Price of Anarchy for Auction Revenue*

Jason Hartline Darrell Hoy Sam Taggart

Northwestern University

April 25, 2014

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 *price of anarchy* results for social welfare and revenue. Our approach separates the standard smoothness framework [e.g., 16] into two distinct parts. The first part, *value covering*, employs best-response analysis to individually relate each agent's expected price for allocation and welfare in *any Bayes-Nash equilibrium.* The second part, *revenue covering*, uses properties of an auction's rules and feasibility constraints to relate the revenue of the auction to the agents' expected prices for allocation (not necessarily in equilibrium). Because value covering holds for any equilibrium, proving an auction is revenue covered is a sufficient condition for approximating optimal welfare, and under the right conditions, the optimal revenue. In mechanisms with reserve prices, our welfare results show approximation with respect to the optimal mechanism with the same reserves.

As a center-piece 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 [13] the auction's revenue is at least a $\frac{2e}{e-1} \approx 3.16$ approximation to the revenue of the optimal auction revenue. We also give bounds for matroid auctions with first price or all-pay semantics, and the generalized first price position auction. Finally, we give an extension theorem for simultaneous composition, i.e., when multiple auctions are run simultaneously, with single-valued and unit demand agents.

*We thank Vasilis Syrgkanis for comments on a prior version of this paper for which simultaneous composition did not hold, suggesting study of the simultaneous composition setting 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. Methods from the price of anarchy provide 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 analyzing the price of anarchy for revenue.

Equilibrium requires that each agent's strategy be a best response to the strategies of others. A typical price-of-anarchy analysis obtains a bound on the social welfare (the sum of the revenue and all agent utilities) from a lower bound 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 looks at identifying a mechanism with optimal performance in equilibrium. Optimal mechanisms tend to be complicated and impractical; consequently, a recent branch of mechanism design has looked at quantifying the loss between simple mechanisms and optimal mechanisms. These simple (designed) mechanisms have carefully constructed equilibrium (typically, the truthtelling equilibrium). The restriction to truthtelling equilibrium, though convenient in theory, is problematic in practice [1]. 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 price of anarchy literature instead considers the analysis of 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 higher chance of winning (which is good) but higher payments (which is bad). This first-price auction is the most 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 then first-price equilibria are well-behaved: the symmetry of the setting enables the easy solving for equilibrium [10], the equilibrium is unique [3, 11, 12], and the highest valued agent always wins (i.e., the social welfare is maximized). When the agents' values are non-identically distributed, analytically solving for equilibrium is notoriously difficult. For example, Vickrey [18] 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 [7].

The intractibility of solving analytically for equilibrium is foremost a problem of theory. It does not rule out BNE as a practical concept: agents can reach equilibrium by playing learning strategies, numerically solving the differential equations implied by equilibrium, etc. Free from the demands of theoretical analysis, agents may use these heuristic techniques, may focus on specific instances of their optimization problem, and may employ algorithmic techniques such as those developed by Jiang and Leyton-Brown [6], rather than pursuing a general, analytical characterization.

Price-of-anarchy analysis allows us to make general statements about equilibrium nonetheless. For example, a recent analysis of Syrgkanis and Tardos [16] shows 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 price-of-anarchy analysis sidesteps the intractability of solving for equilibrium and instead derives its bounds from simple best-response arguments.

1.1 Methods

Our analysis breaks down the problem of analyzing welfare and revenue into two parts. The first part, *value covering*, considers each agent individually and requires that an agent's contribution to BNE welfare and the expected price for allocation she faces combine to approximate her contribution to the welfare in the optimal mechanism. It uses only properties of BNE. The second part, *revenue covering*, captures the relevant mechanism-specific details and considers the auction rules in aggregate across the agents. It requires that the auction's expected revenue approximately covers the effective prices for an allocation across agents. The two parts combine to give a price of anarchy bounds for welfare. More importantly, proving these welfare bounds in this manner allows us to extend the same approximation with reserve prices and to revenue.

Our analysis begins by translating the payments in any auction into *equivalent bids*: the first-price bids or payments if the payment rule of the mechanism used first-price semantics. Beyond first-price auctions, it allows us to simplify the action space and the optimization problem a bidder faces into effectively the same problem a bidder in the first price auction faces. From this standard viewpoint we show that an agent's welfare in Bayes-Nash equilibria of *any auction* and her equivalent bid thresholds, combine to cover an $(e-1)/e$ fraction of her welfare in the optimal mechanism. Intuitively, either the agent's welfare is high, or the price she has to pay for allocation is high relative to her value for service.

We then make use of the characterization of revenue in Bayes-Nash equilibrium of Myerson [13] to reduce revenue to welfare. Value covering has a direct analog in terms of positive virtual values. Combined with revenue covering, this implies an approximation result for the virtual welfare for agents with positive virtual values. We provide several ways to then prove that a revenue covered mechanism has approximately optimal revenue for bidders with regular distributions.

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 price of anarchy results with both first-price and all-pay payment semantics. The first-price 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. These results are compatible with reserve prices. 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.

Welfare. In first-price auctions, we show that the price of anarchy for welfare is at most $2e/(e-1)$, with or without reserves. These results also extend to the generalized first-price position auction. For all-pay auctions in the above environments, the price of anarchy for welfare is $3e/(e-1)$. Tighter versions of these results with no reserves are known via Syrgkanis and Tardos [16]; the results with reserves are new.

Revenue. For first-price auctions with monopoly reserves in regular, single-parameter environments, we show that the price of anarchy for revenue is at most $2e/(e-1)$. 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 price of anarchy for revenue in first-price auctions is at most $3e/(e-1)$; in all-pay auctions, at most $4e/(e-1)$.

Simultaneous Composition. We also show via an extension theorem that the above bounds 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 [14] and the smooth mechanisms extension of Syrgkanis and Tardos [16]. Using this framework, one can show many properties based on a simple, full-information property, smoothness. In addition, Syrgkanis and Tardos [16] show that the smoothness guarantees hold under sequential and simultaneous composition. Our framework differs from smoothness in three notable ways. First, we decompose smoothness into two components, value- and revenue-covering, and argue about individual agents approximating their contribution to the optimal welfare and revenue. Second, we only consider the Bayesian setting, which allows us to use the BNE characterization for revenue, and allows us to relate other auctions to the first-price auction via equivalent bids. Third, equivalent bids allow us to eschew the deviations in the definitions of smoothness.

There have been a number of papers looking at revenue guarantees for the welfare-optimal Vickrey-Clarke-Groves (VCG) mechanism in asymmetric settings. Hartline and Roughgarden [5] show that VCG with monopoly reserves or duplicate bidders achieves revenue that is a constant approximation to the revenue optimal auction. Dhangwatnotai et al. [4] show that the single-sample mechanism, which is essentially VCG with a reserve sampled from all of the distributions of the bidders, achieves approximately optimal revenue in broader settings. Roughgarden et al. [15] 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, a number of papers have explored properties of asymmetric first-price auctions. Kirkegaard [8] shows that understanding the ratios of expected payoffs in equilibrium can be easier than understanding equilibrium and lead to insights about equilibria. Kirkegaard [9] shows that some properties of distributions can be used to compare revenue of the first price auction to revenue of the second price auction. Lebrun [11] and Maskin and Riley [12] 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 utility. Each agent has a private value for service, v_i, drawn independently from a distribution F_i over V_i, the agent's valuation space. We write $\mathbf{F} = \prod_i F_i$ and $\mathbf{V} = \prod_i V_i$ to denote the joint value distribution and space of value profiles, respectively. A *mechanism* consists of an allocation rule $\tilde{\mathbf{x}}$ and a payment rule $\tilde{\mathbf{p}}$, mapping actions of agents to allocations and payments respectively. Each agent i draws their private value v_i from F_i and selects an action according to some strategy $s_i : V_i \to A_i$, where A_i is the set of possible actions for 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 $\tilde{\mathbf{x}}(\mathbf{a})$ and $\tilde{\mathbf{p}}(\mathbf{a})$. Each agent's utility is $\tilde{u}_i(\mathbf{a}) = v_i\tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})$.

Mechanisms typically operate with constraints on permissible allocations. A *feasibility environment* is a set of feasible allocation vectors. Mechanisms for a feasibility environment choose only allocations from the feasible set. The simplest example is a single-item auction, in which at most one person at a time can be served. This paper assumes feasibility environments are *downward-closed*: if $(x_1, \ldots, x_k, \ldots, x_n)$ is feasible, so is $(x_1, \ldots, 0, \ldots, x_n)$. We will often consider the special case of *matroid* environments, in which the set of feasible allocations correspond to the independent sets of a matroid set system.

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 $\tilde{x}_i(a_i) = \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$, with $\tilde{p}_i(a_i)$ and $\tilde{u}_i(a_i)$ defined analogously. Given $\mathbf{s}$, 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 can 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': $\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{s}(\mathbf{v}))] \geq \mathbf{E}_{\mathbf{v}_{-i}}[u_i(a', \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$.

We consider only mechanisms where agents can gain from participation, regardless of their value - that is, we require that mechanisms be *interim individually rational*. We implement this by assuming each agent can withdraw from the mechanism. Specifically, define a withdraw action as any action w_i such that $\tilde{x}_i(w_i, \mathbf{a}_{-i}) = 0$ and $\tilde{p}_i(w_i, \mathbf{a}_{-i}) = 0$ for any value of $\mathbf{a}_{-i}$. We assume all mechanisms have at least one such action for each agent. In any BNE, each agent has the option to withdraw and must therefore get nonnegative utility.

Myerson [13] characterizes the interim allocation and payment rules that arise in BNE. These results are summarized in the following theorem.

Theorem 1 (13)**.** *For any mechanism and value distribution* $\mathbf{F}$*,*

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 $\mathbf{x}(\mathbf{v})$ must therefore have the same interim payments as well.*

Mechanism Design Objectives We consider the problem of maximizing two main objectives in BNE: expected welfare and expected revenue. The revenue of a mechanism M is the total payment of all agents. Mechanism M's expected revenue for $\mathbf{v} \sim \mathbf{F}$ in a given Bayes-Nash equilibrium $\mathbf{s}$ is denoted $\textsc{Rev}(M) = E_{\mathbf{v}}[\sum_i p_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\sum_i \phi_i(v_i)x_i(\mathbf{v})]$. The welfare of a mechanism M is the total utility of all participants including the auctioneer; its expected welfare is denoted $\textsc{Welfare}(M) = \textsc{Rev}(M) + E_{\mathbf{v}}[\sum_i v_i x_i(\mathbf{v}) - p_i(\mathbf{v})] = E_{\mathbf{v}}[v_i x_i(\mathbf{v})]$. We will also refer to welfare throughout the paper as surplus.

Our welfare benchmark is pointwise-optimal feasible allocation. That is, we seek to approximate $\textsc{Welfare}(\textsc{Opt}) = \mathbf{E}_{\mathbf{v}}[\max_{\mathbf{x}^*} \sum_i v_i x_i^*]$. This can be implemented via the Vickrey-Clarke-Groves (VCG) mechanism. We measure a mechanism M's welfare performance by the *Bayesian price of anarchy for welfare*, given by $\max_{\mathbf{F},\mathbf{s}\in \mathrm{BNE}(M,\mathbf{F})} \textsc{Welfare}(\textsc{Opt})/\textsc{Welfare}(M)$, where $\mathrm{BNE}(M,\mathbf{F})$ is the set of BNE for M under value distribution $\mathbf{F}$.

To understand revenue, we rely on the alternate characterization derived in Myerson [13]:

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

Using this result, Myerson [13] 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 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 $\sum_i \phi_i(v_i)x_i(\mathbf{v})$. We will seek to minimize the *Bayesian price of anarchy for revenue*, $\max_{\mathbf{F}\in\mathcal{R},\mathbf{s}\in\mathrm{BNE}(M,\mathbf{F})} \textsc{Rev}(\textsc{Opt}_{\mathbf{F}})/\textsc{Rev}(M)$, where $\mathcal{R}$ is the set of regular distributions and $\textsc{Opt}_{\mathbf{F}}$ is the Bayesian revenue-optimal mechanism for value distribution $\mathbf{F}$.

3 Single-Item First Price Auction with Reserves

We begin by analyzing the single-item first price auction with per-bidder reserves, and show that it approximates the welfare of the optimal mechanism with the same reserves. With zeroed reserves, this result implies that the welfare of the first-price auction with no reserves approximates the welfare of the welfare optimal auction. We will then connect these results to revenue approximation results, and taking the reserves to be the monopoly reserves gives revenue (and welfare) approximation results with respect to the revenue optimal mechanism.

3.1 Welfare

We now aim to show that the welfare and revenue of the first price auction together approximate the welfare of the optimal auction:

$$\textsc{Welfare}(\mathrm{FPA}_{\mathbf{r}}) + \textsc{Rev}(\mathrm{FPA}_{\mathbf{r}}) \geq \tfrac{e-1}{e}\textsc{Welfare}(\textsc{Opt}_{\mathbf{r}}). \tag{1}$$

Our proof will proceed by first analyzing the optimization problem of the bidder, then relating that optimization problem to welfare and revenue. We will conclude with the following theorem.

Theorem 3. *The welfare in any BNE of the first price auction with reserves* $\mathbf{r}$ *is at least a* $\frac{2e}{e-1}$*-approximation to the welfare of the welfare optimal mechanism that serves no agent with* $v_i < r_i$.

Note that equation (1) is quite similar to the inequality in the smooth games and mechanism frameworks [16, 14]. It differs primarily in that we are not defining a specific deviation but

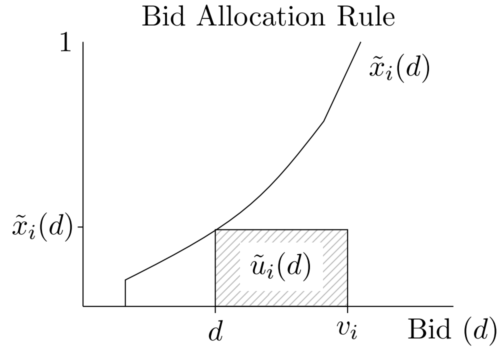

Figure 1: For any bid d, the area of a rectangle between $(d, \tilde{x}_i(d))$ and $(v_i, 0)$ on the bid allocation rule is the expected utility $\tilde{u}_i(d)$. The BNE bid b_i is chosen to maximize this area.

deriving bounds explicitly from BNE. Moreover we show that (a variant of) equation (1) holds in the interim for every bidder and realized valuation, rather than only in aggregate.

A bidder's optimization problem. Consider the optimization problem faced by a bidder i with value v_i in the first price auction. A bidder's expected utility over possible bids d is $\tilde{u}_i(d) = (v_i - d)\tilde{x}_i(d)$, where $\tilde{x}_i(d)$ is the interim bid allocation rule faced by the bidder. Let b_i be her best response bid given her value v_i. If we plot the bid allocation rule $\tilde{x}_i(d)$ for any alternate bid d, then $\tilde{u}_i(b_i)$ is precisely the area of the rectangle in the lower right of Figure 1.

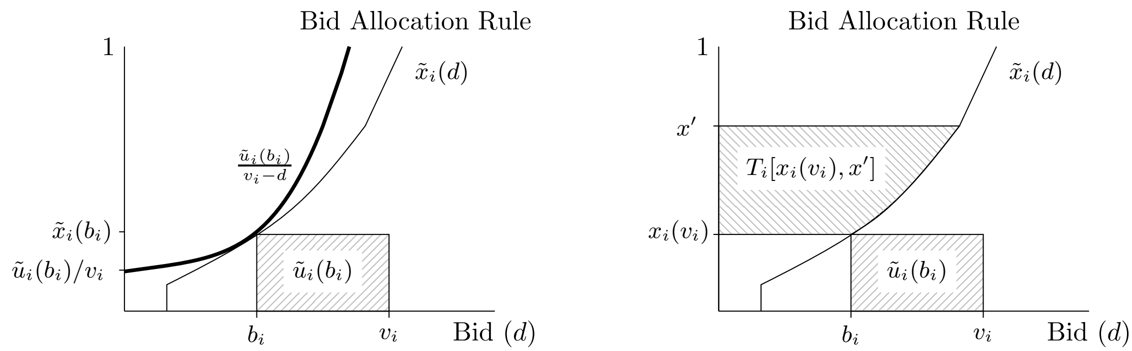

(a) As b_i is a best-response to the actions of other agents, the indifference curve $\tilde{u}_i(b_i)/(v_i - d)$ upper bounds $\tilde{x}_i(d)$.

(b) The additional threshold bid $T_i[x_i(v_i), x']$ prevents bidders from bidding to receive allocation x'.

Figure 2

When other bidders have realized values and submitted bids, there is a minimum or threshold bid a bidder must make to win, $\tau_i(\mathbf{v}_{-i}) = \max(r_i, \max(b_{-i}(\mathbf{v}_{-i})))$, the maximum of a player's reserve and the bids of all other bidders. We call this bidder i's *pointwise threshold bid.* As we are in the Bayesian setting, a bidder is not reacting to this pointwise threshold, but is acting in expectation over the types and actions of her competitors. These actions induce a distribution over threshold bids. The cumulative distribution function of threshold bids for a bidder i is precisely her bid allocation rule $\tilde{x}_i$.

We will also refer to thresholds using the probability of allocation that they represent achieving. Let $\tau_i(x)$ refer to the smallest bid that achieves allocation of at least x, hence $\tau_i(x) = \min\{\ b \mid \tilde{x}_i(b) \geq x\}$. Let $B_{-i}(b)$ be the cumulative distribution function of the highest

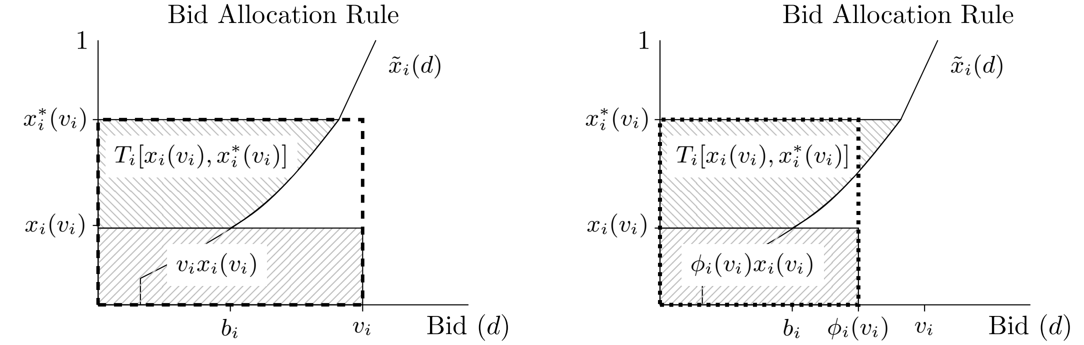

(a) Lemma 4 shows the shaded areas cover a $(e-1)/e$ fraction of the dashed box, bidder i's contribution to the optimal welfare ($v_i x_i^*(v_i)$).

(b) Lemma 7 shows the shaded areas cover an $(e-1)/e$ fraction of i's contribution to the optimal virtual welfare, $\phi_i(v_i)x_i^*(v_i)$. The proof can be done by comparison with Figure 3a.

Figure 3

bids from other bidder. Then $\tau_i(x)$ is either the reserve r_i or the bid required to beat the highest bid from other agents a x fraction of the time, $\tau_i(x) = \max(r_i, B_{-i}^{-1}(x))$.[1]

For an alternate allocation probability x', the additional threshold $T_i[x_i, x'] = \int_{x_i}^{x'} \tau_i(z)\, dz$ will be used as a measure of how much more expensive it is for a bidder to get allocation $x' > x_i$. This is illustrated in Figure 2b.

Relating Contribitions to First-Price and Optimal Welfare. Let $\mathbf{x}^*$ be the allocation rule from the welfare optimal mechanism that serves no agent with value $v_i < r_i$, $\textsc{Opt}_{\mathbf{r}}$. Thus $\textsc{Welfare}(\textsc{Opt}_{\mathbf{r}}) = \sum_i \mathbf{E}_{v_i}[v_i x_i^*(v_i)]$, and we can view $v_i x_i^*(v_i)$ as a bidder's contribution to the optimal welfare. We will now aim to approximate each bidder's contribution individually, using the bidder's contribution to welfare in the first-price auction, i.e., $v_i x_i(v_i)$, and a fraction of the revenue in the first-price auction:

Our proof proceeds in two steps:

1. *Value Covering*: A bidder's contribution to welfare in the $\text{FPA}_{\mathbf{r}}$ and *additional threshold* together approximate her contribution to welfare in any alternate allocation. (Lemma 4)
2. *Revenue Covering*: The revenue of the $\text{FPA}_{\mathbf{r}}$ approximates the additional threshold for all agents. (Lemma 5)

The final approximation result follows by summing the value covering condition across agents, taking expectation over values, and combining with revenue covering.

Lemma 4 (Value Covering)**.** *For any bidder i with value $v_i \geq r_i$ in a BNE of the* $\text{FPA}_{\mathbf{r}}$ *and alternate feasible allocation x',*

$$v_i x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e} v_i x'. \qquad (2)$$

When value covering is used to approximate the welfare induced by an allocation rule x_i^*, the alternate allocation x' used for every bidder and value will be precisely $x' = x_i^*(v_i)$.

Proof. We will prove value covering in two steps: first, by developing a lower bound $\underline{T}$ on the additional threshold T; second, by optimizing the lower bound to get the right side of (2).

[1] If B_{-i} is not invertible, then define $B_{-i}^{-1}(x)$ to be the function $B_{-i}^{-1}(x) = \inf\{b \mid \tilde{x}_i(b) \geq x\}$.

Lowerbounding T In best responding, bidder i chooses an action which maximizes her utility. If b_i is a best response bid, then for any alternate bid d, $\tilde{u}_i(b_i) \geq (v_i - d)\tilde{x}_i(d)$, hence $\tilde{x}_i(d) \leq \frac{\tilde{u}_i(b_i)}{v_i - d}$. With equality, this bound gives an indifference curve for bidder i; it is the alternate bid allocation rule that would lead to her being indifferent over all reasonable bids (see Figure 2a). Call $\underline{T}_i[x_i(v_i), x']$ the expected threshold bid from the indifference curve, then $\underline{T}_i[x_i(v_i), x'] = \int_{x_i(v_i)}^{x'} \max(0, v_i - u_i(v_i)/z)\ dz$ and

$$v_i x_i(v_i) + T_i[x_i(v_i), x'] \geq v_i x_i(v_i) + \underline{T}_i[x_i(v_i), x'] \geq u_i(v_i) + \underline{T}_i[u_i(v_i)/v_i, x']. \tag{3}$$

The last inequality followed because $b_i(v_i)x_i(v_i) \geq \underline{T}_i[u_i(v_i)/v_i, x_i(v_i)]$ if $x_i(v_i) > u_i(v_i)/v_i$.

Optimizing $\underline{T}_i$ Evaluating the integral for $\underline{T}_i[u_i(v_i)/v_i, v_i]$ gives $u_i(v_i) + \underline{T}_i[u_i(v_i)/v_i, x'] = v_i x' + u_i(v_i) \ln \frac{u_i(v_i)}{v_i x'}$. Holding $v_i x'$ fixed and minimizing with respect to $u_i(v_i)$ yields a minimum at $u_i(v_i) = \frac{v_i x'}{e}$, hence

$$u_i(v_i) + \underline{T}_i[u_i(v_i)/v_i, x'] \geq \tfrac{e-1}{e} v_i x'. \tag{4}$$

Combining (3) and (4) gives exactly our desired result, (2). □

We now show that in the first price auction, the expected revenue is greater than the additional threshold bids for any alternate feasible allocation $\mathbf{x}'$, which we can then combine with value covering to give a welfare approximation result. While value covering depended critically on equilibrium (or at least on bidders best responding), revenue covering will only depend on the form of the first price auction, and will thus hold for arbitrary (not necessarily BNE) bidding strategies that satisfy a light participation requirement (that is always satisfied in BNE). We call a bidding strategy *participatory* if bidders always bid at least their reserve unless no bid in $[r_i, v_i]$ gives positive probability of allocation.

Lemma 5 (Revenue Covering). *For any participatory bidding strategies* $\mathbf{s}$ *and alternative allocation* $\mathbf{x}'$,

$$\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \sum_{i\,:\,v_i \geq r_i} T_i[\tilde{x}_i(s_i(v_i)), x'_i]. \tag{5}$$

Proof. It suffices to show that for bidder i with value above her reserve,

$$\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq T_i[\tilde{x}_i(s_i(v_i)), 1]. \tag{6}$$

Once (6) is shown, multiplying by x'_i, summing over agents, and observing that $T_i[\tilde{x}_i(s_i(v_i)), z]$ is convex in z concludes the lemma:

$$\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \textstyle\sum_i x'_i T_i[\tilde{x}_i(s_i(v_i)), 1] \geq \textstyle\sum_i T_i[\tilde{x}_i(s_i(v_i)), x'_i]. \tag{7}$$

We now show (6). By the participatory assumption, a bidder i with value above her reserve bids above her reserve r_i. Thus the additional threshold when playing $s_i(v_i)$ is bounded by the additional threshold when bidding the reserve, hence $T_i[\tilde{x}_i(s_i(v_i)), 1] \leq T_i[\tilde{x}_i(r_i), 1]$. Then, using the definition of T and the monotonicity of $B_{-i}^{-1}(x)$, we have:

$$T_i[\tilde{x}_i(s_i(v_i)), 1] \leq T_i[\tilde{x}_i(r_i), 1] = \int_{\tilde{x}_i(r_i)}^{1} \max(r_i, B_{-i}^{-1}(z))\ dz = \int_{\tilde{x}_i(r_i)}^{1} B_{-i}^{-1}(z)\ dz. \tag{8}$$

As the revenue of a first price auction is the expected highest bid, and $B_{-i}^{-1}(z)$ is the inverse of the cumulative distribution function of highest bid from bidders aside from i, $\int_0^1 B_{-i}^{-1}(z)\, dz = \mathbf{E}_{\mathbf{v}}[\max_{j\neq i,\ s_i(v_j)\geq r_j} s_i(v_j)] \leq \textsc{Rev}(\text{FPA}_{\mathbf{r}})$. Combining this with monotonicity of $B_{-i}^{-1}(z)$ gives

$$T_i[\tilde{x}_i(s_i(v_i)), 1] \leq \int_{\tilde{x}_i(r_i)}^{1} B_{-i}^{-1}(z)\, dz \leq \int_0^1 B_{-i}^{-1}(z)\, dz \leq \textsc{Rev}(\text{FPA}_{\mathbf{r}}). \tag{9}$$

Chaining (8) and (9) gives $\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq T_i[\tilde{x}_i(s_i(v_i)), 1]$. □

We now combine value and revenue covering to attain an approximation to the optimal welfare.

Proof of Theorem 3. We apply value covering and revenue covering with $\mathbf{x}' = \mathbf{x}^*(\mathbf{v})$. Taking expectation of (2) over all players and values and combining with (5) gives $\textsc{Welfare}(\text{FPA}_{\mathbf{r}}) + \textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \frac{e-1}{e}\textsc{Welfare}(\textsc{Opt}_{\mathbf{r}})$. As $\textsc{Welfare}(\text{FPA}_{\mathbf{r}}) \geq \textsc{Rev}(\text{FPA}_{\mathbf{r}})$, $\textsc{Welfare}(\text{FPA}_{\mathbf{r}})$ is then a $2e/(e-1)$ approximation to $\textsc{Opt}_{\mathbf{r}}$. □

The following are the main ideas and differences between the proof above and the proof of Syrgkanis and Tardos [16] that enables treatment of reserve prices. The Syrgkanis and Tardos [16] result can be viewed as combining value covering and revenue covering in one step (via the smoothness definition). Their equation has bidders' utilities where we have the bidders' surpluses and they have the full expected threshold where we have the additional threshold. The thresholds that a bidder faces that correspond to bids of other bidders translate to revenue and can be thus bounded by a revenue covering argument. Reserve prices, however, induce thresholds that do not correspond to bids of other bidders. A participatory bidder, however, will bid above the reserve when her value is above the reserve. Therefore, this bidder's payment will always compensate for the part of the threshold distribution that corresponds to the reserve price. Because we use surplus instead of utility our analysis loses a factor of two on the no-reserves bound of Syrgkanis and Tardos [16].[2]

3.2 Revenue

In the tradition of Bayesian mechanism design, we will prove the revenue approximation result by reducing to the welfare approximation above. Let x_i^* now denote the allocation rule from the revenue optimal auction, given by Theorem 1. For revenue, we will instead approximate each bidders contribution to the optimal *virtual welfare*, $\phi_i(v_i)x_i^*(v_i)$ when $\phi_i(v_i) \geq 0$. In a regular environment, monopoly reserves at $r_i = \phi_i(v_i)$ for each bidder will result in no bidder being served with a negative virtual value. Thus approximating the optimal virtual surplus using only agents with positive virtual values will be sufficient to approximate the expected surplus of the optimal auction.

Theorem 6. *In any BNE of the first price auction with monopoly reserves (*$\text{FPA}_{\mathbf{r}}$*) in a regular environment, the revenue is at least a* $\frac{2e}{e-1}$*-approximation to revenue of the optimal auction.*

Recall that in the welfare proof, the expected threshold bid plus BNE welfare approximated each bidders contribution to optimal welfare (Lemma 5). For revenue, we will use the expected threshold bid and each bidders BNE virtual welfare to approximate their virtual welfare in $\textsc{Opt}$.

[2]The bound of $e/(e-1) \approx 1.58$ is not known to be tight for the single-item first-price auction. An example is provided in Appendix A that shows the worst known price of anarchy of 1.15.

Lemma 7 (Virtual Value Covering)**.** *In any BNE of the* $\text{FPA}_{\mathbf{r}}$*, for any bidder* i *with value* $v_i \geq r_i$*,* $\phi_i(v_i) \geq 0$ *and alternate allocation* x'*,*

$$\phi_i(v_i)x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e}\phi_i(v_i)x'. \tag{10}$$

Proof. This follows directly from value covering (Lemma 4) — see Figure 3b for an illustration. Combining $0 \leq \phi_i(v_i) \leq v$ with (12) gives

$$\tfrac{1}{\phi_i(v_i)x'}\left(\phi_i(v_i)x_i(v_i) + T_i[x_i(v_i), x']\right) \geq \tfrac{1}{v_i x'}\left(v_i x_i(v_i) + T_i[x_i(v_i), x']\right) \geq \tfrac{e-1}{e}.$$

Multiplying through by $\phi_i(v_i)x'$ gives our desired result. □

Proof of Theorem 6. As no agents with negative virtual values are served, the revenue is larger than both terms on the left side of (10) when summed over all agents and values, so

$$\begin{aligned}
\textsc{Rev}(\text{FPA}_{\mathbf{r}}) + \textsc{Rev}(\text{FPA}_{\mathbf{r}}) &\geq \sum\nolimits_i \mathbf{E}_{v_i}\left[T_i[x_i, x_i^*(v_i)]\right] + \sum\nolimits_i \mathbf{E}_{v_i}\left[\phi_i(v_i)x_i(v_i)\right] \\
&\geq \tfrac{e-1}{e}\sum\nolimits_i \mathbf{E}_{v_i}\left[\phi_i(v_i)x_i^*(v_i)\right] \\
&= \tfrac{e-1}{e}\textsc{Rev}(\textsc{Opt})
\end{aligned} \tag{11}$$

Thus $\textsc{Rev}(\text{FPA}_{\mathbf{r}})$ is at least a $\frac{2e}{e-1}$ approximation to $\textsc{Rev}(\textsc{Opt})$. □

4 Framework

In equilibria of the single-item first price auction, we observed that agents with low expected utility had high expected threshold bids. Because high thresholds were connected to high payments, we could conclude that the first price auction is both approximately welfare- and revenue-optimal. The goal for this section is to build up a framework for making this same argument for mechanisms with different payment semantics, such as all-pay auctions. In particular, we seek to prove results about behavior in Bayes-Nash equilibrium while *ignoring* the particular payment semantics of each auction. We begin by defining *equivalent bids*, which connect the optimization problem a bidder faces to the problem they would face in a first-price auction. This allows us to reuse machinery and construct a framework to understand how single-dimensional agents act in BNE for any mechanism.

4.1 Equivalent Bids

Utility-maximizing agents must balance two objectives: getting allocated frequently, and getting allocated cheaply. In a first-price auction, agents bid to explicitly specify the tradeoff they are willing to make: their bid is the price they pay per unit of allocation. In general mechanisms, for any agent i and any action a_i, define the *equivalent bid* for an action a_i to be $\beta_i(a_i) = \tilde{p}_i(a_i)/\tilde{x}_i(a_i)$; this can be thought of as the price per unit of allocation for that action. For first-price auctions, this is exactly the bid. For mechanisms with different payment semantics, $\beta_i(a_i)$ can still be thought of as an equivalent first-price bid for action a_i.

Equivalent Threshold Bid In proving Theorem 3, we noted that b is the minimum payment necessary to get the allocation probability $\tilde{x}_i(b)$. We used this property to bound the distribution of other agents' bids. For auctions where this relationship is less clear, we think of agents partitioning the actions in their choice set by interim allocation probability, then for each probability consider only the cheapest such action in terms of price per unit of allocation. For each allocation probability z, define $\alpha_i(z)$ to be the action which minimizes $\beta_i(\alpha_i(z))$ subject to $x_i(\alpha_i(z)) \geq z$, and the *equivalent threshold bid* $\tau_i(z)$ to be the value of $\beta_i(\alpha_i(z))$.[3] This captures exactly what we need: $\tau_i(z)$ is the minimum price per unit needed to get allocation probability at least z - exactly the notion that the inverse CDF of the first-price bid allocation rule satisfied in the proof of Theorem 3. Note that $\tau_i(z)$ depends on $\mathbf{s}$. For notational convenience, we suppress the strategy profile as an argument.

Cumulative Equivalent Threshold Bid We can now use $\tau_i(z)$ to track the expense an agent faces from increasing their allocation. Specifically, assume an agent is playing some action a_i and seeks to increase their allocation probability to x'. The barrier to i doing so is the collection of equivalent threshold bids in $[\tilde{x}_i(a_i), x']$. We can use this notion to measure i's expense for additional allocation. Define the *expected equivalent threshold bid* as $T_i[\tilde{x}_i(a_i), x'] = \int_{\tilde{x}_i(a_i)}^{x'} \tau_i(z)\, dz$.

If $x' \leq \tilde{x}_i(a_i)$, then define $T_i[\tilde{x}_i(a_i), x'] = 0$. This quantity will function identically to its counterpart in Section 3, trading off against i's surplus as in Lemma 4, and translating into revenue Lemma 5. Note that because $\tau_i(z)$ is nondecreasing in z, $T_i[\tilde{x}_i(a_i), x']$ is convex in x'.

4.2 Covering Conditions and the Price of Anarchy

With our machinery developed, we can now quantify the tradeoff between an agent's surplus and price of allocation just as we did in Lemma 4:

Lemma 8 (Value Covering)**.** *Consider a mechanism M in BNE with induced allocation and payment rules $(\mathbf{x}, \mathbf{p})$, and an agent i with value v_i. For any $x' \in [0, 1]$,*

$$v_i x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e} v_i x'. \tag{12}$$

The proof can now be done by reduction to the single-item first-price auction (Lemma 4) because bidders now face effectively the same optimization problem as in a single-item first-price auction. The proof is left to the full version of the paper. To prove an approximation result for welfare or revenue, the *only* mechanism-specific detail which remains is specifying the relationship between T_i and the mechanism's revenue.

Intuitively, we saw in Section 3 that if there is a relationship between revenue and the difficulty an agent faces in increasing their allocation once they have chosen to participate in the mechanism, then value covering allows us to show a welfare bound. To make this relationship concrete, we extend the definition of Lemma 5

Definition 9. *A mechanism M is* μ-revenue covered *if for any (implicit) strategy profile* $\mathbf{s}$*, feasible allocation* $\mathbf{x}'$*, and action profile* $\mathbf{a}$*,* $\mu \textsc{Rev}(M) \geq \sum_i T_i[\tilde{x}_i(a_i), x'_i]$.

Note that Definition 9 makes no mention of BNE. It must hold for any strategy profile. This is a stronger condition than Lemma 5, as it is not restricted to bidders with values above a set of reserves or bidders playing only participatory strategies.

[3] This minimization problem has a solution in the mechanisms this paper treats.

As we already saw, revenue covering has a number of important consequences. First is a welfare bound.

Theorem 10. *If a mechanism is μ-revenue covered, then in any BNE, it is a $(1+\mu)\frac{e}{e-1}-$approximation to the welfare of the optimal mechanism.*

Proof. Let $\mathbf{x}^*$ be the welfare-optimal allocation rule, and consider some value profile $\mathbf{v}$. Lemma 8 with $x' = x_i^*(\mathbf{v})$ yields that for each $\mathbf{x}$ and value v_i,

$$v_i x_i(v_i) + T_i[x_i(v_i), x_i^*(\mathbf{v})] \geq \tfrac{e-1}{e} v_i x_i^*(\mathbf{v}).$$

Summing over agents and using revenue covering yields $\sum_i v_i x_i(v_i) + \mu\textsc{Rev}(M) \geq \frac{e-1}{e}\sum_i v_i x_i^*(\mathbf{v})$. Taking expectation with respect to $\mathbf{v}$, we get $\textsc{Welfare}(M) + \mu\textsc{Rev}(M) \geq \frac{e-1}{e}\textsc{Welfare}(\textsc{Opt})$ and hence $(1+\mu)\frac{e}{e-1}\textsc{Rev}(M) \geq \frac{e}{e-1}\textsc{Welfare}(\textsc{Opt})$. □

4.3 Restricted Revenue Covering

Smoothness approaches hinge on proving price of anarchy bounds in a restricted way. The restricted proofs imply extensions to broader environments. Our framework operates in this spirit, and to obtain a reserves extension, we impose similar restrictions.

Restrictions for Reserves Revenue covering is useful for bounding revenue and welfare because it allows the mechanism to make up in revenue for the fact that in BNE, agents might find it too expensive increase their allocation probability to that which they would receive from an optimal allocation rule. That is, the agents for whom revenue covering matters are specifically those being served by the optimal mechanism. It therefore makes sense that if we seek to approximate the welfare or revenue of a mechanism that restricts those it serves, for example, with reserves, we only need revenue covering to pertain to agents served by the benchmark mechanism.

Reserves pose one additional problem for arguing about revenue covering. As we saw in Section 3, adding reserves to an auction changes the threshold bids an agent faces. Whereas with no reserves a threshold corresponded directly to revenue, a threshold in a mechanism with reserves may also correspond to the reserve itself and hence not revenue. We need a way of discerning which thresholds are useful for revenue and which are not.

We thus introduce revenue covering restricted to certain agents. This will allow us to prove approximation results for only a certain set of bidders — for example, the set of bidders with values above their reserves. We will use a function $S(\mathbf{v})$ to specify which such bidders are revenue covered. In the case of individual reserves $\mathbf{r} = (r_1, \ldots, r_n)$, $S(\mathbf{v}) = \{v_i \mid v_i \geq r_i\}$. It is these agents whose optimal welfare or virtual welfare we seek to approximate, and so they are the ones for whom we would like revenue covering to hold.

So long as such agents are bidding above their reserves if they have any chance of winning, revenue covering will hold. We call such actions *participatory*. Given a strategy profile $\mathbf{s}$ and value profile $\mathbf{v}$, define an action a_i to be participatory for $\mathbf{s}$ and $\mathbf{v}$ if $\beta_i(a_i) \leq v_i$ and either $\tilde{x}_i(a_i) > 0$ or there's no a_i' such that $\tilde{x}_i(a_i') > 0$ and $\beta_i(a_i') \leq v_i$.

Definition 11. *A mechanism M is μ-revenue covered with respect to S if for all alternate allocations $\mathbf{x}'$, value profiles $\mathbf{v}$, and profiles of participatory actions $\mathbf{a}$,*

$$\mu\textsc{Rev}(M) \geq \sum\nolimits_{i \in S(\mathbf{v})} T_i[\tilde{x}_i(a_i), x_i'].$$

Using our restriction to reserves, we can compare the welfare of a revenue-covered mechanism to that of the optimal mechanism with reserves. In Section 5, we show how to use these welfare results with reserves to bound revenue as well. The welfare theorem below, stated for general S, is most intuitive when S selects the agents in $\mathbf{v}$ who are above vector of reserves $\mathbf{r} = (r_1, \ldots, r_n)$.

Theorem 12. *For any S mapping value profiles to sets of agents, if a mechanism M is μ-revenue covered with respect to S, then the welfare of M is a $(1 + \mu)e/(e - 1)$-approximation to the welfare of the optimal mechanism which only serves agents in $S(\mathbf{v})$ for every $\mathbf{v}$.*

Proof. Let $\mathbf{x}^*$ be the welfare-optimal allocation rule, and consider some value profile $\mathbf{v}$. Lemma 8 with $x' = x_i^*(\mathbf{v})$ yields that for each $\mathbf{x}$ and value v_i,

$$v_i x_i(v_i) + T_i[x_i(v_i), x_i^*(\mathbf{v})] \geq \tfrac{e-1}{e} v_i x_i^*(\mathbf{v}).$$

Note that in BNE, agents always play participatory actions. It follows that summing over all agents in $S(\mathbf{v})$ and using revenue covering yields:

$$\sum\nolimits_{i \in S(\mathbf{v})} v_i x_i(v_i) + \mu \text{R}{\small\text{EV}}(M) \geq \tfrac{e-1}{e} \sum\nolimits_{i \in S(\mathbf{v})} v_i x_i^*(\mathbf{v}).$$

Taking expectation with respect to $\mathbf{v}$ and noting that $\mathbf{x}^*$ doesn't serve agents not in $S(\mathbf{v})$, we get $\text{W}{\small\text{ELFARE}}(M) + \mu \text{R}{\small\text{EV}}(M) \geq \frac{e-1}{e} \text{W}{\small\text{ELFARE}}(\text{O}{\small\text{PT}})$ and hence $(1 + \mu)\frac{e}{e-1} \text{R}{\small\text{EV}}(M) \geq \frac{e}{e-1} \text{W}{\small\text{ELFARE}}(\text{O}{\small\text{PT}})$. $\square$

Covering-Preserving Reserves In many environments, it is possible to add reserves to a revenue covered mechanism and preserve revenue covering in the sense of Definition 11. As a result, the BNE welfare of the reserves mechanism approximates the welfare of the optimal mechanism with the same reserves. In Appendix C, we provide a general set of conditions under which reserves preserve revenue covering in this manner. These conditions hold, for example, in first-price matroid and position auctions, as well as under simultaneous composition. These mechanisms with reserves consequently meet the conditions of Lemma 12, yielding a welfare approximation, and as we show in the next section, a revenue approximation.

5 Revenue Approximation

Recall that by Myerson's characterization of Bayes-Nash equilibrium (Lemma 2), the expected revenue can be viewed as the expected virtual welfare of agents served. We will consider the task of approximating the revenue of the optimal auction in two parts: showing that the virtual welfare from positive virtual-valued agents approximates the optimal revenue, and demonstrating a few methods to ensure that the virtual welfare from agents with negative virtual values does not hurt revenue too much.

5.1 Positive Virtual Value Approximation

In Theorem 6 of Section 3, we showed that the first-price auction with monopoly reserves had approximately optimal revenue, via a reduction to the welfare approximation result. We show in this section that the same approach suffices to show that for any μ-revenue covered mechanism, the revenue accounted for by positive virtual valued agents approximates the optimal revenue.

Definition 13. *Let the positive and negative virtual values for an agent be* $\phi_i^+(v_i) = \max(\phi_i(v_i), 0)$ *and* $\phi_i^-(v_i) = \min(\phi_i(v_i), 0)$ *respectively. Define the positive and negative virtual welfare of a mechanism to be* $\text{Rev}^+(M) = \sum_i \mathbf{E}_{v_i}[\phi_i^+(v_i)x_i(v_i)]$ *and* $\text{Rev}^-(M) = -\sum_i \mathbf{E}_{v_i}[\phi_i^-(v_i)x_i(v_i)]$.

By Theorem 1, $\text{Rev}(M) = \text{Rev}^+(M) - \text{Rev}^-(M)$. Our primary result in this section is that $\text{Rev}^+(M)$ is a constant approximation to the revenue of the optimal mechanism if M is μ-revenue covered. Thus, bounding the loss from Rev^- as a fraction of Rev^+ is sufficient to show approximately optimal revenue.

Theorem 14. *In any BNE of a* μ*-revenue covered, single-parameter mechanism* M*, the positive virtual welfare* $\text{Rev}^+(M)$ *is a* $(\mu+1)\frac{e}{e-1}$ *approximation to the revenue of the optimal mechanism. More precisely,* $\text{Rev}^+(M) + \mu\text{Rev}(M) \geq \frac{e-1}{e}\text{Rev}(\text{Opt})$.

Recall that the approximation bound for μ-revenue covered auctions (Theorem 10) relied on showing that the surplus from any agent in any alternate allocation was approximated by that player's contribution to BNE surplus and a fraction of the additional threshold.

We begin by showing virtual-value covering, an analogue of value covering for virtual welfare, holds in BNE directly via a reduction to value-covering (Lemma 8).

Lemma 15 (Virtual-Value Covering)**.** *Consider a mechanism* M *in BNE and an agent* i *with value* v_i*. For any* $x' \in [0, 1]$*,*

$$\phi_i^+(v_i)x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e}\phi_i(v_i)x'. \tag{13}$$

The proofs of Lemma 15 and Theorem 14 follow precisely as in Lemma 7 via a reduction to value-covering (Lemma 8), so the details are omitted.

Now that the positive virtual welfare of a mechanism approximates the optimal, the only thing left is to bound the loss due to serving bidders wiht negative virtual values. The subsequent sections discuss methods for mitigating the virtual welfare lost from to serving negative virtual valued agents.

5.2 Reserve Prices

The standard approach to prevent service to agents with negative virtual values is to set reserves such that no negative virtual valued agent is served. As long as the virtual value $\phi_i(v_i)$ is non-decreasing in v_i — equivalently, the distribution is regular — setting monopoly reserves $\mathbf{r}^*$ s.t. $r_i^* = \phi_i^{-1}(0)$ in a first-price auction for every agent will eliminate all negative virtual valued agents. If a (general) mechanism can implement such reserves and serve no agent with $\phi_i(v_i) < 0$, then it too will approximate the revenue of the optimal mechanism:

Lemma 16. *In any BNE of a* μ*-revenue covered mechanism* $M_{\mathbf{r}^*}$ *with monopoly reserves* $\mathbf{r}^*$ *in a regular environment, the revenue of* $M_{\mathbf{r}^*}$ *is a* $(\mu+1)e/(e-1)$ *approximation to the revenue of the optimal mechanism.*

The proof is straightforward — as $M_{\mathbf{r}^*}$ serves no agent with $\phi_i(v_i) < 0$, $\text{Rev}^-(M_{\mathbf{r}^*}) = 0$. By Theorem 14, $M_{\mathbf{r}^*}$ is then a $(\mu+1)e/(e-1)$ approximation to the revenue optimal mechanism. Thus if it is possible to add monopoly reserves to a mechanism, doing so gives approximately optimal revenue.

In a first price auction it is always feasible to implement reserves by restricting the bid-space. In an all-pay auction however, we cannot reliably implement value space monopoly reserves. The willingness of a player to outbid an all-pay reserve depends on the allocation probability as well as the reserve, and as such there is no easy correspondence between all-pay and value space reserves.

5.3 Duplicate bidders

Another approach to mitigating the impact of negative virtual-valued agents is to ensure each agent faces adequate competition. Bulow and Klemperer [2] show that this intuition guarantees approximately optimal revenue in regular, symmetric, single-item settings.

We show the same intuition holds for μ-revenue covered mechanisms: if each bidder must compete with at least $k-1$ other bidders with values drawn from her same distribution and bidders play by identical strategies, revenue is approximately optimal compared to the revenue optimal mechanism (including the duplicate bidders). We say such an auction satisfies *k-duplicates*, and show in Appendix D that both the first-price and all-pay auctions with at least k bidders from each distribution satisfy it.

Lemma 17. *In any BNE of a mechanism M with k-duplicates behaving by identical strategies and values drawn from regular distributions, the virtual surplus lost due to serving agents with negative virtual values is at most $1/k$ the virtual surplus from positive virtual valued agents.*

The proof is included in Appendix D. Two such auctions are the first-price and all-pay auctions.

6 Revenue Covering

In this section we prove that several commonly-used and well-studied mechanisms are revenue covered, implying new revenue results for each. All proofs are included in Appendix E.

6.1 First Price Matroid Auctions

In our discussion of the single-item case (Section 3), we saw that when an agent has trouble getting allocated in a first price auction (that is, T_i is high), it is because other agents submit high bids. These competing bids translate into revenue, implying that the first-price auction is 1-revenue covered. With one extra step, this reasoning extends to first-price auctions where the feasible allocations form a matroid. An agent's threshold bid doesn't precisely correspond to a competing bid, but matroid properties provide a sufficiently close analog, implying revenue covering revenue covering.

Lemma 18. *The first-price auction is 1-revenue covered in any matroid feasibility environment.*

Theorem 10 and Lemma 16 respectively imply welfare and revenue approximations of $2e/(e-1)$ with reserves.

6.2 Position Auctions

In first-price position auctions (ie generalized first-price auction, or GFP), arguments similar to those in the matroid case yield analogous welfare and revenue guarantees.

Formally, a position auction is an auction in which agents can be allocated one of m positions; each of which is valued by an agent at $\alpha_j v_i$. In advertising auctions, these are slots on a webpage to fill, each of which sees worse and worse click-through rates. Order the positions such that $\{\alpha_j\}$ is decreasing in j (hence slot 1 is best).

In the GFP, agents submit bids b_i, and positions are allocated in order of bid. Each agent pays their bid scaled by the quality of the slot: $\alpha_j b_i$. Equivalently, they pay their bid when they are served, which occurs with probability α_j for position j.

Theorem 19. *The generalized first price (GFP) auction is 1-revenue covered.*

As in the matroid case, Theorem 10 and Lemma 16 respectively imply welfare and revenue approximations of $2e/(e-1)$ with reserves.

6.3 All-Pay Auctions

By translating the all-pay auction into first-price semantics, the covering framework can be applied to yield welfare and revenue results. The welfare result is weaker than that in Syrgkanis and Tardos [16], but illustrates the applicability of the framework beyond first-price auctions. The revenue results we derive are new.

Lemma 20. *The all-pay matroid auction is 2-revenue covered.*

Theorem 10 implies a welfare bound of $3e/(e-1)$. It is not feasible to add reserves to an all-pay auction, but if every bidder must compete against at least one duplicate bidder (as discussed briefly in Section 5), the all-pay auction is a $4e/(e-1)$-approximation to the optimal auction.

6.4 The Second-Price Auction

Finally, we illustrate a mechanism which is not revenue-covered. The second-price auction solicits sealed bids and charges the highest bidder the second-highest bid. Consider the two-agent setting with deterministic value distribution where $v_1 = 1$ and $v_2 = \epsilon$. Assume agent 1 bids 1 and agent 2 bids ϵ. The revenue is ϵ, but $T_2[0,1]$ is 1, so the second-price auction can't be revenue covered. Intuitively, the problem comes from a disconnect between payments and allocations: agent 2 is losing not because agent 1 is making large payments to the mechanism to win, as in revenue covered mechanisms. In fact, almost the opposite holds - agent 2's low bids are both the reason she doesn't get allocated and the reason 1 pays so little.

7 Extension: Simultaneous Composition

In this section we prove that if a mechanism satisfies a stronger version of revenue covering when operated in isolation, then it is similarly covered when many instances of the mechanism are simultaneously being run. Specifically, we assume agents are unit-demand and single-valued. In this setting, we define simultaneous composition as:

Definition 21. *Let mechanisms $M_1, \ldots, M_m$ have allocation and payment rules $(\mathbf{x}^j, \mathbf{p}^j)$ and individual action spaces spaces $A_i^1, \ldots, A_i^m$ for each agent i. Let S_i be a non-empty subset of $\{1, \ldots, m\}$ for each i. The* simultaneous composition *of $M_1, \ldots, M_m$ is defined to have:*

- *Action space $\prod_j A_i^j$ for each agent. That is, each agent participates in the global mechanism by participating in each composed mechanism individually.*

- *Allocation rule* $\tilde{x}_i(\mathbf{a}) = \max_{j \in S_i} \tilde{x}_i^j(\mathbf{a}^j)$*. In other words, each agent choose their best allocation from among the composed mechanisms that interest them.*

- *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.*

Given a strategy $\mathbf{s}$ in the composed mechanism, let $\mathbf{s}^j$ denote the strategy profile in mechanism j defined by the element of each agent's strategy profile corresponding to M_j. Given $\mathbf{s}^j$, define τ_i^j, and T_i^j to be the analogous values of τ_i, and T_i in M_j under $\mathbf{s}^j$. In the composed mechanism, let A_i^j be the set of actions comprised of an arbitrary action in mechanism j and withdrawing from all other mechanisms. Further let $A'_i = \bigcup_j A_i^j$, and $\mathbf{A}' = \prod_i A'_i$.

With this notation defined, we present the main theorem - revenue covered mechanisms are closed under simultaneous composition.

Lemma 22. *Let M be the simultaneous composition mechanisms M_j for $j = 1, \ldots, m$, and let the individual mechanisms be μ-revenue covered. Then M is μ-revenue covered with respect to $\mathbf{A}'$.*

The proof is included in Appendix F. This immediately yields a revenue approximation result and when combined with restrictions to handle reserves yields a revenue approximation result.

8 Conclusion

We have shown a framework for proving price of anarchy 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 in broad environments.

We split this framework in two very distinct parts that isolate the contribution of the mechanism and the contribution of Bayes-Nash Equilibrium. The first part, value covering, depends only on Bayes-Nash Equilibrium and relates an agents surplus and expected threshold price for allocation with her optimal surplus.

The second, revenue-covering, depends only on properties of a mechanism over individually rational strategy profiles and feasible allocations. This is especially helpful when equilibria are hard to characterize or understand analytically, as is the case with the first-price auction in asymmetric settings. This has been a barrier in the past to proving results about the behavior of non-truthful auctions in asymmetric settings: we hope this framework will aid broadly in understanding properties of equilibria in auctions well beyond the confines of symmetric settings.

Extensions. We used the characterization of Bayes-Nash equilibrium in a few very 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 [13] to account for revenue via virtual values. This is the crucial part that allows us to relate the allocation a bidder receives to their contribution to revenue. Extensions beyond single-parameter, risk-neutral, private-valued agents will likely need at least an approximate virtual-value equivalent.

References

[1] Ausubel, L. M. and Milgrom, P. (2006). The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1*. MIT Press.

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

[3] Chawla, S. and Hartline, J. (2013). Auctions with Unique Equilibria. In *ACM Conference on Electronic Commerce*.

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

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

[6] Jiang, A. X. and Leyton-Brown, K. (2010). Bayesian action-graph games. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A., editors, *Advances in Neural Information Processing Systems 23*, pages 991–999. Curran Associates, Inc.

[7] 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.

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

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

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

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

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

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

[14] Roughgarden, T. (2009). Intrinsic robustness of the price of anarchy. In *STOC*.

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

[16] Syrgkanis, V. and Tardos, E. (2013). Composable and Efficient Mechanisms. In *STOC*.

[17] Talwar, K. (2003). The price of truth: Frugality in truthful mechanisms. In *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science*, pages 608–619.

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

A Examples of FPA Equilibria

A.1 Single-item FPA PoA ≥ 1.15

Consider a setting with one high bidder with a fixed value of 10, and n small bidders with values drawn from some distribution with value always less than 10. Serving the high valued agent is always the welfare-optimal allocation. We parameterize the expected utility of the high bidder as u_H.Assume the low bidders will bid such that the highest of their bids is distributed according to the CDF $F_L(b) = \frac{u_H}{10-b}$. With this distribution, player H achieves utility u_h for any bid in the range $[0, 10 - u_h]$.

The high player plays a mixed strategy according to the bid CDF $B_H(b) = \sqrt{b/(10 - u_H)}$. The competing bid CDF for each low bidder is $F_c(b) = B_H(b) \cdot B_L(b)^{(n-1/n)}$.

With $u_H = 5.7$, solving for the first order conditions in the first price auction tells us that for any low player bidding b, $v = b + F_c(b)/F_c'(b)$; solved numerically it is approximately $v(b) = \frac{15b-0.5b^2}{5+0.5b}$. Solving numerically gives welfare of 8.69; thus the price of anarchy for welfare is approximately 1.15.

This example is almost be tight against the expected cumulative threshold lowerbound $\underline{T}$ used in the proof of the value covering lemma (Lemma 8). However, the $\frac{e}{e-1}$ price of anarchy proof ignores the bid from the agent allocated in the optimal allocation and the utility of the agents allocated in FPAbut not OPT. Both of these quantities are non-zero, which leads to the 1.15 figure being reasonably far from the $\frac{e}{e-1}$. Bounding these quantities is a likely required step for improving the $\frac{e}{e-1}$ bound for single-item settings.

B Framework Proofs

Lemma 8 (Restatement). *Consider a mechanism M in BNE with induced allocation and payment rules $(\mathbf{x}, \mathbf{p})$, and an agent i with value v_i. If $x_i(v_i) \geq \underline{x}_i$, then for any $x' \in [0, 1]$,*

$$v_i x_i(v_i) + T_i[\underline{x}_i, x'] \geq \frac{e-1}{e} v_i x'. \tag{12}$$

Proof of Lemma 8. By the definition of BNE, i chooses an action which maximizes utility. It follows that

$$u_i(v_i) \geq v_i x_i(\alpha_i(z)) - p_i(\alpha_i(z)) = \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) x_i(\alpha_i(z)) \geq \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) z. \tag{14}$$

Rearranging (14) yields

$$v_i - \frac{u_i(v_i)}{z} \leq \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))} = \tau_i(z). \tag{15}$$

This bound is meaningful as long as $v_i - \frac{u_i(v_i)}{z} \geq 0$, or alternatively $z \geq u_i(v_i)/v_i$. It follows that

$$
\begin{aligned}
v_i x_i(v_i) + T_i[\underline{x}_i, x'] &\geq v_i x_i(v_i) + \int_{\underline{x}_i}^{x'} \max\left(0, v_i - \frac{u_i(v_i)}{z}\right) dz \\
&\geq v_i x_i(v_i) + \int_{u_i(v_i)/v_i}^{x'} v_i - \frac{u_i(v_i)}{z}\, dz \\
&= v_i x' + u_i(v_i) \ln \frac{u_i(v_i)}{x' v_i}.
\end{aligned} \tag{16}
$$

Holding $x'v_i$ fixed and minimizing this quantity as a function of $u_i(v_i)$ yields a minimum at $u_i(v_i) = \frac{x'v_i}{e}$, and at that point assumes value $(1 - 1/e)x'v_i$. This is precisely the righthand side of (12), implying the lemma. □

C Reserves Extension

We require an additional, further restricted version of revenue covering to obtain results in complicated settings such as simultaneous compositions of mechanisms. Using this restriction, we derive a definition of reserves which holds for first price matroid auctions, position auctions, and composed first-price auctions, as well as possibly other mechanisms.

Action Restrictions A mechanism is revenue covered if its revenue approximates a portion of its expected equivalent threshold bids. Each threshold bid solves an optimization problem: the agent seeks allocation probability at least z, and selects the lowest equivalent bid among actions which attain this allocation probability. Further restricting the feasible actions in this optimization problem can only produce a higher threshold. It follows that if a mechanism's revenue approximates expected equivalent threshold bids derived from a restricted action set, then it also approximates those from the unrestricted action set. This suggests a strengthening of revenue covering. We first refine the framework to incorporate restricted action sets. The following are analogous to the original definitions, but with a restricted action set:

Definition 23. *The equivalent threshold bid with respect to a restricted action set $A'_i \subseteq A_i$, denoted $\tau_i^{A'_i}(z)$, is defined as $\min_{a_i \in A'_i} \beta_i(a_i)$ subject to $\tilde{x}_i(a_i) \geq z$.*

Definition 24. *Given an action a_i, the expected equivalent threshold bid with respect to a restricted action set $A'_i \subseteq A_i$, denoted $T_i^{A'_i}[\tilde{x}_i(a_i), x'_i]$, is defined as $\int_{\tilde{x}_i(a_i)}^{x'} \tau_i^{A'_i}(z)\, dz$.*

Note that the lower limit of the integral is $\underline{x}_i$, still defined with respect to the full action set A_i.

Definition 25. *A mechanism M is μ-revenue covered with respect to action set restrictions $A'_1 \subseteq A_1, \ldots, A'_n \subseteq A_n$ (with product space $\mathbf{A}'$) if for every strategy profile $\mathbf{s}$, alternate feasible allocation $\mathbf{x}'$, and profile of general actions $\mathbf{a}$, $\mu\textsc{Rev}(M) \geq \sum_i T_i^{A'_i}[\tilde{x}_i(a_i), x'_i]$.*

A mechanism which is revenue covered in the original sense is revenue covered in the sense of Definition 25 with respect to the full action set $\mathbf{A}$. Moreover, a mechanism which is revenue covered with respect to the restricted action set $\mathbf{A}'$ is also revenue covered in the original sense. This follows from the fact that $\tau_i^{A'_i}(z)$ is the objective value to the same minimization problem as

$\tau_i(z)$, but on a smaller feasible region, so $\tau_i^{A'_i}(z) \geq \tau_i(z)$ for all i and z. Integrating, we see that $T_i^{A'_i}[\tilde{x}_i(a_i), x'_i] \geq T_i[\tilde{x}_i(a_i), x'_i]$ for all i and x', so $\text{Rev}(M) \geq \sum_i T_i^{A'_i}[\tilde{x}_i(a_i), x'_i] \geq \sum_i T_i[\tilde{x}_i(a_i), x'_i]$. These definitions can be further adapted to handle restricted action sets which depend on the strategy profile, i.e. $\mathbf{A}'(\mathbf{s}) = A'_1(\mathbf{s}), \ldots, A'_n(\mathbf{s})$.

We can combine the above modification to revenue covering with Definition 11 in the obvious way to yield the following definition:

Definition 26. *Let S be a function which takes a value profile $\mathbf{v}$ and outputs a set of agents, and action set restrictions $A'_1(\mathbf{s}) \subseteq A_1, \ldots, A'_n(\mathbf{s}) \subseteq A_n$ (with product space $\mathbf{A}'(\mathbf{s})$). A mechanism M is μ-revenue covered with respect to S and $\mathbf{A}'(\mathbf{s})$ if for all alternate allocations $\mathbf{x}'$, value profiles $\mathbf{v}$, and profiles of participatory actions $\mathbf{a}$,*

$$\mu \text{Rev}(M) \geq \sum_{i \in S(\mathbf{v})} T_i^{A'_i}[\tilde{x}_i(a_i), x'_i].$$

Revenue covering in the sense of Definition 26 implies that in the sense of Definition 11.

C.1 Covering-Preserving Reserves

Consider adding arbitrary bidspace reserves $\mathbf{r}$ to a first-price auction. Deriving a revenue result with our framework required two properties: (1) reserves prevented profitable allocation to low-valued agents and (2) we preserved the structure of the auction for bids above the reserve (and therefore preserved revenue covering). Property (1), along with regularity, implied virtual value covering, which combined with the revenue covering from (2) to produce a revenue bound. We seek to capture these two properties when defining reserves for general mechanisms.

Definition 27. *For a mechanism M, let $M^{\mathbf{r}}$ be M restricted to an action space $A^{\mathbf{r}}$. We say $M^{\mathbf{r}}$* implements M with covering-preserving reserves $\mathbf{r} = (r_1, \ldots, r_n)$ *if the following hold:*

1. *For any strategy profile $\mathbf{s}$, and any bidder i,*
   - (a) *Every action $a_i \in A_i^{\mathbf{r}}$ has $\beta_i(a_i) \geq r_i$.*
   - (b) *Either there is a* reserve action *$a_i^{r_i} \in A_i^{\mathbf{r}}$ with $\beta_i(a_i^{r_i}) = r_i$ or there is no action with positive probability of winning and equivalent bid at most r_i in M.*
2. *There exists a set $\mathbf{A}'$ such that M is μ-revenue covered restricted to the action sets $\mathbf{A}^{\mathbf{r}} \cap \mathbf{A}'$ for bidders with values $\mathbf{v} \geq \mathbf{r}$. We say the reserves respect $\mathbf{A}'$*

Note that this definition does not naturally capture the addition of reserves in such mechanisms as the second-price auction. In particular, for that auction, the addition of reserves increases expected payments for agents, so condition revenue covering isn't obviously preserved in the reserves mechanism. Second-price auctions are not revenue covered, as we show in Section 6.4, so this is less concerning. The definition does, however, capture the addition of reserves in auctions with first-price payment semantics.

Adding arbitrary reserves $\mathbf{r}$ to a μ-revenue covered mechanism in the manner of Definition 27 preserves revenue covering. Formally:

Lemma 28. *Let $M^{\mathbf{r}}$ implement a mechanism M with covering-preserving reserves $\mathbf{r}$ which respect restricted actions $\mathbf{A}'$, and let $S^{\mathbf{r}}$ map value profiles to agents for which $v_i \geq r_i$. If M is μ-revenue covered with respect to $\mathbf{A}'$, then $M^{\mathbf{r}}$ is μ-revenue covered with respect to $S^{\mathbf{r}}$.*

Lemma 28 implies that revenue covered mechanisms with reserves added are still revenue covered. We can therefore use Theorem 12 to bound the welfare of the reserves mechanism.

Theorem 29. *Let $M^{\mathbf{r}}$ implement a mechanism M with individual reserves $\mathbf{r}$. If M is μ-revenue covered, then in any BNE of $M^{\mathbf{r}}$, the welfare of $M^{\mathbf{r}}$ is a $(1+\mu)\frac{e}{e-1}$-approximation to the welfare of any other mechanism which only serves agents with $v_i \geq r_i$.*

In particular, we have that the first-price auction with reserves $\mathbf{r}$ in bidspace approximates the welfare of the optimal mechanism which only serves agents with $v_i \geq r_i$, which is VCG with reserves $\mathbf{r}$. Furthermore, we show in Section 7 that the simultaneous composition of several first-price auctions is revenue covered.

D Revenue Extension Proofs

D.1 Duplicates Environment

We first formally define the k-duplicates environment.

Definition 30 (k-duplicates)**.** *An auction has* k-duplicates *if there is a partition of bidders into groups $\{B_1, \ldots, B_p\}$ each of size $n_j \geq k$ such that:*

- *the value of every bidder in each group is drawn from the same distribution F_j,*
- *at most one bidder from any group can feasibly be served, and*
- *each bidder in a group is treated identically by the auction.*

This is the same as the duplicates environment of Hartline and Roughgarden [5]; however, our approximation results will hold with respect to the optimal mechanism in the duplicates environment, which is always better than the optimal auction in the same environment without duplicates. If duplicates play by identical strategies in BNE, this guarantees that a bidder can only be served if she has the highest value among her duplicates.

Additionally, we assume pointwise monotonicity in the allocation rule. Note that this is lightly stronger than the monotonicity given by the BNE characterization of Myerson [13], but will be satisfied by any bid-based mechanism like the first-price or all-pay auctions.

We assume for this proof that the allocation rule is pointwise monotonic for any bidder. This is stronger than the monotonicity assumption of the BNE characterization, but is satisfied in any bid-based mechanism (first-price, all-pay, etc.) with monotonicity of bids. A slightly more intricate analysis of the revenue curves of the duplicates as a group extends the argument here to the weaker monotonicity assumption.

We first 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 [2].

Let $\text{SPA}_R(B)$ be a second price auction run among agents in group B with a random reserve drawn according to R.

Lemma 31. *There exist reserve value distributions $R_1, R_2 \ldots R_p$ such that in any mechanism M with k-duplicates,*

$$\textsc{Rev}(M) = \sum\nolimits_j \textsc{Rev}(\text{SPA}_{R_j}(B_j)), \tag{17}$$

$$\textsc{Rev}^+(M) = \sum\nolimits_j \textsc{Rev}^+(\text{SPA}_{R_j}(B_j)). \tag{18}$$

Proof. By the pointwise monotonicity of the allocation rule, fixing the values and actions of bidders outside a group j results in a threshold value for the top ranking member of a group. Let the distribution of such thresholds be R_j; then a second price auction among the group members with reserve drawn precisely from R_j will induce exactly the same allocation rule for all members of the group. By revenue equivalence (Part 3 of Theorem 1), the revenue from the group in mechanism M will be the same as $\textsc{Rev}(\text{SPA}_{R_j}(B_j))$. The same argument holds for $\textsc{Rev}^+(M)$ and $\textsc{Rev}^+(\text{SPA}_{R_j}(B_j))$. $\square$

Proof of Lemma 17. A second-price auction within a group is now a symmetric setting, and thus we can now use the work of Bulow and Klemperer [2] to relate (17) and (18). By Bulow and Klemperer [2], if $k \geq 2$, $\textsc{Rev}(\text{SPA}_{R_j}(B_j)) \geq \frac{k-1}{k}\textsc{Rev}^+(\text{SPA}_{R_j}(B_j))$ and hence:

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

Thus $\textsc{Rev}^- \leq \frac{1}{k}\textsc{Rev}^+(M)$, exactly our desired result. $\square$

Combining Lemma 17 with Theorem 14 thus ensures approximately optimal revenue:

Corollary 32. *In any BNE of a μ-revenue covered auction in a regular environment with k-duplicates behaving by identical strategies, the revenue is a $\left(\frac{k}{k-1} + \mu\right)\frac{e}{e-1}$-approximation to the revenue of the optimal mechanism.*

In both first price and all-pay auctions, duplicates will play by identical strategies and thus each will give approximately optimal revenue. Chawla and Hartline [3] show that in a single-item setting, all bidders in a class that includes first-price and all-pay auctions (rank-and-bid based allocation rules, and bid-based payments) will behave symmetrically in BNE. If a mechanism has k-duplicates with such a payment rule, then for any group of duplicates, competing for allocation appears like a single-item auction, since at most one bidder of the group can be served. Thus, Theorem 3.1 of Chawla and Hartline [3] will imply that agents in the same group behave by identical strategies:

Corollary 33 (of Theorem 3.1, Chawla and Hartline [3])**.** *In any BNE of an auction with k-duplicates, rank-and-bid based allocation and bid-based payment, for any group B_j of agents, all agents in the group play by identical strategies everywhere except on a measure zero set of values.*

Thus Corollary 32 will hold for first-price and all-pay auctions, after proving they are revenue covered.

E Revenue Covering Proofs

E.1 First Price Matroid Auctions

To prove revenue covering for matroids, we first make note of the following lemma, due to Talwar [17], which holds for any auction which selects a basis maximizing the sum of bids, regardless of the payment semantics.

Lemma 34. *Let $M = (\mathbf{x}, \mathbf{p})$ be a mechanism which allocates to the basis which maximizes the sum of the bids of the allocated agents. For each agent i, let $\bar{s}_i(\mathbf{v}_{-i})$ be the threshold bid for i in the realized value profile $\mathbf{v}$ in the (implicit) strategy profile $\mathbf{s}$, and let $s_i(v_i)$ be i's bid in that same strategy profile, and let $\mathbf{x}'$ be any feasible allocation. Then*

$$\sum_i s_i(v_i)x_i(\mathbf{v}) \geq \sum_i \bar{s}_i(\mathbf{v}_{-i})x_i'$$

To prove the lemma, we require the following property of matroids:

Lemma 35 (Replacement Property)**.** *Let S_1 and S_2 be independent sets of size k in a matroid $\mathcal{M}$. Then there is a bijective function $f : S_2 \setminus S_1 \to S_1 \setminus S_2$ such that, for every $i \in S_2 \setminus S_1$, the set $(S_1 \setminus \{f(i)\}) \cup \{i\}$ is independent in $\mathcal{M}$.*

Proof of Lemma 34. Because subsets of feasible allocations are feasible, threshold bids are nonnegative, so we only need consider allocations $\mathbf{x}'$ which are bases. Let S and S' be sets served by $\mathbf{x}$ and $\mathbf{x}'$, respectively. Since bids are nonnegative, it follows that S and S' are the same size. By Lemma 35, there exists a bijection f from $S' \setminus S$ to $S \setminus S'$ with the replacement property in the lemma.

For each $i \in S' \setminus S$, $b_{f(i)}(\mathbf{v}_{f(i)}) \geq \bar{s}_i(\mathbf{v}_{-i})$, as if i bids above $b_{f(i)}$, then $(S \setminus \{f(i)\}) \cup \{i\}$ would be optimal and therefore i would be allocated in BNE. For each $i \in S' \cap S$, i was served in $x(\mathbf{v})$, it must be that $b_i(v_i) \geq \bar{s}_i(\mathbf{v}_{-i})$. The result follows by summing over i. □

With a relationship between threshold bids and revenue established, it remains to connect the threshold bids to T_i. We already saw in the proof of Lemma 5 that this is simple. With first-price semantics, $\tau_i(z)$ is simply the z-quantile of threshold bids. It follows that $T_i[0, 1]$ is i's expected threshold bid. Using this relationship, we get:

Lemma 18 (Restatement). *The first-price auction is 1-revenue covered in any matroid feasibility environment.*

Proof. Consider some alternate allocation $\mathbf{x}'$. By the mechanism's payment scheme and Lemma 34,

$$\textsc{Rev}(M) = \mathbf{E}_{\mathbf{v}}\left[\sum_i s_i(v_i)x_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i \bar{s}_i(\mathbf{v}_{-i})x_i'\right] = \sum_i \mathbf{E}_{\mathbf{v}}\left[\bar{s}_i(\mathbf{v}_{-i})\right]x_i'.$$

For the first-price matroid auction with no reserves, $\underline{x}_i = 0$. It follows that $\mathbf{E}_{\mathbf{v}}[\bar{s}_i(\mathbf{v}_{-i})] = T_i[\underline{x}_i, 1]$. Using this fact, we get $\textsc{Rev}(M) \geq \sum_i T_i[\underline{x}_i, 1]x_i'$. Finally, the convexity of T_i yields that $\sum_i T_i[\underline{x}_i, 1]x_i' \geq \sum_i T_i[\underline{x}_i, x_i']$, which proves the lemma. □

Combining Lemma 18 with Theorem 10 and yields

Corollary 36. *For the first price matroid auction with arbitrary reserves, the welfare of any BNE is a $2e/(e-1)$-approximation to that of any other mechanism with those same reserves.*

Moreover, using Lemma 16 and Lemma 17, we get

Corollary 37. *For the first price matroid auction with monopoly reserves and regular bidders, the revenue of any BNE is a $2e/(e-1)$-approximation to that of any other mechanism.*

Corollary 38. *For the first price matroid auction with regular bidders and at least 2 duplicates, the revenue of any BNE is a $3e/(e-1)$-approximation to that of any other mechanism.*

E.2 All-pay

As a warm-up, consider the single-item case. For any agent i, a conservative lower bound on the mechanism's revenue is i's highest competing bid. In the first-price auction, we simply noted that the expected competing bid was exactly $T_i[0,1]$. In the all-pay auction, however, this translation isn't so simple. In particular, we need to relate the expected competing bid $\mathbf{E}_{\mathbf{v}}[\bar{s}_i(\mathbf{v}_{-i})]$ to the expected *equivalent* threshold bid $T_i[0,1]$ - that is, we need to go from payments to payments divided by allocation. As we show momentarily, we only lose a factor of two in making this switch. Consequently:

$$\text{Rev}(M) \geq \mathbf{E}_{\mathbf{v}}\left[\bar{s}_i(\mathbf{v}_{-i})\right] \geq \frac{1}{2} T_i[0,1].$$

Noting that any alternate allocation simply selects an agent i, and that $\underline{x}_i = 0$ for all-pay auctions without reserves yields 2-revenue covering. If we formalize and generalize this argument, we get:

Lemma 20 (Restatement). *The all-pay matroid auction is 2-revenue covered.*

Proof. The proof follows the structure of the informal argument above - we first relate revenue to threshold bids, using Lemma 34. We then translate threshold bids to equivalent threshold bids, losing a factor of two in the process. As above, these two steps prove the result.

Revenue to Threshold Bids By the payment semantics of the mechanism,

$$\text{Rev}(M) = \mathbf{E}_{\mathbf{v}}\left[\sum\nolimits_i s_i(v_i)\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum\nolimits_i s_i(v_i) x_i(\mathbf{v})\right].$$

Now let $\bar{s}_i(\mathbf{v}_{-i})$ be the threshold bid for i in realized value profile $\mathbf{v}_{-i}$ under strategy profile $\mathbf{s}$ (without index i). Because the served agents are the basis which maximizes the sum of bids, Lemma 34 implies that

$$\mathbf{E}_{\mathbf{v}}\left[\sum\nolimits_i s_i(v_i) x_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum\nolimits_i \bar{s}_i(\mathbf{v}_{-i}) x'_i\right] = \sum\nolimits_i \mathbf{E}_{\mathbf{v}}\left[\bar{s}_i(\mathbf{v}_{-i})\right] x'_i. \tag{19}$$

Threshold Bids to Equivalent Threshold Bids We have bounded M's revenue in terms of $E_{\mathbf{v}_{-i}}[\bar{s}_i(\mathbf{v}_{-i})]$, i's expected threshold bid. To prove revenue covering, we need to bound $E_{\mathbf{v}_{-i}}[\bar{s}_i(\mathbf{v}_{-i})]$ in terms of i's expected *equivalent* threshold bid - the lowest equivalent first-price bid required to get allocated with probability z. We do so by comparing $\tau_i(z)$ to the the z-quantile i's threshold bids.

To this end, let $a_i(z)$ be the z-quantile of i's competing bids. That is, $a_i(z) = \arg\min_{a_i}$ $\tilde{p}_i(a_i)$ subject to $\tilde{x}_i(a_i) \geq z$. By the definition of τ_i,

$$\frac{\tilde{p}_i(a_i(z))}{\tilde{x}_i(a_i(z))} \geq \tau_i(z).$$

Rearranging and noting that in an all-pay auction, $\tilde{p}_i(a_i(z)) = a_i(z)$, we obtain

$$a_i(z) \geq \tau_i(z)\tilde{x}_i(a_i(z)) \geq \tau_i(z) z. \tag{20}$$

This yields the following sequence of inequalities:

$$\mathbf{E}_{\mathbf{v}_{-i}}\left[\bar{s}_i(\mathbf{v}_{-i})\right] = \int_0^1 a_i(z)\, dz \geq \int_0^1 \tau_i(z) z\, dz \geq \frac{1}{2}\int_0^1 \tau_i(z)\, dz = T_i[0,1], \tag{21}$$

where the first equality follows from noting that expected value can bee computed by integrating over quantiles, the first inequality from (20), and the second inequality from the fact that τ_i is an increasing function and basic calculus. Substituting (21) into the lower bound (19) on revenue yields

$$2\text{REV}(M) \geq \sum_i T_i[0,1]x'_i. \tag{22}$$

By the convexity of T_i, $T_i[0,1]x'_i \geq T_i[0, x'_i]$. In the all-pay matroid auction without reserves, $\underline{x}_i = 0$. This implies that the all-pay auction is 2-revenue covered.

□

E.3 GFP

In deterministic mechanisms, we used the pointwise equivalent bid threshold for allocation $\tau_i(\mathbf{v}_{-i})$, or the required bid to be allocated when other agents have values $\mathbf{v}_{-i}$. In a randomized mechanism like a position auction, fixing the actions of other results not in a single threshold but a number of thresholds — the actions of others induce an allocation rule that in the case of position auctions, is piecewise constant (a "stair" function).

We will make use of the threshold that is induced by this action profile in proving that GFP is revenue-covered. Let $\tau_i^{\mathbf{a}_{-i}}(z) = \beta_i(\alpha_i(z, \mathbf{a}_{-i}), \mathbf{a}_{-i})$ be the smallest equivalent bid of an action for bidder i which achieves at least allocation of z when other bidders play actions $\mathbf{a}_{-i}$. Let $T_i^{\mathbf{a}}[x_i(\mathbf{a}), x'] = \int_{x_i(\mathbf{a})}^{x'} \tau_i^{\mathbf{a}}(z)\, dz$ denote the expected additional threshold for agent i when other bidders play $\mathbf{a}_{-i}$.

To prove GFP is revenue covered for all strategy profiles, we will show first that GFP satisfies a pointwise variant of revenue covering; then, that pointwise revenue covering implies revenue covering.

Definition 39. *A mechanism M is pointwise μ-revenue covered if for any participatory actions $\mathbf{a}$ and alternate feasible allocation x',*

$$\mu REV(M(\mathbf{a})) \geq \sum_i T_i^{\mathbf{a}}[x_i(\mathbf{a}), x'_i]. \tag{23}$$

Lemma 40. *GFP is pointwise 1-revenue covered.*

Proof. Consider the bid-based allocation rule of an agent in GFP, $\tilde{x}_i(b_i, b_{-i})$. For any bid b_i, $\tilde{x}_i(b_i, b_{-i})$ is the position weight of the best slot such that the current resident of the slot is bidding less than b_i. So, $\tilde{x}_i(b_i)$ will be a stair function, with a stair corresponding to each position. The area above the curve between allocation probabilities $\tilde{x}_i(b_i, b_{-i})$ and $\mathbf{x}'_i$, $T_i^{\mathbf{a}}[\tilde{x}_i(b_i, b_{-i}), x']$, is a lower bound on the actual payment made by the bidder in the slot, because it assumes that the current winner is paying his bid only for the extra marginal clicks, not the clicks across all the slots. Denote by b^j the winning bid for each position j; then

$$T_i^{\mathbf{a}}[\tilde{x}_i(b_i, b_{-i}), \alpha_j] = \sum_{i=j}^{m} (\alpha_i - \alpha_{i+1}) b^j. \tag{24}$$

The revenue in GFP given a set of bids is $\sum_j \alpha_j b^j$. For any slot j, the threshold amount for the bidder allocated j in the alternate allocation is less than payment of the bidder who won the slot j: $\alpha_j b^j \geq \sum_{i=j}^{m} (\alpha_i - \alpha_{i+1}) b^j$. Summing over all bidders gives that $REV(\mathbf{a}) \geq \sum_i T_i^{\mathbf{a}}[\tilde{x}_i(b_i, b_{-i}), x'_i]$, our desired result. □

We now prove (general) revenue covering for the generalize first price auction:

Theorem 19 (Restatement). *The generalized first price (GFP) auction is* 1*-revenue covered.*

Proof. Let $\mathbf{x}'$ be a feasible allocation. We begin by applying pointwise revenue covering for every set of realized actions; thus

$$REV(GFP) \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i T_i^{\mathbf{s}(\mathbf{v})}[x_i(\mathbf{a}), x_i']\right]. \tag{25}$$

We now have expectation with a fixed allocation probability, but this means a different action for every specific set of actions played by the other bidders. To translate this to a fixed action and a varying allocation probability across the actions played by other bidders, we will use properties of $T[\underline{x}_i,]$. By the convexity of $T^{\mathbf{s}(\mathbf{v})}[\underline{x}_i, x_i']$ in x_i' and because $\tau_i(x_i') = \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(\tau_i(x_i'))]$,

$$\mathbf{E}_{\mathbf{v}}\left[\sum_i T_i^{\mathbf{s}(\mathbf{v})}[x_i(\mathbf{s}(\mathbf{v})), x_i']\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i T_i^{\mathbf{s}(\mathbf{v})}[x_i(\mathbf{s}(\mathbf{v})), \tilde{x}_i(\tau_i(x_i'))]\right] = \sum_i T_i[x_i(\mathbf{s}(\mathbf{v})), x_i']. \tag{26}$$

Chaining with (25) gives our desired result, $REV(GFP) \geq \sum_i T_i[x_i(v_i), x_i']$. ☐

F Simultaneous Composition Proofs

Proof of Lemma 22. Let $\mathbf{x}'$ be a feasible allocation for the global mechanism. By the way we defined composition, $\mathbf{x}'$ is feasible only if we can construct a matching between agents and mechanisms such that for any j, there is a feasible allocation for M_j that allocates each i matched to j according to x_i'. Define $x_{i,j}'$ be x_i' if i and j are matched, and 0 otherwise. Note that for each agent i, $x_{i,j}' > 0$ for at most one j, with $x_{i,j}' = 0$ for all j if $x_i' = 0$. By downward closure, $\mathbf{x}'_{\cdot,j}$ is feasible for M_j.

Now by the definition of the composed mechanism, $\mu\textsc{Rev}(M) = \mu \sum_j \textsc{Rev}(M_j)$, where $\textsc{Rev}(M_j)$ is taken with respect to $\mathbf{s}^j$ for each j. Because each M_j is μ-revenue covered, it follows that $\mu \sum_j \textsc{Rev}(M_j) \geq \sum_j \sum_i T_i^j[\underline{x}_i^j, x_{i,j}']$. Moreover, for all j, $T_i^j[\underline{x}_i^j, x_{i,j}'] = T_i^{A_i^j}[\underline{x}_i, x_{i,j}'] \geq T_i^{A_i'}[\underline{x}_i, x_{i,j}']$. But for each agent i, $x_{i,j}' > 0$ for at most one j, so $T_i^j[\underline{x}_i^j, x_{i,j}'] > 0$ for at most one j as well, again with $T_i^j[\underline{x}_i^j, x_{i,j}'] = 0$ for all j if $x_i' = 0$. The same also holds for $T_i^{A_i'}[\underline{x}_i, x_{i,j}']$. It follows that $\sum_j \sum_i T_i^j[\underline{x}_i^j, x_{i,j}'] \geq \sum_i T_i^{A_i'}[\underline{x}_i, x_i']$, which implies the result. ☐