Price of Anarchy for Auction Revenue*

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

[1]Northwestern University
[2]Oberlin College

January 16, 2019

Abstract

This paper develops tools for welfare and revenue analyses of Bayes-Nash equilibria in asymmetric auctions with single-dimensional agents. We employ these tools to derive approximation results for social welfare and revenue. Our approach separates the smoothness framework of, e.g., Syrgkanis and Tardos (2013), into two distinct parts, isolating the analysis common to any auction from the analysis specific to a given auction. The first part relates a bidder's contribution to welfare in equilibrium to their contribution to welfare in the optimal auction using the price the bidder faces for additional allocation. Intuitively, either an agent's utility and hence contribution to welfare is high, or the price she has to pay for additional allocation is high relative to her value. We call this condition *value covering*; it holds in every Bayes-Nash equilibrium of any auction. The second part, *revenue covering*, uses the auction rules and feasibility constraints to relate the revenue of the auction to the prices bidders face for additional allocation. Combining the two parts gives approximation results to the optimal welfare, and, under the right conditions, the optimal revenue.

As a centerpiece result, we analyze the single-item first-price auction with individual monopoly reserves (the price that a monopolist would post to sell to that agent alone; these reserves are generally distinct for agents with values drawn from distinct distributions). When each distribution satisfies the regularity condition of Myerson (1981), the auction's revenue is at least a $2e/(e-1) \approx 3.16$ approximation to the revenue of the optimal auction. We also give bounds for matroid auctions with winner-pays-bid or all-pay semantics, the generalized winner-pays-bid position auction, and winner-pays-bid single-minded combinatorial auctions. Finally, we give an extension theorem for simultaneous composition, i.e., when multiple auctions are run simultaneously, with single-valued, unit-demand agents.

*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.

1 Introduction

The first step of a classical microeconomic analysis is to solve for equilibrium. Consequently, such analysis is restricted to settings for which equilibrium is analytically tractable; these settings are often disappointingly idealistic. Worst-case equilibrium analysis (often referred to as "price of anarchy analysis") provides an alternative approach. Instead of solving for equilibrium, properties of equilibrium can be quantified from consequences of best response. These methods have been primarily employed for analyzing social welfare. While welfare is a fundamental economic objective, there are many other properties of economic systems that are important to understand. This paper gives methods for deriving worst-case bounds for auction revenue.

Equilibrium requires that each agent's strategy be a best response to the strategies of others. A typical worst-case approximation analysis obtains a bound on the social welfare (the sum of the revenue and all agent utilities) from a lower bound on an agent's utility implied by best response. Notice that the agents themselves are each directly attempting to optimize a term in the objective. This property makes social welfare special among objectives. Can simple best-response arguments be used to quantify and compare other objectives? This paper considers the objective of revenue, i.e., the sum of the agent payments. Notice that each agent's payment appears negatively in her utility and, therefore, she prefers smaller payments; collectively the agents prefer smaller revenue.

The agenda of this paper parallels a recent trend in mechanism design. Mechanism design typically seeks to identify a mechanism with optimal performance in equilibrium. Optimal mechanisms tend to be complicated and impractical; consequently, a recent branch of mechanism design has sought to quantify the loss between simple mechanisms and optimal mechanisms. These simple (designed) mechanisms have carefully constructed equilibria (typically, the truthtelling equilibrium). The restriction to truthtelling equilibrium, though convenient in theory, is problematic in practice (Ausubel and Milgrom, 2006). In particular, this truthtelling equilibrium is specific to an ideal agent model and tends to be especially non-robust to out-of-model phenomena. The worst-case equilibrium analysis instead considers the performance of simple mechanisms absent a carefully constructed equilibrium.

As an example, consider the single-item first-price auction, in which agents place sealed bids, the auctioneer selects the highest bidder to win, and the winner pays her bid. The fundamental tradeoff faced by the agents in selecting a bidding strategy is that higher bids correspond to a higher probability of winning (which is beneficial) but higher payments (which is detrimental). This first-price auction is a fundamental auction in practice and it is the role of auction theory to understand its performance. When the agents' values for the item are drawn independently and identically, first-price equilibria are well-behaved: the symmetry of the setting enables the easy computation of equilibrium (Krishna, 2009), the equilibrium is unique (Chawla and Hartline, 2013; Lebrun, 2006; Maskin and Riley, 2003), and the highest-valued agent always wins (hence, social welfare is maximized). When the agents' values are non-identically distributed, analytically solving for equilibrium is notoriously difficult. For example, Vickrey (1961) posed the question of solving for equilibrium with two agents with values drawn uniformly from distinct intervals; this problem was finally resolved half a century later by Kaplan and Zamir (2012).

Worst-case approximation analysis allows us to make general statements about behavior in equilibrium without requiring an explicit characterization. For example, Syrgkanis and Tardos (2013) show that the first-price auction's social welfare in equilibrium is at least an $e/(e-1) \approx 1.58$ approximation to the optimal social welfare, and moreover, this bound continues to hold if multiple items are sold simultaneously by independent first-price auctions. Importantly, this analysis sidesteps the intractability of solving for equilibrium and instead derives its bounds from simple best-response arguments.

1.1 Approach

Our analysis comprises two main arguments. The first, *value covering*, encapsulates the welfare consequences of best response. For any given agent, either their expected utility is high, or they are unable to get allocated cheaply because their threshold prices are high. Specifically, we show that for any agent in BNE, the sum of their utilty and expected threshold bid is at least a $\frac{e-1}{e}$-fraction of their value and, therefore, of their contribution to the expected welfare of any mechanism. The second argument, *revenue covering*, captures the mechanism-specific details that affect equilibrium welfare. An auction is revenue covered if whenever allocation is difficult to achieve (i.e. threshold bids are high), this difficulty translates into revenue for the auctioneer. Combining these two properties yields welfare results akin to those proven in Syrgkanis and Tardos (2013).

By decomposing welfare analysis into this modular framework, we are able to extend the arguments to revenue analysis. The Bayes-Nash equilibrium characterization of Myerson (1981) reduces revenue optimization to welfare maximization in the space of "virtual values." We adopt a similar approach: using the value covering argument from our welfare analysis, we show that an agent's positive virtual surplus plus their expected threshold bid approximates their virtual value. If equilibrium revenue is not too badly reduced by negative virtual surplus, this aforementioned *virtual value covering* argument, combined with revenue covering, implies that the first price auction's revenue approximates that of the optimal mechanism. To control the negative virtual surplus from the first-price auction, we use two common methods. First, we extend value and revenue covering to mechanisms with reserves. This analysis shows that adding individual monopoly reserves to the first-price auction prevents the allocation of agents with negative virtual values, which implies a revenue bound. The second method we use begins with an extension of a theorem of Bulow and Klemperer (1996). They argue that with sufficient competition in a second-price auction with i.i.d. agents, the revenue loss from negative virtual surplus is small. We extend this argument to asymmetric first-price and all-pay auctions. As a corollary, we obtain a revenue bound for first-price auctions with sufficient competition.

We present the above approach in Section 3, Section 4, and Section 5 in the context of auctions with winner-pays-bid semantics. In Section 6, we show how to extend the above approach to mechanisms with other payment semantics. We show that value covering is a property of BNE of single-dimensional mechanisms in general. Consequently, all that is required to show approximation results is revenue covering for the mechanism being considered. As an example, we show that the all-pay auction is revenue covered, yielding that in equilibrium the welfare of the all-pay auction approximates the optimal welfare, and that with sufficient competition, the equilibrium revenue approximates the optimal revenue as well. Finally, we show that revenue covering is robust to the simultaneous composition of mechanisms. Consequently, our welfare (and some revenue results) extend to simultaneously run mechansisms as well.

1.2 Results

For single-item and matroid auctions (where the feasibility constraint is given by a matroid set system), we give welfare and revenue results with both winner-pays-bid and all-pay payment formats. The winner-pays-bid variants of these auctions (a) solicit bids, (b) choose an outcome to optimize the sum of the reported bids of served agents, and (c) charge the agents that are served their bids. The all-pay variants of these auctions (a) solicit bids, (b) choose an outcome to optimize the sum of the reported bids of served agents, and (c) charge all agents their bids. Our analyses of these auctions are compatible with reserve prices.

Welfare. In first-price auctions and winner-pays-bid matroid auctions, we derive a welfare approximation bound of $e/(e-1)$. These results also extend to the generalized first-price position auction. For all-pay auctions in the above environments, we use the same proof framework to show a welfare bound of $2e/(e-1)$. The all-pay result is not the best-known bound, but we show how to modify our methods to obtain the best-known welfare bound of 2. While these results do not improve on the best-known bounds, our proofs are notable in that they can be extended to revenue.

Revenue. For winner-pays-bid single-item and matroid auctions with monopoly reserves and regular distributions, we show that the equilibrium revenue is at most a factor of $2e/(e-1)$ from optimal. The same bound holds in the generalized first-price position auction with monopoly reserves. If instead of reserves each bidder must compete with at least one duplicate bidder, the approximation bound for revenue in first-price auctions is at most $3e/(e-1)$; in all-pay auctions, at most 6.

Simultaneous Composition. We also show via an extension theorem that the above welfare bounds (and revenue results for winner-pays-bid auctions with reserves) hold when auctions are run simultaneously if agents are *unit-demand* and *single-valued* across the outcomes of the auctions.

1.3 Related Work

Understanding welfare in games without solving for equilibrium is a central theme in the smooth games framework of Roughgarden (2009) and the smooth mechanisms extension of Syrgkanis and Tardos (2013). A core principle of smoothness is that restricting the arguments used in proving the smoothness property dictate how broadly the result extends. One way to view our work is that we limit our proofs a way that allows for extensions to revenue approximations.

Our framework refines the smoothness framework for Bayesian games in three notable ways. First, we decompose smoothness into two components, separating the the consequences of best-response (value covering) from the specifics of a mechanism (revenue covering). Second, because we focus on the optimization problem that individual bidders are facing, we can attain results that only hold for certain bidders — for instance, bidders with values above their reserve prices. Third, we only consider the Bayesian setting, which allows us to employ the BNE characterization of Myerson (1981) to approximate revenue and to relate other formats of auctions to winner-pays-bid formats via a framework of equivalent bids.

We note two subsequent works with strong connections to our decomposition of smoothness into value covering and revenue covering. First, Dütting and Kesselheim (2015) show that revenue covering, 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 value and revenue covering hold, rather than inferring agents' true values.

A number of papers have derived revenue guarantees for the welfare-optimal Vickrey-Clarke-Groves (VCG) mechanism in asymmetric settings. 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 bid as a reserve, achieves approximately optimal revenue in broader settings. Roughgarden et al. (2012) showed that in broader environments, including matching settings, limiting the supply of items in relation to the number of bidders gives a constant approximation to the optimal auction.

In the economics literature, Kirkegaard (2009) shows that understanding the ratios of expected payoffs in equilibria of asymmetric auctions can lead to insights into equilibrium structure. Kirkegaard (2012) considers properties of distributions on which the revenue of the first price auction exceeds that of the second price auction, and vice versa. Lebrun (2006) and Maskin and Riley (2003) establish equilibrium uniqueness in the asymmetric setting with some assumptions on the distributions of agents.

2 Preliminaries

Bayesian Mechanisms. This paper considers mechanisms for n single-dimensional agents with linear utilities. Each agent has a private value for service, v_i, drawn independently from a distribution F_i over valuation space V_i. We write $\mathbf{F} = \prod_i F_i$ and $\mathbf{V} = \prod_i V_i$ to denote the joint value distribution and space of valuation profiles, respectively, where A_i is the set of possible actions for i. A *mechanism* consists of an action space $\mathbf{A} = \prod_i A_i$, a bid allocation rule $\tilde{\mathbf{x}}$, and a bid payment rule $\tilde{\mathbf{p}}$, mapping actions of agents to probabilistic allocations and payments respectively. Each agent i draws their private value v_i from F_i and selects an action according to some strategy $s_i : V_i \to A_i$. We write $\mathbf{s} = (s_1, \ldots, s_n)$ to denote the vector of agents' strategies. Given the actions $\mathbf{a} = (a_1, \ldots, a_n)$ selected by each agent, the mechanism computes $\tilde{\mathbf{x}}(\mathbf{a}) \in [0, 1]^n$ and $\tilde{\mathbf{p}}(\mathbf{a})$. Each agent's utility is $\tilde{u}_i(\mathbf{a}) = v_i\tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})$.

Typically mechanisms operate with constraints on permissible allocations. Examples include single-item environments, $\tilde{\mathbf{x}}(\mathbf{a})$ must satisfy $\sum_i \tilde{x}_i(\mathbf{a}) \leq 1$ for all actions $\mathbf{a} \in \mathbf{A}$, or matroid environments, where $\tilde{\mathbf{x}}(\mathbf{a})$ must be the membership vector for an underlying matroid (see Section 5.1). We will denote the set of feasible allocations for a given environment by $\mathcal{F}$.

Given a strategy profile $\mathbf{s}$, we often consider the expected allocation and payment an agent faces from choosing some action $a_i \in A_i$, with expectation taken with respect to other agents' values and actions induced by $\mathbf{s}$. We treat $\mathbf{s}$ as implicit and write $\tilde{x}_i(a_i) = \mathbb{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$, with $\tilde{p}_i(a_i)$ and $\tilde{u}_i(a_i)$ defined analogously. Given $\mathbf{s}$ implicitly, we also consider values as inducing payments and allocations. We write $\mathbf{x}(\mathbf{v}) = \tilde{\mathbf{x}}(\mathbf{s}(\mathbf{v}))$ and $\mathbf{p}(\mathbf{v}) = \tilde{\mathbf{p}}(\mathbf{s}(\mathbf{v}))$, respectively. Furthermore, we will denote agent i's interim allocation probability and payment by $x_i(v_i) = \tilde{x}_i(s_i(v_i))$ and $p_i(v_i) = \tilde{p}_i(s_i(v_i))$. We define $u(\mathbf{v})$ and $u_i(v_i)$ similarly. In general, we use a tilde to denote outcomes induced by actions, and omit the tilde when indicating outcomes induced by values. We refer to $\tilde{\mathbf{x}}$ as the *bid allocation rule*, to distinguish it from $\mathbf{x}$, the *allocation rule*. We adopt a similar convention with other notation.

Bayes-Nash Equilibrium. A strategy profile $\mathbf{s}$ is in *Bayes-Nash equilibrium* (BNE) if for all agents i, $s_i(v_i)$ maximizes i's interim utility, taken in expectation with respect to other agents' value distributions $\mathbf{F}_{-i}$ and their actions induced by $\mathbf{s}$. That is, for all i, v_i, and alternative actions a': $\mathbb{E}_{\mathbf{v}_{-i}}[\tilde{u}_i(\mathbf{s}(\mathbf{v}))] \geq \mathbb{E}_{\mathbf{v}_{-i}}[\tilde{u}_i(a', \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$.

Myerson (1981) characterizes the interim allocation and payment rules that arise in BNE. These results are summarized in the following theorem.

Theorem 1 (Myerson, 1981)**.** *For any mechanism with $p_i(0) = 0$ and any value distribution $\mathbf{F}$, BNE implies the following:*

1. *(monotonicity) The interim allocation rule $x_i(v_i)$ for each agent is monotone non-decreasing in v_i.*
2. *(payment identity) The interim payment rule satisfies $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)dz$.*
3. *(revenue equivalence) Mechanisms and equilibria which result in the same interim allocation rule $\mathbf{x}$ must therefore have the same interim payments as well.*

Mechanism Design Objectives. We consider the problem of maximizing two primary objectives in expectation at BNE: utilitarian welfare and revenue. The revenue of a mechanism M is the total payment of all agents. Given a mechanism M and a distribution over action profiles $\mathbf{G}$, the revenue of M under $\mathbf{G}$ is given by $\textsc{Rev}(M, \mathbf{G}) = \mathbb{E}_{\mathbf{a}\sim\mathbf{G}}[\sum_i \tilde{p}_i(\mathbf{a})]$. Alternatively, a strategy profile $\mathbf{s}$ and value distribution $\mathbf{F}$ jointly determine a distribution over action profiles. We may therefore also write the revenue of a mechanism M under $\mathbf{s}$ and $\mathbf{F}$ as $\textsc{Rev}(M, \mathbf{s}, \mathbf{F}) = \mathbb{E}_{\mathbf{v}\sim\mathbf{F}}[\sum_i \tilde{p}_i(\mathbf{s}(\mathbf{v}))]$. The welfare of a mechanism M under a strategy profile $\mathbf{s}$ and value distribution $\mathbf{F}$ is the total utility of all participants including the auctioneer; denoted $\textsc{Welfare}(M, \mathbf{s}, \mathbf{F}) = \textsc{Rev}(M, \mathbf{s}, \mathbf{F}) + \mathbb{E}_{\mathbf{v}\sim\mathbf{F}}[\sum_i \tilde{u}_i(\mathbf{s}(\mathbf{v}))] = \mathbb{E}_{\mathbf{v}\sim\mathbf{F}}[\sum_i v_i \tilde{x}_i(\mathbf{s}(\mathbf{v}))]$ We will also refer to welfare as "surplus." For both welfare and revenue, we will suppress $\mathbf{G}$, $\mathbf{s}$, and $\mathbf{F}$ when context makes the distributions of bids and values clear.

Our welfare benchmark is the outcome that always serves the highest valued feasible agents. That is, we seek to approximate $\textsc{Welfare}(\textsc{Opt}) = \mathbb{E}_{\mathbf{v}}[\max_{\mathbf{x}^*\in\mathcal{F}} \sum_i v_i x_i^*]$. This can be implemented via the Vickrey-Clarke-Groves (VCG) mechanism. We measure a mechanism M's welfare performance by its worst-case approximation ratio, given by

$$\max_{\mathbf{F}\in\text{Indep};\,\mathbf{F},\mathbf{s}\in\text{BNE}(M,\mathbf{F})} \frac{\textsc{Welfare}(\textsc{Opt}, \mathbf{F})}{\textsc{Welfare}(M, \mathbf{s}, \mathbf{F})},$$

where $\text{BNE}(M, \mathbf{F})$ is the set of BNE for M under value distribution $\mathbf{F}$.

For revenue, we will make extensive use of the characterization of revenue in Myerson (1981) that follows from Theorem 1:

Lemma 2. *In any BNE $\mathbf{s}$ for distributions $\mathbf{F}$, the ex ante expected payment of an agent is $\mathbb{E}_{v_i}[p_i(v_i)] = \mathbb{E}_{v_i}[\phi_i(v_i)x_i(v_i)]$, 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. It follows that $\textsc{Rev}(M) = E_{\mathbf{v}}[\sum_i p_i(\mathbf{v})] = \mathbb{E}_{\mathbf{v}}[\sum_i \phi_i(v_i)x_i(\mathbf{v})]$.*

Using this result, Myerson (1981) derives the revenue-optimal mechanism for any value distribution $\mathbf{F}$. This mechanism is parameterized by the value distribution $\mathbf{F}$, and the optimality is in expectation over $\mathbf{v} \sim \mathbf{F}$. We specifically consider distributions with no point masses where $\phi_i(v_i)$ is monotone in v_i for each i. Such distributions are said to be *regular*. If each agent has a regular distribution, then the revenue-optimal mechanism selects the allocation which maximizes $\sum_i \phi_i(v_i)x_i(\mathbf{v})$. For revenue, we will again measure the performance of a mechanism by its worst-case approximation ratio:

$$\max_{\mathbf{F}\in\text{Reg};\,\mathbf{s}\in\text{BNE}(M,\mathbf{F})} \frac{\textsc{Rev}(\textsc{Opt}_{\mathbf{F}}, \mathbf{F})}{\textsc{Rev}(M, \mathbf{s}, \mathbf{F})},$$

where Reg is the set of regular distributions and $\textsc{Opt}_{\mathbf{F}}$ is the Bayesian revenue-optimal mechanism for value distribution $\mathbf{F}$.

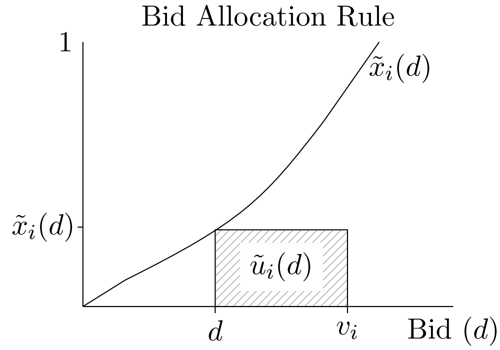

Figure 1: For any bid d (and implicit value v_i), the expected utility $\tilde{u}_i(d)$ is the area of a rectangle between $(d, \tilde{x}_i(d))$ on the bid allocation rule and $(v_i, 0)$. The best-response bid a_i is chosen to maximize this area.

3 Single-Item First Price Auction

We motivate our framework by analyzing the welfare of the first-price auction, showing that it always approximates the welfare of the welfare-optimal mechanism. This result has been known since the work of Syrgkanis and Tardos (2013), but our proof will lend itself to extension and generalization. The rest of the paper will use this proof structure as a template.

Theorem 3. *The welfare in any BNE of the first price auction is at least an $\frac{e}{e-1}$-approximation to the welfare of the welfare-optimal mechanism.*

Our proof proceeds in two steps. First, we analyze the interim optimization problem faced by every bidder. We quantify an intuitively obvious tradeoff: either the bidder can get allocated cheaply, attaining high utility, or allocation is expensive for that bidder to obtain. Second, we note that allocation is only expensive to obtain if the mechanism's revenue is high. This yields a tradeoff between revenue (seller welfare) and utilities (buyer welfare):

$$\sum_i \textsc{Util}_i(\text{FPA}) + \textsc{Rev}(\text{FPA}) \geq \tfrac{e-1}{e}\textsc{Welfare}(\textsc{Opt}). \tag{1}$$

This equation directly implies the theorem.

A Bidder's Optimization Problem: We now develop ideas needed to make this analysis formal. Consider the problem faced by a bidder i with value v_i in the first price auction. Her expected utility for a possible bid d is $\tilde{u}_i(d) = (v_i - d)\tilde{x}_i(d)$, where $\tilde{x}_i(d)$ is the interim bid allocation rule she faces in BNE. If we plot the bid allocation rule $\tilde{x}_i(d)$ for any alternate bid d, then $\tilde{u}_i(d)$ is precisely the area of the rectangle in the lower right; see Figure 1. Let a_i be her best response bid given her value v_i. It must be that a_i maximizes $\tilde{u}_i(d)$ and therefore the area of the rectangle under $\tilde{x}_i(d)$.

When other bidders have realized values and submitted bids, bidder i wins only if her bid exceeds the bids of other players. Consequently the price a bidder must pay to win is $\tau_i(\mathbf{v}_{-i}) = \max_{j \neq i} s_j(v_j)$; we will refer to it as her *threshold bid.* In a Bayesian setting, a bidder reacts not to a deterministic threshold, but rather views $\tau_i(\mathbf{v}_{-i})$ as a random variable, with the bid allocation rule $\tilde{x}_i(\cdot)$ as its CDF.

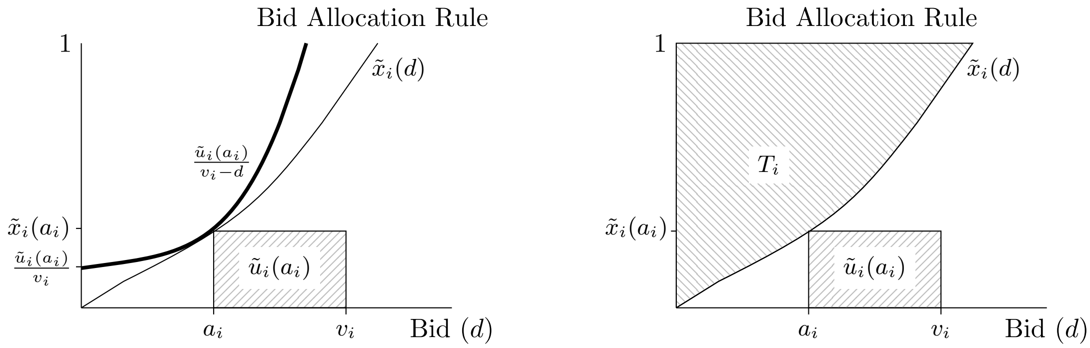

(a) As a_i is a best-response to the actions of other agents, the indifference curve $\tilde{u}_i(a_i)/(v_i - d)$ upper bounds $\tilde{x}_i(d)$.

(b) The expected threshold T_i is the area above the allocation rule.

Figure 2

The expected threshold bid $T_i = \mathbb{E}_{\mathbf{v}_{-i}}[\tau_i(\mathbf{v}_{-i})]$ of an agent is a rough measure of how hard it is for agent i to receive allocation, and will be the focal quantity of our analysis. The expected value of a nonnegative random variable is the area above its CDF - in other words, agent i's expected threshold bid is the area above $\tilde{x}_i(\cdot)$, which is $\int_0^1 1 - \tilde{x}_i(z)\ dz$. It will be convenient to compute this quantity by integrating the inverse of $\tilde{x}_i(\cdot)$, which is given by $t_i(x) = \min\{a \mid \tilde{x}_i(a) \geq x\}$. The inverse allocation function $t_i(x)$ is the amount agent i must bid to ensure allocation probability x. In terms of $t_i(\cdot)$, the expected threshold bid is $T_i = \int_0^1 t_i(z)\ dz$.

Relating Contributions to First-Price and Optimal Welfare: We will now approximate each bidder's contribution to the optimal welfare individually, using the bidder's utility in the first-price auction and a fraction of the revenue in the first-price auction. In these terms, the steps to prove Theorem 3 are:

1. *Value Covering*: Each bidder's utility and expected threshold together approximate her value. (Lemma 4)
2. *Revenue Covering*: The revenue of the first price auction upperbounds the expected thresholds of all agents. (Lemma 5)

Lemma 4 (Value Covering)**.** *In any BNE of the first-price auction, for any bidder i with value v_i,*

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{2}$$

Proof. We will prove value covering in two steps: first, by developing a lower bound $\underline{T}$ on the expected threshold T_i; second, by optimizing to get the worst such bound. The first-price bid deviation approach of Syrgkanis and Tardos (2013) gives the same result.

Lowerbounding T_i. In best responding, bidder i chooses an action which maximizes her utility. Hence for an agent with value v_i and any bid d, her equilibrium utility u_i satisfies $u_i \geq (v_i - d)\tilde{x}_i(d)$. We may write the righthand side in terms of the inverse allocation function $t_i(\cdot)$ to get $u_i \geq (v_i - t_i(x))x$. Rearranging this inequality yields a bound on $t_i(x)$ for any x:

$$t_i(x) \geq v_i - \tfrac{u_i}{x}.$$

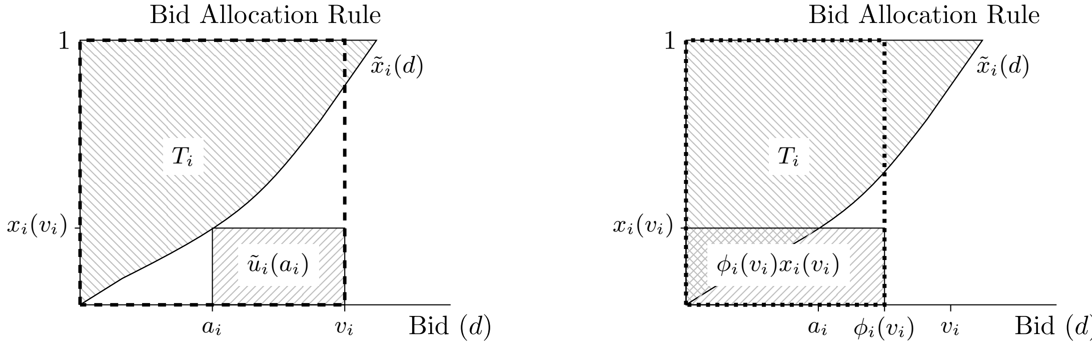

(a) Lemma 4 shows the shaded areas cover a $(e-1)/e$ fraction of the dashed box, bidder i's value v_i and hence maximum contribution to the optimal welfare.

(b) Lemma 6 shows the shaded areas cover an $(e-1)/e$ fraction of i's virtual value $\phi_i(v_i)$.

Figure 3

Note that this bound is meaningful as long as the righthand side is nonnegative - that is, as long as $x \geq u_i/v_i$. Otherwise, note that $t_i(x) \geq 0$. To derive a lower bound on the expected threshold T_i, we use the definition of T_i and integrate $t_i(x)$ over all relevant values of x: from u_i/v_i to 1. Hence:

$$T_i \geq \int_{\frac{u_i}{v_i}}^{1} v_i - \tfrac{u_i}{x}\, dx = \underline{T}_i.$$

Worst-case $\underline{T}_i$. Evaluating the integral for $\underline{T}_i$ gives $\underline{T}_i = v_i - u_i(1 - \ln \frac{u_i}{v_i})$, hence $u_i + \underline{T}_i = v_i + u_i \ln \frac{u_i}{v_i}$. Holding v_i fixed and minimizing with respect to u_i yields a minimum at $u_i = v_i/e$, hence $u_i + \underline{T}_i \geq \frac{e-1}{e} v_i$. The lemma follows. □

Note that this analysis depended only on the fact that bidder i was best responding to a bid distribution. We later will derive a nearly identical condition for every single-dimensional mechanism in BNE using this same logic.

We now show that expected thresholds lowerbound revenue, which will combine with value covering to produce the welfare result. While value covering depended only on equilibrium bidding behavior, revenue covering will only depend on the form of the first price auction, and will thus hold for arbitrary (not necessarily BNE) bidding strategies.

Lemma 5 (Revenue Covering[1])**.** *Fix an arbitrary bid distribution* $\mathbf{G}$ *for the first price auction. For any feasible allocation* $\mathbf{y}$*,*

$$\textsc{Rev}(\text{FPA}, \mathbf{G}) \geq \sum_i T_i y_i. \tag{3}$$

Proof. The revenue of a first price auction is the expected highest bid, and T_i is the expected highest bid from all agents except i. Hence for any agent i, $\textsc{Rev}(\text{FPA}, \mathbf{G}) \geq T_i$. Since single-item feasible allocations sum to at most 1, equation (3) follows. □

We now combine value and revenue covering to approximate the optimal welfare.

[1]To prove our welfare result, it suffices to show $\textsc{Rev}(\text{FPA}) \geq T_i$ for every agent i. We use the more complicated statement to parallel the statement more general feasibility environments.

Proof of Theorem 3. We begin by summing the value covering in equality for each agent in an arbitrary value profile $\mathbf{v}$:

$$\sum_i u_i(v_i) + \sum_i T_i \geq \frac{e-1}{e} \sum_i v_i.$$

Let $\mathbf{x}^*(\mathbf{v})$ be the allocation of the optimal mechanism for $\mathbf{v}$. Since $x_i^*(\mathbf{v}) \leq 1$ for each agent i, and since $u_i(v_i) \geq 0$, we obtain:

$$\sum_i u_i(v_i) + \sum_i T_i x_i^*(\mathbf{v}) \geq \frac{e-1}{e} \sum_i v_i x_i^*(\mathbf{v}).$$

Applying revenue covering and taking expectation with respect to $\mathbf{v}$ shows that $\textsc{Util}(\text{FPA}) + \textsc{Rev}(\text{FPA}) \geq \frac{e-1}{e}\textsc{Welfare}(\textsc{Opt})$. Since welfare is the sum of agent utilities and revenue, the welfare of the first price auction is an $e/(e-1)$ approximation to $\textsc{Opt}$. □

3.1 Welfare Lower Bounds

The approximation results we have given in this section for the single-item first-price auction are not known to be tight. In Appendix A, we present the best-known lower bound, with an approximation factor of 1.15. Note the large gap between this lower bound and the upper bound of $\frac{e}{e-1} \approx 1.58$ from Theorem 3 and Syrgkanis and Tardos (2013).

Beyond a single auction, Christodoulou et al. (2013) have shown that the $\frac{e}{e-1}$ bound is tight for the simultaneous composition of item auctions when bidders have submodular valuations.

4 Revenue Approximation

Our welfare result hinges on the complementary relationship between the utility of a bidder and the bids of other bidders in the mechanism. Using this relationship to directly bound revenue is not as straightforward. The results of Myerson (1981), however, provide another method of accounting for each bidder's impact on revenue, their *virtual value.* Using virtual surplus in place of utilities allows us to adapt our method for proving welfare guarantees to the objective of revenue.

4.1 Revenue

The welfare of a mechanism can be expressed as the expected total value of agents who are served. Myerson (1981) demonstrated a similar characterization of revenue in terms of the expected total virtual value, reducing the problem of revenue maximiation to welfare maximization. We adopt a similar approach, analyzing revenue using tools developed for welfare.

We will begin by showing the analogue of value covering, *virtual value covering*, in which each bidder's positive contribution to equilibrium virtual welfare and expected threshold bid combine to approximate her virtual value, which by Myerson (1981) is upperbounds her contribution to the optimal revenue.

Lemma 6 (Virtual Value Covering)**.** *In any BNE of the first price auction, for any bidder i with value v_i such that $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + T_i \geq \tfrac{e-1}{e}\phi_i(v_i). \tag{4}$$

Proof. First note that surplus is an upper bound on utility, i.e.: $v_i x_i(v_i) \geq u_i(v_i)$. Combined with Lemma 4, this implies that

$$v_i x_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{5}$$

By the definition of virtual value as $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$, we have that $\phi_i(v_i) \leq v_i$. Substituting $\phi_i(v_i)$ for v_i in (5) therefore only weakens the inequality, which implies the result. □

See Figure 3b for an illustration. Intuitively, value covering captures the idea that the expected threshold makes up the difference between an agent's utility and their value. The difference between virtual surplus and virtual value is proportionally smaller, so the expected threshold can cover that gap as well.

Lemma 6 does not immediately imply a revenue approximation result, as virtual value covering only implies an approximation for agents with positive virtual value. A revenue approximation result requires the revenue impact of agents with negative virtual value to be mitigated as well. In Section 4.2, we show that reserve prices suffice for this purpose. In Section 4.3, we prove that sufficient competition also reduces negative virtual surplus and implies an approximation result.

4.2 Reserve Prices

In this section, we show how to adapt the framework of value and revenue covering to accomodate auctions with reserves. The framework was driven by two key arguments. Value covering showed that either an agent was receiving high utility or faced large impediments to obtaining allocation. Revenue covering captured the argument that when the agent could not get allocated easily, the mechanism must be obtaining high revenue. Reserve prices complicate this second argument: an agent with a high reserve might face difficulty winning because of their reserve, which, unlike other agents' bids, does not generate revenue for the mechanism (if the agent loses). We solve this problem below by proving relaxed versions of virtual value and revenue covering. When combined with monopoly reserves (i.e. $r_i = \phi_i^{-1}(0)$) for each agent, they will combine to produce a revenue approximation result, albeit with a slightly larger constant.

The relaxed version of the value and revenue covering framework will use as its pivotal quantity a restricted version of an agent's expected threshold bid. As before, $t_i(x) = \min\{a \mid \tilde{x}_i(a) \geq x\}$ denotes the smallest bid that achieves allocation of at least x. Now, however, note that below $\tilde{x}_i(r_i)$, it no longer corresponds to the inverse cumulative distribution function of the highest bids from all other agents. For $x \leq \tilde{x}_i(r_i)$, $t_i(x) = r_i$ - the threshold comes from the reserve price. Above this point, the threshold comes from the highest bid from other agents, as before. For any bid b, we therefore define the *expected threshold above* b to be $T_i^b = \int_{\tilde{x}_i(b)}^1 t_i(z)\, dz$. With $b = r_i$, the expected threshold above r_i is precisely the portion of the expected threshold generated by bids. See Figure 4a for an illustration.

We now prove weaker notions of value and virtual value covering using $T_i^{r_i}$ instead of T_i. Because $T_i^{r_i} \leq T_i$, value covering as stated in Lemma 4 no longer holds. To solve this problem, we increase the lefthand side by including an agent's expected payments. Meanwhile, revenue covering still holds with $T_i^{r_i}$ instead of T_i. See Figure 4b for an illustration. Formally:

Lemma 7 (Value Covering with Reserves)**.** *In any BNE of* $\text{FPA}_{\mathbf{r}}$*, for any bidder i with value $v_i \geq r_i$,*

$$v_i x_i(v_i) + T_i^{r_i} \geq \tfrac{e-1}{e} v_i. \tag{6}$$

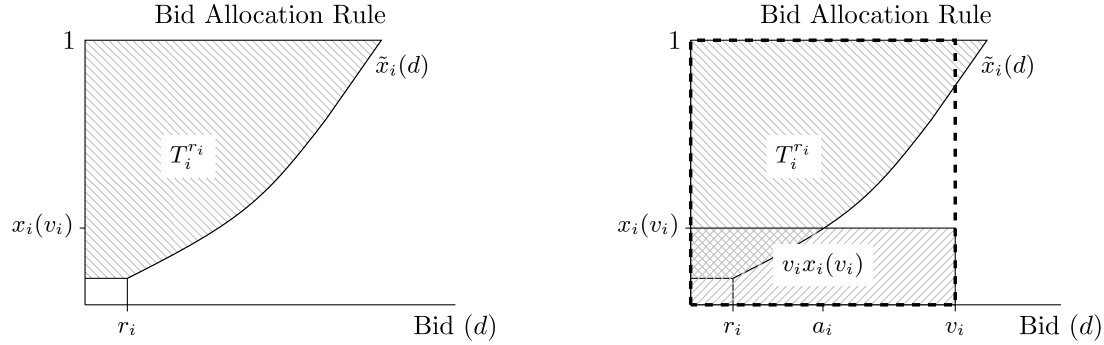

(a) In a first-price auction with reserve r_i, the threshold above r_i, $T_i^{r_i}$, only includes the thresholds when greater than r_i, which corresponds to the case that the threshold is from a bid from another agent rather than the reserve.

(b) Lemma 7 shows that $v_i x_i(v_i)$ and $T_i^{r_i}$ cover an $(e-1)/e$ fraction of i's value v_i, because $v_i x_i(v_i) - u_i(v_i) = a_i x_i(v_i)$ covers $T_i - T_i^{r_i}$

Figure 4

Proof. In BNE, if an agent's value is above the reserve, it is a best response to bid at least the reserve. That is, $a_i(v_i) \geq r_i$. If the reserve price is ever binding for agent i, we have $T_i^{r_i} + p_i(v_i) = T_i^{r_i} + a_i(v_i)x_i(v_i) \geq T_i^{r_i} + r_i\tilde{x}_i(r_i) = T_i$. Otherwise, the reserve never binds, and hence $T_i^{r_i} = T_i$. The result now follows from applying Lemma 4 and the definition of bidder utility. □

Lemma 8 (Virtual Value Covering with Reserves)**.** *In any BNE of* $\mathrm{FPA}_{\mathbf{r}}$*, for any bidder i with value $v_i \geq r_i$ such that $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + T_i^{r_i} \geq \tfrac{e-1}{e}\phi_i(v_i). \tag{7}$$

Proof. Because $v_i \geq \phi_i(v_i)$ for all values v_i, the result follows from Lemma 7. □

Adapting revenue covering to accomodate reserves is simple. As previously mentioned, the thresholds for bidder i above r_i correspond to bids from other agents. It follows that this portion of i's expected threshold corresponds to revenue. We can formalize this with the following lemma:

Lemma 9 (Revenue Covering with Reserves)**.** *Fix an arbitrary bid distribution* $\mathbf{G}$ *of the first price auction with reserves* $\mathbf{r}$*. For any feasible allocation* $\mathbf{y}$*,*

$$\textsc{Rev}(\mathrm{FPA}_{\mathbf{r}}, \mathbf{G}) \geq \sum_i T_i^{r_i} y_i. \tag{8}$$

Proof. For any agent i, recall that $T_i^{r_i} = \int_{\tilde{x}_i(r_i)}^1 t_i(z)\,dz$. That is, $T_i^{r_i}$ is the contribution to i's expected threshold bid from other agents' bids above r_i. Whenever i faces a threshold from such a bid, the mechanism makes at least as much revenue, as it charges the highest bid. Hence for every agent i, $\textsc{Rev}(\mathrm{FPA}_{\mathbf{r}}, \mathbf{G}) \geq T_i^{r_i}$. The result follows. □

With regular value distributions, adding the monopoly reserves $r_i^* = \phi_i^{-1}(0)$ to the auction excludes exactly the agents with negative virtual values. It follows that for such reserves, (7) holds whenever $v_i \geq r_i^*$. Moreover, the optimal mechanism for revenue allocates the item to the agent with the highest positive virtual value. To approximate the optimal revenue, it therefore suffices to approximate this agent's expected virtual value. By adapting revenue covering and virtual value covering to the first price auction with reserves, we are able to treat this quantity much as we treated welfare, yielding the following:

Theorem 10. *The revenue in any BNE of the first price auction with monopoly reserves and agents with regularly distributed values is at least a $2e/(e-1)$-approximation to revenue of the optimal auction.*

Proof. We begin by summing inequality (7) for each agent with $v_i \geq r_i^*$ in an arbitrary value profile $\mathbf{v}$:

$$\sum_{i:v_i \geq r_i^*} \phi_i(v_i)x_i(v_i) + \sum_{i:v_i \geq r_i^*} T_i^{r_i} \geq \frac{e-1}{e} \sum_{i:v_i \geq r_i^*} \phi_i(v_i).$$

Let $\mathbf{x}^*(\mathbf{v})$ be the allocation of the optimal mechanism on $\mathbf{v}$. Since $x_i^*(\mathbf{v}) \leq 1$ for each agent i, and since $\phi_i(v_i)x_i(v_i) \geq 0$, we obtain:

$$\sum_{i:v_i \geq r_i^*} \phi_i(v_i)x_i(v_i) + \sum_{i:v_i \geq r_i^*} T_i^{r_i} x_i^*(\mathbf{v}) \geq \frac{e-1}{e} \sum_{i:v_i \geq r_i^*} \phi_i(v_i)x_i^*(\mathbf{v}). \tag{9}$$

Both the first-price auction with monopoly reserves and the optimal auction exclude agents with $v_i < r_i^*$. Taking expectations of (9), the first term is the expected revenue of the first-price auction with monopoly reserves $\textsc{Rev}(\text{FPA}_{\mathbf{r}^*})$, and the sum on the right-hand side is the optimal revenue, $\textsc{Rev}(\textsc{Opt})$. Applying Lemma 9 to the second term on the left-hand side, therefore, yields $2\textsc{Rev}(\text{FPA}_{\mathbf{r}^*}) \geq \frac{e-1}{e}\textsc{Rev}(\textsc{Opt})$, as desired. □

4.3 Duplicate bidders

Another approach to mitigating the impact of negative virtual-valued agents is to ensure each agent faces adequate competition. Bulow and Klemperer (1996) show that this intuition guarantees approximately optimal revenue in regular, symmetric, single-item settings. In particular, their results can be interpreted as showing that the second price auction with any reserve, even one which allows agents with negative virtual values to be allocated, cannot have its revenue too badly diminished by the contributions of low-valued agents.

Formally, for any mechanism M, let $\textsc{Rev}^+(M) = \sum_i \mathbb{E}_{v_i}[\max(0, \phi_i(v_i))x_i(v_i)]$ denote the expected positive virtual surplus of M. Given a symmetric randomized reserve distribution R, let SPA_R denote the second price auction with reserve R. A simple reinterpretation of Bulow and Klemperer (1996) shows the following:

Theorem 11 (Bulow and Klemperer, 1996)**.** *For any symmetric randomized reserve with distribution R and any single-item environment with n i.i.d. regular bidders, the following inequality holds:*

$$\textsc{Rev}(\text{SPA}_R) \geq \tfrac{n-1}{n}\textsc{Rev}^+(\text{SPA}_R). \tag{10}$$

We show the same intuition holds for first-price and all-pay auctions in asymmetric settings: if each bidder must compete with at least $k-1$ other bidders with values drawn from her same distribution, revenue is approximately optimal compared to the revenue-optimal mechanism (including the duplicate bidders). We say such a setting satisfies *k-duplicates.* Formally:

Definition 12. *A single-item environment satisfies* k-duplicates *if the set of agents can be partitioned into groups $B_1, \ldots, B_p$ for some positive integer p such that $|B_j| \geq k$ and the agents in B_j are identically distributed, for each j in $\{1, \ldots, p\}$.*

We will generalize the analysis of Bulow and Klemperer (1996) to the first-price auction (FPA) and all-pay auction (APA) with k-duplicates. The first-price analysis will combine with the value and revenue covering framework to produce a revenue approximation result. Moreover, in Section 6.2, we extend the framework to include the all-pay auction, which will yield a revenue result for that mechanism as well.

Lemma 13. *In any single-item setting with k-duplicates and regular value distributions, the following inequalities hold:*

$$\text{Rev}(\text{FPA}) \geq \tfrac{k-1}{k}\text{Rev}^{+}(\text{FPA})$$
$$\text{Rev}(\text{APA}) \geq \tfrac{k-1}{k}\text{Rev}^{+}(\text{APA})$$

The proof reduces analyzing the allocation rule for each group of duplicates to analyzing that of a second-price auction with a randomized reserve generated by bidders outside the group. We may then apply Theorem 11 and sum the virtual surplus from the different groups. The full proof is included in Appendix C.

We can combine Lemma 13 with revenue covering and value covering to derive a revenue bound for BNE of the first price auction, which we state below:

Theorem 14. *In any single-item environment with k-duplicates and regular value distributions, the revenue in any BNE of the first price auction (*FPA*) is at least a $\frac{k}{k-1}\frac{2e}{e-1}$ approximation to the revenue of the optimal auction.*

The proof is included in Appendix C. We discuss all-pay auctions in Section 6.2 and derive similar revenue bounds.

4.4 Revenue Lower Bounds

For revenue, the approximation ratio of the first-price auction with monopoly reserves can be at least as bad as 2. The same result was shown by Hartline and Roughgarden (2009) for the second-price auction with monopoly reserves with the following two-agent example. One agent has a deterministic value of 1, the other agent has value drawn according to the equal revenue distribution with support over $[1, H]$ for some large H, with a light perturbation of the distribution so the monopoly price is 1. Assuming ties go to bidder 2, an equilibrium exists where both players bid 1, giving revenue of 1. The optimal auction however can set a reserve of H for the second bidder and sell to the first bidder at price 1 if the reserve is not met, achieving a revenue of 2 as H grows.

5 General Winner-Pays-Bid Mechanism

We now extend the framework developed in Sections 3 and 4 to more general feasibility environments, including matroids and single-minded combinatorial auctions. We will use our framework to prove welfare and revenue results for the analogues of the first-price auction, namely highest-bids-win winner-pays-bid mechanisms. As we will see, most of the proofs from the previous sections extend with little to no modification.

Two main ideas drove the single-item welfare (Theorem 3) and revenue (Theorem 10) results. The first idea, value covering (resp. virtual value covering), captured the tradeoff between an agent's threshold bid and their utility (resp. virtual surplus). This idea depends only on a bidder's interim optimization problem, which is the same in every winner-pays-bid mechanism. We can extend the single-item proof to get:

Lemma 15 (Pay-Your-Bid Value Covering)**.** *In any BNE of any winner-pays-bid mechanism, for any bidder i with value v_i,*

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{11}$$

The second key ingredient is revenue covering, which captured the correspondence between threshold bids and mechanism revenue. Revenue covering is a mechanism-specific property, and must hold under every distribution of bids. In what follows, we will use a parameterized version of revenue covering in which the revenue needs only approximate the threshold bids.

Definition 16 (μ-Revenue Covering)**.** *A mechanism M is μ-*revenue covered *if for any bid distribution $\mathbf{G}$ and feasible allocation $\mathbf{y}$,*

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

Together, revenue covering and value covering provide a framework for deriving approximation results. We can combine revenue and value covering using the logic that drove the proof of Theorem 3. Summing inequality (11) over all bidders, applying revenue covering, and taking expectations with respect to $\mathbf{v}$ yields the following:

Theorem 17. *The welfare of any μ-revenue covered winner-pays-bid mechanism is a $\mu\frac{e}{e-1}$-approximation to the welfare of any other mechanism.*[2]

The value and virtual value covering results for single-item auctions with reserves also hold without modification in general winner-pays-bid auctions. They require only value covering, winner-pays-bid semantics, and Myerson's virtual value characterization, all of which are agnostic to feasibility constraints. For example:

Lemma 18. *In any BNE of any winner-pays-bid auction with reserves $\mathbf{r}$, for any bidder i with value $v_i \geq r_i$ and $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + T_i^{r_i} \geq \tfrac{e-1}{e}\phi_i(v_i).$$

Revenue analysis requires a way to translate thresholds into revenue. Revenue covering with reserves extends beyond single-item environments in the natural way. Specifically:

Definition 19 (Revenue Covering with Reserves)**.** *Given a vector of reserves $\mathbf{r}$, a mechanism M is μ-revenue covered with reserves $\mathbf{r}$ if for any bid distribution $\mathbf{G}$ and feasible allocation $\mathbf{y}$,*

$$\mu\text{Rev}(M^{\mathbf{r}}, \mathbf{G}) \geq \sum_i T_i^{r_i} y_i. \tag{12}$$

If a winner-pays-bid mechanism is μ-revenue covered without reserves, then the addition of any vector of reserves produces a mechanism which is μ-revenue covered with those reserves. Formally:

Lemma 20. *If a winner-pays-bid mechanism M is μ-revenue covered, then adding any vector of individual reserves $\mathbf{r}$ produces a mechanism $M^{\mathbf{r}}$ that is μ-revenue covered with reserves $\mathbf{r}$.*

[2] Note that a tighter analysis of value-covering making use of the parameter μ can give a bound of $\frac{\mu}{1-e^{-\mu}}$ (see Syrgkanis and Tardos (2013)). We do not include the analysis because it does not extend to results for revenue.

Proof. Consider an arbitrary action profile $\mathbf{a}$, and fix a vector of reserves $\mathbf{r}$. In what follows, let $\tau_i(\cdot)$ denote the threshold of agent i under M with no reserves. Now consider constructing an alternate bid profile $\mathbf{a}'$ by setting the bid of any agent with $a_i < r_i$ to 0. Since $\mathbf{a}'$ is a degenerate bid distribution, revenue covering implies:

$$\mu \textsc{Rev}(M, \mathbf{a}') \geq \sum\nolimits_i \tau_i(\mathbf{a}'_{-i}) y_i.$$

Note that the revenue of M under $\mathbf{a}'$ is the total bid of the bid-maximizing set of feasible agents whose bids exceed their reserves. In other words, it is exactly the revenue of $M^{\mathbf{r}}$ under $\mathbf{a}$. Hence we have, for any bid distribution $\mathbf{G}$:

$$\mathbb{E}_{\mathbf{a}\sim\mathbf{G}}\left[\textsc{Rev}(M, \mathbf{a}')\right] = \textsc{Rev}(M^{\mathbf{r}}, \mathbf{G}).$$

Next, consider the quantity $\mathbb{E}_{\mathbf{a}'_{-i}\sim\mathbf{G}_{-i}}[\tau_i(\mathbf{a}'_{-i})]$. This is the expected threshold of agent i after the reserves of every other player have been applied, but with a reserve of 0 for i. In other words, $\mathbb{E}_{\mathbf{a}'_{-i}\sim\mathbf{G}_{-i}}[\tau_i(\mathbf{a}'_{-i})]$ is the expected threshold bid of agent i in $M^{(0,\mathbf{r}_{-i})}$ under action distribution $\mathbf{G}$. Note that considering only the contribution from threshold bids above r_i, i.e. $\mathbb{E}_{\mathbf{a}'_{-i}\sim\mathbf{G}_{-i}}[\tau_i(\mathbf{a}'_{-i})\mathbb{1}_{\tau_i(\mathbf{a}'_{-i})>r_i}]$, produces a smaller quantity. Moreover, note that this latter quantity, the contribution to the i's expected threshold of thresholds above r_i in $M^{(0,\mathbf{r}_{-i})}$, does not change if we consider $M^{\mathbf{r}}$ instead of $M^{(0,\mathbf{r}_{-i})}$. This latter quantity is $T_i^{r_i}$. Hence for any alternate allocation $\mathbf{y}$, we have:

$$\begin{aligned}
\mu \textsc{Rev}(M^{\mathbf{r}}, \mathbf{G}) &= \mu\, \mathbb{E}_{\mathbf{a}\sim\mathbf{G}}\left[\textsc{Rev}(M, \mathbf{a}')\right] \\
&\geq \sum\nolimits_i \mathbb{E}_{\mathbf{a}'_{-i}\sim\mathbf{G}_{-i}}\left[\tau_i(\mathbf{a}'_{-i})\right] y_i \\
&\geq \sum\nolimits_i \mathbb{E}_{\mathbf{a}'_{-i}\sim\mathbf{G}_{-i}}\left[\tau_i(\mathbf{a}'_{-i})\mathbb{1}_{\tau_i(\mathbf{a}'_{-i})>r_i}\right] y_i \\
&= \sum\nolimits_i T_i^{r_i} y_i.
\end{aligned}$$

□

Once μ-revenue covering with monopoly reserves has been established, it is a matter of extending the logic used to prove Theorem 10 to produce a more general revenue result. We can sum, apply revenue covering, and take expectations to obtain:

Theorem 21. *The revenue of any μ-revenue covered winner-pays-bid mechanism with regular bidders and monopoly reserves is a $(\mu + 1)\frac{e}{e-1}$-approximation to the revenue of the optimal mechanism.*

In what follows, we show how to derive revenue covering results for environments beyond single-item auctions. In Section 5.1, we study general single-parameter feasibility environments, and show that greedy approximation algorithms are revenue covered. In Section 5.2 we discuss the revenue covering of the generalized-first-price position auction, and in Section 5.3 we discuss winner-pays-bid auctions with a discretized bid space. In each section, welfare and revenue results will follow as corollaries to our analysis.

5.1 Greedy Auctions

In a single-item auction, the bid of the unique winner gives both the revenue and the losers' threshold bids. This immediately implied 1-revenue covering. For more general feasibility environments, the relationship between total revenue and a particular agent's threshold bid is less straightforward, and the revenue covering parameter depends on the algorithm used to select winners. In this section we show that mechanisms which allocate agents greedily have revenue covering parameter equal to their approximation ratio.

We first formally define greedy algorithms. Intuitively, greedy algorithms order agents, and then make a single pass over the ordering, allocating each agent in turn if doing so maintains feasibility. A greedy algorithm is therefore distinguished by the way it selects an ordering.

Definition 22. *The* greedy by priority *algorithm 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 $\psi_i(a_i)$.*
2. *Initialize the set of winners $S = \emptyset$.*
3. *For each agent i in sorted order: if $S \cup \{i\}$ is feasible, $S = S \cup \{i\}$.*
4. *Return S.*

For example, the greedy-by-bid algorithm is given by priority functions $\psi_i(a_i) = a_i$ for all i. Our results will hold for general single-parameter feasibility environments, but we first highlight two settings where greedy algorithms are of particular interest: matroids and single-minded combinatorial environments. We define each below.

Definition 23. *A feasibility environment $\mathcal{F}$ is a* matroid *if the following two properties hold:*

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

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

Notable examples of matroids include *k-unit environments*, where the seller has k identical items to sell to buyers who demand at most one each, and *transversal matroids*, which are matchable subsets of vertices on one side of a bipartite graph.

Definition 24. *A* single-minded combinatorial auction *feasibility environment is defined by m indivisible items, n agents that each desire a bundle of items, and the constraint that no item can be allocated more than once. Agent i desires the set of items $\mathcal{S}_i$, she receives value v_i for receiving any superset of $\mathcal{S}_i$ and value 0 otherwise. An allocation vector $\mathbf{x} \in \{0,1\}^n$ is feasible if and only if for all agents $i \neq i'$, simultaneous allocation $x_i = x_{i'} = 1$ implies disjoint demands $\mathcal{S}_i \cup \mathcal{S}_{i'} = \emptyset$.*

Maximizing the total value of winners is equivalent to the problem of weighted set packing, which is NP-hard. Moreover, and even for special cases of the problem that are computationally tractable, equilibria of highest-bid-wins winners-pay-bid mechanism can be a factor m from optimal (see Borodin and Lucier, 2010, and Lemma 32, below). As shown by Borodin and Lucier (2010) and rederived in our framework below, greedy algorithms are both computationally tractable and give better welfare guarantees in equilibrium.

Definition 25. *An algorithm $\tilde{\mathbf{x}}$ is an α-approximation for a feasibility environment $\mathcal{F}$ if for any bid profile $\mathbf{a}$ and feasible allocation $\mathbf{y}$, we have:*

$$\sum\nolimits_i a_i \tilde{x}_i(\mathbf{a}) \geq \frac{1}{\alpha} \sum\nolimits_i a_i \mathbf{y}.$$

For matroids, the greedy-by-bid algorithm is optimal, i.e., a 1-approximation. For single-minded combinatorial environments, a greedy algorithm that takes the size of each agent's desired bundle into account achieves a $\sqrt{m}$-approximation. We summarize below.

Lemma 26. *For any matroid feasibility environment, greedy by priority with $\psi_i(a_i) = a_i$ for all agents i is a 1-approximation.*

Lemma 27 (Lehmann and Shoham, 2002)**.** *For any single-minded combinatorial auction environment, greedy by priority with $\psi_i(a_i) = a_i/\sqrt{|\mathcal{S}_i|}$ for all agents i is a $\sqrt{m}$-approximation.*

Borodin and Lucier (2010) show that any winner-pays-bid auction which allocates using a greedy α-approximation achieves at least a $(\alpha + O(\alpha^2/e^\alpha))^{-1}$-fraction of the optimal welfare in equilibrium. This result holds over a broad range of multi-parameter settings. For single-parameter settings, we show that this result is a simple consequence of revenue covering.

Theorem 28. *For any feasibility environment $\mathcal{F}$ and greedy α-approximation algorithm $\mathcal{A}$ for $\mathcal{F}$, the winner-pays-bid mechanism which allocates according to $\mathcal{A}$ is α-revenue covered.*

For matroid and single-minded combinatorial environments, this immediately implies the following corollaries:

Corollary 29. *For any matroid feasibility environment, the winner-pays-bid mechanism with greedy by priority allocation with $\psi_i(a_i) = a_i$ for all agents i is at least a $\frac{e}{e-1}$-approximation to the welfare of the optimal auction in any BNE. With monopoly reserves and regular value distributions, the mechanism is at least a $\frac{2e}{e-1}$-approximation to the revenue of the optimal auction.*

Corollary 30. *For any single-minded combinatorial environment, the winner-pays-bid mechanism with greedy by priority allocation with $\psi_i(a_i) = a_i/\sqrt{|\mathcal{S}_i|}$ for all agents i is at least a $\sqrt{m}\frac{e}{e-1}$-approximation to the welfare of the optimal auction in any BNE. With monopoly reserves and regular value distributions, the mechanism is at least a $(\sqrt{m}+1)\frac{e}{e-1}$-approximation to the revenue of the optimal auction.*

Before proving Theorem 28, we motivate the analysis with an example where a non-greedy algorithm with good performance in the absence of incentives fails to obtain high welfare in equilibrium, and therefore fails to be meaningfully revenue-covered.

Definition 31. *The* winner-pays-bid highest-bid-wins mechanism *allocates the feasible set of bidders with the highest total bid, and charges winners their bid.*

Lemma 32. *There exists a single-minded combinatorial auction environment where the highest-bids-win mechanism is not μ-revenue covered for any $\mu < m$.*

Proof. Consider a setting with m items and $m + 2$ bidders. The first m bidders each want a single item: bidder j wants item j, each with a value for allocation of 1, deterministically. The final two bidders, meanwhile, each want the grand bundle of all m items, with values of $1 + \epsilon$, again deterministically. With appropriate tiebreaking, it is a BNE for each of the first m agents to bid 0, while the final two bidders bid $1 + \epsilon$. The optimal social welfare and revenue are both attained by selling to the first m bidders, for welfare and revenue of m, yielding a factor of m loss in both welfare and revenue.

This equilibrium also lower bounds the revenue-covering parameter for the highest-bids-win mechanism. First, note that the total revenue of the mechanism is $1 + \epsilon$. Next, consider the feasible allocation $\mathbf{y}$ which allocates the first m bidders. For these bidders, they must bid at least $1+\epsilon$ to get allocated, at which point they get allocated with probability 1. It follows that for such bidders, $T_i = 1+\epsilon$, and hence $\sum_i T_i y_i = (1+\epsilon)m$, m times the mechanism's revenue. □

In the example, the high bids of the $(m+1)$th and $(m+2)$th bidders discouraged participation from the other bidders - individually, each bidder would have needed to bid $1+\epsilon$ to win. As a group, though, the losing bidders could have won by increasing each of their bids by a tiny amount. Greedy algorithms lack this pathology. On the other hand, for any greedy allocation algorithm, we could increase the bids of every losing agent to their threshold without changing the outcome. We formalize this property as follows.

Definition 33. *Allocation rule* $\tilde{\mathbf{x}}$ *is* coalitionally non-bossy *if: for any profiles of bids* $\mathbf{a}$ *and* $\mathbf{a}'$ *where the bids in* $\mathbf{a}'$ *are the same as* $\mathbf{a}$ *for winners under* $\mathbf{a}$ *and at most their critical prices for losers under* $\mathbf{a}$*, i.e., if* $\tilde{x}_i(\mathbf{a}) = 0$ *then* $a'_i \leq \tau_i(\mathbf{a}_{-i})$*; then the allocations of* $\tilde{\mathbf{x}}$ *under* $\mathbf{a}$ *and* $\mathbf{a}'$ *are the same* $\tilde{\mathbf{x}}(\mathbf{a}) = \tilde{\mathbf{x}}(\mathbf{a}')$.

Lemma 34. *Any greedy by priority allocation rule is coalitionally non-bossy.*

Proof. Imagine changing $\mathbf{a}$ to $\mathbf{a}'$ by increasing one loser's bid at a time. Each time we increase a bid, say, of bidder i, two things remain true: (1) i still loses: as long as $a'_i \leq \tau_i(\mathbf{a}_{-i})$, i is passed over as infeasible when she is reached by the greedy algorithm; 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. □

Proof of Theorem 28. We argue that for any action profile $\mathbf{a}$ and alternate allocation $\mathbf{y}$,

$$\sum\nolimits_i a_i \tilde{x}_i(\mathbf{a}) \geq \tfrac{1}{\alpha} \sum\nolimits_i \tau_i(\mathbf{a}_{-i}) y_i.$$

Taking the expectation of both sides yields the desired inequality.

Let $\mathbf{a}'$ be a vector of bids where losers under $\mathbf{a}$ bid $\tau_i(\mathbf{a}_{-i})$, while winners bid as before. The following inequalities hold, with justifications after.

$$\begin{aligned}
\sum\nolimits_i a_i \tilde{x}_i(\mathbf{a}_{-i}) &= \sum\nolimits_i a'_i \tilde{x}_i(\mathbf{a}_{-i}) \\
&= \sum\nolimits_i a'_i \tilde{x}_i(\mathbf{a}') \\
&\geq \tfrac{1}{\alpha} \sum\nolimits_i a'_i y_i \\
&\geq \tfrac{1}{\alpha} \sum\nolimits_i \tau_i(\mathbf{a}_{-i}) y_i.
\end{aligned}$$

The first line holds because $\mathbf{a}'$ differs from $\mathbf{a}$ only on the bids of losing agents. The second follows from Lemma 34, and the third from Lemma 27. The last line follows from the fact that $\mathbf{a}'$ doesn't change the bids of winners under $\mathbf{a}$, and for those agents, $a_i \geq \tau_i(\mathbf{a}_{-i})$. □

The conclusions of Theorem 28 and its analysis are twofold. First, comparing the greedy-by-priority algorithm from Lemma 27 to the highest-bids-win algorithm of Lemma 32 reveals that welfare loss in equilibrium can stem from two sources. Ignoring computational constraints, the highest-bids-win algorithm is optimal in the absence of incentives. Hence, the factor of $\Omega(m)$ welfare loss in the example of Lemma 32 is entirely due to incentives. The greedy algorithm, on the other hand, produces suboptimal allocations without incentives. However, Corollary 30 states that the loss from the introduction of incentives is limited to at most a factor of $e/(e-1)$. Hence, the welfare loss is primarily due the algorithm's performance without incentives. Second, this comparison suggests the value of revenue covering as objective in algorithm design. Indeed, Dütting and Kesselheim (2015) subsequently study the design of algorithms with low revenue covering parameters. For single-minded combinatorial auctions, they show that no algorithm is $o(\sqrt{m})$-revenue covered. Hence, the greedy algorithm is optimal with respect to this objective.

5.2 Position Auctions

In this section, we derive welfare and revenue results for winner-pays-bid position auctions, a.k.a., the generalized first-price auction (GFP). While these results are driven by the same fundamental notions as before, value covering and revenue covering, the precise formulations of these two ideas must be modified to fit environments such as position auctions where allocation is inherently probabilistic. We give the appropriate definitions of value covering and revenue covering that yield the results, and relegate the proof details to Appendix D.1.

Formally, a position auction is an auction in which agents can be allocated one of m positions; position j is valued by an agent at $\alpha_j v_i$. In advertising auctions, these are positions on a webpage to fill where lower positions receive fewer clicks. The positions are ordered such that $\{\alpha_j\}$ is decreasing in j (hence position 1 is best). In GFP, agents submit bids a_i, and positions are allocated in order of bid. Each agent pays their bid scaled by the quality of the position: $\alpha_j a_i$. Equivalently, they pay their bid when they are served, which occurs with probability α_j for position j.

In a position auction, an agent can win one of several different positions, and consequently the agent's threshold bid for guaranteed allocation is not useful for analysis. We use instead a finer-grained measure of an agent's threshold bid that is parameterized by allocation probability. In particular, we consider $T_i(y_i) = \int_0^{y_i} t_i(z)\, dz$, which we refer to $T_i(y_i)$ as the *partial threshold.*

Definition 35. *For any (implicit) strategy profile* $\mathbf{s}$ *and any allocation probability* $y_i \in [0,1]$, *the* partial threshold *for* y_i, *denoted* $T_i(y_i)$, *is given by* $T_i(y_i) = \int_0^{y_i} t_i(z)\, dz$.

With partial thresholds defined, we can now state revenue covering for partial thresholds:

Definition 36 (Revenue Covering with Partial Thresholds)**.** *A mechanism is* μ*-revenue covered with partial thresholds if for any (implicit) strategy profile* $\mathbf{s}$*, and feasible allocation* $\mathbf{y}$*,*

$$\mu\textsc{Rev}(M) \geq \sum\nolimits_i T_i(y_i). \tag{13}$$

Theorem 37. *The generalized first price auction is* 1*-revenue covered with partial thresholds.*

The proof is included in Appendix D.1. Note that revenue covering with partial thresholds is weaker than the general revenue covering (Definition 16) condition. This follows from the fact that $T_i(y_i)$ is convex in y_i, and hence $T_i(y_i) \geq T_i y_i$. With a weaker notion of revenue covering, we need a stronger value covering condition with partial thresholds.

Lemma 38 (Value Covering for Partial Thresholds)**.** *In any BNE of a winner-pays-bid auction, for any bidder* i *with value* v_i*,*

$$u_i(v_i) + T_i(y_i) \geq \tfrac{e-1}{e} y_i v_i. \tag{14}$$

The proof is omitted; it proceeds almost identically to Lemma 15 but with $T_i(y_i)$ in place of T_i and $y_i v_i$ in place of v_i. The worst case u_i becomes $y_i v_i/e$, rather than v_i/e. Combining revenue and value covering then gives a welfare approximation:

Theorem 39. *The welfare of any BNE of* GFP *is an* $\frac{e}{e-1}$*-approximation to the welfare of the optimal auction.*

All analysis prior to this section could be performed with partial thresholds rather than full thresholds $T_i y_i$. We chose not to do so as the partial threshold definition is only necessary in probabilistic feasibility settings. In particular, the value and revenue covering framework with partial thresholds naturally extends to virtual values and reserves. This extension yields the following revenue result:

Theorem 40. *For regular environments, the revenue in any BNE of* GFP *with monopoly reserves is a* $\frac{2e}{e-1}$*-approximation to the revenue of the optimal auction.*

5.3 Discretized Bids

In practice, the bid space in an auction is often discretized for convenience or feasibility. We note here that discretizing the bid space of a winner-pays-bid auction only results in a small additive loss to the bounds proved via the revenue covering framework. In particular, consider a winner-pays-bid mechanism with unrestricted bid space. Restricting the bid space to integral multiples of δ for some $\delta > 0$ preserves μ-revenue covering up to an additive term. Formally:

Lemma 41. *For any winner-pays-bid mechanism M, let M^δ be M with bid space restricted to integral multiples of δ. Let T_i^δ be agent i's expected threshold in M^δ for some bid distribution $\mathbf{G}$ over multiples of δ. If M is μ-revenue covered, then the following inequality holds for any feasible allocation $\mathbf{y}$:*

$$\textsc{Rev}(M^\delta, \mathbf{G}) \geq \sum\nolimits_i T_i^\delta y_i - n\delta.$$

Proof. Assume agent i bidding a_i in M against $\mathbf{G}_{-i}$ gets allocation probability $\tilde{x}_i(a_i)$. To get allocation at least $\tilde{x}_i(a_i)$ in M^δ when facing $\mathbf{G}_{-i}$, i will need to bid at most the next-highest increment of M. Hence, if we define $t_i^\delta(x)$ the smallest bid that achieves allocation of at least x in M^δ, we have $t_i^\delta(x) \leq t_i(x) + \delta$ for all $x \in [0, 1]$. Integrating over all $x \in [0, 1]$ for each agents and applying μ-revenue covering yields the result. □

The additive $n\delta$ loss in revenue covering translates directly to an additive $n\delta$ loss in revenue. If δ is small compared with the revenue of the original mechanism, this makes effectively no difference: if δ is very large relative to revenue, it could make a large difference.[3]

6 Beyond Winner-pays-bid Auctions

In this section, we extend our framework for proving worst-case approximation bounds to auctions other than those with winner-pays-bid semantics. We do this by showing that a bidder's decision problem in any auction is equivalent to choosing how to bid in a winner-pays-bid auction. Consequently, a form of value covering holds for *any* auction for single-dimensional agents with independent value distributions in BNE. We then apply this winner-pays-bid reduction to analysis of welfare and revenue in all-pay auctions (Section 6.2), as well as to the simultaneous composition of revenue-covered auctions (Section 6.4).

[3] In a setting where at most m agents may be allocated simultaneously, the loss in Lemma 41 becomes $m\delta$.

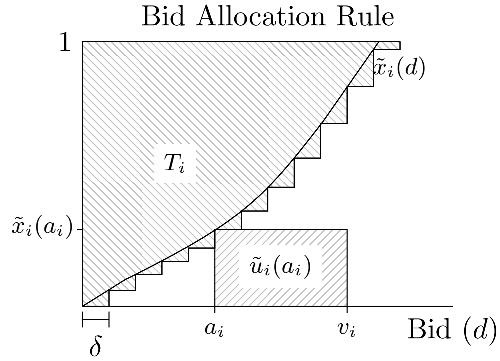

Figure 5: When bids are only allowed in multiples of δ, the expected threshold expands, but only by at most an additive amount of δ.

6.1 Extending the Framework

To instantiate the revenue covering framework in mechanisms with general payment semantics, we first write the utility $\tilde{u}_i(\cdot)$ of an agent i as:

$$\tilde{u}_i(a_i) = v_i\tilde{x}_i(a_i) - \tilde{p}_i(a_i) = \left(v_i - \frac{\tilde{p}_i(a_i)}{\tilde{x}_i(a_i)}\right)\tilde{x}_i(a_i).$$

Note the resemblance to the utility of an agent in a winner-pays-bid mechanism: the term $\frac{\tilde{p}_i(a_i)}{\tilde{x}_i(a_i)}$ plays the same role as the bid in a pay-your-bid auction: it is the price per unit of allocation. For this reason, we will refer to it as the equivalent bid for action a_i.

Definition 42. *The* equivalent bid *for an action a_i, denoted $\tilde{\beta}_i(a_i)$, is given by $\tilde{\beta}_i(a_i) = \frac{\tilde{p}_i(a_i)}{\tilde{x}_i(a_i)}$.*

Note that unlike an agent's winner-pays-bid bid, the equivalent bid is generally a function of other agents' bid distributions. We will now define expected thresholds, revenue covering, and value covering with equivalent bids playing the role of winner-pays-bid bids.

Equivalent Threshold Bids The key quantity in our winner-pays-bid proof framework is the threshold bid - the lowest value an agent must bid to get allocated. Threshold bids are an ex-post notion, while equivalent bids as given in Definition 42 are only well-defined in the interim. Recall that in Section 3 we computed the expected threshold bid by integrating the inverse of its cumulative distribution function. We employ a similar approach for our more general framework.

Recall that the bid threshold $t_i(z)$ for winner-pays-bid mechanisms was the lowest bid which achieved allocation at least z. We extend this to general mechanisms by using the lowest equivalent bid required to achieve allocation at least z:

Definition 43. *Fixing an action distribution $\mathbf{G}$ and allocation probability z, let $\alpha_i(z) = \arg\min_{a_i:\tilde{x}_i(a_i)\geq z} \tilde{\beta}_i(a_i)$ be the action inducing allocation probability at least z with the lowest equivalent bid. The* equivalent threshold bid *for z is then given by $t_i(z) = \tilde{\beta}_i(\alpha_i(z))$.*

In the first-price auction, we used value covering and revenue covering to relate expected threshold bids to the welfare and revenue of the first-price auction and the optimal auction. By taking the expectation of the threshold bid, we were able to aggregate all of the possible allocation versus price-per-unit tradeoffs an agent might consider into a single parameter. In the more general setting, this expectation is not defined, so we instead recall the alternative formulation: we integrated the inverse of the bid allocation rule, $t_i(z)$. For general mechanisms, we will aggregate equivalent threshold bids in a similar manner. Formally:

Definition 44. *Fixing an action distribution* $\mathbf{G}$*, let the* cumulative equivalent threshold bid*, or simply* cumulative threshold *for agent i to be* $T_i = \int_0^1 t_i(z)\,dz$.

We use the term "cumulative threshold" rather than "expected threshold" in the framework for general mechanisms because unlike in winner-pays-bid mechanisms, T_i does not necessarily correspond to an expectation.

Covering Conditions and the Price of Anarchy. We now show how equivalent bids reduce the optimization problem of a bidder in a general auction to that of a bidder in a first-price auction. In particular, value covering, which quantifies the tradeoff between utility and expected threshold, still holds with cumulative thresholds:

Lemma 45 (Value Covering)**.** *Consider a mechanism M in BNE and interim utility function $u_i(v_i)$ for agent i with value v_i and cumulative equivalent threshold bid T_i. Then*

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{15}$$

The proof (included in Appendix B) is identical to that of winner-pays-bid value covering (Lemma 4) because agent i is making the same tradeoff between allocation probability and price per unit of allocation in each.

We define μ-revenue covering in the same way using the generalized definition of cumulative equivalent threshold T_i:

Definition 46 (μ-Revenue Covering)**.** *A mechanism M is μ-*revenue covered *if for any (implicit) distribution of actions* $\mathbf{G}$ *and feasible allocation* $\mathbf{y}$*,*

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

As before, value covering and revenue covering together imply a welfare result:

Theorem 47 (Extension of Theorem 17)**.** *If a mechanism is μ-revenue covered, then in any BNE it is a $\mu\frac{e}{e-1}$−approximation to the welfare of the optimal mechanism.*

Finally, virtual value covering, which held as a corollary of value covering, did not require any properties of T_i, and consequently holds for general mechanisms as well:

Lemma 48 (Extension of Lemma 6)**.** *In any BNE of the first price auction, for any bidder i with value v_i such that $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + T_i \geq \tfrac{e-1}{e}\phi_i(v_i). \tag{16}$$

Note that Lemma 48 does not directly give a revenue result, as the techniques for controlling negative virtual surplus (e.g. reserves) may not apply for all single-dimensional mechanisms. In Section 6.2, however, we will show that with sufficient competition, i.e. k-duplicates, the single-item all-pay auction approximates the optimal revenue.

6.2 All-Pay Mechanisms

In an *all-pay* mechanism, each agent pays their bid, regardless of whether they win or lose. We now show, by giving a reduction to the analysis of winner-pays-bid mechanisms, that all-pay auctions have approximately optimal welfare and, with sufficient competition, revenue. Since value covering holds for any mechanism in BNE, we must simply show that all-pay mechanisms are revenue covered, i.e., the expected revenue of the mechanism must approximate the cumulative equivalent thresholds. Note that the equivalent bid corresponding to an all-pay bid can be found simply by dividing by the allocation probability:

$$\tilde{\beta}_i(a_i) = a_i/\tilde{x}_i(a_i). \tag{17}$$

With a factor two loss we will reduce revenue covering in all-pay mechanisms to revenue covering in winner-pays-bid mechanisms with the same bid-allocation rule. We conclude below that for single-item and matroid environments, highest-bids-win all-pay mechanisms satisfy revenue covering with $\mu = 2$ and, thus, are $2e/(e-1)$ approximately efficient.

To make the reduction clear we denote the cumulative effective thresholds of all-pay and winner-pays-bid mechanisms with allocation rule $\tilde{\mathbf{x}}$, respectively, as:

$$T_i^{\mathrm{AP}} = \int_0^1 t_i^{\mathrm{AP}}(z)\, dz,$$
$$T_i^{\mathrm{WPB}} = \int_0^1 t_i^{\mathrm{WPB}}(z)\, dz,$$

where $t_i^{\mathrm{AP}}(z)$ and $t_i^{\mathrm{WPB}}(z)$ are the inverses of the respective effective bid allocation rules. Note that the bid allocation rule and the effective bid allocation rule are identical for winner-pays-bid mechanisms but distinct for all-pay mechanisms.

Lemma 49. *For a common bid distribution, the cumulative equivalent thresholds of all-pay and winner-pays-bid mechanisms satisfy $T_i^{\mathrm{AP}} \leq 2T_i^{\mathrm{WPB}}$.*

Proof. Fix a common distribution of bids. The effective thresholds of all-pay and winner-pays-bid format mechanisms are related by equation (17) as

$$t_i^{\mathrm{AP}}(z) = z/t_i^{\mathrm{WPB}}(z). \tag{18}$$

This relation yields the following sequence of inequalities:

$$T_i^{\mathrm{WPB}} = \int_0^1 t_i^{\mathrm{WPB}}(z)\, dz = \int_0^1 z\, t_i^{\mathrm{AP}}(z)\, dz \geq \frac{1}{2}\int_0^1 t_i^{\mathrm{AP}}(z)\, dz = \frac{1}{2}T_i^{\mathrm{AP}} \tag{19}$$

where the second equality follows from equation (18) and the inequality from Chebyshev's sum inequality and the fact that t_i^{AP} is a non-decreasing function. □

Theorem 50. *For bid allocation rule $\tilde{\mathbf{x}}$, if the winner-pays-bid mechanism is μ-revenue covered, then the all-pay mechanism is 2μ-revenue covered.*

Proof. For a fixed common distribution of bids and common bid allocation rule $\tilde{\mathbf{x}}$, the revenue of the all-pay mechanism exceeds the revenue of the winner-pays-bid mechanism. The former is the sum of all the bids and the latter is the only the sum of the winning bids. On the other hand, Lemma 49 shows that the cumulative effective threshold of the all-pay mechanism is at most twice that of the winner-pays-bid mechanism for the common distribution of bids. Thus, if μ-revenue covering holds for the winner-pays-bid mechanism, 2μ-revenue covering holds for the all-pay mechanism. □

Combining revenue covering with Theorem 47 for highest-bids-win all-pay mechanisms in single-item and matroid environments gives a welfare bound of $2e/(e-1)$. For revenue, using the techniques of Section 4.3 and ensuring that at least two bidders have values drawn from each distribution (i.e. 2-duplicates) gives a $4e/(e-1)$-approximation to the revenue of the optimal auction.[4]

The bounds can be improved to 2 and 6 for welfare and revenue, respectively, by adapting the value-covering condition to the all-pay format, shown in Appendix E.

6.3 The Second-price Auction

Not all mechanisms are revenue covered. In the second-price auction, agents submit sealed bids, the highest bidder wins and is charged the second-highest bid. This auction lacks a direct connection between bidders' threshold bids and the revenue of the auction, which is required for revenue covering. To illustrate, consider a two-agent setting, and assume agent 1 bids 10 and agent 2 bids 0, deterministically. The revenue is 0, but T_2 is 10, so the second-price auction cannot be revenue covered. Moreover, the welfare of this equilibrium can be far from the optimal welfare, for example, when agent 2 has value 1 and agent 1 has value 9.

If agents are assumed to never bid above their values, the bidders' threshold bids are a lower bound on the welfare of the auction, which yields a "welfare covering" property which functions similarly to revenue covering in our framework. This results in welfare guarantees for no-overbidding equilibria that are analogous to the no-overbidding analyses of Syrgkanis and Tardos (2013) and Caragiannis et al. (2014).

6.4 Simultaneous Composition

In this section, we prove that μ-revenue covering is closed under simultaneous composition. In other words, running several μ-revenue covered auctions simultaneously yields an auction which is also μ-revenue covered. Consequentially, welfare bounds proved for individual mechanisms via revenue covering extend to collections of such composed mechanisms.

For simplicity we assume that each component mechanism has a "withdraw" action which induces zero allocation and payment in that mechansim, akin to bidding 0 in a first-price auction. Furthermore, we make two assumptions on agent utilities. First, they are *unit-demand* in the sense that allocation from more than one mechanism gives provides the same surplus as being allocated from exactly one mechanism. Second, they are *single-valued*, meaning agent i has the same value v_i for allocation, regardless of the mechanism whose allocation agent i receives. Formally, the simultaneous composition of m mechanisms for single-dimensional agents is the following:

Definition 51. *Let mechanisms $M^1, \dots, M^m$ have allocation and payment rules $(\mathbf{x}^j, \mathbf{p}^j)$ for $j \in \{1, \dots, m\}$ and individual action spaces spaces $A_i^1, \dots, A_i^m$ for each agent i. The* simultaneous composition *M of $M^1, \dots, M^m$ is defined to have:*

[4] We do not discuss a revenue approximation bound for all-pay mechanisms with reserves as exogenously imposed reserves in bid space map to endogenous reserves in value space. This phenomenon makes imposing monopoly reserves in value space difficult with all-pay mechanisms.

- *Action space $A_i = \prod_j A_i^j$ for each agent. That is, each agent participates in the global mechanism by participating individually in each composed mechanism.*
- *Vector-valued allocation rule $\left(\tilde{x}_i^1(\mathbf{a}^1), \ldots, \tilde{x}_i^m(\mathbf{a}^m)\right)$. In other words, the mechanism gives each agent their allocated bundle from each mechanism. The induced single-dimensional allocation rule of the global mechanism is $\tilde{x}_i(\mathbf{a}) = \max_j \tilde{x}_i^j(\mathbf{a})$.*
- *Payment rule $\tilde{p}_i(\mathbf{a}) = \sum_j \tilde{p}_i^j(\mathbf{a}^j)$. That is, agents make payments to every composed mechanism.*

Agent utilities are therefore of the form $\tilde{u}_i(\mathbf{a}) = v_i \cdot (\max_j \tilde{x}_i^j(\mathbf{a})) - \tilde{p}_i(\mathbf{a}) = v_i\, \tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})$. Now using $\tilde{\mathbf{x}}$, define t_i and T_i as discussed in Section 6.1. We can now state the main theorem of the section.

Theorem 52. *Let M be the simultaneous composition of μ-revenue covered mechanisms $M^1, \ldots M^m$ with unit-demand, single-valued agents; then M is μ-revenue covered.*

The intuition driving the proof of Theorem 52 is that (a) the cumulative threshold of the composite mechanism is smaller than the cumulative threshold of each of the individual mechanisms (it is only easier for an agent to get allocated when there are more mechanisms to bid in) while (b) the revenue of the composite mechanism is equal to the sum of the revenues of the individual mechanisms.

Each agent's vector of bids in the global mechanism yields a bid for each individual mechanism. Therefore, for each individual mechanism j, any global bid distribution $\mathbf{G}$ induces a local bid distribution $\mathbf{G}^j$. The distribution $\mathbf{G}^j$ in mechanism j induces local versions of the equivalent bid $\tilde{\beta}_i^j$, threshold τ_i^j, and expected threshold T_i^j.

We can now formalize the above intuition that individual expected thresholds are larger than the expected thresholds in the global mechanism:

Lemma 53. *For any (implicit) bid distribution $\mathbf{G}$ in the composite mechanism, the cumulative thresholds satisfy $T_i \leq T_i^j$.*

Proof. Fix an individual mechanism j. The threshold function for the global mechanism is defined as $t_i(z) = \min_{a_i : \tilde{x}_i(a_i) \geq z} \tilde{\beta}_i(a_i)$, and for the local mechanism as $t_i^j(z) = \min_{a_i^j : \tilde{x}_i^j(a_i^j) \geq z} \tilde{\beta}_i^j(a_i^j)$. For every action a_i^j in the local mechanism, there is a corresponding action a_i in the global mechanism where i takes action a_i^j in mechanism j and withdraws from every other mechanism. Since $\tilde{\beta}_i(a_i) = \tilde{\beta}_i^j(a_i^j)$, for this j, a_i, and a_i^j it follows that $t_i(z) \leq t_i^j(z)$. This inequality only improves if we allow any action in the composite mechanism. Integrating over all values of z implies the lemma. □

Proof of Theorem 52. An allocation $\mathbf{y}$ is feasible in the global mechanism if and only if there exists a feasible allocation profile $\mathbf{y}^j$ for each component mechanism j such that for each agent i, $y_i = \sum_j y_i^j$. The theorem now follows from the following sequence of inequalities:

$$\begin{aligned}
\mu \text{REV}(M) &= \mu \sum\nolimits_j \text{REV}(M_j) \\
&\geq \sum\nolimits_j \sum\nolimits_i T_i^j y_i^j \\
&\geq \sum\nolimits_i T_i \sum\nolimits_j y_i^j \\
&= \sum\nolimits_i T_i y_i
\end{aligned}$$

The first equality follows from the definition of simultaneous composition. The second line follows revenue covering of mechanism j. The third is a reordering of sums and application of Lemma 53. The final inequality follows from the fact that $y_i = \sum_j y_i^j$. □

Theorem 52 implies that our welfare bounds extend to simultaneous compositions of revenue-covered mechansisms. Specifically:

Corollary 54. *Let M be the simultaneous composition of m μ-revenue covered mechansisms for unit-demand, single-valued agents. Then the welfare of M is a $\mu\frac{e}{e-1}-$approximation to the optimal welfare for the composite environment.*

It is possible to extend the proof of Theorem 52 to hold for winner-pays-bid auctions with reserves. In this case, the expected threshold for the global mechanism with reserves $\mathbf{r}$ should be defined as $T_i^{r_i} = \int_{\underline{x}_i(r_i)}^1 t_i(z)\,dz$, for $\underline{x}_i(r_i)$ defined as $\max_{a_i^j:\tilde{\beta}_i^j(a_i^j)\leq r_i} \tilde{x}_i^j(a_i^j)$. Here $\underline{x}_i(r_i)$ serves the role of $\tilde{x}_i(r_i)$ in the winner-pays-bid definition of $T_i^{r_i}$. Under the new definition of $T_i^{r_i}$, one can show that revenue and value covering with reserves extend from the single-item case to the simultaneous composition. Since the simultaneous composition of first-price auctions is revenue covered with reserves, we may conclude the following:

Theorem 55. *Let M be the simultaneous composition of m first-price auctions with monopoly reserves $\mathbf{r}^*$, and unit-demand, single-valued agents with regular value distributions. The revenue of M is a $\frac{2e}{e-1}-$approximation to the revenue of the optimal global mechanism.*

7 Conclusion

We have given a framework for proving worst-case approximation results for welfare and revenue in Bayes-Nash equilibrium. This framework enabled us to prove both welfare and new revenue approximation results for non-truthful auctions in asymmetric settings, including first price and all-pay auctions.

This framework has two distinct parts that isolate the analysis of Bayes-Nash equilibrium from the analysis of the specific mechanism. The first part, value covering, depends only on Bayes-Nash equilibrium and relates an agent's surplus and expected threshold price to her value. The second, revenue covering, is a property of the mechanism which must hold for every bid distribution. This framework is especially helpful when equilibria are hard to characterize or understand analytically, as is the case with the first-price auction in asymmetric environments and we expect this framework to aid broadly in understanding properties of equilibria in auctions well beyond the confines of classical analyses.

We invoked the characterization of Bayes-Nash equilibrium in a few specific places in our proofs. For value covering and virtual value covering, it is only important that an agent be best responding to the expected actions of other bidders. For the revenue approximation results, we do rely on the characterization of equilibrium by Myerson (1981) to account for revenue via virtual values. This crucially allows us to relate the allocation a bidder receives to their contribution to revenue. Extensions beyond single-parameter, risk-neutral, private-valued agents will be challenging without a virtual-value equivalent.

References

Ausubel, L. M. and Milgrom, P. (2006). The lovely but lonely vickrey auction. *Combinatorial auctions*, pages 17–40.

Borodin, A. and Lucier, B. (2010). Price of anarchy for greedy auctions. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 537–553.

Bulow, J. and Klemperer, P. (1996). Auctions Versus Negotiations. *The American Economic Review*, 86(1):180–194.

Caragiannis, I., Kaklamanis, C., Kyropoulou, M., Lucier, B., Paes Leme, R., and Tardos, E. (2014). Bounding the inefficiency of outcomes in generalized second price auctions. pages 1–45.

Chawla, S. and Hartline, J. D. (2013). Auctions with unique equilibria. In *ACM Conference on Electronic Commerce*, pages 181–196.

Christodoulou, G., Kovács, A., Sgouritsa, A., and Tang, B. (2013). Tight bounds for the price of anarchy of simultaneous first price auctions. *arXiv preprint arXiv:1312.2371*.

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

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

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*.

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.

Kirkegaard, R. (2009). Asymmetric first price auctions. *Journal of Economic Theory*, 144(4):1617–1635.

Kirkegaard, R. (2012). A Mechanism Design Approach to Ranking Asymmetric Auctions. *Econometrica*, 80(5):2349–2364.

Krishna, V. (2009). *Auction Theory*. Academic Press/Elsevier.

Lebrun, B. (2006). Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior*, 55(1):131–151.

Lehmann, Daniel, I. O. L. and Shoham, Y. (2002). Truth Revelation in Approximately Efficient Combinatorial Auctions. *Journal of the ACM*, 49(5):577–602.

Maskin, E. and Riley, J. (2003). Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45(2):395 – 409.

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

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

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

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

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37.

A First-Price Welfare Approximation Lower Bound

In this appendix, we describe an equilibrium of the single-item first-price auction with independent but non-identically distributed values in which the equilibrium welfare is less than that of the optimal welfare by a factor of approximately 1.15. Our example will have $n+1$ bidders. Bidders $1,\ldots,n$ will be designated *low-valued bidders*, with an identical value distribution to be determined shortly. Bidder $n+1$ will be the *high-valued bidder*, with value deterministically 1. Allocating to bidder $n+1$ in all value profiles yields a lower bound on the optimal welfare of 1. Our constructed allocation will misallocate to low-valued bidders, yielding an expected welfare of approximately .869.

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

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

All that remains is to select a mixed strategy for the high-valued bidder, their expected utility parameter, u_H, and a number of low-valued bidders n. To produce an equilibrium with low welfare, we must navigate a tradeoff. If G_H is too aggressive, then the high-valued bidder will win frequently, yielding high welfare. If G_H is too weak, then noting the formula for the low-valued bidders' values, we see that these values will generally be high. A similar tradeoff applies in selecting u_H. Numerical experimentation shows that choosing $G_H(a) = \sqrt{a/(1-u_H)}$ and $u_H = .57$ yields low welfare. Given these choices, one can compute the expected welfare in equilibrium as approximately .869 for very large n. As mentioned, allocating the high-valued bidder yields a lower bound of 1 on the optimal social welfare. This implies the desired approximation ratio of 1.15.

B Framework Proofs

Lemma 45 (Restatement). *Consider a mechanism M in BNE with interim utility function $u_i(v_i)$ for agent i with value v_i. Then*

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{15}$$

Proof of Lemma 45. Note that by the definition of BNE, i chooses an action which maximizes utility. It follows that

$$u_i(v_i) \geq v_i x_i(\alpha_i(z)) - p_i(\alpha_i(z)) = \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) x_i(\alpha_i(z)) \geq \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) z. \tag{20}$$

Rearranging (20) yields

$$v_i - \frac{u_i(v_i)}{z} \leq \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))} = t_i(z). \tag{21}$$

The remainder of the proof is identical to the proof of Lemma 4. □

C Revenue Extension Proofs

C.1 Auctions with Duplicate Bidders

We first prove our extension of the analysis of Bulow and Klemperer (1996) to first-price and all-pay auctions with asymmetric distributions, which we restate below:

Lemma 13 (Restatement). *In any single-item setting with k-duplicates and regular value distributions, the following inequalities hold:*

$$\textsc{Rev}(\text{FPA}) \geq \tfrac{k-1}{k}\textsc{Rev}^+(\text{FPA})$$
$$\textsc{Rev}(\text{APA}) \geq \tfrac{k-1}{k}\textsc{Rev}^+(\text{APA})$$

To prove the lemma, assume k-duplicates holds, and consider a partition of the agents into groups $B_1, B_2, \ldots B_p$ such that each group has size at least k and all agents in each group B_i have values drawn from the distribution F_i.

First note that in any BNE of a first-price or all-pay auction with k-duplicates, agents of the same group will play symmetric strategies. This follows from Theorem 3.1 of Chawla and Hartline (2013), which gives that any two agents when competing against a reserve distribution will behave identically. We obtain the following:

Corollary 56 (of Theorem 3.1, Chawla and Hartline (2013))**.** *In any BNE of a first-price or all-pay auctions with k-duplicates, for any group B_j of agents who have identically distributed values, all agents in the group play by identical strategies everywhere except on a measure zero set of values.*

We now relate the revenue from each group of bidders to the revenue from a symmetric second price auction with reserves among only the bidders within the group of duplicates, allowing us to use the symmetric auction approximation results of Bulow and Klemperer (1996). Let $\text{SPA}_R(B)$ be a second price auction run among agents in group B with a random reserve drawn according to the distribution R.

Lemma 57. *In any first-price or all-pay auction with k-duplicates, denoted M_k, and duplicate groups $B_1, \ldots, B_p$, there exist reserve distributions $R_1, R_2 \ldots R_p$ such that*

$$\text{Rev}(M_k) = \sum\nolimits_j \text{Rev}(\text{SPA}_{R_j}(B_j)), \tag{22}$$

$$\text{Rev}^+(M_k) = \sum\nolimits_j \text{Rev}^+(\text{SPA}_{R_j}(B_j)). \tag{23}$$

Proof. Fix the values and actions of bidders outside a group j in the first-price or all-pay auction, as well as the strategies of bidders in group j. Fixing these results in a value $\underline{v}_j$ such that the highest bidder in group j will get allocated if and only if their value is above $\underline{v}_j$. Let R_j be the distribution of $\underline{v}_j$ induced by the random values of bidders outside group j. A second price auction among the members of B_j with reserve price drawn from R_j will induce exactly the same allocation rule as BNE of the first-price or all-pay auction for all members of the group. By revenue equivalence (Part 3 of Theorem 1), the revenue from members of group j in M_k will be the same as $\text{Rev}(\text{SPA}_{R_j}(B_j))$. The same argument holds for $\text{Rev}^+(M_k)$ and $\text{Rev}^+(\text{SPA}_{R_j}(B_j))$. □

A second-price auction within a group is now a symmetric setting, and thus we can now use the work of Bulow and Klemperer (1996) to relate (22) and (23).

Proof of Lemma 13. By Theorem 11, if $k \geq 2$, $\text{Rev}(\text{SPA}_{R_j}(B_j)) \geq \frac{k-1}{k}\text{Rev}^+(\text{SPA}_{R_j}(B_j))$ and hence for any first-price or all-pay auction M:

$$\begin{aligned}\text{Rev}(M) &= \sum\nolimits_j \text{Rev}(\text{SPA}_{R_j}(B_j)) \\ &\geq \sum\nolimits_j \frac{k-1}{k}\text{Rev}^+(\text{SPA}_{R_j}(B_j)) \\ &= \frac{k-1}{k}\text{Rev}^+(M).\end{aligned}$$

□

We now prove our revenue result for the first-price auction with duplicates.

Proof of Theorem 14. Let $\mathbf{x}^*(\cdot)$ be the revenue-optimal allocation rule. For any value profile $\mathbf{v}$, Lemma 6 implies:

$$\sum_i \phi_i(v_i)x_i(v_i) + \sum_i T_i x_i^*(\mathbf{v}) \geq \frac{e-1}{e}\sum_i v_i x_i^*(\mathbf{v}).$$

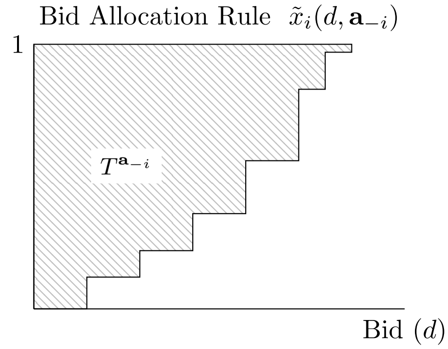

(a) In GFP, when other players play actions $\mathbf{a}_{-i}$, there is not just one threshold for allocation, but a threshold for allocation in each slot.

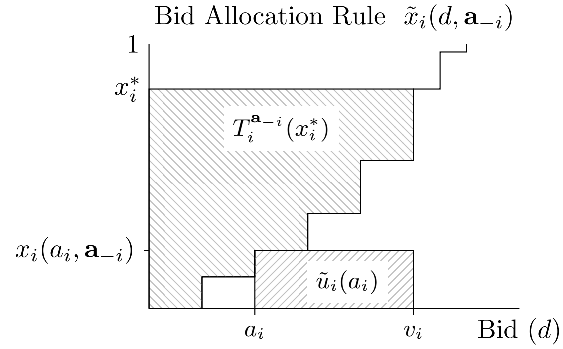

(b) For revenue covering in position auctions, the partial threshold $T_i^{\mathbf{a}_{-i}}(y_i)$ is used in place of $y_i T_i^{\mathbf{a}_{-i}}$. Note that by the convexity of $T_i^{\mathbf{a}_{-i}}(y_i)$ in y_i, $T_i^{\mathbf{a}_{-i}}(y_i) \leq y_i T_i^{\mathbf{a}_{-i}}$.

Figure 6

Taking expectations over $\mathbf{v}$ and using revenue covering yields

$$\frac{2e}{e-1}\textsc{Rev}^{+}(\text{FPA}) \geq \textsc{Rev}(\textsc{Opt}).$$

By Lemma 13, we may replace positive virtual values with the first price auction's expected revenue while only losing a $\frac{k}{k-1}$-factor. This proves the theorem. □

D Revenue Covering Proofs

D.1 GFP

In this section, we show that the generalized first-price auction satisfies revenue covering with partial thresholds given in Definition 36. For any alternate allocation $\mathbf{y}$, the revenue will cover for each player the partial threshold up to their alternate allocation amount $\mathbf{y}$. See Figure 6b for an illustration.

The proof has two main steps: first we show that GFP is revenue covered with partial thresholds if bidders play deterministic bids; next, that revenue-covering with deterministic actions implies revenue covering with general distributions over bids.

Proposition 58. *The generalized first-price position auction is* 1*-revenue covered with partial thresholds when agents play deterministic bids.*

Proof. Consider the bid-based allocation rule of an agent in GFP, $\tilde{x}_i(a_i, \mathbf{a}_{-i})$. For any bid a_i, if a_i would be the jth highest bid (e.g., $j-1$ other bidders are bidding above a_i), then $\tilde{x}_i(a_i, \mathbf{a}_{-i})$ is the position weight of slot j. So, $\tilde{x}_i(a_i, \mathbf{a}_{-i})$ is a stair function, with a stair corresponding to each position (like in Figure 6a). Furthermore, for a given allocation probability $z \in [0, 1]$, let $\tau^{\mathbf{a}_{-i}}(z) = \tilde{x}_i^{-1}(z, \mathbf{a}_{-i})$ be the threshold price for obtaining allocation probability z. Note that $\tau^{\mathbf{a}_{-i}}(\cdot)$ is also a stair function.

Let $T_i^{\mathbf{a}_{-i}}(y) = \int_0^y \tau^{\mathbf{a}_{-i}}(z)\, dz$ be the partial threshold when the other agents bid $\mathbf{a}_{-i}$. Denote by b^j and b^j_{-i} the jth highest bid from all bidders including and excluding i, respectively; then

$$T_i^{\mathbf{a}_{-i}}(\alpha_j) = \sum_{i=j}^{m} (\alpha_i - \alpha_{i+1}) b^j_{-i}. \tag{24}$$

The revenue in GFP is $\textsc{Rev}(M(\mathbf{a})) = \sum_j \alpha_j b^j \geq \sum_j \alpha_j b^j_{-i}$. For any slot j, the partial threshold amount for the bidder allocated j in the alternate allocation is less than payment of the bidder who won the slot j: $\alpha_j b^j \geq \sum_{i=j}^{m} (\alpha_i - \alpha_{i+1}) b^j_{-i}$. Summing over all bidders gives $\textsc{Rev}(M(\mathbf{a})) \geq \sum_i T_i^{\mathbf{a}_{-i}}(x'_i)$, our desired result. □

We now show that GFP is revenue-covered with general distributions over bids, by showing that partial thresholds with distributions over bids are less than the expectation of the partial thresholds with deterministic actions:

$$T_i(y_i) \leq \mathbb{E}_{\mathbf{a}_{-i}} \left[T_i^{\mathbf{a}_{-i}}(y_i) \right]. \tag{25}$$

Proof of Lemma 37. Fix a distribution over bids $\mathbf{G}$. Consider the problem of a bidder who can react to the bids of other bidders, and wants to bid so as to minimize her expected threshold while getting allocation at least y_i. She solves the following optimization problem:

$$\begin{aligned} \min_{\pi(\cdot)} \; & \mathbb{E}_{\mathbf{a}_{-i}} \left[T_i^{\mathbf{a}_{-i}}(x_i(\pi(\mathbf{a}_{-i}), \mathbf{a}_{-i})) \right] \\ \text{s.t. } & \mathbb{E}_{\mathbf{a}_{-i}} \left[x_i(\pi(\mathbf{a}_{-i}), \mathbf{a}_{-i}) \right] \geq y_i \end{aligned} \tag{26}$$

One strategy is to always bid to get allocation y_i, no matter what the cost. This strategy, however, is not optimal - the optimal strategy comes from equating marginal costs, and always buying allocation up to a price such that the average allocation purchased is at least y_i. That price is exactly $t_i(y_i)$. Thus, always bidding $t_i(y_i)$ is exactly the expected threshold minimizing strategy that solves (26), and hence comparing with always bidding to get allocation y_i gives

$$\mathbb{E}_{\mathbf{v}} \left[T_i^{\mathbf{a}_{-i}}(\tilde{x}_i(t_i(y_i), \mathbf{a}_{-i})) \right] \leq \mathbb{E}_{\mathbf{v}_{-i}} \left[T_i^{\mathbf{a}_{-i}}(y_i) \right]. \tag{27}$$

Equation (27) now allows us to derive general revenue covering from revenue covering with deterministic bids, Proposition 58. Noting that the left side of Equation (27) is $T_i(y_i)$ and taking expectation over partial thresholds with deterministic actions gives our desired result,

$$\begin{aligned} T_i(y_i) &= \mathbb{E}_{\mathbf{a}_{-i}} \left[T_i^{\mathbf{a}_{-i}}(\tilde{x}_i(t_i(y_i), \mathbf{a}_{-i})) \right] && (28) \\ &\leq \mathbb{E}_{\mathbf{a}_{-i}} \left[T_i^{\mathbf{a}_{-i}}(y_i) \right] && (29) \\ &\leq \mathbb{E}_{\mathbf{a}_{-i}} \left[\textsc{Rev}(M, \mathbf{G}) \right] && (30) \\ &= \textsc{Rev}(M, \mathbf{G}). \end{aligned}$$

□

E All-Pay Mechanisms

The proof of Theorem 50 lost a factor of 2 translating all-pay bids into their first-price equivalents. By moving from a first-price centric version of the value covering and revenue covering framework to one which works directly in terms of all-pay bids, we can match the welfare results of Syrgkanis and Tardos (2013). We can also derive a tighter revenue result with duplicates than was possible in the first-price-based framework.

The key quantity in our framework was the expected first-price threshold bid. The key step to deriving an all-pay native version of our framework is to switch to expected all-pay threshold bids. To this end, let $\tilde{t}_i(z)$ be the inverse of the CDF of agent i's all-pay threshold bid. Formally, $\tilde{t}_i(z) = \min\{b \,|\, \tilde{x}_i(b) \geq z\}$. As in the first-price auction, we can compute the expected value of this threshold bid as $\tilde{T}_i = \int_0^1 \tilde{t}_i(z)\, dz$. For the first-price auction, we used the payment semantics to derive a distribution of threshold bids for which i would be indifferent between all bids less than v_i. We can do the same thing for the all-pay auction and get the following result:

Lemma 59 (All-Pay Value Covering)**.** *For any BNE of an all-pay auction and agent i with value v_i,*

$$u_i(v_i) + \tilde{T}_i \geq \frac{v_i}{2}.$$

Proof. The proof parallels that of Lemma 4 - we lower bound $\tilde{T}_i$ using the payment semantics of the all-pay auction, then minimize the lower bound. As with Lemma 4, the deviation-based approach of Syrgkanis and Tardos (2013) also suffices.

Lowerbounding $\tilde{T}_i$. Bidder i chooses a best response bid a_i which maximizes her utility, $\tilde{u}_i(a_i) = v_i\tilde{x}_i(a_i) - a_i$. It follows that for any other deviation bid d, $\tilde{u}_i(a_i) \geq v_i\tilde{x}_i(d) - d$. Rearranging, we get $\tilde{x}_i(d) \leq \frac{\tilde{u}_i(a_i)+d}{v_i}$. Since $\tilde{x}_i(d)$ is the CDF of i's threshold bid, we can lower bound $\tilde{T}_i$ by integrating above the curve $\frac{\tilde{u}_i(a_i)+d}{v_i}$. In other words: $\tilde{T}_i \geq \int_0^1 \max(0, v_i z - \tilde{u}_i(a_i))\, dz$. Call the latter quantity $\underline{\tilde{T}_i}$.

Optimizing $\underline{\tilde{T}_i}$. Evaluating the integral for $\underline{\tilde{T}_i}$ gives $\underline{\tilde{T}_i} = v_i/2 - \tilde{u}_i(a_i) + \tilde{u}_i(a_i)^2/2v_i$, hence $u_i + \tilde{u}_i(a_i) = v_i/2 + \tilde{u}_i(a_i)^2/2v_i$. Holding v_i fixed and minimizing with respect to $\tilde{u}_i(a_i)$ yields a minimum at $u_i(a_i) = 0$, hence $\tilde{u}_i(a_i) + \underline{\tilde{T}_i} \geq v_i/2$. Using the facts that $\tilde{u}_i(a_i) = u_i(v_i)$ and $\tilde{T}_i \geq \underline{\tilde{T}_i}$ yields the result. ☐

As in the original framework, value covering characterizes the tradeoff between an agent's utility and the difficulty they face getting allocated. Now, however, the latter quantity is represented by $\tilde{T}_i$, which comes from all-pay rather than equivalent first-price bids. In proving revenue covering, we can therefore skip the translation from all-pay bids to equivalent first-price bids, yielding revenue covering with $\mu = 1$:

Lemma 60 (All-Pay Revenue Covering)**.** *For any distribution over actions* **G** *and agent i in the all-pay auction, the expected revenue is at least $\tilde{T}_i$.*

Proof. The revenue of the all-pay auction is expected sum of all bids. This is at least the expected highest bid from all agents except i, which is exactly $\tilde{T}_i$. ☐

We may combine our revenue and value covering lemmas in the manner used to prove Theorem 3 to produce the welfare bound of Syrgkanis and Tardos (2013):

Theorem 61. *The welfare in any BNE of the all-pay auction is at least a 2-approximation to the welfare of the welfare optimal mechanism.*

Furthermore, from Lemma 59, we can derive a virtual value covering result:

Lemma 62 (All-Pay Virtual Value Covering)**.** *For any BNE of the all-pay auction, any agent i with value v_i such that $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + \tilde{T}_i \geq \frac{\phi_i(v_i)}{2}.$$

Proof. Since $v_i x_i(v_i) \geq u_i(v_i)$, we have $v_i x_i(v_i) + \tilde{T}_i \geq \frac{e-1}{e} v_i$. Using the fact that $\phi_i(v_i) \leq v_i$ produces the desired inequality. □

For revenue, combining this with revenue covering and Lemma 13 as we did for the first-price auction in Section 4.3 yields:

Theorem 63. *The revenue in any BNE of the all-pay auction with at least* 2 *bidders from each distribution is at least a 6-aproximation to the revenue of the optimal mechanism.*

Finally, note that Lemma 60 (and therefore Theorem 61) can be extended to greedy auctions in general single-parameter environments using the approach used to derive Theorem 28.