Robust Analysis of Auction Equilibria*

Jason Hartline[1], Darrell Hoy , and Sam Taggart[2]

[1]Northwestern University
[2]Oberlin College

May 17, 2022

Abstract

Equilibria in auctions can be very difficult to analyze, beyond the symmetric environments where revenue equivalence renders the analysis straightforward. This paper takes a robust approach to evaluating the equilibria of auctions. Rather than identify the equilibria of an auction under specific environments, it considers min-max analysis where an auction is evaluated according to the worst environment and worst equilibrium in that environment. It identifies a non-equilibrium property of auctions that governs whether or not their worst-case equilibria are good for welfare and revenue. This property is easy to analyze, can be refined from data, and composes across markets where multiple auctions are run simultaneously.

1 Introduction

Equilibria in auctions can be very difficult to analyze, beyond the symmetric environments where revenue equivalence renders the analysis straightforward. This paper takes a robust approach to evaluating the equilibria of auctions. Rather than identify the equilibria of an auction under specific environmental conditions, it considers worst-case analysis where an auction is evaluated according to the worst environment and worst equilibrium in that environment. It identifies a non-equilibrium property of auctions that governs whether or not their worst-case equilibria are good for welfare and revenue. This property is easy to analyze, can be refined from data, and composes across markets where multiple auctions are run simultaneously.

Classical economic analyses identify three main drivers of inefficiency: (a) externalities, (b) incomplete information, and (c) market power. The analysis of this paper decomposes the performance of auctions into two terms: *competitive efficiency* quantifies the degree to which externalities cause losses in performance and *individual efficiency* quantifies the degree to which incomplete information causes losses in performance. The decomposition does not ascribe further loss of efficiency to market power.

The quantity of interest for this paper is the *robust efficiency* of mechanisms, i.e., the fraction of the optimal welfare or revenue that is attained in any equilibrium and under any informational

*This paper provides an economic interpretation on results presented in an extended abstract under the title "Price of Anarchy for Auction Revenue" at the fifteenth ACM Conference on Economics and Computation (Hartline et al., 2014). We thank Vasilis Syrgkanis for comments on a prior version of this paper for which simultaneous composition did not hold, for suggesting the study of simultaneous composition, and for perspective on price-of-anarchy methodology.

model. By measuring the efficiency of a mechanism as the ratio of its performance to the optimal performance we obtain bounds that are invariant with respect to the relative magnitudes of the environment. By normalizing the performance with respect to the optimal performance, the robust analysis framework requires that performance is good when good performance is possible. In contract, absolute evaluation of robust performance gives result that are primarily determined by environments where the optimal performance is bad.

The analysis of the paper shows that competitive efficiency is a central determinant of whether an auction is good or bad; while individual efficiency is always relatively high. Competitive efficiency can be low when externalities are significant and mechanisms that are more competitively efficient reduce the impact of externalities. On the other hand, individual efficiency which quantifies the impact of incomplete information is always relatively high, though mechanisms with the winner-pays-bid payment format can be seen as better than those with the all-pay payment format. Applying the perspective of Hartline (2013) to these two concerns: externalities are a critical feature to be treated carefully in mechanism design while incomplete information is more of a detail.

Introducing the notion of competitive efficiency by example, consider any set of bids in a first-price auction. For each agent and this fixed set of bids, there is a minimum bid that serves as a threshold for whether the agent wins or loses. Consider two quantities: (a) the revenue of the auction for the given bids and (b) the optimal revenue of an auction if each agent instead bid their threshold. Quantity (b) is a measure of the level of competition in the auction. The *competitive efficiency* μ is the worst case over bids of the ratio of (a) to (b), this number is at most 1 and large numbers are more efficient. For the first-price auction this ratio is $\mu = 1$. The revenue is the highest bid and the thresholds of all losers are equal to the highest bid. Thus, the revenue from thresholds and the revenue from bids are the same.

Even mechanisms that are competitively efficient may be less efficient due to the inability of the agents to precisely respond to thresholds, i.e., due to incomplete information. Returning to the example of first price auctions, with deterministic equilibrium concepts like pure Nash equilibrium, an agent's threshold is deterministic, and the agent can respond efficiently. One the other hand, in stochastic equilibrium concepts like Bayes-Nash equilibrium, the threshold an agent faces is stochastic, and the agent cannot respond efficiently. Specifically, when the threshold is below the agent's value the agent would prefer to bid just above the threshold; however, this threshold is stochastic and the agent must place a single bid. Viewing this threshold as the seller's outside option, the inability for an agent to respond precisely results in an additional loss of performance. This *individual efficiency* η is a property of the best-response problem faced by the agents and depends on the equilibrium notion (e.g., pure Nash or Bayes-Nash) and payment format (e.g., winner-pays-bid or all-pay).

This paper shows that robust efficiency can be bounded by a combination of the individual efficiencies of the agents and competitive efficiency of the mechanism. The main welfare analysis of this paper proves that, broadly, the fraction of the optimal welfare of an auction in equilibrium is at least the product of the mechanism's competitive efficiency and the agents' individual efficiencies. For example, we will show that the individual efficiency of the agent's response for pure Nash equilibria in winner-pays-bid mechanisms is $\eta = 1$. Combined with the competitive efficiency of $\mu = 1$ for the first-price auction, we see that pure Nash equilibria in the first-price auction are fully efficient, i.e., they obtain a $\eta\mu = 1$ fraction of the optimal welfare. (More precisely, pure Nash tend not to exist in the first-price auction, but the same result approximately holds for approximate pure Nash which exist; details given subsequently.) The main revenue analysis of the paper shows that the fraction of the optimal revenue of an auction with appropriate reserve prices in equilibrium is at least half the product of its competitive efficiency and the agents' individual efficiency.

Individual efficiency is a property of the agents' best response problem that takes into account the incompleteness of information and the payment format. Thus, it is sufficient to analyze a few canonical models of incomplete information and payment formats. We provide the following individual efficiency results:

- $\eta = (1 - \epsilon)$ for winner-pays-bid mechanisms in pure $(1 - \epsilon)$-Nash equilibria (for any $\epsilon \geq 0$). This result is written for approximate Nash equilibrium rather than exact Nash equilibrium because in winner-pays-bid auctions the latter tends not to exist. This is a complete information model.

- $\eta = 1 - 1/e \approx 0.63$ for winner-pays-bid mechanisms in Bayes-Nash equilibria. This is an incomplete information model.

- $\eta = 1/2$ for all-pay mechanisms in Bayes-Nash equilibria. This is an incomplete information model.

These results quantify the role of incomplete information. When information is compete (in payoffs, actions of opponents, and rules of the mechanism), an agent's best response problem is fully efficient. On the other hand we see that with incomplete information the agent's best response problem can be inefficient, but that inefficiency is bounded by a constant. As mentioned above, the best response problem of winner-pays-bid payment formats mitigates informational inefficiencies more than that of the all-pay payment formats.

Combining these individual efficiency bounds with the competitive efficiency of the first-price auction ($\mu = 1$), the equilibrium welfare is at least an $\eta\mu = 0.63$ fraction of the optimal welfare in equilibrium. The revenue of the first-price auction with per-agent monopoly reserves a $\eta\mu/2 = 0.31$ fraction of the optimal revenue in equilibrium. Recall that these auctions are optimal for welfare and revenue when the agents values are identically distributed, the above robust welfare guarantee holds for all correlated distributions and the above robust revenue guarantee holds for all non-identical product distributions satisfying a standard regularity property. We see from this analysis that, while the first-price auction can be inefficient, it can never be extremely inefficient.

Competitive efficiency is a property of the rules of the auction that map bids to winners, a.k.a., bid allocation rules. Given the individual efficiencies of standard best response problems of the agents, the robust efficiency of a mechanism is approximately governed by its competitive efficiency. Therefore, analysis of robust welfare and revenue of an auction approximately reduces to analysis of its competitive efficiency.

To aid in the analysis of the competitive efficiency we develop a number of closure properties.

- Competitive efficiency is closed under reserve pricing, i.e., the competitive efficiency of mechanism without reserve prices is equal to the competitive efficiency of the mechanism with the worst reserve prices.

- Competitive efficiency is closed under randomizations of mechanism and selection, i.e., the competitive efficiency of any randomized mechanism on any randomized sets of agents is equal to the competitive efficiency of the worst combination of mechanism and set of agents.

- Competitive efficiency is closed under simultaneous composition (when bids are independently distributed), i.e., the competitive efficiency of the worst composite mechanism where a set of mechanisms are run in parallel is equal to the competitive efficiency of the worst mechanism in the set. (In this composition, agents are assumed to be unit-demand, but can bid in multiple mechanisms at once if it is in their best interest.)

These closure properties imply that it is generally sufficient to analyze competitive efficiency of deterministic mechanisms with deterministic selection. Such analyses are fairly straightforward compared to more technically involved analyses of stochastic equilibrium concepts.

The following mechanisms have competitive efficiency $\mu = 1$. Intuitively, in these environments the externalities are only those of one-for-one substitution:

- Single-item multi-unit unit-demand highest-bids-win mechanisms, i.e., there are k identical units for sale and they are allocated to the k-highest bidders. (Note that the first-price auction is the special case where $k = 1$.)

- Rank-by-bid position auctions, i.e., there are k positions with descending weights, and the highest k bidders are assigned to these k positions in order of bid. The agent in the jth highest position receives a unit with probability equal to the jth position weight. (The k-unit auction is a special case where the position weights are all 1.) The position auction model was popularized by the studies of Varian (2007a) and Edelman et al. (2007a) auctions for advertising on Internet search engines.

- (Single-bid) highest-bids-win matching markets, i.e., there are m items and n bidders who each desire a single item from known subset of the items. The highest-bid-wins rule selects the bidders to match to maximize the sum of the matched bids.

- (Multiple-bid) per-item highest-bids-win matching markets, i.e., there are m items and n bidders who each desire a single item from a unknown subset of items, bids are submitted for each item, and the highest bidder for each wins. (For this result, bids are required to be independently distributed in equilibrium.)

These results all follow from a single analysis of the competitive efficiency of bid allocation rules based on greedy algorithms and the above closure properties. Greedy algorithms order the agents by a function of their bids and then allocate to each one in turn, if doing so is feasible. Notice that the k-unit highest-bids-win allocation rule is given by the greedy-by-bid algorithm. For complicated optimization problems, greedy algorithms are not generally optimal. Our main analysis of competitive efficiency shows that the non-strategic efficiency of the greedy algorithm, i.e., the fraction of the optimal welfare it obtains in non-strategic environments, is equal to its competitive efficiency. The results listed above follow because the greedy algorithm is non-strategically efficient for these environments. The result for rank-by-bid position auctions additionally views the position auction as a convex combination of multi-unit auctions and invokes the closure of competitive efficiency under convex combinations. The result for multiple-bid per-item matching markets follows from the closure under simultaneous composition of the first-price auction.

The competitive efficiency of mechanisms can be very bad when externalities between agents are ones of many-for-one complementarities. The classic environment where agents are complements is the single-minded combinatorial auction. Here there are n agents and m items and agents each desire the entirety of a known subset of items. The competitive efficiency of highest-bids-win single-minded combinatorial auctions is $1/m$ and indeed winner-pays-bid highest-bids-win auctions possess equilibria that are $1/m$ efficient. A classical result from the computer science literature on algorithm design, however, shows that there is a greedy algorithm for the single-minded combinatorial allocation problem that is $1/\sqrt{m}$ efficient;[1] thus the winner-pays-bid auction with this greedy bid allocation rule $1/\sqrt{m}$ competitively efficient. Explicitly designing mechanisms to maximize competitive efficiency can significantly reduce the impact of externalities.

1.1 Related Work

In symmetric environments in the canonical independent private value model second-price, first-price, and all-pay auctions with or without reserves are both welfare and revenue equivalent (Myerson, 1981). In asymmetric models, they are not equivalent, and the equilibria, even under the simplifying assumption that the agents' values are uniformly (but asymmetrically with different supports) distributed, are very difficult to solve for (Kaplan and Zamir, 2012).

An approach from computer science for understanding potentially complex equilibria in mechanisms is to bound is to give robust bounds on equilibria that do not require exactly identifying an equilibrium. This literature analyzes the robust efficiency, a.k.a., the price of anarchy, of games and mechanisms. Within this literature, this paper builds on the smooth games framework of Roughgarden (2009) and the smooth mechanisms extension of Syrgkanis and Tardos (2013). In this context, this paper refines the smoothness framework for Bayesian games in two notable ways. First, it decomposes smoothness into two components, separating the consequences of best-response (individual efficiency) from the specifics of a mechanism (competitive efficiency). The former quantifies losses due to incomplete information and the latter quantifies losses due to externalities between agents. Second, the framework is compatible with the analysis of auction revenue by Myerson (1981) and allows for robust bounds on the equilibrium revenue of auctions.

There are two subsequent works with strong connections to our decomposition of smoothness into competitive efficiency and individual efficiency. First, Dütting and Kesselheim (2015) show that competitive efficiency, which they call "permeability," is in fact a necessary condition (we show sufficiency) for the equilibrium welfare of a mechanism to be proven to be good via the smoothness framework. Second, Hoy et al. (2015) show how to derive empirical welfare bounds by measuring the degree to which competitive efficiency and individual efficiency hold, without needing to infer the agents' true values.

A number of papers have derived revenue guarantees for the welfare-optimal Vickrey-Clarke-Groves (VCG) mechanism in asymmetric environments. Hartline and Roughgarden (2009) show that VCG with monopoly reserves, a carefully chosen anonymous reserve, or duplicate bidders achieves revenue that is a constant approximation to the revenue optimal auction. Dhangwatnotai et al. (2010) show that the single-sample mechanism, essentially VCG using a single sample from the distribution as a reserve, achieves approximately optimal revenue in broader environments. Roughgarden et al. (2012) showed that in broader environments, including matching environments, limiting the supply of items in relation to the number of bidders gives a constant approximation to the optimal auction. See Hartline (2013) for a survey of results in this area.

[1]Computer scientists study greedy algorithms like this one in part because computing the allocation that maximizes the sum of winning bids is believed to be computationally intractable. Thus, we see that greedy algorithms should be preferred as allocation rules for auctions both for their computational tractability and for there competitive efficiency.

Bergemann et al. (2017) consider a robust analysis of the revenue in a first-price auction with respect to the knowledge of the agents. They derive tight lower bounds on this revenue based on a characterization of the minimum distribution over winning bids. In contrast, though our welfare bounds are tight, our revenue bounds are lose. Our bound, however, apply more broadly than single-item auctions. Bergemann et al. (2019) give worst-case, over information structure, revenue rankings of standard auction formats.

Robust analyses in the economics literature have focused on mechanism design context where a principal faces ambiguity with respect to an aspect of the model and aims to maximize the worst-case revenue with respect to the ambiguity. For example, Carroll (2017) considers a revenue maximizing seller with multiple items and a buyer with values distributed with known marginals. The correlation structure is ambiguous. He shows that the max-min mechanism is linear pricing. Brooks and Du (2021) consider the design of revenue optimal common-value auctions that are robust to information structures. A key difference between these works and this paper is that this paper considers robust performance relative to optimal performance.

2 Preliminaries

This paper studies mechanisms for single-parameter agents. A mechanism consists of action spaces A_i for each agent i (and joint action space $A = \prod_i A_i$), an allocation rule $\tilde{\mathbf{x}}$, and a payment rule $\tilde{\mathbf{p}}$. Given the profile $\mathbf{a}$ of actions selected by each agent, the mechanism computes an allocation level $\tilde{x}_i(\mathbf{a}) \in [0, 1]$ and payment $\tilde{p}_i(\mathbf{a}) \in \mathbb{R}$ for each agent i, with $\tilde{\mathbf{x}}(\mathbf{a})$ and $\tilde{\mathbf{p}}(\mathbf{a})$ describing the full profiles of allocations and payments, respectively. Each agent i has a value v_i for service, and linear utility function $\tilde{u}_i(\mathbf{a}) = v_i\tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})$.

We study n-agent mechanisms under both complete and incomplete information. A mechanism's allocation is constrained by a *feasibility environment* $\mathcal{X}$. For example, in selling a single item, $\mathcal{X} = \{\mathbf{x} \in [0, 1]^n \mid \sum_i x_i \leq 1\}$. Settings we consider will be *downward-closed*, in the sense that for any $\mathbf{x} \in \mathcal{X}$ and any index i, $(0, \mathbf{x}_{-i}) \in \mathcal{X}$. Even fixing a mechanism's payment format (e.g. first-price), richer feasibility settings admit many possible allocation rules. In our analysis framework, competitive efficiency (Section 3) captures the consequences of this variation.

Values are drawn according to a common (possibly degenerate or correlated) prior $\mathcal{F}$, with marginals $\{\mathcal{F}_i\}_i$. Throughout, we maintain the interpretation that each agent i is randomly selected to participate from a population with value distribution $\mathcal{F}_i$, with the joint selection governed by $\mathcal{F}$. A strategy function s_i maps values for agent i to actions, with the interpretation that the value v_i agents in population i all play $s_i(v_i)$ when selected to participate. Denote a profile of strategies by $\mathbf{s}$, and for an action a_i, denote by $\tilde{x}_i(a_i) = \mathbb{E}_{\mathbf{v}_{-i} \mid v_i}[\tilde{x}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$ (resp. $\tilde{p}_i(a_i)$, $\tilde{u}_i(a_i)$) the interim allocation (resp. payment, utility) rule induced by $\mathbf{s}$. For mechanism M, prior $\mathcal{F}$, and strategy profile $\mathbf{s}$, our two objectives of interest are revenue, given by $\textsc{Rev}(M, \mathbf{s}, \mathcal{F}) = \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\sum_i \tilde{p}_i(\mathbf{s}(\mathbf{v}))]$, and welfare (alternatively, "surplus"), given by $\textsc{Welfare}(M, \mathbf{s}, \mathcal{F}) = \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\sum_i \tilde{u}_i(\mathbf{s}(\mathbf{v}))] + \textsc{Rev}(M, \mathbf{s}, \mathcal{F}) = \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\sum_i v_i\tilde{x}_i(\mathbf{s}(\mathbf{v}))]$.

The mechanisms we study are non-truthful and depend only minimally on priors. The standard theory of revenue maximization requires detailed knowledge of the prior $\mathcal{F}$, and the families of mechanisms we consider often have equilibria with suboptimal welfare. We therefore pursue robust, or worst-case, approximation analyses, and ask how far from optimal a mechanism can be, quantified over equilibria in a family. Formally, given prior $\mathcal{F}$, denote the optimal expected welfare by $\textsc{Welfare}(\textsc{Opt}, \mathcal{F}) = \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\max_{\mathbf{x}^* \in \mathcal{X}} \sum_i v_i x_i^*]$. Fixing a family of distributions (e.g. degenerate, product, or unrestricted) $\mathbb{F}$ and family of equilibria $\mathrm{Eq}_{\mathcal{F}}(M)$ for each $\mathcal{F} \in \mathbb{F}$ and mechanism M (e.g. ϵ-Nash, Bayes-Nash), we study the *approximation ratio*

$$\min_{\mathcal{F}\in\mathbb{F},\mathbf{s}\in\mathrm{Eq}_{\mathcal{F}}(M)} \frac{\textsc{Welfare}(M,\mathbf{s},\mathcal{F})}{\textsc{Welfare}(\textsc{Opt},\mathcal{F})}.$$

When the value of the approximation ratio is at least ρ, we say the welfare of M is a ρ-*approximation* to the optimal welfare over $\mathbb{F}$ and $\mathrm{Eq}_{\mathcal{F}}(M)$. A large value of ρ (close to 1) indicates that M is always nearly efficient, whereas ρ very small suggests a pathology that might rule out M in practice. By exposing the structural characteristics that influence worst-case performance, our framework can inform robust design. We also show that several commonly-observed formats such as the single-item first-price auction can be explained by their low approximation ratio.

For revenue, we assume independently distributed values across agents and argue with respect to the Bayesian optimal revenue for $\mathcal{F}$. Per Myerson (1981), the optimal mechanism can depend intricately on the prior. Denote the expected revenue of the optimal mechanism $\textsc{Opt}_{\mathcal{F}}$ for prior $\mathcal{F}$ by $\textsc{Rev}(\textsc{Opt}_{\mathcal{F}},\mathcal{F})$. We study mechanisms which are prior-independent except for *monopoly reserves* (defined formally in Section 4), which depend on much less fine-grained information than the form of the Bayesian optimal mechanism. We study the ratio

$$\min_{\mathcal{F}\in\mathbb{F},\mathbf{s}\in\mathrm{BNE}_{\mathcal{F}}(M_{\mathcal{F}})} \frac{\textsc{Rev}(M_{\mathcal{F}},\mathbf{s},\mathcal{F})}{\textsc{Rev}(\textsc{Opt}_{\mathcal{F}},\mathcal{F})},$$

where $M_{\mathcal{F}}$ denotes mechanism M endowed with monopoly reserve prices for $\mathcal{F}$ and $\mathrm{BNE}_{\mathcal{F}}(M_{\mathcal{F}})$ the family of Bayes-Nash equilibria for $M_{\mathcal{F}}$ under $\mathcal{F}$.

3 Competitive Efficiency

In this section, we formally define competitive efficiency. Competitive efficiency depends only on the bid allocation rule of the mechanism. Importantly, it does not depend on preferences, beliefs, or strategies of agents; nor does it depend on relationships between these. As we will see in Section 4, this parameter governs the extent to which equilibria in the mechanism obtain good welfare and revenue, and is robust to specifics that are typically critical for equilibria such as preferences, beliefs, and strategies.

For the time being we restrict our analysis to *single-bid* rules, for which each agent i's action is a single real-valued bid b_i. We extend the analysis to general mechanisms in Section 5. Given single-bid allocation rule $\tilde{\mathbf{x}}$, two salient payment formats are the winner-pays-bid mechanism for $\tilde{\mathbf{x}}$ which has payment rule $\tilde{p}_i(\mathbf{b}) = b_i\tilde{x}_i(\mathbf{b})$, and the all-pay mechanism for $\tilde{\mathbf{x}}$, with payment rule $\tilde{p}_i(\mathbf{b}) = b_i$. Competitive efficiency compares two quantities: *threshold surplus*, which quantifies the level of competition each agent faces, and the total bid of the winners, i.e. *bid surplus*. In winner-pays-bid and all-pay mechanisms, the bid surplus lower bounds total payments. Hence, competitive efficiency measures the extent to which competition translates into revenue. We first consider deterministic allocation rules without reserve prices under full information, where this comparison is particularly straightforward. Competition can be easily quantified by the *threshold bid* that each agent's bid must exceed to win.

Definition 1. *In an (implicit) deterministic single-bid rule $\tilde{\mathbf{x}}$ and a profile of bids $\mathbf{b}$, we summarize agent i's competition by the* threshold bid $\hat{b}_i(\mathbf{b}_{-i}) = \inf\{b_i \,|\, \tilde{x}_i(b_i,\mathbf{b}_{-i}) = 1\}$ *that the agent must outbid to win. Denote the* threshold surplus *for allocation y_i by $T_i(y_i) = \hat{b}_i\, y_i$.*

Definition 2. *The* competitive efficiency *of a deterministic allocation rule* $\tilde{\mathbf{x}}$ *for a deterministic environment is the largest* μ *such that, for any profile of bids* $\mathbf{b}$ *and any feasible allocation* $\mathbf{y}$*, the bid surplus is at least a* μ *fraction of the* threshold surplus*.*

$$\sum\nolimits_i b_i \tilde{x}_i(\mathbf{b}) \geq \mu \sum\nolimits_i T_i(y_i). \tag{1}$$

In the case of deterministic single-bid rules, threshold surplus is a simple linear function of the allocation, i.e., $\sum_i T_i(y_i) = \sum_i \hat{b}_i\, y_i$. As already described in the introduction, the highest-bid-wins rule has competitive efficiency $\mu = 1$. Given bids, the threshold bids of losers are the highest bid; the threshold bid of the winner is the second highest bid. Thus, both the optimal threshold surplus and the bid surplus are equal to the highest bid (see the formal treatment in Section 3.1). Later in this section, we will identify a number of deterministic rules of interest that have easy-to-analyze competitive efficiencies.

It is common in practice to augment a winner-pays-bid mechanism with a profile of individualized reserve prices. One generic way to incorporate a profile of reserves $\mathbf{r}$ into a mechanism with single-bid rule $\tilde{\mathbf{x}}$ to produce a rule $\tilde{\mathbf{x}}^{\mathbf{r}}$ is: (1) Solicits a bid b_i from each agent i. (2) For each agent i, if $b_i < r_i$, sets $\tilde{b}_i = 0$, else $\tilde{b}_i = b_i$. (3) Allocates according to $\tilde{\mathbf{x}}(\tilde{\mathbf{b}})$. We typically consider adding reserves to rules where bidding 0 guarantees an agent goes unallocated, i.e. $\tilde{x}_i(0, \mathbf{b}_{-i}) = 0$.

The definition of competitive efficiency extends naturally to allocation rules with reserves, as well as to randomization in the auction environment. Such randomization could come from feasible allocations being randomized, randomization in the allocation rule itself, or uncertainty over the participants of the auction, captured by the prior distribution. The general definition of competitive efficiency will enable quantification of welfare and revenue over the more complicated mechansisms and equilibria arising from reserves and randomness. Importantly, though, extension theorems given subsequently will allow us to characterize competitive efficiency in these environments via analysis of simpler mechanisms with full information and no reserves.

For randomized allocation rules or environments, threshold bids are not deterministic, and generally, the competition faced by an agent depends on the desired level of allocation. We now give a natural generalization of threshold surplus that quantifies the competition an agent faces in the presence of randomness as a function of the desired level of allocation. The definition discounts reserve prices, which make it harder for an agent to receive allocation, but do not stem from competition. This distinction is critical for analysis of mechanisms with reserves. Intuitively, an agent i faces strong competition under a single-bid allocation rule $\tilde{\mathbf{x}}$ with reserves $\mathbf{r}$, strategy profile $\mathbf{s}$, and prior $\mathcal{F}$ if bids above their reserve price generally yield low allocation probabilities. To this end, recall we define i's interim allocation rule by $\tilde{x}_i(b \,|\, v_i) = \mathbb{E}_{\mathbf{v}_{-i} \,|\, v_i}[\tilde{x}_i(b, \mathbf{s}_{-i}(\mathbf{v}_{-i})]$, and define its inverse as $t_i(x \,|\, v_i) = \inf\{b \,|\, \tilde{x}_i(b \,|\, v_i) \geq x\}$.

Definition 3. *For an (implicit) single-bid allocation rule* $\tilde{\mathbf{x}}$ *with (implicit) reserves* $\mathbf{r}$*, strategy profile* $\mathbf{s}$*, and prior* $\mathcal{F}$*, The competition faced by agent* i *with value* v_i *for obtaining allocation* y_i *is summarized by the* threshold surplus with discounted reserve*, defined as* $T_i(y_i \,|\, v_i) = \int_{\tilde{x}_i(r_i \,|\, v_i)}^{y_i} t_i(x \,|\, v_i)\, dx$*. With the reserve price denoted explicitly, define* $T_i^{r_i}(y_i \,|\, v_i)$*.*

For deterministic environments, Definition 2 defined competitive efficiency in terms of feasible, deterministic allocations. Under a general prior, we now consider feasible interim allocations. Taking the view of the prior as defining populations of agents, an interim allocation level defines a feasible allocation probability for each agent in each population. Formally, let $\mathbf{y}$ map value profiles to allocations in the feasible set $\mathcal{X}$. Denote the interim allocations for $\mathbf{y}$ by $y_i(v_i) = \mathbb{E}_{\mathbf{v}}[y_i(\mathbf{v}) \,|\, v_i]$.

Definition 4. *The* competitive efficiency *of a (randomized) allocation rule* $\tilde{\mathbf{x}}$ *for prior* $\mathcal{F}$ *is the largest* μ *such that, for any strategy profile* $\mathbf{s}$ *and any feasible allocation function* $\mathbf{y}$ *mapping value profiles to* $\mathcal{X}$*, the expected bid surplus is at least a* μ *fraction of the* threshold surplus*:*

$$\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}\left[\sum\nolimits_i s_i(v_i)\tilde{x}_i(\mathbf{s}(\mathbf{v}))\right] \geq \mu \sum\nolimits_i \mathbb{E}_{v_i\sim\mathcal{F}_i}\left[T_i(y_i(v_i)\,|\,v_i)\right]. \tag{2}$$

The simple definition of competitive efficiency (Definition 2) quantified over all deterministic bid profiles. In Definition 4 a strategy profile picks a bid for every agent in every population. Hence, Definition 4 generalizes the simpler definition to bid profiles over populations of participants. Furthermore, though Definition 4 depends on the prior $\mathcal{F}$, we show via closure properties that a rule's competitive efficiency on all priors can be analyzed by considering any single prior, including a degenerate one. It will therefore often suffice to work with the simpler Definition 2 to characterize the competitive efficiency of Definition 4. Conceptually, this shows that competitive efficiency depends only on the allocation rule of the mechanism.

We now derive two closure properties of competitive efficiency. Our first is that competitive efficiency is closed with respect to reserve prices, i.e., if a rule has competitive efficiency μ without reserves, then its competitive efficiency with reserves is μ. Then, we consider the impact of two types of randomization on competitive efficiency. First, we consider randomization over allocation rules. In other words, if the lowest competitive efficiency of any rule in a family of allocation rules is μ then the competitive efficiency of any convex combination of rules in the family is at least μ. Second, we study closure under mixture over priors. We start with closure under reserves.

Lemma 5. *The competitive efficiency is closed under reserve pricing, i.e., an allocation rule with competitive efficiency* μ *on* $\mathcal{F}$ *(without reserves) has competitive efficiency* μ *on* $\mathcal{F}$ *with reserves.*

Proof of Lemma 5. Let $\tilde{\mathbf{x}}^{\mathbf{0}}$ be a rule that has competitive efficiency μ without reserves (equivalently with reserves $\mathbf{0}$) on prior $\mathcal{F}$. We will show that adding any profile of reserves $\mathbf{r}$ yields a rule $\tilde{\mathbf{x}}^{\mathbf{r}}$ with competitive efficiency μ with reserves $\mathbf{r}$ on $\mathcal{F}$. The converse direction in the lemma is immediate.

For a strategy profile $\mathbf{s}$, let $\mathbf{s}^{\mathbf{r}}$ denote the strategy profile obtained by setting to 0 all bids $s_i(v_i) < r_i$ failing to meet the reserves $\mathbf{r}$. The main ideas of the proof are that (a) outcomes are equivalent for $\tilde{\mathbf{x}}^{\mathbf{0}}$ on $\mathbf{s}^{\mathbf{r}}$ and $\tilde{\mathbf{x}}^{\mathbf{r}}$ on $\mathbf{s}$, and (b) fixing the bid allocation rules, reserves only lower threshold surplus. Thus, a competitive efficiency without reserves implies the same competitive efficiency with reserves. We adopt the following notation that makes the allocation rule $\tilde{\mathbf{x}}$ and strategy profile $\mathbf{s}$ explicit in our notation for threshold surplus with discounted reserve $T_i^{r_i}(x, \tilde{\mathbf{x}}, \mathbf{s}_{-i}\,|\,v_i) = T_i^{r_i}(x\,|\,v_i)$. Zero reserves will be explicitly designated as such.

Consider the following analysis, with subsequent discussion:

$$\begin{aligned}
\mu\,\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i s_i(v_i)\tilde{x}_i^{\mathbf{r}}(\mathbf{s}(\mathbf{v}))\right] &= \mu\,\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i s_i^{\mathbf{r}}(v_i)\tilde{x}_i^{\mathbf{0}}(\mathbf{s}^{\mathbf{r}}(\mathbf{v}))\right] && (3)\\
&\geq \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^0(y_i(v_i), \tilde{\mathbf{x}}^{\mathbf{0}}, \mathbf{s}^{\mathbf{r}}_{-i}\,|\,v_i)\right] && (4)\\
&= \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^0(y_i(v_i), \tilde{\mathbf{x}}^{(0,\mathbf{r}_{-i})}, \mathbf{s}_{-i}\,|\,v_i)\right] && (5)\\
&\geq \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^{r_i}(y_i(v_i), \tilde{\mathbf{x}}^{(0,\mathbf{r}_{-i})}, \mathbf{s}_{-i}\,|\,v_i)\right] && (6)\\
&= \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^{r_i}(y_i(v_i), \tilde{\mathbf{x}}^{\mathbf{r}}, \mathbf{s}_{-i}\,|\,v_i)\right] && (7)
\end{aligned}$$

Equations (3) and (5) follow by the equivalence of outcomes from reserves in the allocation rule and reserves in the bids. Equation (4) follows from the assumed competitive efficiency of the rule without reserves $\tilde{\mathbf{x}}^{\mathbf{0}}$. Equation (6) follows from the definition of threshold surplus with discounted reserves; i.e., discounting reserves lowers the threshold surplus. Equation (7) follows because the thresholds considered are only above the reserve. Combining the sequence of inequalities we observe that $\tilde{\mathbf{x}}^{\mathbf{r}}$ has competitive efficiency μ on strategy profile $\mathbf{s}$. □

Closure under convex combination of both allocation rules and priors will follow from a single convexity argument. The definitions of each are below, followed by a proof encompassing both.

Definition 6. *Let $\theta \sim U[0,1]$ be a uniform random variable indexing over allocation rules $\tilde{\mathbf{x}}^\theta$ and their feasibility environments $\mathcal{X}^\theta$. The convex combination of these rules is $\tilde{\mathbf{x}}(\mathbf{b}) = \mathbb{E}_\theta[\tilde{\mathbf{x}}^\theta(\mathbf{b})]$. The corresponding feasibility environment is $\mathcal{X} = \{\mathbf{x} = \mathbb{E}_\theta[\mathbf{x}^\theta] : \forall \theta \in [0,1],\ \mathbf{x}^\theta \in \mathcal{X}^\theta\}$.*

Definition 7. *Let $\omega \sim U[0,1]$ be a uniform random variable indexing over priors $\mathcal{F}^\omega$. To draw from the convex combination of these distributions $\mathcal{F}$, draw $\omega \sim U[0,1]$ then draw $\mathbf{v} \sim \mathcal{F}^\omega$.*

Lemma 8. *The competitive efficiency is closed under convex combination of allocation rules and of priors, i.e. (i) if a rule $\tilde{\mathbf{x}}$ has competitive efficiency μ on $\mathcal{F}^\omega$ for all ω, it also has competitive efficiency μ on the convex combination $\mathcal{F}$; (ii) if for all θ, $\tilde{\mathbf{x}}^\theta$ has competitive efficiency μ on $\mathcal{F}$, then the convex combination $\tilde{\mathbf{x}}$ also has competitive efficiency μ on $\mathcal{F}$.*

Proof. Let $\tilde{\mathbf{x}}$ be the convex combination of rules indexed by θ (with corresponding feasibility settings $\mathcal{X}^\theta$ and combination $\mathcal{X}$). Further let $\mathcal{F}$ be a convex combination of distributions, indexed by ω. Assume for all θ and ω, $\tilde{\mathbf{x}}^\theta$ has competitive efficiency μ on $\mathcal{F}^\omega$. We will argue that $\tilde{\mathbf{x}}$ has competitive efficiency μ on $\mathcal{F}$. This implies both stated claims, as either convex combination could be trivial.

Let $\mathbf{y}$ map value profiles to allocations feasible for the convex combination environment $\mathcal{X}$. Then for each profile $\mathbf{v}$, we can write $\mathbf{y}(\mathbf{v}) = \mathbb{E}_\theta[\mathbf{y}^\theta(\mathbf{v})]$ for some collection allocation functions $\mathbf{y}^\theta$ respectively feasible for $\mathcal{X}^\theta$. Moreover, for an agent i with value v_i the interim allocation $y_i(v_i) = \mathbb{E}_{\mathbf{v}}[y_i(\mathbf{v}) \,|\, v_i]$ satisfies $y_i(v_i) = \mathbb{E}_{\theta,\omega}[y_i^{\theta,\omega}(v_i) \,|\, v_i]$, where $y_i^{\theta,\omega}(v_i)$ is the interim allocation with respect to $\mathbf{y}^\theta$ under distribution $\mathcal{F}^\omega$. Now consider an agent i with value v_i, and consider an allocation level $y_i(v_i) = \mathbb{E}_{\theta,\omega}[y_i^{\theta,\omega}(v_i) \,|\, v_i]$. The following inequalities show that $\mathbb{E}_{\theta,\omega}[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \,|\, v_i) \,|\, v_i] \geq T_i(y_i(v_i) \,|\, v_i)$, explained after their statement:

$$
\begin{aligned}
T_i(y_i(v_i) \,|\, v_i) &= \int_0^\infty \max(y_i(v_i) - \tilde{x}_i(z \,|\, v_i), 0)\, dz \\
&= \int_0^\infty \max\left(\mathbb{E}_{\theta,\omega}\left[y_i^{\theta,\omega}(v_i) - \tilde{x}_i^{\theta,\omega}(z \,|\, v_i) \,\middle|\, v_i \right], 0 \right) dz \\
&\leq \mathbb{E}_{\theta,\omega}\left[\int_0^\infty \max(y_i^{\theta,\omega}(v_i) - \tilde{x}_i^{\theta,\omega}(z \,|\, v_i), 0)\, dz \,\middle|\, v_i \right] \\
&= \mathbb{E}_{\theta,\omega}\left[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \,|\, v_i) \,\middle|\, v_i \right].
\end{aligned}
$$

The first and last equalities are another way to write the integral defining the threshold surplus. The second equality is from the definitions of convex combination, and linearity of expectation. The inequality follows from convexity of the function $\max(\cdot, 0)$ and linearity of integration.

Now we write out the definition of the competitive efficiency for feasible allocation $\mathbf{y}$:

$$\begin{aligned}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_i s_i(v_i)\tilde{x}_i(\mathbf{s}(\mathbf{v}))\right] &= \mathbb{E}_{\theta,\omega}\left[\mathbb{E}_{\mathbf{v}\sim\mathcal{F}^\omega}\left[\sum_i s_i(v_i)\tilde{x}_i^\theta(\mathbf{s}(\mathbf{v}))\right]\right] \\ &\geq \mathbb{E}_{\theta,\omega}\left[\mu\sum_i \mathbb{E}_{v_i\sim\mathcal{F}_i^\omega}\left[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i)\,|\,v_i)\right]\right] \\ &= \mu\sum_i \mathbb{E}_{v_i\sim\mathcal{F}}\left[\mathbb{E}_{\theta,\omega\,|\,v_i}\left[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i)\,|\,v_i)\right]\right] \\ &\geq \mu\sum_i \mathbb{E}_{v_i\sim\mathcal{F}}\left[T_i(y_i(v_i)\,|\,v_i)\right].\end{aligned}$$

The first line follows from the definitions of convex combination and revenue. The second follows from applying competitive efficiency. The third line follows from properties of expectation, and the final line follows because $\mathbb{E}_{\theta,\omega}[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i)\,|\,v_i)\,|\,v_i] \geq T_i(y_i(v_i)\,|\,v_i)$, as argued above. We conclude that M has competitive efficiency μ on $\mathcal{F}$. □

As stated earlier, an important consequence of Lemma 8 is that competitive efficiency does not depend on the prior, but rather is inherent to the allocation rule.

Corollary 9. *An allocation rule $\tilde{\mathbf{x}}$ has competitive efficiency μ on a degenerate prior if and only if it has competitive efficiency μ on all priors.*

Proof. Assume $\tilde{\mathbf{x}}$ has competitive efficiency μ on a degenerate prior. Under a degenerate prior, Definition 4 is equivalent to the statement that for any bid profile $\mathbf{b}$ and fixed allocation profile $\mathbf{y} \in \mathcal{X}$,

$$\sum_i b_i\tilde{x}_i(\mathbf{b}) \geq \mu\sum_i T_i(y_i),$$

where degeneracy of the prior allows us to omit conditioning from $T_i(y_i)$. This inequality does not depend on the prior, and hence $\tilde{\mathbf{x}}$ has competitive efficiency μ on every degenerate prior. Now consider some nondegenerate prior $\mathcal{F}$. Since $\mathcal{F}$ can be written as a mixture over degenerate priors, the forward direction follows from Lemma 8. The reverse direction is immediate. □

Definition 10. *We say an allocation rule $\tilde{\mathbf{x}}$ has competitive efficiency μ if it has competitive efficiency μ on all priors.*

In the remainder of the section, we analyze the competitive efficiency of a variety of well-studied single-bid allocation rules. In Section 3.1, we consider the highest-bid-wins rule for single-item auctions, and prove a competitive efficiency of 1. For contrast, we then turn the multi-item setting of single-minded combinatorial auctions in Section 3.2. We show that for multiple items, highest-bids-win, which is welfare-optimal in the absence of incentives, has an undesirable competitive efficiency. In Section 3.3, we consider the competitive efficiency of greedy allocation rules. We observe that they generally lack the pathology of highest-bids-win single-minded combinatorial auctions. Finally, we demonstrate the usefulness of closure under convex combination by considering position auctions in Section 3.4. In Section 4, we will show how these competitive efficiency bounds imply a variety of worst-case welfare and revenue results for winner-pays-bid and all-pay mechanisms under several standard notions of equilibrium.

3.1 Single-item Multi-unit Auctions

This section considers the n-agent single-item and multi-unit highest-bids-win allocations. Under the k-unit highest-bids-win rule, each agent i submits a bid b_i, the and the k highest bidders win a unit. This generalizes the standard allocation rule for single-item auctions. By Lemma 9, it suffices to analyze degenerate priors. Since the highest-bids-win rule and k-unit environment are both deterministic, we may consider the simpler Definition 2. We first give the proof for single-item environments.

Theorem 11. *The highest-bids-win rule has competitive efficiency* 1 *in single-item environments.*

Proof. The bid surplus under bid profile $\mathbf{b}$ is the highest bid, i.e., $\sum_i b_i \tilde{x}_i(\mathbf{b}) = \max_i b_i$. Each agent's threshold bid $\hat{b}_i(\mathbf{b}_{-i})$ is at most the highest bid. Thus, for any feasible allocation $\mathbf{y}$, i.e., with $\sum_i y_i \leq 1$, we can bound the threshold surplus by the bid surplus.

$$\begin{aligned}\sum\nolimits_i \hat{b}_i(\mathbf{b}_{-i})\, y_i &\leq \max\nolimits_i b_i \sum\nolimits_i y_i \\ &\leq \sum\nolimits_i b_i \tilde{x}_i(\mathbf{b}).\end{aligned}$$ □

A very similar proof shows that the highest-bids-win rule for selling k-units to unit-demand agents has competitive efficiency 1. Rather than give the elementary proof here, we will observe it as a corollary of Theorem 23, given subsequently.

Theorem 12. *For any single-item or multi-unit environment, the highest-bids-win winner-pays-bid rule has competitive efficiency 1.*

3.2 Single-Minded Combinatorial Auctions: Highest Bids Win

In the Section 3.1, we saw that in single-item and k-unit settings, allocating the highest bidders has competitive efficiency 1. We now present a negative example, and show that in another natural multi-item setting, competition under the generalization of the highest-bids-win rule corresponds much less directly to bid surplus. Our multi-item setting and allocation rule are as follows:

Definition 13. *A* single-minded combinatorial auction *feasibility environment is defined by m indivisible items, n agents that each desire a bundle of items, and the constraint that no item can be allocated more than once. Agent i desires the set of items $\mathcal{S}_i$, she receives value v_i for receiving any superset of $\mathcal{S}_i$ and value* 0 *otherwise. An allocation vector $\mathbf{x} \in \{0,1\}^n$ is feasible if and only if for all agents $i \neq i'$, simultaneous allocation $x_i = x_{i'} = 1$ implies disjoint demands $\mathcal{S}_i \cap \mathcal{S}_{i'} = \emptyset$.*

Definition 14. *The* highest-bids-win rule *allocates the feasible set of bidders with the highest total bid.*

In a single-minded combinatorial auction under the highest-bids-win rule, a single bidder may exclude many others simultaneously. As an example, consider an $m+1$-agent environment, where agents $1, \ldots, m$ desire only the item with their same index, and agent $m+1$ desires the grand bundle of all items. Agent $m+1$ is mutually exclusive with any subset of the other agents. The impact of agent $m+1$ on the competition experienced by agents $1, \ldots, m$ far exceeds their impact on bid surplus. Under the bid profile $\mathbf{b} = (0, \ldots, 0, 1)$, the bid surplus is 1, while each agent $i \in \{1, \ldots, m\}$ now faces a threshold bid of $\hat{b}_i(\mathbf{b}_{-i}) = 1$. Since the allocation vector $\mathbf{y} = (1, \ldots, 1, 0)$ is feasible, this immediately implies:

Lemma 15. *There exists a single-minded combinatorial auction environment where the highest-bids-win rule does not have competitive efficiency μ for any $\mu > 1/m$.*

3.3 Greedy Auctions

The previous section demonstrates the inability of the highest-bids-win rule to effectively convert competition into bid surplus in multi-item settings. We now show that other rules may manage this relationship more effectively. In particular, this section considers allocation rules which are *greedy*, defined below.

Definition 16. *The* greedy by priority *rule is given by a profile* $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_n)$ *of nondecreasing priority functions mapping bids for each agent* i *to real numbers. It proceeds in the following way:*

1. *Sort agents in nonincreasing order of priority* $\psi_i(b_i)$.
2. *Initialize the set of winners* $S = \emptyset$.
3. *For each agent* i *in sorted order: if* $S \cup \{i\}$ *is feasible,* $S = S \cup \{i\}$.
4. *Return* S.

For example, the greedy by bid rule is given by priority functions $\psi_i(b_i) = b_i$ for all i. Greedy by priority rules may be defined in any feasibility environment. In many settings, including the single-minded combinatorial auction, greedy rules will be suboptimal in the absence of incentives — they may not select a set of winners with highest total bid. We will show that this suboptimality completely governs the competitive efficiency of greedy auctions. Since the design of approximately optimal greedy algorithms is well-studied, we obtain several competitive efficiency bounds as immediate corollaries. We first define our measure of approximate optimality.

Definition 17. *A bid allocation rule* $\tilde{\mathbf{x}}$ *is an* α-approximation *for a feasibility environment* $\mathcal{X}$ *if for any bid profile* $\mathbf{b}$ *and feasible allocation* $\mathbf{y}$*, we have:*

$$\sum\nolimits_i b_i \tilde{x}_i(\mathbf{b}) \geq \alpha \sum\nolimits_i b_i y_i.$$

We formalize the relationship between approximation and the competitive efficiency as follows.

Theorem 18. *For any feasibility environment* $\mathcal{X}$*, every greedy* α*-approximation* $\tilde{\mathbf{x}}$ *for* $\mathcal{X}$ *has competitive efficiency at least* α.

For single-minded combinatorial auctions, a greedy $1/\sqrt{m}$-approximation is well-known.

Lemma 19 (Lehmann et al., 2002)**.** *For any single-minded combinatorial auction environment, greedy by priority with* $\psi_i(b_i) = b_i/\sqrt{|\mathcal{S}_i|}$ *for all agents* i *is a* $1/\sqrt{m}$*-approximation.*

Another family of settings where greedy allocation rules are of particular interest are *matroids*, defined below, where the greedy by bid rule is known to be optimal i.e. a 1-approximation, absent incentives. Notable examples of matroids include k-unit environments, discussed in Section 3.1, and *transversal matroids*, which are matchable subsets of vertices on one side of a bipartite graph.

Definition 20. *A feasibility environment* $\mathcal{X}$ *is a* matroid *if the following two properties hold:*

i. *(Downward Closure) For any* $S \in \mathcal{X}$ *and* $i \in S$*,* $S \setminus \{i\} \in \mathcal{X}$.

ii. *(Augmentation Property) For any $S_1, S_2 \in \mathcal{X}$ with $|S_1| > |S_2|$, there exists $i \in S_1 \setminus S_2$ such that $S_2 \cup \{i\} \in \mathcal{X}$.*

Lemma 21. *For any matroid feasibility environment, greedy by priority with $\psi_i(b_i) = b_i$ for all agents i is a 1-approximation.*

We therefore obtain:

Theorem 22. *For any single-minded combinatorial auction environment, greedy-by-priority rule with priority function $\psi_i(b_i) = b_i/\sqrt{|\mathcal{S}_i|}$ for agent i, has competitive efficiency at least $1/\sqrt{m}$.*

Theorem 23. *For any matroid environment, the highest-bids-win rule has competitive efficiency 1.*

To prove Theorem 18, it is helpful to compare the behavior of greedy rules to the non-greedy highest bids win rule of Section 3.2, which had a poor competitive efficiency. In the example which proved Lemma 15, the high bid of the $(m+1)$th agent discouraged participation from the others - individually, each agent would have needed to bid $1+\epsilon$ to win. As a group, though, the losing agents could have won by increasing each of their bids by a tiny amount. Greedy rules lack this pathology. For any greedy allocation rule, we could increase the bids of every losing agent to their threshold without changing the outcome. We formalize this property as follows.

Definition 24. *Bid allocation rule $\tilde{\mathbf{x}}$ is* coalitionally non-bossy *if: for any profiles of bids $\mathbf{b}$ and $\mathbf{b}'$ where the bids in $\mathbf{b}'$ are the same as $\mathbf{b}$ for winners under $\mathbf{b}$ and at most their critical prices for losers under $\mathbf{b}$, i.e., if $\tilde{x}_i(\mathbf{b}) = 0$ then $b'_i \leq \hat{b}_i(\mathbf{b}_{-i})$; then the allocations of $\tilde{\mathbf{x}}$ under $\mathbf{b}$ and $\mathbf{b}'$ are the same, i.e. $\tilde{\mathbf{x}}(\mathbf{b}) = \tilde{\mathbf{x}}(\mathbf{b}')$.*

Lemma 25. *Any greedy by priority allocation rule is coalitionally non-bossy.*

Proof. Imagine changing $\mathbf{b}$ to $\mathbf{b}'$ by increasing one loser's bid at a time. Each time we increase a bid, say, of agent i, two things remain true: (1) i still loses: as long as $b'_i \leq \hat{b}_i(\mathbf{b}_{-i})$, i is passed over as infeasible when she is reached by the greedy rule; and (2) the threshold of every other losing agent i' remains unchanged: each losing agent's threshold is only determined by the bids of the agents who win. □

Lemma 26. *Any coalitionally non-bossy allocation rule has competitive efficiency at least its approximation ratio.*

Proof. Let $\mathbf{y}$ be a feasible allocation, $\mathbf{b}$ a profile of bids, and let $\mathbf{b}'$ be a vector of bids where losers under $\mathbf{b}$ bid $\hat{b}_i(\mathbf{b}_{-i})$, while winners bid as before. The following inequalities hold, with justifications after.

$$\begin{aligned}\sum\nolimits_i b_i \tilde{x}_i(\mathbf{b}) &= \sum\nolimits_i b'_i \tilde{x}_i(\mathbf{b}) \\ &= \sum\nolimits_i b'_i \tilde{x}_i(\mathbf{b}') \\ &\geq \alpha \sum\nolimits_i b'_i y_i \\ &\geq \alpha \sum\nolimits_i \hat{b}_i(\mathbf{b}_{-i}) y_i.\end{aligned}$$

The first line holds because $\mathbf{b}'$ differs from $\mathbf{b}$ only on the bids of losing agents. The second follows from the coalitional non-bossiness of greedy rules, and the third from the assumption that the greedy rule is an α-approximation. The last line follows from the fact that $\mathbf{b}'$ doesn't change the bids of winners under $\mathbf{b}$, and for those agents, $b_i \geq \hat{b}_i(\mathbf{b}_{-i})$. □

Theorem 18 follows from combining Lemma 25 with Lemma 26.

3.4 Position Auctions

The allocation rules and environments in Section 3.1-Section 3.3 have been deterministic. We now consider the canonical and inherently randomized allocation environment of position auctions and show that the most natural single-bid rule has competitive efficiency 1. This result follows directly from the competitive efficiency of multi-unit highest-bids-win allocation (Theorem 12) and closure of the competitive efficiency under convex combination (Lemma 8).

Position environments are a standard model for internet advertising auctions, e.g., Varian (2007b) and Edelman et al. (2007b). Advertisers (agents) compete for ad placement in positions in a list on a webpage. Each position has an associated clickthrough probability, and an agent is considered allocated when clicked. Feasible allocations are assignments of ads to positions, with the possibility for the auctioneer to exclude any agent from the auction.

Definition 27. *A* position environment *is given by position allocation probabilities* $1 \geq \alpha_1 \geq \ldots \geq \alpha_n \geq 0$*. An allocation vector* $\tilde{\mathbf{x}}$ *is feasible if there exists a permutation* π *over* $\{1, \ldots, n\}$ *such that for all agents* i $\tilde{x}_i \in \{0, \alpha_{\pi(i)}\}$*.*

The natural extension of highest-bids-win to position environments is the following:

Definition 28. *The* rank-by-bid *rule for position auctions assigns agents to slots in order of their bid. The agent in slot* j *wins with probability* α_j*.*

The following interpretation of position auctions as a convex combination of multi-unit auctions is well known in the literature, e.g., Devanur et al. (2015). The subsequent theorem combines this interpretation with Lemma 8 and Theorem 12.

Lemma 29. *The allocation rule of the generalized first price auction for position weights* $\alpha_1 \geq \ldots \geq \alpha_n$ *is equivalent the convex combination of single-item multi-unit highest-bids-win allocation rules where* k*-units are sold with probability* $\alpha_k - \alpha_{k+1}$ *(with* $\alpha_{n+1} = 0$*).*

Theorem 30. *For any position environment, the rank-by-bid rule has competitive efficiency* 1*.*

4 Welfare and Revenue Analysis

In the previous section, we studied the way mechanisms manage inter-agent competition. We focused on the mechanisms' allocation rules in isolation, agnostic to the choice of payment format and behavioral assumptions on participating agents. This section considers allocations and payments together, and turns the focus to agents' behavior. In equilibrium, agents respond to a range of allocations at different prices. Behavioral assumptions such as best response dictate the choice from this range, which in turn governs the agent's contribution to social welfare. In what follows, we define *individual efficiency*, which quantifies how an agent's solution to their bidding problem impacts welfare. Individual efficiency will depend on the shape of the allocation rule, the payment format, and the way the agent selects a bid (e.g. the degree of approximate best response). We then show how individual efficiency and competitive efficiency combine to imply robust equilibrium performance guarantees.

Equilibrium induces for each agent a single-agent interim mechanism, with a bid allocation rule $\tilde{x}$ and payment rule $\tilde{p}$. For each bid b, the bid allocation rule and payment rule induce a utility $\tilde{u}(b) = v\tilde{x}(b) - \tilde{p}(b)$. For the winner-pays bid payment rule, an agent with value v receives utility $\tilde{u}(b) = (v - b)\tilde{x}(b)$, and for an all-pay rule, we have $\tilde{u}(b) = v\tilde{x}(b) - b$. The agent's competition is summarized by the bid allocation rule $\tilde{x}$ they face. In Section 3 we measured the strength of this competition by the threshold surplus $T(\cdot)$. Individual efficiency measures the tradeoff between this competition and utility.

Definition 31. *Let $M = (\tilde{x}, \tilde{p})$ be a single-agent mechanism and $b : \mathbb{R}_+ \to \mathbb{R}_+$ a bid function mapping agent values to bids. The* individual efficiency *of $b(\cdot)$ in M is the largest η such that for all $z \in [0,1]$ and values $v \geq 0$, $\tilde{u}(b(v)) + T(z) \geq \eta v z$.*

Definition 31 can be motivated by considering winner-pays-bid mechanisms in the following single-agent problem. A seller faces a choice between allocating a buyer with fixed value v or seeking value from an outside option. The outside option's value v_0 is distributed with CDF equal to $\tilde{x}$. The welfare-maximizing outcome allocates the buyer if $v \geq v_0$, taking the outside option otherwise. The buyer's contribution to social welfare is therefore vz^*, where z^* is the probability that $v \geq v_0$. Individual efficiency η measures the portion of the buyer's surplus captured by a winner-pays-bid mechanism. The first term on the left, $\tilde{u}(b(v))$, captures the contribution from the buyer's utility. The second term, $T(z^*)$ is equal to $\mathbb{E}[v_0 \mid v \geq v_0] \Pr[v \geq v_0]$. Since the seller's revenue is always at least v_0, this lower bounds the seller's contribution to the welfare.

To enable robust analysis, we will consider individual efficiency taken in the worst case over broad families of single-agent mechanisms likely to arise in equilibrium. For example, in Nash equilibria of a single-item first-price auction, each agent faces a single-agent mechanism with a 0-1 allocation rule stepping up at a threshold. The model of agent behavior (e.g. best response) in turn governs the bid functions that arise in response to allocation rules. Taking these two together yields the following.

Definition 32. *Given a set of single-agent mechanisms $\mathcal{M}$ and a behavioral model B mapping single-agent mechanisms M to bid functions b_M, the individual efficiency of B on $\mathcal{M}$ is the infimum of the individual efficiency of b_M in M, taken over all $M \in \mathcal{M}$.*

In Section 4.1, we consider individual efficiency for three pairings of mechanism family (winner-pays-bid/all-pay) and behavioral model (approximate/exact best response). We then connect the single-agent analyses to performance guarantees in auction equilibria in Section 4.2. If equilibrium induces single-agent mechanisms with individual efficiency η in a mechanism with competitive efficiency μ, then we show in Section 4.2.1 that the mechanism has welfare approximation $\mu\eta$. Furthermore, we show in Section 4.2.2 that analogous guarantees hold for the objective of revenue in Bayes-Nash equilibrium of winner-pays-bid mechanisms with carefully-selected reserve prices. Finally, we trace out the limits of our approach by exhibiting examples with high loss in Section 4.3.

4.1 Individual Efficiency Analyses

Individual efficiency quantifies the performance of a behavioral model when faced with a family of single-agent mechanisms. To study equilibria of winner-pays-bid and all-pay auctions, we consider the single-agent mechanisms that arise from those auctions. For winner-pays-bid mechanisms, approximate best response in full-information environments generates ϵ-Nash equilibria, which induce 0-1 (i.e. deterministic) single-agent mechanisms (Section 4.1.1). Similarly, best response under incomplete information generates winner-pays bid single-agent mechanisms with possibly randomized allocations (Section 4.1.2). Finally, we study single-agent, randomized all-pay mechanisms under best response (Section 4.1.3).

4.1.1 Deterministic Rules with Approximate Best Response

This section characterizes the individual efficiency of $(1+\epsilon)$-best response bids against deterministic winner-pays-bid mechanisms, i.e. those with allocations in $\{0,1\}$.

Definition 33. *A bid b is an $(1-\epsilon)$-*best response *for an agent with value v if it maximizes utility up to a factor of $(1-\epsilon)$: for any alternative bid b', $v\tilde{x}(b) - \tilde{p}(b) \geq (1-\epsilon)(v\tilde{x}(b') - \tilde{p}(b'))$.*

Note that taking $\epsilon = 0$ yields exact best response. We give a tight characterization as follows.

Lemma 34. *The individual efficiency of $(1-\epsilon)$-best response in winner-pays-bid mechanisms with deterministic allocation rules is $(1-\epsilon)$.*

Proof. We give an argument lower bounding the individual efficiency, and then exhibit a particular example where our bound holds with equality. Let $\tilde{x}$ step up to 1 at $\hat{b} \in [0,\infty)$. Then $T(z) = \hat{b}z$. For any $\delta > 0$, the agent could bid $\hat{b}+\delta$ and win. Hence, for any agent value v, a $(1-\epsilon)$ best response $b(v)$ satisfies $\tilde{u}(b(v)) \geq (1-\epsilon)(v - \hat{b} - \delta)$. Since δ is arbitrary, we obtain $\tilde{u}(b(v)) + \hat{b} \geq (1-\epsilon)v$. We may further weaken this to obtain the desired inequality for all $z \in [0,1]$:

$$\tilde{u}(b(v)) + \hat{b}z \geq (1-\epsilon)vz.$$

This is the best possible worst-case bound. An example for which it holds with equality is $v = (1+\epsilon)$, $z = 1$, $\hat{b} = 0$, and $b(v) = \epsilon$. □

Note that this analysis relied heavily on the allocation rule being deterministic. We present an example demonstrating this assumption to be necessary in Section 4.3.2.

4.1.2 Winner-Pays-Bid With Randomized Allocation

Randomized allocation rules present a richer bidding problem, as agents may now select from many different allocation levels and associated payments. Randomness may stem from incomplete information with respect to other agent's values or from randomness in the mechanism or feasibility environment. Below, we give a tight single-agent analysis under exact best response.

Lemma 35. *The individual efficiency of exact best response in winner-pays-bid mechanisms with randomized allocation rules is $(e-1)/e$.*

Proof. We give a lower bounding argument, followed by an example for which the analysis is tight. Let allocation rule $\tilde{x}$ and value v be given. An agent's best response bid $b(v)$ must maximize $\tilde{u}(b) = (v-b)\tilde{x}(b)$. For any allocation probability x, the agent could get allocation probability at least x by bidding arbitrarily close to $t(x) = \inf\{b \,|\, \tilde{x}(b) \geq x\}$. Hence, $\tilde{u}(b(v)) \geq (v - t(x))x$ for all $x \in [0,1]$. We may rearrange this as $t(x) \geq v - \tilde{u}(b(v))/x$. Since we also have $t(x) \geq 0$, we may write:

$$\begin{aligned} T(z) \geq \int_0^z \max(v - \tfrac{\tilde{u}(b(v))}{x}, 0)\, dx &= \int_{\tilde{u}(b(v))/v}^z v - \tfrac{\tilde{u}(b(v))}{x}\, dx \\ &= vz - \tilde{u}(b(v))(1 - \ln \tfrac{\tilde{u}(b(v))}{vz}). \end{aligned}$$

We therefore have $\tilde{u}(b(v)) + T(z) \geq vz - \tilde{u}(b(v))\ln(\tilde{u}(b(v))/vz)$. Holding v fixed and minimizing the righthand side as a function of $\tilde{u}(b(v))$ yields the inequality $\tilde{u}(b(v)) + T(z) \geq vz(e-1)/e$. This lower bounds individual efficiency by $e/(e-1)$. An example exhibiting tightness can be produced by choosing $v = 1$, $\tilde{x}(b) = (e(1-b))^{-1}$ for $b \in [0, 1-1/e]$, and $b = 0$. □

4.1.3 All-Pay Mechanisms

For all-pay mechanisms, we consider only randomized allocation rules. Pure approximate Nash equilibria do not exist, rendering an analysis of deterministic mechanisms unnecessary.

Lemma 36. *The individual efficiency of exact best response in all-pay mechanisms with randomized allocation rules is* $1/2$*.*

Proof. To lower bound the individual efficiency, note that the agent selects their best response bid b to maximize their utility $\tilde{u}(b) = v\tilde{x}(b) - b$. For any allocation probability $x \in [0, 1]$, agent i could choose to get allocation probability at least x by bidding their interim threshold $t(x)$. Hence, $\tilde{u}(b(v)) \geq vx - t(x)$ for all $x \in [0, 1]$. We may rearrange this as $t(x) \geq vx - \tilde{u}(b(v))$. Since we also have $t(x) \geq 0$, we may write:

$$T(z) \geq \int_0^z \max(vx - \tilde{u}(b(v)), 0)\, dx = \int_{\tilde{u}(b(v))/v}^z vx - \tilde{u}(b(v))\, dx$$

$$= \tfrac{v}{2} - \tilde{u}(b(v)) + \tfrac{\tilde{u}(b(v))^2}{2v}.$$

We therefore have $\tilde{u}(b(v)) + T(z) \geq v/2 + \tilde{u}(v)^2/2v$. Holding v fixed and minimizing the righthand side as a function of $\tilde{u}(b(v))$ yields a lower bound of $vz/2$, as desired. To exhibit tightness, choose $v = 1$, $\tilde{x}(b) = b$ for $b \in [0, 1]$, and $b = 0$. □

4.2 Performance Guarantees

Competitive efficiency and individual efficiency together yield performance guarantees for auction equilibria. Section 4.2.1 gives the welfare consequences. In Section 4.2.2, we then use the reduction from revenue maximization to welfare maximization of Myerson (1981) to show that similar guarantees apply to winner-pays-bid mechanisms with reserves when values are independent.

4.2.1 Robust Welfare Guarantees

Individual efficiency was motivated in terms of the welfare properties of a single-buyer allocation problem, where a seller chose between the buyer and an outside option. In multi-buyer mechanisms, the seller's outside options are endogenously generated by competition; competitive efficiency measures the way competition translates into revenue. Hence, individual efficiency and competitive efficiency combine to bound welfare in equilibrium. This discussion can be formalized as follows. For a mechanism $M = (\tilde{\mathbf{x}}, \tilde{\mathbf{p}})$ and prior $\mathcal{F}$, any strategy profile $\mathbf{s}$ induces an *interim mechansim* for agent i with value v_i with interim allocation rule $\tilde{x}_i(b \,|\, v_i) = \mathbb{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(b, \mathbf{s}_{-i}(\mathbf{v}_{-i}) \,|\, v_i]$ and interim payment rule $\tilde{p}_i(b \,|\, v_i) = \mathbb{E}_{\mathbf{v}_{-i}}[\tilde{p}_i(b, \mathbf{s}_{-i}(\mathbf{v}_{-i}) \,|\, v_i]$. The interim mechanisms induced by winner-pays-bid and all-pay mechanisms will have those same payment formats.

Theorem 37. *Given mechanism* $M = (\tilde{\mathbf{x}}, \tilde{\mathbf{p}})$ *let* $\tilde{\mathbf{x}}$ *have competitive efficiency* $\mu \leq 1$*. Let* $\mathcal{M}$ *be a set of single-agent mechanisms, and* B *a behavioral model mapping values to bids for every mechanism with individual efficiency* η *on* $\mathcal{M}$*. Then for any prior* $\mathcal{F}$ *and strategy profile* $\mathbf{s}$ *for* M *where each agent's interim mechanism is in* $\mathcal{M}$ *and each player's bid is selected according to* B*, the expected welfare is a* $\mu\eta$*-approximation to the optimal welfare.*

Proof. For any value v_i for agent i, let $x_i^*(v_i)$ denote agent i's interim allocation probability under the welfare-optimal allocation rule. The individual efficiency of B on $\tilde{X}$ implies:

$$u_i(v_i) + T_i(x_i^*(v_i) \mid v_i) \geq \eta v_i x_i^*(v_i).$$

Summing over all agents and taking expectation over $\mathbf{v}$ yields:

$$\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i u_i(v_i)\right] + \mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i T_i(x_i^*(v_i) \mid v_i)\right] \geq \eta\, \mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i v_i x_i^*(v_i)\right].$$

The righthand side is the optimal expected welfare. The second term on the left can be rewritten as $\mathbb{E}_{\mathbf{v}}[\sum_i T_i(x_i^*(v_i) \mid v_i)] = \sum_i \mathbb{E}_{v_i}[T_i(x_i^*(v_i) \mid v_i)]$. We may therefore apply competitive efficiency to obtain:

$$\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i u_i(v_i)\right] + \tfrac{1}{\mu}\textsc{Rev}(M, \mathbf{s}, \mathcal{F}) \geq \eta \textsc{Welfare}(\textsc{Opt}, \mathcal{F}).$$

the result then follows from noting that $\mu \leq 1$ and that the welfare of M is the sum of the expected utilities and revenue. □

Combining Theorem 37 with the individual efficiency guarantees of Lemmas 34 and 35 to obtain the following welfare bounds:

Corollary 38. *Any winner-pays-bid mechanism with deterministic allocation rule and competitive efficiency $\mu \leq 1$ has worst-case welfare approximation $\mu(1-\epsilon)$ in ϵ-Nash equilibrium.*

Corollary 39. *Any winner-pays-bid mechanism with competitive efficiency $\mu \leq 1$ has worst-case welfare approximation $\mu(e-1)/e$ in Bayes-Nash equilibrium.*

Corollary 40. *Any all-pay mechanism with competitive efficiency $\mu \leq 1$ has worst-case welfare approximation $\mu/2$ in Bayes-Nash equilibrium.*

4.2.2 Revenue in Bayes-Nash Equilibrium

Individual efficiency quantifies the way a buyer chooses to trade off their utility and the seller's utility (via competitive efficiency) against their value. We now extend these ideas to study the objective of seller revenue in Bayes-Nash equilibrium with independently distributed values. For winner-pays-bid mechanisms with suitably chosen reserve prices, we will obtain a robust revenue approximation of $\mu(e-1)/2e$ for winner-pays-bid mechanisms with competitive efficiency μ.

Given a Bayes-Nash equilibrium $\mathbf{s}$ for product distribution $\mathcal{F}$, Myerson (1981) shows that the ex ante expected payment of an agent i is $\mathbb{E}_{v_i}[\phi_i(v_i)x_i(v_i)]$, where $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ is the *virtual value* for value v_i and F_i (resp. f_i) the cumulative distribution function (resp. probability density function) of i's value distribution. It follows that $\textsc{Rev}(M, \mathbf{s}, \mathcal{F}) = \mathbb{E}_{\mathbf{v}}[\sum_i \phi_i(v_i)\tilde{x}_i(\mathbf{s}(\mathbf{v}))]$. We refer to an atomless distribution with $\phi_i(v_i)$ nondecreasing in v_i as *regular*. For regular distributions, the revenue-optimal mechanism chooses the allocation with the highest virtual surplus $\sum_i \phi_i(v_i)x_i$. For downward-closed settings, agents with negative virtual value are excluded via *monopoly reserves*, given by $r_i^* = \inf\{v_i \mid \phi_i(v_i) \geq 0\}$.

The above characterization rewrites equilibrium revenue in terms of buyer surplus in a transformed value space. Revenue maximization then amounts to excluding agents with negative virtual values, and maximizing virtual welfare among those that remain. The following lemma extends Lemma 35 to reason about an agent's contribution to revenue under monopoly reserves. We assume agents who are indifferent between a bid of 0 and a positive bid choose the latter.

Lemma 41. *Let F be a regular value distribution with monopoly reserve r^*. Further let M be a single-agent winner-pays-bid mechanism with allocation rule $\tilde{x}$, with $\tilde{x}(b) = 0$ for all $b \leq r^*$. Then for any v, best response bid $b(v)$ and $z \in [0, 1]$,*

$$\phi(v)\tilde{x}(b(v)) + T^{r^*}(z) \geq \tfrac{e-1}{e}\phi(v)z. \tag{8}$$

Proof. Since $\mathcal{F}_i$ is regular, $\phi_i(v_i) < 0$ if and only if $v_i < r_i^*$, in which case i is excluded by both the optimal mechanism and $M^{\mathbf{r}^*}$, i.e. $\tilde{x}_i(\mathbf{s}(\mathbf{v})) = x_i^*(\mathbf{v}) = 0$. For such agents, (10) holds trivially.

For an agent with value $v \geq r^*$, all best response bids are at least r^*. Hence for any best response $b(v)$, we have $\tilde{p}(b(v)) \geq r^*\tilde{x}(r^*)$. Furthermore, if $x \leq \tilde{x}(r^*)$, then $t(x) \leq r^*$. It follows that $\int_0^{\tilde{x}(r^*)} t(x)\,dx \leq r^*\tilde{x}(r^*)$. We therefore obtain the following sequence of inequalities:

$$\tilde{p}(b(v)) \geq r^*\tilde{x}(r^*) \geq \int_0^{\tilde{x}(r^*)} t(x)\,dx \geq T(z) - T^{r^*}(z).$$

Combining with Lemma 35, we obtain a tradeoff between an agent's contribution to surplus and their threshold surplus:

$$v\tilde{x}(b(v)) + T^{r^*}(z) = \tilde{u}(b(v)) + \tilde{p}(b(v)) + T^{r^*}(z) \geq \tilde{u}(b(v)) + T(z) \geq \tfrac{e-1}{e}vz. \tag{9}$$

The lemma then follows from noting that $\phi(v) \leq v$, and hence when $\phi(v) \geq 0$, the inequality (8) is a weakening of (9). □

We may apply Lemma 41 to obtain a revenue guarantee for Bayes-Nash equilibria.

Theorem 42. *In any Bayes-Nash equilibrium of a winner-pays-bid mechanism with competitive efficiency $\mu \geq 1$ and monopoly reserves $\mathbf{r}^*$ for regular product distribution $\mathcal{F}$, the expected revenue is a $\mu(e-1)/2e$-approximation to that of the optimal mechanism.*

Proof. Fix a value v_i for agent i. If $v_i \geq r_i^*$, Lemma 41 implies

$$\phi_i(v_i)x_i(v_i) + T_i^{r_i^*}(x_i^*(v_i)) \geq \tfrac{e-1}{e}\phi_i(v_i)x_i^*(v_i), \tag{10}$$

where $x_i^*(v_i)$ denotes the interim allocation probability for i under allocation rule of the revenue-optimal mechanism $\textsc{Opt}_{\mathcal{F}}$ for $\mathcal{F}$. We omit conditioning from $T_i^{r_i^*}(x_i^*(v_i))$, as $\mathcal{F}$ is a product distribution. We may sum (10) over all agents and take expectations to obtain

$$\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i \phi_i(v_i)x_i(v_i)\right] + \mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i T_i^{r_i^*}(x_i^*(v_i))\right] \geq \tfrac{e-1}{e}\,\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i \phi_i(v_i)x_i^*(v_i)\right].$$

Applying competitive efficiency and noting that $\mu \geq 1$ yields:

$$\tfrac{1}{\mu}\Big(\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i \phi_i(v_i)x_i(v_i)\right] + \textsc{Rev}(M^{\mathbf{r}^*}, \mathbf{s}, \mathcal{F})\Big) \geq \tfrac{e-1}{e}\,\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i \phi_i(v_i)x_i^*(\mathbf{v})\right].$$

Since a mechanism's expected revenue is equal to its expected virtual surplus, we obtain the desired revenue guarantee:

$$2\textsc{Rev}(M^{\mathbf{r}^*}, \mathbf{s}, \mathcal{F}) \geq \mu\tfrac{e-1}{e}\textsc{Rev}(\textsc{Opt}_{\mathcal{F}}, \mathcal{F}).$$

□

4.3 Lower Bounds

To conclude the section, we explore the limits of our approach. We do so by exhibiting three sets of examples. In Section 4.3.1, we give the worst-known examples for welfare and revenue loss in the single-item first-price auction in Bayes-Nash equilibrium. Section 4.3.2 gives a Nash equilibrium of a mechanism with randomized allocations, competitive efficiency 1, and a factor of $(e-1)/e$ welfare loss. This shows that the restriction to deterministic allocation rules in our analysis of ϵ-Nash equilibrium was necessary. Finally, Dütting and Kesselheim (2015) study the extent to which low competitive efficiency is a necessary condition for a good robust welfare guarantee. We briefly discuss their results in Section 4.3.3.

4.3.1 Single-Item Lower Bounds

The competitive efficiency and individual efficiency provide an general framework for robust welfare and revenue analysis. Whether the resulting welfare guarantees are best possible will vary across mechanisms and solution concepts. What follows are three equilibria of the single-item first-price auction which are the worst known for their type. We begin with a Bayes-Nash equilibrium with correlated values which meets our welfare guarantee exactly, due to Syrgkanis (2014).

Example 43. Our example has three agents. Agents 1 and 2 have values perfectly correlated, and drawn according to the distribution with CDF $(e(1-v))^{-1}$ for $v \in [0, 1-1/e]$. Agent 3 has value deterministically 1. If we break ties in favor of agent 3, it is a Bayes-Nash equilibrium for agents 1 and 2 to bid their value and agent 3 to bid 0. In the optimal allocation, agent 3 always wins. A straightforward calculation shows the equilibrium welfare to be $(e-1)/e$.

Example 43 provides a canonical setting in which our welfare analysis is tight. When values are independently distributed, however, Hoy et al. (2018) exploit independence to improve the welfare approximation guarantee from $(e-1)/e \approx .63$ to $\approx .74$. We now give the give the worst-known example for welfare under independent distributions, with a multiplicative loss of $\approx .869$.

Example 44. Our example will have $n+1$ agents. Agents $1, \ldots, n$ will be designated *low-valued agents*, with an identical value distribution to be determined shortly. Agent $n+1$ will be the *high-valued agent*, with value deterministically 1. Allocating to agent $n+1$ in all value profiles yields a lower bound on the optimal welfare of 1. Our constructed equilibrium will occasionally misallocate to low-valued agents, yielding an expected welfare of approximately .869.

We will design the bid distribution of the low-valued agents to make the high-valued agent indifferent over an interval of bids. This will allow us to select a mixed strategy for the high-valued bidder supported on this interval. To do so, fix in advance the expected utility $u_H \in [0, 1]$ of the high-valued agent. The utility u_H will be a parameter which defines a family of examples constructed as below. Let G_L denote the CDF of the bid distribution of an individual low-valued agent. Then the CDF of the distribution of the highest-bidding low-valued agent is G_L^n. Note that if $G_L^n(b) = u_H/(1-b)$, then any bid $b \in [0, 1-u_H]$ for the high-valued bidder yields an expected utility of exactly u_H (breaking ties in favor of agent $n+1$). We will therefore take $G_L(b) = (u_H/(1-b))^{1/n}$.

We have not yet derived a value distribution for the low-valued agents, and we have not derived a bid distribution for the high-valued agent. Given a bid distribution G_H for the high-valued agent, the value distribution for the low-valued agents can be derived from first-order conditions. In other words, for any individual low-valued agent $i \in 1, \dots, n$, agent i is facing the distribution of highest competing bid given by $G_C(b) = G_L^{n-1}(b) G_H(b)$. Agent i bids to maximize $(v_i - b) G_C(b)$. For any $b_i \in (0, 1 - u_H)$, first-order conditions imply that $v_i = b_i + G_C(b_i)/g_C(b_i)$, where $g_C(b) = G'_C(b)$ is the density of agent i's competing bid distribution at b. This mapping immediately implies a value distribution for the low-valued agents.

All that remains is to select a mixed strategy for the high-valued agent, their expected utility parameter, u_H, and a number of low-valued agents n. To produce an equilibrium with low welfare, we must navigate a tradeoff. If G_H is too aggressive, then the high-valued agent will win frequently, yielding high welfare. If G_H is too weak, then noting the formula for the low-valued agents' values, we see that these values will generally be high. A similar tradeoff applies in selecting u_H. Numerical experimentation shows that choosing $G_H(b) = \sqrt{b/(1 - u_H)}$ and $u_H = .57$ yields low welfare. Given these choices, one can compute the expected welfare in equilibrium as approximately .869 for very large n. As mentioned, allocating the high-valued agent yields a lower bound of 1 on the optimal social welfare. This implies the desired welfare ratio.

Finally, we give the worst-known example for revenue in Bayes-Nash equilibrium with monopoly reserves, with independent, regularly-distributed values. As with welfare under independence, there is a gap between our example and our robust guarantee.

Example 45. Our equilibrium will have two agents. The first agent will have value deterministically 1, and the second will have value CDF $F_2(v_2) = 1 - 1/v_2$. The monopoly reserve for agent 1 is trivially 1. For agent 2, any price above 1 has virtual value 0, so the monopoly reserve is ambiguous. Perturb F_2 slightly such that the monopoly reserve is 1. In equilibrium, agent 1 bids their value, 1. If we break ties in favor of agent 2, then agent 2 also bids 1, yielding an expected revenue of 1. This obtains less revenue than posting a price of $H > 1$ to agent 2, and upon rejection selling to agent 1, which has expected revenue $2 - 1/H$. Taking $H \to \infty$ yields the desired multiplicative loss of 1/2.

4.3.2 Welfare Loss Under Randomized Mechanisms

Section 4.1.1 gave an improved individual efficiency guarantee for deterministic allocation rules when compared to randomized rules. This translated to an improved welfare guarantee. We now demonstrate that the restriction to deterministic rules is necessary for this improvement. We do so by giving a mechanism with competitive efficiency 1, randomized allocations, and Nash equilibrium with multiplicative welfare loss $(e - 1)/e$.

Example 46. Define a *partial allocation environment* to be an n-agent feasibility environment given by a vector $(z_1, \dots, z_n)$ of maximum allocations. A feasible allocation selects an agent i and assigns them up to z_i units of allocation. Given a value profile $\mathbf{v}$, the welfare-optimal allocation in a partial allocation environment selects the agent maximizing $v_i z_i$. The highest bids win allocation rule takes a profile of bids $\mathbf{b}$ and allocates the agent maximizing $b_i z_i$. This mechanism has competitive efficiency 1.

Now consider the following convex combination of partial allocation environments. Draw a parameter θ according to $G(\theta) = (e(1 - \theta))^{-1}$, and consider the three-agent partial allocation environment with maximum allocations $(\theta, \theta, 1)$. By Lemma 8 competitive efficiency is closed under convex combination, so the winner-pays-bid, highest-bids-wins mechanism for this environment has competitive efficiency 1. We will now construct a Nash equilibrium for this mechanism with welfare approximation $e/(e-1)$. Let values be $(1, 1, 1)$. If we allocate to break ties to favor agent 3, then it is a Nash equilibrium for agents 1 and 2 to bid 1, and agent 3 to bid 0. Equilibrium welfare is therefore the expected value of θ, which is $(e-1)/e$, rather than the 1 that could be achieved by allocating agent 3.

4.3.3 Necessity of Low Competitive Efficiency

We have shown that low competitive efficiency is a sufficient condition for robust performance guarantees. We now briefly examine whether it is also a necessary condition. In other words, we study the extent to which a mechanism with poor competitive efficiency must also possess equilibria with welfare far from optimal. A formal treatment of this question appears in Dütting and Kesselheim (2015). For brevity, we simply overview the main ideas.

In Section 3.2, we exhibited the highest-bids-win rule for single-minded combinatorial auctions as a mechanism with poor competitive efficiency. We gave a bundle structure with m items and an agent i desiring item i for each $i \in \{1, \ldots, m\}$, along with an agent $m+1$ who desires the grand bundle. The bid profile $(0, \ldots, 0, 1)$ then refuted any competitive efficiency better than m. Note, however, that under the winner-pays-bid format, this bid profile is not an equilibrium for any positive value for agent $m + 1$, as this agent would always prefer to lower their bid. If we allow ourselves to augment the setting with an additional bidder, however, we may produce an equilibrium with welfare approximation m. Specifically, add a second grand bundle bidder, $m + 2$. The value profile $(1, \ldots, 1)$ and bid profile $(0, \ldots, 0, 1, 1)$ is a Nash equilibrium for any tiebreaking: the only unilateral deviations in which bidders $1, \ldots, m$ can win involve bidding at least their value, and if bidders $m + 1$ or $m + 2$ bid less than 1, they are guaranteed to lose.

The approach of duplicating bidders can be extended to most winner-pays-bid mechanisms of interest for single-minded combinatorial auctions. In particular, consider any winner-pays bid mechanism M which, given two or more bidders with identical desired bundles, always allocates the bidder among the group with the highest bid. Let bid profile $\mathbf{b}$ and alternate allocation $\mathbf{y}$ exhibit a competitive efficiency at most μ. For each winner i under $\mathbf{b}$ in M, add a duplicate i' desiring the same bundle. The following value and bid profiles are then a Nash equilibrium with welfare approximation μ in the augmented setting: for each winner i under $\mathbf{b}$, give i and i' value equal to their equilibrium bid equal to b_i. For each agent i winning under $\mathbf{y}$ but not winning under $\mathbf{b}$, give that agent value $\hat{b}_i(\mathbf{b}_{-i})$ and bid b_i. Give all other agents value and bid 0. Under appropriate tiebreaking, this is an equilibrium for the same reasons as in the previous example: winners (and their duplicates) under b_i cannot reduce their bids without losing, and losers under $\mathbf{b}$ cannot win without overbidding. Moreover, the welfare approximation is equal to the competitive efficiency exhibited by $\mathbf{b}$ and $\mathbf{y}$, μ.

Dütting and Kesselheim (2015) formalize the above approach for arbitrary single-parameter environments. For any setting which can be augmented with duplicates, and any winner-pays-bid mechanism which handles duplicates sensibly, any instance with competitive efficiency μ implies the existence of a related instance and equilibrium for that related instance with welfare approximation μ. For single-minded combinatorial auctions in particular, this implies that up to a constant factor, the greedy mechanism discussed in Section 3.3 is the optimal winner-pays-bid mechanism, with respect to the objective of robust welfare approximation. It also pinpoints the highest-bids-wins mechanism's mismanagement of inter-bidder competition as the source of its worst-case inefficiency. More generally, it further justifies the study of competitive efficiency as a design objective in itself.

5 Beyond Single-Bid Mechanisms

Our analysis thus far has considered only single-bid mechanisms. This section adapts our measure of competition to arbitrary mechanisms for single-parameter agents. By generalizing the idea of competitive efficiency and our individual efficiency analysis, we show that analogous relationships continue to hold. We then apply our generalized framework to the sale of identical items via simultaneous auctions to unit-demand agents. Agents may participate in any number of auctions, and are served if they win at least one. Under our new definition, if all component mechanisms mechanism have competitive efficiency μ, we show that their simultaneous composition must as well.

5.1 Generalized Framework

The robust analysis of single-bid mechanisms consisted of three steps. First, we quantified an agent's competition in terms of their threshold bids. Second, competitive efficiency measured the way a mechanism managed this competition across agents, independent of incentives. Finally, individual efficiency related this competition to each agent's bidding problem: in winner-pays-bid (resp. all-pay) mechanisms, bids determine the price per unit (resp. price) of allocation. In general single-parameter mechanisms, each agent i submits an action a_i, which may be a richer object than a real-valued bid. Nonetheless, competitive efficiency and individual efficiency generalize with a suitable new measure of competition. To illustrate, in this section we measure competition directly using the *price per unit* associated with different levels of allocation. This will generalize our results for winner-pays-bid mechanisms from the previous sections, e.g. to simultaneous first-price auctions. A similar approach, instead measuring competition with *price* of allocation, would suffice to generalize our all-pay results.

For a general single-parameter mechanism M and strategy profile $\mathbf{s}$ for prior $\mathcal{F}$, each action a_i obtains for agent i with value v_i an interim allocation $\tilde{x}_i(a_i) = \mathbb{E}_{\mathbf{v}_{-i}\,|\,v_i}[\tilde{x}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i})\,|\,v_i]$ and interim payment $\tilde{p}_i(a_i) = \mathbb{E}_{\mathbf{v}_{-i}\,|\,v_i}[\tilde{p}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i})\,|\,v_i]$. Dividing the former quantity by the latter yields the *price per unit* associated with a_i. For winner-pays-bid mechanisms, the price per unit of a bid b_i is exactly b_i itself.

Definition 47. *The* price per unit *for action a_i, denoted $\beta_i(a_i\,|\,v_i)$, is given by $\beta_i(a_i\,|\,v_i) = \tilde{p}_i(a_i\,|\,v_i)/\tilde{x}_i(a_i\,|\,v_i)$.*

When an agent faced strong competition in a winner-pays-bid mechanism, this was reflected as a high threshold bid. In other words, we said competition was strong when high allocation came with a high price per unit of allocation. We adopt this same perspective for general mechanisms by studying actions on the Pareto frontier between allocation probability and price per unit of allocation.

Definition 48. *Given prior $\mathcal{F}$, strategy profile $\mathbf{s}$, value v_i for agent i, and target allocation probability $x \in [0, 1]$, an agent i's* interim threshold price *is given by $\tau_i(x\,|\,v_i) = \inf_{a_i\,:\,\tilde{x}_i(a_i\,|\,v_i)\geq x} \beta_i(a_i\,|\,v_i)$.*

General single-parameter mechanisms need not support a natural notion of a reserve price. One exception is simultaneous first-price auctions, where an individualized reserve of r_i in each mechanism for each agent i manifests as a minimum price per unit of r_i for agent i. This excludes agents with value below r_i. The following definition will enable Bayes-Nash revenue analysis under these circumstances.

Definition 49. *Given strategy profile* $\mathbf{s}$*, prior* $\mathcal{F}$*, desired allocation probability* $x \in [0,1]$*, and minimum price per unit* r_i*, an agent* i*'s* interim threshold price with discounted reserve *is given by* $\tau_i^{r_i}(x \mid v_i) = \tau_i(x \mid v_i)$ *if* $\tau_i(x \mid v_i) \geq r_i$ *and* 0 *otherwise.*

Finally, we will aggregate threshold prices in a manner analogous to Definition 3.

Definition 50. *Given strategy profile* $\mathbf{s}$*, prior* $\mathcal{F}$*, allocation probability* $y_i \in [0,1]$*, and minimum price per unit* r_i*, agent* i*'s* generalized threshold surplus with discounted reserve *for* y_i *is given by* $\mathcal{T}_i^{r_i}(y_i \mid v_i) = \int_0^{y_i} \tau_i^{r_i}(x \mid v_i)\, dx$*. When* $r_i = 0$*, we omit* r_i *and simply write* $\mathcal{T}_i(y_i \mid v_i)$*.*

We now state our generalized definition of competitive efficiency. Without real-valued bids, a mechanism's bid surplus is no longer well-defined. We instead directly compare the threshold surplus to the mechanism's revenue. Note, however, that this new definition still only depends on the rules of the mechanism, and not on incentives.

Definition 51. *The generalized competitive efficiency of a mechanism* M *for prior* $\mathcal{F}$ *is the smallest* μ *such that, for any strategy profile* $\mathbf{s}$ *and any feasible allocation function* $\mathbf{y}$ *mapping value profiles to* $\mathcal{X}$*, the revenue is at least a* μ *fraction of the threshold surplus:*

$$\textsc{Rev}(M, \mathbf{s}, \mathcal{F}) \geq \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[\mathcal{T}_i(y_i(v_i) \mid v_i) \right].$$

For winner-pays-bid mechanisms, the price per unit of a bid is the bid itself. The mechanism's revenue is exactly its bid surplus. Hence:

Lemma 52. *Let* M *be a winner-pays-bid mechanism and* $\mathcal{F}$ *a prior. If* M *has competitive efficiency* μ *on* $\mathcal{F}$*, then* M *also has generalized competitive efficiency* μ *on* $\mathcal{F}$*.*

A single-bid mechanism's competitive efficiency was the same for all priors, including those with correlation. However, simultaneously-composed auctions are known to possess inefficient equilibria under correlation (Feldman et al., 2013). To obtain a definition which will extend to simultaneous mechanisms, we weaken competitive efficiency in two ways. First, we only consider priors which are product distributions. Second, we restrict to alternate allocations $\mathbf{y}$ which do not depend on the value profile. Formally:

Definition 53. *A mechanism* M *has* weak competitive efficiency *if for any product distribution* $\mathcal{F}$*, any strategy profile* $\mathbf{s}$ *and any feasible allocation* $\mathbf{y} \in \mathcal{X}$*, the revenue is at least a* μ *fraction of the threshold surplus:*

$$\textsc{Rev}(M, \mathbf{s}, \mathcal{F}) \geq \mu \sum \mathcal{T}_i(y_i).$$

Note two notational simplifications for the threshold surplus: first, we omit conditioning from $\mathcal{T}_i(y_i)$, due to the independence of the prior. Second, the deterministic feasible allocation $\mathbf{y}$ allows us to further omit expectation over v_i. The following is then immediate from definitions.

Lemma 54. *Let* M *be a winner-pays-bid mechanism and* $\mathcal{F}$ *a product distribution. If* M *has generalized competitive efficiency* μ *on* $\mathcal{F}$*, then* M *also has weak competitive efficiency* μ *on* $\mathcal{F}$*.*

Next, we generalize individual efficiency. Section 4 built up a modular definition which depended on the type of single-agent mechanism an agent faced and the behavioral assumptions governing the agent's actions. For our application, it will suffice to consider unrestricted single-agent mechanisms under exact best response. That is, we assume our agent faces an arbitrary allocation rule $\tilde{x}$ and payment rule $\tilde{p}$, and chooses their action to maximize $\tilde{u}(a) = v\tilde{x}(a) - \tilde{p}(a)$.

We may extend our single-bid analysis in the following way. Consider the Pareto frontier between allocation and price per unit: $\hat{x}(b) = \sup_{a\,:\,\beta(a)\leq b} \tilde{x}(a)$. Best responding to $\tilde{x}$, the agent will only select actions a inducing (equivalent bid, allocation) pairs (b, x) if $x = \hat{x}(b)$. Note that such b would be a best response bid under $\hat{x}$, treated as a winner-pays-bid single-bid mechanism. Now for any minimum price per unit r^*, consider the threshold surplus T^{r^*} defined according to Definition 3 with respect to $\hat{x}$ and the generalized threshold surplus $\mathcal{T}^{r^*}$ with respect to $\tilde{x}$. For any $z \in [0, 1]$, $T^{r^*}(z) = \mathcal{T}^{r^*}(z)$. Using these observations, we may apply Lemmas 35 and 41 to conclude:

Lemma 55 (Generalized Individual Efficiency)**.** *Let $\tilde{x} : A \to [0, 1]$ and $\tilde{p} : A \to \mathbb{R}_+$ be arbitrary single-agent allocation and payment rules mapping actions to allocation levels and payments. Further let $a : \mathbb{R}_+ \to A$ be a function mapping agent values to best response actions. Then for every value v and every $z \in [0, 1]$,*

$$\tilde{u}(a(v)) + \mathcal{T}(z) \geq \tfrac{e-1}{e} vz. \tag{11}$$

Lemma 56. *Let F be a regular value distribution with monopoly reserve r^*. Further let $\tilde{x} : A \to [0, 1]$ and $\tilde{p} : A \to \mathbb{R}_+$ be arbitrary single-agent allocation and payment rules mapping actions to allocation levels and payments, and assume that for any action a with $\beta(a) \leq r^*$, $\tilde{x}(a) = 0$. Then for every v, best response $a(v)$ and $z \in [0, 1]$,*

$$\phi(v)\tilde{x}(a(v)) + \mathcal{T}^{r^*}(z) \geq \tfrac{e-1}{e} vz. \tag{12}$$

To derive welfare and revenue guarantees from a mechanism's weak competitive efficiency aggregate inequalities (11) and (12) over all agents i and all type profiles $\mathbf{v}$, taking $z = x_i^*(\mathbf{v})$ to be the optimal allocation in the welfare- or revenue-optimal mechanism. This implies:

Theorem 57. *Let M be a mechanism with weak competitive efficiency $\mu \leq 1$. Then in any Bayes-Nash equilibrium of M with independently distributed values, the expected welfare is a $\frac{e-1}{e}\mu$-approximation to that of the optimal mechanism.*

Theorem 58. *Let M be a mechanism for a downward-closed setting with weak competitive efficiency $\mu \leq 1$ which excludes agents below the monopoly reserves $\mathbf{r}^*$. Then in any Bayes-Nash equilibrium of M with independently distributed values, the expected revenue is a $\frac{e-1}{2e}\mu$-approximation to that of the optimal mechanism.*

5.2 Simultaneous Composition

We now apply our generalized framework and study the impact of local mechanisms' properties on the equilibrium performance of a larger system. Consider n single-parameter agents seeking abstract service. Agents may participate in one or more of m separate mechanisms. Each of the mechanisms are run simultaneously; each agent submits a profile of actions, one per mechanism. If the agent wins in any mechanism, they are considered served. They make the assigned payments to all mechanisms. In this section we show that under these assumptions, if each individual mechanisms has weak competitive efficiency μ, then so too does the aggregate mechanism induced for agents by simultaneous participation as described above. We refer to the aggregate mechanism as the *simultaneous composition* of the individual component mechanisms.

Formally, a simultaneous composition of mechanisms consists of m separate feasibility environments $\mathcal{X}^1, \ldots, \mathcal{X}^m$, one per component mechanism. Each component mechanism M^j is comprised of a bid allocation rule $\tilde{\mathbf{x}}^j$ and a bid payment rule $\tilde{\mathbf{p}}^j$, mapping a profile of actions $\mathbf{a}^j$ to an allocation in $\mathcal{X}^j$ and a nonnegative payment, respectively. We assume each mechanism has a withdraw action $\perp$ which guarantees zero allocation and payments in M^j. Define the simultaneous composition of mechanisms $M^1, \ldots, M^m$ in the following way:

Definition 59. *Let mechanisms $M^1, \ldots, M^m$ have bid allocation and bid payment rules $(\tilde{\mathbf{x}}^j, \tilde{\mathbf{p}}^j)$ and individual action spaces spaces $A_i^1, \ldots, A_i^m$ for each agent i. The* simultaneous composition *of $M^1, \ldots, M^m$ is defined to have:*

- *Action space $\prod_j A_i^j$ for each agent. That is, each agent participates in the global mechanism by participating in each composed mechanism individually. Given a profile of actions $\mathbf{a}$ for each agent, we denote by $\mathbf{a}^j$ the profile of actions restricted to mechanism j.*
- *Allocation rule $\tilde{x}_i(\mathbf{a}) = \max_j \tilde{x}_i^j(\mathbf{a}^j)$. That is, each agent is served at their highest level across all component mechanisms.*
- *Payment rule $\tilde{p}_i(\mathbf{a}) = \sum_j \tilde{p}_i^j(\mathbf{a}^j)$. That is, agents make payments to every composed mechanism.*

Note that as a consequence of Definition 59, we may define the composed feasibility environment as the set of allocation levels induced by the component mechanisms, i.e.

$$\mathcal{X} = \{(\max_j x_1^j, \ldots, \max_j x_n^j) \,|\, \mathbf{x}^1, \ldots, \mathbf{x}^m \in \mathcal{X}^1, \ldots, \mathcal{X}^m\}.$$

Having defined the composed mechanism and feasibility environment, we may state the main result of this section.

Theorem 60. *Let M be a simultaneous composition of mechanisms $M^1, \ldots M^m$. If $M^1, \ldots, M^m$ all have weak competitive efficiency μ with minimum per-unit prices $\mathbf{r}$, then so does M.*

The theorem will follow from two observations. First, agents' threshold prices are lower in the composition of mechanisms than in any individual component mechanism. In other words, it is easier for agents to secure allocation with more mechanisms to participate in. Second, the aggregate revenue is the sum of the component mechanisms' revenues.

To formalize the first observation, let $\mathbf{a}$ be a strategy profile under a degenerate prior, i.e. a profile of actions in the composed mechanism. We may treat each component mechanism M^j in isolation, and define $\mathbf{a}^j$ to be the profile of actions taken by agents in M^j. With respect to $\mathbf{a}^j$, we may define for any action a^j the payment $\tilde{p}_i^j(a^j)$, allocation $\tilde{x}_i^j(a^j)$, and price per unit $\beta_i^j(a^j)$. For any allocation $x \in [0, 1]$, we may also define agent i's threshold price $\tau_i^j(x)$, threshold price with discounted reserve $\tau_i^{j,r_i}(x)$ and generalized threshold surplus with discounted reserve $\mathcal{T}_i^{j,r_i}(x)$ for mechanism M^j according to Definitions 48-50. We omit the superscript j when discussing these quantities in the context of the composed mechanism. In these terms, our first observation can be stated as:

Lemma 61. *For any profile of actions $\mathbf{a}$, any $z \in [0, 1]$, any component mechanism M^j, and any minimum price per unit r_i, $\mathcal{T}_i^{r_i}(z) \leq \mathcal{T}_i^{j,r_i}(z)$.*

Proof. The interim threshold prices for the composed mechanism and M^j are defined as $\tau_i(x) = \inf_{a:\tilde{x}_i(a) \geq x} \beta_i(a)$ and $\tau_i^j(x) = \inf_{a^j:\tilde{x}_i^j(a^j) \geq x} \beta_i^j(a^j)$, respectively, where a denotes a profile of actions for agent i in all m mechanisms. For any action a^j in M^j, there exists an action a' in the composed mechanism where agent i plays a^j in M^j and $\perp$ in all other mechanisms. This action has price per unit $\beta(a') = \beta_i^j(a^j)$. This implies $\tau_i(x) \leq \tau_i^j(x)$. For any minimum price per unit r_i, we then obtain $\tau_i^{r_i}(x) \leq \tau_i^{j,r_i}(x)$ and $\mathcal{T}_i^{r_i}(z) \leq \mathcal{T}_i^{j,r_i}(z)$ from the definitions. □

Proof of Theorem 60. Let $\mathbf{a}$ be an action profile in the composed mechanism, and $\mathbf{y}$ a feasible allocation in the composed environment. Let $\mathbf{a}^j$ denote the restriction of $\mathbf{a}$ to actions in component mechanism j, and let $\mathbf{y}^1 \ldots, \mathbf{y}^m \in \mathcal{X}^1, \ldots, \mathcal{X}^m$ denote a profile of allocations for each component mechanism certifying the feasibility of $\mathbf{y}$, i.e. $y_i = \max_j y_i^j$ for all i. For each agent i, let δ_{ij} be an indicator taking value 1 if j is the lowest index such that $y_i = y_i^j$, and 0 otherwise. We obtain the following sequence of inequalities, explained after their statement:

$$\begin{aligned}
\tfrac{1}{\mu}\textsc{Rev}(M, \mathbf{a}) &= \sum\nolimits_j \tfrac{1}{\mu}\textsc{Rev}(M^j, \mathbf{a}^j) \\
&\geq \sum\nolimits_j \sum\nolimits_i \mathcal{T}_i^{j,r_i}(y_i^j) \\
&\geq \sum\nolimits_j \sum\nolimits_i \mathcal{T}_i^{r_i}(y_i^j) \\
&\geq \sum\nolimits_j \sum\nolimits_i \delta_{ij} \mathcal{T}_i^{r_i}(y_i) \\
&= \sum\nolimits_i \mathcal{T}_i^{r_i}(y_i).
\end{aligned}$$

The first inequality follows from the definition of payments in a composed mechanism as the sum of the revenues of the component mechanisms. The second line comes from the assumption that each component mechanism has generalized competitive efficiency μ with minimum per-unit prices $\mathbf{r}$. The third line follows from Lemma 61. The remaining lines follow from the definition of feasibility in the composed mechanism. □

Theorem 60 and the results of Sections 3 and 4 imply that the simultaneous composition of winner-pays-bid mechanisms inherits the worst-case welfare properties of the component mechanisms. Moreover, if all mechanisms have monopoly reserves, the same can be said for revenue.

References

Bergemann, D., Brooks, B., and Morris, S. (2017). First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143.

Bergemann, D., Brooks, B., and Morris, S. (2019). Revenue guarantee equivalence. *American Economic Review*, 109(5):1911–29.

Brooks, B. and Du, S. (2021). Optimal auction design with common values: An informationally robust approach. *Econometrica*, 89(3):1313–1360.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488.

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

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.

Dütting, P. and Kesselheim, T. (2015). Algorithms against anarchy: Understanding non-truthful mechanisms. In *16th ACM Conference on Economics and Computation*.

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

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

Feldman, M., Fu, H., Gravin, N., and Lucier, B. (2013). Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 201–210.

Hartline, J., Hoy, D., and Taggart, S. (2014). Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 693–710.

Hartline, J. D. (2013). Mechanism design and approximation. *Book draft. October*, 122:1.

Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234.

Hoy, D., Nekipelov, D., and Syrgkanis, V. (2015). Robust data-driven efficiency guarantees in auctions. In *EC Workshop on Algorithmic Game Theory and Data Science*.

Hoy, D., Taggart, S., and Wang, Z. (2018). A tighter welfare guarantee for first-price auctions. In *STOC'18—Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 132–137. ACM, New York.

Kaplan, T. R. and Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, 50(2):269–302.

Lehmann, D., Oćallaghan, L. I., and Shoham, Y. (2002). Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)*, 49(5):577–602.

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

Roughgarden, T. (2009). Intrinsic robustness of the price of anarchy. In *ACM Symposium on Theory of Computing*, pages 513–522.

Roughgarden, T., Talgam-Cohen, I., and Yan, Q. (2012). Supply-limiting mechanisms. In *ACM Conference on Electronic Commerce*, pages 844–861.

Syrgkanis, V. (2014). *Efficiency of mechanisms in complex markets*. PhD thesis, Cornell University.

Syrgkanis, V. and Tardos, E. (2013). Composable and efficient mechanisms. In *ACM Symposium on Theory of Computing*, pages 211–220.

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

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