Optimal Type-Dependent Liquid Welfare Guarantees for Autobidding Agents with Budgets*

Riccardo Colini-Baldeschi† Sophie Klumper‡ Twan Kroll‡ Stefano Leonardi§
Guido Schäfer‡ Artem Tsikiridis¶

Abstract. Online advertising systems have recently transitioned to autobidding, allowing advertisers to delegate bidding decisions to automated agents. Each advertiser directs their agent to optimize an objective function subject to *return-on-investment (ROI)* and *budget constraints*. Given their practical relevance, this shift has spurred a surge of research on the *liquid welfare price of anarchy (POA)* of fundamental auction formats under autobidding, most notably *simultaneous first-price auctions (FPA)*. One of the main challenges is to understand the efficiency of FPA in the presence of *heterogeneous* agent types. We introduce a *type-dependent smoothness framework* that enables a unified analysis of the POA in such complex autobidding environments. In our approach, we derive type-dependent smoothness parameters which we carefully balance to obtain POA bounds. This balancing gives rise to a *POA-revealing mathematical program*, which we use to determine tight bounds on the POA of coarse correlated equilibria (CCE). Our framework is versatile enough to handle heterogeneous agent types and extends to the general class of fractionally subadditive valuations. Additionally, we develop a novel reduction technique that transforms budget-constrained agents into budget-unconstrained ones. Combining this reduction technique with our smoothness framework enables us to derive tight bounds on the POA of CCE in the general hybrid agent model with both ROI and budget constraints. Among other results, our bounds uncover an intriguing threshold phenomenon showing that the POA depends intricately on the smallest and largest agent types. We also extend our study to FPAs with reserve prices, which can be interpreted as predictions of agents' values, to further improve efficiency guarantees.

1 Introduction.

Over the past decade, online advertising systems have undergone a major shift with the emergence of autobidding. This shift allows advertisers to delegate complex bidding decisions to automated agents that take various factors into account such as ad performance, campaign constraints, and market dynamics. As a result, advertisers can manage their campaigns more efficiently and aim to maximize their return on investment. Autobidding is now the dominant paradigm: over 80% of online ad traffic is managed by autobidding agents [22]. This widespread adoption has important implications for the behavior of advertisers, publishers, and ad exchanges.

In the autobidding world, advertisers specify high-level constraints for their campaigns—most notably *return-on-investment (ROI) constraints* and *budget constraints*. Basically, the ROI constraint caps the cost per conversion or impression, while the budget constraint limits the total spend for a given campaign. Unlike the traditional view, where agents have intrinsic values for outcomes, autobidding agents operate under ROI and budget constraints reflecting performance goals and financial limitations. In addition to these constraints, agents may differ in the objective they seek to optimize. The autobidding literature typically considers two types: *value maximizers*, who aim to maximize outcome value subject to budget and ROI constraints, and *utility maximizers*, who aim to maximize value minus payment. These two types represent the extremes of a spectrum capturing agents' trade-offs between value and payment. In this work, we allow agents' types to lie anywhere along this spectrum, modeling diverse behaviors in a unified way.

---

*The full version of the paper [17] can be accessed at https://arxiv.org/abs/2506.20908.

†Meta, UK (rickuz@meta.com).
‡Centrum Wiskunde & Informatica (CWI) and Institute for Logic, Language and Computation (ILLC), University of Amsterdam, the Netherlands ({s.j.klumper,t.kroll,g.schaefer}@cwi.nl).
§Sapienza University of Rome, Italy (leonardi@diag.uniroma1.it).
¶Technical University of Munich, Germany (artem.tsikiridis@tum.de).

Alongside the rise of autobidding, the online advertising industry has also undergone a major shift in auction formats—most notably, the move from second-price to *first-price auctions (FPA)*. This transition has been especially pronounced in display ad markets, culminating in Google Ad Exchange's adoption of first-price auctions in 2019 [40]. The combination of constraint-driven autobidding agents and first-price payment schemes raises fundamental questions about the performance of FPA in this new environment. This paper is driven by one central question: *What is the efficiency of simultaneous first-price auctions in the presence of heterogeneous autobidding agents?*

The study of the inefficiency of equilibria for simultaneous first-price auctions in the autobidding setting was initiated by Liaw et al. [35]. To evaluate the efficiency of FPAs in the autobidding setting, we study the *Price of Anarchy (POA)* ([33]), which compares the optimal liquid welfare to the welfare achieved in the worst-case equilibrium. *Liquid welfare* [24] is a widely used metric that extends standard social welfare to settings with budget constraints by capping an agent's contribution at their available budget. This provides a more meaningful benchmark for evaluating the social welfare achieved by a mechanism when agents face budget constraints. The POA of first-price auctions under autobidding was recently studied in [22], both in the model with only value maximizers and in the *mixed agent model* with both value and utility maximizers. They establish tight bounds of 2 and 2.188 on the POA of mixed Nash equilibria (MNE) for value maximizers and the mixed agent model, respectively. However, their analysis is limited to agents subject to ROI constraints only. Their proofs rely on structural properties of MNE and are tailored to their specific setting. Liaw et al. [36] studied budget-constrained value maximizers and derived a bound of 2 on the POA of pure Nash equilibria. To the best of our knowledge, this is also the only work that analyzes the inefficiency of simultaneous FPAs in the autobidding setting under both ROI and budget constraints.

A powerful framework for proving POA bounds is the *smoothness technique*, first introduced by Roughgarden [41] for strategic games and later extended by Syrgkanis and Tardos [45] to composable auctions. This framework allows POA bounds established for a 'base mechanism' (e.g., single-item first-price auction) to extend to more complex compositions (e.g., simultaneous first-price auctions) and to broader equilibrium concepts, including correlated and coarse correlated equilibria. Its ease of application combined with the strong, general guarantees it provides has made smoothness the technique of choice for studying the price of anarchy. Despite its success, applying the smoothness framework to the autobidding setting—where constraints and heterogeneous types play a central role—remains a major open challenge. This paper closes that gap.

1.1 Our Contributions.

Building on the challenges posed by heterogeneous autobidding agents and the shift to first-price auctions, we introduce a *type-dependent smoothness framework* that enables a unified analysis of the POA in this complex environment. Our contributions extend prior results in several important directions:

1. We consider a general *hybrid agent model* in which agents differ in type and are subject to both ROI and budget constraints. The *type* $t \in [0, 1]$ of an agent reflects their aversion to payments. In particular, $t = 0$ corresponds to a value maximizer, and $t = 1$ to a utility maximizer. By allowing arbitrary $t \in [0, 1]$, our model captures the full spectrum of agent behavior between these two extremes. This model strictly generalizes the mixed agent model in [22] and has been suggested in [7, 1].
2. We significantly broaden the class of valuation functions that can be handled in the autobidding context by analyzing *fractionally subadditive* (also known as *XOS* [34]) valuation functions. This class notably includes monotone submodular functions as a special case. These functions are particularly relevant in multi-platform autobidding environments, where advertisers use multiple platforms and experience diminishing returns as their ads are shown across them (see, e.g., [4]). To the best of our knowledge, prior work on autobidding has focused exclusively on additive valuations.
3. We extend our analysis beyond pure and mixed Nash equilibria and derive (tight) POA bounds for more general solution concepts, including *correlated equilibria (CE)* and *coarse correlated equilibria (CCE)*, even in the full generality of our hybrid agent model. CCE are particularly relevant in autobidding settings when agents use regret-minimizing algorithms (see, e.g., [37] and references therein). Building on insights from [32], we show that CCE induced by such learning dynamics satisfy additional structural properties essential for our framework to be applicable.
4. We also study the POA of simultaneous first-price auctions with reserve prices. Prior work by Balseiro et al. [7] and Deng et al. [22] studied this setting in the absence of budgets and under additive valuations.

Table 1.1: Overview of POA upper bounds for CCEs of simultaneous FPA for fractionally subadditive valuations (top table) and POA lower bounds of simultaneous FPA for additive valuations (bottom table). Here, $t_{\max}$ and $t_{\min}$ refer to the largest and smallest agent types, respectively. The function P is defined as $P(t) = 1 + t(1 + W_0(-e^{-t-1}))^{-1}$ for $t \geq \theta$ and $P(t) = 2$ for $t < \theta$, where W_0 is the principal branch of the Lambert W function and θ is a threshold value defined as $\theta = 1 + \frac{1}{2}W_0(-2e^{-2}) \approx 0.797$. Furthermore, the function Q is defined as $Q(t) = \frac{e}{e-1}$ for $t \geq \frac{e-1}{e}$, $Q(0) = 2$, and $Q(t) = 1 - \frac{(1-t)\ln(1-t)}{t}$, otherwise. Lastly, $\beta \approx 0.741$ and is the solution to $\beta = 1 - e^{-\frac{1}{\beta}}$

Upper Bounds

Agent Types	**Budgets**	**POA (UB)**	**Statements**
$t_{\max} \in [0, 1]$	✓	$P(t_{\max})$	Thm. 4.6
$t_{\min} = t_{\max} = t$	✗	$Q(t)$	Thm. 4.4
$t_{\min} \geq \beta$	✗	$t_{\min}\left(1 - e^{-\frac{1}{t_{\min}}}\right)$	Thm. 4.13

Lower Bounds

Agent Types	**Budgets**	**POA (LB)**	**Equilibrium class**	**Statements**
any	✓	2	*MNE*	Thm. 4.12
$t_{\min} = t_{\max} = t > \theta$	✓	$P(t)$	*CCE*	Thm. 4.8
$t_{\min} = t_{\max} = t \geq \frac{e-1}{e}$	✗	$\frac{e}{e-1}$	*CCE*	Cor. 4.5
$t_{\min} = 0$, $t_{\max} \in [0, 1]$	✗	$P(t)$	*MNE*	Thm. 4.10

We extend their analysis to encompass fractionally subadditive valuations, both with and without budget constraints. As observed by Balseiro et al., reserve prices can be interpreted as *predictions* of the agents' values (e.g., derived from historical data through machine-learning techniques), which can be used to improve efficiency guarantees. By leveraging such predictions to set reserve prices, we obtain improved POA bounds. This connects to the broader agenda of *mechanism design with predictions* (see also [28, 15]).

1.2 Overview and Significance of Our Results. An overview of the POA upper bounds that we derive for CCE of simultaneous FPA with fractionally subadditive valuations and the corresponding lower bounds that we derive can be found in Table 1.1. Below, we highlight some key implications of our results. Our bounds depend on the set of agent types $T \subseteq [0, 1]$, where $t_{\min} = \min(T)$ and $t_{\max} = \max(T)$ denote the smallest and largest agent types, respectively.

Among other results, our bounds uncover an intriguing threshold phenomenon: In the hybrid agent model with budgets and heterogeneous types, the POA is exactly 2 when the largest agent type satisfies $t_{\max} \leq \theta \approx 0.797$. When $t_{\max} > \theta$ and value maximizers are present (i.e., $t_{\min} = 0$), the POA bound increases smoothly with $t_{\max}$, following the function $P(t_{\max})$, from 2 up to 2.188 as $t_{\max}$ approaches 1; see Figure 1.1a (blue curve). This bound is tight even for mixed Nash equilibria. This result unifies and generalizes the state-of-the-art POA bounds in [22] and [36], both of which address special cases of our model. The same bound also applies when agents are budget-constrained and have a *uniform* (or *homogeneous*) type t, i.e., $t_{\min} = t_{\max} = t$[1].

Interestingly, we obtain the exact same POA bounds without budget constraints, as long as value maximizers are present. As will become clear below, this is not a coincidence. In the budget-free setting with uniform agent type, we derive strictly improved bounds, illustrated in Figure 1.1a (red curve). This yields a natural separation result for uniform agents: for every $t > 0$, the POA for budget-constrained agents (Figure 1.1a (blue curve)) is *strictly worse* than that for budget-free agents (Figure 1.1a (red curve)). Our results also reveal a second threshold

[1]When agents are utility maximizers i.e., $t = 1$, our objective coincides with the notion of *effective welfare* introduced in [45] (see also [13]). Note that, for budget-constrained agents, Syrgkanis and Tardos [45] establish a guarantee of $\frac{e}{e-1}$ for the weaker benchmark of the ratio between optimal effective welfare and the utilitarian social welfare at equilibrium, see also Further Related Work.

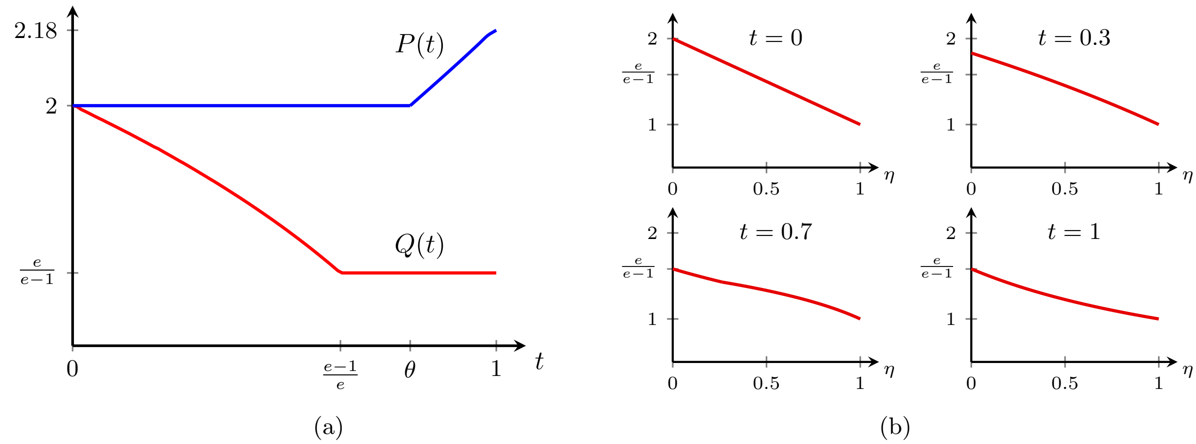

Figure 1.1: Illustration of POA bounds of simultaneous FPA with fractionally subadditive valuations. (a) Bounds for the setting without reserve prices as a function of the type t. (b) Bounds for the setting with reserve prices as a function of η (quality of the reserve prices) for $t \in \{0, 0.3, 0.7, 1\}$.

phenomenon: In the budget-free setting, if agents are sufficiently close to utility maximizers, i.e., $t_{\min} \geq \beta \approx 0.741$, the POA is in guaranteed to lie in the range $[e/(e-1), 1/\beta^2]$ (as a function of $t_{\min}$), regardless of type heterogeneity; see Theorem 4.13 for details.

We briefly comment on our results for simultaneous first-price auctions with reserve prices. In this setting, the POA bounds additionally depend on a parameter $\eta \in [0, 1)$, which quantifies the quality of the reserve prices. Intuitively, higher values of η correspond to better reserve prices; formally η captures the worst-case ratio between the reserve price and the maximum extractable payment over all auctions (see below for a precise definition). We derive new bounds for this setting (not included in the overview table). For budget-free agents with a uniform type t, the POA bound is shown in Figure 1.1b. As expected, the POA improves with the quality of the reserve prices, i.e., it decreases as η increases and converges to 1 as η goes to 1. The bound is tight for both $t = 0$ and $t = 1$.

In our setting with reserve prices, our framework requires that all items are sold in equilibrium outcomes—a property we refer to as *well-supported equilibria*. We show that, for additive valuations, this property holds for equilibria up to and including correlated equilibria (CE), but fails for coarse correlated equilibria (CCE). However, building on [32], we prove that in repeated single-item first-price auctions with reserve prices, CCE arising from regret-minimizing agents are well-supported regardless of the agents' type distribution. This result implies that CCE produced by such learning dynamics inherently possess the well-supported property.

1.3 Our Techniques. At the heart of our analysis lies a novel, *type-dependent smoothness framework*—a non-trivial generalization of the original framework by Syrgkanis and Tardos [45] for utility maximizers to heterogeneous agent types. A key technical challenge is incorporating these diverse agent types. Prior smoothness-based works assume that agents are 'alike', allowing the (λ, μ)-smoothness parameters of the base mechanism to lift directly to composed mechanisms. In contrast, our setting requires handling heterogeneous types, and applying the original approach directly fails to yield meaningful bounds in the autobidding context.

We overcome this by proving a smoothness inequality for each type separately, yielding type-specific (λ_t, μ_t) parameters. The core insight is to balance these parameters—crucially leveraging the ROI constraint—through carefully chosen *calibration vectors* to obtain the best POA bound. This balancing leads to an optimization problem over feasible choices of the smoothness parameters, giving rise to what we term the *POA-revealing mathematical program* (POA-RMP). Bounding the objective of this program then yields upper bounds on the POA. With this machinery in place, we then prove smoothness for a single-item first-price auction with reserve prices across different agent types. By the Extension Theorem, we derive (often tight) POA bounds for coarse correlated equilibria of simultaneous first-price auctions with reserve prices and fractionally subadditive valuations.

To handle budget constraints, we introduce a novel reduction technique that transforms instances with budgets into equivalent *budget-free proxy instances*. A key insight is that any budget-constrained agent can be simulated 'almost perfectly' by a budget-unconstrained agent with a *budget-capped valuation function*, i.e., one that caps their original valuation at their budget. This transformation, however, comes with a caveat: it fails when an agent's valuation exceeds their budget. We resolve this by introducing, for each such agent, a budget-unconstrained *value maximizer* with the corresponding budget-capped valuation function. The resulting instance is our *budget-free proxy instance*.

In a nutshell, our strategy is to reduce an instance with budget-constrained agents to a budget-free proxy instance, and then apply our type-dependent smoothness framework to bound the POA of the proxy instance. Since the transformation preserves the POA, this yields bounds for the original instance as well. This approach, however, requires special care: (i) The proxy instance includes budget-capped valuation functions. For example, if the original valuations are additive, the capped versions become submodular. (ii) The transformation introduces heterogeneity in agent types. Even if the original instance had a single type $t \neq 0$, the proxy instance includes both type t and type 0 agents. Addressing both (i) and (ii) relies critically on the full power of our type-dependent smoothness framework. For (i), the framework supports XOS valuations, and crucially, budget-capped XOS functions remain XOS. For (ii), our framework is explicitly designed to accommodate heterogeneous agent types. Notably, any approach incapable of handling either aspect would render the reduction technique infeasible.

Finally, the augmented type set in (ii) explains a key pattern observed in our results: all POA bounds for instances with budgets and type set T match those of their budget-free proxy instance, whose type set becomes $T^+ = T \cup \{0\}$.

We demonstrate the power of our type-dependent smoothness technique by analyzing the POA of simultaneous first-price auctions. However, the technique is broadly applicable, and we believe it extends to a wide range of autobidding environments—and potentially even beyond. Indeed, we already have evidence that it can be used to bound the POA of multi-unit auctions under autobidding, further underscoring its broader impact.

1.4 Further Related Work. Aggarwal et al. [2] initiated the study of the inefficiency of equilibria for auctions in autobidding environments. Their result implies that the liquid welfare price of anarchy for pure Nash equilibria of the second-price auction is 2. This upper bound was later generalized in [23], which considered a more general autobidding environment and the VCG mechanism, while [20, 21] obtained POA bounds for the Generalized Second Price auction (GSP).

The inefficiency of equilibria for simultaneous FPAs was first studied in [35], which showed that when all agents are value maximizers constrained only by ROI and have additive valuations, the POA of pure Nash equilibria is also 2. This result was then extended in [22] to MNE and ROI-constrained agents, which also introduced the mixed-agent model, i.e., the setting where agents can be either utility or value maximizers, which we capture as a special case. Liaw et al. [36] studied the inefficiency of simultaneous FPAs for agents who are both ROI-constrained and budget-constrained, focusing on pure Nash equilibria and showing a POA of 2. To the best of our knowledge, [36] is the only work aside from ours that studies the inefficiency of simultaneous FPAs under both ROI and budget constraints.

Beyond simultaneous compositions of simple classical auction formats, other autobidding settings that have been considered in the literature include the inefficiency of randomized auction mechanisms (see, e.g., [39, 35]) and scenarios in which the platform is allowed to "boost" the budgets of agents and implement reserve prices (see, e.g., [23, 7]), with the latter having an interpretation as *machine-learned advice*. We refer the reader to Section 5 for a discussion of this perspective. Finally, beyond the inefficiency of equilibria, other directions relevant to autobidding include the study of optimal bidding from the perspective of the agent (see, e.g., [2, 8]), online learning (see, e.g., [10, 27, 14, 3, 37]), auction design (see, e.g., [29, 9, 38]), and multi-platform (multi-channel) autobidding (see, e.g., ([18, 43, 4])). For further details, we refer the interested reader to the recent survey of Aggarwal et al. [1].

For the standard setting in which all agents are utility maximizers, the inefficiency of the first-price auction has been studied in the economics literature since the seminal work of Vickrey [46]. Naturally, due to its simplicity, it has also been considered for simultaneous simple auctions. As shown in [45], the *POA* of *CCE* is at most $\frac{e}{e-1}$ for simultaneous auctions with XOS valuations. This bound is known to be tight even for a single auction [44], and for *MNE* in simultaneous auctions with submodular valuations [16]. Beyond the smoothness framework, Feldman et al. [25] showed a *POA* upper bound of 2 for simultaneous auctions with subadditive valuations, while *CE*, *CCE* and their properties have been considered more closely in [26]. More recently, it was shown that the *Bayesian POA* for the single-item first-price auction is exactly $\frac{e^2}{e^2-1}$ [31], a breakthrough result. For an overview of classical results regarding the POA of auctions for utility maximizers, including compositions of other simple auctions, we refer to the survey by Roughgarden et al. [42].

Finally, we remark that budgeted settings have been considered for utility maximizers from a POA perspective prior to the emergence of autobidding. Since our model captures the setting where all agents are utility maximizers as a special case, our work can also be viewed as a follow-up to this line of research. In this context, the three most closely related works are [45, 6], and [13]. While in [13] the focus is on the proportional mechanism, simultaneous first-price auctions with XOS valuations are considered in both [45] and [6]. However, Syrgkanis and Tardos [45] focus on the ratio of the expected *social* welfare at equilibrium to the optimal liquid welfare, a weaker benchmark than the one we consider in this work (see also Section 4.6). On the other hand, Azar et al. [6] analyze the *ex post* liquid welfare, which is a stronger benchmark, but their results require the items to be divisible into discretely many parts.[2] Note that Caragiannis and Voudouris [13] adopt the same benchmark as we do and call it *Effective Welfare*.

2 Preliminaries. We study *simultaneous first-price auctions*, where a set $N = [n]$ of $n \geq 2$ agents simultaneously participate in a set $M = [m]$ of $m \geq 1$ single-item auctions. Each auction $j \in M$ implements a *first-price auction (FPA) with reserve price*, as detailed below. We use j to denote both the auction and the respective item interchangeably. Each agent $i \in N$ submits a bid $b_{ij} \in \mathbb{R}_{\geq 0}$ to each auction $j \in M$. We use $\boldsymbol{b}_i = (b_{ij})_{j \in M}$ to denote the bid profile of agent i, and $D_i = \mathbb{R}^m_{\geq 0}$ to refer to the set of all bid profiles of i. The aggregated bid profile of all agents is denoted by $\boldsymbol{b} = (\boldsymbol{b}_i)_{i \in N} \in D = \times_{i \in N} D_i$.

We focus on simultaneous first-price auctions with reserve prices. More specifically, each auction $j \in M$ handles a reserve price $r_j \in \mathbb{R}_{\geq 0}$ that must be met to sell item j; we use $FPA(r_j)$ to refer to this auction. Given the bid profile $\boldsymbol{b}_j = (b_{ij})_{i \in N}$ submitted to auction j, $FPA(r_j)$ allocates the item to the highest bidder i meeting the reserve price r_j, i.e., $b_{ij} \geq r_j$, and charges their respective bid b_{ij} for the item. The agent who wins the item (if any) is called the *actual winner*, denoted by $\mathsf{aw}(j) \in \arg\max_{i \in N : b_{ij} \geq r_j} b_{ij}$. In case of ties, the actual winner is chosen according to an arbitrary but fixed tie-breaking rule. If the reserve price r_j is not met (i.e., $b_{ij} < r_j$ for all $i \in N$), we define $\mathsf{aw}(j) = \emptyset$. Let $\boldsymbol{x}_j(\boldsymbol{b}) = (x_{ij}(\boldsymbol{b}))_{i \in N}$ and $\boldsymbol{p}_j(\boldsymbol{b}) = (p_{ij}(\boldsymbol{b}))_{i \in N}$ be the respective allocation and payments of $FPA(r_j)$, i.e., for $i = \mathsf{aw}(j)$ we have $x_{ij}(\boldsymbol{b}) = 1$ and $p_{ij}(\boldsymbol{b}) = b_{ij}$, and for $i \neq \mathsf{aw}(j)$ we have $x_{ij}(\boldsymbol{b}) = 0$ and $p_{ij}(\boldsymbol{b}) = 0$.[3] We write $\boldsymbol{x}_j(\boldsymbol{b}) \neq \boldsymbol{0}$ to indicate that item j is sold.

Our global mechanism, denoted by $\mathcal{M}$, implements the above mechanisms with reserve prices $\boldsymbol{r} = (r_j)_{j \in M}$ simultaneously. That is, given a bid profile $\boldsymbol{b}$, the outcome $\mathcal{M}(\boldsymbol{r}, \boldsymbol{b}) = (\boldsymbol{x}(\boldsymbol{b}), \boldsymbol{p}(\boldsymbol{b}))$ is determined by the allocation $\boldsymbol{x}(\boldsymbol{b}) = (\boldsymbol{x}_j(\boldsymbol{b}))_{j \in M}$ and the payments $\boldsymbol{p}(\boldsymbol{b}) = (\boldsymbol{p}_j(\boldsymbol{b}))_{j \in M}$ obtained by running the m auctions (i.e., $FPA(r_j)$ for each $j \in M$) simultaneously. We write $\boldsymbol{x}(\boldsymbol{b}) \neq \boldsymbol{0}$ to indicate that all items are sold under $\boldsymbol{b}$, i.e., $\boldsymbol{x}_j(\boldsymbol{b}) \neq \boldsymbol{0}$ for all $j \in M$.[4] We use $\boldsymbol{x}_i(\boldsymbol{b}) = (x_{ij}(\boldsymbol{b}))_{j \in M} \in \{0,1\}^m$ and $p_i(\boldsymbol{b}) = \sum_{j \in M} p_{ij}(\boldsymbol{b})$ to denote the allocation and total payment of agent i, respectively. Further, we define $\boldsymbol{X}$ as the set of all feasible allocations, i.e., $\boldsymbol{X} = \{\boldsymbol{x} = (\boldsymbol{x}_i)_{i \in N} \in \{0,1\}^{m \times n} \mid \sum_{i \in N} x_{ij} \leq 1 \; \forall j \in M\}$. We slightly overload notation and use $\boldsymbol{x}_i(\boldsymbol{b})$ also to refer to the set of items assigned to i, i.e., $\boldsymbol{x}_i(\boldsymbol{b}) = \{j \in M \mid x_{ij}(\boldsymbol{b}) = 1\} \subseteq M$. Additionally, we sometimes omit the argument $\boldsymbol{b}$ when it is clear from context.

We use $FPA(m, \boldsymbol{r})$ and $FPA(m)$, respectively, to refer to m simultaneous first-price auctions with reserve prices $\boldsymbol{r}$ and without reserve prices. We use $FPA(r)$ to indicate that we consider a single-item first-price auction (i.e., $m = 1$) with reserve price r. If $m = 1$, we drop the auction index $j = 1$ from all our notation.

Valuation Functions. Each agent $i \in N$ has a valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$ over the subsets of the items, where $v_i(S)$ specifies the value that i obtains when receiving the items in $S \subseteq M$. We assume w.l.o.g. that $v_i(\emptyset) = 0$. Also, we assume that v_i is *monotone*, i.e., $v_i(S) \leq v_i(T)$ for all $S \subseteq T \subseteq M$. We use $\mathcal{V}_i$ to denote the class of valuation functions of agent i and let $\mathcal{V} = \times_{i \in N} \mathcal{V}_i$ be the set of all valuation functions of the agents. We use $\boldsymbol{v} = (v_i)_{i \in N} \in \mathcal{V}$ to refer to the profile of valuation functions of the agents. We consider different classes of valuation functions:

DEFINITION 2.1. *Let $v_i : 2^M \to \mathbb{R}_{\geq 0}$ be a valuation function.*

- v_i *is* additive *if there exist additive valuations $(v_{ij})_{j \in M} \in \mathbb{R}^m_{\geq 0}$ such that for every subset $S \subseteq M$, it holds that $v_i(S) = \sum_{j \in S} v_{ij}$.*

[2]The value of the ex post benchmark is unbounded for our setting of simultaneous first-price auctions, see Theorem C.2. in [5].
[3]Both $\boldsymbol{x}_j$ and $\boldsymbol{p}_j$ only depend on the input profile $\boldsymbol{b}_j$. However, we often use $\boldsymbol{b}$ as the argument for notational convenience.
[4]Note that we slightly abuse notation as $\boldsymbol{x}(\boldsymbol{b}) \neq \boldsymbol{0}$ here indicates that there is exactly one 1 in each row of $\boldsymbol{x}(\boldsymbol{b}) \in \{0,1\}^{m \times n}$.

- v_i *is* submodular *if* $v_i(S \cup \{j\}) - v_i(S) \geq v_i(T \cup \{j\}) - v_i(T)$ *for all* $S \subseteq T \subseteq M$.
- v_i *is* fractionally subadditive *(or,* XOS*), if there exists a class* $\mathcal{L}_i = \{(v^\ell_{ij})_{j \in M} \in \mathbb{R}^m_{\geq 0}\}$ *of additive valuations such that for every subset* $S \subseteq M$*, it holds that* $v_i(S) = \max_{\ell \in \mathcal{L}_i} \sum_{j \in S} v^\ell_{ij}$.

Let $\mathcal{V}_{\text{ADD}}$, $\mathcal{V}_{\text{SUB}}$ and $\mathcal{V}_{\text{XOS}}$ refer to the set of additive, submodular and fractionally subadditive (XOS) valuation functions, respectively. It is well-known (see e.g., [34]) that $\mathcal{V}_{\text{ADD}} \subset \mathcal{V}_{\text{SUB}} \subset \mathcal{V}_{\text{XOS}}$.

Random Bid Profiles. Each agent i can randomize over their deterministic (or pure) bid profiles $\boldsymbol{b}_i$ in D_i. We define Δ_i as the space of random bid profiles of i over D_i. Let π be a probability distribution over the set of bid profiles in D; we use Δ to refer to the set of all such probability distributions. We use $\boldsymbol{B} \sim \pi$ to denote a random bid profile that is drawn from π; we often omit the reference to π and identify $\boldsymbol{B}$ with π. We use $f_{\boldsymbol{B}}$ and $F_{\boldsymbol{B}}$ to refer to the probability density function (PDF) and cumulative distribution function (CDF) of $\boldsymbol{B}$, respectively. The support of $\boldsymbol{B}$ refers to the set of bid profiles that have positive density, i.e., $\text{supp}(\boldsymbol{B}) = \{\boldsymbol{b} \in D \mid f_{\boldsymbol{B}}(\boldsymbol{b}) > 0\}$. If $\text{supp}(\boldsymbol{B}) = \{\boldsymbol{b}\}$ then $\boldsymbol{B}$ chooses $\boldsymbol{b}$ deterministically and we write $\boldsymbol{B} = \boldsymbol{b}$. We use $\text{supp}_i(\boldsymbol{B})$ to refer to the set of bid profiles $\boldsymbol{b}_i$ of agent i that have positive density under $\boldsymbol{B}$. The *marginal* $\boldsymbol{B}_{-i}$ of $\boldsymbol{B}$ is defined by the following PDF:

$$\forall \boldsymbol{b}_{-i} \in D_{-i}: \qquad f_{\boldsymbol{B}_{-i}}(\boldsymbol{b}_{-i}) = \int_{D_i} f_{\boldsymbol{B}}(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) d\boldsymbol{b}_i.$$

Given a bid profile $\boldsymbol{b}'_i$ of agent i, we denote by $(\boldsymbol{b}'_i, \boldsymbol{B}_{-i})$ the random bid profile that we obtain from $\boldsymbol{B}$ when agent i bids $\boldsymbol{b}'_i$ deterministically and the other agents bid according to the marginal $\boldsymbol{B}_{-i}$. We say that a bid profile $\boldsymbol{B}$ is *well-supported* with respect to reserve prices $\boldsymbol{r}$ if the items are always sold under $\boldsymbol{B}$, i.e., for each $\boldsymbol{b} \in \text{supp}(\boldsymbol{B})$, $\boldsymbol{x}(\boldsymbol{b}) \neq \boldsymbol{0}$.

Hybrid Agent Model. We consider the general *hybrid agent model* (see, e.g., [7, 1]), where each agent $i \in N$ maximizes their *gain function* $g_i : D \to \mathbb{R}$ defined as

$$g_i(\boldsymbol{b}) = v_i(\boldsymbol{x}_i(\boldsymbol{b})) - \sigma_i \cdot p_i(\boldsymbol{b}). \tag{2.1}$$

Here, $\sigma_i \in [0, 1]$ defines the *type* of agent i. Intuitively, σ_i represents i's sensitivity to payments: a higher value indicates that i is more negatively affected by payments. In particular, agent i is a value maximizer if $\sigma_i = 0$, and a utility maximizer if $\sigma_i = 1$. Our model thus allows us to capture a large spectrum of agents' types, ranging from value maximizers to utility maximizers. Most previous works focus on the special case of the *mixed agent model* consisting of value and utility maximizers only, i.e., $\sigma_i \in \{0, 1\}$ for all $i \in N$.

Each agent i has a *return-on-investment (ROI) constraint* and a *budget constraint* that must be satisfied (see [1]). Given a bid profile $\boldsymbol{B}$, the ROI constraint of an agent i enforces that the expected total payment of i is at most a factor $\tau_i \in \mathbb{R}_{>0}$ of their expected valuation for the received items, where τ_i is the so-called *target parameter* of i, i.e.,

$$\mathbb{E}\left[p_i(\boldsymbol{B})\right] \leq \tau_i \cdot \mathbb{E}[v_i(\boldsymbol{x}_i(\boldsymbol{B}))]. \tag{2.2}$$

Additionally, the budget constraint of agent i requires that the expected total payment of i is at most $\mathfrak{B}_i \in \mathbb{R}_{>0} \cup \{\infty\}$, i.e.,

$$\mathbb{E}\left[p_i(\boldsymbol{B})\right] \leq \mathfrak{B}_i. \tag{2.3}$$

For an agent i, we define $\mathcal{R}_i$ as the set of bid profiles $\boldsymbol{B}$ that satisfy both the ROI constraint (2.2) and the budget constraint (2.3). It is not hard to see that we can assume w.l.o.g. that $\tau_i \sigma_i \leq 1$ for all agents $i \in N$; we refer to the full version for more details.

Formally, we use $I = (N, M, \boldsymbol{r}, \boldsymbol{v}, \boldsymbol{\sigma}, \boldsymbol{\tau}, \boldsymbol{\mathfrak{B}})$ to denote an instance. For ease of notation, we omit explicit references to N and M and simply write $I = (\boldsymbol{r}, \boldsymbol{v}, \boldsymbol{\sigma}, \boldsymbol{\tau}, \boldsymbol{\mathfrak{B}})$. We say that an instance I is *budget-free* if $\mathfrak{B}_i = \infty$ for all agents $i \in N$. The set of *agent types* of an instance I is defined as $T(I) = \{t \mid \exists i \in N \text{ with } \sigma_i = t\}$. Given an instance I, we use $N_t(I) \subseteq N$ to refer to the subset of agents having type t, i.e., $N_t(I) = \{i \in N \mid \sigma_i = t\}$. The following notion will turn out to be useful below. Given a type set $T \subseteq [0, 1]$, we define the *augmented type set* T^+ as follows: for instances with budgets, we define $T^+ = T \cup \{0\}$ as the type set obtained from T by adding the value-maximizing type 0; for budget-free instances, we define $T^+ = T$ simply.

We use $\mathcal{I}$ to refer to a class of instances. For $\text{VAL} \in \{\text{ADD}, \text{SUB}, \text{XOS}\}$, we use $\mathcal{I}_{\text{VAL}}$ to refer to the set of instances with valuation functions from $\mathcal{V}_{\text{VAL}}$, i.e., $v_i \in \mathcal{V}_{\text{VAL}}$ for all $i \in N$. We use $\mathcal{I}^{\infty}$ to refer to the set of all budget-free instances. Finally, we also define a class of instances whose type set is restricted. Given a set of types $T \subseteq [0,1]$, we use $\mathcal{I}^T$ to denote the set of instances I whose type set is restricted to T, i.e., $T(I) = T$. Note that the above restrictions will also be combined. For example, $\mathcal{I}_{\text{VAL}}^{T,\infty}$ denotes the class of all budget-free instances with valuations functions from $\mathcal{V}_{\text{VAL}}$ and agent types being restricted to T.

Agent's Problem and Equilibrium Notions. The objective of each agent i is to determine a random bid profile $\boldsymbol{B}_i$ that, given the bid profile $\boldsymbol{B}_{-i}$ of the other agents, maximizes their gain subject to their ROI and budget constraints. More formally, each agent i solves the following *agent's problem*:

$$\max_{\boldsymbol{B}_i} \mathbb{E}[g_i(\boldsymbol{B}_i, \boldsymbol{B}_{-i})] \quad \text{subject to} \quad (\boldsymbol{B}_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i.$$

The resulting bid profile $\boldsymbol{B}$ constitutes an equilibrium if for each agent i only the deviations satisfying the ROI and budget constraints are considered. We consider the following equilibrium notions in this paper.

DEFINITION 2.2. *Let $\boldsymbol{B} \in \Delta$ be a bid profile satisfying $\boldsymbol{B} \in \mathcal{R}_i$ for each agent $i \in N$.*

- *$\boldsymbol{B}$ is a* coarse correlated equilibrium *(CCE) if for every agent $i \in N$ we have:*

$$\mathbb{E}[g_i(\boldsymbol{B})] \geq \mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] \qquad \forall (\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i. \tag{2.4}$$

- *$\boldsymbol{B}$ is a* correlated equilibrium *(CE) if for every agent $i \in N$ and every swapping function $h : \text{supp}_i(\boldsymbol{B}) \mapsto \Delta_i$ we have:*

$$\mathbb{E}[g_i(\boldsymbol{B})] \geq \mathbb{E}[g_i(h(\boldsymbol{B}_i), \boldsymbol{B}_{-i})] \qquad \forall (h(\boldsymbol{B}_i), \boldsymbol{B}_{-i}) \in \mathcal{R}_i. \tag{2.5}$$

- *$\boldsymbol{B}$ is a* mixed Nash equilibrium *(MNE) if $\boldsymbol{B} = \prod_{i \in [n]} \boldsymbol{B}_i$ and for every agent $i \in N$ we have:*

$$\mathbb{E}[g_i(\boldsymbol{B})] \geq \mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] \qquad \forall (\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i. \tag{2.6}$$

Given an instance I, we use *MNE*(I), *CE*(I) and *CCE*(I) to refer to the sets of mixed, correlated and coarse correlated equilibria of I, respectively. Note that *MNE*$(I) \subseteq$ *CE*$(I) \subseteq$ *CCE*(I); a more elaborate discussion on the equilibrium hierarchy can be found in the full version. We use *EQ* as a generic placeholder for an equilibrium notion with *EQ* $\in \{$*MNE*, *CE*, *CCE*$\}$.

Liquid Price of Anarchy. We use *liquid welfare* as the social welfare objective, which is also the standard benchmark in the autobidding literature (see, e.g., [1]). Intuitively, the liquid welfare measures the maximum amount of payments one can extract from the agents. We refer the reader to the full version [17] for a more detailed discussion of the liquid welfare objective. Given an instance I and a bid profile $\boldsymbol{B}$, the liquid welfare is defined as

$$LW(I, \boldsymbol{B}) = \sum_{i \in N} \min(\mathbb{E}[\tau_i v_i(\boldsymbol{x}_i(\boldsymbol{B}))], \mathfrak{B}_i).$$

Given an instance I, an optimal allocation $\boldsymbol{x}^*(I) \in \boldsymbol{X}$ maximizes the liquid welfare over all feasible allocations:

$$\boldsymbol{x}^*(I) \in \arg\max_{\boldsymbol{x} \in \boldsymbol{X}} \sum_{i \in N} \min(\tau_i v_i(\boldsymbol{x}_i), \mathfrak{B}_i).$$

We use $OPT(I) = \sum_{i \in N} \min(\tau_i v_i(\boldsymbol{x}_i^*(I)), \mathfrak{B}_i)$ to denote the optimal liquid welfare.

In this paper, we study the *price of anarchy* as introduced by Koutsoupias and Papadimitriou [33] with respect to the liquid welfare objective: the price of anarchy is defined as the worst-case ratio of the optimal liquid welfare over the expected liquid welfare of any equilibrium. More formally, given a set of instances $\mathcal{I}$ and an equilibrium notion *EQ* $\in \{$*MNE*, *CE*, *CCE*$\}$, we define the *price of anarchy with respect to EQ* as:

$$EQ\text{-}POA(\mathcal{I}) = \sup_{I \in \mathcal{I}} \sup_{\boldsymbol{B} \in EQ(I)} \frac{OPT(I)}{LW(I, \boldsymbol{B})}.$$

It is not hard to observe that when studying the liquid price of anarchy, we can assume without loss of generality that $\tau_i = 1$ for all $i \in N$ (we include a discussion in the full version of our work). Subsequently, we therefore omit the target parameters (with the understanding that $\boldsymbol{\tau} = \boldsymbol{1}$) and refer to an instances as $I = (\boldsymbol{r}, \boldsymbol{v}, \boldsymbol{\sigma}, \mathfrak{B})$.

Reserve Prices. Balseiro et al. [7] and Deng et al. [22] studied the effect of reserve prices on the price of anarchy for budget-free instances and additive valuations. Their bounds depend on a parameter η that measures the relative gap between the reserve price and the highest valuation of an agent over all auctions. We extend their model to instances with fractionally subadditive valuations (with or without budgets). Let $I \in \mathcal{I}_{\text{XOS}}$ be such an instance with reserve prices $\boldsymbol{r}$. We can then choose *opt-induced additive representatives* $(v_{ij})_{j \in M}$ for each agent i with respect to their allocation $\boldsymbol{x}_i^*$ in the *optimal solution* $\boldsymbol{x}^*(I)$ (see Section 3.2 for more details). An agent i is said to be the *rightful winner* of auction j, denoted by $\mathsf{rw}(j)$, if i is an agent with maximum valuation for the item, i.e., $\mathsf{rw}(j) \in \arg\max_{i \in N} v_{ij}$. In case of ties, we let $\mathsf{rw}(j)$ denote the winner of auction j in the considered optimal allocation. For each auction $j \in M$, define the relative gap $\eta_j \in [0, 1)$ such that $r_j = \eta_j v_{\mathsf{rw}(j)j}$, and let $\eta = \min_{j \in M} \eta_j$ be the smallest relative gap. Note that, as in [7] and [22], we only consider reserve prices $\boldsymbol{r}$ that satisfy $\eta_j \in [0, 1)$ for each auction $j \in M$. In fact, it is not hard to see that otherwise the POA is unbounded.

The Lambert W Function. In order to derive POA bounds analytically, we use the Lambert W function, which is the multivalued inverse of the function $f(z) = ze^z$. In this work, we exclusively use the *principal branch* of the Lambert W function and denote it by W_0; more details are given in the appendix.

Some proofs are omitted from the main text and are provided in the appendix. All missing proofs and additional discussions are available in the full version of our work [17].

3 Type-Dependent Smoothness Framework.

We introduce a new type-dependent smoothness framework that enables us to bound the POA of coarse correlated equilibria in simultaneous first-price auctions, in the full generality of our hybrid model. We begin by formalizing the type-dependent smoothness notion and deriving corresponding smoothness lemmas. To handle budget constraints, we apply a reduction technique that transforms instances with budgets into budget-free proxy instances. Our Extension Theorem then leverages this framework to establish upper bounds on the POA of coarse correlated equilibria. A key technical challenge lies in balancing the type-dependent smoothness parameters to obtain the best possible bounds. To this end, we formulate a mathematical program that facilitates the analysis of the price of anarchy.

3.1 Type-Dependent Smoothness.

In this section, we focus on single-item instances (i.e., $m = 1$) and thus omit the auction index $j = 1$ from the notation. Note that in this case, fractionally subadditive valuation functions reduce to additive ones. In particular, each agent i has a single additive value v_i for the item (see Definition 2.1).

We need the notion of ROI-restricted bid profiles. Let $B'_i \in \Delta_i$ be a bid profile of agent i. We say that B'_i is *ROI-restricted* if for each $\boldsymbol{b}_{-i} \in D_{-i}$, it holds that $\mathbb{E}[p_i(B'_i, \boldsymbol{b}_{-i})] \le \mathbb{E}[v_i(\boldsymbol{x}_i(B'_i, \boldsymbol{b}_{-i}))]$. We can now introduce our new type-dependent smoothness notion:

Definition 3.1 (Type-Dependent Smoothness). *Let $I = (r, \boldsymbol{v}, \boldsymbol{\sigma}, \mathfrak{B})$ be a single-item instance with reserve price r. Let the rightful winner $i = \mathsf{rw}$ be of type t. Then, FPA(r) is* (λ_t, μ_t)-smooth *for type t with $\lambda_t, \mu_t > 0$, if there is a ROI-restricted bid $B'_i = B'_i(\boldsymbol{v}) \in \Delta_i$ such that for every well-supported bid profile $\boldsymbol{b}$ we have*

$$\mathbb{E}[g_i(B'_i, \boldsymbol{b}_{-i})] \ge \lambda_t v_i - \mu_t p_{\mathsf{aw}(\boldsymbol{b})}(\boldsymbol{b}). \tag{3.1}$$

We remark that, crucially, the random deviation B'_i of the rightful winner i may depend on the valuations $\boldsymbol{v}$ but *not* on the bid profile $\boldsymbol{b}$. Note that (3.1) needs to hold only for bid profiles $\boldsymbol{b}$ that are well-supported, i.e., when the item is sold; clearly, this condition is redundant in the setting without reserve prices, i.e., $r = 0$.

The following two lemmas establish type-dependent smoothness of *FPA*(r) for different types $t \in [0, 1]$. We start with a simple smoothness lemma for type-0 agents (i.e., value maximizers). Recall that $\eta \in [0, 1)$ measures the relative gap between the reserve price and the valuation of the rightful winner, i.e., $r = \eta v_{\mathsf{rw}}$.

Lemma 3.2. *Consider a single-item instance $I = (r, \boldsymbol{v}, \boldsymbol{\sigma}, \mathfrak{B})$ and let the rightful winner $i = \mathsf{rw}$ be of type $t = 0$. Then, FPA(r) is (λ_t, μ_t)-smooth for type t with $\lambda_t = \mu_t = \mu$ for every $\mu \in (0, (1 - \eta)^{-1}]$.*

A key insight used in the proof is that the reserve price allows us to increase the probability mass of the random deviation for larger bids. This provides a better trade-off in terms of the smoothness parameters. Our smoothness lemma for agent types $t \in (0, 1]$ follows the same approach, but is technically more challenging.

Lemma 3.3. *Consider a single-item instance $I = (r, \boldsymbol{v}, \boldsymbol{\sigma}, \mathfrak{B})$ and let the rightful winner $i = \mathsf{rw}$ be of type*

$t \in (0, 1]$. *Then, FPA(r) is* (λ_t, μ_t)*-smooth for type t with*

$$\text{(3.2)} \qquad \lambda_t = \frac{\mu}{t}\left(1 - \frac{1 - t\eta}{e^{t/\mu}}\right) \quad \textit{and} \quad \mu_t = \mu \quad \textit{for every} \quad \begin{cases} \mu \geq t\left(\ln\left(\frac{1-t\eta}{1-t}\right)\right)^{-1}, & \textit{if } t < 1, \\ \mu > 0, & \textit{if } t = 1. \end{cases}$$

3.2 Extension Theorem. We present our Extension Theorem to derive bounds on the price of anarchy. Our bounds depend on the set of available agent types $T \subseteq [0, 1]$. We consider the class of instances $\mathcal{I}^T_{\text{xos}}$ with fractionally subadditive valuations and type set T.

We first introduce the notion of *calibration vectors*, which will be crucial in the proof below. As we show in Lemma A.2 below, the set of calibration vectors $\mathcal{C}(\boldsymbol{\mu}, T)$ is always non-empty.

DEFINITION 3.4 (Calibration Vectors). *Let T be a set of agent types and let $\boldsymbol{\mu} = (\mu_t)_{t\in T}$ be such that $\mu_t > 0$ for each $t \in T$. We define the set of* calibration vectors $\mathcal{C}(\boldsymbol{\mu}, T)$ *as follows:*

$$\text{(3.3)} \qquad \mathcal{C}(\boldsymbol{\mu}, T) = \left\{ \boldsymbol{\delta} \in (0, 1]^{|T|} \;\middle|\; \max_{t\in T}(\delta_t \mu_t) + \max_{t\in T}(\delta_t(1 - t)) \leq 1 \right\}.$$

We can now state the main result of this section. Recall that T^+ is the augmented type set of T, where $T^+ = T \cup \{0\}$ for instances with budgets and $T^+ = T$ for budget-free instances.

THEOREM 3.5 (Extension Theorem). *Let $\mathcal{I}^T_{\text{xos}}$ be the class of instances with fractionally subadditive valuations and type set T. Assume that FPA(r) is (λ_t, μ_t)-smooth for each type $t \in T^+$. Then, the price of anarchy of well-supported coarse correlated equilibria bounded by*

$$\textit{CCE-POA}(\mathcal{I}^T_{\text{xos}}) \leq \left(\max_{\boldsymbol{\delta}\in\mathcal{C}(\boldsymbol{\mu},T^+)} \min_{t\in T^+} \delta_t \lambda_t \right)^{-1}.$$

The remainder of this section is devoted to the proof of Theorem 3.5.

3.2.1 Opt-Induced Additive Representatives and Budget-Free Proxy Instances. We first introduce the notion of *opt-induced additive representatives*. Intuitively, these representatives allow us to treat fractionally subadditive valuations as additive ones in the analysis. Let $I \in \mathcal{I}_{\text{xos}}$ be an instance with XOS valuation functions $\boldsymbol{v} = (v_i)_{i\in N}$. Fix an optimal allocation $\boldsymbol{x}^* := \boldsymbol{x}^*(I)$. By Definition 2.1, for each $i \in N$, there exist additive representatives $(v^*_{ij})_{j\in M}$ with respect to the optimal allocation $\boldsymbol{x}^*_i$.[5] We refer to these representatives as *opt-induced additive representatives*. We define v^*_i as the additive valuation function obtained from these representatives, i.e., $v^*_i(\boldsymbol{x}_i) := \sum_{j\in M} v^*_{ij} x_{ij}$ for any allocation $\boldsymbol{x}_i \subseteq M$. The following two properties follow directly from Definition 2.1: **(XOS1)** $v_i(\boldsymbol{x}^*_i) = v^*_i(\boldsymbol{x}^*_i)$. **(XOS2)** For any allocation $\boldsymbol{x}_i \subseteq M$, $v_i(\boldsymbol{x}_i) \geq v^*_i(\boldsymbol{x}_i)$.[6]

Next, we define *budget-capped valuations* that account for the budget constraints in (2.3). For every $i \in N$, the $\mathfrak{B}_i$*-capped valuation* $v^{\mathfrak{B}_i}_i : 2^M \mapsto \mathbb{R}_{\geq 0}$ is defined as $v^{\mathfrak{B}_i}_i(S) = \min(v_i(S), \mathfrak{B}_i)$ for all $S \subseteq M$. A crucial observation that we use below is that capped XOS valuation functions remain XOS.

PROPOSITION 3.6 (Lemma C.7 in [45]). $v_i \in \mathcal{V}_{\text{xos}} \Rightarrow v^{\mathfrak{B}_i}_i \in \mathcal{V}_{\text{xos}}$.

We can now formalize the notion of *budget-free proxy instances.*

DEFINITION 3.7 (Budget-Free Proxy Instance). *Given an instance $I = (\boldsymbol{r}, \boldsymbol{v}, \boldsymbol{\sigma}, \mathfrak{B})$ and a bid profile $\boldsymbol{B} \in \Delta$, the* budget-free proxy instance *of I and $\boldsymbol{B}$ is defined as $\hat{I}(I, \boldsymbol{B}) = (\boldsymbol{r}, \boldsymbol{v}^{\mathfrak{B}_i}, \hat{\boldsymbol{\sigma}}(\boldsymbol{B}), \boldsymbol{\infty})$ with*

$$\hat{\sigma}_i(\boldsymbol{B}) := \begin{cases} 0 & \textit{if } \mathfrak{B}_i < \mathbb{E}[v_i(x_i(\boldsymbol{B}))], \\ \sigma_i & \textit{otherwise.} \end{cases}$$

Intuitively, the budget-free proxy instance $\hat{I}$ simply replaces the valuation function of each agent i by its budget-capped counterpart and leaves i's type intact, unless i's valuation under $\boldsymbol{B}$ exceeds the budget $\mathfrak{B}_i$. In

[5] Note that these representatives simply coincide with the input valuations if the valuation functions $\boldsymbol{v}$ are additive.

[6] To see this, recall that the additive representatives $(v_{ij})_{j\in M}$ of agent i with respect to $\boldsymbol{x}_i$ are chosen as maximizers from the class $\mathcal{L}_i \ni (v^*_{ij})_{j\in M}$, and thus $v_i(\boldsymbol{x}_i) = \sum_{j\in M} v_{ij} x_{ij} \geq \sum_{j\in M} v^*_{ij} x_{ij}$.

the latter case, i's type is mapped to 0 and i becomes a value maximizer instead. As we show below, it suffices to focus on these proxy instances to bound the price of anarchy of instances with budgets.

Note that for every budget-free instance $I \in \mathcal{I}^\infty$ it holds that $\hat{I}(I, \boldsymbol{B}) = I$ for all $\boldsymbol{B} \in \Delta$. Moreover, it is not hard to show that the optimal solutions of an instance and each of its budget-free proxies coincide.

PROPOSITION 3.8. *Let $I = (\boldsymbol{r}, \boldsymbol{v}, \boldsymbol{\sigma}, \boldsymbol{\mathfrak{B}})$ and $\boldsymbol{B} \in \Delta$. Then, $OPT(\hat{I}(I, \boldsymbol{B})) = OPT(I)$.*

For the remainder of Section 3.2, given a pair $(I, \boldsymbol{B})$, we use $(\hat{g}_i)_{i\in N}$ to denote the agents' gain functions for the proxy instance $\hat{I}(I, \boldsymbol{B})$, i.e., for every $i \in N$ and every $\boldsymbol{b} \in D$, we define $\hat{g}_i(\boldsymbol{b}) := v_i^{\mathfrak{B}_i}(x_i(\boldsymbol{b})) - \hat{\sigma}_i(\boldsymbol{B}) p_i(\boldsymbol{b})$.

LEMMA 3.9. *Consider an instance $I \in \mathcal{I}_{\text{XOS}}$ and let $\boldsymbol{B} \in CCE(I)$. Then, for every agent $i \in N$, for every $\boldsymbol{B}'_i$ with $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$ and every $\delta \in [0, 1]$ it holds*

$$\min\left(\mathbb{E}[v_i(x_i(\boldsymbol{B}))], \mathfrak{B}_i\right) \geq \delta \cdot \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] + (1 - \delta + \delta\hat{\sigma}_i(\boldsymbol{B})) \cdot \mathbb{E}\left[p_i(\boldsymbol{B})\right].$$

3.2.2 Proof of Theorem 3.5. The following Lifting Lemma for budget-free instances provides the final ingredient for the proof of our Extension Theorem. Basically, for any type t, it lifts the smoothness property of $FPA(r)$ for a single auction of type t (i.e., where the rightful winner is of type t) to all auctions having the same type. The corollary below will be useful for its proof.

COROLLARY 3.10. *Consider a budget-free instance $I \in \mathcal{I}^\infty_{\text{XOS}}$. Fix an agent $i \in N$ and consider a bid profile $\boldsymbol{B}'_i \in \Delta_i$ that is ROI-restricted and let $\boldsymbol{B}_{-i} \in \Delta_{-i}$ be arbitrary. Then, $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$.*

LEMMA 3.11 (Lifting Lemma). *Let T be a set of types. Consider an instance $I \in \mathcal{I}^{T,\infty}_{\text{XOS}}$ and let $\boldsymbol{B} \in \Delta$ be a well-supported bid profile. Let (v^*_{ij}) be some opt-induced additive representatives. Assume that $FPA(r)$ is (λ_t, μ_t)-smooth for each type $t \in T$. Then, there exists a bid $\boldsymbol{B}'_i$ for every $i \in N$ satisfying $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$. Further, for each $t \in T$, it holds*

$$\sum_{i \in N_t(I)} \mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] \geq \sum_{j \in M : rw(j) \in N_t(I)} \lambda_t v^*_{\mathsf{rw}(j)j} - \mu_t \mathbb{E}[p_{\mathsf{aw}(j)j}(\boldsymbol{B})]. \tag{3.4}$$

Proof. Consider some agent $i \in N_t(I)$. If i is not the rightful winner of any auction, we define $\boldsymbol{B}'_i$ such that $B'_{ij} = 0$ deterministically. Clearly, $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$ holds. Otherwise, let i is the rightful winner of auction j, i.e., $i = \mathsf{rw}(j)$. By assumption, $FPA(r)$ is (λ_t, μ_t)-smooth for t. Thus, for each auction $j \in M$, there exists an ROI-restricted bid B'_{ij} such that, for each well-supported bid profile $\boldsymbol{b}_j$, we have:

$$\mathbb{E}[g_i(B'_{ij}, (\boldsymbol{b}_j)_{-i})] \geq \lambda_t v^*_{\mathsf{rw}(j)j} - \mu_t p_{\mathsf{aw}(j)j}(\boldsymbol{b}_j), \tag{3.5}$$

We define the random deviation $\boldsymbol{B}'_i$ of agent i for the global mechanism $\mathcal{M}$ simply by drawing a bid b'_{ij} for each auction $j \in M$ independently according to B'_{ij} if $i = \mathsf{rw}(j)$, and letting $B'_{ij} = b'_{ij} = 0$ deterministically if $i \neq \mathsf{rw}(j)$. For each bid profile $\boldsymbol{b}_{-i}$, we have:

$$\begin{aligned}\mathbb{E}[p_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i})] &= \mathbb{E}\left[\sum_{j \in M} p_{ij}(B'_{ij}, (\boldsymbol{b}_j)_{-i})\right] \leq \mathbb{E}\left[\sum_{j \in M} v^*_{ij} x_{ij}(B'_{ij}, (\boldsymbol{b}_j)_{-i})\right] \\ &= \mathbb{E}[v^*_i(x_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}))] \leq \mathbb{E}[v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}))].\end{aligned}$$

Here the inequality holds because B'_{ij} is ROI-restricted for each j (which also holds trivially for all auctions j with $i \neq \mathsf{rw}(j)$). The second equality follows by the definition of v^*_i, and the last inequality follows from property **XOS2**. We conclude that $\boldsymbol{B}'_i$ is ROI-restricted for each agent i. Therefore, by Corollary 3.10, $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$, proving the first part of the lemma.

We continue with the second part. Fix a type $t \in T$. Given any bid profile $\boldsymbol{b} \in \text{supp}(\boldsymbol{B})$, for every $i \in N_t(I)$ we have

$$\mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i})] = \sum_{j \in M} \mathbb{E}\left[g_i\left(B'_{ij}, (\boldsymbol{b}_j)_{-i}\right)\right] \geq \sum_{j \in M : \mathsf{rw}(j) = i} \lambda_t v^*_{\mathsf{rw}(j)j} - \mu_t p_{\mathsf{aw}(j)j}(\boldsymbol{b}). \tag{3.6}$$

Here the equality follows by linearity of expectation. The inequality follows by applying (3.5) to all auctions j such that $i = \mathsf{rw}(j)$, and using that the expected gain of i is non-negative for all j with $i \neq \mathsf{rw}(j)$. Note that $\boldsymbol{B}$ is well-supported and thus $\boldsymbol{x}(\boldsymbol{b}) \neq 0$ for each $\boldsymbol{b} \in \text{supp}(B)$. Taking expectations over $\boldsymbol{B}$ on both sides of (3.6) and summing over all $i \in N_t(I)$ yields (3.4). The claim follows. ∎

We can now present the proof of our Extension Theorem.

Proof of Theorem 3.5. Consider an instance $I \in \mathcal{I}_{\text{XOS}}^T$ and let $\boldsymbol{B}$ be a well-supported CCE of I. Let $\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, T^+)$ be an arbitrary calibration vector; note that such a vector exists by Lemma A.2 (given below). Let $\boldsymbol{x}^*(I)$ be an optimal allocation and (v_{ij}^*) the respective opt-induced additive representatives. Let $\hat{I}(I, \boldsymbol{B})$ be the budget-free proxy instance of $(I, \boldsymbol{B})$ as defined in Definition 3.7. For ease of notation, we write $\hat{N}_t := N_t(\hat{I}(I, \boldsymbol{B}))$ for every $t \in T^+$.

The instance $\hat{I}$ is budget-free by construction. Further, since $\hat{I}$ is a proxy of $I \in \mathcal{I}_{\text{XOS}}$, it follows from Proposition 3.6 that $\hat{I} \in \mathcal{I}_{\text{XOS}}^{\infty}$. Thus, we can apply Lemma 3.11 to $\hat{I} \in \mathcal{I}_{\text{XOS}}^{\infty}$ and $\boldsymbol{B}$, and obtain (from its first statement) that there exists a bid $\boldsymbol{B}_i'$ for each agent $i \in N$ such that $(\boldsymbol{B}_i', \boldsymbol{B}_{-i})$ satisfies the ROI constraint (2.2) for $\hat{I}$. Using this and the definition of budget-capped valuations, we obtain for each agent i:

$$\mathbb{E}[p_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})] \leq \mathbb{E}\left[v_i^{\mathfrak{B}_i}(x_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i}))\right] = \mathbb{E}\left[\min\left(v_i(x_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})), \mathfrak{B}_i\right)\right].$$

In particular, this shows that each $(\boldsymbol{B}_i', \boldsymbol{B}_{-i})$ satisfies the ROI constraint (2.2) and the budget constraint (2.3) for instance I, i.e., $(\boldsymbol{B}_i', \boldsymbol{B}_{-i}) \in \mathcal{R}_i$.

By exploiting (3.4) for instance $\hat{I}$, we obtain that for every $t \in T^+$ we have

$$\sum_{i \in \hat{N}_t} \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right] \geq \sum_{j \in M : \mathsf{rw}(j) \in \hat{N}_t} \lambda_t v_{\mathsf{rw}(j)j}^* - \mu_t \mathbb{E}\left[p_{\mathsf{aw}(j)j}(\boldsymbol{B})\right]. \tag{3.7}$$

Note that we exploit here that $\boldsymbol{x}^*(I)$ is also an optimal solution for $\hat{I} \in \mathcal{I}_{\text{XOS}}^{\infty}$ by Proposition 3.8. Thus, the opt-induced additive representatives with respect to $\boldsymbol{x}^*(I)$ of I and $\hat{I}$ are the same.

We therefore have:

$$\begin{aligned}
LW(I, \boldsymbol{B}) &\geq \sum_{t \in T} \sum_{i \in \hat{N}_t} \delta_t \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right] + (1 - \delta_t + \delta_t \hat{\sigma}_i(\boldsymbol{B}))\, \mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \sum_{t \in T^+} \sum_{i \in \hat{N}_t} \delta_t \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right] + (1 - \delta_t + \delta_t t)\, \mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&\geq \sum_{t \in T^+} \sum_{i \in \hat{N}_t} \delta_t \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right] + \left(1 - \max_{t \in T^+} \left(\delta_t(1 - t)\right)\right) \mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \sum_{t \in T^+} \sum_{i \in \hat{N}_t} \delta_t \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right] + \left(1 - \max_{t \in T^+} \left(\delta_t(1 - t)\right)\right) \mathbb{E}\left[\sum_{i \in N} p_i(\boldsymbol{B})\right]. \qquad (3.8)
\end{aligned}$$

The first inequality follows from applying Lemma 3.9 to each agent $i \in N$, using the deviation $\boldsymbol{B}_i'$ given by Lemma 3.11. Then, the first equality holds by the definition of T^+.

We now lower bound the first term in (3.8). Using Lemma 3.11, we obtain:

$$\begin{aligned}
\sum_{t \in T} \sum_{i \in \hat{N}_t} \delta_t \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right] &\geq \sum_{t \in T^+} \sum_{j \in M : \mathsf{rw}(j) \in \hat{N}_t} \delta_t \lambda_t v_{\mathsf{rw}(j)j}^* - \delta_t \mu_t \mathbb{E}\left[p_{\mathsf{aw}(j)j}(\boldsymbol{B})\right] \\
&\geq \sum_{t \in T^+} \sum_{j \in M : \mathsf{rw}(j) \in \hat{N}_t} \min_{t \in T^+} (\delta_t \lambda_t)\, v_{\mathsf{rw}(j)j}^* - \max_{t \in T^+} (\delta_t \mu_t)\, \mathbb{E}\left[p_{\mathsf{aw}(j)j}(\boldsymbol{B})\right] \\
&= \min_{t \in T^+} (\delta_t \lambda_t)\, OPT(I) - \max_{t \in T^+} (\delta_t \mu_t)\, \mathbb{E}\left[\sum_{i \in N} p_i(\boldsymbol{B})\right], \qquad (3.9)
\end{aligned}$$

where the last equality follows from property **XOS1**.

Substituting (3.9) into (3.8), we obtain:

$$LW(I, \boldsymbol{B}) \geq \min_{t \in T^+} (\delta_t \lambda_t)\, OPT(I) + \left(1 - \max_{t \in T^+} (\delta_t \mu_t) - \max_{t \in T^+} \left(\delta_t(1 - t)\right)\right) \mathbb{E}\left[\sum_{i \in N} p_i(\boldsymbol{B})\right] \geq \min_{t \in T^+} (\delta_t \lambda_t)\, OPT(I),$$

where the second inequality holds because $\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, T^+)$ (see Definition 3.4). By selecting a calibration vector $\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, T^+)$ that maximizes $\min_{t\in T^+}(\delta_t\lambda_t)$, we finally obtain

$$LW(I, \boldsymbol{B}) \geq \max_{\boldsymbol{\delta}\in\mathcal{C}(\boldsymbol{\mu},T^+)} \min_{t\in T^+} \delta_t\lambda_t OPT(I). \tag{3.10}$$

Since (3.10) holds for every instance $I \in \mathcal{I}_{\text{XOS}}$ and every well-supported $\boldsymbol{B} \in CCE(I)$, the proof follows. □

Note that, in the proof above, the whole purpose of our calibration vector was to eventually lower bound the total payments in the final expression by 0. This also explains the specific definition of the feasibility constraint of $\mathcal{C}(\boldsymbol{\mu}, T)$ in (3.3).

3.3 POA-Revealing Mathematical Program. We can now present our *POA-revealing mathematical program (POA-RMP)*, which facilitates bounding the price of anarchy as stated in our Extension Theorem (Theorem 3.5). The program is parameterized by the set T of available agent types. Recall that $\eta = \min_{j\in M}\eta_j \in [0, 1)$.

$$\text{POA-RMP}(T) \;=\; \max\; \min\left\{ \min_{t\in T}\lambda_t, \left(\max_{t\in T}\left(\frac{\mu_t}{\lambda_t}\right) + \max_{t\in T}\left(\frac{1-t}{\lambda_t}\right)\right)^{-1}\right\}$$

$$\text{s.t.}\quad \lambda_t = \mu_t\left(1 - \frac{1-\eta}{e^{1/\mu_t}}\right) \qquad \mu_t > 0 \qquad \forall t \in T\cap\{1\} \tag{3.11}$$

$$\lambda_t = \frac{\mu_t}{t}\left(1 - \frac{1-t\eta}{e^{t/\mu_t}}\right) \qquad \mu_t \geq t\left(\ln\left(\frac{1-t\eta}{1-t}\right)\right)^{-1} \qquad \forall t \in T\cap(0, 1) \tag{3.12}$$

$$\lambda_t = \mu_t \qquad \mu_t \in \left(0, \frac{1}{1-\eta}\right] \qquad \forall t \in T\cap\{0\} \tag{3.13}$$

A crucial building block in deriving our mathematical program (POA-RMP) is the characterization of optimal calibration vectors. By exploiting this characterization, the task of finding the best upper bound on the POA reduces to solving the above program. We summarize this result in the following theorem.

THEOREM 3.12. *Let $\mathcal{I}^T_{\text{XOS}}$ be the class of instances with fractionally subadditive valuations and type set T. Assume that FPA(r) is (λ_t, μ_t)-smooth for each type $t \in T^+$. Then, the price of anarchy of well-supported coarse correlated equilibria is upper bounded by POA-RMP$(T^+)^{-1}$.*

We use (POA-RMP) in the next section to derive (tight) bounds on the POA of CCE.

4 Liquid Welfare Guarantees Without Reserve Prices. In this section, we focus on simultaneous first price auctions without reserve prices, i.e., we assume that $\eta = 0$ in (3.11)–(3.13). In Section 4.1, we develop lower bounds on the optimal value of POA-RMP(T), which lead to POA upper bounds for various sets of types. Then, in Section 4.2, we present these liquid welfare guarantees upper bounds, and in multiple cases, complement our positive results with matching lower bounds.

4.1 Bounding POA-RMP(T) by Partitioning Agent Types. In this section, we characterize a feasible solution to POA-RMP(T) for a given set of types T. The main technical challenge is to identify an analytical solution that yields strong POA upper bounds. Below, we describe a policy for defining a solution vector $\boldsymbol{\mu}$ for POA-RMP(T).

Given a set of types T, the main idea is to partition them into two classes H_ω (high) and L_ω (low), where ω is a separation parameter. We then define μ_t depending on the class each $t \in T$ belongs to. Intuitively, H_ω contains agent types that are structurally close to utility maximizers, while L_ω contains agent types that are structurally close to value maximizers.

DEFINITION 4.1. *Given $\omega \in (0, 1)$ and a set of types T, define $H_\omega(T) = \{t \in T \mid t \geq \omega\}$ and $L_\omega(T) = \{t \in T \mid t < \omega\}$. Define $\boldsymbol{\mu}^*(\omega, T) \in \mathbb{R}^{|T|}_{>0}$ such that, for each $t \in T$,*

$$\mu^*_t(\omega, T) = \begin{cases} \frac{t}{-\ln(1-\omega)}, & \text{if } t \in H_\omega(T),\\ \frac{t}{-\ln(1-t)}, & \text{if } t \in L_\omega(T)\cap(0,1),\\ 1, & \text{if } t \in L_\omega(T)\cap\{0\}. \end{cases} \tag{4.1}$$

The following corollary is easy to verify.

COROLLARY 4.2. *Given $\omega \in (0,1)$ and a set of types T, $\boldsymbol{\mu}^*(\omega, T)$ is a feasible solution of POA-RMP(T).*

Lemma 4.3 will be useful when proving some of the price anarchy bounds that follow.

LEMMA 4.3. *Let T be a set of agent types. If $\max(T) > 0$, then for every $\omega \in (0, \max(T)] \cap (0,1)$,*

$$\textit{POA-RMP}(T) \geq \min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega + \max(T)}\right\}.$$

4.2 Price of Anarchy Bounds

4.2.1 Budget-Free Instances and Common Type. We first investigate the price of anarchy for budget-free instances with XOS valuations, assuming agents have a single type $t \in [0,1]$ i.e., the class $\mathcal{I}_{\text{XOS}}^{\{t\},\infty}$. In Theorem 4.4, we establish a liquid welfare guarantee for $\mathcal{I}_{\text{XOS}}^{\{t\},\infty}$, which interpolates smoothly from $e/e-1 \approx 1.58$ when $t = 1$ (all agents are utility maximizers) to 2 when $t = 0$ (all agents are value maximizers). This interpolation is illustrated in Figure 1.1(a).

THEOREM 4.4. *Let $\mathcal{I}_{\textit{XOS}}^{\{t\},\infty}$ be the class of budget-free instances with fractionally subadditive valuations and a single type $t \in [0,1]$. Then,*

$$\textit{CCE-POA}\left(\mathcal{I}_{\textit{XOS}}^{\{t\},\infty}\right) \leq \begin{cases} \frac{e}{e-1}, & \text{if } t \in [1 - 1/e, 1], \\ 1 - \frac{(1-t)\ln(1-t)}{t}, & \text{if } t \in (0, 1 - 1/e), \\ 2, & \text{if } t = 0. \end{cases}$$

Proof. Set $\omega = 1 - 1/e$. Let $\mu := \mu^*(\omega, \{t\})$ be as prescribed by Definition 4.1. We distinguish three cases based on the value of t.

Case 1: $t = 0$. In this case, $(\lambda, \mu) = (1,1)$ (by (3.13)) and POA-RMP($\{t\}$) $\geq \min\left(\lambda, \frac{\lambda}{\mu+1}\right) = \frac{1}{2}$.

Case 2: $t \in (0, 1 - 1/e)$. Similarly, $(\lambda, \mu) = \left(t(-\ln(1-t))^{-1}, t(-\ln(1-t))^{-1}\right)$ (by (3.12) and we obtain

$$\text{POA-RMP}(\{t\}) \geq \min\left(\lambda, \frac{\lambda}{\mu + 1 - t}\right) = \frac{\lambda}{\mu + 1 - t} = \frac{t}{t - (1-t)\ln(1-t)}.$$

Here, the first equality follows since $\mu = \frac{t}{-\ln(1-t)} \geq t$ holds for all $t \in (0, 1 - 1/e)$.

Case 3: $t \in [1 - 1/e, 1]$. We have that $(\lambda, \mu) = \left(\omega(-\ln(1-\omega))^{-1}, t(-\ln(1-\omega))^{-1}\right) = (1 - 1/e, t)$ (by (3.12) for $t < 1$ and (3.11) for $t = 1$). Furthermore,

$$\text{POA-RMP}(\{t\}) \geq \min\left(\lambda, \frac{\lambda}{\mu + 1 - t}\right) = \lambda = 1 - \frac{1}{e}.$$

Finally, by Theorem 3.12, we have that *CCE-POA*$(\mathcal{I}_{\text{XOS}}^{\{t\},\infty}) \leq (\text{POA-RMP}(\{t\}))^{-1}$. Combining this fact with the lower bounds on POA-RMP($\{t\}$) we obtained for each of the three cases above, the claim follows. ∎

Theorem 4.4 has a few important implications. Note that for $t = 1$, i.e., when agents are utility maximizers, the upper bound of Theorem 4.4 recovers the best possible price of anarchy bound of $\frac{e}{e-1}$ (due to [45]). We show that the same bound holds as long as $t \geq 1 - 1/e$. That is, somewhat surprisingly, the CCE-POA for this range of types does not get worse; in fact, this is true even for single-item first price auctions. The proof of Corollary 4.5 is implied by Theorem 5.3 (for $\eta = 0$) in Section 5.2.

COROLLARY 4.5. *Let $\mathcal{I}_{\textit{ADD}}^{\{t\},\infty}$ be the class of budget-free instances with additive valuations and a single type $t \in [0,1]$. If $t \geq 1 - 1/e$ then, CCE-POA$(\mathcal{I}_{\textit{ADD}}^{\{t\},\infty}) \geq \frac{e}{e-1}$ and the bound holds even for single-item auctions.*

Finally, for instances with value-maximizers only i.e., for $t = 0$, Theorem 4.4 extends the best possible upper bound of 2 to coarse correlated equilibria and XOS valuation functions [35, 22].

4.2.2 Budget-Constrained Agents and Heterogeneous Types. In Theorem 4.6, we make use of all the technical tools developed so far in this section to obtain a liquid welfare guarantee for coarse correlated equilibria and budget-constrained agents with XOS valuations for arbitrarily heterogeneous agent types. Recall that we denote by W_0 the principal branch of the Lambert W function Let $P : [0,1] \mapsto \mathbb{R}_{\geq 0}$ be defined as

$$P(z) = \begin{cases} 1 + \frac{z}{1+W_0(-e^{-z-1})}, & \text{if } z > 1 + \frac{W_0(-2e^{-2})}{2}, \\ 2, & \text{otherwise.} \end{cases} \tag{4.2}$$

Theorem 4.6. *Let $\mathcal{I}_{\text{XOS}}^T$ be the class of instances with fractionally subadditive valuations and type set $T \subseteq [0,1]$. Then, CCE-POA $\left(\mathcal{I}_{\text{XOS}}^T\right) \leq P(\max(T))$.*

Theorem 4.6 reveals an intriguing threshold phenomenon: the POA for a type set T remains at most 2 when $\max(T) < 0.797$, and increases from 2 to 2.1885 as $\max(T)$ approaches 1 (see also Figure 1.1(a)). Note that our liquid welfare guarantee unifies and generalizes two state-of-the-art POA bounds. Specifically, Theorem 4.6 recovers the upper bound of 2.1885 due to [22], for budget-free instances with additive valuation functions and mixed Nash equilibria under the mixed-agent model (i.e., $T = \{0, 1\}$). It also recovers the upper bound of 2 due to [36], for budget-constrained value maximizers (i.e., $T = \{0\}$) with additive valuations. Theorem 4.6 generalizes both results to coarse correlated equilibria and budget-constrained agents with XOS valuations, while simultaneously extending the type set to the general model in [1].

We now prove Theorem 4.6 using the following fact.

Fact 4.7. *Let $f(z) = 1 + W_0\left(-e^{-z-1}\right)$. For every $z \in \left(1 + W_0(-2e^{-2})/2, 1\right]$, we have $f(z) < z$ and $f(z) + z = -\ln(1 - f(z))$.*

Proof of Theorem 4.6. Observe that $\max(T^+) = \max(T \cup \{0\}) = \max(T)$. We distinguish three cases, based on the value of $\max(T)$.

Case 1: $\max(T) = 0$. In this case $T^+ = T = \{0\}$. We have:

$$\text{POA-RMP}(T^+) = \text{POA-RMP}(\{0\}) \geq \min\left\{\lambda_t, \left(\frac{\mu_t}{\lambda_t} + \frac{1}{\lambda_t}\right)^{-1}\right\} = \min\left\{1, (1+1)^{-1}\right\} = \frac{1}{2}. \tag{4.3}$$

Here, the second equality is due to (3.13) and (4.1).

Case 2: $\max(T) \in \left(0, 1 + W_0(-2e^{-2})/2\right]$. Set $\omega = \max(T)$. Then $\omega \in (0, \max(T)] \cap (0,1)$ (since $\omega = \max(T) < 1$), and we can invoke Lemma 4.3 for T^+ with $\omega = \max(T)$. Hence, we obtain that the value of the objective function of POA-RMP(T^+) is at least:

$$\min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega + \max(T)}\right\} = \min\left\{\frac{\max(T)}{-\ln(1-\max(T))}, \frac{1}{2}\right\} = \frac{1}{2}. \tag{4.4}$$

Here, the first and second equality follow by choice of $\omega = \max(T)$ as $\frac{z}{-\ln(1-z)} \geq \frac{1}{2}$ for all $z \leq 1 + W_0(-2e^{-2})/2$.

Case 3: $\max(T) \in (1 + W_0(-2e^{-2})/2, 1]$. Set $\omega = 1 + W_0\left(-e^{-\max(T)-1}\right)$. Using the first statement of Fact 4.7, we get that $\omega = f(\max(T)) < \max(T)$ whenever $\max(T) \in (1 + W_0(-2e^{-2})/2, 1]$. Therefore, $\omega \in (0, \max(T)] \cap (0,1)$ as $\max(T) \leq 1$ and, similarly to **Case 2**, we can invoke Lemma 4.3 for T^+ with $\omega = 1 + W_0\left(-e^{-\max(T)-1}\right)$. Hence, we obtain that the value of the objective function of POA-RMP(T^+) is at least:

$$\min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega + \max(T)}\right\} = \frac{\omega}{\omega + \max(T)} = \frac{1 + W_0\left(-e^{-\max(T)-1}\right)}{1 + W_0\left(-e^{-\max(T)-1}\right) + \max(T)}. \tag{4.5}$$

Here, the first equality holds by the second statement of Fact 4.7 and the second equality follows by our choice of ω.

We therefore conclude that *CCE-POA* $\left(\mathcal{I}_{\text{XOS}}^T\right) \leq \left(\text{POA-RMP}\left(T^+\right)\right)^{-1} \leq P(\max(T))$. The first inequality follows by Theorem 3.12 and the second inequality by (4.3), (4.4) and (4.5). ∎

We now show that the liquid welfare guarantee of Theorem 4.6 is best possible by providing two matching lower bounds for simple classes of instances. The lower bound of Theorem 4.8 matches the upper bound of Theorem 4.6 for CCE, even when the agents have a single type $t \in [0,1]$ and additive valuation functions.

THEOREM 4.8. *Let $\mathcal{I}_{ADD}^{\{t\}}$ be the class of instances with additive valuations and a single type $t \in [0,1]$. Then, CCE-POA$(\mathcal{I}_{ADD}^{\{t\}}) \geq P(t)$.*

Theorem 4.8 has one additional implication for the inefficiency of CCE of simultaneous first price auctions and fractionally subadditive valuations. By combining Theorem 4.8 and Theorem 4.4, we obtain a *separation* in terms of liquid welfare guarantees for environments with a single type $t > 0$. Namely, we show that budget-constraints make the price of anarchy *strictly worse* in this case (see Figure 1.1(a)).

COROLLARY 4.9. *For every type $t \in T$, let $\mathcal{I}_{XOS}^{\{t\}}$ be the class of instances with fractionally subadditive valuations and let $\mathcal{I}_{XOS}^{\{t\},\infty}$ be the subclass of budget-free instances. For $t > 0$, it holds that CCE-POA$(\mathcal{I}_{XOS}^{\{t\}}) >$ CCE-POA$(\mathcal{I}_{XOS}^{\{t\},\infty})$.*

We continue with our second negative result. Here, we show that the liquid welfare guarantee in Theorem 4.6 is also tight for the class of budget-free instances, additive valuations and budget-free instances, even for mixed Nash equilibria, as long as value maximizers are included in the type set T. This lower bound is a generalization of the one in [22] for a mixed agent model with value maximizers and a different type.

THEOREM 4.10. *Let $\mathcal{I}_{ADD}^{\{0,t\},\infty}$ be the class of budget-free instances with additive valuations and a set of agent types $\{0,t\}$ for $t \in (0,1]$. Then, MNE-POA$(\mathcal{I}_{ADD}^{\{0,t\},\infty}) \geq P(t)$.*

Note that Theorem 4.10 together with Theorem 4.6 settle the price of anarchy of all equilibrium classes and all budget-constrained instances for agent type sets that include the value-maximizing type $t = 0$.

COROLLARY 4.11. *For VAL $\in \{$ADD, SUB, XOS$\}$, let $\mathcal{I}_{VAL}^{T}$ be the class of instances with valuations in $\mathcal{V}_{VAL}$ and let T be a set of types. If $0 \in T$, then, for EQ $\in \{$MNE, CE, CCE$\}$ we have EQ-POA$(\mathcal{I}_{VAL}^{T}) = P(\max(T))$.*

We conclude this section with Theorem 4.12, which states a slightly weaker POA lower bound for MNE with budget-constrained agents and additive valuation functions. However, this bound holds for an arbitrary type set T (not necessarily including $t = 0$).

THEOREM 4.12. *For every type set T, it holds that MNE-POA$(\mathcal{I}_{ADD}^{T}) \geq 2$.*

Closing the gap between 2 and $P(\max(T))$ for MNE, for every type set T, is an intriguing open question.

4.2.3 Bounded Minimum Type. A qualitative interpretation of the upper bound in Theorem 4.6 is that equilibria become more inefficient in the presence of agent types whose types resemble utility maximizers—i.e., as t approaches 1, the POA worsens. One contributing factor to this phenomenon is that, in the worst case, the competitors of these agents may, on the contrast, be structurally similar to value maximizers. A natural question is to study mixtures of types in which the minimum type is bounded away from 0. In Theorem 4.13, we present a second threshold phenomenon revealed by our framework: whenever a type set T satisfies $\min(T) \geq 0.74$, the liquid welfare guarantees for CCE of budget-free instances improve, no matter how heterogeneous the set of types is. In fact, for such type sets, *CCE-POA*$(\mathcal{I}_{XOS}^{T,\infty}) \leq 1.83$!

THEOREM 4.13. *Let β be the solution to $\beta = 1 - e^{-\frac{1}{\beta}}$, i.e., $\beta \approx 0.741$. Let $\mathcal{I}_{XOS}^{T,\infty}$ be the class of budget-free instances with fractionally subadditive valuations and type set T. If $\min(T) \geq \beta$, then:*

$$\text{CCE-POA}(\mathcal{I}_{XOS}^{T,\infty}) \leq \left(\min(T)\left(1 - e^{-\frac{1}{\min(T)}}\right)\right)^{-1} \in \left[\frac{e}{e-1}, \frac{1}{\beta^2}\right]$$

5 Improved Liquid Welfare Guarantees with Reserve Prices. Our type-dependent smoothness framework and the POA-RMP developed in Section 3 allow us to study the inefficiency of instances with fractionally subadditive valuations and their induced well-supported equilibria when the auctioneer implements reserve prices for each individual auction (see also Section 2). Namely, as already hinted by our formulation of POA-RMP(T) in Section 3.3, our liquid welfare guarantees for instances with reserve prices depend on the minimum relative gap $\eta \in [0,1)$.

We remark that, as observed by Balseiro et al. [7], a reserve price r_j can be interpreted as a *prediction* (see also [28, 15, 19]) of the value of the rightful winner of an auction $j \in M$. These predictions can then be used to set reserve prices accordingly, with the goal of improving liquid welfare guarantees. In this context, the parameter η can be viewed as an error measure of the prediction: $\eta = 0$ indicates a completely uninformative prediction,[7] whereas as $\eta \to 1$, the reserve prices approach the actual valuations of the rightful winners in all auctions.

Section 5 is structured as follows. In Section 5.1, we devise upper bounds on the POA that are functions of η for well-supported equilibria of a given class. Then, in Section 5.2, we examine when such equilibria are guaranteed to exist. Finally, in Section 5.3, we demonstrate that when budget-free agents repeatedly participate in a first-price auction with reserve price using regret-minimizing algorithms, they converge to CCE that are well-supported.

5.1 POA Bounds as Functions of the Minimum Relative Gap. In the presence of the η parameter, deriving analytical POA bounds for a general type set T becomes significantly more challenging. In order to obtain upper bounds on the POA as functions of η, we can no longer rely on our approach from Section 4.1. Therefore, in this section, we focus on more tractable settings, such as the budget-free single-type environment and the mixed-agent model with budget-constrained agents.

5.1.1 Budget-Free Agents with One Type. We present our POA upper bound for the single type environment $\{t\}$ for every $t \in [0, 1]$, fractionally subadditive valuations and well-supported CCE in Theorem 5.1. As $\eta \to 1$, the liquid welfare of all such CCE tends to optimality; see Figure 1.1b for an illustration.

THEOREM 5.1. *Let $\mathcal{I}_{\text{XOS}}^{\{t\},\infty}$ be the class of budget-free instances with fractionally subadditive valuations, a single type $t \in [0, 1]$, and let $\eta \in [0, 1)$ be the smallest relative gap of the reserve prices. Also, let:*

$$P_t(\eta) = \begin{cases} \frac{e}{e-1+t\eta}, & \text{if } t \in (1 - 1/e, 1] \text{ and } \eta \in \left[0, \frac{1-e(1-t)}{t}\right), \\ 1 + \frac{1}{t}\left(\ln\left(\frac{1-t\eta}{1-t}\right)(1-t)\right), & \text{if } t \in (0, 1) \text{ and } \eta \in \left(\max\left(0, \frac{1-e(1-t)}{t}\right), 1\right), \\ 2 - \eta, & \text{if } t = 0. \end{cases} \tag{5.1}$$

For well-supported coarse correlated equilibria, it holds that CCE-POA($\mathcal{I}_{\text{XOS}}^{\{t\},\infty}$) $\leq P_t(\eta)$. Furthermore, P_t is non increasing in $[0, 1)$ with $\lim_{z\to 1} P_t(z) = 1$ for every $t \in [0, 1]$.

We now present two lower bounds for CCE in budget-free instances with reserve prices in single-type environments. In Theorem 5.2, we show that the bound of Theorem 5.1 is tight for all $\eta \in [0, 1)$ in auctions with value maximizers only, i.e., when $t = 0$. Then, in Theorem 5.3, we establish a negative result for a different restricted range of (η, t), including the case of utility maximizers ($t = 1$) for which we prove that the bound of Theorem 5.1 is tight for all $\eta \in [0, 1)$. Interestingly, the worst-case instances used in the proof of these theorems involve a two-item and single-item auction.

THEOREM 5.2. *Let $\mathcal{I}_{\text{ADD}}^{t,\infty}$ be the class of budget-free instances with additive valuations for type $t = 0$ and let $\eta \in [0, 1)$ be the smallest relative gap from reserve prices in $\mathcal{I}_{\text{ADD}}^{\{0\},\infty}$. Then, CCE-POA($\mathcal{I}_{\text{ADD}}^{\{0\},\infty}$) $\geq 2 - \eta$ for well-supported coarse correlated equilibria.*

THEOREM 5.3. *Let $\mathcal{I}_{\text{ADD}}^{\{t\},\infty}$ be the class of budget-free instances with additive valuations for type $t \in [1 - 1/e, 1]$ and let $\eta \in [0, 1)$ be the smallest relative gap from reserve prices. Then, if $\eta \leq \frac{1-(e-1)t}{t}$, CCE-POA($\mathcal{I}_{\text{ADD}}^{\{0\},\infty}$) $\geq \frac{e}{e-1+\eta}$ for well-supported coarse correlated equilibria.*

COROLLARY 5.4. *Let $\eta \in [0, 1)$ be the smallest relative gap from reserve prices. Then, for $t \in \{0, 1\}$ it holds that CCE-POA($\mathcal{I}_{\text{XOS}}^{\{t\},\infty}$) = CCE-POA($\mathcal{I}_{\text{ADD}}^{\{t\},\infty}$) = $P_t(\eta)$.*

[7]Uninformative for improving liquid welfare guarantees: almost all η_j's may be close to 1, but one auction j with $\eta_j = 0$ can prevent improved guarantees (see also the proof of Theorem 5.2).

5.1.2 Budget-Constrained Agents in the Mixed Agent Model. We now focus on liquid welfare guarantees for $T = \{0, 1\}$, i.e., the mixed-agent model in which agents are either utility or value maximizers, under budget constraints and with reserve prices. This setting has previously been studied with budget-free, mixed Nash equilibria, and additive valuations [22]. In Theorem 5.5, we establish a POA upper bound that depends on the parameter η for budget-constrained agents, CCE, and fractionally subadditive valuations. As the parameter η increases from 0 to 1, the resulting bound interpolates between our liquid welfare guarantee of 2.1885 from Theorem 4.6 and the optimal value of 1.

THEOREM 5.5. *Let $\mathcal{I}_{\text{XOS}}^{\{0,1\},\infty}$ be the class of instances with fractionally subadditive valuations for the set of types $\{0, 1\}$. Let $\eta \in [0, 1)$ be the smallest relative gap from reserve prices. Then, for well-supported equilibria, CCE-POA$(\mathcal{I}_{\text{XOS}}^{\{0,1\}}) \le Q(\eta)$, where*

$$Q(\eta) = (1 - \eta) \cdot \frac{2 - \eta + W_0\left(-(1 - \eta)^2 e^{\eta-2}\right)}{1 - \eta + W_0\left(-(1 - \eta)^2 e^{\eta-2}\right)}.$$

Furthermore, $Q(\eta)$ is non-increasing in $[0, 1)$ with $Q(0) \approx 2.1885$ and $\lim_{z\to1} Q(z) = 1$.

5.2 On the Existence of Well-Supported Equilibria. Recall that for an instance I with reserve prices $\boldsymbol{r}$, a bid profile $\boldsymbol{B} \in \Delta$ is well-supported if, for each $\boldsymbol{b} \in \text{supp}(\boldsymbol{B})$, it holds that $\boldsymbol{x}_j(\boldsymbol{b}) \ne \boldsymbol{0}$ for every item $j \in M$. Throughout our work, we have applied this refinement on equilibria when considering liquid welfare guarantees for instances with reserve prices, as, in the presence of reserve prices, it is a crucial precondition of our Extension Theorem (Theorem 3.12) in Section 3. In this section, we explore when such equilibria are guaranteed to exist.

5.2.1 Budget-Free Instances and Additive Valuations. We begin with a positive result for the class $\mathcal{I}_{\text{ADD}}^{T,\infty}$ given a type set T: namely, in Theorem 5.6, we show that all *correlated* equilibria induced on instances of this class are well-supported.

THEOREM 5.6. *Let T be a set of agent types and let $I \in \mathcal{I}_{\text{ADD}}^{T,\infty}$. Then, every $\boldsymbol{B} \in$ CE(I) is well-supported.*

Intuitively, for agents with additive valuations, the rightful winner i of an auction j always has "room" for additional gain by bidding above the reserve price r_j and competing for the item (unless of course some other agent has already submitted a sufficiently high bid above r_j). In Theorem 5.6, we confirm this intuition for correlated equilibria. Interestingly, we observe that it does not necessarily hold for coarse correlated equilibria; in Theorem 5.7, we construct a simple single-item auction with two utility maximizers and a feasible reserve price, along with a CCE that is not well-supported.

THEOREM 5.7. *Let $\mathcal{I}_{\text{ADD}}$ be the class of budget-free instances with additive valuation functions and reserve prices. There exists an instance $I \in \mathcal{I}_{\text{ADD}}$ and a $\boldsymbol{B} \in$ CCE(I) which is not well-supported for I.*

The intuition behind why such CCE exist even in simple settings is that it can be more cost-effective for the highest-valued agent to remain *coordinated* in their bidding (through the CCE probability distribution) which might mean that they sometimes bid below the reserve price with positive probability, rather than deviating unilaterally and always bidding at least the reserve price which can lead to a higher expected payment. Note that the instance in Theorem 5.7 also implies a lower bound on the POA of $\frac{e}{e-1}$ for CCE that are not necessarily well-supported. We therefore conclude that our inability to devise improved liquid welfare guarantees as parameters of $\eta \in [0, 1)$ for such equilibria (similar to those of Theorem 5.1 and Theorem 5.5) is not an artifact of our analysis (e.g., in Theorem 3.5), but rather an unavoidable structural property of CCE.

COROLLARY 5.8. *Let $\mathcal{I}_{\text{ADD}}$ be the class of instances with additive valuations and let $\eta \in [0, 1)$ be the smallest relative gap from reserve prices. Then, CCE-POA$(\mathcal{I}_{\text{ADD}}) \ge \frac{e}{e-1}$.*

5.2.2 Beyond Additive Valuations. We conclude the section with another negative result. We show that, for the class of budget-free instances with submodular valuations, even *mixed Nash equilibria* are not guaranteed to be well-supported. Moreover, this fact paints a bleak picture for their liquid welfare guarantees: without refining the set of MNE to well-supported equilibria only, the POA for MNE for instances with reserve prices is unbounded.

THEOREM 5.9. *Let $\mathcal{I}_{\text{SUB}}$ be the class of budget-free instances with submodular valuation functions and reserve prices. There exists an instance $I \in \mathcal{I}_{\text{SUB}}$ and a $\boldsymbol{B} \in$ MNE(I) which is not well-supported for I. Furthermore, POA-MNE$(\mathcal{I}_{\text{SUB}}) = \infty$.*

5.3 Mean-Based Algorithms Converge to Well-Supported CCE. It is well known that regret-minimization dynamics in auctions lead to CCE (see, e.g., [30, 11, 47]). In this section, motivated by the negative result in Theorem 5.7, which shows that there exist CCE that are not well-supported even in single-item first-price auctions with agents with no budget constraints, we address the question of whether the CCE reached through such dynamics are well-supported.

We consider a model similar to the repeated auction setting studied in [32]. Specifically, we consider two budget-free agents[8] with arbitrary types who repeatedly participate in a first-price auction with a feasible reserve price r, where ties are broken uniformly at random. Each agent i is assumed not to overbid, i.e., $b_i \leq v_i$ (as we can still assume w.l.o.g. that $\tau = 1$). The agents' values and types remain fixed across all repetitions of the auction. We assume that values, bids, and the reserve price r are all integer multiples of a minimum increment $\varepsilon > 0$. Each agent aims to maximize their cumulative gain g_i over time.

Agents use regret-minimization algorithms, where the regret of agent i after T rounds, given bids $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_T)$, is defined as

$$R_i^T = \sum_{t=1}^{T} \max_b g_i(b, \boldsymbol{b}_{-i}^t) - g_i(b_i^t, \boldsymbol{b}_{-i}^t),$$

with b denoting the optimal fixed bid in hindsight. We introduce the definition of a weakly dominated action below.

Definition 5.10. *Let A_1 and A_2 be subsets of the action spaces of two agents in a two-agent game. An action $i \in A_1$ of agent 1 is* weakly dominated *in A_2 by another action $i' \in A_1$ if:*

1. *$\forall j \in A_2 : g_1(i, j) \leq g_1(i', j)$, and*
2. *$\exists j \in A_2 : g_1(i, j) < g_1(i', j)$.*

Kolumbus and Nisan [32] introduce a specific subclass of CCE called *co-undominated*. In such equilibria, no action in an agent's support is weakly dominated relative to the other agent's support.

Definition 5.11 ([32]). *Let $\boldsymbol{B}$ be a CCE of a (finite) two-agent game with action spaces I_1 and I_2. Let its support be (A_1, A_2), where $A_1 = \{i \in I_1 \mid \exists j \in I_2$ such that $B_{ij} > 0\}$ and $A_2 = \{j \in I_2 \mid \exists i \in I_1$ such that $B_{ij} > 0\}$. The CCE is* co-undominated *if, for every $i \in A_1$ and every $i' \in I_1$, action i is not weakly dominated in A_2 by i', and similarly for A_2.*

If regret-minimizing algorithms[9] converge, Kolumbus and Nisan [32] show that they converge to co-undominated CCE. This convergence result also holds in our setting of a first-price auction with a feasible reserve price and heterogeneous agent types. Moreover, co-undominated CCE possess the desired property of being well-supported, as we show in the next theorem.

Theorem 5.12. *Consider a first-price auction with feasible reserve price among two agents with arbitrary types and action spaces I_1 and I_2. Then, any co-undominated CCE is well-supported.*

Proof. Let $\boldsymbol{B}$ be a co-undominated CCE of a $FPA(r)$ with a feasible reserve price among two agents with arbitrary types and action spaces I_1 and I_2. For contradiction, assume that $\boldsymbol{B}$ is not well-supported, i.e., the item is not sold with probability 1. Denote the support (A_1, A_2) of $\boldsymbol{B}$ by $A_1 = \{i \in I_1 \mid \exists j \in I_2$ such that $B_{ij} > 0\}$ and $A_2 = \{j \in I_2 \mid \exists i \in I_1$ such that $B_{ij} > 0\}$. Assume w.l.o.g. that $v_1 \geq v_2$.

Since the item is not sold with probability 1, there exist actions $i \in A_1$ and $j \in A_2$ such that $i < r$, $j < r$, and $B_{ij} > 0$. In this case, we observe that action i of agent 1 is weakly dominated in A_2 by action r, since for all $j' \in A_2$ it holds that $0 = g_1(i, j') \leq g_1(r, j')$, and $0 = g_1(i, j) < g_1(r, j) = v_1 - \sigma_1 r$, as $r < v_1$ and $\sigma_1 \leq 1$. However, by assumption, $i \in A_1$ i.e., it is in the support of the CCE $\boldsymbol{B}$. This contradicts Definition 5.12 and the claim follows. ∎

Theorem 5.12 indicates that, when autobidding agents converge to a CCE using regret-minimizing algorithms, they reach a well-supported CCE. Hence, for each such CCE, the liquid welfare guarantees we obtained in Section 5.1 apply.

[8]We focus on two agents for simplicity, though the result extends to more than two agents.

[9]In particular, Kolumbus and Nisan [32] focus on a family of algorithms called *mean-based*; see [12] for a definition.

Appendix

A Missing Material of Section 3

A.1 Proof of Lemma 3.2

LEMMA 3.2. *Consider a single-item instance* $I = (r, \boldsymbol{v}, \boldsymbol{\sigma}, \mathfrak{B})$ *and let the rightful winner* $i = \mathsf{rw}$ *be of type* $t = 0$. *Then, FPA*(r) *is* (λ_t, μ_t)*-smooth for type* t *with* $\lambda_t = \mu_t = \mu$ *for every* $\mu \in (0, (1-\eta)^{-1}]$.

Proof of Lemma 3.2. Assume that the rightful winner i is of type $t = 0$. We need to show that there exists a ROI-restricted random bid B'_i such that for every well-supported bid profile $\boldsymbol{b}$ and $\mathsf{aw} = \mathsf{aw}(\boldsymbol{b})$ the actual winner, it holds that:

$$\mathbb{E}\left[g_i\left(B'_i, \boldsymbol{b}_{-i}\right)\right] \geq \mu v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}). \tag{A.1}$$

Let $B'_i = B'_i(\boldsymbol{v})$ be a random unilateral deviation of i drawn from $[\eta v_i, v_i]$ with CDF $F(z) = F_{B'_i}(z) = \mu z/v_i + 1 - \mu$. Note that the domain is well-defined as $\eta \in [0, 1)$, and it is easy to verify that $F(\cdot)$ is non-negative and increasing over $[\eta v_i, v_i]$ and $F(v_i) = 1$. Also, B'_i is ROI-restricted as the condition is even pointwise satisfied, i.e., for every $z \in [\eta v_i, v_i]$ it holds that $p_i(z, \cdot) \leq v_i(\boldsymbol{x}_i(z, \cdot))$.

It remains to show that B'_i satisfies (A.1). Note that the expected gain of i is always non-negative, as i bids above v_i with 0 probability. Thus, (A.1) holds trivially if $v_i \leq p_{\mathsf{aw}}(\boldsymbol{b})$. Therefore, assume that $v_i > p_{\mathsf{aw}}(\boldsymbol{b})$ and define $\theta_i := \max(\eta v_i, \max_{j \neq i} b_j)$. For every $z \geq \theta_i$, i wins the item under the bid profile $(z, \boldsymbol{b}_{-i})$ and pays $p_i(z, \boldsymbol{b}_{-i}) = z$. As the item is sold under the bid profile $\boldsymbol{b}$ by assumption, the actual winner under $\boldsymbol{b}$ either pays the reserve price or their maximum bid, i.e., $p_{\mathsf{aw}}(\boldsymbol{b}) = \max(\eta v_i, \max_j b_j)$. We obtain:

$$\theta_i = \max\left(\eta v_i, \max_{j \neq i} b_j\right) \leq \max\left(\eta v_i, \max_{j} b_j\right) = p_{\mathsf{aw}}(\boldsymbol{b}) < v_i.$$

This leads to the desired result as:

$$\begin{aligned}\mathbb{E}\left[g_i\left(B'_i, \boldsymbol{b}_{-i}\right)\right] &= v_i(1 - F(\theta_i)) = v_i\left(1 - \left(\frac{\mu\theta_i}{v_i} + 1 - \mu\right)\right) \\ &= \mu v_i - \mu\theta_i \geq \mu v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}).\end{aligned}$$

Note that the first equality holds because the sensitivity of i is $\sigma_i = t = 0$. ◻

A.2 Proof of Lemma 3.3

LEMMA 3.3. *Consider a single-item instance* $I = (r, \boldsymbol{v}, \boldsymbol{\sigma}, \mathfrak{B})$ *and let the rightful winner* $i = \mathsf{rw}$ *be of type* $t \in (0, 1]$. *Then, FPA*(r) *is* (λ_t, μ_t)*-smooth for type* t *with*

$$\lambda_t = \frac{\mu}{t}\left(1 - \frac{1 - t\eta}{e^{t/\mu}}\right) \quad \textit{and} \quad \mu_t = \mu \quad \textit{for every} \quad \begin{cases}\mu \geq t\left(\ln\left(\frac{1-t\eta}{1-t}\right)\right)^{-1}, & \textit{if } \ t < 1, \\ \mu > 0, & \textit{if } \ t = 1.\end{cases} \tag{3.2}$$

We elaborate on the expression of (3.2) in Lemma 3.3. First, note that $t \in (0, 1]$, as $t > 0$ and $t \leq 1$ by assumption. The $\ln(\cdot)$ expression decreases as η increases and converges to 0 (from above) as $\eta \to 1$; the lower bound restriction on μ thus increases as η increases. As $\eta < 1$, the $\ln(\cdot)$ expression is well-defined for all combinations of t and η, except when $t = 1$. In the latter case, we only impose the restriction that $\mu > 0$.

Given μ, we define a parameter $\gamma = \gamma(\mu)$ as follows:

$$\gamma(\mu) = \frac{1}{t}\left(1 - \frac{1 - t\eta}{e^{t/\mu}}\right), \tag{A.2}$$

which will be useful in the proof of Lemma 3.3. Note that γ is well-defined because $t > 0$ by assumption.

The following corollary is an immediate consequence of the definitions above.

COROLLARY A.1. *Let* μ *satisfy* (3.2) *and let* γ *be defined as in* (A.2). *Then,* $\gamma \in [\eta, 1]$.

Proof. Note that the interval $[\eta, 1]$ is well-defined because $\eta \in [0, 1)$ by assumption. We first prove the lower bound on γ. Note that $e^{t/\mu} > 1$ as $\mu > 0$ and $t > 0$, and therefore:

$$\gamma = \frac{1}{t}\left(1 - \frac{1 - t\eta}{e^{t/\mu}}\right) > \frac{1}{t} \cdot t\eta = \eta.$$

For the upper bound on γ, we have that:

$$\gamma = \frac{1}{t}\left(1 - \frac{1 - t\eta}{e^{t/\mu}}\right) \leq \frac{1}{t}\left(1 - \frac{1 - t\eta}{e^{\ln\left(\frac{1-t\eta}{1-t}\right)}}\right) = \frac{1}{t} \cdot t = 1,$$

where the inequality follows from (3.2) and because e^x is non-decreasing in x. ∎

We now continue with the proof of Lemma 3.3.

Proof of Lemma 3.3. Assume that the rightful winner i is of type $t \in T$ with $\sigma_i = t > 0$. We need to show that there exists a ROI-restricted random bid B'_i such that for every well-supported bid profile $\boldsymbol{b}$ and $\mathsf{aw} = \mathsf{aw}(\boldsymbol{b})$ the actual winner, it holds that:

$$\mathbb{E}\left[g_i\left(B'_i, \boldsymbol{b}_{-i}\right)\right] \geq \mu\gamma v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}). \tag{A.3}$$

Let $B'_i = B'_i(\boldsymbol{v})$ be a random unilateral deviation of i drawn from $[\eta v_i, \gamma v_i]$ with PDF $f(z) = f_{B'_i}(z) = \mu/(v_i - tz)$. Note that the domain is well-defined as $\gamma \in [\eta, 1]$ by Corollary A.1, and that $f(\cdot)$ is non-negative. Also note that:

$$\int_{\eta v_i}^{\gamma v_i} f(z)dz = \int_{\eta v_i}^{\gamma v_i} \frac{\mu}{v_i - tz}dz = \mu \int_{\eta v_i}^{\gamma v_i} \left(\frac{-\ln(v_i - tz)}{\sigma}\right)' dz$$
$$= \frac{\mu}{t}\ln\left(\frac{1 - t\eta}{1 - t\gamma}\right) = \frac{\mu}{t}\ln\left(e^{t/\mu}\right) = 1,$$

where the fourth equality follows from the definition of γ in (A.2). Furthermore, B'_i is ROI-restricted as the condition is even pointwise satisfied, i.e., for $z \in [\eta v_i, \gamma v_i]$ it holds that $p_i(z, \cdot) \leq \gamma v_i(\boldsymbol{x}_i(z, \cdot)) \leq v_i(\boldsymbol{x}_i(z, \cdot))$, as $\gamma \in [\eta, 1]$ by Corollary A.1.

It remains to show that B'_i satisfies (A.3). Note that the expected gain of i is always non-negative, as i bids above v_i with 0 probability and $t \leq 1$. Thus, (A.3) holds trivially if $\gamma v_i \leq p_{\mathsf{aw}}(\boldsymbol{b})$. Therefore, assume that $\gamma v_i > p_{\mathsf{aw}}(\boldsymbol{b})$ and define $\theta_i := \max(\eta v_i, \max_{j \neq i} b_j)$. Then, for every $z \geq \theta_i$, i wins the item under bid profile $(z, \boldsymbol{b}_{-i})$ and pays $p_i(z, \boldsymbol{b}_{-i}) = z$. As the item is sold under bid profile $\boldsymbol{b}$ by assumption, the actual winner under $\boldsymbol{b}$ either pays the reserve price or their maximum bid, i.e., $p_{\mathsf{aw}}(\boldsymbol{b}) = \max(\eta v_i, \max_j b_j)$. We obtain:

$$\theta_i = \max\left(\eta v_i, \max_{j \neq i} b_j\right) \leq \max\left(\eta v_i, \max_j b_j\right) = p_{\mathsf{aw}}(\boldsymbol{b}) < \gamma v_i.$$

This leads to the desired result as:

$$\mathbb{E}\left[g_i\left(B'_i, \boldsymbol{b}_{-i}\right)\right] = \int_{\theta_i}^{\gamma v_i} (v_i - tp_i(z, \boldsymbol{b}_{-i}))f(z)dz = \int_{\theta_i}^{\gamma v_i} (v_i - tz)f(z)dz$$
$$= \int_{\theta_i}^{\gamma v_i} \mu dz = \mu\gamma v_i - \mu\theta_i \geq \mu\gamma v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}).$$

Note that the first equality holds because the sensitivity of i is $\sigma_i = t$ by the precondition of the Lemma. ∎

A.3 Proofs of Proposition 3.8, Lemma 3.9, and Corollary 3.10

PROPOSITION 3.8. *Let $I = (\boldsymbol{r}, \boldsymbol{v}, \boldsymbol{\sigma}, \boldsymbol{\mathfrak{B}})$ and $\boldsymbol{B} \in \Delta$. Then, $OPT(\hat{I}(I, \boldsymbol{B})) = OPT(I)$.*

Proof of Proposition 3.8. Using Definition 3.7 and the definition of budget-capped valuations, we obtain:

$$OPT(\hat{I}(I, \boldsymbol{B})) = \max_{\boldsymbol{x} \in \boldsymbol{X}} \sum_{i \in N} v_i^{\mathfrak{B}_i}(\boldsymbol{x}_i) = \max_{\boldsymbol{x} \in \boldsymbol{X}} \sum_{i \in N} \min\left(v_i(\boldsymbol{x}_i), \mathfrak{B}_i\right) = OPT(I).$$

∎

LEMMA 3.9. *Consider an instance $I \in \mathcal{I}_{\text{XOS}}$ and let $\boldsymbol{B} \in CCE(I)$. Then, for every agent $i \in N$, for every $\boldsymbol{B}'_i$ with $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$ and every $\delta \in [0,1]$ it holds*

$$\min\left(\mathbb{E}[v_i(x_i(\boldsymbol{B}))], \mathfrak{B}_i\right) \geq \delta \cdot \mathbb{E}\left[\hat{g}_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] + (1 - \delta + \delta\hat{\sigma}_i(\boldsymbol{B})) \cdot \mathbb{E}\left[p_i(\boldsymbol{B})\right].$$

Proof of Lemma 3.9. Fix an agent $i \in N$. We distinguish two cases.

Case 1: $\mathfrak{B}_i < \mathbb{E}[v_i(x_i(\boldsymbol{B}))]$. We have:

$$\begin{aligned}
\min\left(\mathbb{E}[v_i(x_i(\boldsymbol{B}))], \mathfrak{B}_i\right) &= \delta\mathfrak{B}_i + (1-\delta)\mathfrak{B}_i \\
&\geq \delta\,\mathbb{E}\left[\min\left(v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})), \mathfrak{B}_i\right)\right] + (1-\delta)\mathfrak{B}_i \\
&= \delta\,\mathbb{E}\left[v_i^{\mathfrak{B}_i}(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}))\right] + (1-\delta)\mathfrak{B}_i \\
&\geq \delta\,\mathbb{E}\left[v_i^{\mathfrak{B}_i}(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}))\right] + (1-\delta)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \delta\,\mathbb{E}\left[v_i^{\mathfrak{B}_i}(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})) - \hat{\sigma}_i(\boldsymbol{B})p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] + (1-\delta+\hat{\sigma}_i(\boldsymbol{B})\delta)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \delta\,\mathbb{E}\left[\hat{g}_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] + (1-\delta+\delta\hat{\sigma}_i(\boldsymbol{B}))\mathbb{E}\left[p_i(\boldsymbol{B})\right].
\end{aligned}$$

The first equality follows by the definition of **Case 1**, and the second equality follows by the definition of budget-capped valuations. The second inequality follows since, by assumption, $\boldsymbol{B}$ is a CCE for I and therefore satisfies the budget constraint in (2.3) for instance I. Finally, the third equality follows since $\hat{\sigma}_i(\boldsymbol{B}) = 0$ holds by Definition 3.7.

Case 2: $\mathfrak{B}_i \geq \mathbb{E}[v_i(x_i(\boldsymbol{B}))]$. In this case, we have:

$$\begin{aligned}
\min\left(\mathbb{E}[v_i(x_i(\boldsymbol{B}))], \mathfrak{B}_i\right) &= \delta\,\mathbb{E}[v_i(x_i(\boldsymbol{B}))] + (1-\delta)\,\mathbb{E}[v_i(x_i(\boldsymbol{B}))] \\
&\geq \delta\,\mathbb{E}[v_i(x_i(\boldsymbol{B}))] + (1-\delta)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \delta\,\mathbb{E}[g_i(\boldsymbol{B})] + (1-\delta+\delta\sigma_i)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&\geq \delta\,\mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] + (1-\delta+\delta\sigma_i)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \delta\,\mathbb{E}[v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})) - \sigma_i p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] + (1-\delta+\delta\sigma_i)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&\geq \delta\,\mathbb{E}[\min\left(v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})), \mathfrak{B}_i\right) - \sigma_i p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] + (1-\delta+\delta\sigma_i)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \delta\,\mathbb{E}\left[v_i^{\mathfrak{B}_i}(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})) - \sigma_i p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] + (1-\delta+\delta\sigma_i)\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \delta\,\mathbb{E}\left[v_i^{\mathfrak{B}_i}(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})) - \hat{\sigma}_i(\boldsymbol{B})p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] + (1-\delta+\delta\hat{\sigma}_i(\boldsymbol{B}))\,\mathbb{E}\left[p_i(\boldsymbol{B})\right] \\
&= \delta\,\mathbb{E}\left[\hat{g}_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] + (1-\delta+\delta\hat{\sigma}_i(\boldsymbol{B}))\mathbb{E}\left[p_i(\boldsymbol{B})\right].
\end{aligned}$$

Here, the first equality follows by the definition of **Case 2**, and the first inequality holds since, by assumption, $\boldsymbol{B}$ is a CCE of I and therefore satisfies the ROI constraint in (2.2). Similarly, the second inequality follows from (2.4), since $\boldsymbol{B}$ is a CCE of I, and it holds by assumption that $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$. The fourth equality holds by the definition of budget-capped valuations, while the fifth follows since, by Definition 3.7, $\hat{\sigma}_i(\boldsymbol{B}) = \sigma_i$ holds. ∎

COROLLARY 3.10. *Consider a budget-free instance $I \in \mathcal{I}^{\infty}_{\text{XOS}}$. Fix an agent $i \in N$ and consider a bid profile $\boldsymbol{B}'_i \in \Delta_i$ that is ROI-restricted and let $\boldsymbol{B}_{-i} \in \Delta_{-i}$ be arbitrary. Then, $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$.*

Proof. We have:

$$\begin{aligned}
\mathbb{E}[p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] &= \int_{D_{-i}} f_{\boldsymbol{B}_{-i}}(\boldsymbol{b}_{-i}) \cdot \mathbb{E}[p_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i})] d\boldsymbol{b}_{-i} \\
&\leq \int_{D_{-i}} f_{\boldsymbol{B}_{-i}}(\boldsymbol{b}_{-i}) \cdot \mathbb{E}[v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}))] d\boldsymbol{b}_{-i} = \mathbb{E}[v_i(\boldsymbol{x}_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}))],
\end{aligned}$$

where the inequality follows because $\boldsymbol{B}'_i$ is ROI-restricted. We have thus shown that (2.2) is satisfied. Since I is budget-free (and therefore (2.3) is trivially true), the proof follows. ∎

A.4 Proof of Theorem 3.12.

THEOREM 3.12. *Let $\mathcal{I}_{\text{XOS}}^T$ be the class of instances with fractionally subadditive valuations and type set T. Assume that FPA(r) is (λ_t, μ_t)-smooth for each type $t \in T^+$. Then, the price of anarchy of well-supported coarse correlated equilibria is upper bounded by POA-RMP$(T^+)^{-1}$.*

A crucial building block in deriving our mathematical program (POA-RMP) is the characterization of optimal calibration vectors.

LEMMA A.2. *Let T be a set of types, $\boldsymbol{\mu} = (\mu_t)_{t\in T} \in \mathbb{R}_{>0}^{|T|}$, and $\boldsymbol{\lambda} = (\lambda_t)_{t\in T} \in \mathbb{R}_{>0}^{|T|}$. Then,*

$$\max_{\boldsymbol{\delta}\in\mathcal{C}(\boldsymbol{\mu},T)} \min_{t\in T} \lambda_t\delta_t = O, \qquad \text{where} \qquad O = \min\left\{\min_{t\in T}\lambda_t, \left(\max_{t\in T}\left(\frac{\mu_t}{\lambda_t}\right) + \max_{t\in T}\left(\frac{1-t}{\lambda_t}\right)\right)^{-1}\right\}. \tag{A.4}$$

Proof. Define $\boldsymbol{\delta'}$ such that $\delta'_t = O/\lambda_t$ for each $t \in T$. We first show that $\boldsymbol{\delta'} \in \mathcal{C}(\boldsymbol{\mu}, T)$ and $\min_{t\in T} \lambda_t\delta'_t = O$. First, note that $\boldsymbol{\delta'} \in (0, 1]^{|T|}$, since for each $t \in T$ it holds that:

$$\delta'_t = \frac{O}{\lambda_t} \le \frac{\min_{t\in T}\lambda_t}{\lambda_t} \le 1,$$

and $\delta'_t > 0$ as $O > 0$. Furthermore, $\boldsymbol{\delta'}$ satisfies:

$$\max_{t\in T}(\delta'_t\mu_t) + \max_{t\in T}(\delta'_t(1-t)) = \max_{t\in T}\left(\frac{O}{\lambda_t}\mu_t\right) + \max_{t\in T}\left(\frac{O}{\lambda_t}(1-t)\right) = O\cdot\left(\max_{t\in T}\left(\frac{\mu_t}{\lambda_t}\right) + \max_{t\in T}\left(\frac{1-t}{\lambda_t}\right)\right) \le 1,$$

where the last inequality follows from (A.4). Hence, $\boldsymbol{\delta'} \in \mathcal{C}(\boldsymbol{\mu}, T)$, and therefore $\min_{t\in T}\lambda_t\delta'_t = O$.

To complete the proof, we need to show that $\max_{\boldsymbol{\delta}\in\mathcal{C}(\boldsymbol{\mu},T)}\min_{t\in T}\lambda_t\delta_t = O$. Towards a contradiction, assume that there exists $\bar{\boldsymbol{\delta}} \in \mathcal{C}(\boldsymbol{\mu}, T)$ with $\min_{t\in T}\lambda_t\bar{\delta}_t > O$. We distinguish two cases for the value of O.

Case 1: $\min_{t\in T}\lambda_t \le \left(\max_{t\in T}\left(\frac{\mu_t}{\lambda_t}\right) + \max_{t\in T}\left(\frac{1-t}{\lambda_t}\right)\right)^{-1}$. In this case, we conclude that

$$\min_{t\in T}\lambda_t = O < \min_{t\in T}\lambda_t\bar{\delta}_t \le \min_{t\in T}\lambda_t. \tag{A.5}$$

The equality holds by the definition of **Case 1**, and the first inequality holds by assumption. The second inequality follows since $\bar{\boldsymbol{\delta}} \in \mathcal{C}(\boldsymbol{\mu}, T) \subseteq (0, 1]^{|T|}$, and therefore $\bar{\delta}_t \le 1$ holds for all $t \in T$. Thus, our analysis in (A.5) implies that **Case 1** cannot occur, and we move on to **Case 2**.

Case 2: $\min_{t\in T}\lambda_t > \left(\max_{t\in T}\left(\frac{\mu_t}{\lambda_t}\right) + \max_{t\in T}\left(\frac{1-t}{\lambda_t}\right)\right)^{-1}$. Let $\hat{t} := \arg\max_{t\in T}\mu_t/\lambda_t$ and $\tilde{t} := \arg\max_{t\in T}(1-t)/\lambda_t$. In this case,

$$\begin{aligned} 1 < \left(\frac{\mu_{\hat{t}}}{\lambda_{\hat{t}}} + \frac{1-\tilde{t}}{\lambda_{\tilde{t}}}\right)\cdot\min_{t\in T}\lambda_t\bar{\delta}_t &= \frac{\mu_{\hat{t}}}{\lambda_{\hat{t}}}\cdot\min_{t\in T}\lambda_t\bar{\delta}_t + \frac{1-\tilde{t}}{\lambda_{\tilde{t}}}\cdot\min_{t\in T}\lambda_t\bar{\delta}_t \\ &\le \mu_{\hat{t}}\bar{\delta}_{\hat{t}} + (1-\tilde{t})\bar{\delta}_{\tilde{t}} \le \max_{t\in T}(\mu_t\bar{\delta}_t) + \max_{t\in T}((1-t)\bar{\delta}_t) \le 1. \end{aligned}$$

Here, the first inequality holds by the definition of **Case 2**, and the last inequality holds since $\bar{\boldsymbol{\delta}} \in \mathcal{C}(\boldsymbol{\mu}, T)$. However, our analysis implies that **Case 2** also cannot occur.

As neither **Case 1** nor **Case 2** can occur, we have arrived at a contradiction. This concludes the proof. ∎

Proof of Theorem 3.12. We can now use Lemma A.2 together with our Smoothness Lemmas (Lemma 3.2 and Lemma 3.3) to derive the POA-revealing mathematical program (POA-RMP) as defined in Section 3.3. To obtain a bound on the POA, we determine a vector $\boldsymbol{\mu} = (\mu_t)_{t\in T}$ that maximizes the expression in (A.4) subject to the constraints (3.11)–(3.13). This concludes the proof. ∎

B Missing Material of Section 4

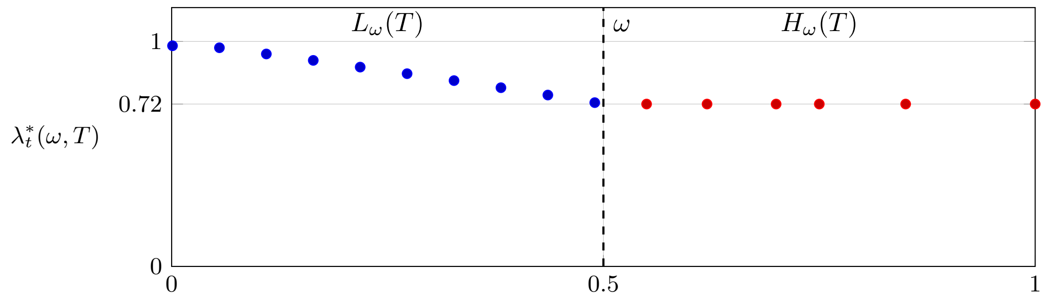

Figure B.1: Illustration of $\boldsymbol{\lambda}^*(\omega, T)$ for $\omega = \frac{1}{2}$ and the partitioning of agent type set T into $L_\omega(T)$ (blue) and $H_\omega(T)$ (red). For all $t \in H_\omega(T)$, the value $\lambda_t^*(\omega, T)$ in Lemma B.1 is given by $\lambda_t^*(\omega, T) = \frac{\omega}{-\ln(1-\omega)} = \frac{1}{2\ln 2} \approx 0.72$. For all $t \in L_\omega(T)$, the value $\lambda_t^*(\omega)$ satisfies $\lambda_t^*(\omega) \geq \frac{\omega}{-\ln(1-\omega)}$.

B.1 Properties of the Feasible Solution.

COROLLARY 4.2. *Given $\omega \in (0, 1)$ and a set of types T, $\boldsymbol{\mu}^*(\omega, T)$ is a feasible solution of POA-RMP(T).*

Proof of Corollary 4.2. Clearly, $\boldsymbol{\mu}^*(\omega, T)$ satisfies (3.11) for all utility maximizers $t \in H_\omega(T) \cap \{1\}$. Also, $\boldsymbol{\mu}^*(\omega, T)$ satisfies (3.12) for each $t \in H_\omega(T) \cap (0, 1)$ because $t \geq \omega$ and because the function $f(z) = -\ln(1-z)$ is non-negative and non-decreasing on $(0, 1)$. Further, it satisfies (3.12) with equality for all types $t \in L_\omega \cap (0, 1)$. Finally, (3.13) holds for all value maximizing types $t \in L_\omega \cap \{0\}$. ∎

Given $\omega \in (0, 1)$ and a set of types T, recall the feasible solution $\boldsymbol{\mu}^*$ in Definition 4.1. (see also Figure B.1).

LEMMA B.1. *The following properties hold for every set of agent types T with $\max(T) > 0$ and every $\omega \in (0, \max(T)] \cap (0, 1)$:*

(i) $\lambda_t^*(\omega, T) = \begin{cases} \frac{\omega}{-\ln(1-\omega)}, & \text{if } t \in H_\omega(T), \\ \frac{t}{-\ln(1-t)}, & \text{if } t \in L_\omega(T) \cap (0, 1), \\ 1, & \text{if } t \in L_\omega(T) \cap \{0\}. \end{cases}$

(ii) $\min_{t \in T \cap (0,1]} \lambda_t^*(\omega, T) = \frac{\omega}{-\ln(1-\omega)}$.

(iii) $\max_{t \in T \cap (0,1]} \frac{1-t}{\lambda_t^*(\omega, T)} = \begin{cases} \frac{-\ln(1-\omega)(1-\min(T\setminus\{0\}))}{\omega}, & \text{if } \min(T \setminus \{0\}) \geq \omega, \\ \frac{-\ln(1-\min(T\setminus\{0\}))(1-\min(T\setminus\{0\}))}{\min(T\setminus\{0\})}, & \text{if } \min(T \setminus \{0\}) < \omega. \end{cases}$

(iv) $\max_{t \in T \cap (0,1]} \frac{\mu_t^*(\omega, T)}{\lambda_t^*(\omega, T)} = \frac{\max(T)}{\omega}$.

Proof of Lemma B.1. For notational convenience, we use $\lambda_t := \lambda_t^*(\omega, T)$ and $\mu_t := \mu_t^*(\omega, T)$ for each $t \in T$. Also, we use $L := L_\omega(T)$ and $H := H_\omega(T)$. Note that H is guaranteed to be non-empty by the range of ω. All statements follow from the definition of $\boldsymbol{\mu}^*(\omega, T)$ in (4.1) and elementary calculus.

(i) For every $t \in L \cap \{0\}$, $\lambda_t = \mu_t = 1$ holds. Then, for all $t \in L \cap (0, 1)$, we have that:

$$\lambda_t = \frac{\mu_t}{t}\left(1 - e^{-\frac{t}{\mu_t}}\right) = \frac{1 - e^{\ln(1-t)}}{-\ln(1-t)} = \frac{t}{-\ln(1-t)}.$$

Similarly, for all $t \in H$ we have that:

$$\lambda_t = \frac{\mu_t}{t}\left(1 - e^{-\frac{t}{\mu_t}}\right) = \frac{1 - e^{\ln(1-\omega)}}{-\ln(1-\omega)} = \frac{\omega}{-\ln(1-\omega)}.$$

(ii) First, assume that $L \cap (0, 1) = \emptyset$. Then, by identity *(i)* we have:

$$\min_{t \in T \cap (0,1]} \lambda_t = \min_{t \in H} \lambda_t = \frac{\omega}{-\ln(1-\omega)}.$$

Next, assume that $L \cap (0,1) \neq \emptyset$. By the definition of L, the fact that $f(x) = \frac{x}{-\ln(1-x)}$ is non-increasing, and statement *(i)*, we obtain:

$$\min_{t \in L \cap (0,1)} \lambda_t = \min_{t \in L \cap (0,1)} \frac{t}{-\ln(1-t)} \geq \frac{\omega}{-\ln(1-\omega)} = \min_{t \in H} \lambda_t.$$

Therefore:

$$\min_{t \in T \cap (0,1]} \lambda_t = \min \left\{ \min_{t \in H} \lambda_t, \min_{t \in L \cap (0,1)} \lambda_t \right\} = \min_{t \in H} \lambda_t = \frac{\omega}{-\ln(1-\omega)}.$$

(iii) First, consider the case of $\min(T \setminus \{0\}) \geq \omega$. Then $L \cap (0,1) = \emptyset$. By statement *(i)* and since $\min H = \min(T \setminus \{0\})$, it holds that:

$$\max_{t \in T: \cap (0,1]} \frac{1-t}{\lambda_t} = \max_{t \in H} \frac{1-t}{\lambda_t} = \frac{-\ln(1-\omega)(1-\min(T \setminus \{0\}))}{\omega}.$$

Next, assume that $\min(T \setminus \{0\}) < \omega$. Then $L \cap (0,1) \neq \emptyset$. We have that:

$$\begin{aligned} \max_{t \in H} \frac{1-t}{\lambda_t} &= \frac{-\ln(1-\omega) \max_{t \in H}(1-t)}{\omega} = \frac{-\ln(1-\omega)\,(1-\min(H))}{\omega} \\ &\leq \frac{-\ln(1-\omega)\,(1-\omega)}{\omega} \leq \max_{t \in L \cap (0,1)} \frac{-\ln(1-t)\,(1-t)}{t} = \max_{t \in L \cap (0,1)} \frac{1-t}{\lambda_t}. \end{aligned} \tag{B.1}$$

The first inequality holds by the definition of H. The second inequality holds by the definition of L and due to the fact that the function $h(x) = \frac{-\ln(1-x)(1-x)}{x}$ is non-increasing in $(0,1)$. Using (B.1), we conclude that:

$$\begin{aligned} \max_{t \in T: \cap (0,1]} \frac{1-t}{\lambda_t} &= \max \left\{ \max_{t \in H} \frac{1-t}{\lambda_t}, \max_{t \in L \cap (0,1)} \frac{1-t}{\lambda_t} \right\} = \max_{t \in L: \cap (0,1)} \frac{1-t}{\lambda_t} \\ &= \max_{t \in L: \cap (0,1)} \frac{-\ln(1-t)(1-t)}{t} = \frac{-\ln(1-t_{\min_+})(1-t_{\min_+})}{t_{\min_+}}. \end{aligned}$$

The last equality holds because the function $h(x)$, as defined above, is non-increasing in $(0,1)$, and thus the maximum is attained for $\min(T \setminus \{0\})$.

(iv) If $L = \emptyset$, the statement follows by statement *(i)*. Otherwise, we have that:

$$\max_{t \in T \cap (0,1]} \frac{\mu_t}{\lambda_t} = \max \left\{ \max_{t \in H} \frac{\mu_t}{\lambda_t}, \max_{t \in L \cap (0,1)} \frac{\mu_t}{\lambda_t} \right\} = \max \left\{ \max_{t \in H} \frac{t}{\omega}, 1 \right\} = \frac{\max(T)}{\omega},$$

where, the last equality follows from the definition of H and because $\omega \leq \max(T)$ holds by assumption. ∎

B.2 Proof of Lemma 4.3.

Lemma 4.3. *Let T be a set of agent types. If $\max(T) > 0$, then for every $\omega \in (0, \max(T)] \cap (0,1)$,*

$$\textit{POA-RMP}(T) \geq \min \left\{ \frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega + \max(T)} \right\}.$$

Proof of Lemma 4.3. Let $\mu_t^*(\omega, T)$ and $\lambda_t^*(\omega, T)$ be as previously defined in (4.1) and Lemma B.1, respectively. For notational convenience, we use $\lambda_t := \lambda_t^*(\omega, T)$ and $\mu_t := \mu_t^*(\omega, T)$ for each $t \in T$. As argued in Corollary 4.2, $(\mu_t)_{t \in T}$ is a feasible solution to POA-RMP(T). We bound the two terms of the min expression in the objective function of POA-RMP(T) separately.

First, we show that:

$$\min_{t \in T} \lambda_t = \frac{\omega}{-\ln(1-\omega)}. \tag{B.2}$$

When $\min(T) > 0$, (B.2) immediately follows from property *(ii)* of Lemma B.1. Otherwise, when $\min(T) = 0$, we have that

$$\min_{t\in T}\lambda_t = \min\left\{\min_{t\in T:t>0}\lambda_t, \min_{t\in T:t=0}\lambda_t\right\} = \min\left\{\frac{\omega}{-\ln(1-\omega)}, 1\right\} = \frac{\omega}{-\ln(1-\omega)}.$$

The last equality holds since $\frac{z}{-\ln(1-z)} < 1$, for all $z \in (0,1)$. We obtain (B.2).

Next, we show that:

$$\max_{t\in T}\frac{\mu_t}{\lambda_t} + \max_{t\in T}\frac{1-t}{\lambda_t} \leq \frac{\max(T)}{\omega} + 1. \tag{B.3}$$

When $\min(T) > 0$, we have:

$$\begin{aligned}\max_{t\in T}\frac{\mu_t}{\lambda_t} + \max_{t\in T}\frac{1-t}{\lambda_t} &= \max_{t\in T\cap(0,1]}\frac{\mu_t}{\lambda_t} + \max_{t\in T\cap(0,1]}\frac{1-t}{\lambda_t} = \frac{\max(T)}{\omega} + \max_{t\in T\cap(0,1]}\frac{1-t}{\lambda_t}\\ &= \frac{\max(T)}{\omega} + \max\left\{\frac{-\ln(1-\omega)\,(1-\omega)}{\omega}, \max_{t\in L\cap(0,1)}\frac{-\ln(1-t)\,(1-t)}{t}\right\} \leq \frac{\max(T)}{\omega} + 1.\end{aligned}$$

Here, the second and third equality follow from properties *(iv)* and *(i)* of Lemma B.1, respectively. Then, the inequality holds by the fact that $\frac{-\ln(1-z)(1-z)}{z} < 1$ holds for all $z \in (0,1)$.

Similarly, when $\min(T) = 0$, we obtain:

$$\begin{aligned}\max_{t\in T}\frac{\mu_t}{\lambda_t} + \max_{t\in T}\frac{1-t}{\lambda_t} &= \max\left\{\max_{t\in T:\cap(0,1]}\frac{\mu_t}{\lambda_t}, \frac{\mu_0}{\lambda_0}\right\} + \max\left\{\max_{t\in T\cap(0,1]}\frac{1-t}{\lambda_t}, \frac{1}{\lambda_0}\right\}\\ &= \max\left\{\frac{\max(T)}{\omega}, 1\right\} + \max\left\{\max_{t\in T:\cap(0,1]}\frac{1-t}{\lambda_t}, 1\right\} = \frac{\max(T)}{\omega} + 1,\end{aligned}$$

where the second equality holds by properties *(i)* and *(iv)* of Lemma B.1, i.e., for $t = 0$, it holds that $\lambda_0 = 1$. The last equality holds because $\omega \leq \max(T)$ (by definition) and by property *(iii)* of Lemma B.1, as $\frac{-\ln(1-z)(1-z)}{z} < 1$ for all $z \in (0,1)$. We thus obtain (B.3).

By combining (B.2) and (B.3), we obtain that the optimal value of POA-RMP(T) is at least:

$$\min\left\{\min_{t\in T}\lambda_t, \left(\max_{t\in T}\frac{\mu_t}{\lambda_t} + \max_{t\in T}\frac{1-t}{\lambda_t}\right)^{-1}\right\} \geq \min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega + \max(T)}\right\}. \qquad \square$$

Acknowledgements S. Klumper was supported by the Dutch Research Council (NWO) through its Open Technology Program, proj. no. 18938. A. Tsikiridis, who was at CWI during the time of this work, was also partially supported by NWO through the Gravitation Project NETWORKS, grant no. 024.002.003 and by the European Union under the EU Horizon 2020 Research and Innovation Program, Marie Skłodowska-Curie Grant Agreement, grant no. 101034253.

References

[1] G. Aggarwal, A. Badanidiyuru, S. R. Balseiro, K. Bhawalkar, Y. Deng, Z. Feng, G. Goel, C. Liaw, H. Lu, M. Mahdian, J. Mao, A. Mehta, V. Mirrokni, R. P. Leme, A. Perlroth, G. Piliouras, J. Schneider, A. Schvartzman, B. Sivan, K. Spendlove, Y. Teng, D. Wang, H. Zhang, M. Zhao, W. Zhu, and S. Zuo, *Auto-bidding and auctions in online advertising: A survey*, SIGecom Exchanges, 22 (2024), p. 159–183.

[2] G. Aggarwal, A. Badanidiyuru, and A. Mehta, *Autobidding with constraints*, in International Conference on Web and Internet Economics, Springer, 2019, pp. 17–30.

[3] G. Aggarwal, G. Fikioris, and M. Zhao, *No-regret algorithms in non-truthful auctions with budget and roi constraints*, in Proceedings of the ACM on Web Conference 2025, WWW '25, New York, NY, USA, 2025, Association for Computing Machinery, p. 1398–1415.

[4] G. Aggarwal, A. Gupta, X. Tan, and M. Zhao, *Multi-platform autobidding with and without predictions*, in Proceedings of the ACM on Web Conference 2025, WWW '25, New York, NY, USA, 2025, Association for Computing Machinery, p. 2850–2859.

[5] Y. Azar, M. Feldman, N. Gravin, and A. Roytman, *Liquid price of anarchy*, 2015, https://arxiv.org/abs/1511.01132, https://arxiv.org/abs/1511.01132.

[6] Y. Azar, M. Feldman, N. Gravin, and A. Roytman, *Liquid price of anarchy*, in Algorithmic Game Theory: 10th International Symposium, SAGT 2017, L'Aquila, Italy, September 12–14, 2017, Proceedings, Berlin, Heidelberg, 2017, Springer-Verlag, p. 3–15.

[7] S. Balseiro, Y. Deng, J. Mao, V. Mirrokni, and S. Zuo, *Robust auction design in the auto-bidding world*, in Proceedings of the 35th International Conference on Neural Information Processing Systems, NIPS '21, Red Hook, NY, USA, 2021, Curran Associates Inc.

[8] S. R. Balseiro, O. Besbes, and G. Y. Weintraub, *Repeated auctions with budgets in ad exchanges: Approximations and design*, Management Science, 61 (2015), pp. 864–884.

[9] S. R. Balseiro, Y. Deng, J. Mao, V. S. Mirrokni, and S. Zuo, *The landscape of auto-bidding auctions: Value versus utility maximization*, in Proceedings of the 22nd ACM Conference on Economics and Computation, EC '21, New York, NY, USA, 2021, Association for Computing Machinery, p. 132–133.

[10] S. R. Balseiro and Y. Gur, *Learning in repeated auctions with budgets: Regret minimization and equilibrium*, Management Science, 65 (2019), pp. 3952–3968.

[11] A. Blum and Y. Monsour, *Learning, regret minimization, and equilibria*, (2007).

[12] M. Braverman, J. Mao, J. Schneider, and M. Weinberg, *Selling to a no-regret buyer*, in Proceedings of the 2018 ACM Conference on Economics and Computation, EC '18, New York, NY, USA, 2018, Association for Computing Machinery, p. 523–538.

[13] I. Caragiannis and A. A. Voudouris, *Welfare guarantees for proportional allocations*, Theor. Comp. Sys., 59 (2016), p. 581–599.

[14] M. Castiglioni, A. Celli, and C. Kroer, *Online learning with knapsacks: the best of both worlds*, in Proceedings of the 39th International Conference on Machine Learning, K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, eds., vol. 162 of Proceedings of Machine Learning Research, PMLR, 17–23 Jul 2022, pp. 2767–2783.

[15] Q. Chen, N. Gravin, and S. Im, *Strategic facility location via predictions*, arXiv preprint arXiv:2410.07497, (2024).

[16] G. Christodoulou, A. Kovács, A. Sgouritsa, and B. Tang, *Tight bounds for the price of anarchy of simultaneous first-price auctions*, ACM Trans. Econ. Comput., 4 (2016).

[17] R. Colini-Baldeschi, S. Klumper, T. Kroll, S. Leonardi, G. Schäfer, and A. Tsikiridis, *Optimal type-dependent liquid welfare guarantees for autobidding agents with budgets*, arXiv preprint arXiv:2506.20908, (2025).

[18] Y. Deng, N. Golrezaei, P. Jaillet, J. C. N. Liang, and V. Mirrokni, *Multi-channel autobidding with budget and roi constraints*, in International Conference on Machine Learning, PMLR, 2023, pp. 7617–7644.

[19] Y. Deng, N. Golrezaei, P. Jaillet, J. C. N. Liang, and V. Mirrokni, *Individual welfare guarantees in the autobidding world with machine-learned advice*, in Proceedings of the ACM Web Conference 2024, WWW '24, New York, NY, USA, 2024, Association for Computing Machinery, p. 267–275.

[20] Y. Deng, M. Mahdian, J. Mao, V. Mirrokni, H. Zhang, and S. Zuo, *Efficiency of the generalized second-price auction for value maximizers*, in Proceedings of the ACM Web Conference 2024, WWW '24, New York, NY, USA, 2024, Association for Computing Machinery, p. 46–56.

[21] Y. Deng, J. Mao, V. Mirrokni, H. Zhang, and S. Zuo, *Autobidding auctions in the presence of user costs*, in Proceedings of the ACM Web Conference 2023, WWW '23, New York, NY, USA, 2023, Association for Computing Machinery, p. 3428–3435.

[22] Y. Deng, J. Mao, V. Mirrokni, H. Zhang, and S. Zuo, *Efficiency of the first-price auction in the autobidding world*, in Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24, Red Hook, NY, USA, 2025, Curran Associates Inc.

[23] Y. Deng, J. Mao, V. Mirrokni, and S. Zuo, *Towards efficient auctions in an auto-bidding world*, in Proceedings of the Web Conference 2021, WWW '21, New York, NY, USA, 2021, Association for Computing Machinery, p. 3965–3973.

[24] S. Dobzinski and R. P. Leme, *Efficiency guarantees in auctions with budgets*, in International Colloquium on Automata, Languages, and Programming, Springer, 2014, pp. 392–404.

[25] M. Feldman, H. Fu, N. Gravin, and B. Lucier, *Simultaneous auctions without complements are (almost) efficient*, Games and Economic Behavior, 123 (2020), pp. 327–341.

[26] M. Feldman, B. Lucier, and N. Nisan, *Correlated and coarse equilibria of single-item auctions*, in Web and Internet Economics: 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings 12, Springer, 2016, pp. 131–144.

[27] Z. Feng, S. Padmanabhan, and D. Wang, *Online bidding algorithms for return-on-spend constrained advertisers*, in Proceedings of the ACM Web Conference 2023, WWW '23, New York, NY, USA, 2023, Association for Computing Machinery, p. 3550–3560.

[28] V. Gkatzelis, K. Kollias, A. Sgouritsa, and X. Tan, *Improved price of anarchy via predictions*, in Proceedings of the 23rd ACM Conference on Economics and Computation, EC '22, New York, NY, USA, 2022, Association for Computing Machinery, p. 529–557.

[29] N. Golrezaei, I. Lobel, and R. Paes Leme, *Auction design for roi-constrained buyers*, in Proceedings of the Web Conference 2021, WWW '21, New York, NY, USA, 2021, Association for Computing Machinery, p. 3941–3952.

[30] J. Hannan, *Approximation to bayes risk in repeated play*, Contributions to the Theory of Games, 3 (1957), pp. 97–139.

[31] Y. Jin and P. Lu, *First price auction is* $1 - 1/e^2$ *efficient*, J. ACM, 70 (2023).

[32] Y. Kolumbus and N. Nisan, *Auctions between regret-minimizing agents*, in Proceedings of the ACM Web Conference 2022, WWW '22, New York, NY, USA, 2022, Association for Computing Machinery, p. 100–111.

[33] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in Annual symposium on theoretical aspects of computer science, Springer, 1999, pp. 404–413.

[34] B. Lehmann, D. Lehmann, and N. Nisan, *Combinatorial auctions with decreasing marginal utilities*, Games and Economic Behavior, 55 (2006), pp. 270–296.

[35] C. Liaw, A. Mehta, and A. Perlroth, *Efficiency of non-truthful auctions in auto-bidding: The power of randomization*, in Proceedings of the ACM Web Conference 2023, WWW '23, New York, NY, USA, 2023, Association for Computing Machinery, p. 3561–3571.

[36] C. Liaw, A. Mehta, and W. Zhu, *Efficiency of non-truthful auctions in auto-bidding with budget constraints*, in Proceedings of the ACM Web Conference 2024, WWW '24, New York, NY, USA, 2024, Association for Computing Machinery, p. 223–234.

[37] B. Lucier, S. Pattathil, A. Slivkins, and M. Zhang, *Autobidders with budget and roi constraints: Efficiency, regret, and pacing dynamics*, in Proceedings of Thirty Seventh Conference on Learning Theory, S. Agrawal and A. Roth, eds., vol. 247 of Proceedings of Machine Learning Research, PMLR, 30 Jun–03 Jul 2024, pp. 3642–3643.

[38] H. Lv, X. Bei, Z. Zheng, and F. Wu, *Auction design for bidders with ex post roi constraints*, in Web and Internet Economics: 19th International Conference, WINE 2023, Shanghai, China, December 4–8, 2023, Proceedings, Berlin, Heidelberg, 2024, Springer-Verlag, p. 492–508.

[39] A. Mehta, *Auction design in an auto-bidding setting: Randomization improves efficiency beyond vcg*, in Proceedings of the ACM Web Conference 2022, WWW '22, New York, NY, USA, 2022, Association for Computing Machinery, p. 173–181.

[40] R. Paes Leme, B. Sivan, and Y. Teng, *Why do competitive markets converge to first-price auctions?*, in Proceedings of The Web Conference 2020, WWW '20, New York, NY, USA, 2020, Association for Computing Machinery, p. 596–605.

[41] T. Roughgarden, *Intrinsic robustness of the price of anarchy*, J. ACM, 62 (2015).

[42] T. Roughgarden, V. Syrgkanis, and E. Tardos, *The price of anarchy in auctions*, J. Artif. Int. Res., 59 (2017), p. 59–101.

[43] F. Susan, N. Golrezaei, and O. Schrijvers, *Multi-platform budget management in ad markets with non-ic auctions*, arXiv preprint arXiv:2306.07352, (2023).

[44] V. Syrgkanis, *Efficiency of mechanisms in complex markets*, Cornell University, 2014.

[45] V. Syrgkanis and E. Tardos, *Composable and efficient mechanisms*, in Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC '13, New York, NY, USA, 2013, Association for Computing Machinery, p. 211–220.

[46] W. Vickrey, *Counterspeculation, auctions, and competitive sealed tenders*, The Journal of finance, 16 (1961), pp. 8–37.

[47] H. P. Young, *Strategic Learning and its Limits*, Oxford University Press, 11 2004.