6
Bayes-Nash Approximation

This text primarily focuses on the design of incentive compatible mechanisms, i.e., ones where truth telling is an equilibrium. This focus is justified in theory by the revelation principle (Section 2.10 on page 47) which suggests that if there is a mechanism with a good equilibrium then there is one where truth telling is a good equilibrium. Thus, nothing "good" is lost by the restriction. In practice, though, mechanisms are rarely incentive compatible, and undoing the revelation principle is not straightforward. It is not always an easy task to identify a practical mechanism with the same Bayes-Nash equilibrium outcome as a designed Bayesian incentive compatible mechanisms. This chapter focuses on the analysis of mechanisms that are not incentive compatible, and in design criteria for them.

In the design of Bayes-Nash (i.e., non-incentive-compatible) mechanisms there will be less fine grained control over the exact equilibrium selected by the mechanism, instead we will look to identify properties of mechanisms from which we can guarantee that any equilibrium is approximately optimal.

Our motivating example is the first-price auction with agents with independent but non-identically distributed values. Recall that with identically distributed values the first-price auction possesses a unique equilibrium in which the highest valued agent always wins the item (see Section 2.9 on page 42). This outcome is optimal from the perspective of social surplus. Moreover, the first-price auction with the monopoly reserve price, for values drawn i.i.d. from a regular distribution, is revenue optimal in equilibrium. For asymmetric distributions, however, the first-price auction is neither optimal for surplus nor revenue. We will show that the first-price auction is an $e/e-1 \approx 1.58$ approximation for social surplus, and the first price auction with asymmetric monopoly reserves is a $^{2e}/_{e-1} \approx 3.16$ approximation for revenue. Neither of these bounds are tight.

One of the reasons analysis of Bayes-Nash mechanisms is important is that the ideal setting of incentive compatible mechanism design, where a mechanism is being run in a closed system, is rare. In many practical applications of mechanism design, agents may have the option to participate in many mechanisms, simultaneously or in sequence. Incentive compatibility of these individual mechanisms does not imply incentive compatibility of the composition of mechanisms. An important development of this chapter is a theory of composition for mechanisms. Via this theory we will show that simultaneous first-price auctions for multiple items (albeit for single-dimensional agents) have the same performance guarantees stated above for the first-price auction in isolation.

The conventional approach to the analysis of Bayes-Nash equilibrium, as a first step, explicitly solves for the Bayes-Nash equilibrium. For asymmetric environments such an analysis would require the solution to analytically intractable differential equations. The approximation-based approach presented herein circumvents solving for BNE by decomposing the analysis into the following two parts. The first part isolates the best-response property of Bayes-Nash equilibrium and formalizes the intuition that either an agent gets good utility or must be facing fierce competition. The second part identifies a *revenue covering* property, that revenue exceeds an aggregate measure of the competition faced by each agent, as a criteria to be approximated. With bounds on utility and revenue, we get approximation bounds on the social surplus (i.e., the sum of utility and revenue).

The bounds we derive on the social surplus and revenue of auctions in Bayes-Nash equilibrium are parameterized by the extent to which revenue covering approximately holds. This observation then gives clear direction for optimization in mechanism design. A Bayes-Nash mechanism's performance is proportional to its approximation with respect to revenue covering. Bayes-Nash mechanisms should be designed to minimize this approximation.

6.1 Social Surplus of Winner-pays-bid Mechanisms

The first-price auction with asymmetric value distributions does not maximize social surplus in Bayes-Nash equilibrium. For two agents and the uniform distribution on distinct supports, the differential equations that govern Bayes-Nash equilibrium can be solved; Example 6.1.1, below, gives the solution to one special case. For more general distributions, more than two agents, and more complex auction formats, equilibrium is analytically intractable. A main result of this section will be the following bound on the social surplus of any Bayes-Nash equilibrium of the first-price auction for any product distribution on agent values.

Theorem 6.1.1. *For (non-identical) independent, single-item environments, the expected social surplus of the first-price auction in Bayes-Nash equilibrium is an* $\frac{e}{e-1} \approx 1.58$ *approximation to the optimal surplus.*

Example 6.1.1. *The equilibrium* $\boldsymbol{b}$ *of Alice (agent 1) and Bob (agent 2) in the first-price auction with* $\mathsf{v}_1 \sim U[0,1]$ *and* $\mathsf{v}_2 \sim U[0,2]$ *is*

- $b_1(\mathsf{v}) = \frac{2}{3\mathsf{v}}\left(2 - \sqrt{4 - 3\mathsf{v}^2}\right)$ *and*
- $b_2(\mathsf{v}) = \frac{2}{3\mathsf{v}}\left(\sqrt{4 + 3\mathsf{v}^2} - 2\right)$.

The asymmetry of strategies implies that the highest-valued agent does not always win, i.e., the auction is inefficient. Bob wins when $\mathsf{v}_2 > (\mathsf{v}_1^{-2} + 3/4)^{-1/2} > \mathsf{v}_1$. *See Figure 6.1.*

In this section we will consider generalizations of the first-price auction where the mechanism selects an allocation based on the bids and all the winners pay their bids. We consider mechanisms for general single-dimensional environments with feasibility constraint given by $\mathcal{X} \subset \{0,1\}^n$ (see Section 3.1 on page 55). In the context of the subsequent definition, the first-price auction is the winner-pays-bid highest-bids-win mechanism for single-item environments.

Definition 6.1.1. *A* winner-pays-bid *mechanism*

(i) *solicits bids,*
(ii) *selects a feasible set of winners, and*
(iii) *charges each winner her bid.*

Chapter 6: Topics Covered.
- The geometry of best response for single-dimensional agents.
- Revenue covering, a criterion for Bayes-Nash optimization.
- Analysis of welfare and revenue in Bayes-Nash equilibrium.
- Analysis of the simultaneous composition of mechanisms.
- Reserve prices in Bayes-Nash mechanisms.

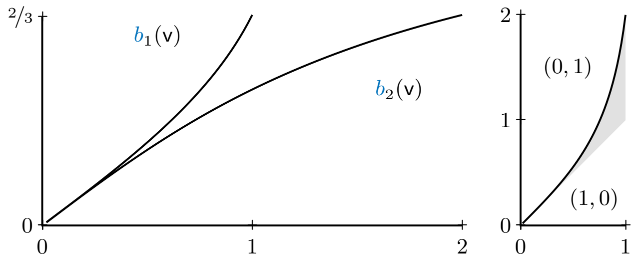

(a) Asymmetric bid strategies (b) Allocation

Figure 6.1. In (a) the Bayes-Nash bid strategies in the first-price auction with the asymmetric value distribution of Example 6.1.1 are depicted. In (b) the Bayes-Nash allocation $\boldsymbol{x}(\mathbf{v})$ is depicted with v_1 on the horizontal axis and v_2 on the vertical axis. In the shaded gray area the BNE allocates to agent 1 while $\mathsf{v}_2 > \mathsf{v}_1$ and this allocation is inefficient.

In the winner-pays-bid highest-bids-win *mechanism, the winners selected in Step* (ii) *are the feasible set of agents with the highest sum of bids.*

6.1.1 The Geometry of Best Response

Consider an agent, Alice, in a winner-pays-bid mechanism. Alice wins the auction when her bid b exceeds a *critical bid* $\hat{\mathsf{b}}$ which is given by the bids of others and the rules of the mechanism. For the first-price auction this critical bid is the maximum of the other agents' bids. Denote the interim *bid allocation rule*, which maps Alice's bid to her probability of winning, as given the distribution of other agents' bids and the mechanism's rules, by $\tilde{x}(\mathsf{b}) = \mathbf{Pr}_{\hat{\mathsf{b}}}[\mathsf{b} > \hat{\mathsf{b}}]$. Notice that this bid allocation rule is precisely to the cumulative distribution function of Alice's critical bid. The expected critical bid Alice faces is a measure of the level of competition in the auction, and is given by the area above its cumulative distribution function (equivalently, the area above the bid allocation rule).[1] By the definition of the auction, Alice's utility with value v for any bid b is given by $\tilde{u}(\mathsf{v}, \mathsf{b}) = (\mathsf{v} - \mathsf{b})\,\tilde{x}(\mathsf{b})$. Our analysis will relate Alice's utility, her expected critical bid, and her value; each of which can be compared geometrically (Figure 6.2).

- Alice's expected utility for any bid, denoted $\tilde{u}(\mathsf{v}, \mathsf{b}) = (\mathsf{v} - \mathsf{b})\,\tilde{x}(\mathsf{b})$, is

[1] Recall, the expected value of a non-negative random variable $\mathsf{v} \sim F$ is given by $\mathbf{E}[\mathsf{v}] = \int_0^\infty (1 - F(z))\,\mathrm{d}z$, cf. Section A.3 on page 426.

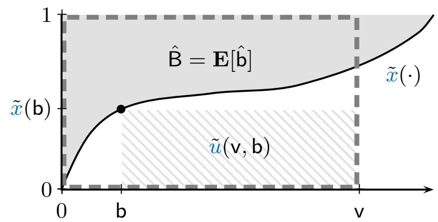

Figure 6.2. Geometry of best-response in the first price auction. The expected critical bid is the region (gray solid) above the bid allocation rule (thin solid line). The utility from a bid is given by the a rectangle (gray striped) below the bid allocation rule. The value of the agent can be depicted by the area of a rectangle (thick dashed outline).

given by a rectangle below the bid allocation rule. Alice's utility in Bayes-Nash equilibrium, denoted $u(\mathsf{v})$, is the largest such rectangle.

- Alice's expected critical bid, denoted $\hat{\mathsf{B}} = \mathbf{E}[\hat{\mathsf{b}}]$, is given by the area above the bid allocation rule, i.e., $\hat{\mathsf{B}} = \int_0^\infty (1 - \tilde{x}(\mathsf{b}))\,\mathrm{d}\mathsf{b}$.
- Alice's value v can be compared geometrically to the above quantities as the area of the rectangle of width v and height 1.

Intuitively, either Alice's utility or her expected critical bid is a large fraction of her value.

6.1.2 Utility Approximates Value

We formalize the geometric intuition that either Alice's utility or expected critical bid is large compared to her value in the following theorem.

Theorem 6.1.2. *In any Bayes-Nash equilibrium of any winner-pays-bid mechanism and for any agent, the expected sum of her utility and her critical bid is an* $e/e-1 \approx 1.58$ *approximation to her value; i.e.,*

$$u(\mathsf{v}) + \hat{\mathsf{B}} \geq {}^{e-1}/_{e}\,\mathsf{v}.$$

One way to prove Theorem 6.1.2 is via a best-response argument. In particular, BNE utility is at least the utility $\tilde{u}(\mathsf{v}, \mathsf{b})$ for any value v and any deviation bid b; a careful selection of deviation gives the desired bound. As a warm up, consider deviating to $\mathsf{b} = \mathsf{v}/2$ and observe that

$$\tilde{u}(\mathsf{v}, \mathsf{v}/2) + \hat{\mathsf{B}} \geq {}^{1}/_{2}\,\mathsf{v}. \qquad (6.1.1)$$

Fix any critical bid $\hat{\mathsf{b}}$. If $\hat{\mathsf{b}} \geq \mathsf{v}/2$, non-negativity of BNE utility implies

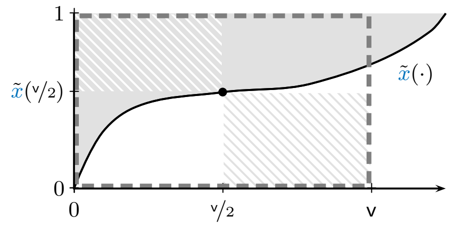

Figure 6.3. Geometry of equation (6.1.1): $\tilde{u}(v, v/2) + \hat{B} \geq 1/2\, v$. The utility $\tilde{u}(v, b)$ from deviating to $b = v/2$ is depicted by the bottom-right gray striped rectangle; the expected critical bid $\hat{B}$ is depicted by the gray solid region; the lower bound on $\hat{B}$ is the top-left white striped rectangle.

inequality (6.1.1). On the other hand, if $\hat{b} \leq v/2$, then Alice wins by bidding $b = v/2$ and her utility is $v - v/2 = v/2$ as required by (6.1.1). Taking expectations of these inequalities over $\hat{b}$ gives equation (6.1.1). This argument is depicted geometrically in Figure 6.3.

Proof of Theorem 6.1.2. Fix the critical bid $\hat{b}$, and consider the utility from deviating to a random bid b drawn from the distribution G on support $[0, {}^{e-1}/_{e}\, v]$ with density function $g(z) = {}^{1}/_{v-z}$. If $\hat{b} \geq {}^{e-1}/_{e}\, v$ the inequality of the theorem holds. Otherwise, the utility from such a deviation is $v - b$ when $b \geq \hat{b}$, zero otherwise.

$$\begin{aligned} u(v) &\geq \tilde{u}(v, b) \\ &\geq \int_{\hat{b}}^{e-1/e\, v} (v - b)\, g(b)\, db \geq \int_{\hat{b}}^{e-1/e\, v} 1\, db \\ &= {}^{e-1}/_{e}\, v - \hat{b}. \end{aligned}$$

Thus, $u(v) + \hat{b} \geq {}^{e-1}/_{e}\, v$. The deviation strategy is independent of $\hat{b}$, so taking expectation over $\hat{b}$ yields the theorem. □

6.1.3 Revenue Covering Approximation

The next part of the analysis is to bound the sum of the expected critical bids for any feasible subset of the agents, as given by $\mathscr{X}$, by the expected revenue of the auction. This analysis is performed for any distribution of bids and does not, in particular, assume that these bids are in equilibrium. For winner-pays-bid mechanisms, we may as well perform this bound pointwise for all bid profiles. Such a pointwise non-equilibrium analysis is both easy and versatile.

Definition 6.1.2. *The* revenue covering approximation μ *of winner-pays-bid mechanism* $\mathscr{M}$ *is the smallest* μ *such that, for any profile of bids* $\mathbf{b}$ *and any feasible allocation* $\mathbf{y} \in \mathscr{X}$*,*

$$\textit{Revenue}(\mathbf{b}) \geq \tfrac{1}{\mu} \sum\nolimits_i \hat{\mathsf{b}}_i \, \mathsf{y}_i, \tag{6.1.2}$$

where $\textit{Revenue}(\mathbf{b})$ *is the revenue of* $\mathscr{M}$ *on bids* $\mathbf{b}$ *and* $\hat{\mathsf{b}}_i$ *is the critical bid of agent* i *given bids* $\mathbf{b}_{-i}$*. Mechanism* $\mathscr{M}$ *is* revenue covered *if* $\mu = 1$.

Of course by taking expectations of both sides of inequality (6.1.2) for and distribution of bids $\mathbf{b} \sim \boldsymbol{G}$, the expected revenue is at least the surplus of expected critical bids of any feasible allocation $\mathbf{y} \in \mathscr{X}$, i.e.,

$$\mathbf{E}_{\mathbf{b}}[\text{Revenue}(\mathbf{b})] \geq \tfrac{1}{\mu} \sum\nolimits_i \hat{\mathsf{B}}_i \, \mathsf{y}_i,$$

where $\hat{\mathsf{B}}_i = \mathbf{E}_{\mathbf{b}}[\hat{\mathsf{b}}_i]$ for mechanism $\mathscr{M}$ as previously discussed.

Notice that in the definition of revenue covering the revenue of the auction and the critical bids are given by the bid profile and the definition of the rules of the mechanism. The allocation $\mathbf{y}$ is unrelated to the bid profile and the rules of the mechanisms, it is only constrained by the feasibility constraint $\mathscr{X} \subseteq \{0,1\}^n$ of the single-dimensional allocation problem. Analysis of the revenue covering approximation μ of any given mechanism in any given environment is generally straightforward.

Theorem 6.1.3. *The first-price auction in single-item environments is revenue covered, i.e.,* $\mu = 1$.

Proof. First, recall that for the first-price auction the critical bid faced by an agent is equal to the highest of the other agents' bids. The revenue of the auction is the highest bid over all. Thus, the revenue of the auction $\text{Revenue}(\mathbf{b}) = \max_j \mathsf{b}_j$ is at least the critical bid $\hat{\mathsf{b}}_i = \max_{j \neq i} \mathsf{b}_j$ of any agent j. Second, feasibility of $\mathbf{y}$ in single-item environments requires that $\sum_i \mathsf{y}_i \leq 1$. Combining these observations,

$$\begin{aligned}\sum\nolimits_i \hat{\mathsf{b}}_i \, \mathsf{y}_i &\leq \sum\nolimits_i \text{Revenue}(\mathbf{b}) \, \mathsf{y}_i \leq \text{Revenue}(\mathbf{b}) \sum\nolimits_i \mathsf{y}_i \\ &\leq \text{Revenue}(\mathbf{b}).\end{aligned}$$ □

This result and proof can be easily extended from single-item environments to multi-unit and matroid environments (Section 4.6 on page 135) for which the winner-pays-bid highest-bids-win mechanism continues to be revenue covered. This analysis is deferred to Section 6.6.2 on page 203.

Importantly, many mechanisms are not revenue covered. Below, Example 6.1.2 shows that the winner-pays-bid highest-bids-win mechanism for the m-item single-minded combinatorial auction environment, a canonical downward-closed environment, has revenue-covering approximation $\mu = m$. The winner-pays-bid highest-bids-win mechanism for the routing environment discussed in Section 1.1.3 on page 14 similarly has revenue covering approximation that is linear in the diameter of the graph (see Exercise 6.2).

Example 6.1.2. *In the* single-minded combinatorial auction *environment there are n agents and m items. Agent i has value v_i for bundle $\mathsf{S}_i \subseteq \{1, \ldots, m\}$ (and no value for any other bundle of items). Item j may be sold to at most one agent. It is assumed that the bundles are known and the values are each agent's private information. An allocation $\mathbf{x}$ is feasible for a single-minded combinatorial environment if no items are allocated to multiple agents, i.e., $\mathsf{x}_i = \mathsf{x}_{i^\dagger} = 1$ only if $\mathsf{S}_i \cup \mathsf{S}_{i^\dagger} = \emptyset$.*

Consider the winner-pays-bid highest-bids-win mechanism for the single-minded combinatorial auction environment. We exhibit an environment (given by bundles $\mathbf{S} = (\mathsf{S}_1, \ldots, \mathsf{S}_n)$, a bid profile $\mathbf{b}$, and feasible allocation $\mathbf{y}$ such that the inequality of Definition 6.1.2 is satisfied only by $\mu = m$.

Consider the environment with $n = m + 1$ agents with:

- *agent $m + 1$ demanding $\mathsf{S}_{m+1} = \{1, \ldots, m\}$, the grand bundle, and*
- *agent $i \neq n$ demanding $\mathsf{S}_i = \{i\}$, a singleton bundle.*

Consider the bid profile $\mathbf{b} = (0, \ldots, 0, 1)$ and feasible (but not highest-bids-win) allocation $\mathbf{y} = (1, \ldots, 1, 0)$. In other words, the agent $m + 1$ demanding the grand bundle bids $\mathsf{b}_{m+1} = 1$ and is not served ($\mathsf{y}_1 = 0$), while singleton agents $i \neq n$ bid $\mathsf{b}_i = 0$ and are served ($\mathsf{y}_i = 1$).

The highest-bids-win mechanism's revenue for bids $\mathbf{b}$ is only 1 as it selects feasible outcome $\mathbf{x} = \mathbf{b} = (0, \ldots, 0, 1)$ that serves only the grand-bundle agent $m + 1$. Each singleton agent $i \neq m + 1$ faces a critical bid of $\hat{\mathsf{b}}_i = 1$ as she must beat agent $m+1$. As the sum of the feasible critical bids is $\sum_i \hat{\mathsf{b}}_i \, \mathsf{y}_i = m$, the auction is not a revenue covering approximation for any $\mu < m$. On the other hand, tt is an easy exercise to see that the auction has revenue covering approximation at most m. Thus, $\mu = m$.

A similar example shows that worst-case equilibria can match the bound implied by this revenue covering approximation (Exercise 6.3).

In the next sections we will see that the approximation of social surplus (and revenue, with monopoly reserves) of a Bayes-Nash mechanism is proportional to its revenue covering approximation μ. Thus, to design a good Bayes-Nash mechanism it is suffices to design a mechanism with small revenue covering approximation.

6.1.4 Social Surplus in Bayes-Nash Equilibrium

Theorem 6.1.2, which shows that utility approximates value, and revenue covering approximation combine to give a bound on the Bayes-Nash equilibrium surplus relative to any feasible outcome, including the one that maximizes social surplus.

Theorem 6.1.4. *For any winner-pays-bid mechanism that has a revenue covering approximation of $\mu \geq 1$, the expected social surplus in Bayes-Nash equilibrium is an $\mu e/e-1$ approximation to the optimal social surplus.*

Proof. Consider a valuation profile $\mathbf{v}$ and agent i. By Theorem 6.1.2 in BNE,

$$u_i(\mathsf{v}_i) + \hat{\mathsf{B}}_i \geq {}^{e-1}/_{e}\, \mathsf{v}_i.$$

Denote by $y^\star(\mathbf{v}) = \operatorname{argmax}_{\mathsf{y}} \sum_i \mathsf{v}_i\, \mathsf{y}_i$ the surplus optimizing allocation. Thus, $\sum_i \mathsf{v}_i\, y_i^\star(\mathbf{v}) = \mathrm{REF}(\mathbf{v})$, the optimal social surplus. Notice that $y_i^\star(\mathbf{v}) \in [0,1]$; thus,

$$u_i(\mathsf{v}_i) + \hat{\mathsf{B}}_i\, y_i^\star(\mathbf{v}) \geq {}^{e-1}/_{e}\, \mathsf{v}_i\, y_i^\star(\mathbf{v}).$$

Sum over all agents i and invoke μ revenue covering:

$$\sum\nolimits_i u_i(\mathsf{v}_i) + \mu\, \mathbf{E}[\text{Revenue}] \geq {}^{e-1}/_{e}\ \mathrm{REF}(\mathbf{v}).$$

Take expectation over values $\mathbf{v}$ from the distribution $\boldsymbol{F}$ and use $\mu \geq 1$:

$$\mu \left(\mathbf{E}[\text{Utilities}] + \mathbf{E}[\text{Revenue}]\right) \geq {}^{e-1}/_{e}\, \mathbf{E}[\mathrm{REF}(\mathbf{v})]\,.$$

The theorem follows from observing that the surplus of the mechanism APX is equal to sum of the utilities of the agents and the mechanism's revenue. □

As is evident from the statement of Theorem 6.1.4, to show that a winner-pays-bid auction has good welfare in Bayes-Nash equilibrium it suffices to show that it has a good revenue covering approximation. As we saw above, the first-price auction has revenue covering approximation of $\mu = 1$ (Theorem 6.1.3); thus, it is a $e/e-1 \approx 1.58$ approximation to social surplus in any Bayes-Nash equilibrium. In other words, Theorem 6.1.1 is proved. More generally, winner-pays-bid highest-bids-win matroid mechanisms are also a 1.58 approximation to social surplus (by Corollary 6.6.3 and Theorem 6.1.4, see Section 6.6.2).

6.2 Beyond Winner-pays-bid Mechanisms

In this section we will extend the analysis of the preceding section to mechanisms that do not have winner-pays-bid semantics. This extension will allow straightforward generalization to all-pay mechanisms and mechanisms with complex action spaces. A motivating example of a mechanism with a complex action space is the simultaneous first-price auction for single-dimensional constrained matching markets, cf. Section 4.6.1 on page 137.

In constrained matching markets there are n agents and m items. Each agent i has a value v_i for any of the items in bundle $\mathsf{S}_i \subset \{1, \ldots, m\}$, but desires at most one of these items. In the simultaneous first-price auction, each agent selects which items to bid on and how much to bid, the agents submit bids simultaneously to the auctions, and each item is sold at the highest bid to the highest bidder. Importantly, an agent who bids in more than one auction may win more than one item even though she only has value for one item.

While the simultaneous first-price auction allows multi-dimensional bids, it is still a single-dimensional game, see Section 2.4 on page 29. Just as the Bayes-Nash equilibrium characterization for single-dimensional games is expressed by allocation and payment rules in terms of each agent's valuation (Theorem 2.5.1 on page 31), revenue equivalence suggests that we can equally well express the allocation and payment rule of the BNE of any mechanism in terms of the winner-pays-bid implementation of the BNE.

Definition 6.2.1. *Consider any mechanism $\mathscr{M}$, an agent with action space A in $\mathscr{M}$, and any distribution of other agents' actions.*

- *The (interim)* action allocation rule $\tilde{y} : \mathsf{A} \to [0, 1]$ *maps any action* $\mathsf{a} \in \mathsf{A}$ *to a probability of allocation.*
- *The (interim)* action payment rule $\tilde{p} : \mathsf{A} \to [0, 1]$ *maps any action to an expected payment.*
- *The (interim, effective)* winner-pays-bid allocation rule, *denoted* $\tilde{x}(\cdot)$, *is the smallest monotone function that upper bounds the pointset given by:*[2]

$$\left\{ \left(\tilde{p}(\mathsf{a}) / \tilde{y}(\mathsf{a}),\ \tilde{y}(\mathsf{a}) \right) \,:\, \mathsf{a} \in \mathsf{A} \right\}. \tag{6.2.1}$$

- *The (interim, effective)* expected critical bid $\hat{\mathsf{B}}$ *is the area above* $\tilde{x}(\cdot)$.

[2] Note, if the mechanism $\mathscr{M}$ is individually rational then $(0, 0)$ is always in the pointset.

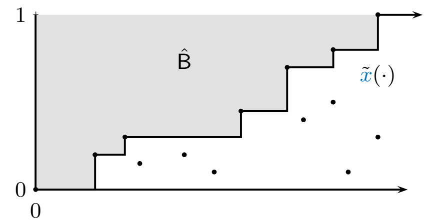

Figure 6.4. The pointset of equation (6.2.1) is depicted. The effective winner-pays-bid allocation rule $\tilde{x}(\cdot)$ is the smallest monotone function that upper bounds the point set. The points strictly below $\tilde{x}(\cdot)$ are dominated and correspond to actions that will never be taken. The expected critical bid $\hat{\mathsf{B}}$ is given by the shaded (light gray) area.

As notated in Definition 6.2.1, this section and the subsequent section will distinguish between allocation rules in (effective) winner-pays-bid bid space and in the action space defined by the mechanism. The former are denoted as $\tilde{x}$ (for consistency with interim winner-pays-bid bid-allocation rules in the rest of this chapter); the latter are denoted as $\tilde{\mathscr{y}}$ and $\tilde{y}$ for the ex post and interim action-allocation rules, respectively.

Especially for auctions like the all-pay auction, is is not well-defined to talk about the (interim, effective) winner-pays-bid allocation rule without imposing a distribution over the actions of the other agents. The following definition generalizes revenue covering approximation to mechanisms that do not have winner-pays-bid semantics.

Definition 6.2.2. *A mechanism $\mathscr{M}$ has* revenue covering approximation *μ if, for any product distribution on action profiles $\mathbf{a} \sim \boldsymbol{G}$ and any feasible allocation $\mathbf{y}$,*

$$\mathbf{E}_{\mathbf{a}}[Revenue(\mathbf{a})] \geq \tfrac{1}{\mu} \sum\nolimits_i \hat{\mathsf{B}}_i\, \mathsf{y}_i,$$

where, for action profile $\mathbf{a} \sim \boldsymbol{G}$ and mechanism $\mathscr{M}$, $Revenue(\mathbf{a}) = \sum_i \tilde{\mathscr{p}}_i(\mathbf{a})$ is the mechanism's revenue and $\hat{\mathsf{B}}_i$ is the (effective) expected critical bid of agent i (from her effective winner-pays-bid allocation rule; Definition 6.2.1).

The developments of the previous section; specifically Theorem 6.1.2 and Theorem 6.1.4; extend without modification to non-winner-pays-bid mechanisms via Definition 6.2.1 and Definition 6.2.2. The following theorem summarizes.

Theorem 6.2.1. *For any individually-rational mechanism that has a* *revenue covering approximation of $\mu \geq 1$, the expected social surplus in Bayes-Nash equilibrium is an $\mu e/e-1$ approximation to the optimal social surplus.*

For example, the following theorem can be shown. From it and Theorem 6.2.1 we conclude that the all-pay auction is a $2e/e-1 \approx 3.16$ approximation to social welfare.[3]

Theorem 6.2.2. *In single-item environments the all-pay auction is 2 revenue covered.*

Proof. See Exercise 6.4. □

6.3 Simultaneous Composition

In this section we consider the simultaneous composition of revenue covered mechanisms and show that the composite mechanism is itself revenue covered. An example to have in mind is the simultaneous first-price auction for single-dimensional constrained matching markets that was described at the onset of the preceding section. We impose three assumptions on the environment of these mechanisms:

(i) The agents are *unit-demand* with respect to simultaneous allocation across several mechanisms. In other words, an agent is considered served if she is served by any of the individual mechanisms in the composition, and she has no additional value for being served by multiple mechanisms over being served in a single mechanism. She must pay for each mechanisms in which she is served.
(ii) Each mechanism is *individually rational*. This assumption requires that each agent has an action that gives non-negative utility. In particular, an agent with value zero must have an action with zero (expected) payment; we may as well assume that such an agent will also not be served. This action effectively enables an agent to abstain from participation in each mechanism.
(iii) The individual environments in the composition are *downward closed* and the composite environment is their *union environment*. In other words, if $\mathbf{x}^1, \ldots, \mathbf{x}^m \in \{0,1\}^n$ are deterministic feasible outcomes for $\mathscr{M}^1, \ldots, \mathscr{M}^m$, respectively; then $\mathbf{x}$ with $x_i = \max_j x_i^j$ is feasible for $\mathscr{M}$.

[3] An improved analysis of the BNE surplus of the all-pay auction is available by proving a version of Theorem 6.1.2 for all-pay-format payment rules. See Exercise 6.1.

Definition 6.3.1. *Given m mechanisms $\mathscr{M}^1, \ldots, \mathscr{M}^m$; the* simultaneous composite *mechanism $\mathscr{M}$ for unit-demand agents is the following:*

- *Agent i's action space in $\mathscr{M}$ is $\mathsf{A}_i = \mathsf{A}_i^1 \times \cdots \times \mathsf{A}_i^m$ where A_i^j is agent i's action space for mechanism $\mathscr{M}^j$.*
- *On action profile $\mathbf{a} = (\mathbf{a}^1, \ldots, \mathbf{a}^m)$ with $\mathbf{a}^j = (\mathsf{a}_1^j, \ldots, \mathsf{a}_n^j)$, the outcome of the mechanism is $\mathscr{M}(\mathbf{a}) = (\mathscr{M}^1(\mathbf{a}^1), \ldots, \mathscr{M}^m(\mathbf{a}^m))$.*
- *The action allocation rule is $\tilde{\boldsymbol{y}}(\mathbf{a})$ with $\tilde{y}_i(\mathbf{a}) = \max_j \tilde{y}_i^j(\mathbf{a}^j)$.*
- *The action payment rule is $\tilde{\boldsymbol{p}}(\mathbf{a})$ with $\tilde{p}_i(\mathbf{a}) = \sum_j \tilde{p}_i^j(\mathbf{a}^j)$.*

Theorem 6.3.1. *Revenue covering approximation is closed under simultaneous composition; i.e., if mechanisms $\mathscr{M}^1, \ldots, \mathscr{M}^m$ are downward closed, individually rational, and have revenue covering approximation μ; then their simultaneous composite mechanism $\mathscr{M}$ has revenue covering approximation μ.*

The following two lemmas, implied by downward closure and individual rationality, respectively, enable the proof of Theorem 6.3.1. Agent i's interim effective winner-pays-bid bid allocation rules, as notated in Definition 6.2.1, corresponding to $\mathscr{M}$ and each $\mathscr{M}^j$ are $\tilde{x}_i$ and $\tilde{x}_i^j$, respectively.

Lemma 6.3.2. *For the union environment of m downward-closed environments, allocation $\mathbf{x}$ is feasible if and only if there exists $\mathbf{x}^1, \ldots, \mathbf{x}^m$ feasible for the individual environments that satisfy $\mathsf{x}_i = \sum_j \mathsf{x}_i^j$ for all i and j.*

Proof. By definition of feasibility in the union environment, if $\mathbf{x}^1, \ldots, \mathbf{x}^m$ are feasible for the environment of $\mathscr{M}^1, \ldots, \mathscr{M}^m$, respectively, then

$$\mathsf{x}_i = \max_j \mathsf{x}_i^j \tag{6.3.1}$$

is feasible for the union environment of $\mathscr{M}$. Moreover, by downward closure of each individual mechanism $\mathscr{M}^j$ if $\mathbf{x}$ is feasible, then there exists $\mathbf{x}^1, \ldots, \mathbf{x}^m$ with each $\mathbf{x}^j$ feasible for $\mathscr{M}^j$ and

$$\mathsf{x}_i = \sum\nolimits_j \mathsf{x}_i^j \tag{6.3.2}$$

for all i and j. We are able to replace the maximization in equation (6.3.1) with the summation in equation (6.3.2) because downward closure allows the summation to be reduced to the maximum by removing service from an agent in all but at most one of the individual mechanisms. □

Lemma 6.3.3. *For the simultaneous composite mechanism* $\mathscr{M}$ *of* m *individually rational mechanisms* $\mathscr{M}^1, \ldots, \mathscr{M}^m$*, any agent* i*, and any effective winner-pays-bid* $\mathsf{b} \in \mathbb{R}$*,*

(i) *Agent* i*'s allocation probability with effective winner-pays-bid* b *is greater in* $\mathscr{M}$ *than in* $\mathscr{M}^j$ *for any* j*, i.e.,* $\tilde{x}_i(\mathsf{b}) \geq \tilde{x}_i^j(\mathsf{b})$*.*

(ii) *Agent* i*'s expected critical bid is smaller in* $\mathscr{M}$ *than in* $\mathscr{M}^j$ *for any* j*, i.e.,* $\hat{\mathsf{B}}_i \leq \hat{\mathsf{B}}_i^j$*.*

Proof. Fix any agent i. The pointset of equation (6.2.1) that defines the winner-pays-bid allocation rule for i in $\mathscr{M}$ contains that of $\mathscr{M}^j$ for all j as one allowable bid in $\mathscr{M}$ is to bid only in $\mathscr{M}^j$ (by individual rationality of the other mechanisms). As such, the smallest monotone function that contains this pointset is higher for $\mathscr{M}$ than for $\mathscr{M}^j$, i.e., $\tilde{x}_i(\mathsf{b}) \geq \tilde{x}_i^j(\mathsf{b})$ for all b. As $\hat{\mathsf{B}}$ and $\hat{\mathsf{B}}^j$ are defined as the area above winner-pays-bid allocation rules $\tilde{x}$ and $\tilde{x}^j$, the former is smaller than the latter. □

Proof of Theorem 6.3.1. Consider feasible allocation $\mathbf{y}$ for the composite mechanism and the following sequence of inequalities with explanation below.

$$\begin{aligned}
\mu\, \mathbf{E}[\text{Revenue}] &= \sum\nolimits_j \mu\, \mathbf{E}\big[\text{Revenue}_j\big] \\
&\geq \sum\nolimits_j \sum\nolimits_i \hat{\mathsf{B}}_i^j\, \mathsf{y}_i^j \\
&\geq \sum\nolimits_i \hat{\mathsf{B}}_i \sum\nolimits_j \mathsf{y}_i^j \\
&= \sum\nolimits_i \hat{\mathsf{B}}_i\, \mathsf{y}_i.
\end{aligned}$$

The first line follows from the definition of revenue as the sum of payments from all agents in all mechanisms. By Lemma 6.3.2 and the feasibility of $\mathbf{y}$ there exists $\mathbf{y}^1, \ldots, \mathbf{y}^m$ which are feasible for $\mathscr{M}^1, \ldots, \mathscr{M}^m$, respectively, and satisfy $\mathsf{y}_i = \sum_i \mathsf{y}_i^j$. The second line follows from revenue covering of $\mathscr{M}^j$ for each j with respect to $\mathbf{y}^j$. Swapping the order of summation and employing the lower bound of $\hat{\mathsf{B}}_i \leq \hat{\mathsf{B}}_i^j$ from Lemma 6.3.3 for all i and j gives the third line. The fourth line is from the definition of $\mathbf{y}^1, \ldots, \mathbf{y}^m$ in terms of $\mathbf{y}$. We are left with the inequality that shows that $\mathscr{M}$ has revenue covering approximation μ. □

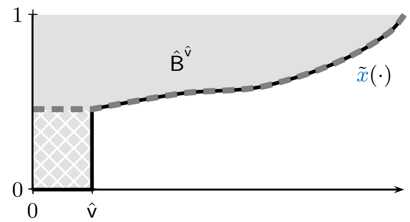

Figure 6.5. Geometry of best-response in a winner-pays-bid auction with reserve $\hat{v}$. The expected critical bid $\hat{B}$ is the area (gray solid and crosshatched) above the bid allocation rule $\tilde{x}(\cdot)$ (thin solid line), the expected critical bid with discounted reserve is the area (gray solid) above its cumulative distribution function $\tilde{x}^{\hat{v}}(\cdot)$ (thick dashed line).

6.4 Reserve Prices

We will shortly be analyzing the revenue of Bayes-Nash mechanisms like the first-price auction. As we understand from Chapter 3 and Chapter 4, reserve prices play an important role in revenue maximization. According to the previous definition of revenue covering approximation (Definition 6.2.2), auctions with reserve prices are not generally approximately revenue covered. Revenue covering arguments stem from relating the critical bid of an agent to potential payments of other agents. For example, in the first-price auction the critical bid of an agent is the maximum bid of the other agents, and if this agent does not bid above this critical bid then this maximum bid of the others is equal to the auction revenue. With a reserve price, an agent's critical bid may come from either bids of others or the reserve price. When the agent does not bid above her reserve price, the reserve price does not translate into auction revenue. See Figure 6.5.

In this section we adapt the framework of analysis to account for reserve prices. As is evident from Figure 6.5, the critical bid $\hat{B}$ as the area above the bid allocation rule $\tilde{x}(\cdot)$ over counts the contribution to revenue from the agent's critical bid. One resolution to this over counting is to explicitly discount the contribution to $\hat{B}$ from the reserve. The following definition captures this idea. Recall the bid allocation rule is equivalently the distribution function for the critical bid; thus, to discount the reserve is to assume the critical bid is zero whenever it would otherwise be the reserve.

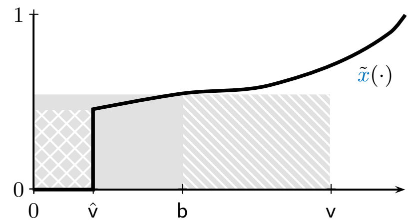

Figure 6.6. Geometric demonstration of equation (6.4.3). The expected discounted reserve $\hat{\mathsf{B}} - \hat{\mathsf{B}}^{\hat{\mathsf{v}}}$ (gray crosshatched region) is at most the expected payment from bid b (gray solid and crosshatched region). The surplus from bid b is the utility (gray striped region) plus the expected payment.

Definition 6.4.1. *The* critical bid with discounted reserve *is*

$$\hat{\mathsf{b}}^{\hat{\mathsf{v}}} = \begin{cases} 0 & \textit{if } \hat{\mathsf{b}} \leq \hat{\mathsf{v}}, \textit{ and} \\ \hat{\mathsf{b}} & \textit{otherwise.} \end{cases}$$

The cumulative distribution function for the critical bid with discounted reserve $\hat{\mathsf{v}}$ *is* $\tilde{x}^{\hat{\mathsf{v}}}(\mathsf{b}) = \tilde{x}(\max(\mathsf{b}, \hat{\mathsf{v}}))$*; see Figure 6.5; its expected value is:*

$$\hat{\mathsf{B}}^{\hat{\mathsf{v}}} = \mathbf{E}[\hat{\mathsf{b}}^{\hat{\mathsf{v}}}] = \int_0^\infty (1 - \tilde{x}(\max(\mathsf{b}, \hat{\mathsf{v}})))\, \mathrm{d}\mathsf{b}.$$

6.4.1 Surplus Approximates Value

We now lift the utility approximation of value result of Theorem 6.1.2 for mechanisms without reserves to mechanisms with reserves.

Theorem 6.4.1. *In any Bayes-Nash equilibrium of any mechanism and for any agent with value* v *exceeding her reserve* $\hat{\mathsf{v}}$*, the expected sum of her surplus and her critical bid with discounted reserve is an* $\frac{e}{e-1} \approx 1.58$ *approximation to her value; i.e.,*

$$\mathsf{v}\, x(\mathsf{v}) + \hat{\mathsf{B}}^{\hat{\mathsf{v}}} \geq {}^{e-1}/_{e}\, \mathsf{v}. \tag{6.4.1}$$

Proof. Theorem 6.1.2 states

$$u(\mathsf{v}) + \hat{\mathsf{B}} \geq {}^{e-1}/_{e}\, \mathsf{v}. \tag{6.4.2}$$

In BNE, an agent with $\mathsf{v} \geq \hat{\mathsf{v}}$ will bid $\mathsf{b} \geq \hat{\mathsf{v}}$ as any lower bid results in zero utility. Recall that the expected payment of an agent with equilibrium bid b is $p(\mathsf{v}) = \mathsf{b}\,\tilde{x}(\mathsf{b})$; geometrically as $\mathsf{b} \geq \hat{\mathsf{v}}$ this payment exceeds the amount of $\hat{\mathsf{B}}$ discounted by the reserve (Figure 6.6). Thus,

$$p(\mathsf{v}) + \hat{\mathsf{B}}^{\hat{\mathsf{v}}} \geq \hat{\mathsf{B}}. \tag{6.4.3}$$

Recall that surplus is utility plus payment, i.e., $\mathsf{v}\,x(\mathsf{v}) = u(\mathsf{v}) + p(\mathsf{v})$. The proof concludes by adding equation (6.4.2) to (6.4.3). □

6.4.2 Revenue Covering Approximation

For the appropriate definition of revenue covering approximation with reserves, revenue covering without reserves implies revenue covering with reserves.

Definition 6.4.2. *A mechanism with reserves has revenue covering approximation μ if the revenue covering approximation condition (Definition 6.2.2) holds with respect to expected critical bids with discounted reserves.*

Theorem 6.4.2. *Revenue covering approximation is closed under reserve pricing; i.e., if a mechanism $\mathscr{M}$ without reserves has revenue covering approximation μ, then with reserves it has revenue covering approximation μ.*

Proof. The revenue covering condition with reserves is only weaker as $\hat{\mathsf{B}}_i^{\hat{\mathsf{v}}} \leq \hat{\mathsf{B}}_i$ for all agents i. □

6.4.3 Social Surplus in Bayes-Nash Equilibrium

Theorem 6.4.3. *For any individually-rational mechanism with reserves that has revenue covering approximation $\mu \geq 1$, the expected social surplus in Bayes-Nash equilibrium is an $e/_{e-1}\,(1+\mu)$ approximation to the optimal surplus with the same reserves.*

Proof. Denote the reserves by $\hat{\mathbf{v}} = (\hat{\mathsf{v}}_1, \ldots, \hat{\mathsf{v}}_n)$. Consider a valuation profile $\mathbf{v}$. By Theorem 6.4.1, in BNE any agent i with $\mathsf{v}_i \geq \hat{\mathsf{v}}_i$ satisfies,

$$\mathsf{v}_i\,x_i(\mathsf{v}_i) + \hat{\mathsf{B}}_i^{\hat{\mathsf{v}}_i} \geq {}^{e-1}/_{e}\,\mathsf{v}_i.$$

Denote by $\boldsymbol{y}^\star(\mathbf{v}) = \operatorname{argmax}_{\mathbf{y}} \sum_{i:\mathsf{v}_i \geq \hat{\mathsf{v}}_i} \mathsf{v}_i\, \mathsf{y}_i$ the surplus optimizing allocation with reserves $\hat{\mathbf{v}}$ (with $y_i^\star(\mathbf{v}) = 0$ for i with $\mathsf{v}_i < \hat{\mathsf{v}}_i$). Thus, $\sum_i \mathsf{v}_i\, y_i^\star(\mathbf{v}) = \mathrm{REF}(\mathbf{v})$, the optimal social surplus with reserves $\hat{\mathbf{v}}$. Notice that $y_i^\star(\mathbf{v}) \in [0,1]$; thus,

$$\mathsf{v}_i\, x_i(\mathsf{v}_i) + \hat{\mathsf{B}}_i^{\hat{\mathsf{v}}_i}\, y_i^\star(\mathbf{v}) \geq {}^{e-1}\!/_e\, \mathsf{v}_i\, y_i^\star(\mathbf{v}).$$

The above equation was derived for agent i with $\mathsf{v}_i \geq \hat{\mathsf{v}}_i$; however, it holds trivially for i with $\mathsf{v}_i < \hat{\mathsf{v}}_i$ as $y_i^\star(\mathbf{v}) = 0$ for such agents. Sum over all agents i and invoke μ revenue covering,

$$\sum\nolimits_i \mathsf{v}_i\, x_i(\mathsf{v}_i) + \mu\, \mathbf{E}[\text{Revenue}] \geq {}^{e-1}\!/_e\ \mathrm{REF}(\mathbf{v}).$$

Take expectation over values $\mathbf{v}$ from the distribution $\boldsymbol{F}$,

$$\mathbf{E}[\text{Surplus}] + \mu\, \mathbf{E}[\text{Revenue}] \geq {}^{e-1}\!/_e\, \mathbf{E}[\mathrm{REF}(\mathbf{v})]\,.$$

The surplus of an individually-rational mechanism always exceeds its revenue; the theorem follows. □

An example consequence of Theorem 6.4.3 is the following. Moreover, analogous corollaries hold for the winner-pays-bid highest-bids-win matroid mechanism and the simultaneous composition of revenue covered mechanisms.

Corollary 6.4.4. *For (non-identical) independent, single-item environments, the expected social surplus of the first-price auction with reserves in Bayes-Nash equilibrium is an* $\frac{2e}{e-1} \approx 3.16$ *approximation to the optimal surplus with the same reserves.*

This approach for treating reserves applies to any mechanism that can be interpreted as having a reserve price. Importantly, our definition of reserves is in value space; while reserves, in the definition of a mechanism, bind in bid space. For the first-price auction and the simultaneous composition thereof, these are the same thing. For all-pay auctions, however, the value at which a bid-based reserve binds is endogenous to the equilibrium. For all-pay auctions, any bid-based reserves and BNE induce value-based reserves for which Theorem 6.4.3 holds.

6.5 Analysis of Revenue

We will adapt the framework for Bayes-Nash analysis of the surplus of mechanisms with reserves to analyze the revenue of Bayes-Nash equilibrium in mechanisms with monopoly reserves. Recall from Chapter 3 that the expected payment in BNE (and thus revenue) from an agent with value $\mathsf{v} \sim F$ satisfies $\mathbf{E}_{\mathsf{v}}[p(\mathsf{v})] = \mathbf{E}_{\mathsf{v}}[\varphi(\mathsf{v})\, x(\mathsf{v})]$ with virtual value function given by $\varphi(\mathsf{v}) = \mathsf{v} - \frac{1-F(\mathsf{v})}{f(\mathsf{v})}$ (see Section 3.3.1 on page 63). The approach will be to adapt Theorem 6.4.1, which bounds an agent's BNE surplus in terms of her value, to bound an agent's BNE virtual surplus in terms of her virtual value. Our analysis is necessarily restricted to regular distributions where the virtual value function $\varphi(\cdot)$ given above is monotone non-decreasing (see Definition 3.3.1 on page 66)

Theorem 6.5.1. *In any Bayes-Nash equilibrium of any mechanism and for any agent with value* v *exceeding her reserve* $\hat{\mathsf{v}}$ *and with non-negative virtual value* $\varphi(\mathsf{v})$*, the expected sum of her virtual surplus and her critical bid with discounted reserve is an* $e/e-1 \approx 1.58$ *approximation to her virtual value; i.e.,*

$$\varphi(\mathsf{v})\, x(\mathsf{v}) + \hat{\mathsf{B}}^{\hat{\mathsf{v}}} \geq {}^{e-1}/_{e}\, \varphi(\mathsf{v}). \tag{6.5.1}$$

Proof. The definition of virtual values for revenue as $\varphi(\mathsf{v}) = \mathsf{v} - \frac{1-F(\mathsf{v})}{f(\mathsf{v})}$ implies that $\mathsf{v} \geq \varphi(\mathsf{v})$ or, in other words, $\varphi(\mathsf{v})/\mathsf{v} \leq 1$. Thus, relative to the surplus and value terms of inequality (6.4.1) of Theorem 6.4.1, the virtual-surplus and virtual-value terms of (6.5.1) are scaled downward. Equivalently, the expected-critical-bid term on the right-hand side is relatively scaled upward. Thus, the inequality (6.5.1) of the present theorem is implied by Theorem 6.4.1. ☐

The following theorem is proved as was Theorem 6.4.3 but with the following key differences. The proof begins with the virtual surplus approximation of virtual value bound of Theorem 6.5.1 instead of the analogous bound of Theorem 6.4.1. It finishes by observing, as virtual surplus and revenue are equal in expectation, that expected virtual surplus plus expected revenue is exactly twice the expected revenue. Additionally, the theorem is stated for monopoly reserves and agents with regular distributions which necessarily excludes from analysis agents with negative virtual value.

Theorem 6.5.2. *For agents with regularly distributed values and any mechanism with monopoly reserves that has revenue covering approximation* $\mu \geq 1$*, the expected revenue in Bayes-Nash equilibrium is an* $e/e-1\,(1+\mu)$ *approximation to the optimal revenue.*

Again, this theorem can be applied to any of the revenue covered mechanisms previously discussed. The following corollary is for the first-price auction, e.g., there are similar corollaries for the winner-pays-bid highest-bids-win matroid mechanism and the simultaneous composition of mechanisms.

Corollary 6.5.3. *For any regular product distribution on values, the first-price auction with monopoly reserves has Bayes-Nash equilibrium revenue that is an* $^{2e}/_{e-1} \approx 3.16$ *approximation to the optimal revenue.*

In Section 5.2 on page 165 we saw that with sufficient competition the surplus maximizing mechanism (without reserves) approximates the revenue optimal mechanism (e.g., Theorem 5.2.4). Similar sufficient competition results extend to revenue covered mechanisms. One such definition of sufficient competition is that there are at least two agents from each distribution that are in direct competition with each other. The following theorem gives an example.

Theorem 6.5.4. *For any regular product distribution on values with at least two agents with values drawn from each distinct distribution, the first-price auction has Bayes-Nash equilibrium revenue that is an* $^{3e}/_{e-1} \approx 4.75$ *approximation to the optimal revenue.*

Proof. See Exercise 6.5. □

6.6 Revenue Covering Optimization

We have seen that the revenue covering approximation of a mechanism governs its Bayes-Nash approximation with respect to both social surplus and revenue. We now consider the problem of optimizing the rules of a mechanism to minimize its revenue covering approximation. The motivating example will be that of single-minded combinatorial auctions. We saw that the winner-pays-bid highest-bids-win mechanism for m-item single-minded combinatorial auctions is only a $\mu = m$ revenue covering approximation (Example 6.1.2). Faced with this negative result, the question remains of identifying a winner-pays-bid mechanism that obtains a non-trivial revenue covering approximation. Importantly, such a mechanism will have to choose a suboptimal, in terms of the surplus of bids, set of winners.

The running example for this section will be a single-minded combinatorial auction environment for n agents and m items. Each agent i has value v_i for obtaining bundle $\mathsf{S}_i \subset \{1, \ldots, m\}$. Two agents that desire the same item, i.e., i and $i^\dagger$ with $\mathsf{S}_i \cup \mathsf{S}_{i^\dagger} \neq \emptyset$, cannot simultaneously be served. The section culminates by showing that a winner-pays-bid mechanism based on a simple greedy heuristic has a revenue covering approximation of $\mu = \sqrt{m}$.

6.6.1 Non-bossiness, Approximation, and Greedy Algorithms

The difficulty of single-minded combinatorial auctions is that one agent can block many other agents that could be simultaneously served. It could be optimal to serve the blocked agents, but in equilibrium the blocking agent bids enough to dissuade any of the blocked agents from individually deviating to win. In Example 6.1.2 this situation was exhibited with one agent demanding the grand bundle $\{1, \ldots, m\}$ and many agents each demanding a single item; the grand-bundle agent then blocked all the singleton agents. When the grand-bundle agent bids 1, and the singleton agents bid 0, then the deviation bid that any singleton agent must make to win is 1. Since their values in the example are 1, this deviation does not improve the singleton agent's utility. Of course, the singleton agents would win if the sum of their bids exceeds the grand-bundle agent's bid of 1. Thus, as one of the singleton agent increases her bid — though, she continues to lose — the critical bids of all other singleton agents are reduced. This bad property is precisely what inhibits revenue covering approximation. The following definition formalizes the non-exhibition of this property.

Definition 6.6.1. *A mechanism is* subcritically non-bossy *if for any bid profile* $\mathbf{b}$*, critical bids* $\hat{\mathbf{b}}$*, and any other bid profile where losers may increase their bids up to their critical bids, i.e.,* $\mathbf{b}^{\dagger}$ *with* $\mathsf{b}_i^{\dagger} \in [\mathsf{b}_i, \max(\mathsf{b}_i, \hat{\mathsf{b}}_i)]$*, the same set of agents win under* $\mathbf{b}$ *and* $\mathbf{b}^{\dagger}$*.*[4]

To solve the combinatorial auction problem we are going to have to replace the highest-bids-win allocation rule with an allocation rule that is does not maximize the sum of the bids of the agents served. There are two potential losses from such an allocation rule. First, there is the direct loss from the fact that the allocation rule chooses a suboptimal set of bids. Even if there is a feasible set of agents with high bid sum, its revenue could be low. Second, there is the indirect loss from strategization on the

[4] This definition adopts the convention that ties in the bid allocation rule, when any loser increases her bid to equal her critical bid, are broken in favor of the current winners. The arguments below can be made without this tie-breaking convention by considering $\mathbf{b}^{\dagger}$ with losers bidding $\mathsf{b}_i^{\dagger} \in [\mathsf{b}_i, \max(\mathsf{b}_i, \hat{\mathsf{b}}_i - \epsilon)]$ for an arbitrarily small ϵ.

part of the agents. The highest-bids-win allocation rule suffers no direct losses, but prohibitively in indirect losses. On the other hand, the first-price auction for the grand bundle, i.e., where only one agent ever wins her desired bundle, suffers prohibitive direct losses but, as the first-price auction is revenue covered, suffers no indirect losses with respect to the optimal mechanism that only serves one agent). Ideally both direct and indirect losses should be kept small. The following definition formalizes a bound on the direct loss in terms of approximation.

Definition 6.6.2. *A mechanism* (APX) *with ex post bid allocation rule* $\tilde{\boldsymbol{x}}(\mathbf{b})$*, which maps a profile of bids to an allocation, is a* β approximation *to highest-bids-win* (REF) *if for all* $\mathbf{b}$*:*

$$\mathrm{APX}(\mathbf{b}) = \sum_i \mathsf{b}_i\, \tilde{x}_i(\mathbf{b}) \geq {}^1\!/\!_\beta \max_{\mathbf{x} \in \mathcal{X}} \mathsf{b}_i\, \mathsf{x}_i = {}^1\!/\!_\beta\, \mathrm{REF}(\mathbf{b}).$$

We now show that in a subcritically non-bossy mechanism the only loss in surplus is the direct loss from the non-optimality of the bid allocation rule, i.e., there is no indirect loss.

Theorem 6.6.1. *A winner-pays-bid subcritically non-bossy mechanism that is a* β *approximation to highest-bids-win has a revenue covering approximation of* $\mu = \beta$*.*

Proof. Fix a bid profile $\mathbf{b}$, the critical bid profile $\hat{\mathbf{b}}$, and any feasible allocation $\mathbf{y} \in \mathcal{X} \subset \{0,1\}^n$. Denote the bid allocation rule by $\tilde{\boldsymbol{x}}(\mathbf{b}) = (\tilde{x}_1(\mathbf{b}), \ldots, \tilde{x}_n(\mathbf{b})) \in \mathcal{X}$. Denote the maximum subcritical bid profile by $\mathbf{b}^\dagger$ with $\mathsf{b}_i^\dagger = \max(\mathsf{b}_i, \hat{\mathsf{b}}_i)$. Subcritical non-bossiness requires allocation to be unchanged if all losers increase their bids to their critical values, i.e., $\tilde{\boldsymbol{x}}(\mathbf{b}^\dagger) = \tilde{\boldsymbol{x}}(\mathbf{b})$.

The following sequence of equations implies that the mechanism has revenue covering approximation $\mu = \beta$; formal justification for each equation follows.

$$\begin{aligned}
\mathrm{Revenue}(\mathbf{b}) &= \sum_i \mathsf{b}_i\, \tilde{x}_i(\mathbf{b}) \\
&= \sum_i \mathsf{b}_i^\dagger\, \tilde{x}_i(\mathbf{b}) \\
&= \sum_i \mathsf{b}_i^\dagger\, \tilde{x}_i(\mathbf{b}^\dagger) \\
&\geq {}^1\!/\!_\beta \sum_i \mathsf{b}_i^\dagger\, \mathsf{y}_i \\
&\geq {}^1\!/\!_\beta \sum_i \hat{\mathsf{b}}_i\, \mathsf{y}_i.
\end{aligned}$$

The first equation is by definition of winner-pays-bid mechanisms. The second equation is the equality of b_i and $\mathsf{b}_i^\dagger = \max(\mathsf{b}_i, \hat{\mathsf{b}}_i)$ where winning

($\tilde{x}_i(\mathbf{b}) = 1$) implies $\mathsf{b}_i \geq \hat{\mathsf{b}}_i$. The third equation is by subcritical non-bossiness, as discussed above. The fourth equation follows by the β-approximate optimality of $\tilde{\boldsymbol{x}}(\cdot)$ on $\mathbf{b}^\dagger$. The fifth and final equation follows from the definition of $\mathsf{b}_i^\dagger = \max(\mathsf{b}_i, \hat{\mathsf{b}}_i) \geq \hat{\mathsf{b}}_i$. We conclude that the mechanism has a revenue covering approximation of $\mu = \beta$. □

Theorem 6.6.1 shows that to find winner-pays-bid mechanisms that are revenue covered it suffices to find a subcritically non-bossy mechanism that is a good approximation to highest-bids-win. A greedy algorithm is one that considers the agents sorted according to some priority and serves each agent if it is feasible to do so given the agents previously served by the algorithm. Greedy algorithms are a standard design methodology in the field of approximation algorithms and they have important consequences for mechanism design. For example, we saw in Section 4.6 on page 135 that greedy algorithms are optimal in ordinal environments such as those given by a matroid set system. Subsequently, we will see that greedy algorithms are approximately optimal in some environments and mechanisms based on them are subcritically non-bossy.

Definition 6.6.3. *For any downward-closed environment and any profile of* priority functions $\boldsymbol{\vartheta} = (\vartheta_1, \ldots, \vartheta_n)$*, the* greedy-by-priority algorithm *on bids* $\mathbf{b}$ *is:*

(i) *Sort the agents in decreasing order of priority* $\vartheta_i(\mathsf{b}_i)$ *(and discard all agents with negative priority).*
(ii) *Initialize* $\mathbf{x} \leftarrow \mathbf{0}$ *(the null assignment).*
(iii) *For each agent* i *(in sorted order), set* $\mathsf{x}_i \leftarrow 1$ *if* $(1, \mathbf{x}_{-i})$ *is feasible. (I.e., serve* i *if* i *can be served alongside previously served agents.)*
(iv) *Output allocation* $\mathbf{x}$.

Theorem 6.6.2. *The greedy-by-priority bid allocation rule is subcritically non-bossy.*

Proof. Fix a profile of bids $\mathbf{b}$, critical bids $\hat{\mathbf{b}}$, and maximum subcritical bid profile $\mathbf{b}^\dagger$ with $\mathsf{b}_i^\dagger = \max(\mathsf{b}_i, \hat{\mathsf{b}}_i)$. Consider varying a single losing bid i on the range $[0, \hat{\mathsf{b}}_i]$ and simulating the algorithm. Wherever this bid arises in the sorted order of agents by priority, since $\mathsf{b}_i \leq \hat{\mathsf{b}}_i$, it must be infeasible to serve the agent. Thus, this agent is discarded and all decisions by the algorithm to serve or not to serve any other agents are unaffected. The same holds for all losing agents simultaneously. Thus, the bid allocation rule is subcritically non-bossy. □

6.6.2 Matroid Auctions

This section takes a brief detour to reconsider the winner-pays-bid highest-bids-win mechanism in matroid environments. Matroid environments include multi-unit environments and constrained matching environments. In Section 4.6 on page 135, matroid environments were defined and characterized as those where the greedy-by-bid algorithm optimizes the sum of bids of the agents it selects (Theorem 4.6.3). In other words, matroid feasibility constraints are precisely those where greedy-by-bid $\beta = 1$ approximates highest-bids-wins. We obtain the following corollary to Theorem 6.6.1.

Corollary 6.6.3. *For matroid environments, the winner-pays-bid highest-bids-win mechanism is revenue covered, i.e.,* $\mu = 1$*; its surplus in Bayes-Nash equilibrium is an* $^e/_{e-1}$ *approximation to the optimal surplus; and with monopoly reserves and regular distributions its revenue is an* $^e/_{e-1}$ *approximation to the optimal revenue.*

6.6.3 Single-minded Combinatorial Auctions

We now instantiate the approach of the preceding section to design an m-item single-minded combinatorial auction that has a non-trivial revenue covering approximation. Theorem 6.6.1 and Theorem 6.6.2 imply that, to find a winner-pays-bid mechanism that is revenue covered, it suffices to identify a profile of priority functions $\boldsymbol{\vartheta}$ such that the greedy-by-priority algorithm obtains a good approximation to highest-bids-win. We now consider this task and identify a priority for which greedy-by-priority is a $\beta = \sqrt{m}$ approximation and, thus, has a revenue covering approximation of $\mu = \sqrt{m}$.

We begin by considering two extremal approaches, both of which yield only an m approximation, and then look at trading off these extremes to obtain the desired $\sqrt{m}$ approximation. The first failed approach to consider is *greedy by bid*, i.e., the priority function is the identity $\vartheta_i(\mathsf{b}_i) = \mathsf{b}_i$. This algorithm is bad because it is an m approximation on the following $n = m + 1$ agent environment and bids. Agents i, for $0 \leq i \leq m$, bid $\mathsf{b}_i = 1$ and demand $\mathsf{S}_i = \{i\}$; agent $m + 1$ bids $\mathsf{b}_{m+1} = 1 + \epsilon$ (for some small $\epsilon > 0$) and demands the grand bundle $S_{m+1} = \{1, \ldots, m\}$. See Figure 6.7(a) with $A = 1$ and $B = 1 + \epsilon$. Greedy-by-bid orders agent $m + 1$ first, this agent is feasible and therefore served. All remaining agents are infeasible after agent $m + 1$ is served. Therefore, the algorithm serves only this one agent and has surplus $1 + \epsilon$. Of course highest-bids-win

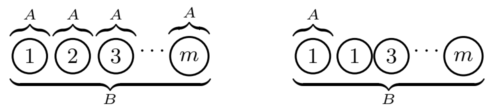

(a) m small agents. (b) 1 small agent.

Figure 6.7. Challenge cases for greedy orderings as a function of bid and bundle size.

serves the m small agents for a total surplus of m. The approximation factor of greedy-by-bid is the ratio of these two performances, i.e., m.

Obviously what went wrong in greedy-by-bid is that we gave preference to an agent with large demand who then blocked a large number of mutually-compatible small-demand agents. We can prevent this situation by instead prioritizing agents by bid-per-size, i.e., $\vartheta(\mathsf{b}_i) = \mathsf{b}_i/|\mathsf{S}_i|$. *Greedy by bid-per-size* also fails on the following $n = 2$ agent input. Agent 1 bids $\mathsf{b}_1 = 1+\epsilon$ and demands $S_1 = \{1\}$, and agent 2 bids $\mathsf{b}_2 = m$ and demands the grand bundle $S_2 = \{1, \ldots, m\}$. See Figure 6.7(b) with $A = 1 + \epsilon$ and $B = m$. Greedy-by-bid-per-size orders agent 1 first, this agent is feasible and therefore served. Agent 2 is infeasible after agent 1 is served. Therefore, the algorithm serves only agent 1 and has surplus $1 + \epsilon$. Of course highest-bids-win serves agent 2 and has surplus of m. The approximation factor of greedy-by-bid-per-size is the ratio of these two performances, i.e., m.

The flaw with this second algorithm is that it makes the opposite mistake of the first algorithm; it undervalues large-demand agents. While we correctly realized that we need to trade off bid for demand size, we have only considered extremal examples of this trade-off. To get a better idea for this trade-off, consider the cases of a single large-demand agent and either m small-demand agents or 1 small-demand agent. We will leave the bids of the two kinds of agents as variables A for the small-demand agent(s) and B for the large-demand agent. Assume, as in our previous examples, that $mA > B > A$. These settings are depicted in Figure 6.7.

Notice that any greedy algorithm that orders by some function of bid and size will either prefer A-bidding or B-bidding agents in both cases. The A-preferred algorithm has surplus Am in the m-small-agents case and surplus A in the 1-small-agent case. The B-preferred algorithm has surplus B in both cases. The highest-bids-win outcome, on the other hand, has surplus mA in the m-small-agents case and surplus

	m small agents	1 small agent	approximation
highest-bids-win	mA	B	1
A-preferred	mA	A	B/A
B-preferred	B	B	mA/B

Figure 6.8. Performances of A- and B-preferred greedy algorithms and their approximation to highest-bids-win in worst-case over the two cases.

B in the 1-small-agent case. Therefore, the worst-case approximation for A-preferred is B/A (achieved in the 1-small-agent case), and the worst-case approximation for B-preferred is mA/B (achieved in the m-small-agents case). These performances and worst-case ratios are summarized in Figure 6.8.

If we are to use the greedy algorithm design paradigm we need to minimize the worst-case ratio. The approach suggested by the analysis of the above cases would be trade off A versus B to equalize the worst-case approximation, i.e., when $B/A = mA/B$. Here m was a stand-in for the size of the large-demand agent. The suggested algorithm is greedy by bid-per-square-root-size which orders the agents by the priority $\vartheta_i(\mathsf{b}_i) = \mathsf{b}_i/\sqrt{|\mathsf{S}_i|}$. The tradeoff above can be observed explicitly in the proof of Theorem 6.6.4, below.

Theorem 6.6.4. *For m-item single-minded combinatorial auctions environments, the greedy by bid-per-square-root-size algorithm is a $\beta = \sqrt{m}$ approximation to highest-bids-win.*

Proof. Let APX represent the greedy by bid-per-square-root-size algorithm and its surplus; let REF represent the optimal algorithm and its surplus. Let J be the set selected by APX and I be the set selected by REF. We will proceed with a *charging argument* to show that if $j \in J$ blocks some set of agents $\mathsf{C}_j \subset I$ then the sum of bids of the blocked agents is not too large relative to the bid of agent j.

Consider the agents sorted (as in APX) by $\mathsf{b}_i/\sqrt{|\mathsf{S}_i|}$. For an agent $i \in I$ not to be served by APX, it must be that at the time it is considered by the greedy algorithm, another agent j has already been selected that blocks i, i.e., the bundles S_j and S_i have non-empty intersection. Intuitively we will charge one such agent j with the loss from not accepting agent i. We define C_j as the set of all $i \in I$ that are charged to j as described above. Of special note, if $i \in J$, i.e., agent i was not yet blocked when considered by APX, we charge it to itself, i.e., $\mathsf{C}_i = \{i\}$. Notice that the sets C_j for winners $j \in J$ of APX partition the winners I of REF.

The theorem follows from the inequalities below. Explanations of each non-trivial step are given afterwards.

$$\begin{aligned} \text{REF} &= \sum\nolimits_{i \in I} \mathsf{b}_i = \sum\nolimits_{j \in J} \sum\nolimits_{i \in \mathsf{C}_j} \mathsf{b}_i && (6.6.1) \\ &\leq \sum\nolimits_{j \in J} \tfrac{\mathsf{b}_j}{\sqrt{|\mathsf{S}_j|}} \sum\nolimits_{i \in \mathsf{C}_j} \sqrt{|\mathsf{S}_i|} && (6.6.2) \\ &\leq \sum\nolimits_{j \in J} \tfrac{\mathsf{b}_j}{\sqrt{|\mathsf{S}_j|}} \sum\nolimits_{i \in \mathsf{C}_j} \sqrt{m/|\mathsf{C}_j|} && (6.6.3) \\ &= \sum\nolimits_{j \in J} \tfrac{\mathsf{b}_j}{\sqrt{|\mathsf{S}_j|}} \sqrt{m\,|\mathsf{C}_j|} && (6.6.4) \\ &\leq \sum\nolimits_{j \in J} \mathsf{b}_j \sqrt{m} = \sqrt{m} \cdot \text{APX}\,. && (6.6.5) \end{aligned}$$

Line (6.6.1) follows because C_j partition I. Line (6.6.2) follows because $i \in \mathsf{C}_j$ implies that j precedes i in the greedy ordering and therefore $\mathsf{b}_i \leq \mathsf{b}_j \sqrt{|\mathsf{S}_i|}/\sqrt{|\mathsf{S}_j|}$. The demand sets S_i of $i \in \mathsf{C}_j$ are disjoint (because they are a subset of I which is feasible and therefore disjoint). Thus, we can bound $\sum_{i \in \mathsf{C}_j} |\mathsf{S}_i| \leq m$. The square-root function is concave and the sum of a concave function applied to terms with a fixed total sum is maximized when each term is equal, i.e., when $|\mathsf{S}_i| = m/|\mathsf{C}_j|$. Therefore, $\sum_{i \in \mathsf{C}_j} \sqrt{|\mathsf{S}_i|} \leq \sum_{i \in \mathsf{C}_j} \sqrt{m/|\mathsf{C}_j|}$ and line (6.6.3) follows. Line (6.6.4) follows from independence of the inner summand on i. Finally, line (6.6.5) follows because the bundle S_i of each agent $i \in \mathsf{C}_j$ is disjoint but contains some demanded item in S_j and, therefore, $|\mathsf{C}_j| \leq |\mathsf{S}_j|$. □

We conclude the section with the following corollary. The first part is a consequence of Theorem 6.6.1, Theorem 6.6.2, and Theorem 10.1.4. The second part is a consequence of the first part and Theorem 6.1.4. The third part is a consequence of the first part and Theorem 6.5.2.

Corollary 6.6.5. *For m-item single-minded combinatorial environments, the winner-pays-bid greedy-by-value-per-square-root-size mechanism has revenue covering approximation $\mu = \sqrt{m}$; its surplus in Bayes-Nash equilibrium is an ${}^{e}/{}_{e-1}\sqrt{m}$ approximation to the optimal surplus; and with monopoly reserves and regular distributions its revenue is an ${}^{e}/{}_{e-1}(1+\sqrt{m})$ approximation to the optimal revenue.*

Exercises

6.1 Consider an all-pay auction and show, analogously to Theorem 6.1.2, that the agents utility and expected critical bid combine to approximate the agent's value in Bayes-Nash equilibrium. Specifically,

$$u(\mathsf{v}) + \hat{\mathsf{B}} \geq {}^{1}/_{2}\,\mathsf{v},$$

where $\hat{\mathsf{B}}$ is the expected critical bid of the agent. Combine this result with revenue covering (with respect to the all-pay bid-allocation rule) to show that the expected social surplus of the all pay auction is a two approximation to the optimal social surplus.

6.2 Consider the single-dimensional routing environment discussed in Section 1.1.3 on page 14 where there is a graph $G = (V, E)$, each agent i has a message to send from source vertex $s_i \in V$ to target vertex $t_i \in V$ (public knowledge) and a private value v_i for sending such a message. A feasible outcome is given by an edge disjoint collection of paths in the graph. Show that the revenue covering approximation μ of the winner-pays-bid highest-bids-win mechanism can be as large as the diameter of the graph, i.e., the maximum over pairs of vertices of the shortest path between the pair.

6.3 Consider the single-minded combinatorial auction problem of Example 6.1.2. Show by example that there is a single-minded combinatorial environment and a Bayes-Nash equilibrium in the winner-pays-bid highest-bids-win mechanism with welfare that is only an m approximation to the optimal welfare. (Hint: Nash equilibrium is a special case of Bayes-Nash equilibrium where the value distributions are deterministic.)

6.4 Prove Theorem 6.2.2: *In single-item environments the all-pay auction is 2 revenue covered.*

6.5 Prove Theorem 6.5.4: *For any regular product distribution on values with at least two agents with values drawn from each distinct distribution, the first-price auction has Bayes-Nash equilibrium revenue that is an* ${}^{3e}/_{e-1} \approx 4.75$ *approximation to the optimal revenue.*

Chapter Notes

Vickrey (1961) posed the question of solving for the equilibrium in the first-price auction and two agents with values drawn from the uniform distribution with asymmetric supports. The solution when the lower bound of the supports is the same, as in the $U[0, 1]$ and $U[0, 2]$ case of Example 6.1.1, was given by Griesmer et al. (1967). The general case of two agents with arbitrary uniform distributions was solved by Kaplan and Zamir (2012).

The quantification of the disutility of equilibrium versus the social surplus maximizing outcome is known as the *price of anarchy*. This topic of study was initially proposed by Koutsoupias and Papadimitriou (1999). It was applied to (full information) congestion games by Roughgarden and Tardos (2002), cf. the routing game of Section 1.1 on page 3. Roughgarden (2012a) abstracted the canonical price of anarchy analysis as what is referred to as the *smoothness framework*. Roughgarden (2012b) and Syrgkanis and Tardos (2013) generalize this smoothness framework to games of incomplete information and auctions, respectively. There has been extensive study of the price of anarchy of specific auction games to which detailed reference is omitted. This text focuses on an adaptation of the smoothness paradigm to single-dimensional agents that was given by Hartline et al. (2014).

The proof that the sum of utility and critical bid approximate an agents value for first-price auctions that is given in this text is from Syrgkanis and Tardos (2013); an alternative geometric argument can be found in Hartline et al. (2014). This approximation analysis of BNE of the first-price auction is not tight; its approximation factor is known to be between 1.15 (Hartline et al., 2014) and 1.35 (Hoy et al., 2018). The improved analysis of the all-pay auction of Exercise 6.1 is based on Syrgkanis and Tardos (2013). A smoothness framework for analyzing the simultaneous composition of auctions was first given by Syrgkanis and Tardos (2013); the analysis given here is the refinement of Hartline et al. (2014) for single-dimensional agents. The analysis of revenue in Bayes-Nash equilibrium is from Hartline et al. (2014).

The relationship between revenue covering approximation and greedy algorithms is a recasting of the main result of Lucier and Borodin (2010) into the analysis framework of Hartline et al. (2014).

The analysis of Syrgkanis and Tardos (2013) is more general than the one presented here primarily in that it allows for multi-dimensional agent preferences. They also give numerous results that are not covered here; one such result is for the sequential composition of mechanisms, i.e., when mechanisms are run one after the other.