text{-almost surely.} \tag{32}$$

Thus a bidder tied at the top with strictly positive surplus must receive the item in the limiting outcome.

We now construct an ordinary bid-only tie rule. A positive-probability tie at a level t under a product bid law requires at least two bid marginals to have atoms at t; there are only countably many such levels. At any level that can be the top bid with positive probability, at most one bidder i can satisfy

$$q_i\{B_i = t, \ V_i > t\} > 0. \tag{33}$$

Otherwise two such bidders i and j, together with the event that all other bids are at most t, would with positive probability both tie at the top and both have positive surplus. Equation (32) would require both $I = i$ and $I = j$, an impossibility.

At each positive-probability top-tie level t, call the unique bidder in (33), if it exists, $h(t)$. At a tie for the highest bid t, let the new rule π award the item to $h(t)$ if that bidder is among those tied; otherwise award it to the lowest-index tied bidder. Away from ties, award it to the unique highest bidder. The exceptional levels are countable, so this is a Borel bid-only rule.

This repair preserves welfare. At a top tie, every tied bidder other than $h(t)$ has value exactly t almost surely. If $h(t)$ is tied and has value above t, Equation (32) says that λ already selects $h(t)$. In every remaining case all tied bidders have value t. Consequently,

$$\mathbb{E}_\nu[W(s)] = W^*. \tag{34}$$

Finally, under π, a bidder with an on-path bid b and positive surplus wins every positive-probability tie at b. Its interim allocation is therefore $G_i(b)$ for q_i-almost every (v, b) with $v > b$. When $v = b$, its payoff is zero and (31) gives $M_i(v) = 0$. Thus the on-path expected utility of almost every type v is $M_i(v)$. Any deviation to t has interim allocation at most $G_i(t)$ and utility at most

$$(v - t)_+ G_i(t) \le M_i(v).$$

The same holds for mixed deviations. Therefore $s \in \mathrm{BNE}(\nu)$ under π. Equations (28) and (34) give (25). □

Remark E.6. *Suppose arbitrarily small regret could coexist with a fixed positive shortfall below the Jin–Lu guarantee. We could then choose a sequence of such auctions with regret tending to zero. Lemma E.5 would turn a subsequence into an exact independent-values equilibrium while preserving the limiting welfare and expected optimum. The positive shortfall would survive, contradicting the Jin–Lu theorem. This proves uniformity because the hypothetical counterexamples were allowed to change their value laws, strategies, and tie rules at every step. The argument gives no explicit rate of convergence, and does not claim that nearby auctions have nearby equilibria.*

With all the ingredients in place, we can finally state the main proof of this section.

Proof of Proposition E.2

Proof. Monotonicity and $0 \le \delta_n(r) \le 1$ are immediate, and the Jin–Lu theorem gives $\delta_n(0) = 0$. If $\delta_n(r) \not\to 0$, some $\varepsilon > 0$ and sequence $r_k \downarrow 0$ admit independent priors, mixed no-overbidding profiles, and bid-only tie rules with maximal regret at most r_k and

$$c_{\mathrm{JL}}\mathbb{E}[V_{(1)}] - \mathbb{E}[W] \ge \varepsilon.$$

Lemma E.5 produces an exact mixed Bayes–Nash equilibrium of an independent-values first-price auction to which both welfare terms converge. The limiting defect remains at least ε, contrary to the Jin–Lu bound. Hence (21) holds. □

E.2 Proof of the main qualitative theorem (Theorem 5).

Proof of Theorem 5. Fix x, set $\mu := \mu_x$, $\nu := \nu_x$, $\tau := \tau(x)$, and $\overline{v} := \overline{v}(x)$. The result is immediate if $\overline{v} = 0$; otherwise run the equilibrium strategy kernels from μ under the product prior ν on their original scale. The one-dimensional value marginals, and hence all permissible own-type deviations, are unchanged.

For a fixed bidder, both on-path utility and the utility from any fixed no-overbidding deviation lie in $[0, \overline{v}]$. Passing either expectation between μ and ν costs at most $\overline{v}\tau$. Exact equilibrium under μ therefore gives

$$\max_i \mathrm{Reg}_i^{\nu}(s) \le 2\overline{v}\tau.$$

Dividing values, bids, and utilities by $\overline{v}$ makes the normalized maximal regret at most 2τ. Proposition E.2, followed by rescaling, implies

$$\mathbb{E}_{\nu}[W] \ge c_{\mathrm{JL}}\mathbb{E}_{\nu}[V_{(1)}] - \overline{v}\delta_n(2\tau).$$

Both W and $V_{(1)}$ lie in $[0, \overline{v}]$. Transferring them back to μ costs at most $\overline{v}\tau$ and $c_{\mathrm{JL}}\overline{v}\tau$, respectively. This proves (8) with the modulus in (22). Applying this argument under the conditional law for almost every public state proves the theorem. □