Robust Analysis of Auction Equilibria*

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

[1]Northwestern University
[2]Oberlin College

August 3, 2025

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 auction equilibria 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.

1 Introduction

This paper takes a robust approach to the analysis of auctions. We consider common non-truthful mechanisms, such as first-price and all-pay auctions. Such mechanisms

*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.

are not typically welfare- or revenue-optimal, and the degree of suboptimality varies with the structure of equilibrium, which in turn depends delicately on the informational and behavioral assumptions placed on agents. To measure equilibrium performance we study the welfare or revenue, as a fraction of the optimal welfare or revenue. Placing minimal assumptions on agent behavior, we define a mechanism's *robust efficiency* to be the fraction of optimal performance it attains in the worst case over all equilibria. Quantifying over all equilibria avoids the analytical difficulties of solving for equilibrium explicitly; namely, the direct analysis approach can be intractable in asymmetric settings, and may not generalize well. Measuring efficiency as a fraction of optimal performance yields bounds that are invariant for environments at different scales.

The robust perspective requires us to identify common sources of inefficiency across environments. Classical economic analyses identify two main drivers of inefficiency: (a) externalities or inter-agent complementarities and (b) incomplete information. The analysis of this paper decomposes the performance of auctions into two terms: *competitive efficiency* quantifies the degree to which inter-agent complementarities cause losses in performance and *individual efficiency* quantifies the degree to which incomplete information causes losses in performance. A mechanism's robust welfare and revenue can be bounded by a combination of these two quantities. Moreover, this guarantee is tight for several natural environments, yielding an exact characterization of robust performance.

The analysis of the paper further characterizes the relative contribution of poor competitive and individual efficiency to performance loss. Competitive efficiency, as we define it, is a central determinant of an auction's performance. Individual efficiency is a secondary concern. 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 it varies with agents' degree of best response or the payment format; e.g. mechanisms with the winner-pays-bid payment format are more individually efficient than those with the all-pay payment format. Applying the philosophy of approximation (Hartline, 2012, 2013) to these two concerns: inter-agent complementarities are a critical feature to

be treated carefully in mechanism design while incomplete information is more of a detail.

We will illustrate competitive efficiency and individual efficiency through the example of the first-price auction, where the framework is simple to describe. Subsequent developments will then generalize to much richer environments and auction formats. To illustrate competitive efficiency, 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 welfare if each agent's value was equal to their threshold. Quantity (b) measures 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 numbers closer to 1 are more efficient. For the first-price auction this ratio is exactly $\mu = 1$. The revenue is the highest bid and the thresholds of all losers are equal to the highest bid. Thus, the two terms, (b) the optimal welfare from thresholds and (a) the revenue from bids, are the same.

Even mechanisms that are competitively efficient may lose efficiency due to the inability of agents to precisely respond to thresholds, i.e., due to incomplete information. Returning to the example of first price auctions, with deterministic concepts like pure Nash equilibrium, an agent's threshold is deterministic, and the agent can respond efficiently. On 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 varies randomly and the agent must place a single bid. Meanwhile, the threshold represents the seller's outside option, generated by other agents in the auction. An agent's inability to respond precisely thus results in performance loss from misallocation. We quantify this loss as the *individual efficiency* η, defined subsequently, which depends on how agents respond to their bidding problem and varies with the equilibrium notion (e.g., pure Nash or Bayes-Nash), degree of best response (e.g. exact or approximate), and payment format (e.g., winner-pays-bid or all-pay).

This paper's main results bound robust performance in terms of the individual efficiencies of the agents and the 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 individual efficiency of the agents' best response problems. 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.) Turning to revenue analysis, the main result 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. Below, we summarize the paper's specific individual and competitive efficiency analyses, starting with individual efficiency.

Individual efficiency is a property of the agents' bidding decisions that takes into account the incompleteness of information and structure of the agents' best response problems. Fixing the mechanism, it is a property of the distribution of single-agent allocation rules an agent faces, the agent's chosen bids, and the correlation between these two. For winner-pays-bid mechanisms, we analyze full- and incomplete-information bidding problems under exact and approximate best response. We also consider all-pay mechanisms under incomplete information. In more detail, we show:

- $\eta = 1$ when agents best respond to a fixed, threshold allocation rule with the winner-pays-bid format. This result is applicable to pure Nash equilibria in deterministic auctions for allocating indivisible goods, such as the first-price auction, when such equilibria exist.

- $\eta = 1 - 1/e \approx 0.63$ when agents best respond to distributions of general allocation rules with the winner-pays-bid format. This result is applicable to Bayes-Nash equilibria and, more generally, Bayes correlated equilibria with arbitrary information structures.

- $\eta = 1/2$ when agents best respond to distributions of general allocation rules

with the all-pay format. This bound only holds with respect to a weakened notion of individual efficiency.

- The above results extend to $(1-\epsilon)$-best response with a multiplicative loss of $(1-\epsilon)$. For example, we obtain $\eta = (1-\epsilon)$ for fixed threshold rules of winner-pays-bid mechanisms. This addresses the issue that exact, pure Nash equilibria tend not to exist in mechanisms such as the first price auction.

These results quantify the role of incomplete information in the loss of efficiency. When information is complete (in payoffs, actions of opponents, and rules of the mechanism), best-response bids are fully efficient. On the other hand, with incomplete information the agent's best-response bids can be inefficient, but that inefficiency is bounded by a constant. As mentioned above, bidding problems with the winner-pays-bid format mitigate informational inefficiencies more than those with the all-pay payment format.

Competitive efficiency is a property of the rules of the auction that map bids to winners and payments. Up to variation due to individual efficiency, discussed above, the robust efficiency of a mechanism is approximately governed by its competitive efficiency. Thus analysis of robust welfare and revenue of an auction approximately reduces to analysis of its competitive efficiency. To aid in the analysis of competitive efficiency we develop a number of closure properties.

- Competitive efficiency in winner-pays-bid mechanisms is closed under reserve pricing, i.e., the competitive efficiency of a 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, agent selection (captured by the prior), and correlated signaling, i.e., the competitive efficiency of any randomized mechanism on any randomized set of agents and with any information structure is at least the competitive efficiency of the worst mechanism in the combination, set of agents, and information structure. As a consequence, the worst-case competitive efficiency is attained at degenerate randomizations over mechanisms and degenerate information structures. To characterize the competitive efficiency of such families of randomizations, it

consequently suffices to analyze the competitive efficiency of simple mechanisms under full information.

- Competitive efficiency in all-pay mechanisms with deterministic allocation rules is at least that in the winner-pays-bid mechanisms with the same allocation rules.

- Competitive efficiency is closed under simultaneous composition (when bids are independently distributed), i.e., the competitive efficiency of a composite mechanism where a set of mechanisms are run in parallel is at least 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 priors. Such analyses are fairly straightforward compared to more technically involved analyses of stochastic equilibrium concepts.

The following winner-pays-bid mechanisms have competitive efficiency $\mu = 1$. Intuitively, agents in these environments are one-for-one substitutes:

- 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 top k position weights are all 1.) The position auction model was popularized by Varian (2007) and Edelman et al. (2007) who study 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 a 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 known subset of items. Bids are submitted for each item, and the highest bidder for each item wins it. (For this result, values must be independently distributed, and the information structure must be independent of values.)

These results all follow from a single analysis of the competitive efficiency of 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. 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, lower bounds its competitive efficiency. The results listed above follow because the greedy algorithm is non-strategically efficient for these environments. Notice that the k-unit highest-bids-win allocation rule is given by the greedy-by-bid algorithm. The highest-bids-win rule for (single-bid) highest-bids-win matching markets can also be implemented greedily. 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 low when externalities between agents are 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 known subsets of items. Feasible allocations serve agents with non-overlapping demand sets. 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 is $1/\sqrt{m}$ 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 the 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 their equilibria, even under the simplifying assumption that the agents' values are uniformly distributed (but asymmetrically with different supports), are very difficult to derive (Kaplan and Zamir, 2012).

An approach from computer science for understanding the potentially complex equilibria of auctions 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, the present 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 this paper's decomposition of robust efficiency into competitive efficiency and individual efficiency. First,

[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 their competitive efficiency. See Hartline and Lucier (2010) and Akbarpour et al. (2021) for analyses of truthful mechanisms based on greedy rules.

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.

Recently, the economics literature has also seen a number of robust treatments of mechanisms. The main distinction between this literature and the present paper is that this literature focuses on absolute performance whereas the present paper (and the computer science literature discussed previously) focuses on robust performance relative to the optimal mechanism. This difference is significant as absolute robust analysis generally focuses on the settings were the optimal performance is the lowest. On the other hand, analysis relative to the optimal performance requires that the performance is close to what is possible in the non-robust problem, and is invariant to the scale. The relative framework should be preferred when the range of the optimal performance varies significantly or robustness is required at all scales.

The robust analyses in economics identify mechanisms with good worse-case revenue when there is ambiguity with respect to key aspects of the models. For example, Carroll (2017) considers a revenue maximizing seller with multiple items and a buyer with values distributed according to known marginals. The correlation structure is ambiguous. He shows that the max-min mechanism is a linear pricing. Brooks and Du (2021) consider the design of revenue-optimal common-value auctions that are robust across information structures.

The most related works from the robust mechanism design literature give lower bounds on the performance of standard mechanisms under permissive assumptions on the beliefs of agents. Bergemann et al. (2017) consider a robust analysis of the revenue in a first-price auction with respect to the information structure 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 loose. Bergemann et al. (2019) give worst-case, over information structure, revenue rankings of standard auction formats. In contrast, the bounds of this paper are relative to the optimal performance and apply broadly

beyond single-item auctions.

Another strand of literature derives 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 for several settings 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.

In the computer science literature a canonical robust mechanism design problem is that of maximizing the ratio of the revenue from a single-item auction to the optimal auction in the worst case over valuation distributions of the agents. This problem was posed by Dhangwatnotai et al. (2015), further considered by Fu et al. (2015) and Allouah and Besbes (2020), and optimally solved by Hartline et al. (2020). These papers aim to design optimal incentive compatible mechanisms. In contrast, the present paper gives robust analyses of standard mechanisms (which are not incentive compatible).

2 Framework Overview

This paper develops a framework for analyzing auction equilibria that is robust to variation in agents' priors, information structures, and degree of best response. Before developing the framework in full generality, this section presents a simplified version. For the simplified version, we restrict attention to deterministic mechanisms for 0-1 allocation environments under full information and pure strategies. Under these assumptions, we define the framework's two components, competitive efficiency and individual efficiency, and show how they together imply welfare guarantees. We also illustrate the framework with three examples.

The simplified framework focuses on single-bid mechanisms: each of n agents i

places a real-valued bid b_i. Given the profile of bids $\mathbf{b} = (b_1, \ldots, b_n)$, the mechanism selects an allocation $\tilde{\mathbf{x}}(\mathbf{b}) = (\tilde{x}_1(\mathbf{b}), \ldots, \tilde{x}_n(\mathbf{b}))$ and nonnegative real-valued payments $\tilde{\mathbf{p}}(\mathbf{b}) = (\tilde{p}_1(\mathbf{b}), \ldots, \tilde{p}_n(\mathbf{b}))$. Agent i has a value v_i, and gets linear utility $\tilde{u}_i(\mathbf{b}) = v_i\tilde{x}_i(\mathbf{b}) - \tilde{p}_i(\mathbf{b})$. For the example of the single-item first-price auction, $\tilde{x}_i(\mathbf{b}) = 1$ if i is the highest bidder (excepting ties) and otherwise 0, and $\tilde{p}_i(\mathbf{b}) = b_i\tilde{x}_i(\mathbf{b})$. Denote by $\mathcal{X}$ the set of feasible allocation vectors: for a single-item auction, $\mathcal{X}$ is the set of vectors in $\{0,1\}^n$ with at most a single 1. We focus in this section on $\mathcal{X} \subseteq \{0,1\}^n$, and *deterministic mechanisms*, i.e. we do not allow nontrivial randomization of $\tilde{\mathbf{x}}(\mathbf{b})$ over $\mathcal{X}$. The social welfare is given by $\sum_i v_i\tilde{x}_i(\mathbf{b})$, which we compare to the welfare of the optimal allocation $\mathbf{x}^* \in \arg\max_{\mathbf{x}\in\mathcal{X}} \sum_i v_i x_i$, denoted by $\textsc{Opt}(\mathbf{v})$.

The two components of the framework are *competitive efficiency* and *individual efficiency*. To define both components, we first establish a measure of agents' competition. In a deterministic single-bid mechanism, we measure competition by considering each agent's threshold bid, defined as the infimum bid they could place and still win, denoted $\hat{b}_i(\mathbf{b}_{-i}) = \inf\{b \,|\, \tilde{x}_i(b, \mathbf{b}_{-i}) = 1\}$. A high threshold implies that allocation is costly, and hence competition is strong. Competitive efficiency then quantifies the extent to which aggregate competition between agents translates into revenue for the auctioneer. We measure aggregate competition by the maximum total threshold of any feasible allocation, $\textsc{Opt}(\hat{\mathbf{b}}) = \max_{\mathbf{y}\in\mathcal{X}} \sum_i \hat{b}_i(\mathbf{b}_{-i})y_i$, and compare this to the total payments $\textsc{Rev}(M, \mathbf{b}) = \sum_i \tilde{p}_i(\mathbf{b})$ to obtain a ratio $\textsc{Rev}(M, \mathbf{b})/\textsc{Opt}(\hat{\mathbf{b}})$, which is at most 1 for the mechanisms we consider. We take this ratio in the worst case over bid profiles to obtain the competitive efficiency. The resulting quantity does not depend on the informational or behavioral assumptions placed on the agents. It can be computed directly from the rules of the auction.

Definition 2.1. *The competitive efficiency of a deterministic single-bid mechanism M for a* 0-1 *allocation environment is the worst-case ratio of revenue to optimal threshold surplus over any bids $\mathbf{b}$ (implicitly inducing threshold bids $\hat{\mathbf{b}}$), i.e.:*

$$\mu = \min_{\mathbf{b}} \textsc{Rev}(M, \mathbf{b})/\textsc{Opt}(\hat{\mathbf{b}}).$$

Individual efficiency, meanwhile, measures how efficiently each agent responds to their competition. The efficiency of an agent's response depends on the agent's information and how closely they maximize their utility. It also depends on the way

they are able to trade off allocation and payment, which is dictated by the auction's payment format: best responses to the same set of bids will be different in an all-pay versus a first-price auction. We abstract these factors away from the inter-agent competition of the mechanism by considering single-agent responses to single-agent bidding problems. The restricted definition below suffices for the full-information setting of this overview.

Definition 2.2. *A* single-agent signal-free bidding outcome $\mathcal{O} = (v, \tilde{x}, \tilde{p}, b)$ *is given by a value v, a bid b, and a single-agent mechanism with allocation rule $\tilde{x}$ (mapping bids to allocation levels in $[0, 1]$) and payment rule $\tilde{p}$ (mapping bids to nonnegative payments). The agent's utility for any bid $b^\dagger$ is $\tilde{u}(b^\dagger) = v\tilde{x}(b^\dagger) - \tilde{p}(b^\dagger)$. The agent's utility for their chosen bid is denoted $u(\mathcal{O}) = \tilde{u}(b)$.*

In full-information equilibrium, each agent's perspective projects down to a single-agent bidding outcome. With full information and pure strategies, bids are deterministic. If tiebreaking is additionally deterministic, then the corresponding single-agent allocation rule $\tilde{x}$ has values in $\{0, 1\}$, and the agent's threshold bid $\hat{b}$ again measures the strength of their competition. For a single-agent bidding outcome to be efficient, it must be that whenever the agent's competition is weak, they obtain high utility. Whenever their utility is low, competition must be strong. The following definition makes this tradeoff precise for the simplified setting.

Definition 2.3. *The* weak individual efficiency *of a single-agent signal-free bidding outcome $\mathcal{O}$ with allocations in $\{0, 1\}$ is the ratio $(u(\mathcal{O}) + \hat{b})/v$. The weak individual efficiency of a family of single-agent bidding outcomes is the smallest individual efficiency of any outcome in the family.*

In Section 6.1 we present a stronger version of individual efficiency which generalizes beyond the simple settings of this overview. Before using competitive efficiency and individual efficiency for welfare analysis, we instantiate each for an example.

Example 2.4 (Single-Item First-Price Auction, Pure Nash Equilibrium)**.** In a single-agent first-price auction, each agent submits a bid b_i. The single item is awarded to the highest bidder, who pays their bid. We consider full-information pure-strategy

equilibria, where each agent i has a commonly-known value v_i. Tiebreaking is deterministic.[2]

Lemma 2.5. *The competitive efficiency of the single-item first-price auction is* 1.

Proof. For first-price auction M on bids $\mathbf{b}$, the following equalities prove the lemma:

$$\textsc{Opt}(\hat{\mathbf{b}}) = \max_{\mathbf{y}\in\mathcal{X}} \sum_i \hat{b}_i(\mathbf{b}_{-i}) y_i = \max_i \hat{b}_i(\mathbf{b}_{-i}) = \max_i \max_{j\neq i} b_j = \max_i b_i = \textsc{Rev}(M, \mathbf{b}).$$

□

We next analyze individual efficiency. With deterministic tiebreaking, the induced single-agent allocation rules step up from 0 to 1 at $\hat{b}$, taking either value at exactly $\hat{b}$. In pure Nash equilibrium, each agent best responds to their threshold bid, and the induced bidding outcomes have full individual efficiency.

Lemma 2.6. *For single-agent winner-pays-bid outcomes where $\tilde{x}$ takes values in $\{0, 1\}$ and b is a best response, the weak individual efficiency is* 1.

Proof. By assumption, b maximizes the agent's utility. If $\tilde{x}(\hat{b}) = 1$, then the lemma holds immediately, as $u(\mathcal{O}) = \max_{b^\dagger} \tilde{u}(b^\dagger) \geq \tilde{u}(\hat{b}) = v - \hat{b}$. More generally, for all $\tilde{x}(\hat{b}) \in \{0, 1\}$, for any $\delta > 0$ we have

$$u(\mathcal{O}) = \max_{b^\dagger} \tilde{u}(b^\dagger) \geq \tilde{u}(\hat{b} + \delta) = v - \hat{b} - \delta.$$

Taking $\delta \to 0$ again implies the stated individual efficiency guarantee. □

Having illustrated the components of the framework, we now combine them to obtain robust welfare guarantees. The welfare in equilibrium is the total value of the winners. Equivalently, it is the utility of the agents plus the utility of the auctioneer, in the form of revenue. Competitive efficiency and individual efficiency together give a tradeoff between these quantities. High individual efficiency implies that agent utilities are only low if competition is high. High competitive efficiency implies that competition is high only if revenue is high. The following theorem makes this precise.

[2]The existence of pure Nash equilibrium of the first-price auction depends on the choice of tiebreaking rule. The results of this section apply whenever equilibrium exists. For example, a sufficient condition for existence under lexicographic tiebreaking is that agent 1 has the highest value. Furthermore, Section 2.2 below demonstrates that the framework is robust to approximate best response and hence implies results for ϵ-Nash equilibria, which always exist in first-price auctions regardless of tiebreaking (Bich and Laraki, 2017).

Theorem 2.7. *Let $M = (\tilde{\mathbf{x}}, \tilde{\mathbf{p}})$ be a deterministic mechanism for a* 0-1 *allocation environment with competitive efficiency $\mu \leq 1$, and let $\mathbf{b}$ be an equilibrium for value profile $\mathbf{v}$. If for each agent i, the induced single-agent bidding outcome $(v_i, \tilde{x}_i(\cdot, \mathbf{b}_{-i}), \tilde{p}_i(\cdot, \mathbf{b}_{-i}), b_i)$ has weak individual efficiency at least η and nonnegative agent utility, then the welfare of M for equilibrium $\mathbf{b}$ is at least a $\mu\eta$-fraction of the optimal surplus:*

$$\sum\nolimits_i v_i \tilde{x}_i(\mathbf{b}) \geq \mu\eta \textsc{Opt}(\mathbf{v}). \tag{1}$$

Proof. Let $\mathbf{x}^* \in \arg\max_{\mathbf{x}\in\mathcal{X}} \sum_i v_i x_i$ be a 0-1 vector maximizing welfare. For any agent i with $x_i^* = 1$, weak individual efficiency implies that

$$\tilde{u}_i(\mathbf{b}) + \hat{b}_i(\mathbf{b}_{-i}) x_i^* \geq \eta v_i x_i^*. \tag{2}$$

Meanwhile for i with $x_i^* = 0$, inequality (2) holds by nonnegativity of agent utilities. Summing over agents and applying competitive efficiency, we obtain:

$$\sum\nolimits_i \tilde{u}_i(\mathbf{b}) + \tfrac{1}{\mu}\textsc{Rev}(M, \mathbf{b}) \geq \sum\nolimits_i \tilde{u}_i(\mathbf{b}) + \sum\nolimits_i \hat{b}_i(\mathbf{b}_{-i}) x_i^* \geq \eta \sum\nolimits_i v_i x_i^*.$$

Since the revenue and utility sum to the welfare $\sum_i v_i \tilde{x}_i(\mathbf{b})$, and since $\mu \leq 1$, we obtain the desired welfare guarantee (1). □

As an example, Theorem 2.7 implies the following well-known result.

Corollary 2.8. *In the single-item first-price auction with deterministic tiebreaking, pure Nash equilibria are efficient whenever they exist (e.g. with lexicographic tiebreaking and $v_1 = \max_i v_i$).*

We conclude the overview with two examples which illustrate the way competitive efficiency and individual efficiency vary across mechanisms and solution concepts.

2.1 Single-Minded Combinatorial Auctions

In multi-item environments, fixing a payment format such as winner-pays-bid, there may be many natural rules for selecting winners. The example below demonstrates how competitive efficiency allows for comparison of these different mechanisms with respect to robust welfare, up to the individual efficiency of the solution concept. We consider two winner-pays-bid mechanisms for the following multi-item environment with complementarities:

Definition 2.9. *A* single-minded combinatorial auction (SMCA) *feasibility environment is defined by m indivisible items, n agents that each desire a known 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 $\mathcal{S}_i \cap \mathcal{S}_{i'} = \emptyset$.*

One natural allocation rule is to maximize the bid surplus of the winners.

Definition 2.10. *The* highest-bids-win rule *allocates the feasible set of bidders with the highest total bid: $\tilde{\mathbf{x}}(\mathbf{b}) = \arg\max_{\mathbf{x}\in\mathcal{X}} \sum_i b_i x_i$.*

The following example demonstrates that the winner-pays-bid highest-bids-win mechanism has low competitive efficiency.

Lemma 2.11. *For any number m of items, there exists an m-item SMCA environment where the winner-pays-bid mechanism with the highest-bids-win allocation rule has competitive efficiency at most $1/m$.*

Proof. Consider an auction with $m+1$ agents. Let $\mathcal{S}_i = \{i\}$ for $i \in \{1,\dots,m\}$ and $\mathcal{S}_{m+1} = \{1,\dots,m\}$. Agents $m+1$ is mutually exclusive with any subset of the other agents. Consider the bid profile $(0,\dots,0,1)$. Agent $m+1$ wins the auction and pays their bid, yielding $\textsc{Rev}(M,\mathbf{b}) = 1$. Meanwhile the thresholds can be computed as $\hat{\mathbf{b}} = (1,\dots,1,0)$. The threshold surplus is maximized by choosing the feasible allocation $\mathbf{y} = (1,\dots,1,0)$, and hence $\textsc{Opt}(\hat{\mathbf{b}}) = m$. Consequently, the competitive efficiency is at most $1/m$. □

We see from the example of Lemma 2.11 that the highest-bids-win rule inefficiently converts competition into revenue. Agent $m+1$ simultaneously competes with many other agents, all of whom could be feasibly allocated, namely $1,\dots,m$. Meanwhile, this agent generates very little revenue for the mechanism. As the following lemma shows, we can modify this example to obtain an equilibrium with low welfare. To do so, we introduce a additional agent to incentivize the desired bidding behavior.

Lemma 2.12. *For any m, there exist pure Nash equilibria of an m-item winner-pays-bid highest-bids-win mechanism with welfare a $1/m$ factor of optimal.*

Proof. Consider an auction with $m+2$ agents. Let $\mathcal{S}_i = \{i\}$ for $i \in \{1, \ldots, m\}$ and $\mathcal{S}_{m+1} = \mathcal{S}_{m+2} = \{1, \ldots, m\}$. When all agents have value $v_i = 1$, and when ties are broken in favor of agent $m+1$, then the bid profile $(0, \ldots, 0, 1, 1)$ is a Nash equilibrium for the highest-bids-win mechanism. Agent $m+1$ wins all items, yielding welfare of 1. The optimal allocation serves agents $1, \ldots, m$ and has welfare m. □

Lemma 2.12 shows that the highest-bids-win mechanism possesses equilibria with very low welfare. This is despite the fact that if agents bid their values, the highest-bids-win mechanism produces optimal allocations; welfare is lost due to poor incentives. We next consider a mechanism which allocates using a greedy heuristic. This heuristic will often produce a suboptimal allocation if agents bid their values. However, the mechanism will lack the pathologies that led to low-welfare equilibria under highest-bids-win allocation. Specifically, we study a *greedy by priority rule:* define the priority of an agent i with bid b_i and desired bundle $\mathcal{S}_i$ to be $\psi_i(b_i) = b_i/\sqrt{|\mathcal{S}_i|}$. The greedy rule considers the agents in nonincreasing order of priority, and allocates them if their desired items are available. In Section 4.3, we analyze the winner-pays-bid mechanism with this rule and derive the following result.

Lemma 2.13. *The competitive efficiency of the winner-pays-bid mechanism for the greedy by priority rule with priority $\psi_i(b_i) = b_i/\sqrt{|\mathcal{S}_i|}$ is at least $1/\sqrt{m}$.*

If agents bid their values, the greedy by priority rule can be as much as a $\sqrt{m}$ factor away from optimal (Lehmann et al., 2002). Theorem 2.7 implies, then, that this welfare loss from suboptimal allocation is in fact the only source of robust performance loss in pure Nash equilibrium.

2.2 Approximate Best Response

Recall that competitive efficiency depends only on the rules of the mechanism, while individual efficiency quantifies the influence of agents' behavior and information on the welfare of equilibria. In particular, individual efficiency measures how efficiently agents solve their individual bidding problems. With deterministic bids and allocations, an agent can respond to their competition fully efficiently. Under approximate rather than exact best response, the efficiency of equilibria will degrade. The analysis

below quantifies the degradation of individual efficiency with the degree of best response, and bounds the corresponding degradation in equilibrium efficiency. We first define approximate best response for mechanisms and single-agent bidding outcomes.

Definition 2.14. *Given bid profile* $\mathbf{b}$*, agent* i*'s bid satisfies* $(1-\epsilon)$-best response *if for all alternative bids* $b_i^\dagger$*,* $\tilde{u}_i(\mathbf{b}) \geq (1-\epsilon)\tilde{u}_i(b_i^\dagger, \mathbf{b}_{-i})$*. Similarly, a single-agent bidding outcome* $\mathcal{O} = (v, \tilde{x}, \tilde{p}, b)$ *satisfies* $(1-\epsilon)$*-best response if* $u(\mathcal{O}) \geq (1-\epsilon)\tilde{u}(b^\dagger)$ *for all alternative bids* $b^\dagger$*.*

The following lemma shows that an agent's individual efficiency degrades linearly with their degree of best response.

Lemma 2.15. *For winner-pays-bid outcomes* $(v, \tilde{x}, \tilde{p}, b)$ *where* $\tilde{x}$ *takes values in* $\{0, 1\}$ *and* b *is a* $(1-\epsilon)$*-best response, the weak individual efficiency is at least* $1-\epsilon$*.*

Proof. The proof generalizes that of Lemma 2.6. By Definition 2.14, the chosen bid b obtains at least a $(1-\epsilon)$-fraction of that from bidding just above the threshold $\hat{b}$. Hence for any $\delta > 0$,

$$u(\mathcal{O}) + \hat{b} \geq (1-\epsilon)\tilde{u}(\hat{b}+\delta) + \hat{b} \geq (1-\epsilon)(\tilde{u}(\hat{b}+\delta) + \hat{b}) \geq (1-\epsilon)(v-\delta).$$

Taking $\delta \to 0$ implies the desired individual efficiency guarantee of $(1-\epsilon)$. □

Lemma 2.15 implies that the welfare of approximate Nash equilibria degrades at worst linearly with the degree of best response of agents. In particular, the bound on the welfare loss has no dependence on the number of agents. Moreover, note that approximate equilibria always exist, in contrast to exact equilibria.

Corollary 2.16. *Let* M *be a deterministic winner-pays-bid mechanism for a* 0-1 *allocation environment with competitive efficiency* $\mu \leq 1$*, and let* $\mathbf{b}$ *be an* ϵ*-Nash equilibrium for value profile* $\mathbf{v}$*. Then the welfare of* M *for* $\mathbf{b}$ *is at least a* $\mu(1-\epsilon)$*-fraction of optimal.*

Corollary 2.17. *The welfare of any* ϵ*-Nash equilibrium of the first-price auction with any deterministic tiebreaking is at least a* $(1-\epsilon)$*-fraction of optimal.*

We conclude the example by noting that we only varied agents' degree of best response, and held fixed the winner-pays-bid payment format and the assumption

of full information. Section 5 further considers the impact of incomplete information over values and bids, more general information structures, and changes to the payment format of the mechanism.

3 Preliminaries

The remainder of the paper extends the framework sketched in Section 2 to more general families of mechanisms and equilibria. Section 2 restricted to deterministic winner-pays-bid mechanisms under full information, and hence deterministic bids and values. This section gives notation necessary for mechanisms and equilibria beyond these restrictions.

A mechanism for single-parameter agents generally consists of action spaces A_i for each agent i (and joint action space $\mathbf{A} = \prod_i A_i$), 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. For mechanisms which randomize over allocations, we overload notation, letting $\tilde{x}_i(\mathbf{a})$ denote i's expected level of service. For indivisible goods, $\tilde{x}_i(\mathbf{a})$ therefore captures agent i's probability of allocation. Each agent i has a value v_i for service, and linear utility $\tilde{u}_i^{v_i}(\mathbf{a}) = v_i\tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})$. We omit the superscript v_i when it is clear from context.

We allow a broad range of solution concepts within the private values model,[3] and define an equilibrium to be a joint distribution over actions and values. Formally, let $\mathbf{V} = V_1 \times \ldots \times V_n \subseteq \mathbb{R}_+^n$ be the space of n-agent value profiles. An equilibrium is a joint distribution $\mathcal{G}$ over $\mathbf{V} \times \mathbf{A}$, which may be correlated across agents' actions, agents' values, or between actions and values. Different solution concepts will impose different best response and independence conditions on these distributions. We use $\mathcal{B}$ and $\mathcal{F}$ to denote the marginal distributions over the action profile $\mathbf{a}$ and value profile $\mathbf{v}$ respectively, with marginal distributions $\mathcal{B}_i$ and $\mathcal{F}_i$ for agent i. Throughout, we maintain the agent-normal form interpretation of the distributions: each agent i is

[3]In particular, values are not interdependent, but may be correlated unless explicitly stated otherwise.

selected from a population with distribution $\mathcal{F}_i$, and given selection $\mathbf{v}$ the value v_i agents in population i play according to $\mathcal{B}_i \mid \mathbf{v}$. Given mechanism M and equilibrium $\mathcal{G}$, the two objectives of interest are revenue, given by $\textsc{Rev}(M, \mathcal{G}) = \mathbb{E}_{\mathbf{a}}[\sum_i \tilde{p}_i(\mathbf{a})]$, and welfare, given by $\textsc{Welfare}(M, \mathcal{G}) = \mathbb{E}_{\mathbf{v},\mathbf{a}}[\sum_i \tilde{u}_i^{v_i}(\mathbf{a})] + \textsc{Rev}(M, \mathcal{G}) = \mathbb{E}_{\mathbf{v},\mathbf{a}}[\sum_i v_i \tilde{x}_i(\mathbf{a})]$.

Given a family of equilibria, we pursue robust, or worst-case, approximation analyses, and ask how far from optimal a mechanism can be, quantified over equilibria for that family. Formally, given value distribution $\mathcal{F}$, 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 value distributions $\mathbb{F}$ (e.g. point mass, product, or unrestricted), and family of equilibria $\text{EQ}_{\mathcal{F}}(M)$ for each $\mathcal{F} \in \mathbb{F}$ and mechanism M (e.g. ϵ-Nash, Bayes-Nash, Bayes coarse correlated), we study the *robust efficiency*

$$\min_{\mathcal{F}\in\mathbb{F}, \mathcal{G}\in\text{EQ}_{\mathcal{F}}(M)} \frac{\textsc{Welfare}(M, \mathcal{G})}{\textsc{Welfare}(\textsc{Opt}, \mathcal{F})}.$$

A high robust efficiency (close to 1) indicates that M is always nearly efficient, whereas a very small value suggests a pathology that might rule out M in practice.

For revenue, we assume independently-distributed values, and compare to the revenue-optimal Bayesian incentive compatible mechanism. Per Myerson (1981), this optimal mechanism can depend intricately on the prior $\mathcal{F}$ over values. 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 6), which depend on much less fine-grained information than the form of the Bayesian optimal mechanism. For these mechanisms, we study the ratio

$$\min_{\mathcal{F}\in\mathbb{F}, \mathcal{G}\in\text{EQ}_{\mathcal{F}}(M_{\mathcal{F}})} \frac{\textsc{Rev}(M_{\mathcal{F}}, \mathcal{G})}{\textsc{Rev}(\textsc{Opt}_{\mathcal{F}}, \mathcal{F})},$$

where $M_{\mathcal{F}}$ denotes mechanism M endowed with monopoly reserve prices for $\mathcal{F}$.

4 Competitive Efficiency

This section defines and analyzes competitive efficiency in full generality. In particular, we give a definition which extends beyond the simple settings of Section 2 to handle uncertainty in the auction environment (which may come from randomized allocations or random selection of participants in the auction, captured by the value

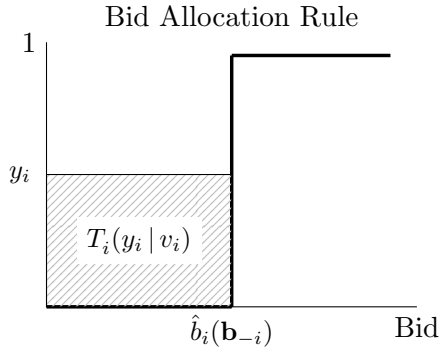
distribution) and reserve prices. Importantly, the general definition will retain the economic interpretation of Definition 2.1: competitive efficiency quantifies a mechanism's ability to translate competition into revenue, and can be measured directly from the rules of the auction without dependence on agents' behavior or beliefs. As in Section 2, we restrict attention to single-bid mechanisms, but our results will extend beyond the winner-pays-bid auctions considered in that section.

We first define the appropriate extension of threshold bids beyond the deterministic, full-information setting of Section 2. For a mechanism with allocation rule $\tilde{\mathbf{x}}$ and a joint distribution $\mathcal{G}$ over bids and values, we measure an agent i's competition conditioned on their value v_i. To this end, define the interim allocation rule for agent i as their probability of allocation for a bid b conditioned on v_i, given by $\tilde{x}_i(b \mid v_i) = \mathbb{E}_{\mathbf{b}_{-i} \mid v_i}[\tilde{x}_i(b, \mathbf{b}_{-i})]$. Define the inverse as $t_i(x \mid v_i) = \inf\{b \mid \tilde{x}_i(b \mid v_i) \geq x\}$. Definition 4.1, presented below with subsequent motivation, provides the desired analog to threshold bids.

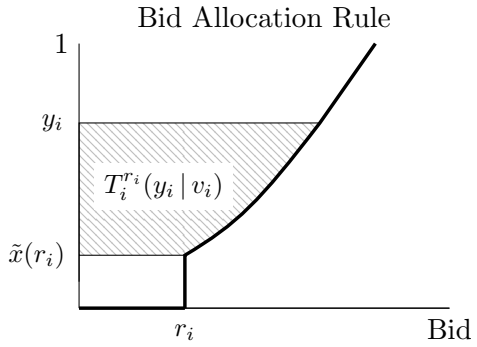
Definition 4.1. *For an (implicit) single-bid allocation rule $\tilde{\mathbf{x}}$ with (implicit) reserves $\mathbf{r}$ and joint distribution $\mathcal{G}$ over bids and values, the competition faced by agent i with value v_i for obtaining allocation y_i is summarized by the* threshold surplus with discounted reserve*, defined as $T_i(y_i \mid v_i) = \int_{\tilde{x}_i(r_i \mid v_i)}^{y_i} t_i(x \mid v_i)\, dx$. With the reserve price denoted explicitly, define $T_i^{r_i}(y_i \mid v_i)$. See Figure 1.*

The threshold surplus of Definition 4.1 generalizes the simpler notion of threshold bids in three senses. First, consider a mechanism with deterministic allocations $\tilde{\mathbf{x}}(\mathbf{b}) \in \{0,1\}^n$. Any distribution $\mathcal{G}$ over bids and values induces a distribution over threshold bids $\hat{b}_i(\mathbf{b}_{-i})$ for each agent i, and the interim allocation rule $\tilde{x}_i(\cdot \mid v_i)$ is the cumulative distribution function of $\hat{b}_i(\mathbf{b}_{-i})$ conditioned on v_i. Then the threshold surplus for $y_i = 1$ is exactly the expected threshold bid for agent i: $T_i(1 \mid v_i) = \int_0^1 t_i(x)\, dx = \mathbb{E}_{\mathbf{b}_{-i} \mid v_i}[\hat{b}_i(\mathbf{b}_{-i})]$. Second, note that the threshold surplus $T_i(y_i \mid v_i)$ is well-defined even for mechanisms with randomized allocation rules, as well as for fractional allocation levels $y_i \in [0,1]$, which will be important for environments where feasible allocations are necessarily randomized. Finally, note that Definition 4.1 discards allocation levels where the reserve price binds: intuitively, barriers to allocation from reserve prices do not correspond to competition between agents, which is the quantity of interest.

Next, we generalize competitive efficiency. Competitive efficiency compares rev-



(a) Threshold surplus for a deterministic allocation rule with deterministic bid profile $\mathbf{b}$. Agent i with value v_i has allocation which steps from 0 to 1 at $\hat{b}_i(\mathbf{b}_{-i})$. For desired allocation level y_i, the threshold surplus is $\hat{b}_i(\mathbf{b}_{-i})y_i$.



(b) Threshold surplus for a randomized allocation rule with a reserve price. Agent i with value v_i receives no allocation for bids below r_i. Threshold surplus is only counted for the fraction of allocations above $\tilde{x}(r_i)$.

Figure 1

enue to threshold surplus for feasible allocations. Definition 2.1 defined competitive efficiency in terms of deterministic allocations. With randomized values, we now consider randomized, *interim* allocations. Taking the view of the value distribution as defining populations of agents, an interim allocation defines a feasible allocation probability for each agent in each population. Formally, let $\mathbf{y}$ be an ex post allocation function, mapping value profiles to feasible allocations $\mathbf{y}(\mathbf{v}) \in \mathcal{X}$. The interim allocation for agent i with value v_i we denote by $y_i(v_i) = \mathbb{E}_{\mathbf{v}_{-i} \mid v_i}[y_i(\mathbf{v})]$. Overloading notation, a profile of interim allocation functions $\mathbf{y} = (y_1(\cdot), \ldots, y_n(\cdot))$ is *feasible* if it is induced by some feasible ex post allocation function (i.e. it is the *reduced form* of such an ex post allocation (c.f. Border, 2007)).

Definition 4.2. *Let $\mathcal{G}$ be a joint distribution over values and bids, with values following marginal distribution $\mathcal{F}$. The* competitive efficiency *of a single-bid mechanism for $\mathcal{G}$ is the largest μ such that for any feasible profile $\mathbf{y}$ of interim allocation functions, the expected revenue is at least a μ fraction of the* threshold surplus*:*

$$\textsc{Rev}(M, \mathcal{G}) \geq \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[T_i(y_i(v_i) \mid v_i) \right]. \tag{3}$$

The competitive efficiency for a family of mechanisms over a family of distributions is the smallest competitive efficiency of any mechanism in the first family over any

distribution in the second. We refer to a mechanism's competitive efficiency over all distributions simply as its competitive efficiency.

Definition 4.2 generalizes the simpler Definition 2.1 to randomized bid profiles over populations of participants. The two definitions coincide when the distribution and mechanism are deterministic and the mechanism has no reserves. Furthermore, we will show subsequently via closure properties that a mechanism's competitive efficiency over all distributions of values and bids is equal to its competitive efficiency over deterministic distributions. Hence it will often suffice to work with the simpler Definition 2.1 to characterize the competitive efficiency of Definition 4.2. Moreover, for deterministic value profiles, computing μ requires only knowledge of the bid distributions. Hence, though values appear in Definition 4.2, a mechanism's competitive efficiency is a property of its rules and not agents' preferences or beliefs.

We now derive two closure properties of competitive efficiency. The first is that competitive efficiency is closed with respect to reserve prices, i.e., if a mechanism has competitive efficiency μ without reserves, then its competitive efficiency with reserves is also μ. It will therefore suffice to analyze mechanisms without reserves. Then, we consider the impact of two types of randomization on competitive efficiency without reserves. First, we consider randomization over mechanisms. In other words, if the lowest competitive efficiency of any mechanism in a family of mechanisms is μ then the competitive efficiency of any convex combination of mechanisms in the family is at least μ. Second, we prove closure under mixture over distributions of bids and values. From this latter closure, we will conclude that it suffices to analyze competitive efficiency on fixed profiles of bids and values. Hence, as discussed above, Definition 2.1 is necessary and sufficient to bound competitive efficiency for mechanisms with deterministic allocations.

We start with closure under reserves. The mechanisms we consider in this paper compute payments formulaically from the allocation rule. For example, in a winner-pays-bid mechanism $\tilde{p}_i(\mathbf{b}) = b_i\tilde{x}_i(\mathbf{b})$. We can therefore incorporate reserves into a mechanism by modifying its allocation rule generically as in the following definition. For the payment formats we consider, this modification of the allocation rule will imply a corresponding modification of the payment rule.

Definition 4.3. *Given a mechanism with single-bid allocation rule $\tilde{\mathbf{x}}$, adding reserves*

$\mathbf{r}$ *yields a new rule* $\tilde{\mathbf{x}}^{\mathbf{r}}$ *that: (1) Solicits a bid* b_i *from each agent* i*. (2) For each agent* i*, if* $b_i < r_i$*, sets* $b_i^{r_i} = 0$*, else* $b_i^{r_i} = b_i$*. (3) Allocates according to* $\tilde{\mathbf{x}}(\mathbf{b}^{\mathbf{r}})$*.*

Throughout, we consider adding reserves to mechanisms where bidding 0 guarantees an agent goes unallocated, i.e. $\tilde{x}_i(0, \mathbf{b}_{-i}) = 0$. Given a mechanism with allocation rule $\tilde{\mathbf{x}}$ and payment rule $\tilde{\mathbf{p}}$, and given reserves $\mathbf{r}$ and bids $\mathbf{b}$, define $\tilde{\mathbf{x}}^{\mathbf{r}}$ and $\tilde{\mathbf{p}}^{\mathbf{r}}$ to be the respective allocation and payment rules of the mechanism with reserves, and define $\mathbf{b}^{\mathbf{r}}$ as in Definition 4.3. To analyze competitive efficiency of $(\tilde{\mathbf{x}}^{\mathbf{r}}, \tilde{\mathbf{p}}^{\mathbf{r}})$ on $\mathbf{b}$, we will compare it to $(\tilde{\mathbf{x}}, \tilde{\mathbf{p}})$ on $\mathbf{b}^{\mathbf{r}}$. By Definition 4.3, $\tilde{\mathbf{x}}^{\mathbf{r}}(\mathbf{b}) = \tilde{\mathbf{x}}(\mathbf{b}^{\mathbf{r}})$. For winner-pays-bid mechanisms, it also holds that $\tilde{\mathbf{p}}^{\mathbf{r}}(\mathbf{b}) = \tilde{\mathbf{p}}(\mathbf{b}^{\mathbf{r}})$. However, this equality need not hold in general: with truthful (second-price) payments, and hence also for mechanisms which combine winner-pays-bid and truthful payments (see Section 4.6), it can be that $\tilde{\mathbf{p}}^{\mathbf{r}}(\mathbf{b}) > \tilde{\mathbf{p}}(\mathbf{b}^{\mathbf{r}})$. (E.g., consider a single-item, truthful highest-bids win mechanism with just one agent, which effectively posts a price equal to the reserve.) Closure under reserve pricing will hold for all mechanisms discussed above, and more generally all those satisfying the following:

Definition 4.4. *Let* $\tilde{\mathbf{p}}$ *and* $\tilde{\mathbf{p}}^{\mathbf{r}}$ *be the payment rules of a mechanism without and with reserves* $\mathbf{r}$*, respectively. For a profile of bids* $\mathbf{b}$*, for each agent* i *let* $b_i^{r_i} = 0$ *if* $b_i < r_i$ *and* $b_i^{r_i} = b_i$ *otherwise. Then the mechanism is* reserve monotone *if for all* i*,* $\tilde{p}_i^{\mathbf{r}}(\mathbf{b}) \geq \tilde{p}_i(\mathbf{b}^{\mathbf{r}})$*.*

We may now state closure under reserve prices.

Lemma 4.5. *Competitive efficiency is closed under reserve pricing for reserve-monotone mechanisms, i.e., given a reserve-monotone mechanism* M *with competitive efficiency* μ *on* $\mathcal{G}$ *(without reserves), competitive efficiency of* M *with reserves on* $\mathcal{G}$ *is also at least* μ*.*

Proof. Let $M = M^{\mathbf{0}}$ be a mechanism with competitive efficiency μ without reserves (equivalently with reserves $\mathbf{0}$) on joint distribution $\mathcal{G}$ over values and bids. We will show that adding any profile of reserves $\mathbf{r}$ yields a mechanism $M^{\mathbf{r}}$ with competitive efficiency at least μ with reserves $\mathbf{r}$ on $\mathcal{G}$. That is, the worst-case competitive efficiency over M with any reserves is achieved by $M^{\mathbf{0}}$. Let $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{x}}^{\mathbf{r}}$ denote the allocation rules of M and $M^{\mathbf{r}}$, respectively.

For joint distribution $\mathcal{G}$, let $\mathcal{G}^{\mathbf{r}}$ denote the distribution obtained by setting to 0 all bids $b_i < r_i$ failing to meet the reserves $\mathbf{r}$. The main ideas of the proof are that (a) revenue for $M^{\mathbf{r}}$ on $\mathcal{G}$ is at least that of $M^{\mathbf{0}}$ on $\mathcal{G}^{\mathbf{r}}$, while allocations are equal, and (b) fixing the bid allocation rules, reserves only lower threshold surplus. Thus, a competitive efficiency without reserves implies only greater competitive efficiency with reserves. The following notation makes the allocation rule $\tilde{\mathbf{x}}$ and distribution $\mathcal{G}$ explicit in the notation for threshold surplus with discounted reserve: $T_i^{r_i}(x, \tilde{\mathbf{x}}, \mathcal{G} \mid v_i) = T_i^{r_i}(x \mid v_i)$. Zero reserves will be explicitly designated as such.

Consider the following analysis, with subsequent discussion:

$$\begin{aligned}
\textsc{Rev}(M^{\mathbf{r}}, \mathcal{G}) &\geq \textsc{Rev}(M^{\mathbf{0}}, \mathcal{G}^{\mathbf{r}}) && (4)\\
&\geq \mu \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^0(y_i(v_i), \tilde{\mathbf{x}}^{\mathbf{0}}, \mathcal{G}^{\mathbf{r}} \mid v_i)\right] && (5)\\
&= \mu \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^0(y_i(v_i), \tilde{\mathbf{x}}^{\mathbf{0}}, \mathcal{G}^{(0,\mathbf{r}_{-i})} \mid v_i)\right] && (6)\\
&= \mu \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^0(y_i(v_i), \tilde{\mathbf{x}}^{(0,\mathbf{r}_{-i})}, \mathcal{G} \mid v_i)\right] && (7)\\
&\geq \mu \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^{r_i}(y_i(v_i), \tilde{\mathbf{x}}^{(0,\mathbf{r}_{-i})}, \mathcal{G} \mid v_i)\right] && (8)\\
&= \mu \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i^{r_i}(y_i(v_i), \tilde{\mathbf{x}}^{\mathbf{r}}, \mathcal{G} \mid v_i)\right]. && (9)
\end{aligned}$$

Inequality (4) follows from reserve monotonicity. Inequality (5) follows from the assumed competitive efficiency of $M^{\mathbf{0}}$. Equation (6) follows from the fact that an agent's threshold surplus depends only on the bids others. Equation (7) follows from the fact that for any bids $\mathbf{b}$, $\tilde{x}_i(\mathbf{b}^{\mathbf{r}}) = \tilde{x}_i^{\mathbf{r}}(\mathbf{b})$. Inequality (8) follows from the definition of threshold surplus with discounted reserves; i.e., discounting reserves lowers the threshold surplus. Equation (9) follows because $T_i^{r_i}(\cdot \mid v_i)$ considers only allocation levels above $\tilde{x}_i(r_i \mid v_i)$. Combining the inequalities we observe that $M^{\mathbf{r}}$ has competitive efficiency at least μ on $\mathcal{G}$. □

Closure under convex combination of both mechanisms and distributions of values and bids will follow from a single convexity argument. We define each notion of convex combination below, then give a proof encompassing both.

Definition 4.6. *Let θ be a random variable indexing over mechanisms M^θ with allocation rules $\tilde{\mathbf{x}}^\theta$, payment rules $\tilde{\mathbf{p}}^\theta$, and feasibility environments $\mathcal{X}^\theta$. The convex combination of these mechanisms has allocation rule $\tilde{\mathbf{x}}(\mathbf{b}) = \mathbb{E}_\theta[\tilde{\mathbf{x}}^\theta(\mathbf{b})]$ and payment*

rule $\tilde{\mathbf{p}}(\mathbf{b}) = \mathbb{E}_\theta[\tilde{\mathbf{p}}^\theta(\mathbf{b})]$. *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 4.7. *Let* ω *be a random variable indexing over joint distributions* $\mathcal{G}^\omega$ *over values and bids. To draw from the convex combination* $\mathcal{G}$*, draw* ω*, then draw* $(\mathbf{v}, \mathbf{b}) \sim \mathcal{G}^\omega$.

Lemma 4.8. *Competitive efficiency with no reserves is closed under convex combination of mechanisms and of distributions, i.e. (i) if a mechanism* M *has competitive efficiency* μ *on a family of distributions* $\mathscr{G}$*, then it has competitive efficiency* μ *on the family of convex combinations over* $\mathscr{G}$*; (ii) if the family of mechanisms* $\mathscr{M}$ *has competitive efficiency* μ *on distribution* $\mathcal{G}$ *over values and bids, then the family of convex combinations over* $\mathscr{M}$ *also has competitive efficiency* μ *on* $\mathcal{G}$.

Proof. Let M be a convex combination of mechanisms indexed by θ (with corresponding allocation rules $\tilde{\mathbf{x}}^\theta$, payment rules $\tilde{\mathbf{p}}^\theta$, feasibility environments $\mathcal{X}^\theta$, and combinations $\tilde{\mathbf{x}}$, $\tilde{\mathbf{p}}$, and $\mathcal{X}$). Further let $\mathcal{G}$ be a convex combination of distributions, indexed by ω. Assume for all θ and ω, M^θ has competitive efficiency at least μ on $\mathcal{G}^\omega$. We will argue that M has competitive efficiency at least μ on $\mathcal{G}$. 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 there exists some collection of allocation functions $\mathbf{y}^\theta$ respectively feasible for $\mathcal{X}^\theta$ such that for each profile $\mathbf{v}$, we can write $\mathbf{y}(\mathbf{v}) = \mathbb{E}_\theta[\mathbf{y}^\theta(\mathbf{v})]$. 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 θ and ω are drawn independently, and where $y_i^{\theta,\omega}(v_i)$ is the interim allocation with respect to $\mathbf{y}^\theta$ under distribution $\mathcal{G}^\omega$.

Now consider an agent i with value v_i, and consider allocation level $y_i(v_i) = \mathbb{E}_{\theta,\omega}[y_i^{\theta,\omega}(v_i) \,|\, v_i]$ for some $y_i^{\theta,\omega}$s. Let $T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \,|\, v_i)$ denote i's threshold surplus with respect to M^θ and $\mathcal{G}^\omega$, and $T_i(y_i(v_i) \,|\, v_i)$ the threshold surplus with respect to M and $\mathcal{G}$. We will show that $\mathbb{E}_{\theta,\omega}[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \,|\, v_i) \,|\, v_i] \geq T_i(y_i(v_i) \,|\, v_i)$ with the

following inequalities, explained after their statement:

$$T_i(y_i(v_i) \mid v_i) = \int_0^\infty \max(y_i(v_i) - \tilde{x}_i(b \mid v_i), 0)\, db \tag{10}$$

$$= \int_0^\infty \max\left(\mathbb{E}_{\theta,\omega}\left[y_i^{\theta,\omega}(v_i) - \tilde{x}_i^{\theta,\omega}(b \mid v_i) \,\middle|\, v_i\right], 0\right) db \tag{11}$$

$$\leq \mathbb{E}_{\theta,\omega}\left[\int_0^\infty \max(y_i^{\theta,\omega}(v_i) - \tilde{x}_i^{\theta,\omega}(b \mid v_i), 0)\, db \,\middle|\, v_i\right] \tag{12}$$

$$= \mathbb{E}_{\theta,\omega}\left[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \mid v_i) \,\middle|\, v_i\right]. \tag{13}$$

The equalities (10) and (13) rewrite $\int_0^{y_i(v_i)} t_i(z \mid v_i)\, dz$ in terms of $\tilde{x}_i(b_i \mid v_i)$ rather than its inverse $t_i(z \mid v_i)$, using a change of variables. The equality (11) is from the definitions of convex combination, and linearity of expectation. The inequality (12) 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 $\mathbf{y}$. Let $\mathcal{F}$ and $\mathcal{F}^\omega$ be the respective marginal distributions over values profiles of $\mathcal{G}$ and $\mathcal{G}^\omega$. Then:

$$\textsc{Rev}(M, \mathcal{G}) = \mathbb{E}_{\theta,\omega}\left[\textsc{Rev}(M^\theta, \mathcal{G}^\omega)\right] \tag{14}$$

$$\geq \mathbb{E}_{\theta,\omega}\left[\mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i^\omega}\left[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \mid v_i)\right]\right] \tag{15}$$

$$= \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}}\left[\mathbb{E}_{\theta,\omega}\left[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \mid v_i) \,\middle|\, v_i\right]\right] \tag{16}$$

$$\geq \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}}\left[T_i(y_i(v_i) \mid v_i)\right]. \tag{17}$$

Line (14) follows from the definitions of convex combination and revenue. Line (15) follows from applying competitive efficiency. Line (16) follows from properties of conditional expectation, and finally line (17) follows because $\mathbb{E}_{\theta,\omega}[T_i^{\theta,\omega}(y_i^{\theta,\omega}(v_i) \mid v_i) \mid v_i] \geq T_i(y_i(v_i) \mid v_i)$, as argued above. We conclude that M has competitive efficiency at least μ on $\mathcal{G}$. □

The closure properties of Lemmas 4.5 and 4.8 simplify competitive efficiency analysis. First, by Lemma 4.5 it suffices to analyze mechanisms with no reserves. Second, by Lemma 4.8 the competitive efficiency of a distribution over deterministic mechanisms on a distribution of value and bid profiles is that of the worst mechanism on the worst profile in those distributions' supports. Hence, for families of distributions containing all such deterministic mechanisms and profiles, the competitive

efficiency computed from Definition 2.1 is tight. For any fixed mechanism, we obtain the following corollary:

Corollary 4.9. *A mechanism's competitive efficiency without reserves is equal to its competitive efficiency on the class of pointmass distributions over values and bids.*

Returning to the agent-normal form interpretation of incomplete information, Corollary 4.9 implies that a mechanism's competitive efficiency is the same whether participants are fixed or drawn from a population, represented by the prior over values. One additional observation is that in Definition 4.2, the prior $\mathcal{F}$ over values serves only to index agents from each population: neither the revenue nor the threshold surplus in (3) depends on the magnitude of the values. In the special case of a deterministic profile $(\mathbf{v}, \mathbf{b})$, the computation of a mechanism's competitive efficiency only uses the bids $\mathbf{b}$, and we may conclude:

Corollary 4.10. *For any mechanism M and any fixed value profile $(\mathbf{v}, \mathbf{b})$, the competitive efficiency of M on $(\mathbf{v}, \mathbf{b})$ does not depend on $\mathbf{v}$.*

Conceptually, Corollary 4.10 confirms the initial intuition that competitive efficiency depends on the way the mechanism maps bids to allocations and payments, and not on behavioral parameters of agents such as their values or information.

The remainder of this section presents competitive efficiency bounds for specific mechanisms. We focus on winner-pays-bid mechanisms throughout, then conclude the section with theorems extending the winner-pays-bid bounds to apply to all-pay and hybrid first/second-price mechanisms. Sections 4.1 and 4.2 review the single-item and single-minded combinatorial auction results for highest-bids-win mechanisms from Section 2. Section 4.3 considers greedy mechanisms, and shows that they lack the pathologies of highest-bids-win mechanisms for settings such as single-minded combinatorial auctions. We analyze position auctions in Section 4.4 and demonstrate the usefulness of closure under convex combination. In Section 4.5 we analyze auctions where agents have multi-unit demands for identical goods. Section 4.6 concludes as promised by extending the section's competitive efficiency results beyond winner-pays-bid mechanisms. Throughout, Corollaries 4.9 and 4.10 allow us to restrict attention to competitive efficiency on degenerate profiles of bids and arbitrary profiles of values in mechanisms without reserves.

4.1 Single-Item Multi-Unit Auctions

This section considers the n-agent single-item and multi-unit highest-bids-win mechanisms for unit-demand agents. Section 2 proved that for the single-unit case and the deterministic, full-information Definition 2.1, the competitive efficiency of the first-price auction was 1. Closure under convex combination, specifically Corollary 4.9, implies that the same bound holds under the more general Definition 4.2, and closure under reserve pricing (Lemma 4.5) extends this to the first price auction with reserves.

Theorem 4.11. *The first-price auction (possibly with reserves) has competitive efficiency* 1 *in single-item environments.*

In the more general *multi-unit* environment, there are k identical units of the item available. Agents demand a single unit. The natural generalization of the first-price auction allocates according to the highest-bids-win rule, allocating the k highest bidders, and charges winners their bids. A very similar proof to that in Section 2 shows a competitive efficiency of 1. Rather than give the elementary proof here, we will observe it as a consequence of Corollary 4.21, given subsequently.

Theorem 4.12. *For any multi-unit environment, the highest-bids-win winner-pays-bid mechanism has competitive efficiency 1.*

4.2 Single-Minded Combinatorial Auctions: Highest Bids Win

Lemma 2.12 demonstrated that there exist examples of single-minded combinatorial auctions (SMCA) where the highest-bids-win, winner-pays-bid mechanism has inefficient pure Nash equilibria. These examples had a small number of bidders winning bundles of items that were demanded by many other agents, yielding a multiplicative welfare loss as bad as $1/m$, where m is the number of items. This implied that competitive efficiency as given by Definition 2.1 was also $1/m$. Since Definition 2.1 and the more general Definition 4.2 coincide on this example, the same bound holds for the more general definition of competitive efficiency as well.

Lemma 4.13. *There exists a single-minded combinatorial auction environment where the highest-bids-win winner-pays-bid mechanism has competitive efficiency at most* $1/m$.

The next section studies greedy mechanisms, including a greedy mechanism for SMCAs. Comparison to highest-bids-win mechanisms illustrates different causes of inefficiency. The highest-bids-win rule is welfare-optimal in the absence of incentives. It was the introduction of incentives — specifically the ability of a few agents to discourage bidding by many others — that drove the inefficiency in Lemma 4.13. The greedy mechanisms of the next section will have suboptimal welfare in the absence of incentives, but will have better incentive properties (i.e. higher competitive efficiency) and hence higher robust welfare.

4.3 Greedy Auctions

The highest-bids-win mechanism is unable to effectively convert competition into revenue in combinatorial multi-item settings. We now show that other mechanisms may manage this relationship more effectively. In particular, this section considers mechanisms with *greedy* allocation rules, defined below for feasibility environments with allocations in $\{0,1\}^n$.

Definition 4.14. *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, set $S = S \cup \{i\}$.*
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 deterministic 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 the measure of approximate optimality for algorithms (ignoring incentives).

Definition 4.15. *An 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 competitive efficiency as follows, and prove the result after discussion.

Theorem 4.16. *For any feasibility environment* $\mathcal{X}$*, if* $\tilde{\mathbf{x}}$ *is a* α*-approximation for* $\mathcal{X}$*, then the winner-pays-bid mechanism for* $\tilde{\mathbf{x}}$ *has competitive efficiency at least* α*.*

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

Lemma 4.17 (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.*

Combining Theorem 4.16 with Lemma 4.17 yields the following:

Corollary 4.18. *For any single-minded combinatorial auction environment, the winner-pays-bid mechanism for the 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* $1/\sqrt{m}$*.*

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 4.1, and *transversal matroids*, which are matchable subsets of vertices on one side of a bipartite graph.

Definition 4.19. *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 4.20 (Kruskal, 1956)**.** *For any matroid feasibility environment, greedy by priority with $\psi_i(b_i) = b_i$ for all agents i is a 1-approximation.*

Corollary 4.21. *For any matroid environment, the highest-bids-win winner-pays-bid mechanism has competitive efficiency 1.*

To prove Theorem 4.16, it is helpful to compare the behavior of greedy mechanisms to the non-greedy highest-bids-win mechanism of Section 4.2, which had poor competitive efficiency. In the example which proved Lemma 4.13, the high bid of the $(m+1)$st 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 4.22. *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 4.23. *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 4.24. *Any winner-pays-bid mechanism with a 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.

$$\sum_i b_i \tilde{x}_i(\mathbf{b}) = \sum_i b'_i \tilde{x}_i(\mathbf{b}) \tag{18}$$

$$= \sum_i b'_i \tilde{x}_i(\mathbf{b}') \tag{19}$$

$$\geq \alpha \sum_i b'_i y_i \tag{20}$$

$$\geq \alpha \sum_i \hat{b}_i(\mathbf{b}_{-i}) y_i. \tag{21}$$

The first line (18) holds because $\mathbf{b}'$ differs from $\mathbf{b}$ only on the bids of losing agents. Line (19) follows from the coalitional non-bossiness of greedy rules, and (20) from the assumption that the greedy rule is an α-approximation (Definition 4.15). The last line (21) 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 4.16 follows from combining Lemma 4.23 with Lemma 4.24.

4.4 Position Auctions

The mechanisms and environments in Section 4.1-4.3 have been deterministic. We now consider the canonical and inherently randomized allocation environment of position auctions and show that the most natural single-bid mechanism has competitive efficiency 1. This result follows directly from the competitive efficiency of multi-unit highest-bids-win mechanism (Theorem 4.12) and closure of competitive efficiency under convex combination (Lemma 4.8).

Position environments are a standard model for internet advertising auctions, e.g., Varian (2007) and Edelman et al. (2007). 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 4.25. *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, *it holds that* $\tilde{x}_i \in \{0, \alpha_{\pi(i)}\}$.

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

Definition 4.26. *The* rank-by-bid *rule for position auctions assigns agents to positions in order of their bid. The agent in position 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 the closure under convex combination of mechanisms (Lemma 4.8) and the competitive efficiency of multi-unit highest-bids-win auctions (Theorem 4.12).

Lemma 4.27 (Devanur et al., 2015)**.** *The rank-by-bid allocation rule 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 $\alpha_{n+1} = 0$).*

It has become increasingly common for online advertising markets to use winner-pays-bid auctions (Paes Leme et al., 2020), with Google adopting the format in 2019 (Bigler, 2019). For the winner-pays-bid mechanism with rank-by-bid allocations, we obtain the following:

Theorem 4.28. *For any position environment, the winner-pays-bid mechanism for the rank-by-bid rule has competitive efficiency* 1.

4.5 Multi-Unit Demands

In this section we consider an allocation environment where agents have heterogeneous, multi-unit demands. This generalizes the multi-unit unit-demand environment of Section 4.1: agents might want more than one copy of the k units available, with their utility being linear in the number of items received. Consequently, agents can be partially served if they receive some but not all of their demanded quantity. The main result of this section is that the winner-pays-bid highest-bids-win mechanism for this setting has competitive efficiency 1. We first analyze the following limit setting with a fully divisible good, before extending the analysis to the setting just described, with indivisible items.

Definition 4.29. *In a* capacitated divisible good *environment, the auctioneer sells a mass of a good, normalized to* 1*. Each agent i desires the good up to a publicly known capacity c_i. The set of feasible allocations is $\mathcal{X} = \{\mathbf{x} \in [0,1]^n \mid x_i \leq c_i, \sum_i x_i \leq 1\}$.*

In this setting, the highest-bids-win rule sorts agents in nonincreasing order by bid, then maximally allocates each in order until the supply is exhausted. Note that the allocation problem for the capacitated divisible good environment is mathematically equivalent to the well-studied *fractional knapsack* problem, where each agent i has a weight w_i and allocation $\mathbf{x} \in [0,1]^n$ is feasible if $\sum_i w_i x_i \leq 1$. For the latter environment the highest-bids-win algorithm sorts algorithms by *bid density* b_i/w_i and allocates maximally in nonincreasing order. An allocation $\mathbf{x}$ in the knapsack environment can be mapped to the divisible good environment by setting $c_i = w_i$ and treating knapsack allocation x_i as the fraction of the demand c_i received by agent i. Moreover, the competitive efficiency of the winner-pays-bid highest-bids-win mechanisms for the capacitated divisible good and fractional knapsack environment are the same.

The competitive efficiency analysis of this section does not follow straightforwardly from the previous analyses: like the position auctions of Section 4.4, optimal allocations may be fractional. However, unlike position auctions, the capacitated digital good environment cannot easily be written as a convex combination of deterministic allocation problems. Instead, the following decomposition of the winner-pays-bid highest-bids-win mechanism's revenue into the prices of marginal units of allocation will be central to the analysis.

Definition 4.30. *Given a profile of bids $\mathbf{b}$, the agent receiving marginal unit $z \in [0,1]$, denoted $\pi(z)$, can be computed as follows. Sort agents in nonincreasing order of bid, and let (j) denote the index of the agent with the jth highest bid. Then $\pi(z)$ is the smallest i such that $\sum_{j=1}^{i} c_{(j)} \geq z$. The revenue from selling the zth marginal unit is thus $b_{\pi(z)}\, dz$.*

Note that $b_{\pi(z)}$ is nonincreasing by definition. It can be interpreted as the price being paid for a particular small fraction of the unit mass being auctioned. Note further that the revenue of the mechanism $\sum_i b_i \tilde{x}_i(\mathbf{b})$ is equal to $\int_0^1 b_{\pi(z)}\, dz$. The competitive efficiency of the highest-bids-win mechanism will follow from comparing the prices for marginal units to the marginal threshold surplus for each agent.

Lemma 4.31. *For any capacitated divisible goods environment, the winner-pays-bid highest-bids-win mechanism has competitive efficiency* 1*.*

Proof. By Corollaries 4.9 and 4.10, we may consider a deterministic profile of bids $\mathbf{b}$ and arbitrary (also deterministic) values. Fixing these, an agent i faces an allocation rule $\tilde{x}_i(\cdot, \mathbf{b}_{-i})$, with inverse $t_i(x) = \inf\{b \,|\, \tilde{x}_i(b, \mathbf{b}_{-i}) \geq x\}$. An agent i's threshold surplus is $T_i(x) = \int_0^x t_i(z)\, dz$. For deterministic profiles, it suffices to show the revenue $\sum_i b_i \tilde{x}_i(\mathbf{b})$ is at least the threshold surplus $\sum_i T_i(y_i)$ for any ex post feasible $\mathbf{y} \in \mathcal{X}$.

Since $b_{\pi(z)}$ is nonincreasing, at least z units in total are allocated to agents with bids less than or equal to $b_{\pi(1-z)}$. Consequently, for any agent i and $z \in [0, c_i]$, if i bid $b_i > b_{\pi(1-z)}$, they would win at least z units of allocation. Hence, $b_{\pi(1-z)} \geq t_i(z)$. We can therefore write the threshold surplus for feasible allocation $\mathbf{y}$ as:

$$\begin{aligned}
\sum\nolimits_i T_i(y_i) &= \sum\nolimits_i \int_0^{y_i} t_i(z)\, dz \\
&\leq \sum\nolimits_i \int_0^{y_i} b_{\pi(1-z)}\, dz \\
&\leq \int_0^1 b_{\pi(1-z)}\, dz \qquad (22) \\
&= \sum\nolimits_i b_i \tilde{x}_i(\mathbf{b}), \qquad (23)
\end{aligned}$$

where line (22) follows from the facts that $\sum_i y_i \leq 1$ and $b_{\pi(z)}$ is nonincreasing in z, and line (23) from the fact that $\int_0^1 b_{\pi(z)}\, dz$ is the revenue of the mechanism. □

An analogous result for the multi-unit indivisible goods environment follows as a corollary. To make the connection apparent, the following definition scales allocations down to increments of $1/k$, where k is the number of units available.

Definition 4.32. *In a* k-unit environment with heterogenenous demands*, the auctioneer sells* 1 *unit of a good, in increments of* $1/k$*. Each agent* i *demands a publicly known number of increments* d_i*. The set of feasible allocations is* $\mathcal{X} = \{\mathbf{x} \in \{0, 1/k, 2/k, \ldots, 1\}^n \,|\, x_i \leq d_i/k, \sum_i x_i \leq 1\}$*.*

The highest-bids-win rule for the (scaled) k-unit environment is identical to that for the divisible good environment with capacities $c_i = d_i/k$. Hence:

Corollary 4.33. *For any* k*-unit environment with heterogeneous demands, the winner-pays-bid highest-bids-win mechanism has competitive efficiency* 1*.*

Proof. Let $\mathcal{X}$ denote the set of feasible allocations for the k-unit environment with demands $\mathbf{d}$, and let $\hat{\mathcal{X}}$ denote that for the divisible good environment with capacities $\mathbf{c} = \mathbf{d}/k$. The highest-bids-win mechanism for the two environments selects the same allocations and yields the same revenue. By Lemma 4.31, the revenue of the highest-bids-win mechanism is at least the threshold surplus for any $\mathbf{y} \in \hat{\mathcal{X}}$. Since $\mathcal{X} \subset \hat{\mathcal{X}}$, the revenue of the highest-bids-win mechanism is also at least the threshold surplus for any $\mathbf{y} \in \mathcal{X}$. □

4.6 Beyond Winner-Pays-Bid Mechanisms

The competitive efficiency analyses of Sections 4.1-4.5 focused on winner-pays-bid mechanisms. This section extends these results to all-pay auctions and hybrid first/second-price auctions, two other common single-bid payment formats. Below, we define each, then give an extension theorem for each which translates the winner-pays-bid competitive efficiency bounds to the other payment format. We first consider all-pay auctions.

Definition 4.34. *An* all-pay *mechanism has payments $\tilde{p}_i(\mathbf{b}) = b_i$ for each agent i.*

Lemma 4.35. *Let M^{WPB} be a winner-pays-bid mechanism with allocation rule $\tilde{\mathbf{x}}$ and competitive efficiency μ. Let M^{AP} be the all-pay mechanism with allocation rule $\tilde{\mathbf{x}}$. Then M^{AP} has competitive efficiency at least μ.*

Proof. Let $\mathcal{G}$ be a distribution over values and bids, with values following marginal $\mathcal{F}$. Let $T_i^{\text{WPB}}(y_i \,|\, v_i)$ and $T_i^{\text{AP}}(y_i \,|\, v_i)$ denote the threshold surplus under $\mathcal{G}$ for i under M^{WPB} and M^{AP}, respectively. Since the two mechanisms have the same allocation rule, $T_i^{\text{WPB}}(y_i \,|\, v_i) = T_i^{\text{AP}}(y_i \,|\, v_i)$. Meanwhile, $\textsc{Rev}(M^{\text{AP}}, \mathcal{G}) = \mathbb{E}[\sum_i b_i] \geq \mathbb{E}[\sum_i b_i \tilde{x}_i(\mathbf{b})] = \textsc{Rev}(M^{\text{WPB}}, \mathcal{G})$. Combining these facts for interim allocation $\mathbf{y}$:

$$\begin{aligned}
\textsc{Rev}(M^{\text{AP}}, \mathcal{G}) &\geq \textsc{Rev}(M^{\text{WPB}}, \mathcal{G}) \\
&\geq \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[T_i^{\text{WPB}}(y_i(v_i) \,|\, v_i) \right] \\
&= \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[T_i^{\text{AP}}(y_i(v_i) \,|\, v_i) \right].
\end{aligned}$$

We conclude that M^{AP} has competitive efficiency at least μ. □

We next consider hybrid mechanisms, that charge a convex combination of first-price and second-price payments. (See Menezes and Monteiro, 2006; van Beek et al., 2022, for an application to corruption in auctions.) The definition below generalizes to multi-item environments and environments with fractional allocation by charging a convex combination of winner-pays-bid and dominant strategy truthful payments.

Definition 4.36. *For nondecreasing allocation rule $\tilde{\mathbf{x}}$ and agent i, denote winner-pays-bid payments by $\tilde{p}_i^{\mathrm{WPB}}(\mathbf{b}) = b_i\tilde{x}_i(\mathbf{b})$ and dominant-strategy truthful payments by $\tilde{p}_i^{\mathrm{DST}}(\mathbf{b}) = b_i\tilde{x}_i(\mathbf{b}) - \int_0^{b_i} \tilde{x}_i(b, \mathbf{b}_{-i})\, db$. The γ-hybrid payment rule for $\tilde{\mathbf{x}}$ is $\tilde{p}_i^{\gamma}(\mathbf{b}) = \gamma\tilde{p}_i^{\mathrm{WPB}}(\mathbf{b}) + (1-\gamma)\tilde{p}_i^{\mathrm{DST}}(\mathbf{b})$.*

For fixed bid profile $\mathbf{b}$, the threshold surplus for a γ-hybrid mechanism does not depend on γ, and is the same as that of a winner-pays-bid mechanism ($\gamma = 1$). However, dominant-strategy truthful payments are generally smaller than winner-pays-bid payments. The competitive efficiency guarantee for γ-hybrid mechanisms therefore degrades naturally with the fraction γ of winner-pays-bid payment. At the extreme with $\gamma = 0$ (e.g. the second-price auction), the existence of fully inefficient equilibria is well-known, and so we should expect $\mu = 0$.

Lemma 4.37. *Let M^{WPB} be a winner-pays-bid mechanism with allocation rule $\tilde{\mathbf{x}}$ and competitive efficiency μ. Let M^{γ} be the γ-hybrid mechanism with allocation rule $\tilde{\mathbf{x}}$. Then M^{γ} has competitive efficiency at least $\gamma\mu$.*

Proof. Let $T_i^{\mathrm{WPB}}(y_i \,|\, v_i)$ and $T_i^{\gamma}(y_i \,|\, v_i)$ denote the threshold surplus under a distribution $\mathcal{G}$ for agent i under M^{WPB} and M^{γ}, respectively. Then $T_i^{\mathrm{WPB}}(y_i \,|\, v_i) = T_i^{\gamma}(y_i \,|\, v_i)$. Further let M^{DST} denote the mechanism which charges dominant-strategy truthful payments. Note that these payments are nonnegative. Then we may lower bound the revenue of M^{γ} as:

$$\begin{aligned}
\textsc{Rev}(M^{\gamma}, \mathcal{G}) &= \gamma\textsc{Rev}(M^{\mathrm{WPB}}, \mathcal{G}) + (1-\gamma)\textsc{Rev}(M^{\mathrm{DST}}, \mathcal{G}) \\
&\geq \gamma\textsc{Rev}(M^{\mathrm{WPB}}, \mathcal{G}) \\
&\geq \gamma\mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i}\left[T_i^{\mathrm{WPB}}(y_i(v_i) \,|\, v_i)\right] \\
&= \gamma\mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i}\left[T_i^{\gamma}(y_i(v_i) \,|\, v_i)\right].
\end{aligned}$$

This implies a competitive efficiency bound of $\gamma\mu$. □

5 Individual Efficiency

The previous section studied competitive efficiency, which quantified the way mechanisms manage inter-agent competition. Competitive efficiency focused on the mechanism's rules in isolation, independent of the behavioral assumptions on participating agents and information structure of the game. This section turns to individual efficiency, which quantifies how an agent's solution to their individual bidding problem impacts welfare. It depends on agents' beliefs and behavior, and abstracts the rules of the larger mechanism down to single-agent bidding problems. The main generalization of this section over the simple definition in Section 2 is that we consider incomplete information, allowing uncertainty and over bids and private values. Individual efficiency then considers agents' interim bidding problems: conditioning on their value, each agent faces a distribution of competing bids, which induces a single-agent mechanism.

The analysis of this section distinguishes between two cases, depending on the structure of the agent's information. We first define individual efficiency in the case where the agent's information is summarized by their value and nothing further. In this case, the agent bids against a single interim mechanism, with fixed allocation and payment rules. We next consider more general information structures, i.e. where the agent observes some signal in addition to their value. In this case, the agent faces a distribution over single-agent mechanisms, and can correlate their bid with this distribution. Importantly, however, we then show that the case of general information reduces to the former signal-free case. In other words, the individual efficiency with additional information is no worse than that without it. It will therefore suffice to consider the signal-free case in the subsequent individual efficiency analyses.

In what follows, we present both definitions of individual efficiency, then analyze the individual efficiency of several common payment formats under varying degrees of agent best response. Section 5.1 presents the signal-free definition of individual efficiency. Section 5.2 provides additional discussion, motivating individual efficiency in the context of mechanisms for selling to a single buyer, when the seller has an outside option. In this setting, individual efficiency directly translates to allocative efficiency, i.e. high welfare. Section 5.3 then gives the full definition for more general information structures. We also give two closure properties of individual efficiency:

first, we prove that individual efficiency is closed under the addition of signals, as previously described. Second, prove closure of individual efficiency under approximate best response, i.e. we show that individual efficiency degrades smoothly as agents go from selecting exactly to approximately optimal bids.

The final three subsections prove concrete individual efficiency guarantees. Section 5.4 generalizes the analysis of step function allocation rules from Lemma 2.6. Section 5.5 shows that winner-pays-bid payments have individual efficiency $1 - 1/e$, and more generally that γ-hybrid mechanisms (Definition 4.36) have individual efficiency that improves smoothly from $1 - 1/e$ to 1 as the probability γ of first-price payments decreases from 1 down to 0. Finally, Section 5.6 considers all-pay mechanisms, which are only individually efficient by a relaxed definition, and even by this definition are less individually efficient than first-price (and hence also hybrid) mechanisms. Taken together, these analyses give a spectrum of individual efficiency across payment formats, with 0-hybrid (truthful) mechanisms being the most individually efficient, and all-pay mechanisms the least.

5.1 Individual Efficiency: Signal-Free Definition

This section considers signal-free single-agent bidding outcomes (Definition 2.2, restated below). The agent has a fixed value v, faces an allocation rule $\tilde{x}$ and payment rule $\tilde{p}$, and places a single bid b. For this setting, we present a definition of individual efficiency that is stronger than Definition 2.3, and that generalizes beyond the $\{0, 1\}$-valued allocation rules considered in Section 2. Section 5.3 will further generalize the set of bidding outcomes considered, but the analysis will reduce to the set of outcomes studied here. Recall the definition of signal-free bidding outcomes:

Definition 2.2. *A* single-agent signal-free bidding outcome $\mathcal{O} = (v, \tilde{x}, \tilde{p}, b)$ *is given by a value v, a bid b, and a single-agent mechanism with allocation rule $\tilde{x}$ (mapping bids to allocation levels in $[0, 1]$) and payment rule $\tilde{p}$ (mapping bids to nonnegative payments). The agent's utility for any bid $b^\dagger$ is $\tilde{u}(b^\dagger) = v\tilde{x}(b^\dagger) - \tilde{p}(b^\dagger)$. The agent's utility for their chosen bid is denoted $u(\mathcal{O}) = \tilde{u}(b)$.*

Individual efficiency quantifies how easily an agent is able to obtain high utility when their competition is weak. For the deterministic single-agent mechanisms of

Section 2, the strength of competition was measured by the threshold bid $\hat{b}$. With the more general allocation rules considered here, we measure the strength of competition by the threshold surplus $T(z) = \int_0^z t(x)\,dx$, where recall $t(x) = \inf\{b \mid \tilde{x}(b) \geq x\}$. We further generalize individual efficiency by quantifying the tradeoff with respect to a (possibly fractional) target level of allocation $z \in [0, 1]$, which can be interpreted as the agent's efficient interim probability of allocation. Quantifying over all such z yields a stronger notion of individual efficiency than Definition 2.3, which implicitly fixed $z = 1$. The definition is as follows.

Definition 5.1. *The* individual efficiency *of single-agent signal-free bidding outcome $\mathcal{O} = (v, \tilde{x}, \tilde{p}, b)$ for target allocation z is the ratio $\eta(z) = (u(\mathcal{O}) + T(z))/vz$. We refer to $\eta = \inf_{z\in[0,1]} \eta(z)$ simply as the individual efficiency of $(v, \tilde{x}, \tilde{p}, b)$. The individual efficiency of a family of signal-free bidding outcomes is the smallest individual efficiency of any outcome in the family.*

5.2 Discussion: Single-Buyer Mechanisms

Individual efficiency captures the impact of an agent's bidding behavior on social welfare. The following discussion defines and analyzes a family single-buyer mechanisms where this intuition is most apparent. In contrast to the rest of the paper, we consider a seller who also has a value for the item being sold. We measure social welfare by the total utility of the buyer and seller, and show below that for single-bid mechanisms in this setting, outcomes with high individual efficiency also have high welfare. The robust efficiency results of Section 6 extend this reasoning to multi-buyer mechanisms.

Formally, we consider a seller with value v_0 for the single item being sold, and a buyer with value v. For this setting, we consider the family of single-bid mechanisms which allocate the item to the buyer whenever their bid b is at least the seller's value v_0. We assume the seller value v_0 is unknown to the buyer, and drawn from some prior. To parallel Definition 2.2 of signal-free bidding outcomes, we assume the buyer's value v is fixed. Under these assumptions, the single-bid mechanisms we consider induce a probability of allocation $\tilde{x}(b) = \Pr[b \leq v_0]$ for each bid b from the buyer. We allow arbitrary payment rules, yielding the following definition.

Definition 5.2. *A* single-buyer instance *is given by buyer value v and distribution F for seller value v_0. A single-bid mechanism for a single-buyer instance solicits a*

bid b from the buyer, allocates if and only if $b < v_0$ (inducing bid allocation rule $\tilde{x}(b) = F(b)$), and charges payments according to payment rule $\tilde{p}(b, v_0)$ (inducing expected payments $\tilde{p}(b) = \mathbb{E}_{v_0}[\tilde{p}(b, v_0)]$).[4]

The single-buyer setting provides additional intuition for Definition 5.1 of individual efficiency. Comparing Definition 2.2 of signal-free bidding outcomes to Definition 5.2, we see that any signal-free bidding outcome with nondecreasing allocation rule $\tilde{x}$ can be mapped to a single-buyer instance by giving the seller's value a CDF equal to $\tilde{x}$, and vice versa. Moreover, the lemma below shows that the threshold surplus $T(1) = \int_0^1 t(x)\,dx$ is exactly the seller's expected value $\mathbb{E}[v_0]$, and more generally the threshold surplus $T(z)$ for $z < 1$ is the contribution of low values (in particular, the lowest z-fraction) to the expectation. Hence, the individual efficiency of a bidding outcome can be understood to measure the tradeoff between the agent's utility and the strength of the seller's value distribution in the implied single-buyer mechanism.

Lemma 5.3. *Given a single-buyer instance with random seller value v_0 and fixed buyer value v, let $z = Pr[v_0 \le v]$. Then it holds that $T(z) = \mathbb{E}[v_0 \,|\, v_0 \le v]\, Pr[v_0 \le v]$.*

Proof. The lemma follows from the definitions of $T(z)$ and conditional expectation. In more detail, integrating over values rather than probabilities yields the identity $T(z) = \int_0^v z - \tilde{x}(b)\,db$. Dividing out by $z = \Pr[v_0 \le v]$ and applying the definition of conditional probability gives the equivalent expression $\int_0^v 1 - \Pr[v_0 \le b \,|\, v_0 \le v]\,db \Pr[v_0 \le v]$. Finally, applying the standard formula for the expectation of a nonnegative random variable, we obtain $\mathbb{E}[v_0 \,|\, v_0 \le v] \Pr[v_0 \le v]$ as desired. □

Now define welfare for a single-buyer instance with buyer bid b to be the utility of the seller, given by $\tilde{p}(b) + \mathbb{E}[v_0 \,|\, v_0 > b] \Pr[v_0 > b]$, plus the utility of the buyer, given by $v\tilde{x}(b) - \tilde{p}(b)$. Equivalently, this welfare is the value of the agent who receives the item, given by $\mathbb{E}[v_0 \,|\, v_0 > b] \Pr[v_0 > b] + v \Pr[v_0 \le b]$. The welfare-optimal allocation allocates the item to the agent with the higher value, yielding welfare $\mathbb{E}[v_0 \,|\, v_0 > v] \Pr[v_0 > v] + v \Pr[v_0 \le v]$. We show below that when a single-buyer instance has individual efficiency η, the welfare is also a η-fraction of optimal. The result requires the following mild condition on payments:

[4] We consider only buyer incentives, and assume the seller's true value is known to the mechanism.

Definition 5.4. *A single-buyer instance satisfies* seller individual rationality (IR) *if the seller never loses utility by participating:* $\tilde{p}(b) \geq \mathbb{E}[v_0 \,|\, v_0 \leq b]\, Pr[v_0 \leq b]$ *for all* b.

The welfare guarantee for single-buyer instances is the following:

Proposition 5.5. *Consider a single-buyer instance with buyer value* v*, induced allocation rule* $\tilde{x}$ *and payment rule* $\tilde{p}$*, and bid* b*. If the instance satisfies seller IR and the corresponding signal-free bidding outcome has individual efficiency* $\eta \leq 1$*, then the single-buyer instance has welfare at least a* η*-fraction of optimal.*

Proof. The optimal social welfare is whichever is higher of the buyer's or seller's values, which has expected value $v \Pr[v_0 \leq v] + \mathbb{E}[v_0 \,|\, v_0 > v] \Pr[v_0 > v]$. Letting $z^* = \Pr[v_0 \leq v]$, we may rewrite the optimal welfare as $vz^* + \mathbb{E}[v_0 \,|\, v_0 > v] \Pr[v_0 > v]$. We will compare this quantity to the welfare of the single-buyer mechanism, i.e. the sum of the buyer utility $\tilde{u}(b) = v\tilde{x}(b) - \tilde{p}(b)$ and seller utility $\tilde{p}(b) + \mathbb{E}[v_0 \,|\, b \leq v_0]\Pr[b \leq v_0]$. The result follows from three observations: first, seller IR guarantees that the seller utility is at least the expected seller value, $\mathbb{E}[v_0]$. Second, Lemma 5.3 implies $\mathbb{E}[v_0 \,|\, v_0 \leq v] \Pr[v_0 \leq v] = T(z^*)$. Finally, individual efficiency implies that $T(z^*)$ and the buyer utility $\tilde{u}(b)$ together approximate vz^*.

Formally, first consider the seller's utility:

$$\begin{aligned}
&\tilde{p}(b) + \mathbb{E}[v_0 \,|\, v_0 > b]\Pr[v_0 > b] \\
&\quad \geq \mathbb{E}[v_0 \,|\, v_0 \leq b]\Pr[v_0 \leq b] + \mathbb{E}[v_0 \,|\, v_0 > b]\Pr[v_0 > b] && (24)\\
&\quad = \mathbb{E}[v_0] && (25)\\
&\quad = \mathbb{E}[v_0 \,|\, v_0 \leq v]\Pr[v_0 \leq v] + \mathbb{E}[v_0 \,|\, v_0 > v]\Pr[v_0 > v] && (26)\\
&\quad = T(z^*) + \mathbb{E}[v_0 \,|\, v_0 > v]\Pr[v_0 > v]. && (27)
\end{aligned}$$

Line (24) follows from seller IR. The equality of lines (24), (25), and (26) holds by the law of total expectation. Line (27) follows from Lemma 5.3. We can now combine all three components of the single-buyer mechanism's welfare to obtain a lower bound in terms of the optimal welfare:

$$\begin{aligned}
&\tilde{u}(b) + \tilde{p}(b) + \mathbb{E}[v_0 \,|\, v_0 \geq b]\Pr[v_0 > b] \\
&\quad \geq u(v, \tilde{x}, \tilde{p}, b) + T(z^*) + \mathbb{E}[v_0 \,|\, v_0 > v]\Pr[v_0 > v] \\
&\quad \geq \eta v z^* + \mathbb{E}\left[v_0 \,|\, v_0 > v\right] \Pr\left[v_0 > v\right] && (28)\\
&\quad \geq \eta(v z^* + \mathbb{E}\left[v_0 \,|\, v_0 > v\right] \Pr\left[v_0 > v\right]). && (29)
\end{aligned}$$

Inequality (28) follows from individual efficiency, and (29) because $\eta \leq 1$. □

5.3 Individual Efficiency: General Definition

We now extend Definition 5.1 to more general information structures. In equilibrium, an agent with value v receives an additional signal, correlated with other agents' bids. This induces a distribution over single-agent interim mechanisms, and the agent may use their signal to correlate their bid with this distribution. The following generalization of single-agent bidding outcomes permits correlation of this form.

Definition 5.6. *A* single-agent bidding outcome $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ *is given by a value $v \geq 0$, a distribution over single-agent mechanisms with allocation rules $\tilde{x}^\theta$ and payment rules $\tilde{p}^\theta$, jointly indexed by random variable θ, and a distribution over bids b^θ, also indexed by θ. The agent's utility for $\mathcal{O}$ is $u(\mathcal{O}) = \mathbb{E}_\theta[v\tilde{x}^\theta(b^\theta) - \tilde{p}^\theta(b^\theta)]$.*

To model an agent with additional information, we may take θ in Definition 5.6 to be the agent's signal. Note that the definition strictly generalizes Definition 2.2 of signal-free bidding outcomes, as we may fix a single allocation rule, payment rule, and bid. For the more general definition, individual efficiency still quantifies the welfare impact of an agent's behavior conditioned on their value v. The expected utility of such an agent is $\mathbb{E}_\theta[v\tilde{x}^\theta(b^\theta) - \tilde{p}^\theta(b^\theta)]$, and their competition is measured with respect to expected allocation rule $\tilde{x}(b) = \mathbb{E}_\theta[\tilde{x}^\theta(b)]$. The definition is again parametrized by target allocation z, representing the efficient interim allocation level when the agent's value is v. The necessary notation and definition are below.

Definition 5.7. *In a single-agent bidding outcome, the* expected allocation rule *at bid b is $\tilde{x}(b) = \mathbb{E}_\theta[\tilde{x}^\theta(b)]$, with inverse $t(x) = \inf\{b \mid \tilde{x}(b) \geq x\}$, and* threshold surplus $T(z) = \int_0^z t(x)\,dx$.

Definition 5.8. *For single-agent bidding outcome $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$, the* individual efficiency *for target allocation z is the ratio $\eta(z) = (u(\mathcal{O}) + T(z))/vz$. We refer to $\eta = \inf_{z\in[0,1]} \eta(z)$ simply as the individual efficiency of $\mathcal{O}$. The individual efficiency of a family of single-agent bidding outcomes is the smallest individual efficiency of any outcome in the family.*

We now prove closure of individual efficiency under information and approximate best response. Before stating the closure properties, we define exact and approximate best response for bidding outcomes without and with signals. Notably, with signals, we consider *coarse best response*, which compares the agent's utility to the best single bid. This will encompass the bidding outcomes arising from (Bayes) coarse correlated equilibria (c.f. Bergemann and Morris, 2016).

Definition 5.9. *A single-agent signal-free bidding outcome $\mathcal{O} = (v, \tilde{x}, \tilde{p}, b)$ satisfies* $(1-\epsilon)$-best response *if for every fixed deviation bid b', we have $u(\mathcal{O}) \geq (1-\epsilon)(v\tilde{x}(b') - \tilde{p}(b'))$. If $\epsilon = 0$, we say $\mathcal{O}$ satisfies* best response.

Definition 5.10. *A single-agent bidding outcome $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ satisfies* coarse $(1-\epsilon)$-best response *if for every fixed deviation bid b', we have $u(\mathcal{O}) \geq (1-\epsilon)\,\mathbb{E}_\theta[v\tilde{x}^\theta(b') - \tilde{p}^\theta(b')]$. If $\epsilon = 0$, we say $\mathcal{O}$ satisfies* coarse best response.

Finally, Definition 5.11, below, will help to state closure under information, and to prove both closure properties. Given a distribution over single-agent mechanisms indexed by a signal, it defines a new mechanism induced by ignoring the signal.

Definition 5.11. *For distribution over single-agent mechanisms with allocation rules $\{\tilde{x}^\theta\}_\theta$ and payment rules $\{\tilde{p}^\theta\}_\theta$, define $\tilde{x}(b) = \mathbb{E}_\theta[\tilde{x}^\theta(b)]$ and $\tilde{p}(b) = \mathbb{E}_\theta[\tilde{p}^\theta(b)]$. Then the* signal-free projection *of $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ is the single-agent mechanism $(\tilde{x}, \tilde{p})$.*

The closure properties are now as follows. Informally, closure under information states that for a best-responding agent, additional information can only make bids more individually efficient: if the individual efficiency of a family of signal-free bidding outcomes is η, then that of the same bidding outcomes with additional information is at least η as well. Meanwhile, closure under approximate best response states that bids that are nearly best responses have nearly the same individual efficiency as exact best responses. Formally:

Lemma 5.12. *Individual efficiency is closed under:*

(i) Additional information: Let $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ denote a joint distribution over allocation and payment rules. For value v, assume that all best responses to the signal-free projection $(\tilde{x}, \tilde{p})$ have individual efficiency η. Then all coarse best responses to $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ have individual efficiency at least η.

(ii) Approximate best response: Let $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ denote a joint distribution over allocation and payment rules. For value v, assume that all coarse best responses to $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ have individual efficiency at least η. Then all coarse $(1-\epsilon)$-best responses to $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ have individual efficiency at least $(1-\epsilon)\eta$.

Lemma 5.12 will follow from a single argument about coarse best responses, given in Lemma 5.13. The definition of coarse (approximate) best response compares the agent's utility to that from the best single bid. Lemma 5.13 therefore compares the individual efficiency of coarse (approximate) best responses to the individual efficiency of that bid as well, for any choice z of target allocation level. In proving Lemma 5.12 and in subsequent applications we will assume that such a best response bid exists. Note that we make this assumption solely to illustrate key concepts clearly. The individual efficiency guarantees we derive can also be proven directly, even in the absence of a best response bid.

Lemma 5.13. *Let $(v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ be a single-agent bidding outcome satisfying coarse $(1-\epsilon)$-approximate best response. If for signal-free projection $(\tilde{x}, \tilde{p})$ of $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ and best response bid b, the individual efficiency of $(v, \tilde{x}, \tilde{p}, b)$ for target allocation z is η, then the individual efficiency of $(v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ for z is at least $(1-\epsilon)\eta$.*

Proof. The lemma follows from two observations. First, since the expected allocation rules of $(v, \tilde{x}, \tilde{p}, b)$ and $(v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ are identical, the threshold surplus $T(z)$ is also identical. Second, coarse (approximate) best response directly compares the utility of $\{b^\theta\}_\theta$ to that of b. More precisely, let u and u^{SF} denote the agent utilities under $(v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ and $(v, \tilde{x}, \tilde{p}, b)$, respectively. The following inequalities hold for any $z \in [0,1]$ and directly imply the lemma:

$$u + T(z) \geq (1-\epsilon)u^{\mathrm{SF}} + T(z) \tag{30}$$

$$\geq (1-\epsilon)(u^{\mathrm{SF}} + T(z)) \tag{31}$$

$$\geq (1-\epsilon)\eta v z. \tag{32}$$

Inequality (30) follows from coarse $(1-\epsilon)$-best response. Inequality (31) holds because $T(z) \geq 0$. Inequality (32) follows from the individual efficiency of $(v, \tilde{x}, \tilde{p}, b)$. □

Proof of Lemma 5.12. Let value v and distribution $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ over allocation and payment rules be given, with signal-free projection $(\tilde{x}, \tilde{p})$. We obtain each closure property by applying Lemma 5.13 to best response bid b^* to $(\tilde{x}, \tilde{p})$.

To prove closure under information (Part (i)), let $\{b^\theta\}_\theta$ denote a coarse best response to $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$, and assume the individual efficiency of $\mathcal{O}^* = (v, \tilde{x}, \tilde{p}, b^*)$ is at least η for every target allocation z. Then by Lemma 5.13, the individual efficiency of outcome $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ is at least η for all z. Since this holds for all z, the individual efficiency of $\mathcal{O}$ is at least that of $\mathcal{O}^*$ in the worst case over z as well, proving the desired individual efficiency guarantee.

To prove closure under approximate best response (Part (ii)), let $\{\hat{b}^\theta\}_\theta$ be a coarse $(1-\epsilon)$-best response to $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$, and assume all coarse (exact) best responses have individual efficiency η. Note that the best response bid b^* to $(\tilde{x}, \tilde{p})$ is also a coarse best response to $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$, and hence has individual efficiency η. Applying Lemma 5.13, we conclude that $\{\hat{b}^\theta\}_\theta$ has individual efficiency η as well. □

The subsequent sections analyze the individual efficiency of common families of bidding outcomes. By closure under additional information, it suffices to analyze signal-free bidding outcomes. By closure under approximate best response, it further suffices to analyze only bidding outcomes where the agent best responds. Both observations significantly simplify the analysis.

5.4 Step Function Allocations

In pure-strategy full-information equilibria of a mechanism M, the induced single-agent mechanisms are signal-free. If M further has allocations in $\{0, 1\}$, then the induced single-agent allocation rules will have such allocations as well. Section 2 analyzed a weaker notion of individual efficiency (Definition 2.3) which restricted to target allocation $z = 1$, and showed a individual efficiency of 1. The following lemma extends this result in three ways. First, it generalizes the result to γ-hybrid auctions. Second, applies to the stronger, full definition of individual efficiency (Definition 5.1), which is worst-case over $z \in [0, 1]$. Finally, it extends to allocation rules which may have fractional allocation at an agent's threshold bid. This latter set of allocation rules arises in mechanisms with allocations in $\{0, 1\}$ except for randomized tiebreaking. We formalize the allocation rules we consider below, then state the result.

Definition 5.14. *An allocation rule $\tilde{x}$ is a* step function *if for threshold $\hat{b}$, it takes value $\tilde{x}(b) = 0$ for $b < \hat{b}$ and $\tilde{x}(b) = 1$ for $b > \hat{b}$. The value $\tilde{x}(\hat{b})$ may be arbitrary.*

Lemma 5.15. *For any single-agent signal-free γ-hybrid bidding outcome $\mathcal{O} = (v, \tilde{x}, \tilde{p}, b)$ where $\tilde{x}$ is a step function and b is a best response, the individual efficiency is $\eta = 1$.*

Proof. The argument below lower bounds the individual efficiency, then exhibits a particular example where the bound holds with equality. By Definition 5.14, the allocation rule $\tilde{x}$ steps up to 1 at some $\hat{b} \in [0, \infty)$, but may take any value in $[0, 1]$ at $\hat{b}$. For target allocation z, we have threshold surplus $T(z) = \hat{b}z$. To lower bound u, note that for any $\delta > 0$, the agent could bid $\hat{b} + \delta$ and win. Their γ-hybrid payments for doing so are at most the winner-pays-bid payments, which are $\hat{b} + \delta$. Hence $u(\mathcal{O}) \geq (v - \hat{b} - \delta)$. Taking $\delta \to 0$, we obtain $u + \hat{b} \geq v$. We may further weaken this to obtain the desired inequality for all $z \in [0, 1]$: $u(\mathcal{O}) + \hat{b}z \geq vz$. This bound is the best possible, as it holds with equality when $v = 1$, $z = 1$, $\hat{b} = 0$, and $b = 0$. □

Applying closure under approximate best response (Lemma 5.12, Part (ii)) extends the result of Lemma 5.15 to approximate best response.

Corollary 5.16. *For any single-agent signal-free γ-hybrid bidding outcome $\mathcal{O}$ where $\tilde{x}$ is a step function and b is a $(1 - \epsilon)$-best response, the individual efficiency is at least $1 - \epsilon$.*

Due to the full-information nature of step function rules, there is no appropriate extension of Lemma 5.15 to general information structures. The next section characterizes the individual efficiency under incomplete information.

5.5 Winner-Pays-Bid and Hybrid Mechanisms

We now consider single-agent bidding outcomes with allocations rules taking fractional values representing different probabilities of service. Such rules can arise due to incomplete information over values (or random selection of participants), randomized mechanisms, or randomized environments such as position auction environments. Solutions to these bidding problems are generally less individually efficient

than those with step function allocations. This section characterizes individual efficiency of winner-pays-bid mechanisms under these general single-agent bidding outcomes, proving an individual efficiency bound of $1 - 1/e \approx .63$.

Lemma 5.17. *The individual efficiency of single-agent winner-pays-bid bidding outcomes under coarse best response is $\eta = 1 - 1/e$.*

Proof. We will lower bound the individual efficiency of information-free bidding outcomes, then invoke closure under information (Lemma 5.12, Part (i)). The lower bound uses best response inequalities to upper bound the allocation rule and thus lower bound the threshold surplus in terms of the utility. Figure 2 illustrates the argument graphically. To prove tightness, we exhibit a matching worst-case example.

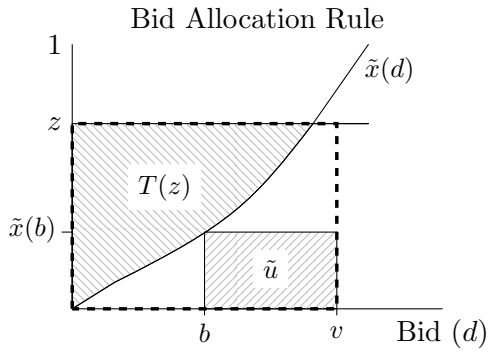
Given signal-free bidding outcome $\mathcal{O} = (v, \tilde{x}, \tilde{p}, b)$, the first step is to rewrite the best response inequalities in terms of $t(x) = \inf\{d \,|\, \tilde{x}(d) \geq x\}$. For brevity, write $u = u(\mathcal{O})$. For an allocation $x \in [0, 1]$, the agent could bid $t(x) + \epsilon$ and get utility at least $(v - t(x) - \epsilon)x$ for any ϵ, and hence the best response utility satisfies $u \geq (v - t(x))x$. Consequently, $u \geq (v - t(x))x$ for all $x \in [0, 1]$, which can be rearranged as $t(x) \geq v - u/x$.

Next, we use the bound on $t(x)$ to lower bound the individual efficiency. Since $t(x) \geq 0$, we may write:

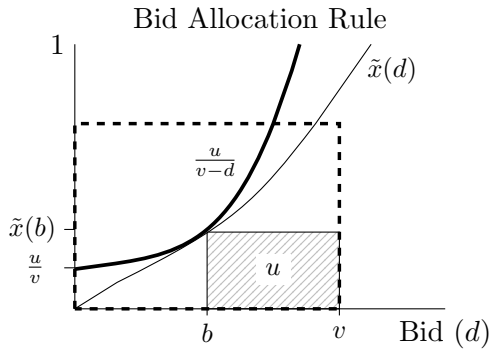
$$\begin{aligned} T(z) \geq \int_0^z \max(v - \tfrac{u}{x}, 0)\, dx &= \int_{u/v}^z v - \tfrac{u}{x}\, dx \\ &= vz - u(1 - \ln \tfrac{u}{vz}). \end{aligned}$$

We therefore have $u + T(z) \geq vz + u \ln(u/vz)$. Holding v fixed and minimizing the righthand side as a function of u yields a worst-case $u = vz/e$, and hence $u + T(z) \geq vz(1 - 1/e)$. This lower bounds individual efficiency by $1 - 1/e$. For a bidding outcome exhibiting the tightness of this bound, take $v = 1$, $z = 1$, allocation rule $\tilde{x}(d) = (e(1 - d))^{-1}$ for $d \in [0, 1 - 1/e]$, and best response bid $b = 0$. □

The lemma below further generalizes this analysis to γ-hybrid mechanisms, of which winner-pays-bid auctions are a special case ($\gamma = 1$). The resulting bound is worst for fully winner-pays-bid mechanisms. As γ decreases, the bound increases, assuming value $\eta = 1$ at $\gamma = 0$.



(a) Individual efficiency analysis for randomized allocation rule $\tilde{x}(d)$. Best response utility u maximizes $\tilde{u}(b) = (v-b)\tilde{x}(b)$ shown as the bottom-right shaded box. Together, u and $T(z)$ approximate the dotted box, of area vz.



(b) The agent's equiutility curve $u/(v-d)$ (dark line) upper bounds their bid-allocation rule, yielding a lower bound on $T(z)$.

Figure 2

Lemma 5.18. *The individual efficiency of single-agent γ-hybrid bidding outcomes under coarse best response is $\eta = 1 - \gamma^{1/(1-\gamma)}$ for $\gamma \in [0, 1)$.*

The proof is deferred to Appendix B. Lemma 5.12, Part (ii) implies a corresponding bound for approximate best response:

Corollary 5.19. *The individual efficiency of single-agent γ-hybrid bidding outcomes under coarse $(1-\epsilon)$-best response is at least $(1-\epsilon)(1-\gamma^{1/(1-\gamma)})$ for $\gamma \in [0, 1)$, and $(1-\epsilon)(1-1/e)$ for $\gamma = 1$.*

5.6 All-Pay Mechanisms

An all-pay mechanism charges the agent $\tilde{p}(b) = b$. This section shows that all-pay bidding outcomes are less individually efficient than other payment formats considered thus far in two concrete senses. First, we show that all-pay bidding outcomes do not have individual efficiency η under best response for any non-trivial $\eta > 0$. We therefore define a weaker version of individual efficiency and show for this version that the individual efficiency of all-pay bidding outcomes is $\eta = 1/2$. In Section 6 we show that this weaker notion of individual efficiency, when combined with competitive

efficiency, is still sufficient to imply welfare guarantees when values are independently distributed. The following lemma implies that all-pay bidding outcomes have trivial individual efficiency:

Lemma 5.20. *There is an information-free all-pay bidding outcome satisfying best response, and with individual efficiency* 0*.*

Proof. We will construct a bidding outcome $\mathcal{O}$ where both the agent's utility is 0 and the threshold surplus $T(z)$ vanishes as the target allocation $z \to 0$. Consider the bidding outcome with $v = 1$, $\tilde{x}(b) = b$, $\tilde{p}(b) = b$, and $b = 0$. Note that for any bid $b \leq 1$, the agent's utility is 0, and hence $b = 0$ is a best response. Now consider target allocation z. The threshold surplus is $T(z) = \int_0^z x\,dx = z^2/2$. Hence, for any $\eta > 0$, there is a z sufficiently small such that $u(\mathcal{O}) + T(z) = z^2/2 < \eta z = \eta v z$, and so the individual efficiency is 0. □

Next, we define a weakening of individual efficiency. In the example of Lemma 5.20, utility and threshold surplus $T(z)$ failed to approximate the target surplus vz when z was small. The weaker version, defined below, only requires this approximation to hold for $z = 1$.

Definition 5.21. *The* weak individual efficiency *of a single-agent bidding outcome is its individual efficiency for $z = 1$. The weak individual efficiency of a family of single-agent bidding outcomes is the smallest weak individual efficiency of any outcome in the family.*

The weak individual efficiency of any bidding outcome is at least its individual efficiency per the stronger Definition 5.8. Under this weaker definition, all-pay bidding outcomes have nontrivial individual efficiency, but are still less individually efficient than γ-hybrid auctions for any choice of γ. Note further that Lemma 5.13 implies closure under information and approximate best response for weak individual efficiency just as it does for ordinary individual efficiency:

Lemma 5.22. *Weak individual efficiency is closed under:*

(i) Additional information: Let $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ denote a joint distribution over allocation and payment rules. For value v, assume that all best responses to the

signal-free projection $(\tilde{x}, \tilde{p})$ *have weak individual efficiency* η*. Then all coarse best responses to* $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ *have weak individual efficiency at least* η*.*

(ii) Approximate best response: Let $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ *denote a joint distribution over allocation and payment rules. For value* v*, assume that all coarse best responses to* $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ *have weak individual efficiency at least* η*. Then all coarse* $(1-\epsilon)$*-best responses to* $(\{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta)$ *have weak individual efficiency at least* $(1-\epsilon)\eta$*.*

We may thus analyze signal-free all-pay bidding outcomes to prove the following:

Lemma 5.23. *The weak individual efficiency of single-agent all-pay bidding outcomes under coarse best response is* $1/2$*.*

Proof. We give an argument lower bounding the weak individual efficiency, then demonstrate tightness of the analysis with an example. Let information-free bidding outcome $(v, \tilde{x}, \tilde{p}, b)$ be given. For any allocation probability $x \in [0,1]$ and any $\epsilon > 0$, the agent could choose to get allocation probability at least x and pay $t(x)+\epsilon$. Hence, $u \geq vx - t(x)$ for all $x \in [0,1]$. We may rearrange this as $t(x) \geq vx - u$. Since we also have $t(x) \geq 0$, we may write:

$$\begin{aligned} T(1) \geq \int_0^1 \max(vx - u, 0)\, dx &= \int_{u/v}^1 vx - u\, dx \\ &= \tfrac{v}{2} - u + \tfrac{u^2}{2v}. \end{aligned}$$

We therefore have $u + T(1) \geq v/2 + u^2/2v \geq v/2$, as desired. To exhibit tightness, take $v = 1$, allocation rule $\tilde{x}(d) = d$ for $d \in [0,1]$, and best response bid $b = 0$. □

Lemma 5.22 implies a corresponding bound for approximate best response:

Lemma 5.24. *The weak individual efficiency of single-agent all-pay bidding outcomes under coarse* $(1-\epsilon)$*-best response is* $(1-\epsilon)/2$*.*

6 Welfare and Revenue Analysis

This section translates competitive efficiency and individual efficiency into robust equilibrium performance guarantees. The results extend the motivation for individual

efficiency in the single-buyer mechanisms of Proposition 5.5. In the single-buyer setting, the seller has an outside option for an item, and individual efficiency implies a welfare guarantee: if the mechanism is individually efficient, the buyer's contribution to welfare can only be small if this outside option's contribution is large. Equilibria of a multi-buyer mechanism induce single-agent bidding problems where the seller's outside option is endogenously generated by other agents. Competitive efficiency then measures the contribution of these endogenous outside options to the seller's utility. These combine to imply a robust welfare guarantee parametrized by the competitive efficiency of the mechanism and the individual efficiency of the induced single-agent bidding outcomes. The reduction from revenue to welfare maximization of Myerson (1981) further implies a revenue guarantee.

Section 6.1 formalizes the connection between equilibria and single-agent mechanisms, and presents the welfare guarantee. In Section 6.2 we extend the result to revenue, assuming independently distributed values and suitably chosen reserve prices. In particular, these sections imply results for winner-pays-bid and hybrid mechanisms. For all-pay mechanisms, we show in Section 6.3 that a similar welfare guarantee holds under independently distributed values, despite the weaker individual efficiency of the all-pay format. Finally, Section 6.4 discusses the tightness of the welfare and revenue guarantees and whether competitive efficiency and individual efficiency are necessary for high robust performance.

6.1 Robust Welfare Analysis

We now formalize the robust analysis of equilibrium welfare. An equilibrium $\mathcal{G}$ is a joint distribution over bid profiles $\mathbf{b}$ and value profiles $\mathbf{v}$. Bids and values may be arbitrarily correlated. For agent i, fix their value v_i. The bid distribution of other agents conditional on v_i induces a distribution over single-agent mechanisms. Depending on their information, i may additionally correlate their bid with this distribution over mechanisms. The resulting single-agent bidding outcome is as follows.

Definition 6.1. *Given equilibrium $\mathcal{G}$ and agent i with value v_i, the* conditional bidding outcome *for i is defined by i's conditional bid distribution and conditional distribution of bid allocation rules. Formally, the distribution of $(\mathbf{v}_{-i}, \mathbf{b}) \sim \mathcal{G} \mid v_i$ induces a joint distribution $\mathcal{D}_i(v_i)$ over bids b_i for i, allocation rules $\tilde{x}_i(\cdot, \mathbf{b}_{-i})$, and payment*

rules $\tilde{p}(\cdot, \mathbf{b}_{-i})$*. The conditional bidding outcome for* i *is* $(v_i, \mathcal{D}_i(v_i))$*.*

The main result of the section is that the individual efficiency of the conditional bidding outcomes and the competitive efficiency of the mechanism together imply a welfare guarantee.

Theorem 6.2. *Let* M *be a mechanism with competitive efficiency* $\mu \leq 1$*, and* $\mathcal{G}$ *an equilibrium for* M *in which the agents' conditional bidding outcomes have individual efficiency* η*. Then the expected welfare in* $\mathcal{G}$ *is a* $\mu\eta$*-approximation to the optimal welfare.*

Proof. The proof has two main steps. First, individual efficiency implies that for any value v_i of agent i the sum of each agent's utility and threshold surplus approximate their interim contribution to the optimal welfare. Second, competitive efficiency implies that threshold surplus is approximated by revenue. Welfare is total utility plus revenue, implying the result.

Formally, 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, and let $u_i(v_i) = \mathbb{E}_{\mathbf{b} \,|\, v_i}[\tilde{u}_i(\mathbf{b})]$ denote agent i's interim expected utility in $\mathcal{G}$. The individual efficiency of i's conditional bidding outcome implies:

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

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

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

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

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

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. □

Using the individual efficiency guarantees of Lemmas 5.15, 5.17, and 5.18, we can instantiate Theorem 6.2. Different equilibrium concepts impose different constraints

on the joint distribution of bids and values. These constraints dictate the relevant individual efficiency guarantee. We first state the relevant equilibrium concept for multi-agent mechanisms. We state the definition in full generality for mechanisms which may have actions that are more complex than single bids.

Definition 6.3. *A joint distribution over values and bids $\mathcal{G}$ satisfies* $(1-\epsilon)$-coarse best response *if for each agent i, value v_i, and deviation action a,*

$$\mathbb{E}_{\mathbf{a}\,|\,v_i}\left[\tilde{u}_i(\mathbf{a})\right] \geq (1-\epsilon)\,\mathbb{E}_{\mathbf{a}_{-i}\,|\,v_i}\left[\tilde{u}_i(a,\mathbf{a}_{-i})\right]. \tag{33}$$

If $\mathcal{G}$ satisfies (33) with $\epsilon = 0$, we say it satisfies coarse best response*.*

Most well-studied equilibrium concepts are special cases of Definition 6.3. For example Bergemann and Morris (2016); Bergemann et al. (2017) impose an information structure on the game, which they model with value-dependent signals for each agent. They assume each agent best responds to their signal. While Bergemann et al. (2017) allow interdependent values, we restrict to private values. However, the set of private-value equilbria we allow is wider, and includes the coarse analog of their Bayes correlated equilibria:

Definition 6.4. *A joint distribution over values and bids $\mathcal{G}$ is a* $((1-\epsilon)$-$)$Bayes coarse correlated equilibrium *if it satisfies ($(1-\epsilon)$-)coarse best response.*

Other concepts which satisfy Definition 6.3 are Bayes-Nash equilibrium and the communication equilibria of Forges (1986). Combining Theorem 6.2 with Lemmas 5.17 and 5.18 yields the following guarantee for exact best response:

Corollary 6.5. *Any winner-pays-bid mechanism with competitive efficiency $\mu \leq 1$ has robust welfare approximation $\mu(1-1/e)$ in Bayes coarse correlated equilibrium. For $\gamma \in [0,1)$, any γ-hybrid mechanism with competitive efficiency $\mu \leq 1$ has robust welfare approximation $\mu(1-\gamma^{1/(1-\gamma)})$ in Bayes coarse correlated equilibrium.*

Stated for single-item auctions, we have the following:

Corollary 6.6. *The single-item first-price auction has robust welfare approximation $1-1/e$ in Bayes coarse correlated equilibrium. For $\gamma \in [0,1)$, the single-item γ-hybrid mechanism has robust welfare approximation $\gamma(1-\gamma^{1/(1-\gamma)})$ in Bayes coarse correlated equilibrium.*

An analogous theorem holds for $(1-\epsilon)$-approximate best response, with an additional factor of $(1-\epsilon)$ loss. Lemma 5.15 further yields an improved bound in pure ϵ-Nash equilibrium, where bids and value are deterministic and satisfy $(1-\epsilon)$-best response.

Corollary 6.7. *Any γ-hybrid mechanism for $\gamma \in [0,1]$ with step function allocation rule and competitive efficiency $\mu \leq 1$ has robust welfare approximation $\mu(1-\epsilon)$ in ϵ-Nash equilibrium. Every single-item highest-bid-wins γ-hybrid auction has robust welfare approximation $(1-\epsilon)$ in ϵ-Nash equilibrium.*

6.2 Robust Revenue Analysis

Individual efficiency quantifies the way a buyer trades off their utility and the seller's utility (via competitive efficiency) against their value. This section extends this tradeoff to the objective of seller revenue under suitable independence conditions on values and bids. With carefully chosen reserve prices, we will obtain a robust revenue approximation of $\mu(1-1/e)/2$ for winner-pays-bid mechanisms with competitive efficiency μ when agents best respond (and more generally $\mu(1-\gamma^{1/(1-\gamma)})/2$ for γ-hybrid mechanisms). We consider equilibria with independently-distributed values satisfying the following additional independence condition on actions, which we state in generality beyond single-bid mechanisms.

Definition 6.8. *An equilibrium $\mathcal{G}$ satisfies* no bidder communication *if for every agent i and value v_i, the action a_i of agent i is independent of $\mathbf{v}_{-i}$ conditioned on v_i.*

Definition 6.8 rules out information structures where, in addition to learning their own values, bidders receive further informative signals about the values of others. Equivalently, this eliminates equilibria where bidders communicate nontrivial information about their values to each other or a mediator, including many Bayes correlated equilibria (Bergemann and Morris, 2016; Bergemann et al., 2017) and communication equilibria (Forges, 1986). It is important to rule these out, as it is possible to construct low-revenue equilbria under these notions of correlation. Definition 6.8 still permits some correlated equilibria, however, namely those with one-way communication from a mediator. In particular, Definition 6.8 permits standard Bayes-Nash

equilibrium, as well as coarse correlated equilibria of the agent-strategic or agent-normal forms of the Bayesian game, and hence learning outcomes under incomplete information (Hartline et al., 2015).

Given an equilibrium $\mathcal{G}$ with independently distributed values and no bidder communication, the analysis of Myerson (1981) implies 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 $\textsc{Rev}(M, \mathcal{G}) = \mathbb{E}_{\mathbf{v},\mathbf{b}}[\sum_i \phi_i(v_i)\tilde{x}_i(\mathbf{b})]$. 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. We will extend the individual efficiency analysis of Lemma 5.17 to reason about an agent's contribution to revenue under monopoly reserves.

We consider only allocation rules which exclude agents below reserves. Consistent with Section 4, we add reserves $\mathbf{r}$ to a mechanism or single-agent allocation rule by setting to 0 all bids $b_i < r_i$ before computing allocations, and we assume a bid of 0 yields allocation 0. For such allocation rules to successfully exclude low types, we require additional assumptions on agents' bids. In particular, we rule out equilibria where agents overbid to beat the reserve, and those where high-valued agents could attain positive utility but bid below the reserve inappropriately. Additionally, we make the standard tie-breaking assumption that agents who are indifferent between bidding below the reserve and above it bid above. These properties hold in winner-pays-bid (and hybrid) mechanisms under many standard solution concepts, including Bayes Nash equilibrium, where they follow from standard best response arguments. The definition below makes these assumptions precise.

Definition 6.9. *An equilibrium $\mathcal{G}$ for mechanism M* respects reserves $\mathbf{r}$ *if for all agents i, $v_i \geq r_i$ if and only if $b_i \geq r_i$. Similarly, a single-agent bidding outcome*

$(v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ *respects reserve* r *if for all* θ, $b^\theta \geq r$ *if and only if* $v \geq r$.

When a bidding outcome respects the monopoly reserve price r^*, we can extend the analysis of individual efficiency to virtual welfare. We trade off an agent's virtual surplus $E_\theta[\phi(v)\tilde{x}^\theta(b^\theta)]$ and threshold surplus with discounted reserve $T^{r^*}(z) = \int_{\tilde{x}(r^*)}^{z} t(x)\,dx$ against their virtual value.

Lemma 6.10. *Let* F *be a regular value distribution with monopoly reserve* r^*. *Further let* $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ *be a single-agent bidding outcome with reserve* r^*. *If* $\mathcal{O}$ *has winner-pays-bid payments, satisfies coarse best response, and respects reserve* r^*, *then for any target allocation* $z \in [0, 1]$,

$$\mathbb{E}_\theta[\phi(v)\tilde{x}^\theta(b^\theta)] + T^{r^*}(z) \geq \eta\phi(v)z, \tag{34}$$

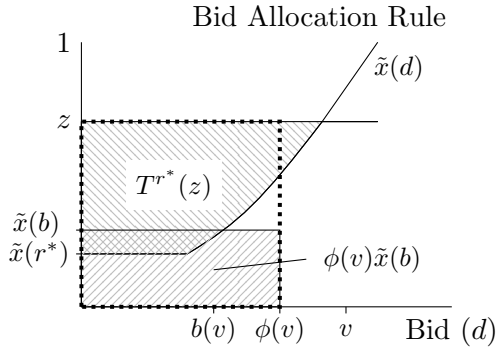
where $\eta = 1 - 1/e$. *If* $\mathcal{O}$ *instead has* γ*-hybrid payments for* $\gamma \in [0, 1)$, *then* (34) *holds with* $\eta = 1 - \gamma^{1/(1-\gamma)}$.

Proof. Since F is regular, $\phi(v) < 0$ if and only if $v < r^*$, in which case the agent bids below r^*. For such agents, $\tilde{x}(b) = 0$, and hence (34) holds trivially. Meanwhile for an agent with value $v \geq r^*$, individual efficiency implies a tradeoff between utility and threshold surplus: $u(\mathcal{O}) + T(z) \geq \eta vz$. The desired result, (34), will follow from this inequality in two steps. First, note that discounting reserves in (34) yields a smaller threshold surplus, $T^{r^*}(z)$. To compensate for this loss, we show below that $\mathbb{E}_\theta[\tilde{p}(b^\theta)] \geq T(z) - T^{r^*}(z)$. The agent's surplus $E_\theta[v\tilde{x}^\theta(b^\theta)]$ is the sum of utility and payments, so we obtain an approximation in terms of this surplus:

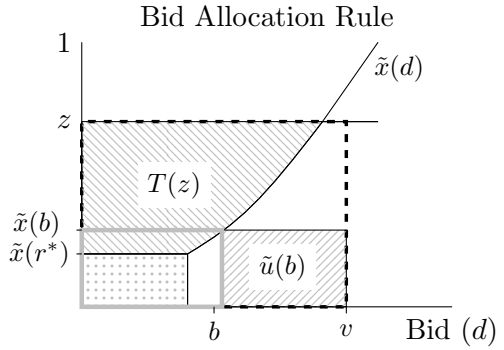
$$\mathbb{E}_\theta[v\tilde{x}^\theta(b^\theta)] + T^{r^*}(z) = u(\mathcal{O}) + \mathbb{E}_\theta[\tilde{p}(b^\theta)] + T^{r^*}(z) \geq u(\mathcal{O}) + T(z) \geq \eta vz. \tag{35}$$

The second step is to note that (35) implies a corresponding approximation for virtual surplus: when $v \geq r^*$, we have $\phi(v) \geq 0$. Moreover, by the definition of virtual values, $\phi(v) \leq v$. Hence, the desired inequality (34) is a weakening of (35). The argument is illustrated geometrically for winner-pays-bid mechanisms in Figure 3.

To complete the proof, we argue that $\mathbb{E}_\theta[\tilde{p}(b^\theta)] \geq T(z) - T^{r^*}(z)$ for any agent with $v \geq r^*$. The argument below holds for winner-pays-bid mechanisms, and we defer the hybrid proof to Appendix A. Figure 3b once again illustrates the winner-pays-bid



(a) Equation (34), illustrated for a deterministic bid $b^\theta = b$. Lemma 6.10 shows the shaded areas cover an $\eta(\gamma)$ fraction of the dashed box, which has area $\phi(v)z$. Note that $T^{r^*}(z)$ omits the area below $\tilde{x}(r^*)$.



(b) Equations (35) and (43), illustrated for deterministic bid $b^\theta = b$. Winner-pays-bid payments $\tilde{p}(b) = b\tilde{x}(b)$ (thick gray rectangle) cover the threshold surplus from the reserve, given by $T(z) - T^{r^*}(z)$ (dotted box). Combining this with individual efficiency of best response, we see that the shaded areas and $\tilde{p}(b)$ together cover a $(e-1)/e$ fraction of dashed box, of area vz.

Figure 3

analysis. For bidding outcomes that respect r^*, bids are at least r^* whenever $v \geq r^*$. We may therefore lower bound winner-pays-bid payments by

$$\mathbb{E}_\theta[\tilde{p}(b^\theta)] = \mathbb{E}_\theta[b^\theta \tilde{x}(b^\theta)] \geq r^* \tilde{x}(r^*) \geq \int_0^{\tilde{x}(r^*)} t(x)\,dx = T(z) - T^{r^*}(z),$$

where the final inequality follows because $t(x)$ is nondecreasing in x. This completes the proof for winner-pays-bid mechanisms. □

We may apply Lemma 6.10 to obtain a revenue guarantee.

Theorem 6.11. *Let $\mathcal{G}$ be a Bayes coarse correlated equilibrium of a mechanism M with monopoly reserves $\mathbf{r}^*$. Assume further that $\mathcal{G}$ satisfies no bidder communication, that values are independently distributed according to regular distributions, and that $\mathcal{G}$ respects reserves $\mathbf{r}^*$. Then if M has winner-pays-bid payments, the expected revenue is a $\mu(1-1/e)/2$-approximation to that of the optimal mechanism, and if M has γ-hybrid payments for $\gamma \in [0,1)$, then the expected revenue is a $\mu(1-\gamma^{1/(1-\gamma)})$-approximation to that of the optimal mechanism.*

Proof. Fix a value v_i for agent i. If $v_i \geq r_i^*$, Lemma 6.10 implies

$$\phi_i(v_i)x_i(v_i) + T_i^{r_i^*}(x_i^*(v_i) \mid v_i) \geq \eta\phi_i(v_i)x_i^*(v_i), \tag{36}$$

where $x_i^*(v_i)$ denotes the interim allocation probability for i under the allocation rule of the revenue-optimal mechanism $\textsc{Opt}_{\mathcal{F}}$ for joint value distribution $\mathcal{F}$, and $x_i(v_i)$ denotes their interim allocation probability in equilibrium. We sum (36) over all agents and take expectations to obtain

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

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

$$\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i \phi_i(v_i)x_i(v_i)\right] + \textsc{Rev}(M^{\mathbf{r}^*}, \mathcal{G}) \geq \mu\eta\,\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i \phi_i(v_i)x_i^*(\mathbf{v})\right].$$

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

$$2\textsc{Rev}(M^{\mathbf{r}^*}, \mathcal{G}) \geq \mu\eta\textsc{Rev}(\textsc{Opt}_{\mathcal{F}}, \mathcal{F}). \qquad \square$$

6.3 Robust Welfare from Weak Individual Efficiency

This section proves robust welfare guarantees for equilibria that only satisfy weak individual efficiency, such as those of all-pay auctions. Under ordinary individual efficiency, each agent's equilibrium utility and threshold surplus approximate their interim contribution to optimal welfare. Optimal interim allocation probabilities are typically fractional, so weak individual efficiency requires a different accounting strategy for optimal welfare. The argument below exploits the independence conditions of the previous section to decompose optimal welfare instead into ex post contributions. In a 0-1 allocation environment $\mathcal{X} \subseteq \{0,1\}^n$, weak individual efficiency then yields a tradeoff with respect to these ex post contributions.

The welfare result for weak individual efficiency requires independently-distributed values and no bidder communication (Definition 6.8). We first spell out the consequences of these assumptions. The first observation is that under these assumptions, bidders' interim allocation rules do not depend on their values. Consequently, threshold surplus is also independent of values. The following notation makes this explicit.

Definition 6.12. *For bid distribution $\mathcal{B}$, let $t_i(\cdot)$ be the inverse of agent i's unconditioned allocation rule $\mathbb{E}_{\mathbf{b}_{-i}}[\tilde{x}_i(\cdot, \mathbf{b}_{-i})]$. Then i's* unconditioned threshold surplus *is $T_i(x) = \int_0^x t_i(z)$.*

With independent values and no bidder communication, the definition of competitive efficiency simplifies. It is possible to further simplify when considering ex post feasible rather than interim feasible allocations. The lemma below captures both simplifications.

Lemma 6.13. *Let M be a mechanism with competitive efficiency μ and $\mathbf{y} \in \mathcal{X}$ be an ex post feasible allocation. Then for any joint distribution $\mathcal{G}$ over bids and values with independently distributed values and no bidder communication,*

$$\textsc{Rev}(M, \mathcal{G}) \geq \mu \sum\nolimits_i T_i(y_i).$$

Proof. Definition 4.2 expresses competitive efficiency in terms of interim-feasible allocations $\hat{\mathbf{y}} = (\hat{y}_1(\cdot), \ldots, \hat{y}_n(\cdot))$, mapping profiles of values to profiles of allocation probabilities. If competitive efficiency holds for all interim-feasible allocations, it must also hold for all constant allocations $\mathbf{y} \in \mathcal{X}$. For such constant $\mathbf{y}$, agent i's threshold surplus $T_i(y_i \,|\, v_i)$ only depends on v_i due to conditioning. Under independent values and no bidder communication, the threshold surplus further satisfies $T_i(y_i \,|\, v_i) = T_i(y_i)$. The definition of competitive efficiency then becomes:

$$\textsc{Rev}(M, \mathcal{G}) \geq \mu \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i(y_i(v_i) \,|\, v_i)\right] = \sum\nolimits_i \mathbb{E}_{v_i}\left[T_i(y_i)\right] = \mu \sum\nolimits_i T_i(y_i). \qquad \square$$

The welfare result for weak individual efficiency is the following.

Theorem 6.14. *Let M be a mechanism for an allocation environment $\mathcal{X} \subseteq \{0,1\}^n$. Assume M has competitive efficiency $\mu \leq 1$. Let $\mathcal{G}$ be an equilibrium with independently distributed values, no bidder communication, nonnegative interim utilities, and in which agents' conditional bidding outcomes have weak individual efficiency η. Then the expected welfare in $\mathcal{G}$ is a $\mu\eta$-approximation to the optimal welfare.*

Proof. Given value profile $\mathbf{v}$, let $\mathbf{x}^*(\mathbf{v}) \in \mathcal{X}$ denote the optimal ex post feasible allocation for $\mathbf{v}$. Since $\mathcal{X} \subseteq \{0,1\}^n$, we have $x_i^*(\mathbf{v}) \in \{0,1\}$ for each agent i. For agents with $x_i^*(\mathbf{v}) = 1$, weak individual efficiency implies:

$$u_i(v_i) + T_i(x_i^*(\mathbf{v})) \geq \eta v_i x_i^*(v_i), \tag{37}$$

where $u_i(v_i) = \mathbb{E}_{\mathbf{b} \mid v_i}[\tilde{u}_i(\mathbf{b})]$ is i's interim utility, which is nonnegative by assumption. For i with $x_i^*(\mathbf{v}) = 0$, inequality (37) holds as well, by the nonnegativity of $u_i(v_i)$. Summing over agents and applying Lemma 6.13, we obtain:

$$\sum\nolimits_i u_i(v_i) + \tfrac{1}{\mu}\textsc{Rev}(M, \mathcal{G}) \geq \eta \sum\nolimits_i v_i x_i^*(v_i),$$

Taking expectations with respect to $\mathbf{v}$ and using the fact that $\mu \leq 1$, we conclude:

$$\mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i u_i(v_i)\right] + \textsc{Rev}(M, \mathcal{G}) \geq \mu\eta\, \mathbb{E}_{\mathbf{v}}\left[\sum\nolimits_i v_i x_i^*(v_i)\right].$$

The first quantity is the expected utility in equilibrium. The last quantity is the expected optimal welfare. Since welfare is utility plus revenue, we obtain the stated approximation guarantee. □

We obtain the following corollary for all-pay mechanisms.

Corollary 6.15. *Let M be an all-pay mechanism for an allocation environment $\mathcal{X} \subseteq \{0,1\}^n$. Assume M has competitive efficiency $\mu \leq 1$. Let $\mathcal{G}$ be a $(1-\epsilon)$-Bayes coarse correlated equilibrium with independently distributed values, no bidder communication, and nonnegative interim utilities. Then the expected welfare in $\mathcal{G}$ is a $\mu(1-\epsilon)/2$-approximation to the optimal welfare.*

6.4 Optimality and Necessity

The earlier results in this section provide sufficient conditions for welfare or revenue in a mechanism to be high in the worst case over broad families of equilibria. We now discuss two senses in which these bounds are often best possible. First, in Section 6.4.1, we discuss whether the bounds derived are quantitatively optimal given our assumptions. We show in particular that the welfare bound for the single-item first-price auction is optimal. Second, in Section 6.4.2, we discuss whether competitive efficiency is necessary for high robust welfare. Dütting and Kesselheim (2015) demonstrate necessity for many natural settings; we briefly overview their results.

6.4.1 First-Price Auctions: Worst-Case Tightness

Given a mechanism (with an associated competitive efficiency) and a family of equilibria, defined by a set of assumptions on information and behavior (with an associated

individual efficiency), this paper's framework outputs a lower bound on the welfare or revenue approximation ratio. Often the performance on individual equilibria in the family is much closer to optimal than the bound produced by the framework. When an equilibrium in the family meets the framework's performance guarantee with equality, it certifies that the guarantee is the best possible for this pairing of mechanism and equilibrium family. For other pairings, it might be possible to improve on the framework's bounds using more tailored arguments. We discuss both cases in more detail for single-item first-price auctions below.

The following example is due to Syrgkanis (2014). It demonstrates the tightness of the $(1 - 1/e)$-approximation of Corollary 6.6. Agents in the example have correlated, private values, and play a Bayes-Nash equilibrium.

Example 6.16. The example has three agents. Agents 1 and 2 have perfectly correlated values, 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 $1 - 1/e$, while the optimal welfare is clearly 1.

Bayes-Nash equilibrium in the first-price auction is notable as a special case of several more general families of mechanisms and equilibria. All the winner-pays-bid mechanism families studied in Section 4 contain the single-item first-price auction. Furthermore, Bayes-Nash equilibrium is a special case of Bayes Correlated and Bayes Coarse Correlated equilibrium (Bergemann and Morris, 2016). Example 6.16 has two consequences for these more general families. First, it implies that the worst-case bound of Theorem 6.2 is tight for all such families. Second, it implies that these general families attain no worse approximations to the optimal welfare than Bayes-Nash equilibria of the first-price auction, despite their greater richness.

It may be possible to impose additional restrictions on equilibria that rule out Example 6.16. The worst-case bounds for these restricted families may improve upon the bounds from Theorem 6.2. For example, when values are independently distributed, Jin and Lu (2022) exploit independence to improve the welfare approximation guarantee for the first-price auction from $1 - 1/e \approx .63$ to $1 - 1/e^2 \approx .865$. They further

exhibit an example where this new bound is tight. As in Example 6.16, the tight example has a single high-valued agent, who should always receive the item in the welfare-optimal allocation, but underbids in equilibrium. Without correlation to coordinate the bid distribution of the low-valued competition, however, Jin and Lu (2022) instead use a large population of identical, independent agents to implement a competing distribution for the high-valued agent.

Finally, tightness of the revenue result for first-price auctions is unknown. Below, we give the worst-known example for revenue in Bayes-Nash equilibrium with monopoly reserves, with independent, regularly-distributed values. Despite being the worst known, the revenue approximation of the example is $1/2$, compared to the bound from Theorem 6.11 of $(1 - 1/e)/2 \approx .31$.

Example 6.17. The equilibrium has two agents. The first agent has value deterministically 1, and the second will have value CDF $F_2(v_2) = 1 - 1/v_2$ on support $[1, \infty)$. 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 $1/2$.

6.4.2 Necessity of High Competitive Efficiency

We have shown that high 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. In settings where this is the case, robust efficiency and competitive efficiency are equivalent design objectives, up to the individual efficiency of the family of equilibria being considered. A formal treatment of necessity appears in Dütting and Kesselheim (2015). For brevity, we simply overview the main ideas.

Recall the analysis of the highest-bids-win mechanism for single-minded combinatorial auctions in Section 2.1. Lemma 2.11 exhibited a bid profile for this mechanism

with low competitive efficiency. The profile had $m+1$ agents, with $\mathcal{S}_i = \{i\}$ for $i \in \{1, \ldots, m\}$ and $\mathcal{S}_{m+1} = \{1, \ldots, m\}$. Ties were broken in favor of agent $m+1$, and the bid profile $\mathbf{b} = (0, \ldots, 0, 1)$ implied a competitive efficiency of $1/m$. Moreover, Lemma 2.12 gave a related Nash equilibrium with low welfare by adding an additional agent, $m+2$, with $\mathcal{S}_{m+2} = \{1, \ldots, m\}$. When all agents had value 1, the bid profile $(0, \ldots, 0, 1, 1)$ was a Nash equilibrium: the presence of agent $m+2$ incentivized the bid of agent $m+1$. The optimal welfare was m and the equilibrium welfare was 1.

This template can be extended to other bundle structures and allocation rules for winner-pays-bid single-minded combinatorial auctions. Given a bid profile $\mathbf{b}$ with low competitive efficiency, we consider adding duplicate bidders to enforce this bid profile in equilibrium. In more detail, consider any winner-pays bid mechanism M which, given two or more bidders with identical desired bundles, allocates to the one with the highest bid. Let bids $\mathbf{b}$ and alternate allocation $\mathbf{y}$ exhibit competitive efficiency at most μ. For each winner i under $\mathbf{b}$ in M, add duplicate i' desiring the same bundle. The following value and bid profiles are a Nash equilibrium with welfare approximation μ in the augmented setting: for each winner i under $\mathbf{b}$, give i and i' value and equilibrium bid equal to b_i. For 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 $\mathbf{y}$.

Dütting and Kesselheim (2015) formalize the above approach for winner-pays-bid mechanisms in 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. This pinpoints the highest-bids-wins mechanism's mismanagement of inter-bidder competition as the source of its worst-case inefficiency. For single-minded combinatorial auctions in particular, they also show that no mechanism has competitive efficiency better than $\sqrt{m}/2$. This implies that up to a constant factor, the greedy

mechanism discussed in Section 4.3 is the optimal winner-pays-bid mechanism, with respect to the objective of robust welfare approximation. More generally, this discussion further justifies the study of competitive efficiency as a design objective in itself.

7 Beyond Single-Bid Mechanisms

Thus far, we have considered only single-bid mechanisms. We now generalize competitive efficiency and individual efficiency to other families of mechanisms, where agents' available actions may be more complex than single bids. The central example of this section is the simultaneous sale of multiple items to single-parameter unit-demand agents: each agent i has value v_i for any single item from a known subset S_i, and each item is sold separately by first-price auction. An agent's action in this composite mechanism is a vector of bids, one per auction. This section develops a framework that accommodates mechanisms with richer action spaces such as these.

The section has two main goals. First, we give generalizations of competitive efficiency and individual efficiency beyond single-bid mechanisms. These generalizations imply robust welfare and revenue guarantees that hold under correlated values and arbitrary information structures, as before. Second, we give a robust welfare and revenue analysis of simultaneous auctions. However, with correlated values, simultaneous first-price auctions are known to have arbitrarily inefficient equilibria (Feldman et al., 2013). Therefore, this section also provides a another definition of competitive efficiency tailored specially to independent value distributions that is closed under simultaneous composition of mechanisms. This will imply robust welfare and revenue guarantees for simultaneous auctions.

Section 7.1 generalizes competitive efficiency, as well as the necessary preliminary definitions. It also defines the weaker notion of competitive efficiency tailored to independent value distributions mentioned above. Section 7.2 then generalizes individual efficiency, and presents the resulting robust welfare and revenue guarantees. Finally, Section 7.3 applies the new framework to simultaneous auctions, which yields robust welfare and revenue guarantees for simultaneous first-price auctions, and robust welfare guarantees for simultaneous all-pay auctions.

7.1 Generalized Competitive Efficiency

Competitive efficiency compares a mechanism's revenue to its threshold surplus, which measures the competition between agents. For single-bid mechanisms, threshold surplus is in turn governed by the agent's threshold bids, or in randomized environments by the agent's tradeoff between bids and allocation. It is thus not defined when agents have richer action spaces than real-valued bids. In this section, we adapt competitive efficiency to more general mechanisms. First, we generalize the notion of threshold surplus, and define a general notion of competitive efficiency. More precisely, we define several threshold surplus functions, each of which is well-defined in general mechanisms, and each of which gives a different way to measure agents' competition. We then define generalized competitive efficiency, which, given a choice of threshold surplus function, compares a mechanism's threshold surplus to its revenue. In addition to the general definition of competitive efficiency, we define a weaker version tailored to independent value distributions, which we require for the simultaneous composition result of Section 7.3.

To generalize threshold surplus, we measure an agent's tradeoff between allocation and the cost of that allocation. For single-bid mechanisms, the cost of allocation was captured by the bid. For mechanisms with richer action spaces, we instead consider a real-valued cost function defined over each agent's actions. Two examples of interest include the action's *price per unit of allocation* and *gross price*, each defined below. In what follows, let $\tilde{x}_i(a_i \mid v_i) = \mathbb{E}_{\mathbf{a}_{-i} \mid v_i}[\tilde{x}_i(a_i, \mathbf{a}_{-i} \mid v_i)]$ denote agent i's interim allocation and $\tilde{p}_i(a_i \mid v_i) = \mathbb{E}_{\mathbf{a}_{-i} \mid v_i}[\tilde{p}_i(a_i, \mathbf{a}_{-i} \mid v_i)]$ their interim payment function for action a_i given value v_i.

Definition 7.1. *Given joint distribution $\mathcal{G}$ over actions and values, the* price per unit *for action a_i, denoted $\beta_i^{PPU}(a_i \mid v_i)$, is given by $\beta_i^{PPU}(a_i \mid v_i) = \tilde{p}_i(a_i \mid v_i)/\tilde{x}_i(a_i \mid v_i)$. The* gross price *of action a_i, denoted $\beta_i^{AP}(a_i \mid v_i)$, is given by $\beta_i^{AP}(a_i \mid v_i) = \tilde{p}_i(a_i \mid v_i)$.*

In general, we will consider cost functions that do not depend on the realized values or bids of other agents, but may depend on the distributions of these quantities, as is the case for price per unit and gross price. Note that in winner-pays-bid mechanisms, holds that $\beta_i^{\mathrm{PPU}}(b_i \mid v_i) = b_i$, and in all-pay mechanisms, $\beta_i^{\mathrm{AP}}(b_i \mid v_i) = b_i$. To define threshold surplus, we then measure an agent's Pareto frontier between a chosen cost

function and allocation probabilities. For each agent i, we further discount costs below a minimum cost r_i. This will enable us to study simultaneous first-price auctions with individualized reserves: if agent i faces a reserve of r_i in each mechanism, this imposes a minimum price per unit of r_i for allocation overall. Definitions 7.2 and 7.3, below, make this precise.

Definition 7.2. *Given joint distribution $\mathcal{G}$ over actions and values, value v_i for agent i, target allocation probability $x \in [0,1]$, and cost function β_i, agent i's* interim threshold cost *is given by $\tau_i(x \mid v_i) = \inf_{a_i : \tilde{x}_i(a_i \mid v_i) \geq x} \beta_i(a_i \mid v_i)$. Given reserve cost r_i, agent i's* interim threshold cost 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 $\tau_i^{r_i}(x \mid v_i) = 0$ otherwise.*

Definition 7.3. *Given joint distribution $\mathcal{G}$ over actions and values, allocation probability $y_i \in [0,1]$, and reserve cost r_i, agent i's* generalized threshold surplus with discounted reserve *for y_i is given by $\mathscr{T}_i^{r_i}(y_i \mid v_i) = \int_0^{y_i} \tau_i^{r_i}(x \mid v_i)\, dx$. When $r_i = 0$, we omit r_i and simply write $\mathscr{T}_i(y_i \mid v_i)$. See Figure 4a for illustration.*

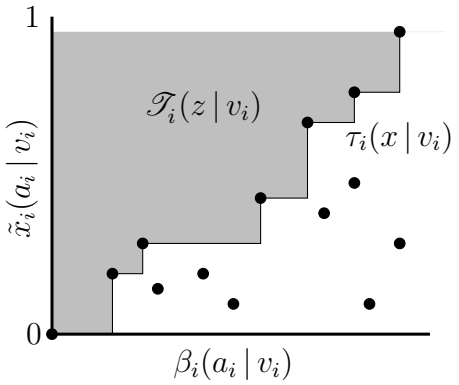
The generalized definition of competitive efficiency compares the threshold surplus to the mechanism's revenue. The definition parallels Definition 4.2, but threshold surplus is now taken with respect to general cost functions.

Definition 7.4. *Let $\mathcal{G}$ be a joint distribution over values and actions, with values following marginal distribution $\mathcal{F}$. For (implicit) cost functions β_i for each agent i, the* generalized competitive efficiency *of a mechanism M and reserve costs $\mathbf{r}$ for $\mathcal{G}$ is the largest μ such that for any feasible profile $\mathbf{y}$ of interim allocation functions,*

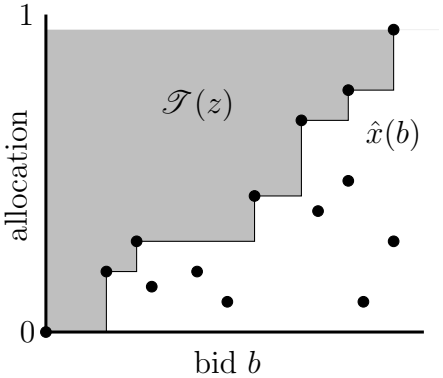
$$\textsc{Rev}(M, \mathcal{G}) \geq \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[\mathscr{T}_i^{r_i}(y_i(v_i) \mid v_i) \right].$$

The competitive efficiency for a family of mechanisms over a family of distributions is the smallest competitive efficiency of any mechanism in the first family over any distribution in the second.

For the mechanisms analyzed in Section 4, we immediately obtain competitive efficiency guarantees under the generalized Definition 7.4. For winner-pays-bid (resp. all-pay) mechanisms, the price per unit (resp. gross price) of a bid is the bid itself. Thus with cost functions β_i^{PPU} (resp. β_i^{AP}), Definitions 7.2, 7.3, and 7.4 coincide



(a) Construction of generalized threshold surplus without reserves. Each point represents some action a_i, and plots cost $\beta_i(a_i \mid v_i)$ against allocation $\tilde{x}_i(a_i \mid v_i)$. The Pareto frontier between cost and allocation is given by $\tau_i(x \mid v_i)$, and the threshold surplus by the area $\mathscr{T}_i(z \mid v_i)$.



(b) Generalized individual efficiency analysis of Lemma 7.10 for price per unit costs $\beta^{\text{PPU}}(a \mid v_i)$. Each action induces a point $(\beta^{\text{PPU}}(a), \tilde{x}(a))$, as in Figure 4a. The Pareto frontier traces a mapping $\hat{x}(\cdot)$ from prices per unit to allocations. This is equivalent to a winner-pays bid mechanism with allocation rule $\hat{x}(\cdot)$.

Figure 4

with their analogs from Section 4. Hence these mechanisms satisfy the generalized definition of competitive efficiency. A generalization of Lemma 4.37 further implies generalized competitive efficiency bounds for hybrid mechanisms with price-per-unit costs. The following lemma formalizes this discussion.

Lemma 7.5. *For winner-pays-bid (resp. γ-hybrid) mechanism M and reserve costs $\mathbf{r}$, if M has competitive efficiency μ, then with price per unit costs, M also has generalized competitive efficiency μ (resp. $\gamma\mu$). For all-pay mechanism M, if M has competitive efficiency μ, then with gross price costs, M also has generalized competitive efficiency μ.*

Proof. We first consider all-pay and winner-pays-bid mechanisms, then extend the winner-pays-bid analysis to hybrid mechanisms. For both winner-pays-bid mechanisms and all-pay mechanisms under price-per-unit and gross price respectively, the cost of a bid is the bid itself. Hence, threshold surplus and generalized threshold surplus are equal, i.e. $T_i^{r_i}(x \mid v_i) = \mathscr{T}_i^{r_i}(x \mid v_i)$ for all x. This implies the desired

generalized competitive efficiency for both winner-pays-bid and all-pay mechanisms:

$$\textsc{Rev}(M, \mathcal{G}) \geq \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[T_i^{r_i}(y_i(v_i) \,|\, v_i) \right] = \mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[\mathscr{T}_i^{r_i}(y_i(v_i) \,|\, v_i) \right].$$

For γ-hybrid mechanism M^γ with allocation rule $\tilde{\mathbf{x}}$, we consider price per unit costs, and compare to the winner-pays-bid mechanism M^{WPB} for $\tilde{\mathbf{x}}$. We superscript relevant quantities by γ or WPB as appropriate. First note that winner-pays-bid payments for bid b_i are always at least the hybrid payments for that same bid. That is, $\tilde{p}_i^{\text{WPB}}(b_i \,|\, v_i) \geq \tilde{p}_i^{\gamma}(b_i \,|\, v_i)$. It follows that the price per unit for b in $\tilde{\mathbf{x}}$ is higher with winner-pays-bid than with hybrid payments, i.e.: $\beta_i^{\text{WPB}}(b_i \,|\, v_i) = \tilde{p}_i^{\text{WPB}}(b_i \,|\, v_i)/\tilde{x}_i(b_i \,|\, v_i) \geq \tilde{p}_i^{\gamma}(b_i \,|\, v_i)/\tilde{x}_i(b_i \,|\, v_i) = \beta_i^{\gamma}(b_i \,|\, v_i)$. Consequently, the same inequality holds for generalized threshold surplus: $\mathscr{T}_i^{r_i,\text{WPB}}(y_i(v_i) \,|\, v_i) \geq \mathscr{T}_i^{r_i,\gamma}(y_i(v_i) \,|\, v_i)$. Meanwhile, note that payments in the γ-hybrid mechanism are at least a γ-fraction of the winner-pays-bid payments, by definition. Hence:

$$\begin{aligned}
\textsc{Rev}(M^\gamma, \mathcal{G}) &\geq \gamma \textsc{Rev}(M^{\text{WPB}}, \mathcal{G}) \\
&\geq \gamma\mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[\mathscr{T}_i^{r_i,\text{WPB}}(y_i(v_i) \,|\, v_i) \right] \\
&\geq \gamma\mu \sum\nolimits_i \mathbb{E}_{v_i \sim \mathcal{F}_i} \left[\mathscr{T}_i^{r_i,\gamma}(y_i(v_i) \,|\, v_i) \right].
\end{aligned}$$

This implies the stated generalized competitive efficiency bound. □

Definition 7.4 allows arbitrarily correlated distributions $\mathcal{G}$ over values and bids. We now present a version of competitive efficiency which is tailored to independent distributions. More precisely, this version is restricted in two ways. First, it quantifies only over distributions with independent values and no bidder communication. For such distributions, an agent's threshold surplus does not depend on their value v_i. We therefore omit the conditioning on v_i, and write $\mathscr{T}_i^{r_i}(x)$ for i's threshold surplus. Second, we consider only ex post feasible allocations in the computation of threshold surplus. The resulting definition is weaker (i.e. easier to satisfy) than Definition 7.4, and parallels Lemma 6.13 for single-bid mechanisms. Formally we have:

Definition 7.6. *Let $\mathcal{G}$ be a joint distribution over values and actions with independently-distributed values and no bidder communication. The* weak competitive efficiency *of a mechanism M for minimum costs $\mathbf{r}$ is the largest μ such that for any ex post feasible allocation $\mathbf{y} \in \mathcal{X}$, the revenue is at least a μ fraction of the threshold surplus:*

$$\textsc{Rev}(M, \mathcal{G}) \geq \mu \sum\nolimits_i \mathscr{T}_i^{r_i}(y_i).$$

We refer to the weak competitive efficiency of a mechanism over all $\mathcal{G}$ with independently-distributed values and no bidder communication simply as its weak competitive efficiency.

The following is then immediate from definitions.

Lemma 7.7. *If a mechanism M has generalized competitive efficiency μ, then it also has weak competitive efficiency μ.*

The analysis in the next two sections focus on weak competitive efficiency and its consequences. Section 7.2 generalizes individual efficiency beyond single-bid mechanisms, and shows how weak competitive efficiency and generalized individual efficiency together imply welfare and revenue guarantees. In Section 7.3 we prove that weak competitive efficiency is preserved when mechanisms are run simultaneously.

7.2 Generalized Individual Efficiency

We next generalize individual efficiency. As in the single-bid framework, individual efficiency measures how well agents trade off their utility against threshold surplus in their single-agent bidding problem. For general mechanisms, we now measure threshold surplus with respect either the to price per unit or gross price of allocation (Definition 7.1). We first define single-agent outcomes and individual efficiency for the general setting. Next, we prove individual efficiency bounds for best response, which now depend on the cost function used to calculate threshold surplus. Finally, we state the robust welfare and revenue guarantees that follow from combining competitive and individual efficiency.

The definitions below generalize individual efficiency. Definition 7.8 first gives the necessary notation for outcomes of general single-agent mechanisms. Definition 7.9 then generalizes ordinary and weak individual efficiency. Both definitions are mechanically identical to their single-bid counterparts, with the addition of considering general cost functions.

Definition 7.8. *A single-agent* outcome $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{a^\theta\}_\theta)$ *is given by a value $v \geq 0$ and a joint distribution over actions a^θ, allocation rules $\tilde{x}^\theta$, and payment rules $\tilde{p}^\theta$, indexed by random variable θ. The agent's utility is given by*

$u(\mathcal{O}) = \mathbb{E}_\theta[v\tilde{x}^\theta(a^\theta) - \tilde{p}^\theta(a^\theta)]$. *The agent's* expected allocation rule *for action* a *is* $\tilde{x}(a) = \mathbb{E}_\theta[\tilde{x}^\theta(a)]$, *and given a cost function* β, *the interim threshold costs are* $\tau(x) = \inf\{\beta(a) \mid \tilde{x}(a) \geq x\}$. *The threshold surplus is* $\mathscr{T}(z) = \int_0^z \tau(x)\,dx$.

Definition 7.9. *The* generalized individual efficiency *of a single-agent outcome* $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{a^\theta\}_\theta)$ *for target allocation* z *is the ratio* $\eta(z) = (u(\mathcal{O}) + \mathscr{T}(z))/vz$. *We call* $\eta = \inf_{z\in[0,1]} \eta(z)$ *the generalized individual efficiency of* $\mathcal{O}$, *and* $\eta(1)$ *the weak generalized individual efficiency. The (weak) generalized individual efficiency of a family of outcomes is the smallest of any outcome in the family.*

We now prove concrete individual efficiency bounds, restricting our analysis to exact best response. Depending on the cost function used to calculate threshold surplus, the quantitative tradeoff between threshold surplus and utility will change. We consider price per unit and gross price, and show that the individual efficiency under these two cost function is equal to the individual efficiency of best response in winner-pays-bid and all-pay mechanisms, respectively. Notably, these guarantees hold *regardless of the payment format of the single-agent mechanism*, as long as the agent best responds. The result follows from identifying, for any single-agent mechanism, a single-agent winner-pays-bid (resp. all-pay) mechanism that is equivalent with respect to both best response utility and price per unit (resp. gross price) threshold surplus. Formally, we have:

Lemma 7.10. *The generalized individual efficiency of outcomes satisfying coarse best response with price per unit costs is* $1 - 1/e$. *The weak generalized individual efficiency of outcomes satisfying coarse best response with gross price costs is* $1/2$.

Proof. We argue the claim for price per unit costs, then explain how to extend the argument to gross price costs. Let $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{a^\theta\}_\theta)$ be a single-agent outcome, with expected allocation rule $\tilde{x}(\cdot) = \mathbb{E}_\theta[\tilde{x}^\theta(\cdot)]$. Then $\tilde{x}(\cdot)$ induces a Pareto frontier between allocation and price per unit: $\hat{x}(b) = \sup_{a\,:\,\beta^{\text{PPU}}(a)\leq b} \tilde{x}(a)$, with inverse $\hat{\tau}(x) = \inf\{b \,:\, \hat{x}(b) \geq x\}$. See Figure 4b for illustration. We may treat $\hat{x}(b)$ and $\hat{p}(b) = b\hat{x}(b)$ as a single-agent winner-pays-bid mechanism and apply the individual efficiency guarantees for single-agent signal-free bidding outcomes. Specifically, let b^* be a single bid maximizing $v\hat{x}(b) - \hat{p}(b)$. Then the single-agent signal-free bidding outcome $(v, \hat{x}, \hat{p}, b^*)$ satisfies best response. Moreover, for any target allocation z,

the winner-pays-bid threshold surplus $\int_0^z \hat{\tau}(x)\,dx$ for $(v, \hat{x}, \hat{p}, b^*)$ is equal to the price per unit surplus $\mathscr{T}(z)$ of $\tilde{x}(\cdot)$, and hence by Lemma 5.17, $v\hat{x}(b^*) - \hat{p}(b^*) + \mathscr{T}(z) \geq vz(1 - 1/e)$. But since the original outcome $(v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{a^\theta\}_\theta)$ satisfies coarse best response, we have $u(\mathcal{O}) \geq v\hat{x}(b^*) - \hat{p}(b^*)$, yielding the stated individual efficiency guarantee. For gross price costs, we apply the same argument to the Pareto frontier between allocation and gross price, and treat this frontier as the allocation rule in an all-pay bidding outcome. The claim then follows from applying coarse best response and the weak individual efficiency of Lemma 5.23 for all-pay bidding outcomes. □

We conclude the section by stating the robust welfare and revenue guarantees implied by competitive efficiency and individual efficiency. We assume independent value distributions and no bidder communication and focus on weak competitive efficiency, so that we may apply these results to simultaneous auctions in the next section. Note however that robust guarantees hold for the general form of competitive efficiency as well. As with the single-bid framework, we study mechanisms satisfying competitive efficiency in equilibria where agents' corresponding single-agent outcomes satisfy individual efficiency. These single-agent outcomes are defined as follows.

Definition 7.11. *Given equilibrium $\mathcal{G}$ and agent i with value v_i, define the* conditional outcome *for i as follows. The distribution of $(\mathbf{v}_{-i}, \mathbf{a}) \sim \mathcal{G} \,|\, v_i$ induces a joint distribution $\mathcal{D}_i(v_i)$ over actions a_i for i, allocation rules $\tilde{x}_i(\cdot, \mathbf{a}_{-i})$, and payment rules $\tilde{p}(\cdot, \mathbf{a}_{-i})$. The conditional outcome for i is $(v_i, \mathcal{D}_i(v_i))$.*

We state the welfare guarantees for the generalized framework below as Theorems 7.12 and 7.13. Theorem 7.12 assumes generalized individual efficiency η, and holds for all mechanisms with weak competitive efficiency μ. Theorem 7.13 assumes the less strict notion of weak generalized individual efficiency, but only holds with 0-1 allocations. Both results require independent values and no bidder communication. They hold by identical arguments to their single-bid analogs, Theorems 6.2 and 6.14.

Theorem 7.12. *Let M be a mechanism with weak competitive efficiency $\mu \leq 1$ for cost function β, and $\mathcal{G}$ an equilibrium for M with independent values and no bidder communication, and in which agents' conditional outcomes have individual efficiency η for β. Then the expected welfare in $\mathcal{G}$ is a $\mu\eta$-approximation to the optimal welfare.*

Theorem 7.13. *Let M be a mechanism for an allocation environment $\mathcal{X} \subseteq \{0,1\}^n$. Assume M has weak competitive efficiency $\mu \leq 1$ for cost function β. Let $\mathcal{G}$ be an equilibrium for M with independent values and no bidder communication, and in which agents' conditional outcomes have weak generalized individual efficiency η for β. Then the expected welfare in $\mathcal{G}$ is a $\mu\eta$-approximation to the optimal welfare.*

Finally, the general framework admits revenue analysis with few modifications. Below, we present the suitable extensions using the price per unit cost function, along with the resulting robust revenue guarantee.

Definition 7.14. *A mechanism M* enforces reserves $\mathbf{r}$ *if for all agents i and actions a_i, $\tilde{x}_i(a_i) = 0$ whenever $\beta_i^{PPU}(a_i) < r_i$. An equilibrium $\mathcal{G}$ for M* respects reserves $\mathbf{r}$ *if for all agents i with value v_i and all actions a_i, $v_i \geq r_i$ if and only if $\beta_i^{PPU}(a_i) \geq r_i$. Define both terms analogously for single-agent outcomes.*

Lemma 7.15. *Let F be a regular value distribution with monopoly reserve r^*, and let $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{a^\theta\}_\theta)$ be a single-agent outcome satisfying coarse best response. If $\mathcal{O}$ enforces and respects r^*, then with price per unit costs, the following holds for all target allocation levels z:*

$$\mathbb{E}_\theta[\phi(v)\tilde{x}^\theta(a^\theta)] + \mathscr{T}^{r^*}(z) \geq (1 - 1/e)\phi(v)z. \tag{38}$$

Theorem 7.16. *Let M be a mechanism with weak competitive efficiency $\mu \leq 1$ with price per unit costs that enforces monopoly reserves $\mathbf{r}^*$. Let $\mathcal{G}$ be a Bayes coarse correlated equilibrium with independently distributed values and no bidder communication, and assume $\mathcal{G}$ respects $\mathbf{r}^*$. Then the expected revenue is a $\mu(1 - 1/e)/2$-approximation to that of the optimal mechanism.*

Proofs of Lemma 7.15 and Theorem 7.16, both omitted, again mimic those of previous results. The proof of Lemma 7.15 parallels that of its winner-pays-bid analog (Lemma 6.10) exactly. The proof of Theorem 7.16 combines those of Theorem 6.11 and Theorem 6.14.

7.3 Simultaneous Composition

This section analyzes the competitive efficiency of simultaneous auctions. We consider a setting with n single-parameter agents seeking an abstract service. Agents

may participate in one or more of m separate mechanisms. The mechanisms are run simultaneously; each agent takes a profile of actions, one per mechanism. Taking allocation levels in $[0, 1]$ to be fractional levels of service, we define an agent's service level in the combined mechanism to be their maximum service level across all mechanisms. Agents make the assigned payments to all mechanisms. We refer to the aggregate mechanism as the *simultaneous composition* of the individual component mechanisms. The main result is that if each individual mechanism has weak competitive efficiency μ, then so too does the aggregate mechanism produced by simultaneous participation as described above. This implies that to understand the robust performance of a larger system, it suffices to understand that of the 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 an allocation rule $\tilde{\mathbf{x}}^j$ and a 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 vector, 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 7.17. *Let mechanisms $M^1, \ldots, M^m$ have allocation and 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 component 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 component mechanism.*

Note that as a consequence of Definition 7.17, we may define the composite feasibility environment as the set of allocation levels induced by the component mechanisms,

i.e.

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

Competitive efficiency holds with respect to cost functions β_i for each i for the composite mechanism. For each component mechanism M^j, we consider projecting β_i down to obtain a cost β_i^j in M^j. This can be done in the following formulaic way. Let a_i^j be an action for agent i in mechanism j, and let $\hat{a}_i^j$ denote the action for agent i in the composite mechanism where i plays a_i^j in mechanism M^j and withdraws from all other mechanisms $j' \neq j$. Now let a_i be an action in the composite mechanism consisting of actions a_i^j in mechanism M^j for each j. Define $\beta_i^j(a_i) = \beta_i(\hat{a}_i^j)$ (where we omit conditioning in $\beta_i(\hat{a}_i^j \,|\, v_i)$ as we will consider only independent value distributions). For the example of price per unit costs β_i^{PPU}, β_i^j corresponds to the price per unit cost in the component mechanism M^j as one would expect. Furthermore, for any allocation $x \in [0,1]$, we may define agent i's threshold cost $\tau_i^j(x) = \inf_{a_i : \tilde{x}_i(a_i) \geq x} \beta_i^j(a_i)$, and similarly define the threshold cost with discounted reserve $\tau_i^{j,r_i}(x)$ and generalized threshold surplus with discounted reserve $\mathscr{T}_i^{j,r_i}(x)$ for mechanism M^j analogously to Definitions 7.2 and 7.3. Define competitive efficiency in the component mechanisms with respect to the projections β_i^j. We can now give the section's main result.

Theorem 7.18. *Let M be the simultaneous composition of mechanisms $M^1, \ldots M^m$. If all of $M^1, \ldots, M^m$ have weak competitive efficiency μ with reserve costs $\mathbf{r}$, then so too does M.*

The theorem follows from two observations. First, revenue in the composite mechanism is the sum of that in the component mechanisms. Competitive efficiency of the component mechanisms thus lower bounds the aggregate revenue in terms of the sum of threshold surpluses of the component mechanisms. Second, agents' threshold costs are lower in the composite mechanisms than in any individual component mechanism. In other words, it is easier for agents to secure allocation with more mechanisms to participate in. We state the second observation as Lemma 7.19 below.

Lemma 7.19. *For any distribution of actions and values $\mathcal{G}$, any $z \in [0,1]$, any component mechanism M^j, and any reserve cost r_i, $\mathscr{T}_i^{r_i}(z) \leq \mathscr{T}_i^{j,r_i}(z)$.*

Proof. By definition, $\tau_i(x) = \inf_{a_i:\tilde{x}_i(a_i)\geq x} \beta(a_i) \leq \inf_{a_i:\tilde{x}_i(a_i)\geq x} \beta_i^j(a_i) = \tau_i^j(x)$, and so $\tau_i^{r_i}(x) \leq \tau_i^{j,r_i}(x)$. Integrating both quantities yields $\mathscr{T}_i^{r_i}(x) \leq \mathscr{T}_i^{j,r_i}(x)$. □

Proof of Theorem 7.18. Let $\mathbf{y}$ be an ex post feasible allocation in the composite environment, where $\mathbf{y}^1 \ldots, \mathbf{y}^m \in \mathcal{X}^1, \ldots, \mathcal{X}^m$ denotes a profile of ex post feasible allocations for each component environment 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. Further let $\textsc{Rev}^j(M, \mathcal{G})$ denote the revenue from component mechanism M^j under $\mathcal{G}$. We obtain the following sequence of inequalities, explained after their statement:

$$\begin{aligned}
\tfrac{1}{\mu}\textsc{Rev}(M, \mathcal{G}) &= \sum\nolimits_j \tfrac{1}{\mu}\textsc{Rev}^j(M, \mathcal{G}) && (39)\\
&\geq \sum\nolimits_j \sum\nolimits_i \mathscr{T}_i^{j,r_i}(y_i^j) && (40)\\
&\geq \sum\nolimits_j \sum\nolimits_i \mathscr{T}_i^{r_i}(y_i^j) && (41)\\
&= \sum\nolimits_j \sum\nolimits_i \delta_{ij} \mathscr{T}_i^{r_i}(y_i)\\
&= \sum\nolimits_i \mathscr{T}_i^{r_i}(y_i).
\end{aligned}$$

Line (39) follows from the definition of payments in a composite mechanism as the sum of the revenues of the component mechanisms. Line (40) comes from the assumption that each component mechanism has weak competitive efficiency μ with minimum costs $\mathbf{r}$. Line (41) follows from Lemma 7.19. The remaining lines follow from the definition of feasibility in the composite mechanism. □

If we consider price per unit costs, then Theorem 7.18 implies that the simultaneous composition of winner-pays-bid and hybrid mechanisms inherits the competitive efficiency of the component mechanisms, and hence their robust welfare guarantees. Moreover, if all mechanisms have monopoly reserves, the same can be said for revenue.

Corollary 7.20. *Let M be the simultaneous composition of winner-pays-bid or hybrid mechanisms, each with weak competitive efficiency at least $\mu \leq 1$. Then in any Bayes coarse correlated equilibrium of M with independent values satisfying no bidder communication, the welfare is at least a $\mu(1 - 1/e)$ fraction of optimal.*

Corollary 7.21. *Let M be the simultaneous composition of winner-pays-bid mechanisms, each with weak competitive efficiency at least $\mu \leq 1$, and each with monopoly*

reserves. Then in any Bayes coarse correlated equilibrium with independent, regularly distributed values, satisfying no bidder communication, and which respects monopoly reserves $\mathbf{r}^*$*, the revenue is at least a* $\mu(1 - 1/e)/2$ *fraction of optimal.*

Theorem 7.13 similarly implies a welfare guarantee for simultaneous all-pay mechanisms.

Corollary 7.22. *Let* M *be the simultaneous composition of all-pay mechanisms for composed feasibility environment* $\mathcal{X} \subseteq \{0,1\}^n$*, each with weak competitive efficiency at least* $\mu \leq 1$*. Then in any Bayes coarse correlated equilibrium with independent values satisfying no bidder communication, the welfare is at least a* $\mu/2$ *fraction of optimal.*

References

Akbarpour, M., Kominers, S. D., Li, S., and Milgrom, P. R. (2021). Investment incentives in near-optimal mechanisms. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 26–26.

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

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.

Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.

Bich, P. and Laraki, R. (2017). On the existence of approximate equilibria and sharing rule solutions in discontinuous games. *Theoretical Economics*, 12(1):79–108.

Bigler, J. (2019). Rolling out first price auctions to google ad manager partners. Google Ad Manager Blog. https://adage.com/article/digital/google-adx-moving-a-price-auction/316894. Accessed: 2024-08-16.

Border, K. (2007). Reduced form auctions revisited. *Economic Theory*, 31(1):167–181.

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.

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

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. (2007). 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.

Forges, F. (1986). An approach to communication equilibria. *Econometrica: Journal of the Econometric Society*, pages 1375–1385.

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

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., Johnsen, A., and Li, Y. (2020). Benchmark design and prior-independent optimization. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 294–305. IEEE.

Hartline, J. and Lucier, B. (2010). Bayesian algorithmic mechanism design. In *ACM Symposium on Theory of Computing*, pages 301–310.

Hartline, J., Syrgkanis, V., and Tardos, E. (2015). No-regret learning in bayesian games. *Advances in Neural Information Processing Systems*, 28.

Hartline, J. D. (2012). Approximation in mechanism design. *American Economic Review*, 102(3):330–336.

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*.

Jin, Y. and Lu, P. (2022). First price auction is 1–1/e2 efficient. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 179–187. IEEE Computer Society.

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.

Kruskal, J. B. (1956). On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50.

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.

Menezes, F. M. and Monteiro, P. K. (2006). Corruption and auctions. *Journal of Mathematical Economics*, 42(1):97–108.

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

Paes Leme, R., Sivan, B., and Teng, Y. (2020). Why do competitive markets converge to first-price auctions? In *Proceedings of The Web Conference 2020*, pages 596–605.

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.

van Beek, A., Brokkelkamp, R., and Schäfer, G. (2022). Corruption in auctions: Social welfare loss in hybrid multi-unit auctions. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 1283–1291.

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

A Proof of Lemma 6.10

This appendix proves Lemma 6.10, which gives an analog to individual efficiency for virtual surplus for winner-pays-bid and hybrid mechanisms. In particular, we prove the hybrid case of the lemma, restated below.

Lemma 6.10. *Let F be a regular value distribution with monopoly reserve r^*. Further let $\mathcal{O} = (v, \{\tilde{x}^\theta\}_\theta, \{\tilde{p}^\theta\}_\theta, \{b^\theta\}_\theta)$ be a single-agent bidding outcome with reserve r^*. If $\mathcal{O}$ has γ-hybrid payments, satisfies coarse best response, and respects reserve r^*, then for any target allocation $z \in [0, 1]$,*

$$\mathbb{E}_\theta[\phi(v)\tilde{x}^\theta(b^\theta)] + T^{r^*}(z) \geq \eta\phi(v)z, \tag{42}$$

with $\eta = 1 - \gamma^{1/(1-\gamma)}$.

Proof. The proof in Section 6.2 establishes the desired inequality (42), assuming that $\mathbb{E}_\theta[\tilde{p}(b^\theta)] \geq T(z) - T^{r^*}(z)$, and shows that this inequality holds for winner-pays-bid payments. All that remains is to prove the same inequality for hybrid mechanisms. To do so, we first prove the inequality for truthful payments. As a consequence, the same inequality will hold for the convex combination of truthful and winner-pays-bid payments. The analysis is as follows:

$$\begin{aligned}
\mathbb{E}_\theta\Big[b^\theta\tilde{x}(b^\theta) - \int_0^{b^\theta} \tilde{x}(b^\dagger)\,db^\dagger\Big] &\geq \mathbb{E}_\theta\Big[r^*\tilde{x}^\theta(r^*) - \int_0^{r^*} \tilde{x}^\theta(b^\dagger)\,db^\dagger\Big] && (43)\\
&= \mathbb{E}_\theta[r^*\tilde{x}^\theta(r^*)] && (44)\\
&= r^*\tilde{x}(r^*) && (45)\\
&= T(z) - T^{r^*}(z), && (46)
\end{aligned}$$

where (43) follows from monotonicity of truthful payments, (44) from the fact that $\tilde{x}^\theta(b^\dagger) = 0$ for all $b^\dagger < r^*$, and (45) and (46) from the definitions of the average allocation rule $\tilde{x}(b) = \mathbb{E}_\theta[\tilde{x}^\theta(b)]$ and threshold surplus, respectively. □

B Proof of Lemma 5.17

This appendix presents the analysis of individual efficiency for γ-hybrid single-agent bidding outcomes for general values of $\gamma \in [0, 1)$, complementing the winner-pays-bid ($\gamma = 1$) analysis in Section 5.5. The bounds we obtain are restated below.

Lemma 5.18. *The individual efficiency of single-agent γ-hybrid bidding outcomes under coarse best response is $\eta = 1 - \gamma^{1/(1-\gamma)}$ for $\gamma \in [0, 1)$.*

By Lemma 5.12, it suffices to analyze signal-free bidding outcomes. To do so, we will express individual efficiency as the solution to an infinite-dimensional linear program, then solve the problem by exhibiting matching primal and dual solutions. For a fixed γ, value v, and target allocation level z, consider the following linear program, denoted $\mathrm{LP}(\gamma, v, z)$, with a decision variable u and variables $t(x)$ for all $x \in [0, 1]$:

$$\begin{aligned} \text{minimize} \quad & u + \int_0^z t(x)\,dx \\ \text{subject to} \quad & u \geq vx - \gamma t(x)x - (1-\gamma)\int_0^x t(y)\,dy && x \in [0,1] \qquad (\mathrm{BR}(x)) \\ & t(x) \geq 0 && x \in [0,1]. \end{aligned}$$

To interpret this program, consider a γ-hybrid bidding outcome with best response utility u, allocation rule $\tilde{x}$, and inverse allocation rule $t(x) = \inf\{b \mid \tilde{x}(b) \geq x\}$ for every $x \in [0, 1]$. We will verify shortly that the pair (u, t) is feasible for $\mathrm{LP}(\gamma, v, z)$, where $\mathrm{BR}(x)$ will hold for each $x \in [0, 1]$ because of the best response inequalities. If we denote the optimal solution to $\mathrm{LP}(\gamma, v, z)$ by $\mathrm{OPT}(\gamma, v, z)$, then it follows that the worst-case individual efficiency for γ-hybrid bidding outcomes is at least $\inf_{v,z} \mathrm{OPT}(\gamma, v, z)/vz$. The individual efficiency guarantee will follow from lower bounding $\mathrm{OPT}(\gamma, v, z)$ for every choice of v and z, then exhibiting a matching single-agent bidding outcome for each v and z.

Verifying Feasibility. Given a γ-hybrid bidding outcome with best response utility u and inverse allocation rule t, we can verify feasibility of (u, t) for $\mathrm{LP}(\gamma, v, z)$ as follows. Nonnegativity of $t(x)$ follows by definition. The constraints $\mathrm{BR}(x)$ for each x will follow from the best response inequalities. In more detail, given $x \in [0, 1]$, if there exists a b such that $\tilde{x}(b) = x$, then $t(x) = b$. The utility for bidding b in this case can be shown to be exactly the righthand side of $\mathrm{BR}(x)$. Otherwise, there is no bid yielding exactly allocation probability x. The argument below verifies that $\mathrm{BR}(x)$ holds even for such x, by showing that there is some bid in the neighborhood of $t(x)$ that yields utility ϵ-close to the righthand side of $\mathrm{BR}(x)$ for any ϵ.

Before making the main argument, note that we may rewrite the righthand side of BR(x) as:

$$\begin{aligned} vx - \gamma t(x)x - (1-\gamma)\int_0^x t(y)\,dy &= (v - t(x))x + (1-\gamma)\int_0^x t(x) - t(y)\,dy \\ &= (v - t(x))x + (1-\gamma)\int_0^{t(x)} \tilde{x}(b^\dagger)\,db^\dagger, \end{aligned} \tag{47}$$

where (47) follows from a change of variables, integrating with respect to bids rather than allocations. Similarly, we may rewrite the utility from a bid b in the hybrid mechanism as:

$$\begin{aligned} \tilde{u}(b) &= \gamma(v-b)\tilde{x}(b) + (1-\gamma)\Big(v\tilde{x}(b) - \Big(b\tilde{x}(b) - \int_0^b \tilde{x}(b^\dagger)\,db^\dagger\Big)\Big) \\ &= (v-b)x + (1-\gamma)\int_0^b \tilde{x}(b^\dagger)\,db^\dagger. \end{aligned} \tag{48}$$

We now consider the utility from bids just above and below $t(x)$. Define $x^- = \lim_{b\to t(x)^-} \tilde{x}(b)$ and $x^+ = \lim_{b\to t(x)^+} \tilde{x}(b)$ to be left- and right-hand limits of the allocation rule at $t(x)$. Recall that we only consider nondecreasing $\tilde{x}$ in hybrid mechanisms, so that there exist dominant-strategy truthful payments. For nondecreasing $\tilde{x}$, the limits x^- and x^+ exist and satisfy $x^- \le x^+$. Moreover, by the definition of $t(x)$, we also have $x \in [x^-, x^+]$. Now consider the utility for bidding just below and just above $t(x)$. Letting $u^- = \lim_{b\to t(x)^-} \tilde{u}(b)$ and $u^+ = \lim_{b\to t(x)^+} \tilde{u}(b)$, we have:

$$\begin{aligned} u^- &= (v - t(x))x^- + (1-\gamma)\int_0^{t(x)} \tilde{x}(b^\dagger)\,db^\dagger \\ u^+ &= (v - t(x))x^+ + (1-\gamma)\int_0^{t(x)} \tilde{x}(b^\dagger)\,db^\dagger. \end{aligned}$$

Let u_x denote the righthand side of (47). Then if $v \ge t(x)$, it holds that $u^- \le u_x \le u^+$. Otherwise, $u^+ \le u_x \le u^-$. Either way, this implies that for any ϵ, there exists a bid b such that $\tilde{u}(b) \ge u_x - \epsilon$. Since this holds for all ϵ, it must be that the best response utility u is at least u_x, and hence BR(x) holds.

Lower Bounding OPT(γ, v, z). To lower bound the objective value of LP(γ, v, z), we exhibit a solution to the dual. Rather than state the dual program, we will describe the dual variables as weights on the constraints, and directly integrate the

constraints to obtain a lower bound. The help define the weights $\psi(x)$ for each $\mathrm{BR}(x)$, let $\underline{x} = z\gamma^{\gamma/(1-\gamma)}$. Then we have:

$$\psi(x) = \begin{cases} 0 & x < \underline{x} \\ \frac{z^{1-1/\gamma}}{\gamma x^{2-1/\gamma}} & \underline{x} \le x < z \\ 0 & z \le x \le 1. \end{cases}$$

We next integrate each term in $\mathrm{BR}(x)$, weighting by $\psi(x)$. The integrals can be computed one by one, starting with the lefthand side:

$$\int_0^1 u\,\psi(x)\,d = \int_{\underline{x}}^z u\,\psi(x)\,dx = u\int_{\underline{x}}^z \frac{z^{1-1/\gamma}}{\gamma x^{2-1/\gamma}}\,dx = u.$$

The first term on the right of $\mathrm{BR}(x)$ integrates to:

$$\int_{\underline{x}}^z vx\,\psi(x)\,dx = \frac{vz^{1-1/\gamma}}{\gamma}\int_{\underline{x}}^z \frac{1}{x^{1-1/\gamma}}\,dx = \frac{vz^{1-1/\gamma}}{\gamma}\gamma(1-\gamma^{1/(1-\gamma)})z^{1/\gamma} = vz(1-\gamma^{1/(1-\gamma)}).$$

Skipping the second term momentarily, we focus on the third term on the right of $\mathrm{BR}(x)$. Let $\Psi(x) = \frac{z^{1-1/\gamma}}{x^{1-1/\gamma}(1-\gamma)}$ denote the antiderivative of ψ. We can integrate the third term as follows:

$$\begin{aligned}
&\int_{\underline{x}}^z (1-\gamma)\int_0^x t(y)\,dy\,\psi(x)\,dx \\
&\qquad = (1-\gamma)\int_{\underline{x}}^z\int_0^x t(y)\,dy\,\psi(x)\,dx \\
&\qquad = (1-\gamma)\Big(\int_0^x t(y)\,dy\Psi(x)\Big|_{x=\underline{x}}^z - \int_{\underline{x}}^z t(x)\Psi(x)\,dx\Big) \qquad (49)\\
&\qquad = (1-\gamma)\Big(\int_0^x t(y)\,dy\frac{z^{1-1/\gamma}}{x^{1-1/\gamma}(1-\gamma)}\Big|_{x=\underline{x}}^z - \int_{\underline{x}}^z t(x)\Psi(x)\,dx\Big) \\
&\qquad = \int_0^z t(x)\,dx - \frac{z^{1-1/\gamma}}{\underline{x}^{1-1/\gamma}}\int_0^{\underline{x}} t(y)\,dy - \int_{\underline{x}}^z t(x)\frac{z^{1-1/\gamma}}{x^{1-1/\gamma}}\,dx, \qquad (50)
\end{aligned}$$

where line (49) follows from integration by parts. We now consider the second term on the right of $\mathrm{BR}(x)$, which simplifies to:

$$\int_{\underline{x}}^z \gamma x t(x)\,\psi(x)\,dx = \int_{\underline{x}}^z \gamma x t(x)\frac{z^{1-1/\gamma}}{\gamma x^{2-1/\gamma}}\,dx = \int_{\underline{x}}^z t(x)\frac{z^{1-1/\gamma}}{x^{1-1/\gamma}}\,dx.$$

Notice that this is equal to the third term in (50). Combining these integrals, we get the following inequality:

$$u \geq vz(1-\gamma^{1/(1-\gamma)}) - \int_0^z t(x)\,dx + \frac{z^{1-1/\gamma}}{\underline{x}^{1-1/\gamma}} \int_0^{\underline{x}} t(y)\,dy.$$

Noting that the final term on the right is nonnegative and rearranging, we get:

$$u + \int_0^z t(x)\,dx \geq vz(1-\gamma^{1/(1-\gamma)}). \tag{51}$$

This implies the stated lower bound on the individual efficiency of γ-hybrid bidding outcomes.

Upper Bounding Individual Efficiency. To show that our lower bound on individual efficiency is the best possible, we exhibit a bidding outcome for each v and z that meets the stated individual efficiency bound exactly. Define the single-agent allocation rule as:

$$\tilde{x}(b) = \begin{cases} \frac{z\gamma^{\gamma/(1-\gamma)}}{(1-b/v)^\gamma} & 0 \leq b < v(1-\gamma^{1/(1-\gamma)}z^{1/\gamma}) \\ 1 & v(1-\gamma^{1/(1-\gamma)}z^{1/\gamma}) \leq b. \end{cases}$$

Again letting $\underline{x} = z\gamma^{\gamma/(1-\gamma)}$, the inverse allocation rule t can be computed as:

$$t(x) = \begin{cases} 0 & 0 \leq x < \underline{x} \\ v(1-\frac{z^{1/\gamma}\gamma^{1/(1-\gamma)}}{x^{1/\gamma}}) & \underline{x} \leq x \leq 1. \end{cases}$$

The agent is indifferent over all bids less than $v(1-\gamma^{1/(1-\gamma)}z^{1/\gamma})$. In particular, from equation (48), we can compute the utility for such a bid as:

$$\begin{aligned} &(v-b)\tilde{x}(b) + (1-\gamma)\int_0^b \tilde{x}(b^\dagger)\,db^\dagger \\ &\quad = (v-b)\frac{z\gamma^{\gamma/(1-\gamma)}}{(1-b/v)^\gamma} + (1-\gamma)\int_0^b \frac{z\gamma^{\gamma/(1-\gamma)}}{(1-b/v)^\gamma}\,db^\dagger \\ &\quad = (v-b)\frac{z\gamma^{\gamma/(1-\gamma)}}{(1-b/v)^\gamma} + (1-\gamma)\frac{z\gamma^{\gamma/(1-\gamma)}v^\gamma}{1-\gamma}((v-b)^{1-\gamma} - v^{1-\gamma}) \\ &\quad = \gamma^{1/(1-\gamma)}vz. \end{aligned} \tag{52}$$

Meanwhile, the threshold surplus is:

$$\begin{aligned}
\int_0^z t(x)\,dx &= \int_{\underline{x}}^z t(x)\,dx \\
&= \int_{\underline{x}}^z v\Big(1 - \frac{z^{1/\gamma}\gamma^{1/(1-\gamma)}}{x^{1/\gamma}}\Big)\,dx \\
&= vx\Big|_{\underline{x}}^z - \frac{vz^{1/\gamma}\gamma^{1/(1-\gamma)}x^{1-1/\gamma}}{1-1/\gamma}\Big|_{\underline{x}}^z \\
&= vx\Big|_{z\gamma^{\gamma/(1-\gamma)}}^z - \frac{vz^{1/\gamma}\gamma^{1/(1-\gamma)}x^{1-1/\gamma}}{1-1/\gamma}\Big|_{z\gamma^{\gamma/(1-\gamma)}}^z \\
&= (1-\gamma^{\gamma/(1-\gamma)})vz - \frac{vz^{1/\gamma}\gamma^{1/(1-\gamma)}x^{1-1/\gamma}}{1-1/\gamma}\Big|_{z\gamma^{\gamma/(1-\gamma)}}^z \\
&= (1-\gamma^{\gamma/(1-\gamma)})vz - \gamma^{1/(1-\gamma)}vz. \qquad (53)
\end{aligned}$$

Combining (52) and (53), we get:

$$u + \int_0^z t(x)\,dx = \gamma^{\gamma/(1-\gamma)}vz + (1-\gamma^{\gamma/(1-\gamma)})vz - \gamma^{1/(1-\gamma)}vz = (1-\gamma^{1/(1-\gamma)})vz.$$

This is matches the lower bound (51) exactly.