Foundations and Trends® in
Theoretical Computer Science
Vol. 8, No. 3 (2012) 143–263

DOI: 10.1561/0400000045

Bayesian Mechanism Design

By Jason D. Hartline

Contents

Foundations and Trends® in
Theoretical Computer Science
Vol. 8, No. 3 (2012) 143–263

DOI: 10.1561/0400000045

Bayesian Mechanism Design

Jason D. Hartline

Northwestern University, Evanston, 60208, USA,
hartline@eecs.northwestern.edu

Abstract

Systems wherein strategic agents compete for limited resources are ubiquitous: the economy, computer networks, social networks, congestion networks, nature, etc. Assuming the agents' preferences are drawn from a distribution, which is a reasonable assumption for small mechanisms in a large system, *Bayesian mechanism design* governs the design and analysis of these systems.

This article surveys the classical economic theory of Bayesian mechanism design and recent advances from the perspective of algorithms and approximation. Classical economics gives simple characterizations of Bayes-Nash equilibrium and optimal mechanisms when the agents' preferences are linear and single-dimensional. The mechanisms it predicts are often complex and overly dependent on details of the model. Approximation complements this theory and suggests that simple and less-detail-dependent mechanisms can be nearly optimal. Furthermore, techniques from approximation and algorithms can be used to describe good mechanisms beyond the single-dimensional, linear model of agent preferences.

1

Introduction

The Bayesian[1] approach to optimization assumes that an input to the optimization problem is drawn from a distribution and requests that an output be found that is good in expectation. This approach is both compatible with the standard worst-case approach to algorithm design and suggests a rich problem space that is relatively unexplored from an algorithmic perspective. The approach is to partition the optimization problem into two stages. In the first stage, a distribution is given and can be preprocessed so as to construct an algorithm tailored to the distribution. In the second stage inputs to the distribution are drawn and the algorithm is run on these inputs. For instance, this approach is similar to data structure problems such as Huffman [64] coding and search problems such as nearest neighbor, cf., Cover and Hart [36].

[1] We say "Bayesian" here instead of "stochastic" to connote settings where *Bayesian updating* is relevant. As an example, if the distribution over inputs is correlated and an optimization algorithm must make decisions before all the input is known, the algorithm should perform Bayesian updating to refine its beliefs about remaining inputs. That said, many of the methods and results described in this survey pertain to the special case that the inputs are independent and the term "Bayesian" is only used for consistency in presenting the result in the context of the greater literature within which they are contained.

Furthermore, the algorithm practitioner's design and analysis framework can be abstracted as being Bayesian. As a caricature, such a practitioner will design and test her algorithm on a class of data sets that she deems relevant. When the attempted improvements to the algorithm on these data sets fail to significantly increase performance, she will deploy the final algorithm. Methodologies for the Bayesian framework for algorithm design can then potentially be applied to this form of practical algorithm design (see, e.g., the *Bayesian reductions* of Section 7).

This survey focuses on a large class of Bayesian optimization problems that exhibit both algorithmic and economic challenges. In particular, in a departure from traditional algorithms, we assume that the desired input to the algorithm is the private information of self-interested agents. These agents can report this information to the algorithm (if requested) but can equally well misreport if it benefits them and the algorithm cannot tell the difference. Our abstract question pertains to how a designer should structure the rules of a system so that in the equilibrium of strategic play, the designer's objective is optimized. This is the question of *mechanism design*.

Mechanism design has broad applications; however, we will focus on potential applications to computer systems. In so far as scarce system resources must be shared amongst parties with diverse and selfish interests, mechanism design governs the proper working of computer systems. When incentives are not properly accounted for, strategic misbehavior is common. A few examples include spam in email systems (see Dwork et al. [44]), link-spam and Sybil networks in the ranking of Internet search results (see Dwork et al. [45] and Gyongyi and Garcia-Molina [55]), and freeloading in file-sharing networks (see Vishnumurthy et al. [92]).

The literature that melds algorithmic and economic issues in mechanism design, a.k.a., *algorithmic mechanism design*, since its inception over a decade ago, has predominantly focused on worst-case mechanisms with dominant strategy equilibria. The strategic agents should have a single best action no matter what the actions of other agents, and the designed mechanism should be good for any preferences the agents may have. As a motivating example, the second-price auction for selling a single item serves the highest bidder at price equal to the second highest bid. Truthtelling is a dominant strategy in the second-price auction; therefore, in dominant strategy equilibrium the bidder with the highest true value wins and, for any values the bidders might possess, the auction maximizes social welfare. In this line of worst-case dominant-strategy mechanism design, Lehmann et al. [70] initiated the merger of algorithmic complexity and economic incentive considerations for the objective of social welfare maximization, Nisan and Ronen [75] initiated the study of mechanisms for non-linear objectives such as makespan in machine scheduling, and Goldberg et al. [50] initiated the study of objectives that depend on agent payments (e.g., revenue). The latter two agendas are related in that, even absent computational constraints, the economic incentives of the agents preclude a pointwise optimal mechanism. As an example, notice that for the single-item auction, a seller would could obtain more revenue by setting a reserve price between the highest and second highest value, but to do so requires knowledge of the agent values. Therefore, for these problems the worst-case optimization question under consideration is inherently one of approximation akin to the competitive analysis of online algorithms (e.g., Borodin and El-Yaniv [16]).

In the classic microeconomic treatment of mechanism design, the non-pointwise-optimality of mechanisms is resolved by formulating the problem as one of Bayesian design. A distribution over the preferences of the agents is given, and the designer seeks to optimize her objective in expectation over this distribution. For any specific distribution, there is such an optimal mechanism. The Bayesian assumption also allows the strategic incentives of the agents to be relaxed. Given that the agents' preferences are drawn from a distribution, instead of requiring a mechanism to have dominant strategies, mechanisms can be considered where agents' actions are best responses to the distribution of actions of other agents.

Both the Bayesian and worse-case (henceforth, *prior-free*) frameworks for mechanism design have merits and, recognizing this, a Bayesian branch of algorithmic mechanism design has emerged. Paramount of study in *Bayesian algorithmic mechanism design* are algorithmic techniques, approximation, and computational issues. The goal of this survey is to discuss the most fundamental of these results in the context of classical Bayesian mechanism design.

1.1 Topics Covered

This survey is organized into the following sections. For environments with agents with linear and single-dimensional preferences, Section 2 characterizes the equilibrium of strategic play in auction-like games, Section 3 characterizes optimal auctions for welfare and revenue, and Section 4 describes simple, practical mechanism that are approximately optimal. For environments where agents have multi-dimensional and non-linear preferences, Section 5 characterizes optimal auctions and Section 6 describes simple approximation mechanisms. Finally, Section 7 gives generic performance-preserving reductions from mechanism design to algorithm design for both single-dimensional and multi-dimensional agent preferences. Mathematical reference is given in Appendix A.

Section 2: Equilibrium. In Bayesian games an agent's strategy maps her private information to an action in the game; the distribution of private information and strategies induce a distribution of actions in the game; and a profile of strategies is in equilibrium if the strategy of each agent is a best response to this distribution of actions. The classical Bayesian approach to mechanism design starts with a characterization of equilibrium. For single-dimensional environments, i.e., when an agent's private preference is given by a single number denoting her value for a single abstract service, Myerson [74] characterized all possible equilibria. This characterization states that an agent's probability of service should be monotonically non-decreasing in her value and it gives a formula for her expected payment. This payment identity implies the *revenue equivalence* of auctions with the same equilibrium outcome. Moreover, it suggests a method for solving for equilibrium.

Section 3: Optimal Mechanisms. For welfare maximization, the well known Vickrey-Clarke-Groves (VCG) mechanism is pointwise optimal [91, 35, 53]. For revenue maximization in single-dimensional environments, Myerson [74] gives a reduction from the non-pointwise problem of maximizing revenue in expectation over the distribution, to the problem of optimizing a pointwise virtual welfare. The reduction works by transforming the agent values to *virtual values*, and then optimizing in the transformed space. In the special case where values are i.i.d. from a distribution from a sufficiently well-behaved class, this optimal mechanism is simply the second-price auction with a suitably chosen reserve price. Bulow and Roberts [20] give a microeconomic reinterpretation Myerson's virtual values as the derivative of an appropriate *revenue curve*, a.k.a., as a *marginal revenue*. Alaei et al. [5] observe that this characterization is based on the "revenue linearity" of optimal mechanisms for single-dimensional agents.

Section 4: Approximation Mechanisms. While the optimal mechanism is often complex and impractical, there is often a simple and practical mechanism that is approximately optimal. For example, the *second-price auction with reserve* is widely prevalent even though it is not optimal beyond the ideal setting of symmetric agents with values drawn from a well-behaved distribution. Of course, even more prevalent are simple posted-pricing mechanisms, e.g., most stores sell goods by posting take-it-or-leave-it while-supplies-last prices on the goods for sale. Approximation can resolve this disconnect between theory and practice. Chawla et al. [29, 30] and Hartline and Roughgarden [62] show that reserve pricing is approximately optimal in many environments. Similarly, Chawla et al. [30], Yan [94], and Chakraborty et al. [27] show that posted pricings (take-it-or-leave-it while-supplies-last prices) are approximately optimal quite broadly.

As discussed previously, for many mechanism design problems there is not a single optimal mechanism. For a given prior distribution, the optimal mechanism generally depends on the prior distribution. Nonetheless, there may still be a single *prior-independent* mechanism that is approximately optimal. I.e., for any distribution, the prior-independent mechanism approximates the optimal mechanism for that distribution. Dhangwatnotai et al. [41] show that a single sample from the distribution gives a sufficient market analysis for obtaining a two approximation to the revenue-optimal auction in many single-dimensional environments; moreover, a single sample can be easily attained on-the-fly as the mechanism is being run. Hartline and Roughgarden [61] provide a post hoc Bayesian justification of the preceding literature on prior-free revenue approximation (e.g., Goldberg et al. [49]) and show that prior-free approximation with respect to an appropriate prior-free benchmark implies prior-independent approximation.

Section 5: Multi-dimensional and Non-linear Preferences. Revenue maximization in multi-dimensional environments is much more complex than in single-dimensional environments. Nonetheless, a series of recent papers has given an algorithmic generalization of the single-dimensional reduction to virtual welfare maximization of Myerson [74] to multi-dimensional preferences. Cai et al. [22] and Alaei et al. [4] do so for environments where welfare is optimized by a greedy algorithm; and Cai et al. [23] extend these results to general environments with additive preferences. These results tie the theory of optimal mechanism design to a natural convex optimization problem. While these approaches result in mechanisms that have polynomial complexity in the number of agents, they rely on brute-force solutions to single-agent pricing problems.

An important special case of multi-dimensional mechanism design is single-agent pricing; the unit-demand pricing problem is a paradigmatic challenge problem. Consider a single agent who desires one of a set of items. This agent is multi-dimensional in that she may have a distinct value for each item. The agent's multi-dimensional preference is drawn from a distribution and, given this distribution, we would like to price items to maximize revenue. Briest et al. [18] show that this problem can be solved in time polynomial in the number of distinct agent types. Unfortunately, when the agent's values for the items are independent, the type space is exponentially big in the number of items.

Section 6: Approximation for Multi-dimensional and Non-linear Preferences. The aforementioned unit-demand pricing problem with independently distributed values can be simplified with approximation. Chawla et al. [29, 30] show that there is a simple two approximation to the optimal item pricing. Moreover, Chawla et al. [32] show that this two approximation to the optimal item pricing is in fact also a four approximation to the optimal mechanism which, in addition to pricing items, may also price lotteries over items. Chawla et al. [30, 32], Bhattacharya et al. [11], Chakraborty et al. [27], and Alaei [2] extend these results quite broadly to show that often for unit-demand auction problems simple to find posted-pricing-based mechanisms are approximately optimal.

As described above, in both single-dimensional and multi-dimensional revenue maximization, the (non-pointwise) Bayesian mechanism design problem reduces to a (pointwise) virtual welfare maximization problem. It should be noted, however, that the single-dimensional case and multi-dimensional case have rather different structure. In the single-dimensional case, the transformation from value space to virtual value space is deterministic and separates across agents; whereas, in the multi-dimensional case, the transformation requires solving a convex optimization problem on all the agents together and it may be stochastic. Alaei et al. [5] show that, in fact, there is a transformation for the multi-dimensional case, with similar structure and economic intuition as in the single-dimensional case, that is approximately optimal.

Section 7: Computation and Approximation Algorithms. It is standard (from the prior-free mechanism design literature) that there is a reduction from exact welfare maximization with incentives to exact welfare maximization without incentives. I.e., if we have an optimal algorithm for maximizing welfare, we can convert that algorithm into an optimal mechanism that, in the dominant-strategy equilibrium of strategic play, maximizes welfare. The resulting mechanism is known as the Vickrey-Clarke-Groves (VCG) mechanism [35, 53, 91]. Lehmann et al. [70] point out, however, this reduction is incompatible with generic approximation algorithms. I.e., from a generic approximation algorithm we cannot instantiate the reduction to obtain a generic approximation mechanism which has an equilibrium with performance comparable to the original algorithm.

While mounting evidence suggests that such generic approximation-compatible reductions (see Chawla et al. [31] and Dobzinski and Vondrák [42]) do not exist for prior-free mechanisms, they do for Bayesian mechanisms. For welfare maximization in single-dimensional environments, Hartline and Lucier [60] give a generic approximation-preserving reduction from Bayesian mechanism design to Bayesian algorithm design. (Notice that, of course, a worst-case algorithm is also a Bayesian algorithm.) Via the reduction from revenue to welfare maximization of Myerson [74], this reduction can be adapted to the revenue objective. Hartline et al. [59] and Bei and Huang [9] generalize the above single-dimensional reduction to multi-dimensional environments. The multi-dimensional approach is brute-force in each agent's type space and it is an open question as to whether a similar reduction exists for large but succinctly represented type spaces.

Appendix A: Mathematical Reference. A number of mathematical constructs play a prominent role in our treatment of Bayesian mechanism design. These are *submodular set functions* which capture the concept of diminishing returns, *matroid set systems* which represent substitutability (e.g., Oxley [76]), and *convex optimization* within which most questions in Bayesian mechanism design reside (e.g., Schrijver [86]).

1.2 Topics Omitted

Having described above the material covered by this survey, we now turn to related material that is not covered. We have omitted discussion of almost all of the literature on prior-free mechanism design. Prior-free mechanism design and recent results relating to computational tractability and approximation are a topic warranting a survey of their own, which to cover adequately, would be even longer than this survey.

Also notably absent from this survey is discussion of non-revelation mechanisms, i.e., mechanisms that do not have truthtelling or otherwise easy-to-find equilibria. The big challenge of analyzing non-revelation mechanisms is that almost any departure from simple symmetric environments renders solving for equilibrium analytically intractable (cf. Section 2.4). To address this challenge, methodologies from the literature on quantifying the *price of anarchy*, i.e., the suboptimality of performance that results from strategic behavior, can be employed. Instead of explicitly solving for equilibrium, a price-of-anarchy analysis considers minimal necessary properties of equilibria to argue that any equilibrium must be pretty good. Of course, this would be a problematic exercise if we believed that, just as we cannot solve for equilibrium, neither can the agents. Fortunately, because the analysis employs only minimal assumptions, a frequent corollary of a price-of-anarchy analysis is that best-response dynamics and no-regret learning algorithms have good performance even if they never reach an equilibrium.

To illustrate the power of the price-of-anarchy approach, we will describe a few recent results. Consider the welfare objective in multi-agent multi-item environment (which we consider in Section 5 and Section 6 for the revenue objective). Suppose instead of running a simultaneous auction that coordinates the sale of the items, we run independent first- or second-price auctions either simultaneously or sequentially. These auctions are strategically complex, e.g., because there is an exposure problem where an agent may win multiple items even if she only wants one. A series of papers, Bikhchandani [13], Christodoulou et al. [34], Bhawalkar and Roughgarden [12], Hassidim et al. [63], Paes Leme and Tardos [78], Paes Leme et al. [79], Roughgarden [83], and Syrgkanis and Tardos [88, 89] showed that in various configurations the price of anarchy of independent auctions for multiple items is often a constant like two. Paes Leme et al. [77] give a short survey of these results.

Price of anarchy approaches have also been applied to an auction known as the *generalized second-price* auction which is used by Internet search engines to sell advertisements that are displayed alongside search results (see, e.g., Fain and Pedersen [46]). In this auction, bidders are ranked by bid and each bidder is charged the bid of her successor. This auction does not have simple equilibria as does the second-price auction. Gomes and Sweeney [51] show that even in symmetric Bayesian settings the generalized-second-price auction may not have any efficient equilibria; subsequent work of Caragiannis et al. [25] bounded the potential inefficiency of its equilibria by a factor of slightly under three.

A final class of results looks at non-revelation mechanisms based on approximation algorithms (because, as mentioned above, there are no general approaches for converting approximation algorithms into mechanisms without Bayesian assumptions). Lucier and Borodin [71] consider multi-dimensional combinatorial auctions based on greedy approximation algorithms. They show that in equilibrium, a mechanism based on a greedy β approximation algorithm is at worst a $(\beta + 1)$ approximation in equilibrium.

The results described above are primarily for the welfare objective, and methodologies from the price of anarchy have seen relatively less success for the revenue objective. One exception is by Lucier et al. [72] who give revenue bounds for the previously discussed generalized-second-price auction.

2

Equilibrium

The goal of equilibrium analysis is to first understand how players behave in strategic environments and then to make a prediction of what will happen in such an environment.

Consider the *second-price* auction which solicits sealed bids, awards the item to the highest bidder, and charges the winner the second highest bid. Suppose Alice has value for winning the item of v, how should she bid in such an auction? This question is relatively easy to answer. Suppose the highest other bid is $\hat{v}$. Notice that if Alice bids above $\hat{v}$ she wins and pays $\hat{v}$ and if she bids below $\hat{v}$ she loses and pays nothing. Her utilities for these two outcomes are $v - \hat{v}$ and 0, respectively. She prefer winning when $v - \hat{v}$ is non-negative, and losing otherwise. Therefore, if $v \geq \hat{v}$ she should bid any $b \geq \hat{v}$ and if $v \leq \hat{v}$ then she should bid any $b \leq \hat{v}$. Of course, bidding $b = v$ satisfies both conditions and furthermore does not depend on $\hat{v}$. Truthful bidding is a *dominant strategy!* If all bidders follow this natural strategy then we can predict that the winner will be the bidder who had the highest value for the item.

Now consider the *first-price* auction which solicits sealed bids, awards the item to the highest bidder, and charges the winner her bid. How can we predict the outcome of such an auction? How should Alice bid in such an auction? If she knows the highest other bid, call it $\hat{v}$, then with value $v > \hat{v}$ she would want to bid $\hat{v} + \epsilon$. There is no dominant strategy equilibrium as knowledge of the bids of others affects her strategy.

In order to analyze the equilibrium of the first-price auction we need to posit the availability of more information. Suppose the bidders in the auction are Alice and Bob and that their values for the item are drawn uniformly at random from the $[0, 1]$ interval, i.e., with cumulative distribution function $F(z) = z$. Further suppose that Alice believes that Bob will bid half of his value. Notice that the distribution of Bob's value and the assumption on his strategy induces a distribution on his bid, i.e., that it is uniform on $[0, 1/2]$. This is enough information for Alice to calculate her optimal bid b for any value v she might have. Her utility u in the auction is $v - b$ if she wins and zero, otherwise. Therefore,

$$\mathbf{E}[u(v, b)] = (v - b) \times \mathbf{Pr}[\text{Alice wins with bid } b].$$

Calculate

$$\begin{aligned}\mathbf{Pr}[\text{Alice wins with bid } b] &= \mathbf{Pr}[\text{Bob's bid} \leq b] = \mathbf{Pr}\left[\tfrac{1}{2}\text{Bob's value} \leq b\right] \\ &= \mathbf{Pr}[\text{Bob's value} \leq 2b] = F(2b) = 2b.\end{aligned}$$

So,

$$\begin{aligned}\mathbf{E}[u(v, b)] &= (v - b) \times 2b \\ &= 2vb - 2b^2.\end{aligned}$$

To optimize Alice's bid, differentiate the function with respect to b and set its derivative equal to zero. The result is $\frac{d}{db}(2vb - 2b^2) = 2v - 4b = 0$ and we can conclude that Alice's optimal bid is $b = v/2$. Of course, if our assumption was that Bob was bidding half his value and this assumption prompts Alice to bid half her value, then by symmetry, Bob will bid half his value if Alice is so doing. In fact the strategy profile where Alice and Bob both bid half of their values is an equilibrium. From these strategies we can again deduce that in equilibrium the bidder with the highest value will win.

If the goal of the auctioneer is to maximize the social surplus, i.e., to give the item to the bidder with the highest value for the item, then both the first-price and second-price auction achieve that goal. What if the auctioneer's goal was instead to maximize revenue. To discuss this question we can calculate the expected profit of both the first- and second-price auctions in our example scenario where Alice and Bob (i.e., two bidders) have values uniformly distributed on $[0, 1]$. A useful fact about uniform distributions is that in expectation a set of uniform random variables evenly divide the interval they are over. Therefore, for two uniform random variables on $[0, 1]$, the expected minimum is $1/3$ and the expected maximum is $2/3$. The revenue of the second-price auction is, by definition, the second highest bid, which in our truthtelling equilibrium is the second highest value, and is therefore $1/3$ in expectation. The revenue of the first-price auction is by definition the highest bid, which in our bid-half-your-value equilibrium is half the highest value, which is $1/3$ in expectation. Both auctions have the same revenue in expectation!

The "revenue equivalence" exhibited by the example above is not a coincidence, in fact it is a general principle. In this section, we will formalize the notion of equilibrium in games of incomplete information, characterize these equilibria and in particular show revenue equivalence, and finally give a simple method for solving for equilibrium in symmetric auction-like games. Chawla and Hartline [28] give a general revenue-equivalence-based method for showing that this symmetric equilibrium is the unique Bayes-Nash equilibrium; however, we will not cover it here.

2.1 Equilibrium

We will assume that each agent has some private information and this information affects the payoff of this agent in the game. We will refer to this information as the agent's *type* and denote it by t_i for agent i. The profile of types for the n agents in the game is $\boldsymbol{t} = (t_1, \ldots, t_n)$.

A *strategy* in a game of incomplete information is a function that maps an agent's type to any of the agent's possible actions in the game (or a distribution over actions for mixed strategies). We will denote by $s_i(\cdot)$ the strategy of agent i and by $\boldsymbol{s} = (s_1, \ldots, s_n)$ a *strategy profile.*

The first- and second-price auctions described above are games of incomplete information where an agent's private type is her value for receiving the item, i.e., $t_i = v_i$. We saw that the second-price auction has a dominant strategy equilibrium of truthtelling.

Definition 2.1. A *dominant strategy equilibrium* (DSE) is a strategy profile $\boldsymbol{s}$ such that for all i, t_i, and $\boldsymbol{b}_{-i}$ (where $\boldsymbol{b}_{-i}$ generically refers to the actions of all players but i), agent i's utility is maximized by following strategy $s_i(t_i)$.

Naturally, many games of incomplete information do not have dominant strategy equilibria. Of course, equilibria in these games are not well defined without an additional assumption. We therefore, make the assumption that the agents' types are drawn from a known distribution.

Definition 2.2. Under the *common prior assumption*, the agent types $\boldsymbol{t}$ are drawn at random from a *prior distribution* $\boldsymbol{F}$ (a joint probability distribution over type profiles) and this prior distribution is *common knowledge*.

The distribution $\boldsymbol{F}$ over $\boldsymbol{t}$ may generally be correlated. Which means that an agent with knowledge of her own type must do *Bayesian updating* to determine the distribution over the types of the remaining agents. We denote this conditional distribution as $\boldsymbol{F}_{-i}\big|_{t_i}$. Of course, when the distribution of types is independent, i.e., $\boldsymbol{F}$ is the *product distribution* $F_1 \times \cdots \times F_n$, then $\boldsymbol{F}_{-i}\big|_{t_i} = \boldsymbol{F}_{-i}$.

A strategy profile is a mapping for each agent of her private type to an action in the game. Notice that a prior $\boldsymbol{F}$ and strategies $\boldsymbol{s}$ induce a distribution over the actions of each of the agents. With such a distribution over actions, the problem each agent faces of optimizing her own action is fully specified.

Definition 2.3. A *Bayes-Nash equilibrium (BNE)* for a game and common prior $\boldsymbol{F}$ is a strategy profile $\boldsymbol{s}$ such that for all i and t_i, $s_i(t_i)$ is a best response when other agents play $\boldsymbol{s}_{-i}(\boldsymbol{t}_{-i})$ for $\boldsymbol{t}_{-i} \sim \boldsymbol{F}_{-i}\big|_{t_i}$.

In Bayesian games it is useful to distinguish between stages of the game in terms of the knowledge sets of the agents. The three stages of a Bayesian game are *ex ante*, *interim*, and *ex post*. The ex ante stage is before values are drawn from the distribution. Ex ante, the agents know this distribution but not their own types. The interim stage is immediately after the agents learn their types, but before playing in the game. In the interim, an agent knows her own type and assumes the other agent types are drawn from the prior distribution conditioned on her own type, i.e., via Bayesian updating. In the ex post stage, the game is played and the actions of all agents are known.

Reasoning about equilibria in Bayesian games can be challenging. For mechanism design, however, the *revelation principle* often implies that it is without loss of generality to restrict attention to mechanisms with truthtelling as an equilibrium. The principle states that if there is a game with a good equilibrium, then there is a mechanism where truthtelling is a good equilibrium. The construction is simple: given a game with a good equilibrium and the profile of equilibrium strategies, consider the *revelation mechanism* that (a) assumes its inputs are the true types of the agents, (b) simulates the strategies of the agents with these types in the original game, and (c) outputs the outcome of this simulation. Clearly, the new mechanism has a truthtelling equilibrium and its outcome in this equilibrium is good.

Definition 2.4. A mechanism is *Bayesian (resp. dominant strategy) incentive compatible* if truthtelling is a Bayes-Nash (resp. dominant strategy) equilibrium.

Proposition 2.1 (Revelation Principle, Myerson [74]). For any mechanism and Bayes-Nash (resp. dominant strategy) equilibrium of the mechanism, there exists a Bayesian (resp. dominant strategy) incentive compatible mechanism with the same equilibrium.

The revelation principle fails to hold in some environments of interest. Two such environments, for instance, are where agents only learn their values over time, or where the mechanism designer does not know the prior distribution (and hence cannot simulate the agent strategies).

2.2 Independent, Single-dimensional, and Linear Utilities

In the classical economic environments for mechanism design it is assumed that agents' types are independent and single dimensional, and that their utilities are linear. Such a single-dimensional type represents an agent's value for an abstract service and her utility is the difference in her value for service and the price she must pay for service. The outcome of a mechanism in such an environment can be specified by allocation and payment vectors $\boldsymbol{x} = (x_1, \ldots, x_n)$ and $\boldsymbol{p} = (p_1, \ldots, p_n)$ where x_i is an indicator for whether agent i is served or not and p_i is agent i's required payment. Denoting agent i's private value for service by v_i, her utility for such an outcome is thus $u_i = v_i x_i - p_i$.

A mechanism and a profile of equilibrium strategies induce *ex post allocation* and *payment rules* that map the type profile $\boldsymbol{v} = (v_1, \ldots, v_n)$ to an allocation $\boldsymbol{x}(\boldsymbol{v})$ and payments $\boldsymbol{p}(\boldsymbol{v})$. Importantly, agents in a game act in the interim stage, i.e., when their own type is known, but the types of the other agents are random from the prior distribution. Because the agents' utilities are linear and their types are independent, we can summarize the relevant *interim allocation* and *payment rules* as $x_i(z) = \mathbf{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}}[x_i(z, \boldsymbol{v}_{-i})]$ and $p_i(z) = \mathbf{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}}[p_i(z, \boldsymbol{v}_{-i})]$. Here $x_i(v_i)$ is the probability that agent i is served when her value is v_i; $p_i(v_i)$ is her expected payment.

Bayes-Nash equilibrium and dominant strategy equilibrium respectively require; for all i, $\boldsymbol{v}$, and z; that agent i has higher utility for her strategy at v_i than following her strategy for value z, i.e.,

$$v_i \cdot x_i(\boldsymbol{v}) - p_i(\boldsymbol{v}) \geq v_i \cdot x_i(z, \boldsymbol{v}_{-i}) - p_i(z, \boldsymbol{v}_{-i}) \tag{DSE}$$

$$v_i \cdot x_i(v_i) - p_i(v_i) \geq v_i \cdot x_i(z) - p_i(z). \tag{BNE}$$

Importantly, (BNE) assumes the types of the agents are independent. A strategy profile is *onto* if for each agent i and each possible action b_i of i in the mechanism, there is a value for which $s_i(v_i) = b_i$. Note that the truthtelling equilibrium of the revelation mechanism where the space of actions is equal to the space of types is onto. For strategy profiles that are onto, the converse also holds: (DSE) and (BNE) imply dominant strategy and Bayes-Nash equilibrium, respectively. The truthtelling equilibrium of an incentive compatible mechanism is onto; therefore, for incentive compatible mechanisms equations (DSE) and (BNE) are "if and only if."

In the examples above, we have discussed the dominant strategy equilibrium of the second-price auction and the Bayes-Nash equilibrium of the first-price auction. Both auctions have ex post allocation rule (in equilibrium):

$$x_i(\boldsymbol{v}) = \begin{cases} 1 & \text{if } \forall j,\ v_i \geq v_j \\ 0 & \text{otherwise.} \end{cases}$$

Moreover, for two players with uniformly distributed values the allocation rule of each player is $v_i(v) = F(v) = v$ (the probability that v exceeds the other player's value). The interim payment rule can be calculated as $p_i(v) = v^2/2$. (In the first-price auction an agent with value v wins with probability v and pays her bid which is $v/2$; in the second-price auction an agent with value v wins with probability v and pays the expected value of the other agent conditioned on being at most v which is $v/2$.)

Both the first- and second-price auction with bidder values drawn uniformly from $[0, 1]$ are examples of independent, single-dimensional, and linear utilities. In these auction problems, the auctioneer also has a feasibility constraint that at most one agent can win the item, i.e., $\sum_i x_i \leq 1$. Notice that the auctioneer's feasibility constraint does not directly play a role in the incentive constraints of the agents.

2.3 Equilibrium Characterization

The characterization of Bayes-Nash equilibria is given below. Analogous characterizations hold for dominant strategy equilibria where $x_i(z)$ and $p_i(z)$ are replaced by $x_i(z, \boldsymbol{v}_{-i})$ and $p_i(z, \boldsymbol{v}_{-i})$. See Figure 2.1 for a graphical depiction of the theorem.

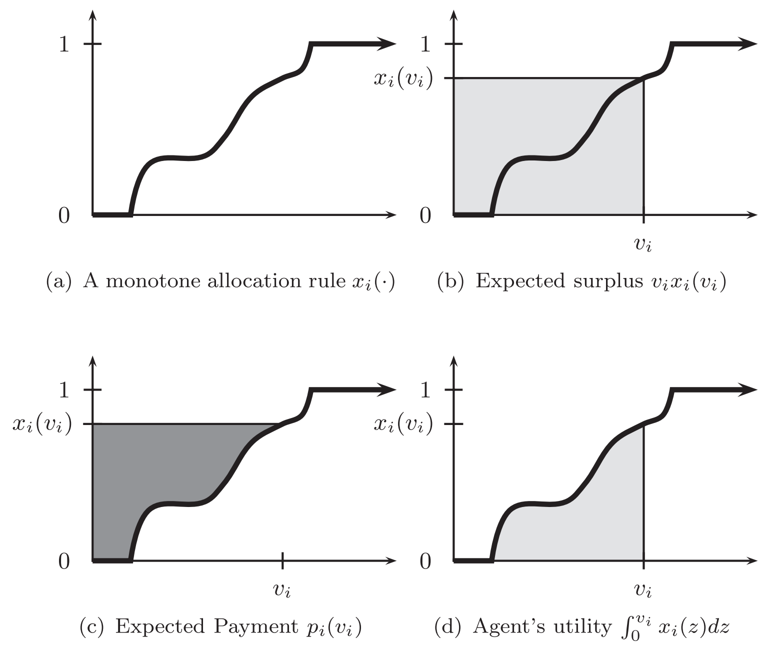

Fig. 2.1 A depiction of the statement of Theorem 2.2. In equilibrium, the auction generates expected surplus $v_i x_i(v_i)$ when agent i has value v_i. This surplus is split between the agent (her utility) and the auctioneer (the payment) as the portion of the surplus that is below and above the allocation rule, respectively.

Theorem 2.2 (Myerson [74]). For agents with independent, single-dimensional, and linear utility, a profile of allocation and payment rules is in Bayes-Nash equilibrium only if for all i,

(a) (monotonicity) $x_i(v_i)$ is monotone non-decreasing, and
(b) (payment identity) $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\, dz + p_i(0)$,

where often $p_i(0) = 0$. If the strategy profile is onto then the converse also holds.

Revenue equivalence is an immediate corollary of the payment identity in the BNE characterization. If two mechanisms have the same allocation in equilibrium, then their expected revenues must be the same.

2.4 Solving for Equilibrium

The payment identity (a.k.a., revenue equivalence) gives a simple way to calculate the Bayes-Nash equilibrium in symmetric auctions. The bid-half-your-value equilibrium in the first-price auction with two players and uniformly distributed values can be derived from this method.

Suppose we are to solve for the BNE strategies of mechanism $\mathcal{M}$. The approach is to obtain and equate two expressions for an agent's payment as a function of the agent's action. The first expression comes from the definition of mechanism and is in terms of the agent's action; the second expression comes from strategically-simple mechanism $\mathcal{M}^{\dagger}$ that is revenue equivalent to $\mathcal{M}$ (usually a "second-price" implementation of $\mathcal{M}$). Setting these terms equal and solving for the agent's action gives the equilibrium strategy.

Notice that in the first- and second-price auctions an agent only pays when she wins. Revenue equivalence then tells us that this expected payment upon winning must be the same in both auctions. We can solve for and equate these two expected payments leaving an agent's strategy in the first-price auction as a variable. From the resulting equation we can solve for the equilibrium. Finally, it is important to double check that the resulting strategies are indeed an equilibrium.

Consider two agents with values drawn from $U[0,1]$. Conditioned on winning with value v_1, agent 1's expected payment in the second-price auction is,

$$\mathbf{E}_{v_2}\big[p_1^{\mathrm{SP}}(\boldsymbol{v}) \mid 1 \text{ wins}\big] = \mathbf{E}_{v_2}[v_2 \mid v_2 < v_1] = v_1/2, \tag{2.1}$$

because of the dominant strategy truthtelling equilibrium and because conditioned on agent 1 winning with value v_1, agent 2's value is $U[0, v_1]$. Of course, in a first-price auction agent 1 pays her bid when she wins,

$$\mathbf{E}_{v_2}\big[p_1^{\mathrm{FP}}(\boldsymbol{v}) \mid 1 \text{ wins}\big] = s_1(v_1). \tag{2.2}$$

Equating (2.1) and (2.2) gives,

$$s_1(v_1) = v_1/2.$$

Indeed, this is an equilibrium. First, if agents follow the suggested symmetric strategies $s_1(z) = s_2(z) = z/2$ then the agent with the highest value wins. Second, bidding over $s_i(1)$ is dominated by bidding $s_i(1)$ as both bids always win, but the later gives a lower payment. This second check is important as the suggested strategies do not map the set of values onto the set of actions (cf. Theorem 2.2).

3

Optimal Mechanisms

Consider the example of selling a single-item to one of two potential buyers each with value drawn uniformly from $[0,1]$. We saw in Section 2 that both the first-price auction and the second-price auction maximize social surplus, i.e., they award the item to the buyer who has the highest value for it. Similarly, both these auctions have the same expected revenue of $1/3$. Is this revenue optimal, or is there another auction that obtains a higher revenue?

Consider the second-price auction with reserve price $1/2$. This auction awards the item to the highest bidder as long as her bid is at least the reserve price. It charges the winner the maximum of the second highest bid and the reserve price. It is relatively easy to calculate the revenue of the second-price auction with reserve $1/2$ for our two buyers with values uniform on $[0,1]$; its revenue is $5/12$.[1] As $5/12 > 1/3$ this results in the perhaps surprising conclusion that a seller can make more

[1] The calculation proceeds as follows with $v_{(i)}$ denoting the ith highest value: There are three cases (i) $1/2 > v_{(1)} > v_{(2)}$, (ii) $v_{(1)} > 1/2 > v_{(2)}$, and (iii), $v_{(1)} > v_{(2)} > 1/2$. Case (i) happens with probability $1/4$ and has no revenue; case (ii) happens with probability $1/2$ and has revenue $1/2$; and case (iii) happens with probability $1/4$ and has expected revenue $\mathbf{E}\big[v_{(2)} \mid \text{case (iii) occurs}\big] = 2/3$. The calculation of the expected revenue in case (iii) follows from the conditional values being $U[1/2,1]$ and the fact that, in expectation, uniform random variables evenly divide the interval they are over. The total expected revenue can then be calculated as $5/12$.

money by sometimes not selling the item even when there is a buyer willing to pay.

In this section we describe a simple method for optimizing over Bayes-Nash equilibria to derive the Bayesian optimal mechanism. This optimization will allow us to conclude that indeed the second-price auction with reserve $1/2$ is optimal for $U[0,1]$ agents.

3.1 Single-dimensional Environments

We have focused thus far on describing the preferences of agents with independent, single-dimensional, and linear utilities. We now consider the preferences and constraints of the mechanism designer. The outcome of a mechanism is denoted by an allocation $\boldsymbol{x} = (x_1, \ldots, x_n)$ and prices $\boldsymbol{p} = (p_1, \ldots, p_n)$. Typically, the designer will have a constraint $\mathcal{X}$ specifying which allocations are feasible. For instance, in the single-item auctions discussed previously, the designer had a constraint $\sum_i x_i \leq 1$.

The designer's constraint is *downward closed* if any subset of a feasible set is feasible. I.e., if $\boldsymbol{x}$ is feasible then $\boldsymbol{x}_{-i} = (x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n)$ is feasible. Downward closure is standard for allocation problems where there are items to be allocated and it is possible to leave some items unallocated.

Interesting classes of mechanism design problems can be described within this framework. Suppose there is a set of n agents and a set of m items. Suppose for each agent i there is a set of items S_i that the agent wants. An agent i with type v_i has "or" preferences if v_i represents her value for any one of the items $j \in S_i$ and she desires at most one; she has "and" preferences if her value is v_i only if she receives the entire bundle S_i of items. The case where all agents have or-preferences is essentially a bipartite matching problem; we will alternatively refer to such scenarios as *(single-dimensional) matching environments*. The case where all agents have and-preferences is essentially the set packing problem; the corresponding auction problem is the *single-minded combinatorial auction*; and we will refer to these environments as such.

Other environments we will discuss are

single-item environments: The feasibility constraint permits at most one agent to be served.

multi-unit environments: The feasibility constraint permits at most a specific number (usually denoted k) of agents to be served.

matroid environments: The feasible sets of agents correspond to the independent sets of a *matroid* set system. Importantly, matroids correspond exactly to the set systems for which the greedy-by-weight algorithm solves weighted optimization subject to feasibility, i.e., $\max_{\boldsymbol{x} \in \mathcal{X}} \sum_i v_i x_i$ for weights $\boldsymbol{v} = (v_1, \ldots, v_n)$. Single-dimensional matching environments will be our motivating example of matroid environments; matroid environments also include multi-unit and single-item environments as special cases. See Section A.2 for further discussion.

downward-closed environments: The feasible sets of agents are constrained only so that subsets of feasible sets are feasible. Single-minded combinatorial environments will be our motivating example of a downward-closed environment. Downward closed environments include everything discussed above; an example of a non-downward-closed environment is that of a *public project* which is given by the *all-or-nothing* set system.

There are two classical economic objectives for mechanism design, *social surplus* and *revenue*. The social surplus of an outcome is the cumulative utility of all participants. For value $\boldsymbol{v}$ and outcome $(\boldsymbol{x}, \boldsymbol{p})$ the social surplus is $\sum_i v_i x_i$; importantly, the payments arise positively in the auctioneer's utility and negatively in the agent's utility and cancel from the objective. The revenue of an outcome $(\boldsymbol{x}, \boldsymbol{p})$ is $\sum_i p_i$. The theory of optimal mechanisms is quite different for these two objectives.

Notice that surplus maximization gives a monotone allocation rule, i.e., if an agent's value is increased, her allocation is not decreased. Therefore, the dominant-strategy version of Theorem 2.2 implies that there is an auction with this outcome as a dominant strategy equilibrium. Furthermore, the payments of this auction are given by the payment identity. Denoting $\mathrm{OPT}(\boldsymbol{v}) = \max_{\boldsymbol{x} \in \mathcal{X}} \sum_i v_i x_i$, the payment of an agent i is:

$$p_i = \begin{cases} \mathrm{OPT}(\boldsymbol{v}_{-i}) - \mathrm{OPT}(\boldsymbol{v}) + v_i & \text{if } x_i = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Notice that when $x_i = 1$, $\mathrm{OPT}(\boldsymbol{v}) - v_i$ can be easily calculated without knowing v_i as its value is just the optimal surplus from agents other than i subject to the constraint that i is served. This *surplus maximizing mechanism* is often referred to as the Vickrey-Clarke-Groves (VCG) mechanism [91, 35, 53]; it is the natural generalization of the second-price auction, previously discussed, to more complex environments.

Proposition 3.1 (Vickrey [91], Clarke [35], Groves [53]). The surplus maximizing mechanism is dominant-strategy incentive compatible and it optimizes social surplus pointwise.

The revenue objective, on the other hand, is more challenging to optimize. Notice that the surplus of a mechanism is given by the allocation on the given valuation profile. The revenue, on the other hand, is given by the sum of the agent payments which, via the payment identity of Theorem 2.2, depend on the allocation rule of each agent (in particular, on $x_i(z)$ for $z \leq v_i$ for agent i). This allocation rule is derived from the specification of what the mechanism does on all valuation profiles. Changes to the allocation rule, therefore, can tradeoff revenue on one valuation profile for revenue on another. This tradeoff can be optimized in expectation for a valuation profile $\boldsymbol{v}$ drawn from distribution $\boldsymbol{F}$. Importantly, though, there is no mechanism that is revenue-optimal pointwise.

To illustrate this absence of a pointwise optimal mechanism, consider a single agent with value v drawn uniformly from $[0,1]$. If her value is 0.2, then it is pointwise optimal to offer her the item at price 0.2. This corresponds to the allocation rule which steps from 0 to 1 at 0.2. Similarly if her value is 0.7, then it is pointwise optimal to offer her the item at price 0.7. Of course, offering a 0.7-valued agent a price of 0.2 or a 0.2-valued agent a price of 0.7 is not optimal. In contrast, given a distribution over the agent's value, we can easily optimize for the mechanism with maximum expected revenue: post the price p that maximizes $p \cdot (1 - F(p))$. For the uniform distribution where $F(p) = p$, this optimal price is $p = 1/2$.

We conclude that social surplus is special among objectives in that there is a pointwise optimal mechanism; in general, a mechanism's performance should be optimized in expectation over a distribution.

3.2 Single-agent Revenue Maximization

Before looking at the problem of optimizing revenue in an multi-agent mechanism, we will look at a few fundamental single-agent optimization problems. A mechanism for a single agent is simply a menu of outcomes where, after the agent realizes her type from the distribution, she chooses the outcome she most prefers. This observation is known as the *taxation principle* (see Wilson [93]). The outcomes on the menu are allowed to be probabilistic, e.g., "with probability 0.5 you are served and you pay 0.25." A single-agent problem is then to find the menu of outcomes that maximizes revenue subject to any imposed constraints.

We will first look at single-agent problems where the constraint is on the ex ante service probability, i.e., the probability that an agent is served given the randomization in her type and her chosen outcome. From these we can define a "revenue curve" as the optimal revenue as a function of the ex ante service probability. We will then look at a more general class of constraints, describe a "revenue-linearity" condition, and show that revenue linearity implies that optimal single-agent mechanisms for these general constraints can be decomposed into optimal single-agent mechanisms for the simpler ex ante constraints. We will see that this decomposition implies that the optimal revenue for general constraints is equal to the cumulative "marginal revenue" (given by the derivative of the revenue curve). Finally, we show that a single-dimensional agent is in fact revenue linear.

The revenue-linearity based approach described here, which is due to Alaei et al. [5], is more involved than the original derivation of the optimal auction by Myerson [74]. In contrast, Myerson starts with the payment identity of Theorem 2.2, takes expectation with respect to the distribution, and simplifies to derive "virtual values" which Bulow and Roberts [20] observe are equivalent to the aforementioned marginal revenue (see the survey of Hartline and Karlin [58]). The advantage of the revenue-linearity based approach, in addition to exposing important structure of optimal mechanisms, is that it is not reliant on the specific form of the payment identity and is therefore generalizable beyond single-dimensional agents (see Section 5).

3.2.1 Inverse Demand Curve

Consider an agent Alice with a single-dimensional linear preference as described in Section 2.2. Alice's preference is described by her value v which is drawn from distribution F. The economic geometry of optimal mechanisms is more apparent when the preferences are indexed by strength relative to the distribution (a.k.a., quantile) instead of value (see Figure 3.1). The definition below formalizes the correspondence between the value and distribution representation and the quantile and inverse demand curve representation for any single-dimensional agent.

Definition 3.1. The *quantile* of a single-dimensional agent with value $v \sim F$ is the measure with respect to F of stronger values, i.e., $q = 1 - F(v)$; the *inverse demand curve* maps an agent's quantile to her value, i.e., $V(q) = F^{-1}(1 - q)$.

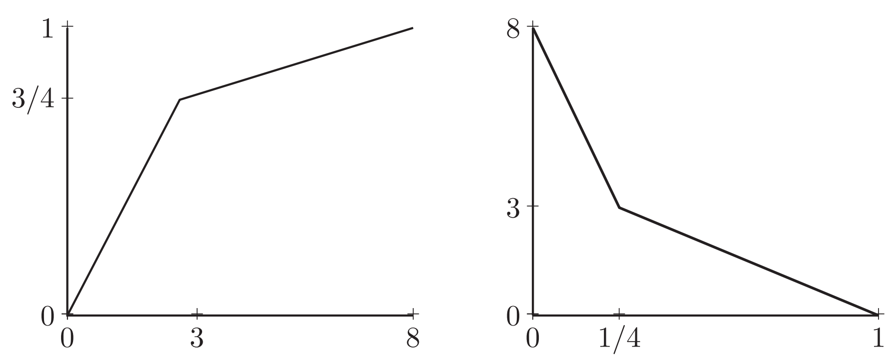

(a) Cumulative distribution function $F(v)$. (b) Inverse demand curve $V(q)$.

Fig. 3.1 Depicted are the cumulative distribution function and inverse demand curve corresponding to a single-dimensional linear agent with value v drawn from distribution with density function $f(v) = 1/4$ for $v \in [0, 3]$ and $f(v) = 1/20$ for $v \in (3, 8]$, i.e., the distribution that is uniform on $[0, 3]$ with probability $3/4$ and uniform on $(3, 8]$ with probability $1/4$. The inverse demand curve is obtained from the cumulative distribution function by rotating it 90 degrees counterclockwise.

3.2.2 Revenue Curves

Given Alice's inverse demand curve $V(\cdot)$ and a target ex ante probability of service $\hat{q}$ we can ask what is the revenue optimal mechanism (i.e., menu) that serves her with the given probability. As we vary $\hat{q}$ we get revenue as a function of probability of sale.

Definition 3.2. The *revenue curve* $R(\hat{q})$ is the optimal revenue from an agent (with implicit distribution F) as a function of her ex ante probability of service $\hat{q}$. The mechanism for ex ante sale probability $\hat{q}$ that has revenue $R(\hat{q})$ is the $\hat{q}$ *ex ante optimal mechanism*.

One way to sell to Alice with ex ante probability $\hat{q}$ is to post a price that she accepts with probability $\hat{q}$, i.e., the price $V(\hat{q})$. The revenue from this *price-posting* approach is $P(\hat{q}) = \hat{q} \cdot V(\hat{q})$. This approach gives a lower bound on the revenue curve, i.e., $R(\hat{q}) \geq P(\hat{q})$. In fact, it is easy to see what price posting is missing in terms of revenue. If $P(\hat{q})$ is not concave then we can get a higher revenue than $P(\hat{q})$ by taking the convex combination of two mechanisms so that the final sale probability is $\hat{q}$. In fact, $R(\hat{q})$ is exactly given by the smallest concave function that upper bounds $P(\hat{q})$. See Figure 3.2(a).

Proposition 3.2. $R(\hat{q})$ is the smallest concave function that upper bounds $P(\hat{q})$.

3.2.3 Allocation Rules and Revenue

In Section 2.2 we defined the allocation rule for an agent as a function of her value as $x(\cdot)$. In Section 2.3 we characterized the allocation rules that can arise in Bayes-Nash equilibrium as the class of monotone non-decreasing functions (of value). Consider the allocation rule instead in quantile space, i.e., $y(q) = x(V(q))$. Since quantile and value are indexed in the opposite direction, $y(\cdot)$ will be monotone non-increasing in quantile.

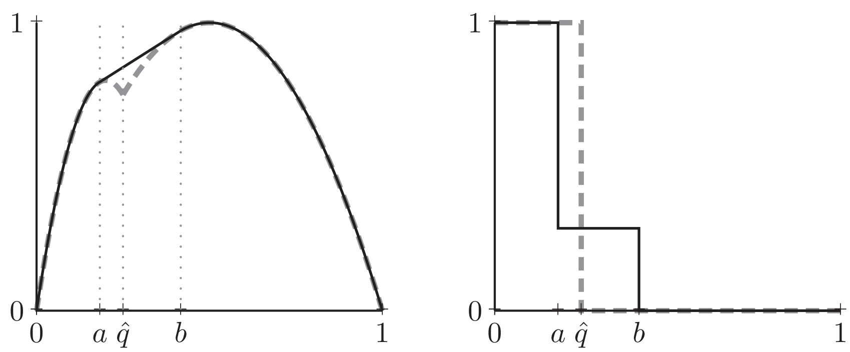

(a) Revenue curves $R(\cdot)$ and $P(\cdot)$.

(b) Allocation rules $\hat{y}(\cdot)$ and $y(\cdot)$.

Fig. 3.2 Depicted are the revenue curve, price-posting revenue curve, and their allocation rules corresponding to ex ante allocation constraint $\hat{q}$ for a single-dimensional linear agent with value v drawn from distribution with density function $f(v) = 1/4$ for $v \in [0,3]$ and $f(v) = 1/20$ for $v \in (3,8]$, i.e., the distribution that is uniform on $[0,3]$ with probability 3/4 and uniform on $(3,8]$ with probability 1/4. For such a distribution the revenue curve $R(\cdot)$ (thin, black, solid line) is obtained from the price-posting revenue curve $P(\cdot)$ (thick, grey, striped line) by averaging on interval $[a,b]$. The allocation rule for posting price $V(\hat{q})$ is the reverse step function at $\hat{q}$ (thick, grey, striped line). For $\hat{q} \in [a,b]$ as depicted, the allocation rule (thin, black, solid line) for the $\hat{q}$ ex ante optimal mechanism is the appropriate convex combinations of the reverse step functions at a and b. The area under both allocation rules is equal to the ex ante service probability $\hat{q}$.

Notice that the allocation rule of the mechanism that posts price $V(\hat{q})$ is simply the reverse step function that starts at 1 and steps from 1 to 0 at $\hat{q}$. The allocation rule of the $\hat{q}$ ex ante optimal mechanism, when it is given by a convex combination of two price postings, is given by the same convex combination of two reverse step functions. Figure 3.2(b) depicts these allocation rules.

Notice that the allocation rule from posting price $V(\hat{q})$ is “strong” in the sense that it allocates with certainty to the highest $\hat{q}$ measure of Alice’s values. The ex ante optimal mechanism, optimizing subject to the same constraint, can choose to use a “weaker” allocation rule. In Figure 3.2(b) the highest $a < \hat{q}$ measure of values is served with certainty, while the remaining $\hat{q} - a$ service probability is spread out among the next $b - \hat{q}$ measure of values. The following proposition is immediate and, though it gives a coarser picture than Proposition 3.2, it will be sufficient for deriving the main results of this section.

Proposition 3.3. For any ex ante service probability, the optimal mechanism has a weaker allocation rule and higher revenue than price posting.

An allocation rule y is a monotone non-increasing function from $[0,1]$ to $[0,1]$. As reverse step functions form a basis for such functions and reverse step functions are the allocation rule for price postings, we can construct a single-agent mechanism with allocation rule y by taking the appropriate convex combination of price postings. Consider the distribution $G^y(z) = 1 - y(z)$, draw $\hat{q} \sim G^y$, and post price $V(\hat{q})$. The resulting mechanism has allocation rule exactly $y(\cdot)$ and its expected revenue is equal to the convex combination of revenues $P(\hat{q})$ from posting price $V(\hat{q})$. This revenue is as follows, via G^y's density function $-y'(q) = -\frac{d}{dq}y(q)$ at q, integration by parts, and the assumption that $P(0) = P(1) = 0$ (there is no revenue from always selling or never selling).

$$\begin{aligned} \mathbf{E}_{\hat{q}\sim G^y}[P(\hat{q})] &= \mathbf{E}_q\left[-y'(q) \cdot P(q)\right] \\ &= \mathbf{E}_q\left[P'(q) \cdot y(q)\right], \end{aligned}$$

where $P'(q) = \frac{d}{dq}P(q)$ is the marginal increase in price-posting revenue for an increase in ex ante allocation probability, a.k.a., the *marginal price-posting revenue* at q. By revenue equivalence (Theorem 2.2), any mechanism with the same allocation rule has the same revenue.

Proposition 3.4. A single-agent mechanism with allocation rule y has expected revenue equal to the cumulative marginal price-posting revenue $\mathbf{E}_q[P'(q) \cdot y(q)]$.

3.2.4 Optimal and Cumulative Marginal Revenue

We now show that optimal single-agent mechanisms can be decomposed into convex combinations of ex ante optimal mechanisms (with various service probabilities). To do so, we consider a more general single-agent mechanism design problem. The ex ante probability that allocation rule $y(\cdot)$ allocates to the strongest $\hat{q}$ measure of quantiles is $Y(\hat{q}) = \int_0^{\hat{q}} y(q)\,dq$; we refer to $Y(\cdot)$ as *cumulative allocation rule* for $y(\cdot)$. The monotonicity of allocation rules implies that cumulative allocation rules are concave. As follows, we can view an allocation rule $\hat{y}(\cdot)$ as a constraint via its cumulative allocation rule. If $y(\cdot)$ satisfies constraint $\hat{y}(\cdot)$, relatively, it has service probability shifted from stronger quantiles to weaker quantiles.

Definition 3.3. Given an allocation constraint $\hat{y}$ with cumulative constraint $\hat{Y}$, the allocation rule y with cumulative allocation rule Y is *weaker* (resp. $\hat{y}$ is *stronger*) if and only if it satisfies $Y(\hat{q}) \leq \hat{Y}(\hat{q})$ for all $\hat{q}$; denote this relationship by $y \preceq \hat{y}$.

In comparison to the $\hat{q}$ ex ante optimal mechanism which is constrained so as to serve all quantiles with total probability at most $\hat{q}$, the allocation constraint $\hat{y}$ gives a more fine-grained control by constraining the ex ante service probability of any $\hat{q}$ measure of quantiles by $\hat{Y}(\hat{q})$. Clearly, $\hat{y}$ is the strongest allocation rule for constraint $\hat{y}$. From this notion of strength we can take an allocation rule as a constraint and consider the optimization question of finding a weaker allocation rule with the highest possible revenue.

Definition 3.4. The optimal revenue subject to allocation constraint $\hat{y}(\cdot)$ is $\mathbf{Rev}[\hat{y}]$, the mechanism that attains it is the $\hat{y}$ *optimal mechanism.*[2]

One approach to find a mechanism that satisfies a given allocation constraint $\hat{y}$ would be to take the convex combination of ex ante optimal mechanisms for $\hat{q}$ drawn from $G^{\hat{y}}$ (cf. the construction of Proposition 3.4). Using the ex ante optimal mechanisms improves on price postings in that for each $\hat{q}$ the optimal mechanism's revenue $R(\hat{q})$ may exceed the price-posting revenue $P(\hat{q})$. Moreover, it produces an

[2] Subsequently, when $\hat{y}$ comes from the interim mechanism for a given agent (i.e., the mechanism that internalizes the other agents' types as random draws from the distribution) we refer to this mechanism as the $\hat{y}$ *interim optimal mechanism*.

allocation rule that is weaker than $\hat{y}$ (e.g., see Figure 3.3(d)) as each ex ante optimal mechanism may have a weaker allocation rule than the reverse step function at $\hat{q}$ (e.g., see Figure 3.2(b)). The optimal revenue for allocation constraint $\hat{y}$ can be no less than the revenue of this convex combination. Therefore, as above, we have:

$$\begin{aligned}\mathbf{Rev}[\hat{y}] &\geq \mathbf{E}_{\hat{q}\sim G^{\hat{y}}}[R(\hat{q})] \\ &= \mathbf{E}_q\big[-\hat{y}'(q)\cdot R(q)\big] \\ &= \mathbf{E}_q\big[R'(q)\cdot \hat{y}(q)\big],\end{aligned}$$

where $R'(q) = \frac{d}{dq}R(q)$ is the *marginal revenue* at q.

Definition 3.5. The *cumulative marginal revenue* of an allocation constraint $\hat{y}$ is $\mathbf{MargRev}[\hat{y}] = \mathbf{E}_q[R'(q)\cdot\hat{y}(q)]$.

3.2.5 Revenue Linearity

The above derivation says the cumulative marginal revenue of an allocation constraint is a lower bound on its optimal revenue. A central dichotomy in optimal mechanism design is given by the partitioning of single-agent problems (given by the agent's type space, distribution, and utility function) into those for which this inequality is tight and those when it is not. Notice that linearity of the revenue operator $\mathbf{Rev}[\cdot]$ implies, by the above derivation, that the optimal revenue and cumulative marginal revenue are equal.

Definition 3.6. A single-agent problem is *revenue linear* if $\mathbf{Rev}[\cdot]$ is linear, i.e., if when $\hat{y} = \hat{y}^\dagger + \hat{y}^\ddagger$ then $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger]$.

Proposition 3.5. For a revenue-linear single-agent problem and any allocation constraint, the optimal revenue is equal to the cumulative marginal revenue, i.e., for all $\hat{y}$, $\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}]$.

Revenue linearity of single-dimensional linear agents stems from two main ingredients: Proposition 3.3 (price posting gives less revenue with a stronger allocation rule than the ex ante optimal mechanism) and Proposition 3.4 (essentially, revenue equivalence). Optimal revenue equaling cumulative marginal revenue for single-dimensional agents is an immediate corollary.

Theorem 3.6 (Alaei et al. [5]). An agent with single-dimensional linear utility is revenue linear.

Proof. Before we begin, notice that for a revenue curve $R(\cdot)$ and allocation rule $y(\cdot)$ the cumulative marginal revenue $\mathbf{MargRev}[y]$ can be equivalently expressed as

$$\mathbf{E}_q\big[-R(q)y'(q)\big] = \mathbf{E}_q\big[R'(q)y(q)\big] = \mathbf{E}_q\big[-R''(q)Y(q)\big] + R'(1)Y(1)$$

via integration by parts (with $R(1) = R(0) = Y(0) = 0$). Two observations:

(a) The first term shows that a pointwise higher revenue curve gives no lower a revenue (as $-y'(\cdot)$ is non-negative). In particular, the cumulative marginal revenue exceeds the cumulative price-posting revenue as $R(q) \geq P(q)$ for all q (by Proposition 3.3).

(b) The last term shows that for concave revenue curves, i.e., where $-R''(\cdot)$ is non-negative, e.g., $R(\cdot)$ not $P(\cdot)$; a stronger allocation rule gives higher revenue. In particular, the allocation rule y obtained by optimizing for $\hat{y}$ has no higher cumulative marginal revenue than does $\hat{y}$.

We have already concluded that the cumulative marginal revenue lower bounds the optimal revenue; to prove the theorem it then suffices to upper bound the optimal revenue by the cumulative marginal revenue. Suppose we optimize for $\hat{y}$ and get some weaker allocation rule y, then y is a fixed point of $\mathbf{Rev}[\cdot]$ (optimizing for allocation constraint y gives back allocation rule y); therefore,

$$\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[y].$$

By revenue equivalence (Proposition 3.4), the revenue of any allocation rule is equal to its cumulative marginal price-posting revenue, so

$$\mathbf{Rev}[y] = \mathbf{E}_q\big[P'(q) \cdot y(q)\big]\,.$$

By observation (a), for y the cumulative marginal revenue is at least the cumulative marginal price-posting revenue,

$$\mathbf{E}_q\big[-y'(q) \cdot P(q)\big] \leq \mathbf{E}\big[-y'(q) \cdot R(q)\big]\,.$$

By observation (b), the cumulative marginal revenue for $\hat{y}$ is at least that of y,

$$\mathbf{E}\big[-R''(q) \cdot Y(q)\big] \leq \mathbf{E}[-R''(q) \cdot \hat{Y}(q)]\,.$$

This sequence of inequalities implies that the cumulative marginal revenue is at least the optimal revenue for $\hat{y}$,

$$\mathbf{Rev}[\hat{y}] \leq \mathbf{MargRev}[\hat{y}]. \qquad \square$$

Corollary 3.7 (Myerson [74], Bulow and Roberts [20]). For an agent with single-dimensional, linear utility, the optimal revenue equals the marginal revenue, i.e.,

$$\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}] = \mathbf{E}\big[R'(q)\hat{y}(q)\big]\,.$$

3.3 Multi-agent Revenue Maximization

We now return to the question of optimal auction design for a set of n agents with values drawn from a product distribution. To solve multi-agent revenue maximization we take a standard economic approach. First, relax the incentive constraints (in particular, the monotonicity requirement of Theorem 2.2). Second, optimize revenue. Third, check to see if the incentive constraints have been violated. If the constraints have not been violated, then the solution found is optimal.

Corollary 3.7 says that the expected revenue of each agent is equal to that agent's cumulative marginal revenue. Given the distribution we can easily map the valuation profile to a quantile profile to a marginal revenue profile. To optimize marginal revenue pointwise on this valuation profile we should serve the agents to maximize the sum of the marginal revenues of served agents subject to feasibility, i.e., with allocation $\boldsymbol{y} \in \mathcal{X}$ to maximize $\sum_i R'(q_i) \cdot y_i$. E.g., for single-item auctions we would serve the agent with the highest positive marginal revenue. Since this maximizes marginal revenue pointwise, it certainly maximizes cumulative marginal revenue (i.e., in expectation).

Reconsidering the incentive constraints, observe as follows that the allocation rule of marginal revenue maximization is monotone. The concavity of revenue curves (Proposition 3.2) implies that the marginal revenues are monotonically non-increasing (in quantile). If an agent increases her value, her marginal revenue does not decrease, therefore the implied allocation rule is monotone. The price each agent is charged is given by (the dominant strategy version of) Theorem 2.2; it is the minimum bid the agent can submit and still win. We refer to this mechanism as the *marginal revenue mechanism*.

Theorem 3.8 (Myerson [74], Bulow and Roberts[20]). For single-dimensional environments, the marginal revenue mechanism is dominant strategy incentive compatible and revenue-optimal in expectation.

Proof. The optimal mechanism induces some profile of interim allocation rules $(y_1, \ldots, y_n)$, the optimal revenue for these allocation rules is given by their cumulative marginal revenues (by Corollary 3.7), the marginal revenue mechanism induces a profile of interim allocation rules with no lower total cumulative marginal revenue. Therefore, the marginal revenue mechanism is optimal. $\square$

It is interesting to consider the commonality between Theorem 3.8 and Proposition 3.1 (the analogous result for surplus maximization). While the surplus maximization mechanism finds the allocation $\boldsymbol{y}$ to maximize $\sum_i V_i(q_i) \cdot y_i$, the revenue maximizing mechanism finds the allocation $\boldsymbol{y}$ to maximize $\sum_i R_i'(q_i) \cdot y_i$. It is useful to view the marginal revenue for a given quantile of an agent as a *virtual value* for which the marginal revenue mechanism can then be viewed as optimizing *virtual surplus* (see Definition 3.8, below). Similarly, we will see in Section 5 that cumulative marginal revenue maximization is not optimal when the single-agent problems are not revenue linear; however, there will still be a virtual-value-based approach for which maximizing virtual surplus is optimal (e.g. Theorem 5.9). Existence of a virtual value function for which virtual surplus maximization is optimal in expectation reduces the revenue maximization problem to a simpler surplus maximization problem.

The characterization of the optimal mechanism in Theorem 3.8 is perhaps a bit mysterious; the next section will serve to demystify it.

3.4 Regular Distributions and Ironing

In the derivation of the optimality of the marginal revenue mechanism above it should be clear that (a) the price-posting revenue curve played an important role and (b) that when the price-posting and optimal revenues are equal then the marginal revenue mechanism and its proof of optimality are less complex. Indeed, in such a case the allocation rule y that is optimal for constraint $\hat{y}$ is essentially $y = \hat{y}$. In this section we delve into the details of this distinction.

Definition 3.7. A distribution F is *regular* if its revenue curve and price-posting revenue curve are equal (equivalently, if the price-posting revenue curve is concave).

Below we will reinterpret the characterization of the optimal auction in terms of standard auctions which are typically described in value space. The following definition translates our derivation of the marginal revenue in quantile space back to value space.

Definition 3.8. The *virtual value* of an agent with $v \sim F$ is $\phi(v) = R'(1 - F(v))$; virtual values have a closed-form formula as $\phi(v) = v - \frac{1-F(v)}{f(v)}$ when the distribution F is regular. The *virtual surplus* of an allocation is the sum of the virtual values of the agents served.[3]

---

[3] In much of the literature, the term "virtual value" is used specifically for the marginal price-posting revenue, the term "ironed virtual value" is used for the marginal revenue,

3.4.1 Regular Distributions

For regular distributions the revenue, marginal revenue, and virtual value functions have simple forms. Consider as an example an agent with value drawn uniformly from $[0,1]$ (distribution function $F(v) = v$; inverse demand curve $V(q) = 1 - q$). The price-posting revenue is calculated as follows. A price of $V(\hat{q}) = 1 - \hat{q}$ is accepted with probability $\hat{q}$. Therefore, $P(\hat{q}) = \hat{q} \cdot (1 - \hat{q})$ as with probability $\hat{q}$ the agent pays $1 - \hat{q}$. This function is concave, the uniform distribution is regular, and $R(\cdot) = P(\cdot)$. The marginal revenue $R'(q) = 1 - 2q$ is strictly decreasing and has a unique inverse. The distribution function enables transformation back and forth from value to quantile space; the virtual value function from this transformation is $\phi(v) = R'(1 - F(v)) = 2v - 1$.

Consider the unconstrained single-agent problem. What is the optimal single-agent auction? The revenue curve $R(\hat{q})$ specifies the optimal revenue from selling with any ex ante probability $\hat{q}$. If we have a single agent and are allowed to not sell the item if we do not want to, then we should pick $\hat{q}$ to maximize $R(\hat{q})$ and run the $\hat{q}$-optimal mechanism. Since $R(\cdot)$ is concave it has a global maximum at some $\hat{q}^\star$ where its derivative satisfies $R'(\hat{q}^\star) = 0$. For our $U[0,1]$ example, $\hat{q}^\star = 1/2$ and the $\hat{q}^\star$-optimal mechanism posts price $\hat{v}^\star = V(\hat{q}^\star) = 1/2$. We refer to $\hat{v}^\star$ as the *monopoly price* as it is the price at which a monopolist would sell to a single agent.

What is the optimal auction for two players with values drawn uniformly from $[0,1]$? By Theorem 3.8, the agent with the highest positive marginal revenue should win. If the quantiles of the agents are $\boldsymbol{q} = (q_1, q_2)$ then agent 1 wins whenever $R_1'(q_1) > \max(R_2'(q_2), 0)$. Because the marginal revenue curve is the same for both agents and strictly decreasing, we can deduce that agent 1 wins whenever $q_1 < \min(q_2, \hat{q}^\star)$, or, back in value space, whenever $v_1 > \max(v_2, \hat{v}^\star)$. This auction is the second-price auction with reserve $\hat{v}^\star$.

---

and the terms coincide for regular distributions. To preserve the semantic definition of virtual value as the quantity to be optimized, we will only ever refer to unironed virtual values as marginal price-posting revenues.

Corollary 3.9 (Myerson [74]). For agents with values drawn i.i.d. from a regular distribution, the optimal single-item auction is the second-price auction with the monopoly reserve price.

Recall our calculation at the section's onset of the expected revenue of the second-price auction with reserve $1/2$ from two agents with uniformly distributed values. It is now apparent that $\hat{v}^{\star} = 1/2$ is the optimal reserve price and that the auction's revenue of $5/12$ is indeed optimal among all mechanisms. It is also readily apparent that the choice of reserve price does not depend on the number of agents.

In fact, Corollary 3.9 generalizes beyond single-item environments. Whenever surplus maximization subject to feasibility is optimized by the greedy-by-value algorithm, the optimal auction maximizes surplus subject to the monopoly reserve price. The greedy-by-value algorithm is optimal if and only if the feasibility constraint is given by a matroid set system (Section A.2, Theorem A.2). Therefore, e.g., the optimal auction for i.i.d., regular, single-dimensional matching environments (a.k.a., or-preferences) is surplus maximization with the monopoly reserve price.

It is worthwhile to reemphasize where regularity is used in the proof of Corollary 3.9. For an irregular distribution two distinct values that exceed the monopoly price may correspond to the same marginal revenue and therefore, maximizing cumulative marginal revenue does not always allocate to the highest-valued agent whose value exceeds the monopoly price. Importantly, as we will discuss in the next section, a tie in marginal revenue must be broken consistently, e.g., randomly or lexicographically, but not as a function of an agent's value.

3.4.2 Irregular Distributions

Now consider the case where the value distribution is irregular, i.e., the revenue curve is strictly larger than the price-posting revenue curve for some interval of quantiles $I = (a, b)$. Since the revenue curve is the smallest concave function that upper bounds the price-posting revenue curve, the marginal revenue on this interval must be constant, i.e., $R''(q) = 0$ for $q \in I$. We refer to I as an *ironed interval* as the non-concavity of $P(\cdot)$ is "ironed" on this interval to obtain the concave revenue curve $R(\cdot)$ (equivalently, the non-monotonicity of $P'(\cdot)$ is ironed to obtain the monotone marginal revenue $R'(\cdot)$).

Since the marginal revenue for q in an ironed interval I is constant, the marginal revenue mechanism (which maximizes cumulative marginal revenue) should treat all quantiles in this interval the same. Since the measure, i.e., probability $q \in I$, of this set is $b - a > 0$ there is a measurable probability of ties in the marginal revenue mechanism. Ties can be broken randomly or arbitrarily; it is instructive to view random tie breaking as the uniform distribution over all lexicographical tie-breaking rules.

Fix a lexicographical tie-breaking rule. For each agent i, if we fix the quantiles $\boldsymbol{q}_{-i}$ of the other agents, i faces a critical quantile $\hat{q}_i$ at which she goes from winning to losing. Furthermore with lexicographical tie-breaking, this quantile $\hat{q}_i$ cannot be within any ironed interval I. Agent i, thus, faces the $\hat{q}_i$-optimal mechanism at a $\hat{q}_i$ that satisfies $R(\hat{q}_i) = P(\hat{q}_i)$ so it is simply the posted pricing of $V(\hat{q}_i)$. Importantly, the $\hat{q}$-optimal mechanisms for $\hat{q}$ strictly contained within some ironed interval I are never required for the construction of an optimal mechanism.

For an alternative viewpoint, reconsider the single-agent problem of optimizing revenue subject to an allocation constraint $\hat{y}$ (i.e., solving $\mathbf{Rev}[\cdot]$). Recall that $\hat{y}$ is a constraint, but the allocation rule y of the optimal mechanism subject to $\hat{y}$ may be generally weaker than $\hat{y}$, i.e., $y \preceq \hat{y}$. Just as we can view the ironing of the price-posting revenue curve on interval I as averaging marginal price-posting revenue on this interval, we can so view the optimization of y subject to $\hat{y}$. To optimize a weakly monotone function $R'(\cdot)$ subject to $\hat{y}$ we should greedily assign low quantiles to high probabilities of service except on ironed intervals where quantiles are assigned to the average probability of service for the ironed interval. Therefore, the desired allocation rule y can be obtained via the *resampling transformation* σ that, for quantile q in some ironed interval I, resamples the quantile from this interval, i.e., as $y(q) = \mathbf{E}_\sigma[\hat{y}(\sigma(q))]$. The cumulative allocation rule Y is exactly

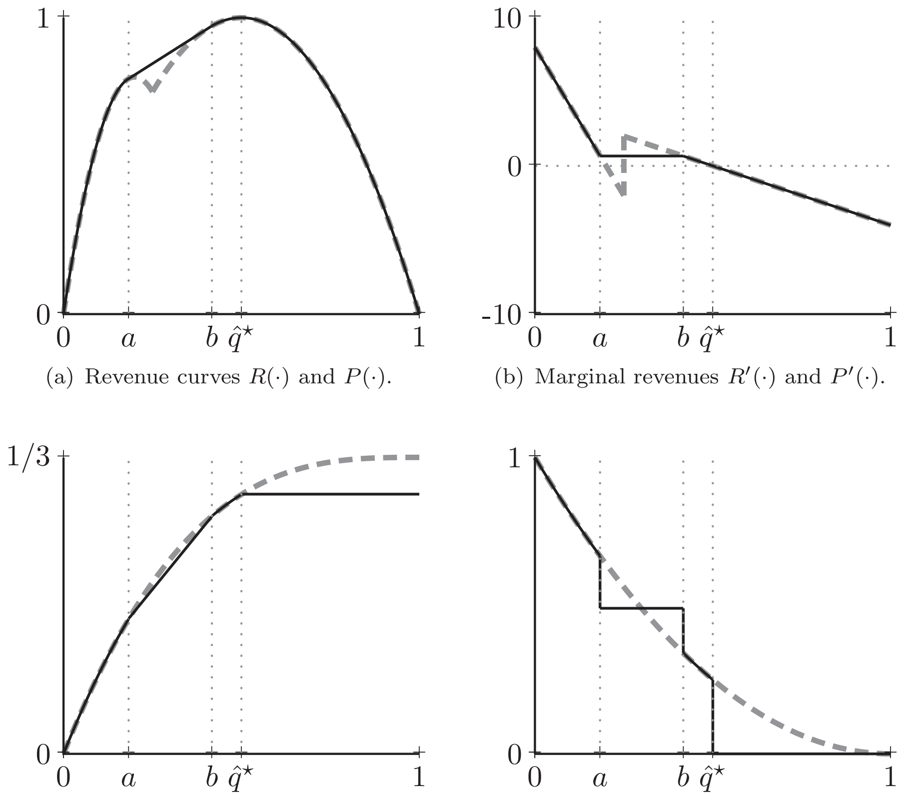

(a) Revenue curves $R(\cdot)$ and $P(\cdot)$.

(b) Marginal revenues $R'(\cdot)$ and $P'(\cdot)$.

(c) Cumulative allocations $\hat{Y}(\cdot)$ and $Y(\cdot)$.

(d) Allocation rules $\hat{y}(\cdot)$ and $y(\cdot)$.

Fig. 3.3 The optimal single-item auction is depicted for three agents with values drawn from irregular distribution defined by density function $f(v) = 1/4$ for $v \in [0,3]$ and $f(v) = 1/20$ for $v \in (3,8]$, i.e., the distribution that is uniform on $[0,3]$ with probability 3/4 and uniform on $(3,8]$ with probability 1/4. The price-posting revenue curve $P(\cdot)$ is depicted by a thick, grey, dashed line in Figure 3.3(a). The revenue curve (thin, black, solid line) is its concave hull. The ironed interval (a,b) where $R(q) > P(q)$ is depicted. The allocation constraint $\hat{y}(q) = (1-q)^2$ (Figure 3.3(d), thick, grey, dashed line) corresponds to lowest-quantile-wins for three agents; the allocation rule $y(q)$ (thin, black, solid line) results from optimizing $\mathbf{Rev}[\hat{y}]$. Simply, ironing corresponds to a line-segment for revenue curves and cumulative allocation rules and to averaging for marginal revenues and allocation rules.

equal to the cumulative allocation constraint $\hat{Y}$ except every ironed interval is replaced with a line segment. In other words, the revenue optimization of $\mathbf{Rev}[\cdot]$ can be effectively solved by superimposing the revenue curve and the allocation constraint on the same quantile axis and then ironing the allocation constraint where the revenue curve is ironed. If the environment is downward-closed, meaning we are free to exclude the agent when her marginal revenue is negative, we further drop y to zero after the quantile $\hat{q}^\star$ of the monopoly price; equivalently Y is flat after $\hat{q}^\star$. Figure 3.3 illustrates this construction.

4

Approximation Mechanisms

In this section we will answer a number of questions about approximation of the optimal mechanism for single-dimensional environments by mechanisms that are simpler and more practical. We consider three main obstacles. First, optimal mechanisms for asymmetric environments, e.g., when agents' values are not identically distributed, are not reserve-price based. Instead they involve careful optimization with respect to the agents' distinct marginal revenue curves. We give several bounds that show that surplus maximization subject to reserve prices yields a good approximation to the optimal auction. Second, we consider the issue that auctions can be a slow and inconvenient way to allocate resources. In many contexts posted pricing is preferred to auctions. The results we present give economic justification to such a preference as posted pricing mechanisms are often approximately optimal. As posted pricing mechanisms are fundamentally collusion resistant, these results imply that good collusion resistant mechanisms exist. Third, we consider the impracticality of the necessity of knowledge of the distribution of agent preferences (a.k.a., the *prior distribution*) in prescription of the Bayesian optimal auction. To address this impracticality, we discuss the design of several prior-independent mechanisms that are good approximations to the optimal mechanism for any prior distribution.

4.1 Reserve Pricing

One of the most intriguing conclusions from the preceding section is that for i.i.d., regular, single-item environments the second-price auction with a reservation price is revenue optimal (Corollary 3.9). This result is compelling as the solution it proposes is quite simple; therefore, making it easy to prescribe. Furthermore, reserve-price-based auctions are often employed in practice so this theory of optimal auctions is also descriptive. Unfortunately, i.i.d., regular, single-item environments are hardly representative of the scenarios in which we would like to design good mechanisms. Furthermore, if any of the assumptions are relaxed, reserve-price-based mechanisms are not optimal.

As an example consider a 2-agent single-item environment with agent 1's value from $U[0,1]$ and agent 2's value from $U[0,2]$. For regular distributions (Definition 3.7) it is often most convenient to work with marginal revenues in value space, i.e., with virtual values (Definition 3.8). For our two agents, the virtual value functions are $\phi_1(v_1) = 2v_1 - 1$ and $\phi_2(v_2) = 2v_2 - 2$. We serve agent 1 whenever $\phi_1(v_1) > \max(\phi_2(v_2), 0)$, i.e., when $v_1 > \max(v_2 - 1/2, 1/2)$. This auction is not the second-price auction with reserves; in particular, agent 2 has an additive handicap of 1/2.

What we see from the above example is that with non-identical distributions the optimal single-item auction indeed needs the exact marginal revenue functions to determine the optimal allocation. This contrasts to the i.i.d., regular case where all we needed was a single number, the monopoly price for the distribution, and reserve pricing with this number is optimal.

The following theorems approximately generalize Corollary 3.9 to the case of asymmetric agents. With asymmetric agents, i.e., with values $v_i \sim F_i$, the monopoly reserve prices $\hat{v}_i^\star$ may be distinct. The second-price auction with distinct reserves first removes any agent whose value does not meet her reserve, then it offers the item to the highest-valued surviving agent (if any) at a price that is the maximum of her reserve price and the second highest value of a surviving agent.

Theorem 4.1 (Chawla et al. [29]). For single-item environments and agents with values drawn independently from regular distributions, the second-price auction with (asymmetric) monopoly reserve prices obtains at least half the revenue of the optimal auction.

The main idea behind the proof of this theorem is that the revenue of the optimal auction can be attributed to two cases (a) that where the two auctions serve the same agent and (b) that where the two auctions serve distinct agents. It is easy to observe that the total revenue from the monopoly-reserves auction upper bounds the optimal revenue of each case. This is because the monopoly-reserves auction's virtual value upper bounds the optimal auction's virtual value in case (a); and because the monopoly reserve auction's payment upper bounds the optimal auction's payment in case (b). The first is because in case (a) the virtual values are equal; the second is because the monopoly-reserves auction serves the highest valued agent at a price of at least the second highest surviving agent value which is an upper bound on what the optimal auction could charge the winner in case (b). Like Corollary 3.9, this result generalizes to single-dimensional environments with a matroid constraint such as single-dimensional matching environments (a.k.a., or-preferences), see Hartline and Roughgarden [62].

While the monopoly-reserves auction (parameterized by n numbers) is significantly less complex than the optimal auction (parameterized by n functions), it is not often used in practice. In practice, even in asymmetric environments, auctions are often parameterized by a single *anonymous* reserve price. The following theorem shows that the second-price auction with an anonymous reserve is a good approximation to the optimal auction.

Theorem 4.2 (Hartline and Roughgarden [62]). For single-item environments and agents with values drawn independently from regular distributions, the second-price auction with a (suitably chosen) anonymous reserve price obtains at least a fourth of the revenue of the optimal auction.

This theorem has several proofs, the most intuitive takes the following steps. First, show that the second-price auction in a *duplicate environment*, i.e., with two agents from each distribution, obtains at least half the optimal revenue of the original environment. (The proof is similar to that of Theorem 4.1.) Viewing one agent from each pair as the original and the other as the duplicate, we can view the duplicate agents as imposing a random (and anonymous) reserve price on the original agents that is distributed according to the maximum value from the distributions. The second-price auction with a random reserve chosen from this distribution obtains at least half the revenue of the second-price auction in the duplicate environment; therefore, it is a four approximation to the optimal auction. Of course the best anonymous reserve is better than this random one.

There are two things to note about this anonymous-reserve four approximation. First, it does not extend beyond single-item auctions as does the monopoly-reserves two approximation, in fact for k-unit auctions its approximation factor is $\Omega(\log k)$. Second, it is not known to be a tight analysis (for the $k = 1$ case); the only known lower bound for the anonymous-reserve auction is two. This lower bound is exhibited by an $n = 2$ agent example where agent 1's value is a point-mass at one and agent 2's value is drawn from the *equal revenue distribution* on $[1, \infty)$, i.e., $F_2(z) = 1 - 1/z$. Importantly, for the equal revenue distribution, posting any price $\hat{v} \geq 1$ gives an expected revenue of one. For this asymmetric setting the revenue of the second-price auction with any anonymous reserve is at most one. On the other hand, an auction could first offer the item to agent 2 at a very high price (for expected revenue of one), and if (with very high probability) agent 2 declines, then it could offer the item to agent 1 at a price of one. The expected revenue of this mechanism in the limit is two.

Both Theorem 4.1 and Theorem 4.2 require that the distributions of the agents be regular; neither generalizes to the case of irregular distributions. For irregular distributions the approximation factor of monopoly and anonymous reserves are $\Omega(\log n)$. Such a logarithmic lower bound is easy to exhibit with a set of distributions based on the same principles as the equal revenue distribution described in the preceding paragraph. For instance, if agent i's value is $1/i$ with probability $1/n^2$ and zero otherwise.

4.2 Posted Pricing

Two problematic aspects of employing auctions to allocate resources are that (a) they require multiple rounds of communication (i.e., they are slow) and (b) they require all agents to be present at the time of the auction. Often both of these requirements are prohibitive. In routing in computer networks a packet needs to be routed, or not, quickly (and if our network is like the Internet then without state in the routers); therefore, auctions are unrealistic for congestion control. In a supermarket where you go to buy lettuce, we should not hope to have all the lettuce buyers in the store at once. Finally, in selling goods on the Internet, eBay has found empirically that posted pricing via the "buy it now" price is more appropriate than a slow (days or weeks) ascending auction.

In a posted pricing, distinct prices can be posted to the agents with first-come-first-served and while-supplies-last semantics. We will describe below two results for posted pricing in single-dimensional environments as well as some consequences. First, we show that *oblivious posted pricing*, where agents arrive and consider their respective prices in any arbitrary order, gives a two approximation to the optimal auction. Second, we show that *sequential posted pricing*, where the mechanism chooses the order in which the agents are permitted to consider their respective posted prices, gives an improved approximation of $\frac{e}{e-1} \approx 1.58$ with respect to the optimal auction. Both results hold for objectives of revenue and social surplus and for any independent distribution on agent values (i.e., regularity is not assumed).

4.2.1 Oblivious Posted Pricings and the Prophet Inequality

The oblivious posted pricing theorem we present is an application of a *prophet inequality* theorem from optimal stopping theory. Consider the following scenario. A gambler faces a series of n games, one on each of n days. Game i has prize v_i distributed independently according to F_i.

The order of the games and distribution of the game prizes are fully known in advance to the gambler. On day i the gambler *realizes* the prize $v_i \sim F_i$ of game i and must decide whether to keep this prize and *stop* or to return the prize and *continue* playing. In other words, the gambler is only allowed to keep one prize and must decide whether or not to keep a given prize immediately on realizing the prize and before any other prizes are realized.

A *threshold strategy* is given by a single threshold $\hat{v}$ and requires the gambler to accept the first prize i with $v_i \geq \hat{v}$. Threshold strategies are clearly suboptimal as even on day n if prize $v_n < \hat{v}$ the gambler will not stop and will, therefore, receive no prize. We refer to the prize selection procedure when multiple prizes are above the threshold as the *tie-breaking rule*. The tie-breaking rule implicit in the specification of the gambler's game is lexicographical, i.e., by "smallest i."

Theorem 4.3 (Samuel-Cahn [85]). There exists a threshold strategy such that the expected prize of the gambler is at least half the expected value of the maximum prize; moreover, for continuous distributions with non-negative support one such threshold strategy is the one where the probability that the gambler receives no prize is exactly 1/2; moreover, the bound is invariant with respect to the tie-breaking rule.

Theorem 4.3 is a *prophet inequality*: it suggest that even though the gambler does not know the realizations of the prizes in advance, she can still do half as well as a prophet who does. Notice that the order of the games makes no difference in the determination of the threshold, and if the distribution above or below the threshold changes, nothing about the bound or suggested strategy is affected. Moreover, the invariance of the performance bound to tie-breaking rule suggests the bound can be applied to other related scenarios. The profit inequality is quite robust.

Proof. Define $\hat{q}_i = 1 - F_i(\hat{v}) = \mathbf{Pr}[v_i \geq \hat{v}]$ as the probability that prize i is above the threshold $\hat{v}$ and $\chi = \prod_i (1 - \hat{q}_i)$ as the probability that the gambler rejects all prizes. The proof follows in three steps. In terms of the threshold $\hat{v}$ and failure probability χ, we get an upper bound on the expected prophet's payoff. Likewise, we get a lower bound on expected gambler's payoff. Finally, assuming that the prize distributions are continuous and non-negative,[1] we plug in $\chi = 1/2$ to obtain the result.

The prophet's expected payoff is

$$\begin{aligned} P = \mathbf{E}[\max_i v_i] &= \hat{v} + \mathbf{E}[\max_i (v_i - \hat{v})] \\ &\leq \hat{v} + \mathbf{E}\big[\max_i (v_i - \hat{v})^+\big] \\ &\leq \hat{v} + \sum\nolimits_i \mathbf{E}\big[(v_i - \hat{v})^+\big] \end{aligned}$$

where $(v_i - \hat{v})^+$ denotes $\max(v_i - \hat{v}, 0)$.

We will split the gambler's payoff into two parts, the contribution from the first $\hat{v}$ units of the prize and the contribution from the remaining $v_i - \hat{v}$ units of the prize. The first part is $G_1 = (1 - \chi) \cdot \hat{v}$. To get a lower bound on the second part we consider only the contribution from the no-tie case. For any i, let $\mathcal{E}_i$ be the event that all other prizes j are below the threshold $\hat{v}$ (but v_i is unconstrained). The bound is:

$$\begin{aligned} G_2 &\geq \sum\nolimits_i \mathbf{E}\big[(v_i - \hat{v})^+ \mid \mathcal{E}_i\big] \, \mathbf{Pr}[\mathcal{E}_i] \\ &\geq \chi \cdot \sum\nolimits_i \mathbf{E}\big[(v_i - \hat{v})^+\big]. \end{aligned}$$

The second line follows because $\chi = \prod_j (1 - \hat{q}_j) \leq \prod_{j \neq i} (1 - \hat{q}_j) = \mathbf{Pr}[\mathcal{E}_i]$ and because the conditioned variable $(v_i - t)^+$ is independent from the conditioning event $\mathcal{E}_i$. Therefore, the gambler's payoff is at least:

$$G \geq (1 - \chi) \cdot \hat{v} + \chi \cdot \sum\nolimits_i \mathbf{E}\big[(v_i - \hat{v})^+\big].$$

Plugging in $\chi = 1/2$ gives the result. □

We now wish to use the prophet inequality to show that there is an oblivious posted pricing that guarantees half the optimal revenue. The main idea is to transform from value space to virtual value space (Definition 3.8) and to consider the implications of the gambler's

[1] The theorem is true, modulo a more sophisticated choice of χ, without the continuity and non-negativity assumptions; we omit the discussion necessary to derive this more general result.

threshold strategy in virtual value space. A *uniform virtual pricing* is $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$ that satisfies $\phi_i(\hat{v}_i) = \phi_j(\hat{v}_j)$ for all i and j. The revenue of an oblivious posted pricing is worst if the agents arrive in increasing order of price, that way if multiple agents can afford their offered price, the one with the minimum price wins. For this ordering, the prophet inequality guarantees that the expected virtual value of the winning agent is at least half the maximum expected virtual value. Corollary 3.7 which says that expected revenue is equal to expected virtual surplus then implies that the oblivious posted pricing revenue approximates the optimal revenue.

Theorem 4.4 (Chawla et al. [30]). For single-item environments and agents with independent values, there is a uniform virtual pricing whose expected revenue under any order of agent arrival is at least half of that of the optimal auction.

One important consequence of the approximate optimality of oblivious posted pricings is their inherent robustness to collusion, see e.g., Goldberg and Hartline [48]. As the prices posted are fixed in advance, losing agents cannot manipulate the price of the winning agent. Analogously, competition between agents is not a crucial driver of revenue.

The oblivious posted pricing theorem extends easily to multi-unit environments and in Section 6.3 we describe its extension to the a bipartite matching environment were agents correspond to edges. It is conjectured that oblivious posted pricing is a constant approximation for any matroid environment; the closest result to this is by Kleinberg and Weinberg [66] who show that, if prices can be adjusted as agents arrive, then there is a posted pricing mechanism that adjusts to any arrival order and is a constant approximation.

4.2.2 Sequential Posted Pricings and Correlation Gap

The sequential posted pricing theorem we present is an application of a *correlation gap* theorem from stochastic optimization. Consider a non-negative real-valued set function g over subsets S of an n element ground set $N = \{1, \ldots, n\}$ and a distribution over subsets given by $\mathcal{D}$.

Let $\hat{q}_i$ be the ex ante[2] probability that element i is in the random set $S \sim \mathcal{D}$ and let $\mathcal{D}^I$ be the distribution over subsets induced by independently adding each element i to the set with probability equal to its ex ante probability $\hat{q}_i$. The *correlation gap* is then the ratio of expected value of the set function for the (correlated) distribution $\mathcal{D}$, i.e., $\mathbf{E}_{S\sim\mathcal{D}}[g(S)]$, to the expected value of the set function for the independent distribution $\mathcal{D}^I$, i.e., $\mathbf{E}_{S\sim\mathcal{D}^I}[g(S)]$.

As pointed out by Yan [94], correlation gap is central to the theory of approximation for sequential posted pricings. As an example consider a single-item auction. Let $\mathcal{D}$ be the distribution by which the optimal auction selects a winner. Importantly, only sets of cardinality at most one are in the support of $\mathcal{D}$. The probability that any individual agent is served by the optimal mechanism is $\hat{q}_i$. Of course, the revenue that the optimal mechanism generates from agent i is upper bounded by the optimal revenue that can be attained from agent i when served with ex ante probability $\hat{q}_i$. This is exactly given by the revenue curve (Definition 3.2) as $R_i(\hat{q}_i)$.

Assume for now that the distribution is regular and that the revenue of $R_i(\hat{q}_i) = \hat{q}_i\hat{v}_i$ is obtained by posting price $\hat{v}_i = V_i(\hat{q}_i)$. An upper bound on the optimal revenue is $\sum_i R_i(\hat{q}_i) = \sum_i \hat{q}_i\hat{v}_i$. This upper bound is equal to the expected value of the *maximum-weight-element* set function $g(S) = \max_{i\in S}\hat{v}_i$ for $S \sim \mathcal{D}$ (the distribution of the winner in the optimal auction).

Consider as an alternative to the optimal auction, posting prices $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$. Clearly our revenue is optimized by arranging these agents in decreasing order of price. Let S denote the set of agents whose values are at least their prices, i.e., $S = \{i : v_i \geq \hat{v}_i\}$. Each agent i is in S independently with probability $\hat{q}_i$. Importantly, S may have cardinality larger than one, but when it does, the ordering of agents by price implies that the agent $i \in S$ with the highest price wins. The revenue of the sequential posted pricing is given by the expected value of the maximum-weight-element set function $g(S)$ on $S \sim \mathcal{D}^I$.

[2] In probability theory, this probability is also known as the marginal probability of $i \in S$; however to avoid confusion with usage of the term "marginal" in economics, we will refer to it via its economic interpretation as an ex ante probability as if S was the feasible set output by a mechanism.

It is immediately clear that the ratio of the sequential posted pricing revenue to the optimal auction revenue (and its upper bound, above) is given by the correlation gap. A typical analysis of correlation gap will consider specific families of set functions g in worst case over distributions $\mathcal{D}$. For instance, when the set function is submodular (Definition A.1) then the correlation gap is $\frac{e}{e-1}$. The maximum-weight-element set function, which corresponds to single-item auctions, is submodular. In fact, it is a special case of the set function given by the maximum weight independent set in a matroid, a.k.a., the matroid weighted rank function, which is submodular (Theorem A.1).

Theorem 4.5 (Agrawal et al. [1]). The correlation gap for a submodular set function and any distribution over sets is $\frac{e}{e-1}$.

Theorem 4.5 is technical and we refer readers to Agrawal et al. [1]. Instead we give a proof of the simpler maximum-weight-element special case.

Lemma 4.6 (Chawla et al. [30]). The correlation gap for the maximum-weight-element set function and any distribution over sets is $\frac{e}{e-1}$.

Proof. This proof proceeds in three steps. First, we argue that it is without loss to consider distributions $\mathcal{D}$ over singleton sets. Second, we argue that it is without loss to consider set functions where the weights are uniform, i.e., the one-or-more set function. Third, we show that for distributions over singleton sets, the one-or-more set function has a correlation gap of $\frac{e}{e-1}$.

(1) We have a set function $g(S) = \max_{i \in S} \hat{v}_i$. Add a dummy element 0 with weight $\hat{v}_0 = 0$; if $S = \emptyset$ then changing it to $\{0\}$ affects neither the correlated value nor the independent value. Moreover, the correlated value $\mathbf{E}_{S \sim \mathcal{D}}[g(S)]$ is unaffected by changing the set to only ever include its highest weight element. This change to the set function only (weakly)

decreases the ex ante probabilities $\hat{\boldsymbol{q}} = (\hat{q}_1, \ldots, \hat{q}_n)$ and the independent value $\mathbf{E}_{S\sim\mathcal{D}^I}[g(S)]$ is monotone increasing in the ex ante probabilities. Therefore, this transformation only makes the correlation gap larger. We conclude that it is sufficient to bound the correlation gap for distributions $\mathcal{D}$ over singleton sets for which the ex ante probabilities sum to one, i.e., $\sum_i \hat{q}_i = 1$.

(2) With set distribution $\mathcal{D}$ over singletons and a maximum-weight-element set function $g(S) = \max_{i\in S} \hat{v}_i$, the correlated value simplifies to $\mathbf{E}_{S\sim\mathcal{D}}[g(S)] = \sum_i \hat{q}_i\hat{v}_i$. Scaling the weights $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$ by the same factor has no effect on the correlation gap; therefore, it is without loss to normalize so that the correlated value is $\sum_i \hat{q}_i\hat{v}_i = 1$. We now argue that among all such normalized weights $\hat{\boldsymbol{v}}$, the ones that give the largest correlation gap are the uniform weights $\hat{v}_i = 1$ for all i. This special case of the maximum-weight-element set function is the one-or-more set function, $g(S) = 1$ if $|S| \geq 1$ and otherwise $g(S) = 0$.

Sort the elements by $\hat{v}_i$ and let $c_i = \prod_{j<i}(1 - \hat{q}_j)$ denote the probability that no element with higher weight than i is in S and, therefore, i's contribution to the independent value is $c_i\hat{q}_i\hat{v}_i$. Let $\delta_i = \hat{q}_i \cdot (\hat{v}_i - 1)$ be the difference in the expected contribution to the correlated value of the ith element with value $\hat{v}_i$ or one. The expected independent value for the maximum-weight-element set function is

$$\sum\nolimits_i c_i\hat{q}_i\hat{v}_i = \sum\nolimits_i c_i \cdot (\hat{q}_i + \delta_i) \geq \sum\nolimits_i c_i\hat{q}_i. \tag{4.1}$$

where the inequality follows from monotonicity of c_i and the fact that $\sum_i \delta_i = 0$. The right-hand side of (4.1) is the expected independent value of the one-or-more set function. As the correlated value is normalized to one, uniform weights give no lower a correlation gap.

(3) The correlation gap of the one-or-more set function on any distribution $\mathcal{D}$ over singletons can be bounded as follows. First, the expected correlated value is one. Second, the

expected independent value is, for $S \sim \mathcal{D}^I$,

$$\begin{aligned}\mathbf{E}[g(S)] &= \mathbf{Pr}[|S| \geq 1] = 1 - \mathbf{Pr}[|S| = 0] = 1 - \prod\nolimits_i (1 - \hat{q}_i) \\ &\geq 1 - (1 - 1/n)^n \geq 1 - 1/e,\end{aligned}$$

where the first inequality follows because the product of a set of positive numbers with a fixed sum is maximized when the numbers are equal. □

We now pause to consider the implications of irregularity of the distribution (Definition 3.7) on the analysis of the approximation factor of sequential posted pricing via the correlation gap. In the regular case, the optimal mechanism for selling to an agent with value $v \sim F$ with some given ex ante probability $\hat{q}$ is attained by posting the price $\hat{v} = F^{-1}(1 - \hat{q})$. In the irregular case, when $\hat{q}$ falls within an interval where values are ironed, the optimal mechanism is to mix between the prices that correspond to the end points of the ironed interval so as to sell to the agent with the desired ex ante probability $\hat{q}$. Recall that we refer to this optimal single-agent mechanism as the $\hat{q}$ optimal mechanism.

If it is unproblematic to randomize the prices posted, then the randomized prices from the $\hat{q}_i$ optimal mechanism can be posted to each of the agents sequenced in order of bang-per-buck, i.e., $R(\hat{q}_i)/\hat{q}_i$. If, however, a deterministic sequential posted pricing is desired, the following approach gives one that weakly improves revenue. First, when given a sequential posted pricing, it always weakly improves revenue to sequence the agents in order of price. Second, when given a randomized sequential posted pricing, there must be a posted pricing in the support of the distribution that has above average revenue. Therefore, it suffices to consider the posted pricings in the support of the distribution, to sort the agents by price in each posted pricing, and then use the posted pricing with the highest expected revenue. Alternatively, it also weakly improves revenue to employ a simple greedy algorithm to derandomize the price of each agent individually while assuming that the other agents are regular.

Corollary 4.7 (Chawla et al. [30], Yan [94]). For single-item environments and agents with independent values, there is a sequential posted pricing with revenue that is a $\frac{e}{e-1}$ approximation to the optimal auction revenue.

Yan [94] generalized Corollary 4.7 beyond single-item auctions to matroid environments (e.g., single-dimensional matching) where, by Theorem 4.5, sequential posted pricings continue to be a $\frac{e}{e-1}$ approximation. The main idea behind the proof of this extension is that, as observed above, the weighted matroid rank function is submodular and, therefore, has a correlation gap of $\frac{e}{e-1}$ by Theorem 4.5. Yan also considered multi-unit auctions and showed that the corresponding correlation gap converges to one with k, the number of units available.

In discussion of sequential posted pricing we have employed the ex ante probabilities of service $\hat{\boldsymbol{q}}$ from the optimal auction. However, our sequential posted pricing is actually approximating something stronger, $\sum_i R_i(\hat{q}_i)$. We may as well optimize $\hat{\boldsymbol{q}}$ for this objective directly. In fact, this optimization problem corresponds exactly to the question of constructing an auction to meet the feasibility constraint in expectation, i.e., ex ante. For a single item auction, the feasibility constraint says that, ex post, for the allocation $\boldsymbol{x}(\boldsymbol{v})$ produced on valuation profile $\boldsymbol{v}$, it must be that $\sum_i x_i(\boldsymbol{v}) \leq 1$. The relaxed ex ante feasibility constraint would be $\sum_i \mathbf{E}_{\boldsymbol{v}}[x_i(\boldsymbol{v})] \leq 1$. In fact, this ex ante optimization problem is exactly the standard microeconomic problem of optimizing the sale of a commodity across distinct markets. It is given by the convex program:

$$\max_{\hat{\boldsymbol{q}}} \sum\nolimits_i R(\hat{q}_i) \qquad \text{(EAMAP)}$$
$$\text{s.t.} \sum\nolimits_i \hat{q}_i \leq 1.$$

As described previously, the marginal revenue interpretation provides a simple method for solving this program.

For matroid environments the ex ante feasibility constraint requires that the sum of the ex ante probabilities of any subset of agents be at most the rank of that subset, i.e., the maximum cardinality of an independent subset. Let rank(S) denote this matroid rank function (see Section A.2). As discussed previously, the matroid rank function is submodular. Therefore, the resulting feasibility constraint is a polymatroid (See Theorem A.3.1).

$$\begin{aligned} \max_{\hat{\boldsymbol{q}}} & \sum\nolimits_i R(\hat{q}_i) && \text{(EAMAP.M)} \\ \text{s.t.} & \sum\nolimits_{i \in S} \hat{q}_i \leq \operatorname{rank}(S), && \forall S \subset N. \end{aligned}$$

Again, the marginal revenue approach enables this program to be optimized via a simple greedy algorithm.[3]

4.3 Prior-independent Approximation

Recall and contrast the results presented previously for revenue and social surplus maximization. For surplus maximization, there is a single optimal mechanism that is not dependent on any distributional assumptions. On the other hand, for revenue maximization, the optimal mechanism depends on the distribution from which values are drawn. In this section we will consider the design of a single mechanism that is simultaneously a good approximation to the optimal mechanism for a large class of distributions.

There are generally two approaches for designing mechanisms without knowledge of a prior. First, a Bayes-Nash mechanism could ask agents to make reports based on their own distributional knowledge. Caillaud and Robert [24], for example, give a simple ascending single-item auction that has a revenue-optimal Bayes-Nash equilibrium. Mechanisms like this rely perhaps too heavily on agents' ability to behave optimally in complex environments; for instance, Bergemann and Morris [10] discuss the inherent non-robustness of this sort of mechanism. A second approach, is to restrict to dominant strategy mechanisms that use agent reports of their own values and competition between agents to ensure a good revenue. We focus on this latter approach in the sections to follow.

[3] Discretize quantile space $[0,1]$ into Q evenly sized pieces. Consider the Q-wise union of the matroid set system (the class of matroid set systems is closed under union). Calculate marginal revenues of each discretized quantile of each agent. Run the greedy-by-marginal-revenue algorithm. Calculate $\hat{q}_i$ as the total quantile of agent i that is served by algorithm, i.e., $1/Q$ times the number of i's discretized pieces that are served.

4.3.1 "Resource" Augmentation

A classic result of Bulow and Klemperer [19] states that in i.i.d., regular, single-item environments the second-price auction with one more agent (from the distribution) obtains a higher revenue than the optimal auction (without the additional agent). This result is often interpreted as a critique on exogenous entry or a statement about competition being better for revenue than reserve prices.[4] Dhangwatnotai et al. [41] give a third interpretation. The Bulow-Klemperer result suggests a prior-independent strategy for approximating the revenue of the optimal mechanism: recruit one more agent.

Theorem 4.8 (Bulow and Klemperer [19]). In i.i.d., regular, single-item environments, the expected revenue of the second-price auction on $n + 1$ agents is at least the expected revenue of the optimal auction on n agents.

Unfortunately the "just add a single agent" result fails to reasonably generalize beyond single-item auctions. When k units of an item are auctioned to $n + 1$ agents with the $(k + 1)$st-price auction, the revenue does not approximate that of the optimal k-unit n-agent auction. In order to beat the optimal auction, k additional agents must be added.

Theorem 4.9. In i.i.d., regular, k-unit environments the expected revenue of $k + 1$st-price auction on $n + k$ agents is at least the expected revenue of the optimal auction on the original n agents.

Consider the extreme case where $k = n$, a.k.a., a *digital good*. Notice that the $n + 1$st-price auction obtains no revenue (the $n + 1$st price is zero). All items are given away for free. Unfortunately, the Bulow-Klemperer result in such an environment seems less actionable as prior-independent approach to mechanism design; to obtain at least the

[4] Recall, that in Section 4.2 we came to (approximately) the opposite conclusion, i.e., that posted pricings, without competition, are enough to guarantee good revenue. Both viewpoints are correct and interesting. The Bulow-Klemperer result provides useful intuition for competitive environments, where as the Chawla et al. result provides useful intuition for environments where collusion is an issue.

revenue of the optimal mechanism we would need to double the size of the market!

4.3.2 Single-sample Mechanisms

Dhangwatnotai et al. [41] show that the Bulow-Klemperer result can be approximately extended to multi-unit environments (with one additional agent) and simultaneously address the question of how large a marketing sample must be in order to to provide good enough statistical information for the design of approximately optimal mechanisms. Their answer: one. Suppose instead of recruiting an additional agent to participate in the mechanism, we find a single agent for market analysis, and then run a second-price auction with this agent's reported value as a reserve price. In i.i.d., regular, multi-unit environments this auction is a 2-approximation.

Definition 4.1. For k-unit environments, the *single-sample auction* draws a single-sample from the distribution and runs the $k+1$st-price auction with the sampled value as a reserve price.

The following lemma can be seen as an equivalent statement to the $n = 1$ agent special case of the Bulow-Klemperer result. We give an alternative geometric proof of it that is due to Dhangwatnotai et al. [41].

Lemma 4.10 (Dhangwatnotai et al. [41]). For a single agent with value drawn from regular distribution F, the revenue from a random take-it-or-leave-it price $\hat{v} \sim F$ is at least half the revenue from the (optimal) monopoly price $\hat{v}^\star$.

Proof. Let $R(q)$ be the revenue curve for distribution F. Let $\hat{q}^\star$ be the quantile corresponding to the monopoly price, i.e., $\hat{q}^\star = \operatorname{argmax}_{\hat{q}} R(\hat{q})$. The expected revenue from such a price is $R(\hat{q}^\star)$. Recall that drawing a random value from the distribution F is equivalent to drawing a uniform quantile $\hat{q} \sim U[0,1]$. The revenue from such a random price is $\mathbf{E}_{\hat{q}}[R(\hat{q})]$. In Figure 4.1 the area of region A is $R(\hat{q}^\star)$. The area of region B is $\mathbf{E}_{\hat{q}}[R(\hat{q})]$. Of course, the area of C is less than the area of B, by

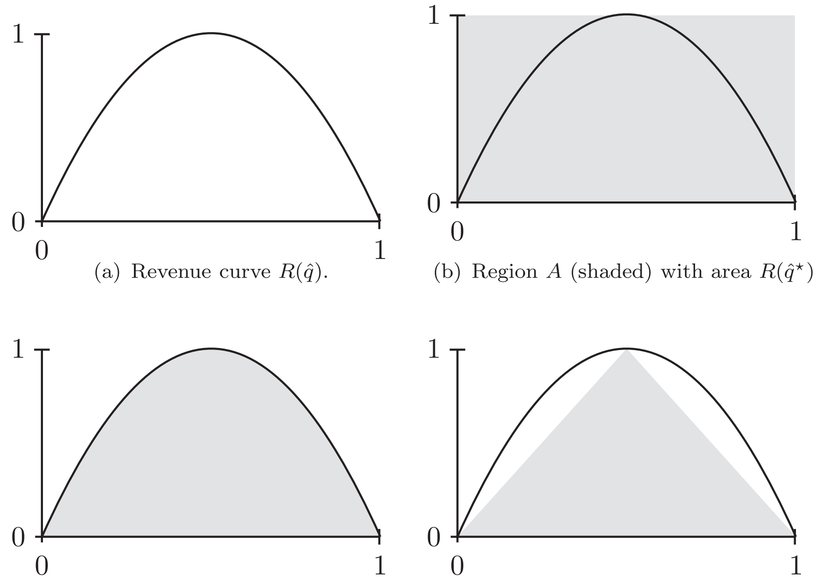

(a) Revenue curve $R(\hat{q})$. (b) Region A (shaded) with area $R(\hat{q}^\star)$.

(c) Region B (shaded) with area $\mathbf{E}[R(\hat{q})]$. (d) Region C (shaded) with area $R(\hat{q}^\star)/2$.

Fig. 4.1 In the geometric proof of the that a random reserve is a 2-approximation to the optimal reserve, the areas of the shaded regions satisfy $A \geq B \geq C = A/2$.

concavity of $R(\cdot)$, but at least half the area of A, by geometry. The lemma follows. □

It is now a simple exercise to generalize Lemma 4.10 to multi-unit auctions to give the following theorem. Essentially, a random reserve is approximately as good as the monopoly reserve. This theorem can also be generalized to matroid and downward-closed environments where the mechanism that is a good approximation is surplus maximization subject to a lazy single-sample reserve. In comparison to reserve-priced mechanisms we have previously discussed, this lazy-reserve mechanism finds the surplus maximizing set then serves only the agents in this set who meet the reserve (clearly, the feasibility of this mechanism requires downward closure).

Theorem 4.11 (Dhangwatnotai et al. [41]). For any i.i.d., regular, multi-unit environment, the single-sample auction is a 2-approximation to the optimal auction.

4.3.3 Prior-independent Mechanisms

To design approximately optimal mechanisms without any prior information, as observed by Goldberg et al. [50], Segal [87], Baliga and Vohra [8], we can use the reports of some agents for market analysis on other agents. For example, for a digital good environment, i.e., where all outcomes are feasible, and agents with i.i.d. regular distributions, the mechanism that pairs the agents and runs a second-price auction on each pair is a 2-approximation. This result follows as each agent faces a random reserve from the distribution and Theorem 4.10 implies that such a reserve is a 2-approximation to the monopoly reserve (which is optimal).

Theorem 4.12. For i.i.d., regular, digital good environments, the *pairing auction*, i.e., that randomly pairs the n agents in $n/2$ second-price auctions, is a 2-approximation to the optimal revenue.

This approach can be extended to the same environments as the single-sample auction as follows. Simulate both the surplus maximizing mechanism and the pairing auction in parallel, but serve only the winners of both mechanisms (at the higher of their prices).

4.3.4 Beyond Regular Distributions

We have exhibited a few simple approaches for designing prior-independent mechanisms. These approaches make strong usage of the regularity of the distribution, the symmetry of identical distributions, and downward-closure of the feasible set system. The main conclusion is that market analysis can be done on-the-fly by the mechanism as it is run; the resulting mechanism is often a good approximation.

The regularity assumption can be relaxed if more samples from the distribution are available. Of course, in i.i.d. environments for each agent there are $n - 1$ other samples from the distribution. The following *random-sampling mechanism* from Goldberg et al. [50] is prior-independent and gives a good approximation fairly generally. Partition the agents into two parts, estimate the distribution for one part, and run the optimal mechanism for the estimated distribution on opposite part. In some environments this can be done symmetrically for both parts. See, e.g., Goldberg et al. [50], Baliga and Vohra [8], and Devanur et al. [40].

The distributional symmetry (of the i.i.d. assumption) can be relaxed. If the agents are a priori distinguishable by publicly observable attributes, and there are at least two agents with each attribute, then agents can be paired with other agents with the same attribute. See, e.g., Balcan et al. [7], Dhangwatnotai et al. [41].

One final note, the viewpoint presented here on prior-independent mechanisms is one where there is a prior but the designer just does not know it. As the mechanisms under discussion are dominant strategy incentive compatible, the prior is not needed for equilibrium. It is possible, then, to dispense with the prior completely; however, in doing so it is not clear to what mechanisms should be compared against for approximation. Goldberg et al. [49] introduced such a prior-free framework for mechanism design and studied the aforementioned digital good environment; their benchmark is the revenue from posting the best single price for the given valuation profile. Hartline and Roughgarden [61] give post hoc justification of this benchmark as the unlimited supply special case of a more general benchmark; specifically, the supremum over i.i.d. distributions of the revenue for the optimal auction for that distribution on the given valuation profile. This benchmark has the nice property that a prior-free mechanism that approximates it simultaneously approximates the optimal mechanism for any i.i.d. Bayesian distribution. Devanur et al. [40] consider a benchmark based on envy freedom which is similar to the Hartline-Roughgarden benchmark in magnitude, but is structurally easier to work with. They show that, with a minor technical adjustment, the partitioning mechanism described above gives a good approximation to this benchmark quite generally.

5

Multi-dimensional and Non-linear Preferences

Discussion up to this point has assumed that agent preferences are single-dimensional and linear, i.e., an agent's utility for receiving a service at a given price is her value minus the price. In many settings of interest, however, agents' preferences are multi-dimensional and non-linear. Common examples include (a) multi-item environments where an agent has different values for each item, (b) agents that are financially constrained, e.g., by a budget, where an agent's utility is her value minus price as long as the price is at most her budget (if the budget is private knowledge of the agent then this agent is multi-dimensional and non-linear), or (c) agents who are risk averse; a common way to model risk averse preferences is to assume an agent's utility is given by a concave function of her value minus price.

The challenge posed by multi-dimensional non-linear preferences is three-fold. First, multi-dimensional type spaces can be large, even optimizing a single-agent problem (like those in Section 3.2) may be analytically or computationally intractable. Second, we should not expect the revenue-linearity condition of Theorem 3.6 to hold when agents have non-linear preferences (in fact, it also does not generally hold for multi-dimensional preferences). Third, often settings with multi-dimensional agents have multiple items and the externality an agent imposes on the other agents when she is served one of these items is, therefore, multi-dimensional as well.

Our approach to multi-dimensional and non-linear preferences will be to address the challenges above in roughly the order given.

5.1 Single-agent Optimization

We will consider two representative examples, (a) a single-dimensional agent with a public budget (a non-linear preference), and (b) a (multi-dimensional) unit-demand agent with linear utility given by the value she obtains for a given item minus the price of the item. For these problems we consider (i) an unconstrained single-agent problem, (ii) the ex ante constrained single-agent problem, and (iii) the interim constrained single-agent problem. In the terminology of Section 3.2 we are looking to understand the single-agent mechanisms that correspond to $R(1)$, $R(q)$, and $\mathbf{Rev}[\hat{y}]$.

5.1.1 A Mathematical Program for the Unconstrained Problem

A general agent has a type t drawn from an abstract type space T according to a distribution F. For finite type spaces, the distribution F can be specified explicitly by a probability mass function f which maps each $t \in T$ to a probability that the agent has this type. A general utility function $u(t, x, p)$ maps the agents type, allocation, and payment to the utility that the agent derives. We can write a mathematical program for the single-agent mechanism design problem by including a variable for the allocation and payment of each type, i.e., $x(t)$ and $p(t)$ for all $t \in \boldsymbol{t}$.

$$
\begin{aligned}
\max_{\boldsymbol{x},\boldsymbol{p}} \ & \textstyle\sum_{t \in T} p(t) f(t) && && \text{(SAP)} \\
\text{s.t. } & u(t, x(t), p(t)) \geq u(t, x(t^\dagger), p(t^\dagger)), && \forall t, t^\dagger \in T, && \text{(SAP.IC)} \\
& u(t, x(t), p(t)) \geq 0, && \forall t \in T, && \text{(SAP.IR)} \\
& x(t) \leq 1, && \forall t \in T. && \text{(SAP.UF)}
\end{aligned}
$$

In this program a single-dimensional budgeted agent's utility, where we interpret t as her value for service, is $u(t,x,p) = tx - p$ as long as $p \leq B$ (where B denotes her budget) and $-\infty$ otherwise. This nonlinearity can be captured in the mathematical program by adding a constraint on payments of $p(t) \leq B$ for all types $t \in T$. Monotonicity of payment (implied by Theorem 2.2) implies that it is sufficient to constrain only the payment of the highest type, denoted $t^{\max}$. I.e.,

$$p(t^{\max}) \leq B. \tag{SAP.B}$$

With this formulation the unconstrained single-agent program is linear and can be solved in time polynomial in the size of the type space, $|T|$.

To express the unit-demand agent's utility we assume that there are m items and write the agent's type t and allocation x as vectors $t = (t_1, \ldots, t_m)$ and $x = (x_1, \ldots, x_m)$ where t_j is the agent's value for item j and x_j is the probability with which she receives item j. The agent's utility is then $u(t,x,p) = t \cdot x - p$ where $t \cdot x = \sum_j t_j x_j$ denotes the vector product. With respect to (SAP.UF), let $x(t) = \sum_j x_j(t)$ be the total probability that t receives any item. While a unit-demand agent does not explicitly have linear utility, i.e., her utility for a bundle of items is the maximum of the item values not the sum of the item values, if we restrict the bundles to be size at most one then her utility is linear. Moreover, given that the agent is unit demand, there is no loss in such a restriction. To be explicit, we add the following constraint.

$$\sum_j x_j(t) \leq 1, \qquad \forall t \in T. \tag{SAP.UD}$$

Briest et al. [18] point out that by this formulation, the unconstrained single-agent problem is linear and can be effectively solved in time polynomial in the size of the type space. It should not be taken for granted, though, that unit-demand type spaces may not be small. For instance, with m items and independent values for each item, an agent would have a type space that is exponentially big in the number of items and direct optimization via the above program would be intractable; we address this issue with approximation in Section 6.1.

Notice that this formulation of the optimal unit-demand single-agent mechanism is inherently randomized. It produces a *lottery* *pricing*, in that the agent faces a priced menu of probabilistic outcomes. Optimizing over deterministic *item pricings* is computationally more difficult and not without loss for unit-demand preferences (see Section 6).

5.1.2 Mathematical Programs for the Ex Ante and Interim Problems

The following formulation of the ex ante and interim feasible single-agent mechanism design problems are from Alaei et al. [4, 5].

To derive the ex ante constrained single-agent problem, consider any single-agent mechanism. In the case that there are multiple services available, we will consider the agent to be served if she receives any of these services and denote this probability by $x(t)$ as above. The ex ante service probability of the mechanism is then, simply, the probability that a random type from the distribution is served. For an ex ante constraint of $\hat{q}$ on the service probability the following constraint can be added to the above unconstrained program (SAP).

$$\sum\nolimits_{t \in T} x(t) f(t) \leq \hat{q}. \tag{SAP.EAF}$$

From the solution to these single-agent problems, i.e., ex ante allocation constraints $\hat{q} \in [0,1]$, we can again define revenue curves as in Definition 3.2 of Section 3.2.

Now consider the single-agent optimization subject to an interim constraint $\hat{y}$ on the allocation rule, i.e., the problem of optimizing $\mathbf{Rev}[\hat{y}]$. The allocation rule of any single-agent mechanism can be derived as follows. Consider the function defined by the measure of types that are served with at least a given service probability, the allocation rule is the inverse of this function, i.e., $y(q) = \sup\{x^\dagger : \mathbf{Pr}_{t \sim F}[x(t) \geq x^\dagger] \geq q\}$. For discrete type spaces this allocation rule can be found by making a rectangle of width $f(t)$ and height $x(t)$ for each type t, and sorting these rectangles in decreasing order of height. The resulting non-increasing piece-wise constant function on $[0,1]$ is the allocation rule y. Any mechanism induces such an allocation rule.

The feasibility constraint in the unconstrained single-agent program (SAP) can be modified for an interim allocation constraint $\hat{y}$

with cumulative constraint $\hat{Y}$ (recall Definition 3.3). Intuitively, $\hat{Y}(\hat{q})$ bounds the probability with which any measure $\hat{q}$ of types can be served.

$$\sum\nolimits_{t \in S} x(t) f(t) \leq \hat{Y}(f(S)), \qquad \forall S \subseteq T. \tag{SAP.IF}$$

From an optimization point of view it is useful to note that, while the interim feasibility constraint is defined across all subsets of the type space, it is easy to optimize. For instance, the concavity of $\hat{Y}(\cdot)$ and linearity of $f(\cdot)$, defined as $f(S) = \sum_{t \in S} f(t)$, imply the submodularity of $\hat{Y}(f(\cdot))$ on subsets of type space and, therefore, that the feasible space is a polymatroid (see Proposition A.4). Moreover, if we consider types ordered by decreasing allocation probability (SAP.IF) is only binding on the linear number of sets corresponding to prefixes of the ordering.

5.1.3 Analytical Solutions for Budgeted and Unit-demand Agents

We now consider analytical solutions to special cases of the interim single-agent problem, i.e., (SAP) with (SAP.IF), with continuous type spaces. In particular, we will sketch constructions for the case of a budgeted agent with a regular distribution (Definition 3.7) of non-decreasing density, and the case of a unit-demand agent with values for the m items distributed uniformly on the unit hypercube.

Theorem 5.1 (Laffont and Robert [68]; Pai and Vohra [80]). For a budgeted agent with value distributed from a regular distribution with non-decreasing density, the optimal single-agent mechanism for allocation constraint $\hat{y}$ is parameterized by quantiles $\hat{q}^\dagger < \hat{q}^\ddagger$ where quantiles in $[0, \hat{q}^\dagger]$ are ironed,[1] quantiles in $q \in (\hat{q}^\dagger, \hat{q}^\ddagger]$ receive the corresponding allocation probability $\hat{y}(q)$, and the remaining quantiles, $(\hat{q}^\ddagger, 0]$, are rejected.

---

[1] I.e., each receives the average allocation probability of $\hat{y}$ for $q \in [0, \hat{q}^\dagger]$; this is average allocation probability is equal to $\hat{Y}(\hat{q}^\dagger)/\hat{q}^\dagger$ (cf. Subsection 3.2.4).

The proof of this theorem follows from the single-dimensional agent analysis where the objective is rewritten as the cumulative marginal price-posting revenue of the agents served. The budget constraint (SAP.B) can be Lagrangian relaxed and incorporated into the objective with parameter λ. The effective price-posting revenue curve becomes $P^{\lambda}(\hat{q}) = (\hat{q} - \lambda)V(\hat{q})$ for $\hat{q} \in (0,1]$ and $P^{\lambda}(0) = 0$. Notice that for $\hat{q} < \lambda$ this quantity is non-positive. Regularity and non-increasing density imply that this Lagrangian price-posting revenue curve is concave on $(0,1]$. The observed form of the optimal auction then follows from observing that the Lagrangian revenue curve, i.e., the smallest concave function that upper bounds the Lagrangian price-posting revenue curve, irons only on some interval $[0, \hat{q}^{\dagger})$ and has negative slope after some $\hat{q}^{\ddagger}$.

Further discussion optimal mechanisms for agents with public budgets and the generalization of Theorem 5.1 to arbitrary distributions is given by Devanur et al. [37].

Theorem 5.2 (Haghpanah [56]). For a unit-demand agent of m items with value uniformly distributed on the m-dimensional hypercube, the optimal single-agent mechanism for allocation constraint $\hat{y}$ projects the agent's type into a single dimension, by considering only the agent's favorite item $j^{\star} = \operatorname{argmax}_j t_j$, and optimally serves according to $\hat{y}$ and the distribution of the agent's value for her favorite item. The single-dimensional value distribution in the projection corresponds to the maximum of m uniform random variables which is a regular distribution; therefore, the allocation rule y is the same as $\hat{y}$ except with an appropriate reserve price.

The proof of this theorem is technical, but at a high-level follows the marginal revenue approach (cf. Subsection 3.2.4). Notice that the marginal revenue approach is a kind of amortized analysis. The marginal revenue of a type encompasses the change in the total revenue when this type is served along with higher types (i.e., it is the gain in revenue from serving this type less the loss in revenue from higher types due to lowering the price so that this type may also be served). The utility of a type is the surplus of the type for the service received less the payment. Analogous to the single-dimensional characterization of Theorem 2.2 where the allocation rule is required to be monotone and the payment is given by the integral over the allocation to lower types; a multi-dimensional mechanism is in Bayes-Nash equilibrium if and only if the utility is convex and the payment is given by a path integral (e.g., from the low type $t^{\min} = (0, \ldots, 0)$; see Rochet [81]). What makes the multi-dimensional setting complex is that, while with a single-dimension there is only one path that connects a given type to the lowest type; in multi-dimensional settings there are many such paths.

The proof of Theorem 5.2 is by the following high-level argument. We will write the expected payment of a type as the difference in surplus and utility (where the surplus can be written in terms of the gradient of the utility) and take expectation over types. In formulating the payment as a path integral, use the Euclidean straight-line path from the origin to the type t, which is interpreted as a point in $T = [0,1]^m$. We can swap the order of integration (from taking expectation over types and from the path integral) and reformulate the problem in terms of a multi-dimensional virtual value times the gradient of the utility function. This virtual value for the coordinate $j^\star$ is exactly the marginal revenue of the distribution that corresponds to the maximum of m uniform random variables (and is a smaller value for coordinates $j \neq j^\star$). Relaxing the incentive constraints (in this case, the convexity of the utility function) and optimizing these virtual values suggests sorting types in probability space by their maximum coordinate and assigning service probability according to $\hat{y}$ to those types that meet a given reserve (given by the monopoly price for the distribution of the maximum of m uniform random variables). Finally, we reconsider the relaxed incentive constraints (convexity of the agent's utility) and confirm that they are satisfied by the mechanism obtained; therefore, it must be optimal.

5.2 Multi-agent Optimization with Revenue Linearity

Consider an environment, like the single-dimensional environments of Section 3, where agents only impose a single-dimensional externality on each other. In these environments each agent would like to receive an abstract service and there is a feasibility constraint over the set of agents who can be simultaneously served. Agents may also have preferences over unconstrained attributes that may accompany service. Payments are one such attribute; for example, a seller of a car can only sell one car, but she can assign arbitrary payments to the agents (subject to the agents' incentives, of course). Likewise, the seller of the car could paint the car one of several colors as it is sold and the agents may have multi-dimensional preferences over colors. Of course, if the car is sold to one agent then it cannot be sold to other agents so, while the color plays an important role in the agents' multi-dimensional incentive constraints, it plays no role in the feasibility constraints. We refer to environments with single-dimensional externalities as service constrained environments. The more general case of multi-dimensional externalities is deferred to Section 5.4.

Definition 5.1. A *service constrained environment* is one where a feasibility constraint restricts the set of agents who can be simultaneously served, but imposes no restriction on how they are served.

We now characterize optimal multi-agent mechanisms in terms of the solution to single-agent problems with the simplifying revenue-linearity property. Revenue linearity immediately implies that the optimal cumulative marginal revenue is equal to the optimal revenue (Proposition 3.5). However, it does not immediately suggest how to implement cumulative marginal revenue maximization as the mapping from type space to quantile space is not implicit in a multi-dimensional type space (as it is in a single-dimensional type space). For example, we will see shortly that an agent with value uniformly distributed on an m-dimensional hypercube (recall Theorem 5.2) is revenue linear. However, it is unclear which type is the stronger, e.g., for $m = 3$, of $(.9, .2, .2)$ or $(.7, .7, .7)$? The definition and theorem below resolve this issue.

Definition 5.2. A single-agent problem given by a type space and distribution is *orderable* if there is an equivalence relation on types and an ordering over equivalence classes, such that for every allocation constraint $\hat{y}$, the optimal mechanism for $\hat{y}$, i.e., solving $\mathbf{Rev}[\hat{y}]$, induces an allocation rule that is greedy by the given ordering with ties broken uniformly at random and with types in a special lowest equivalence class $\perp$ (if any) rejected.

Notice that the single-dimensional budgeted example is not orderable by the above definition. Though, the agent's values give a weak ordering on types, the optimal mechanism for $\hat{y}$ irons so as to meet the budget constraint with equality and the values in ironed intervals depend on the allocation constraint $\hat{y}$ (see Theorem 5.1).

Theorem 5.3 (Alaei et al. [5]). For any single-agent problem, revenue linearity implies orderability.

The main intuition behind Theorem 5.3 comes from observing that the allocation rule y that is obtained from optimizing subject to $\hat{y}$ must be equal to $\hat{y}$ at all quantiles where the revenue curve is strictly concave; the equivalence classes in the theorem statement then correspond to types with equal marginal revenue (which have non-zero measure only on intervals where the revenue curve is linear).

Theorem 5.3 says that while there is not an inherent ordering on type space that is respected by all mechanisms, there is one that is respected by all optimal mechanisms. The following corollary is immediate; from it the marginal revenue mechanism for orderable agents is easy to generalize.

Corollary 5.4 (Alaei et al. [5]). The $\hat{q}$ ex ante optimal single-gent mechanism (Definition 3.2), for any revenue-linear agent and at any $\hat{q}$ for which the revenue curve $R(\cdot)$ is strictly concave, serves exactly the types corresponding to quantiles $q \in [0, \hat{q}]$.

Definition 5.3. The *marginal revenue mechanism for orderable agents* works as follows:

(a) Map the agents' types $\boldsymbol{t}$ to quantiles $\boldsymbol{q}$ via the ordering guaranteed to exist by the orderability assumption.[2]
(b) Calculate marginal revenues for the profile of quantiles.
(c) For each agent i, calculate the supremum quantile $\hat{q}_i$ she could possess and still be in the cumulative marginal revenue maximizing feasible set (S to maximize $\sum_{i \in S} R_i'(q_i)$).
(d) For each agent i, run the $\hat{q}_i$ optimal single-agent mechanism.

Theorem 5.5 (Alaei et al. [5]). The marginal revenue mechanism for revenue-linear agents is (a) dominant strategy incentive compatible, (b) feasible, (c) revenue-optimal, and (d) deterministically selects the set of winners.

Proof. Dominant strategy follows because for any $\hat{q}_i$ for agent i, the $\hat{q}_i$ optimal mechanism is dominant strategy incentive compatible. It is feasible as the set of agents served by the optimal single-agent mechanisms is exactly the set that maximizes cumulative marginal revenue (by Corollary 5.4). It is revenue optimal because it maximizes marginal revenue pointwise, each agent's critical quantile in the interim stage is drawn from an distribution that corresponds to that agent's interim allocation rule. The expected revenue from this agent, therefore, is the expected marginal revenue. By the definition of the mechanism, it is marginal revenue optimal; by revenue linearity, it is revenue optimal. □

Finally, we instantiate and interpret Theorem 5.5 for an interesting multi-dimensional example. Suppose we are selling a car by auction to

[2] For types in equivalence classes defined by Definition 5.2 that have measurable size (with respect to the distribution), we can uniformly pick a quantile from within the measure. This choice ensures that quantiles are uniform (which is a standing assumption); though, this random choice is never relevant to the mechanism as all quantiles in the interval corresponding to this equivalence class are mapped to the same marginal revenue in step (b).

a set of n agents, as the car is sold it can be painted one of m colors. Each agent's type is m-dimensional with the jth coordinate denoting her value for the jth color. This is a service constrained environment. Furthermore, when agent values are uniform on the m-dimensional hypercube Theorem 5.2 implies that the problem is equivalent to the single-dimensional projection of each agent's favorite color. The revenue linearity of this single-dimensional projection implies the revenue linearity of the original unit-demand problem.

Corollary 5.6 (Haghpanah [56]). The single-agent problem for a unit-demand agent of m items with values uniformly distributed on the m-dimensional hypercube is revenue linear.

The implied marginal revenue mechanism is to auction the car with free color choice via the second-price auction with reserve set as the monopoly price for the distribution of the maximum of m uniform random variables.

While it is often assumed that the optimality of the marginal revenue mechanism is special to single-dimensional, linear agents (as in Section 3), we have shown here that the special condition is revenue linearity not single dimensionality.

5.3 Multi-agent Optimization without Revenue Linearity

We now characterize optimal multi-agent mechanisms for service constrained environments (Definition 5.1) in terms of the solution to single-agent problems without the simplifying revenue-linearity property. It is useful to contrast the complexity of the optimal mechanism here with that of the optimal mechanism with revenue linearity discussed previously in Section 5.2.

5.3.1 Interim Feasibility

Consider any multi-agent mechanism $\mathcal{M}$. When the agent types $\boldsymbol{t}$ are drawn from the product distribution $\boldsymbol{F}$, each agent i has an induced *interim* mechanism $\mathcal{M}_i$. This interim mechanism maps the agent's type t_i to a distribution over outcomes (including the service received and non-service-constrained attributes such as payments). Any single-agent mechanism $\mathcal{M}_i$ induces an allocation rule y_i as described in the Section 5.1; since this is an interim mechanism we refer to this allocation rule as the interim allocation rule. Repeating this construction for each of the n agents we obtain a profile of interim allocation rules $\boldsymbol{y}$ that is feasible in the sense that there exists an ex post feasible mechanism (e.g., $\mathcal{M}$) that induces it.[3]

Definition 5.4. A profile of allocation rules $\boldsymbol{y}$ is *interim feasible* if it is induced by some ex post feasible mechanism $\mathcal{M}$ and type distribution $\boldsymbol{F}$.

It is useful to consider some examples of interim allocation profiles that are feasible and infeasible. Consider selling a single item to one of two agents each with one of two interim allocation rules:

$$y^{\dagger}(q) = 1/2, \qquad\qquad y^{\ddagger}(q) = \begin{cases} 1 & \text{if } q \in [0, 1/2], \\ 0 & \text{otherwise.} \end{cases}$$

Notice that both allocation rules have an ex ante probability of 1/2 of allocating (as agent quantiles are always drawn from the uniform distribution). Consider the profile of allocation rules $\boldsymbol{y} = (y^{\dagger}, y^{\dagger})$, i.e., where both agents have interim allocation rule $y^{\dagger}$. This profile is interim feasible as it is the outcome of the *fair-coin-flip* mechanism. Similarly the profile $\boldsymbol{y} = (y^{\ddagger}, y^{\dagger})$ is interim feasible, it is induced by the mechanism that serves agent 1 if she has a high type (i.e., $q_1 \in [0, 1/2]$) and agent 2 otherwise. The profile $\boldsymbol{y} = (y^{\ddagger}, y^{\ddagger})$ is not interim feasible. If both agents have high types, which happens with probability 1/4, the interim allocation rules require that both agents be served, but doing so would not be ex post feasible.

Our goal is to maximize expected revenue over (Bayesian incentive compatible and interim individually rational) mechanisms subject to ex post feasibility. This is equivalent to optimizing expected revenue

[3] In the literature, e.g., Border [14], such a profile of interim allocation rules is often referred to as the *reduced form* of the mechanism.

over profiles of allocation constraints subject to interim feasibility. I.e.,

$$\max_{\hat{\boldsymbol{y}}} \sum\nolimits_i \mathbf{Rev}[\hat{y}_i] \tag{MAP}$$

$$\text{s.t. "}\hat{\boldsymbol{y}}\text{ is interim feasible."} \tag{MAP.IF}$$

To get some intuition for this program, notice that if the agents are i.i.d. then the optimal mechanism design problem is symmetric and there is always a symmetric solution to a symmetric convex optimization problem. For a single-item environment the *lowest-quantile-wins* mechanism induces the allocation constraint (for each of the n agents) of $\hat{y}(q) = (1 - q)^{n-1}$. This is the strongest (Definition 3.3) interim feasible symmetric allocation constraint. From this allocation constraint and the single-agent mechanism that solves $\mathbf{Rev}[\hat{y}]$, the optimal mechanism can be constructed (see Figure 3.3). This observation has explicitly or implicitly enabled most of the derivations of optimal auctions for agents with non-linear preferences, e.g., see Matthews [73], Border [14], and Laffont and Robert [68].

Optimization of this mathematical program in asymmetric environments (e.g., when the agents' distributions are not identical) relies on better understanding the constraint posed by interim feasibility. Consider first interim feasibility in single-item environments. Take any profile of ex ante constraints $\hat{\boldsymbol{q}} = (\hat{q}_1, \ldots, \hat{q}_n)$ and consider the ex ante probability by which each agent i with quantile at most $\hat{q}_i$ is served, i.e., $Y_1(\hat{q}_1), \ldots, Y_n(\hat{q}_n)$. Of course the probability that some agent i with quantile bounded by $\hat{q}_i$ is realized is $1 - \prod_i (1 - \hat{q}_i)$.

Theorem 5.7 (Border [14, 15]). For single-item environments, a profile of allocation rules $\boldsymbol{y}$ (with cumulative allocation profile $\boldsymbol{Y}$) is interim feasible if and only if,

$$\sum\nolimits_i Y_i(\hat{q}_i) \leq 1 - \prod\nolimits_i (1 - \hat{q}_i), \qquad \forall \hat{\boldsymbol{q}} \in [0,1]^n.$$

To understand the interim feasibility condition of Theorem 5.7 note that the left-hand side is the expected number agents with quantiles in their respective ranges that are served and the right-hand side is the probability that an agent with quantile in her requisite range is realized at all. The necessity of this condition is obvious as we cannot serve agents with quantiles in the given ranges more than there is such an agent. The sufficiency follows, for example, from a max-flow-min-cut-style argument. Alaei et al. [4] extend this characterization to matroid environments where the right-hand side becomes the *expected rank*, with short-hand notation $\operatorname{rank}(\hat{\boldsymbol{q}})$ representing $\mathbf{E}_S[\operatorname{rank}(S)]$ where $i \in S$ independently with probability $\hat{q}_i$ (cf. Subsection 4.2.2).

Theorem 5.8 (Alaei et al. [4]). For matroid environments, a profile of allocation rules $\boldsymbol{y}$ is interim feasible if and only if,

$$\sum\nolimits_i Y_i(\hat{q}_i) \leq \operatorname{rank}(\hat{\boldsymbol{q}}), \qquad\qquad \forall \hat{\boldsymbol{q}} \in [0,1]^n.$$

For feasibility constraints beyond matroids, we will not get a succinct representation for the interim feasibility constraint. However, we will nonetheless be able to describe optimization subject to interim feasibility quite naturally.

5.3.2 Optimization and Ex Post Implementation

In the discussion above we expressed interim feasibility in quantile space. The advantage of formulating the interface between the single-agent problem and the multi-agent problem in quantile space is that it is single-dimensional and therefore can be represented succinctly even when the type spaces of the agents are very large. The solution to the single-agent problem needs then to specify both the (stochastic) mapping from types to outcomes and the (stochastic) mapping from types to quantiles. The optimized mechanism then, for a reported type profile, maps types to quantiles, determines the set of winners from the quantile profile, and then samples service outcomes for the winners and non-service outcomes for the losers from the outcome mapping.

For understanding interim feasibility a number of aspects of the original mechanism design problem are not relevant. First, the distinction between quantiles and types is not relevant; both are just names by which we distinguish between the possible values of each coordinate of the input (while our discussion below is given in type space, it could be equally well given in quantile space). Second, the independence of agent types will not be important. An input type profile is $\boldsymbol{t}$, it is drawn from distribution $\boldsymbol{F}$, and we write the marginal probability of agent i having type t as $f_i(t) = \mathbf{Pr}_{\boldsymbol{t}\sim\boldsymbol{F}}[t = t_i]$.

The important quantity for interim feasibility is the ex ante probability that specific types of specific agents are served. To make these quantities more explicit in our notation we employ the following shorthand. Up to this point, we have viewed an ex post allocation rule $\boldsymbol{x}(\boldsymbol{t})$ as a function mapping an n-dimensional type profile to an n-dimensional vector indicating service to agents. It will be useful to view it as a *flattened* vector indicating service to each of the possible types of each agent, i.e., as vector of dimension equal to the sum of the sizes of the agent type spaces where coordinate it indexes $x_{it}(\boldsymbol{t})$ to indicate whether agent i with type t is served. Of course if $t_i \neq t$ in type profile $\boldsymbol{t}$ then $x_{it}(\boldsymbol{t}) = 0$ as a mechanism cannot serve a type that "does not show up". Similarly, we consider a profile of interim allocation rules $\boldsymbol{x}$ as a flattened vector with x_{it} as the (ex ante) probability that i with type t is served by the mechanism. The short-hand notations for flattened ex post and interim allocation vectors, respectively, are,

$$x_{it}(\boldsymbol{t}) = \begin{cases} x_i(\boldsymbol{t}) & \text{if } t = t_i \\ 0 & \text{otherwise,} \end{cases} \qquad\qquad x_{it} = x_i(t)/f_i(t).$$

With such a redundant notation, the interim allocation vector is just the expectation of the ex post allocation vector.

$$\boldsymbol{x} = \mathbf{E}_{\boldsymbol{t}\sim\boldsymbol{F}}[\boldsymbol{x}(\boldsymbol{t})]. \tag{5.1}$$

We are now prepared to make a number of simple observations. First, as we allow randomized mechanisms, both the spaces of ex post feasible allocations $\boldsymbol{x}(\boldsymbol{t})$ and interim feasible allocations $\boldsymbol{x}$ are convex. Moreover, as the ex post constraints are linear and the type space is finite, both the ex post and interim feasibility are given by a polytope. As previously observed, our single-agent optimal mechanisms may also be randomized, therefore, the single-agent revenue (given by the revenue operator $\mathbf{Rev}[\cdot]$) is concave. Therefore, the optimization over interim feasible allocation rules of (MAP) is a convex optimization problem. See Section A.3 for a review of convex optimization.

Suppose instead our objective was linear and given by weights $\boldsymbol{w}$ indexed in our flattened vector space by it; then optimization of these weights in expectation, i.e., $\sum_{it} x_{it} w_{it}$, subject to interim feasibility, is achieved by optimizing them pointwise for each profile of types $\boldsymbol{t}$. Moreover, for any such weights (which we view as a direction in the flattened vector space) we can, via (5.1), estimate the corresponding interim feasible allocation vector $\boldsymbol{x}^{\boldsymbol{w}}$ by sampling. To simplify the discussion we assume an ideal model where we can exactly calculate $\boldsymbol{x}^{\boldsymbol{w}}$ from $\boldsymbol{w}$.

A vertex of a convex polytope can be specified by a direction $\boldsymbol{w}$ that is orthogonal to a separating hyperplane at this vertex; this vertex is given by optimizing, as above, in the given direction. Any interior point in a convex polytope can be represented as a convex combination of vertices of the polytope; therefore, we can implement any interim feasible allocation rule by sampling a direction from an appropriate distribution, and then for the type profile realized, optimizing the weights given by the direction subject to ex post feasibility. Cai et al. [23] point out that this stochastic direction can be viewed as a randomized virtual value in the spirit of Myerson [74]. The following theorem applies to any, even non-downward-closed, ex post feasibility constraint.

Theorem 5.9 (Cai et al. [23], Alaei [3]). For any joint distribution on type profiles and service constrained environment, any interim feasible allocation profile can be ex post implemented as a stochastic virtual surplus maximization (i.e., serving the set of agents with the highest cumulative virtual value) with virtual value functions drawn from an appropriate joint distribution.

In the special case that the set system is a matroid, the ex post feasibility constraint is a polymatroid (and so is the induced interim feasibility constraint, cf. Theorem 5.8). Greedy algorithms solve weighted optimization subject to polymatroid constraints (Theorem A.3.1). Therefore, we get a refinement on Theorem 5.9 as follows (the single-item case is by Border [15]; the generalization to matroids is by Alaei et al. [4]).

Theorem 5.10 (Border [15], Alaei et al. [4]). For any joint distribution on type profiles and any service constrained matroid environment, any interim feasible allocation profile can be ex post implemented by the greedy algorithm on a stochastic ordering of agent types with the ordering drawn from the the appropriate joint distribution (where types in a designated lowest class $\perp$ are always rejected).

We now turn to the question of finding the appropriate mixture over directions for Theorem 5.9. In general, a convex program such as (MAP) can be optimized by the ellipsoid method if there is a *separation oracle*, i.e., an algorithm that determines whether a given *query point* is in the feasible polytope or not, and if it is not, gives a separating hyperplane, a plane that separates the given point from the feasible polytope. Cai et al. [23] give a method for finding a separating hyperplane for interim feasibility; we describe a similar approach from Saeed Alaei's Ph.D. thesis [3]. If the query point is feasible the method proves its feasibility by giving a small (i.e., polynomial in the dimension) set of vertices of the polytope that contain the point in their convex hull. To implement this feasible point, it suffices to solve for it as a linear combination of these vertices. Moreover, these vertices are specified by the direction by which linear optimization yields the vertex. Therefore, the interim feasible point can be implemented by drawing one of these directions according to the convex combination and then optimizing ex post for that direction.

Consider the following method for finding a separating hyperplane for a query point $\boldsymbol{x}$ and the interim feasible polytope. Consider any $\boldsymbol{w}$ in the unit ball (i.e., of length at most one). Given a query point $\boldsymbol{x}$ the direction $\boldsymbol{w}$ gives a separating hyperplane under the following condition (see Figure 5.1(a)). Let $\boldsymbol{x}^{\boldsymbol{w}}$ be the interim feasible point for maximizing in the direction of $\boldsymbol{w}$. Consider the projection of $\boldsymbol{x}$ and $\boldsymbol{x}^{\boldsymbol{w}}$ onto $\boldsymbol{w}$ (i.e., the objective value of both points with respect to the weights). If the objective value for $\boldsymbol{x}$ exceeds that of $\boldsymbol{x}^{\boldsymbol{w}}$ then $\boldsymbol{x}$ cannot be interim feasible: for objective $\boldsymbol{w}$ it exceeds in objective value the optimal interim feasible point $\boldsymbol{x}^{\boldsymbol{w}}$. Indeed, a separating hyperplane is the one orthogonal to weights that passes through $\boldsymbol{x}^{\boldsymbol{w}}$.

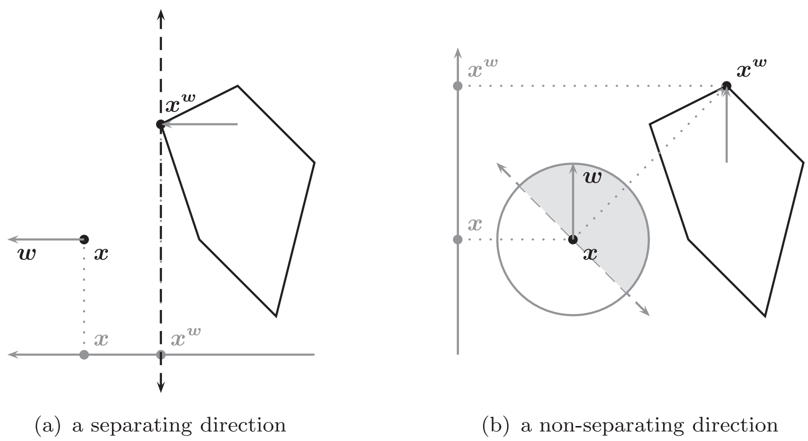

Fig. 5.1 Depicted are a query point $\boldsymbol{x}$, the interim feasible polytope, a direction $\boldsymbol{w}$, and the vertex $\boldsymbol{x}^{\boldsymbol{w}}$ of the interim feasible polytope that is optimal in direction $\boldsymbol{w}$. The projections of $\boldsymbol{x}$ and $\boldsymbol{x}^{\boldsymbol{w}}$ onto $\boldsymbol{w}$ (the gray line) are depicted in gray. On the left, the projection of $\boldsymbol{x}$ exceeds that of $\boldsymbol{x}^{\boldsymbol{w}}$ in the direction of $\boldsymbol{w}$; therefore, the plane through $\boldsymbol{x}^{\boldsymbol{w}}$ orthogonal to $\boldsymbol{w}$ (depicted by a dashed line) separates $\boldsymbol{x}$ from from the interim feasible polytope. On the right, the projection of $\boldsymbol{x}$ does not exceeds that of $\boldsymbol{x}^{\boldsymbol{w}}$ in the direction of $\boldsymbol{w}$; therefore, $\boldsymbol{w}$ does not give a separating hyperplane for the interim feasible polytope. The plane (grey dashed line) orthogonal to $\boldsymbol{x} - \boldsymbol{x}^{\boldsymbol{w}}$ (grey dotted line) separates the direction $\boldsymbol{w}$ (and all directions in the shaded area) from the space of directions that give separating hyperplanes for interim feasibility.

To find such a separating hyperplane for $\boldsymbol{x}$, we must find such a direction $\boldsymbol{w}$. The subset of the unit ball (i.e., directions) that gives separating hyperplanes is convex (see Figure 5.1(b)); therefore, an approach to find a separating hyperplane is to use the ellipsoid method recursively. To do so, for any direction that does not give a separating hyperplane to the interim feasibility problem, we must give a separating hyperplane between it and the convex space of directions that does give separating hyperplanes for the interim feasibility problem.

Suppose that we have a direction $\boldsymbol{w}$ that does not give a separating hyperplane to interim feasibility, see Figure 5.1(b); specifically, the objective value (with respect to $\boldsymbol{w}$) of the interim feasible solution $\boldsymbol{x}^{\boldsymbol{w}}$ that optimizes for $\boldsymbol{w}$ exceeds that of $\boldsymbol{x}$. We get a separating hyperplane for $\boldsymbol{w}$ by asking for what other objectives $\boldsymbol{w}^\dagger$ does $\boldsymbol{x}^{\boldsymbol{w}}$ prove to be non-separating, i.e., where $\boldsymbol{w}^\dagger \cdot \boldsymbol{x}^{\boldsymbol{w}}$ exceeds $\boldsymbol{w}^\dagger \cdot \boldsymbol{x}$. The space of weights ruled out by $\boldsymbol{x}^{\boldsymbol{w}}$, i.e., the grey shaded area in Figure 5.1(b), is given by the hyperplane that is orthogonal to the direction from $\boldsymbol{x}$ to $\boldsymbol{x}^{\boldsymbol{w}}$.

The ellipsoid method, in polynomial time in the dimension and desired precision, will either find a separating hyperplane or the set of directions considered by the method correspond to vertices of the interim feasible polytope that contain the query point $\boldsymbol{x}$ in their convex hull. As described above, decomposing $\boldsymbol{x}$ as a convex combination of these vertices gives an ex post mechanism for implementing $\boldsymbol{x}$: sample vertices (a.k.a., directions) according to the convex combination and optimize the ex post allocation according to the sampled direction.

Theorem 5.11 (Cai et al. [23]). For any joint distribution on type profiles and any feasibility constraint on allocations, maximization of a concave function of interim allocation constraints subject to interim feasibility reduces in polynomial time to weighted maximization subject to ex post feasibility.

To prove Theorem 5.11 we must replace ideal sampling of $\boldsymbol{x}^{\boldsymbol{w}}$ with real sampling which is inexact. This inexactness results in error that is polynomial small. Generally the error must be made exponentially small to employ the ellipsoid method (Theorem A.5); Cai et al. [23] address this difficulty.

5.3.3 Single-item Environments, Revisited

For a single-item environment, Theorem 5.10 states that any interim feasible mechanism can be implemented by the greedy algorithm on a stochastic ordering on the agents' types. Below we give an alternative characterization from Alaei et al. [4] that can be optimized over via a quadratic (in the total number of agent types) linear program.

Consider the following simple token passing procedure. A dummy agent starts with the token. Considering the agents in any fixed order, the next agent steals the token from the current token holder with a fixed probability that is a function of both agents' types and identities. Finally, after the last (nth) agent has had a chance to steal the token, the current holder of the token is given the item. Obviously, any fixed transition probabilities encode an interim feasible allocation profile; non-obviously, any interim feasible allocation profile can be thus encoded. The probabilities in the token passing procedure are given by a quadratic (in the total number of types) variables; and the token passing process can be expressed as a linear program in these variables. We conclude with the following theorem and corollary, the proofs of which follow from a technical max-flow-min-cut-style argument for a generalized formulation of network flow that corresponds to the token passing procedure.

Theorem 5.12 (Alaei et al. [4]). For any independent distribution on type profiles and any single-item environment, any interim feasible allocation profile can be ex post implemented by a token passing procedure.

Corollary 5.13. For single-item environments, interim feasibility can be (a) checked, (b) optimized over, and (c) implemented via the solution to a quadratic-sized program with linear constraints.

Notice that for our example problems of unit-demand preferences or single-dimensional budgeted preferences the single-agent problems can also be expressed with a quadratic-sized (in the size of each agent's type space) linear program; these single-agent programs can be easily combined with interim feasibility program for the token passing procedure to give a single linear program with a quadratic number of variables (in the total number of types). Such a program can be easily and effectively solved.

5.4 Multi-agent Optimization with Multi-dimensional Externalities

We now consider the case where agents have multi-dimensional preferences and the way an agent is served imposes multi-dimensional externality on the other agents via the feasibility constraint. We will work up to providing a solution to the paradigmatic example of multi-dimensional matching environments. In the multi-dimensional matching environment, there are n unit-demand agents and m unit-supply items. The environment exhibits a multi-dimensional externality because when an item j is assigned to agent i then it cannot be assigned to another agent $i^\dagger \neq i$ but other items $j^\dagger \neq j$ can. The approach in this section requires an agent's utility to linearly separate across distinct services in which she is interested.

5.4.1 Multi-item Auctions for Agents with Additive Preferences

We first consider a relaxation of the problem where the agents' preferences are additive, not unit demand. Specifically, multiple items can be assigned to each agent and the agent's value for a bundle of items is the sum of her values for each of the items in the bundle. In particular the unconstrained feasibility is relaxed from (SAP.UD) to

$$x_j(t) \leq 1, \qquad \forall t \in T, j \in \{1, \ldots, m\}. \tag{SAP.A}$$

Cai et al. [22] make the following two observations which enable the problem to be easily solved using the methodology already discussed. First, as each agent's utility is linear then we can express the multi-dimensional optimization problem in terms of the profile of allocation rules corresponding to the items. For single-item allocation constraints $\hat{y}_1, \ldots, \hat{y}_m$ (in the quantile space of each coordinate) define the single-agent revenue operator $\mathbf{Rev}[\hat{y}_1, \ldots, \hat{y}_m]$ as optimizing the single-agent program (SAP) with the interim constraint,

$$\sum\nolimits_{t \in S} x_j(t) f(t) \leq \hat{Y}_j(f(S)), \quad \forall S \subseteq T, j \in \{1, \ldots, m\}. \tag{SAP.IF.A}$$

Second, the multi-agent problem requires resolving interim feasibility for each of the items but this resolution can be performed independently for each item. In particular, a profile of $n \times m$ interim allocation constraints are interim feasible if and only if for each item j the n corresponding interim allocation constraints are interim feasible. This single-item interim feasibility is characterized by Theorem 5.7, can be optimized over by a quadratic sized linear program by Corollary 5.13, and implemented by simple token-passing procedure by Theorem 5.12.

Combining m of these programs with the n quadratic sized linear programs for optimizing the single-agent problems; there is one cubic-sized program for solving the multi-item optimal mechanism design problem for agents with additive preferences.

Theorem 5.14 (combining Cai et al. [22] and Alaei et al. [4]). For multi-item environments with m items and n agents with additive preferences, the optimal auction can be solved for via a linear program of size $m \cdot (\sum_i |T_i|)^2$.

5.4.2 Multi-item Auctions for Agents with Unit-Demand Preferences

We now return to our original multi-item unit-demand unit-supply matching environment, i.e., relative to the previous subsection, we add back the unit-demand constraint (SAP.UD) to each single-agent problem, repeated here for convenience,

$$\sum\nolimits_j x_j(t) \leq 1, \qquad \forall t \in T. \qquad \text{(SAP.UD)}$$

In the context of the previous section, we are no longer free to consider the interim feasibility of each item separately.

The problem is therefore one of determining whether a set of $n \times m$ interim allocation rules, one for each agent-item pair, are interim feasible given the ex post feasibility constraint that requires the assignment of items to agents to be a matching. In fact, the results of Section 5.3 have already solved this problem. Theorem 5.9 implies that the ex post mechanism is given by optimizing a stochastic virtual surplus; Theorem 5.11 implies that the optimization necessary to find the appropriate distribution over virtual values is polynomial time tractable (by reduction to the weighted bipartite matching problem). We can apply these theorems as they are to the multi-dimensional matching problem by pretending that we have *nm representative agents* corresponding to each edge in the matching. Recall that interim feasibility does not require independence of agent types; however, the incentive constraints do. The values that an agent has for distinct items may be correlated; however, the values that distinct agents have must be independent.

Corollary 5.15 (Cai et al. [23]). In unit-demand unit-supply matching environments for agents with independent types, the optimal mechanism is a stochastic virtual surplus maximizer and the virtual value functions can be found in polynomial time.

We summarize by describing the general environments where the approach presented herein can be applied. Optimizing concave (e.g., from the revenue operator $\mathbf{Rev}[\cdot]$) functions over interim feasible allocation constraints is tractable whenever optimizing a linear function over ex post feasible allocations is tractable. Independence of the agent types is not required for feasibility concerns. For the interim approach to mechanism design to be valid for optimizing over Bayes-Nash equilibrium, however, the types of the agents must be independent. This is because optimizing revenue for a type drawn from a distribution subject to an interim allocation constraint requires that the interim constraint not be correlated with the type. The preferences of the agents can be multi-dimensional and non-linear, however, they must linearly separate across distinct services with respect to the feasibility constraint. Here are two positive examples and one negative example: (separable) an agent who only desires a single service (with respect to the feasibility constraint) as in Section 5.3; (separable) an agent who desires several services, has additive values across these services, and a budget representing her total ability to pay; and (non-separable) an agent with additive wealth across services but a concave utility function that maps wealth to her overall utility.

6

Approximation for Multi-dimensional and Non-linear Preferences

As we saw in the preceding section, optimal mechanisms for multi-dimensional and non-linear preferences are generally quite complex. In this section we look at simple mechanisms that are approximately optimal. Our goal will be to show that mechanisms that were good in revenue-linear settings, e.g., like the marginal revenue mechanism (Section 3.3 and Section 5.2) or posted pricings (Section 4.2) continue to be good in non-revenue-linear settings.

Prior-independent approximation is also possible for agents with multi-dimensional and non-linear preferences (cf. Section 4.3). Devanur et al. [39] and Roughgarden et al. [84] give prior-independent approximations for the multi-item matching problem of Section 5.4. Fu et al. [47] discuss prior-independent approximation for risk-averse agents (a specific form of non-linear preference), and Devanur et al. [37] discuss prior-independent approximation for agents with budgets. These results will not be further discussed in this survey.

6.1 Single-agent Approximation

Our first order of business will be to address the complexity of the single-agent problem. As described in Section 5.1, many single-agent problems can be solved by a linear program that is quadratic in the size of the type space. For a number of multi-dimensional problems, however, the size of the type space is exponential in the dimension of the type space. A paradigmatic example is the unit-demand preference over a set of items distributed according to a product distribution.

In the special case of a unit-demand agent with independently distributed values for distinct items, it is not at all clear whether anything interesting can be observed from the linear program for solving the unconstrained unit-demand single-agent problem, i.e., (SAP) with (SAP.UD). Therefore, the first issue we encounter in an attempt to find an approximation for this single-agent problem is that we really do not understand the optimal mechanism very well. The usual approach to such a situation is to identify an upper bound that is analytically tractable.

As discussed in Section 3.2 the taxation principle suggests that any mechanism is given by a menu over, possibly randomized, outcomes. For a unit-demand agent we refer to these menus as *lottery pricings* or, in the special case where the outcomes are all deterministic, as *item pricings*. We desire, then, an upper bound on the revenue of the optimal lottery pricing.

Consider the following thought experiment. We have a seller facing a single unit-demand agent to buy one of m items. Consider the *representative environment* where this unit-demand agent is replaced with m single-dimensional representatives with the unit-demand preference reinterpreted as a feasibility constraint (SAP.UD). Would a seller prefer to be in the original environment or in the representative environment; i.e., in which environment can the seller obtain a higher revenue? Intuition suggests that the seller should prefer the representative environment as competition between representatives should produce a higher revenue. The following theorem shows that this intuition is mostly correct.

Theorem 6.1 (Chawla et al. [29]). For a unit-demand agent with independent values, the revenue of the optimal auction for the representative environment is at least the revenue for the optimal item pricing in the original unit-demand environment.

Proof. It is sufficient to show that for any item pricing $(\hat{t}_1, \ldots, \hat{t}_m)$ for the original environment there is an auction for the representative environment that obtains at least the item pricing's revenue. Notice that the allocation rule of the item pricing is to serve the agent the item $j^\star$ that maximizes $t_{j^\star} - \hat{t}_{j^\star}$ if non-negative, otherwise, nothing. Holding other values t_j for $j \neq j^\star$ fixed, this allocation rule is monotone in each value t_j. Therefore, Theorem 2.2 implies that this allocation rule corresponds to a dominant strategy mechanism in the representative environment. Importantly, in this mechanism, the representative who wins is the one who corresponds to the item that the original unit-demand agent selects. The revenue in the original environment is simply $\hat{t}_{j^\star}$; the revenue in the representative environment is the minimum bid that representative $j^\star$ must make to win. This minimum bid is at least $\hat{t}_{j^\star}$. Therefore, the auction revenue exceeds the item-pricing revenue. □

While revenue of the optimal lottery pricing may exceed the revenue of the optimal item pricing, it is nonetheless interesting to consider whether it is easy to approximate the optimal item pricing (e.g., via an item pricing that is easy to find). Using the optimal representative revenue as an upper bound on the optimal item pricing revenue, we can easily attain such a bound. In fact, the revenue of an item pricing is lower bounded by the revenue we would obtain if the agent buys the cheapest priced item instead of her utility maximizing item, i.e., when her values for several items are above their corresponding prices, if instead of choosing to maximize $t_j - \hat{t}_j$ she chose to minimize $\hat{t}_j$. We discussed the factor by which oblivious posted pricings (i.e., tie-breaking by minimum price) approximate optimal auctions in Subsection 4.2.1; the following is a corollary of Theorem 4.4 and Theorem 6.1.

Corollary 6.2 (Chawla et al. [30]). For a unit-demand agent with independent values, a uniform virtual item pricing is a two approximation to the optimal item-pricing revenue (and the optimal representative revenue).

To address the more general question of approximating the optimal lottery pricing, we will show that twice the optimal representative revenue upper bounds the optimal lottery-pricing revenue. Before doing so, consider the following brief discussion of why our intuition about the increased competition of the single-dimensional environment is not entirely correct. The mechanism designer when facing an agent with a multi-dimensional type that is independent in each dimension can take advantage of concentration bounds related to the law of large numbers. For instance, consider the lottery where the agent is awarded an item uniformly at random, i.e., the probability of obtaining any item is $1/m$. The value that the agent has for such a lottery is an $1/m$ fraction of the sum of her values for the individual items. By the independence of her values, the sum of her values is a concentrated random variable; for instance, if she is offered a price that is equal to its expected value less a standard deviation or two, then she will buy with high probability and pay nearly her full value. An auctioneer facing this agent can potentially link independent agent decisions (i.e., actions that the agent takes for each item) and it is in the auctioneer's interest to do so. See Armstrong [6] for a discussion of grand-bundle pricing for additive preferences and Jackson and Sonnenschein [65] for a general discussion of linking mechanisms. The following theorem upper bounds the benefit from linking for independent unit-demand preferences.

Theorem 6.3 (Chawla et al. [32]). For a unit-demand agent with independent values, the revenue of the optimal auction for the representative environment is at least half the revenue for the optimal lottery pricing in the original unit-demand environment.

Proof. A lottery pricing is a set of lotteries with prices where the agent, for each type she may have, selects her utility-maximizing lottery. It is convenient to index these lotteries by the agent type for which it is preferred. I.e., the lottery that type t would prefer is $(x_1(t),\ldots,x_m(t))$ at price $p(t)$ as defined for the unconstrained single-agent program (SAP) of Subsection 5.1.1. Recall that $\sum_j x_j(t) \leq 1$ as required by (SAP.UD).

To prove the theorem we will give two auctions for the representative environment and show that the sum of these auctions' revenues upper bounds the optimal unit-demand revenue. Of course, optimal representative revenue upper bounds each of these auctions revenue.

The *lottery mimicking auction* considers the profile of values $t = (t_1, \ldots, t_m)$ of the representatives and the lottery $(x_1(t), \ldots, x_m(t))$ that would have been selected by the unit-demand agent with type t. It serves the highest-valued representative $j^\star = \mathrm{argmax}_j\, t_j$ with probability $x_{j^\star}(t)$ and charges her $p(t) - \sum_{j \neq j^\star} t_j x_j(t) + \mu$ where μ is the utility of the unit-demand player with type $t^\dagger$ which is equal to t except with the highest coordinate value lowered to be equal to the second-highest coordinate value. Notice that the utility of the winning representative $j^\star$ in this auction is exactly the same as the unit-demand agent with an additive normalization term μ that is a function only of the values of the other representatives. The utility of the unit-demand agent is monotone in her value for each item. The utility of representative $j^\star$ is then equal to the change in utility of the unit-demand agent as she changes her value for the $j^\star$ item from the second-highest value in $\{t_1, \ldots, t_m\}$ to $t_{j^\star}$. This utility is non-negative because $j^\star$ was the highest valued representative. We conclude that this auction is incentive compatible and has revenue at least $p(t) - \sum_{j \neq j^\star} t_j x_j(t)$ on valuation profile t where $j^\star$ is the highest valued representative. For a given profile of representatives call the second term in the winning representative's payment the *deficit* of the lottery mimicking auction.

The motivation for the next auction is that we want to obtain back the deficit lost by the lottery mimicking auction. This is easy, as $\sum_{j \neq j^\star} x_j(t) \leq 1$ and t_j for $j \neq j^\star$ is at most the second highest value. Therefore, in this instance, the second-price auction obtains a higher revenue than the deficit of the lottery mimicking auction. Taking the expectation over all valuation profiles, the expected deficit is at most the expected revenue of the second-price auction.

We have given two incentive compatible auctions for the representative environment with combined expected revenue exceeding the revenue of the lottery pricing. Therefore, twice the optimal representative revenue is at least the optimal unit-demand revenue. $\square$

Corollary 6.4 (Chawla et al. [32]). For a unit-demand agent with independent values, a uniform virtual item pricing is a four approximation to the optimal lottery pricing.

There is a natural computational question of whether the optimal mechanism can be computed in time polynomial in the number of items. For independent distributions, Cai and Daskalakis [21] give a partial answer to this question by showing that the optimal item pricing can be approximated arbitrarily closely by a computationally efficient dynamic program.[1] The question remains, however, as to whether the best lottery pricing can be approximated arbitrarily closely.

6.2 Multi-agent Approximation in Service Constrained Environments

We now consider approximation of the optimal mechanism for multi-dimensional and non-linear agent preferences in service constrained environments (Definition 5.1). Recall from Section 5.2 that when the single-agent problems are revenue linear, that there is a relatively simple mechanism that is optimal, namely the marginal revenue mechanism (Definition 5.3). Our goals here are to generalize this simple mechanism to non-revenue-linear agents and show that it remains a good approximation to the optimal mechanism.

6.2.1 Approximation via the Marginal Revenue Mechanism

We outline below two general methods for showing that the marginal revenue mechanism is a good approximation to the optimal mechanism. The first is by using properties of the feasibility constraint imposed by the service-constrained environment and the second is by showing that the single-agent problem is approximately linear.

The feasibility-based approach is to show that the optimal ex ante feasible mechanism (which is a marginal revenue mechanism by

[1] Interestingly, Briest [17] show that if we restrict to item prices but the distribution is correlated, then non-trivial computationally efficient approximation algorithms do not exist under reasonable complexity theoretic assumptions.

definition) does not violate interim feasibility too much and that there is a natural way to address its violations of feasibility and construct from it a mechanism that is ex post feasible. In fact, we already saw this approach in subsection 4.2.2. If the optimal ex ante feasible mechanism serves the agents with probabilities $\hat{\boldsymbol{q}}$, its revenue is $\sum_i R_i(\hat{q}_i)$. The sequential posted pricing approach resolves ex post feasibility with the first-come-first-served principle. If an agent arrives and would be served by the mechanism but it is infeasible to do so given the agents previously committed to, then this agent is not served. If we consider the agents sorted by "bang per buck", i.e., $R_i(\hat{q}_i)/\hat{q}_i$, then for single-item environments Corollary 4.7 shows that there is a mechanism based on the ex ante optimal mechanisms for each agent that is an $\frac{e}{e-1}$ approximation (and the same bound extends to matroid environments). The marginal revenue mechanism, of course, is the optimal single-agent mechanism based on the ex ante optimal mechanisms of each agent; the following corollary is immediate.

Corollary 6.5 (Alaei et al. [5]). For matroid service-constrained environments and agents with multi-dimensional non-linear preferences and independent types, the marginal revenue mechanism is an $\frac{e}{e-1}$ approximation to the optimal mechanism.

Alaei et al. [5] generalize the result of this corollary to show that in downward-closed service-constrained environments on n agents that the marginal revenue mechanism is an $O(\log n)$ approximation.

We now turn to showing the marginal revenue mechanism is close to optimal when the optimal revenue (given by $\mathbf{Rev}[\cdot]$) is close to linear.

Proposition 6.6 (Alaei et al. [5]). For downward-closed service-constrained environments and agents with independent types and multi-dimensional non-linear preferences that satisfy $\mathbf{MargRev}[\hat{y}] \geq \frac{1}{\beta}\mathbf{Rev}[\hat{y}]$ for all $\hat{y}$, the marginal revenue mechanism is a β approximation to the optimal mechanism.

Proof. Consider the profile of interim allocation rules $\boldsymbol{y}$ of the optimal mechanism. By the assumption of the proposition, the marginal revenue for this profile is within a β fraction of its (optimal) revenue. The marginal revenue mechanism, by definition, optimizes marginal revenue and therefore can only do better. □

We now instantiate the above proposition for agents with unit-demand values drawn from a product distribution. Consider any of the agents in the mechanism. It suffices to find a linear upper bound on the optimal single-agent revenue (as a function of the interim constraint $\hat{y}$) and show that for any ex ante constraint $\hat{q}$, that optimal single-agent mechanism approximates this upper bound.

For the first part, consider an allocation constraint $\hat{y}$ and consider the representative environment (Section 6.1) where the multi-dimensional agent is replaced with single-dimensional representatives. The optimal revenue for the representative environment is easy to describe. Consider the distribution of the maximum virtual value and serve the agent with the maximum virtual value with the probability specified by $\hat{y}(\cdot)$ for the quantile corresponding to this virtual value. By Theorem 3.8 this mechanism is optimal and by Corollary 3.7 its revenue equals its virtual surplus and therefore it is linear. It is relatively straightforward to generalize Theorem 6.1, which shows that twice the representative revenue upper bounds the optimal unit-demand revenue for unconstrained problems, to single-agent environments with interim constraints.

Theorem 6.7 (Alaei et al. [5]). For a unit-demand agent with independent values and any allocation constraint $\hat{y}$, the revenue of the optimal auction for the representative environment is at least half the revenue for the optimal lottery pricing in the original unit-demand environment; moreover, its revenue is linear in $\hat{y}$.

For the second part, i.e., showing that for every ex ante constraint $\hat{q}$, that the $\hat{q}$ optimal mechanism approximates the upper bound for the ex ante constraint (recall, the ex ante constraint $\hat{q}$ corresponds to the allocation constraint $\hat{y}^{\hat{q}}(\cdot)$ that is a reverse step function from 1 to 0 at $\hat{q}$). In fact, just as we obtained a four approximation to the unconstrained upper bound via the prophet inequality (Theorem 4.3), there is a straightforward adaptation of the prophet inequality to the case where both the prophet and the gambler have an ex ante constraint $\hat{q}$. In this generalization the $\hat{q} = 1$ case (i.e., the original prophet inequality) gives the worst approximation bound of two. As a consequence, via the same argument as Corollary 6.2, with an ex ante allocation constraint a uniform virtual price gives a good approximation to the optimal representative revenue.

Theorem 6.8 (Alaei et al. [5]). For a unit-demand agent with independent values and ex ante allocation constraint $\hat{q}$, a uniform virtual item pricing is a two approximation to the optimal representative revenue.

We conclude that a unit-demand agent with independent values is approximately linear and therefore the marginal revenue mechanism for a collection of such agents is approximately optimal. Notice that for matroid environments, the $\frac{e}{e-1}$ bound via the feasibility constraint is better than the bound we get via the unit-demand single-agent problem; however, for downward-closed environments where the known bound implied by feasibility is logarithmic in the number of agents [5], this approximate linearity gives an improvement.

Corollary 6.9 (Alaei et al. [5]). For unit-demand agents with independent values in a downward-closed service-constrained environment, the marginal revenue mechanism is a four approximation to the optimal mechanism.

The result implied by the statement of Corollary 6.9 is a bit problematic when the agents' type spaces are high dimensional (cf. Section 6.1); for instance, the computational tractability of even the unconstrained single-agent problem is unknown. We can in fact, use the marginal revenue framework with approximations to the ex ante constrained mechanism. Theorem 6.8 shows that setting a uniform virtual price for which an item is bought with probability $\hat{q}$ is a four approximation to the $\hat{q}$ optimal mechanism. Denote by $\tilde{P}(\hat{q})$ the revenue of this approximation mechanism as a function of $\hat{q}$. Denote by $\tilde{R}(\hat{q})$ the smallest concave function that upper bounds $\tilde{P}(\hat{q})$. Refer to $\tilde{R}(\cdot)$ as the *pseudo-revenue curve*. Notice that uniform virtual pricings induce an ordering on types (cf. Definition 5.2) and therefore the *marginal pseudo-revenue mechanism* following Definition 5.3 is well defined. This mechanism, by the same argument as Corollary 6.9, is a four approximation to the optimal mechanism and it is computationally simple.

For the special case of i.i.d. agents and a single-item service-constrained environment the outcome of the marginal pseudo-revenue mechanism is easy to describe. The agent with the overall highest virtual value (for any item) wins and buys her utility maximizing (i.e., value minus price) item under the uniform virtual prices corresponding to the overall highest virtual value of the other agents or a virtual reserve price, whichever is higher.

6.2.2 Implementation of the Marginal Revenue Mechanism

We describe at a high level two approaches for implementing the marginal revenue mechanism without revenue linearity. In particular, the main obstacle that we must overcome is the lack of an implicit ordering on types.

Theorem 6.10 (Alaei et al. [5]). For any service constrained environment and agents with independent types, the marginal revenue mechanism is implementable (by a polynomial time reduction to weighted optimization and the ex ante optimal single-agent problems).

We saw in Section 5.2 that revenue linearity (via orderability) implies that the marginal revenue mechanism is easy to implement. Above, it is evident that the marginal pseudo-revenue mechanism is easy to implement (this is because the ex ante pseudo mechanisms correspond to full lotteries, i.e., the probability that a type is served is either one or zero; cf. Corollary 5.4). The challenge then of implementing the marginal revenue mechanism is in dealing with partial lotteries, i.e., when a type is probabilistically allocated.

When the ex ante optimal mechanisms (i.e., the mechanisms that define the revenue curve) for an agent satisfy a natural monotonicity property, specifically, for any type of the agent, the agent's probability of service in a $\hat{q}$ optimal mechanism is monotone non-decreasing in $\hat{q}$, there is a simple method for implementing the marginal revenue mechanism. At high level we treat this aforementioned monotone function as a distribution function and draw the agent's quantile from it. The remainder is similar to the original marginal revenue mechanism for orderable agents (Definition 5.3); the agents served are the ones that maximize cumulative marginal revenue. This approach gives a dominant strategy mechanism. As an example, this monotonicity property is satisfied for single-dimensional budgeted agents (by Theorem 5.1; see Alaei et al. [5] for further details).

Alaei et al. [5] also give brute-force approach for defining the marginal revenue mechanism in general. Notice that for a given profile of revenue curves $\boldsymbol{R}$ and a profile of quantiles $\boldsymbol{q}$ with each quantile q_i in the profile drawn uniformly at random, the marginal revenue mechanism induces an profile of interim allocation constraints $\hat{\boldsymbol{y}}$ which, by definition, is feasible. Moreover, considering a single agent, the interim mechanism that attains the cumulative marginal revenue for allocation constraint $\hat{y}$ is given simply by the definition of cumulative marginal revenue as the appropriate convex combination of revenues of the ex ante optimal mechanisms. The desired interim mechanism is the same convex combination over ex ante optimal mechanisms. This interim mechanism induces a mapping of types to quantiles (by ordering types by decreasing service probability of the mechanism). We cannot just use this ordering; however, the allocation rule y of the marginal revenue mechanism for $\hat{y}$ satisfies $y \preceq \hat{y}$ (recall Definition 3.3, Section 3.2) but may not be equal to $\hat{y}$. Instead we need to calculate the stationary transformation σ that satisfies $y(q) = \mathbf{E}_{\sigma}[\hat{y}(\sigma(q))]$ for all q and use these to calculate marginal revenues.[2] The *general marginal revenue mechanism*

[2] This is a continuous version of majorization where such a transformation is given by a doubly-stochastic matrix. Such a doubly stochastic matrix can be calculated via techniques of Hardy et al. [57]; however, for our purposes it is sufficient to be able to sample from σ. Alaei et al. [5] give a method for sampling from σ efficiently.

then serves the feasible set of agents S to maximize $\sum_i R'_i(\sigma_i(q_i))$. This mechanism is Bayesian incentive compatible.

6.3 Multi-agent Approximation with Multi-dimensional Externalities

In this section we will generalize the approach of Section 6.1 to give a simple approximation mechanism for the n agent m item matching environment with independent values. It is straightforward to generalize Theorem 6.1 to show that the optimal revenue in the representative environment, i.e., with each of the n unit-demand agents replaced by m single-dimensional representatives, is an upper bound on the optimal deterministic dominant-strategy incentive-compatible auction revenue. The key observation that enables this generalization is that a deterministic dominant-strategy incentive-compatible auction, via the taxation principle, looks to each agent, for the reports of the other agents fixed, like an item pricing.

Theorem 6.11 (Chawla et al. [30]). For unit-demand agents with independent values, the revenue of the optimal auction for the representative environment upper bounds the revenue for the optimal deterministic dominant-strategy incentive-compatible auction for the original unit-demand environment.

We will now show that in the single-dimensional representative environment, that there is an oblivious posted pricing (cf. Subsection 4.2.1) that is a good approximation to the optimal mechanism. The same result then holds for original unit-demand environment. In the environment with nm representatives, the oblivious result implies that if these representatives arrive in any order, and are offered the posted prices as take-it-or-leave-it while-supplies-last, that the approximation bound is obtained. In comparison, in the original unit-demand environment, the n agents can arrive in any order and when offered prices over the items remaining, each chooses to maximize value minus price. This induces an ordering over agent-item pairs that is a valid ordering for the representative environment, and, consequently, the representative environment revenue for the worst-case order can be no better. The proof of the following theorem is deferred to the end of this section.

Theorem 6.12 (adapted from Chawla et al. [30]). For single-dimensional bipartite matching environments (where agents are edges and feasible outcomes are matchings), for any independent distribution of values, there is an oblivious posted pricing that is a nine approximation to the optimal revenue.

Corollary 6.13 (adapted from Chawla et al. [30]). For unit-demand unit-supply matching environments with independently distributed values, there is an oblivious posted pricing that is a nine approximation to the optimal deterministic dominant-strategy incentive-compatible auction.

Of course, we would prefer our auction to approximate the optimal Bayesian-incentive-compatible potentially-randomized auction. Chawla et al. [32] give an analog of Theorem 6.3 that shows that the optimal (randomized) multi-dimensional matching revenue is at most five times that of the single dimensional representative environment.

Theorem 6.14 (adapted from Chawla et al. [32]). For unit-demand unit-supply matching environments with independently distributed values, there is a oblivious posted pricing that is a 45 approximation to the optimal auction.

The construction above can be viewed as an instance of a general methodology that reduces unit-demand pricing problems to an analogous single-dimensional pricing problem. The final step, i.e., the instantiation of the reduction, is the proof of a theorem like Theorem 6.12 that shows that in the single-dimensional representative environment oblivious posted pricing is a good approximation to the optimal auction. The explicit reduction can be found in Chawla et al. [30].

The high-level sketch of the proof of Theorem 6.12 is the following. We start with the ex ante probabilities of service of the optimal representative mechanism; we decrease these probabilities by a constant factor and argue that this decrease implies that, given the independence in the representatives' values, for any representative there is a constant probability that she can be served if she arrives last. By the convexity of revenue curves, the decrease in allocation probability by a constant decreases expected revenue by at most the same constant. Therefore, if these points on the revenue curve correspond to price postings then posting these prices gives a constant approximation.

Of course, the revenue curve of an agent may not be equal to the price-posting revenue curve. Therefore, we may not be able to implement the desired point on a given revenue curve via a deterministic price posting. The lemma below generalizes a result of Devanur and Hartline [38] with $\beta = 1$, and shows that we can always find a posted price that approximates the desired point on the revenue curve.

Lemma 6.15. For any single-dimensional agent, any ex ante probability $\hat{q}$, and any factor $\beta > 1$, there is a posted price $\hat{v}^{\dagger}$ (corresponding to $\hat{q}^{\dagger} \leq \hat{q}/\beta$) for which the price-posting revenue at $\hat{q}^{\dagger}$ is a $\beta + 1$ approximation to the original revenue at $\hat{q}$, i.e., $P(\hat{q}^{\dagger}) \geq R(\hat{q})/(\beta + 1)$.

Proof. We would like to offer a posted price according to $\hat{q}/\beta$. For regular distributions (Definition 3.7) where price-posting and optimal revenue are equal, this price gives a β approximation as $P(\hat{q}/\beta) = R(\hat{q}/\beta) \geq R(\hat{q})/\beta$ (by convexity of $R(\cdot)$). The challenge then is irregular distributions.

Consider $\hat{q}^{\ddagger}$, the highest quantile below $\hat{q}$ for which the price-posting revenue and optimal revenue are equal, i.e., $P(\hat{q}^{\ddagger}) = R(\hat{q}^{\ddagger})$, and let $\hat{v}^{\ddagger}$ be its corresponding price. If $\hat{q}/\beta \leq \hat{q}^{\ddagger}$ then again the price corresponding to $\hat{q}/\beta$ is good; the price posted by $\hat{q}/\beta$ is at least $\hat{v}^{\ddagger}$ so $P(\hat{q}/\beta) \geq \hat{v}^{\ddagger}\hat{q}/\beta \geq R(\hat{q})/\beta$. The challenge then is the case depicted in Figure 6.1 where $\hat{q}^{\ddagger} < \hat{q}/\beta$.

The $\hat{q}^{\dagger}$ in the statement of the lemma is either $\hat{q}^{\ddagger}$ or $\hat{q}/\beta$ whichever has larger price-posting revenue. Partition the optimal revenue of $\hat{q}$

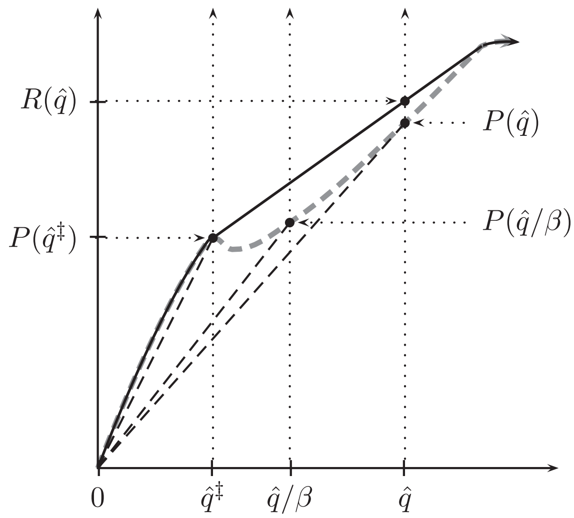

Fig. 6.1 A geometric depiction of the quantities in Lemma 6.15. The revenue curve R (thin, black, solid line) is the convex hull of the price-posting revenue curve P (thick, grey, dashed line). The quantiles $\hat{q}$ and $\hat{q}/\beta$ are within an ironed interval with lower bound $\hat{q}^{\ddagger}$. The dashed lines that connect to points on the price-posting revenue curve have slope equal to the price posted, i.e., $P(\hat{q}) = \hat{v}\hat{q}$ by definition. Importantly for the proof, the price-posting revenue of $\hat{q}/\beta$ is at least a β fraction of the price-posting revenue for $\hat{q}$, i.e., $P(\hat{q}/\beta) \geq P(\hat{q})/\beta$, and the price-posting revenue of $\hat{q}^{\dagger}$ is at least the difference in the optimal revenue and price-posting revenue of $\hat{q}$, i.e., $P(\hat{q}^{\ddagger}) \geq R(\hat{q}) - P(\hat{q})$.

into two pieces, that of the price-posting revenue $P(\hat{q})$ and the remainder $R(\hat{q}) - P(\hat{q})$. The price-posting revenue from $\hat{q}/\beta$ is at least a β fraction of $P(\hat{q}) = \hat{v}\hat{q}$ (as its corresponding price is only higher than $\hat{v}$ and its probability of service is exactly a β factor lower). The slope of the revenue curve at $\hat{q}$, i.e., the marginal revenue $R'(\hat{q})$, is always at most $\hat{v} = V(\hat{q})$. By geometry, then, the price-posting revenue of $\hat{q}^{\ddagger}$ is at least the remainder $R(\hat{q}) - P(\hat{q})$. Therefore, the maximum of the price-posting revenue at $\hat{q}^{\ddagger}$ or $\hat{q}/\beta$ is at least a $1 + \beta$ fraction of the total $R(\hat{q})$. □

We now give the proof of Theorem 6.12.

Proof. As this is entirely a single-dimensional approximation question, we will adopt the notation and terminology of Section 4. However, for convenience we will index the $n \cdot m$ representatives corresponding to edges in the bipartite matching by their ij coordinate. Representative ij's value is v_{ij} drawn independently from distribution F_{ij}, her quantile is q_{ij}, her revenue curve is R_{ij}, etc.

As we did in Subsection 4.2.2 we use the upper bound on the optimal mechanism given by ex ante feasibility; cf. (EAMAP). Let $\hat{\boldsymbol{q}}$ be the revenue-optimal profile of probabilities that is ex ante feasible; i.e., $\hat{\boldsymbol{q}}$ optimizes $\sum_{ij} R_{ij}(\hat{q}_{ij})$ subject to $\sum_i \hat{q}_{ij} \leq 1$ for all j and $\sum_j \hat{q}_{ij} \leq 1$ for all i. Invoking Lemma 6.15 with $\beta = 3$ we obtain a profile of quantiles $\hat{\boldsymbol{q}}^\dagger$ with at most one third the allocation probability, i.e., $\hat{q}^\dagger_{ij} \leq \hat{q}_{ij}/3$, and at least one fourth the revenue, i.e., $P(\hat{q}^\dagger_{ij}) \geq R(\hat{q}_{ij})/4$ for all ij.

By definition, $\hat{\boldsymbol{q}}^\dagger$ satisfies $\sum_{ij} P_{ij}(\hat{q}^\dagger_{ij}) \geq \frac{1}{4} \sum_{ij} R_{ij}(\hat{q}_{ij})$ and $\sum_i \hat{q}_{ij} \leq 1/3$ for all j and $\sum_j \hat{q}_{ij} \leq 1/3$ for all i. Moreover, if ij is feasible at the time ij arrives then the corresponding posted pricing of $\hat{\boldsymbol{v}}^\dagger$ attains revenue $P_{ij}(\hat{q}^\dagger_{ij})$. To get a lower bound on the revenue from ij, imagine ij is the last representative to arrive. By the union bound, the probability of the event that another representative $i^\dagger j$ for $i^\dagger \neq i$ was previously served is at most $1/3$, likewise for the event that another representative $ij^\dagger$ for $j^\dagger \neq j$ was previously served. Independence of these two events implies that the probability neither happens is $(2/3)^2 = 4/9$. Therefore, the revenue we can expect from representative ij under any ordering is at least $4/9 \cdot 1/4 \cdot R_{ij}(\hat{q}_{ij})$. Summing over all representatives gives the desired nine approximation. □

In this section we focused on approximation by oblivious posted pricings. Sequential posted pricings can also be used. Many related approximation bounds based on multi-dimensional sequential posted pricings are given by Chawla et al. [30]. These techniques are refined by Alaei [2] who introduces a multi-dimensional prophet-inequality-like (cf. Subsection 4.2.1) approach and multi-dimensional revenue curves to give very small approximation bounds. For example, his sequential posted pricing improves the matching approximation (with respect to optimal deterministic mechanisms) to four.

7

Computation and Approximation Algorithms

In this section we address computational challenges in mechanism design. To implement a mechanism we must both calculate its outcome, i.e., who wins, and accompanying payments. For many problems of interest, i.e., for combinatorial auctions such as those given by and-preferences, the underlying weighted optimization problem is NP-hard. Unfortunately, however, generic approximation algorithms cannot directly be used in constructing mechanisms because they may fail the monotonicity requirement of Theorem 2.2.

There are a number of approaches to resolving the potential non-monotonicity of generic approximation algorithms. A first approach is to look at classes of approximation algorithms that satisfy monotonicity, or for which there are generic procedures for converting them to satisfy monotonicity. Along these lines Lehmann et al. [70] point out that in single-dimensional environments, algorithms that are greedy by any monotone function of agents values satisfy the (dominant-strategy version of the) monotonicity condition of Theorem 2.2. Recall the setting of single-minded combinatorial auctions (a.k.a., and-preferences, Section 3.1) where there are n agents, m items, and each agent i has a value v_i for obtaining all of the items in demand set S_i. Two agents with demand sets that intersect cannot simultaneously be served; the corresponding optimization problem is *weighted set packing*. The algorithm greedy-by-$\frac{v_i}{\sqrt{|S_i|}}$ is an $O(\sqrt{m})$ approximation which, up to constant factors, is the best possible under standard computational tractability assumptions; the corresponding mechanism is dominant strategy incentive compatible and achieves the same approximation factor [70].

Lavi and Swamy [69] observe that algorithms for solving integer programs (IPs) based on rounding linear programs (LPs) can be made appropriately monotone (even for linear multi-dimensional preferences). The *integrality gap* between an integer program and its linear program relaxation is the maximum ratio between their optimal solutions. Intuitively, we can think of scaling the LP polytope to fit inside the IP polytope and then finding a rounding that is feasible with respect to the original IP and an unbiased estimator of the optimal solution to the scaled linear program. Optimality of the scaled LP solution implies the rounded integral solution is monotone as required for dominant strategy incentive compatibility. For example, via this method Lavi and Swamy [69] generalize the Lehmann et al. [70] $O(\sqrt{m})$ approximation from single- to multi-dimensional agent preferences.

Finally, Dughmi and Roughgarden [43] show that any algorithm that gives a polynomial time approximation scheme for a weighted maximization problem, i.e., for any ϵ there is an algorithm that gives a $1 + \epsilon$ approximation, can be converted into a mechanism. The result exploits an elegant connection between polynomial time approximation schemes and smoothed complexity; the former being approximately optimizing a specific input and the latter being exactly optimizing a perturbed input. In particular, they give a linear perturbation scheme for which optimization of the perturbed preferences of agents for exact feasibility constraints is equivalent to optimization of exact preferences on a perturbed feasibility constraint. As the latter satisfies the multi-dimensional monotonicity requirement for dominant strategy incentive compatibility, so does the mechanism of their construction.

In the rest of this section we give generic approaches for converting any algorithm into a Bayesian incentive compatible mechanism (for maximizing welfare). The first is a simple and economically-intuitive approach for single-dimensional environments that is based on the ironing of non-monotone allocation rules. The second approach generalizes to quasi-linear multi-dimensional environments and is based on optimally solving maximum weight matching. Both approaches can be viewed as a *Bayesian reduction:* Given a good (in terms of expected performance) Bayesian algorithm, a good Bayesian mechanism is constructed.

For both of these sections we assume that an agent's expected value for an outcome of the mechanism (with reports of other agents randomly drawn from the distribution) and any fixed report of hers is given precisely by an oracle. (The oracle can be replaced with statistical estimates from sampling, but then due to errors in sampled estimates the induced mechanism will only be approximately incentive compatible. The gains from misreporting in the resulting mechanism can be bounded in terms of the accuracy of the sampling procedure. There are constructions for removing the oracle and obtaining exact incentive compatibility, which we defer to the papers where they were first presented, see, Hartline and Lucier [60] and Hartline et al. [59].

7.1 Single-dimensional Monotonization via Ironing

Suppose we are given an algorithm $\mathcal{A}$ that obtains a good expected surplus for a given distribution over values $\boldsymbol{F}$, but the induced allocation rules of the algorithm $\boldsymbol{y}$ are not monotone. Meaning, there is some agent i whose interim allocation rule y_i is not monotone. Theorem 2.2 implies that there is no payment rule that can accompany such an allocation rule that would incentivize the agent to truthfully report her preference. Consequentially, the strategic incentives of the agents prevent this algorithm from being implemented.

As in previous sections, understanding and resolving this monotonicity issue is easiest in quantile space rather than value space. Since the transformation between values and quantiles is given by the distribution F (equivalently, by inverse demand curve $V(\cdot)$, Definition 3.1), we without loss assume our algorithm takes its input in quantile space and the agents specify their values also by quantile.

One approach to resolve the monotonicity issue is based on the following observation. Suppose for this agent i with non-monotone allocation rule y_i instead of running the algorithm on the agent's actual quantile, when the quantile q_i falls within given interval $[a,b]$, we redraw a quantile $q_i^\dagger$ from $U[a,b]$ and input the redrawn quantile into the algorithm instead of q_i.[1] Notice that the induced allocation rule $y_i^\dagger$ is, by construction, flat on $[a,b]$ and equal to the expected value of the allocation rule on $[a,b]$, i.e., $y_i^\dagger(q) = \mathbf{E}_{q^\dagger \sim U[a,b]}\big[y_i(q^\dagger)\big] = \frac{Y(b)-Y(a)}{b-a}$ for $q \in [a,b]$. We refer to this process as *ironing* the allocation rule on interval $[a,b]$; it is mathematically analogous to the ironing discussed in Section 3.4. Consider the affect of ironing $[a,b]$ on the cumulative allocation rule Y_i for $\hat{y}_i$. As $Y_i^\dagger(a) = Y_i(a)$, $Y_i^\dagger(b) = Y_i(b)$, and $Y_i^\dagger(q)$ is linear on $q \in [a,b]$; we get $Y_i^\dagger$ from Y_i by connecting the points a and b on Y_i with a line segment.

We make three important observations about this ironing method. First, the new allocation rule $y_i^\dagger$ is weakly monotone on $[a,b]$. Second, as the transformation preserves the uniform distribution of the quantile input into the algorithm, the allocation rules of the other agents are unaffected. Third, the transformation has the potential to improve the surplus of the agent, e.g., if the allocation rule was initially monotone in the opposite direction on the interval and was thereby serving lower valued agents with higher probability than higher-valued agents then flattening it and giving all agents in the interval the same service probability improves surplus. Together these observations suggest that it could be possible to identify intervals in which the allocation rule can be ironed so as to make the agent's entire allocation rule monotone, weakly improve the agent's surplus, and not affect the allocation rule of any other agents.

Not affecting the allocation rule of any other agents is important as it decomposes the problem of making the allocation rules monotone into separate problems for each agent. Furthermore, the algorithm is only guaranteed to have good performance in expectation for the given distribution; if our construction were to change the distribution input into the algorithm then it may not perform well.

[1] Note: if we were implementing this construction in value space and we wanted to flatten the agents allocation rule between on the interval $[a,b]$ (in value space) we would draw a new value $v^\dagger \sim F[a,b]$, the distribution F restricted to the interval $[a,b]$.

We now focus on the single-agent problem and give a construction that closely mirrors that of Section 3. The main idea is to construct allocation rule $\bar{y}$ with cumulative allocation rule $\bar{Y}$ that is constructed from the smallest concave function that upper bounds Y. First, by the above line-segment interpretation of ironing, this allocation rule can be constructed. Second, it is easy to see that it increases surplus (strictly so, given that y was non-monotone to begin with). As $\bar{Y}$ is concave, $\bar{y}$ is monotone as desired.

Proposition 7.1 (Hartline and Lucier [60]). The allocation rule $\bar{y}$ with cumulative allocation rule $\bar{Y}$ equal to the smallest concave upper bound on Y is monotone and has higher expected surplus than y.

Proof. We derive a sequence of inequalities starting from the expected surplus of y expressed in terms of the inverse demand curve $V(q) = F^{-1}(1-q)$ (with derivative $V'(q)$). The expected surplus is,

$$\mathbf{E}_q[V(q) \cdot y(q)] = \Big[V(q) \cdot Y(q)\Big]_0^1 - \mathbf{E}_q\big[V'(q) \cdot Y(q)\big],$$

by integral definition of expectation and integration by parts. Assume the first term on the right-hand side is zero.[2] Observe,

$$\begin{aligned}\mathbf{E}_q\big[-V'(q) \cdot Y(q)\big] &\leq \mathbf{E}_q\big[-V'(q) \cdot \bar{Y}(q)\big] \\ &= \mathbf{E}_q[V(q) \cdot \bar{y}(q)],\end{aligned}$$

which follows because by definition $\bar{Y}(q) \geq Y(q)$ and because "$-V'(q)$" is non-negative (and then again integrating by parts). □

We conclude that the following approach monotonizes the interim allocation rules of any algorithm. For each agent i, identify the intervals in which the cumulative allocation rule Y_i is not equal to its smallest concave upper bound $\bar{Y}_i$. For each agent i, if the agent's quantile falls within such a previously identified interval, redraw the quantile

[2] It is zero is for value distributions supported continuously on interval $[0,h]$ as $V(0) = h$ (importantly: it is finite), $\hat{Y}(0) = 0$, $V(1) = 0$, and $\hat{Y}(1) \leq 1$ (importantly: it is finite).

uniformly from the interval. Run the algorithm on the resulting profile of quantiles and output its outcome. We refer to the algorithm that results from this construction as the *ironed algorithm*.

Theorem 7.2 (Hartline and Lucier [60]). For any algorithm $\mathcal{A}$ the corresponding ironed algorithm is monotone and obtains at least its expected welfare.

A few observations are worthy of note. First, the ironed algorithm only makes a single call to the original algorithm. Suppose the algorithm is the result of an offline optimization and then used repeatedly in an online setting. The ironed intervals can be similarly calculated in the offline stage and then stored in a lookup table. There is essentially, then, no added complexity of running the ironed algorithm in the online stage. Second, this approach can be applied to any objective that separates linearly across the agents, e.g., revenue. The caveat being that we should initially view the algorithm as taking as input the marginal revenues of the agents, and then we replace the quantile to value mapping with the quantile to marginal revenue mapping given by $R_i'(\cdot)$ for each agent i. Notice that, as needed by the proof, the concavity of $R_i(\cdot)$ implies that "$-R_i''(\cdot)$" is non-negative. Third, this reduction, as simple as it is, does not generically extend to non-separable objectives; Chawla et al. [31] give a counter example for the objective of makespan for related machines (i.e., minimizing the maximum completion time of jobs scheduled on strategic machines where the run time of a job j on machine i is given by the product of the job's public length and the machine's private speed). Fourth, this reduction gives a monotone interim allocation rule which is sufficient for Bayesian incentive compatibility as stated by Theorem 2.2; Chawla et al. [31] show that it does not seem to be generalizable to give dominant strategy incentive compatible mechanisms. Fifth and finally, the ironing approach does not seem to have a natural direct generalization to environments with multi-dimensional agent preferences; in the next section we give a similar but general approach.

7.2 Multi-dimensional Monotonization via Matching

Our approach to fixing the Bayesian non-incentive-compatibility of algorithms for agents with quasi-linear, i.e., arbitrary over outcomes but linear in payments, multi-dimensional preferences proceeds similarly to the single-dimensional case. For each agent, we look for resampling transformation σ from types to types that is (a) stationary with respect to the distribution, (b) improves welfare, and (c) for which the composition with the algorithm is Bayesian incentive compatible. As in the single-dimensional case, the restriction to stationary transformations allows us to address the non-incentive-compatibility of each agent independently.

As a warm up, consider an agent with type t drawn uniformly from a finite type space T, i.e., the probability that the agent has any type $t \in T$ is $1/|T|$. The interim *outcome rule* (with other agents' types drawn from their respective distributions) of the algorithm maps each type of this agent to a distribution over outcomes $w(t)$. Non-incentive-compatibility of the algorithm means there is no pricing rule $p(t)$ such that for all $t, t^\dagger \in T$, $u(t, w(t)) - p(t) \geq u(t, w(t^\dagger)) - p(t^\dagger)$. Consider the weighted bipartite graph between types and outcomes with weights between type t and outcome $w(t^\dagger)$ given by the utility function $u(t, w(t^\dagger))$.

Proposition 7.3 (e.g., Rochet [82] or Gul and Stacchetti [54]). An interim outcome rule is incentive compatible if and only if it corresponds to a maximum weight matching in the utility-weighted bipartite graph between agent types and outcomes of the rule. Moreover, the appropriate payments are the dual variables corresponding to each outcome being matched to at most one type in the matching linear program.

The transformation σ to make any non-incentive-compatible algorithm incentive compatible is as follows. With respect to the aforementioned bipartite graph, calculate the maximum weighted matching and set $\sigma(t) = t^\dagger$ if t is matched to $w(t^\dagger)$. Note that (a) the transformation σ is stationary by the fact that it corresponds to a matching and that the distribution on types is uniform, (b) expected welfare is improved as the new expected surplus is a $|T|$ fraction of the maximum matching and the original expected surplus was the same fraction of a non-maximum matching, and (c) the constructed allocation rule $\bar{w}(t) = w(\sigma(t))$ corresponds to a maximum matching and is therefore incentive compatible.

To relax the assumption that the types are uniformly distributed (and from a finite type space) we give a generalization of the construction that preserves two of its implicit properties. First, the outcomes on one side of the matching are not a function of the agent's actual type. Second, if we pick one of these outcomes uniformly at random then the type from which the outcome is obtained is distributed according to the distribution on types.

The construction, which we refer to as the *replica-surrogate-matching transformation*, is as follows. Define a set of m *replica types* that contains the actual agent type and $m-1$ independent draws from the distribution F. Define a set of m *surrogate types* each drawn independently from the distribution F. We will consider the bipartite matching between replica types and surrogate outcomes, i.e., the distribution over outcomes given by the algorithm applied individually to each of the surrogate types (with the types of the other agents drawn from their own distributions as usual). The transformation then will map the real type to the surrogate whose outcome it is matched to in the maximum weighted matching on of the utility-weighted bipartite graph between replica types and surrogate outcomes.

It is easy to see that the constructed algorithm induced by the composition of the replica-surrogate-matching transformation with the original algorithm satisfies two of our three requirements for such a transformation. It is Bayesian incentive compatible and stationary. It does not, however, necessarily preserve surplus. Unlike in the previous case where the outcome of the original algorithm was one of the candidate matchings, because the surrogates are not equal to the replicas, the outcome of the original algorithm is not a candidate outcome for the construction. It should be intuitive, though, that in the limit as the number of types m in the replica surrogate matching approaches infinity, the potential loss in welfare approaches zero.

For finite m we can get a bound on the loss incurred by the replica surrogate matching in terms of a quantity known as the *transportation cost* (see, e.g., Talagrand [90]). The transportation cost for a space T, distribution F, cost metric c, and size m is the average edge cost in the minimum cost matching (according to c) between two sets of size m drawn independently from F.

Transportation cost bounds the performance loss by the replica surrogate matching. One way to get a matching between replica types and surrogate outcomes, is as follows. First, match replicas types to surrogate types to minimize the cost of the matching, where the cost of matching replica type r to surrogate type s is the maximum difference in utilities of the two types for the same outcome, i.e., $\max_w |u(r,w) - u(s,w)|$. Second match surrogate types to surrogate outcomes to maximize the utility of the matching. The surplus of the composition of these two matchings gives a lower bound on the surplus of the optimal replica-surrogate matching. The expected cost of the first matching is the transportation cost for the metric; the expected surplus of the second matching is at least that of the algorithm (as the matching of each surrogate to its own outcome is a candidate matching). As the real type is equally likely to be any of the replicas, the expected loss for the agent is equal to the transportation cost.

Theorem 7.4. The mechanism constructed by applying a replica-surrogate-matching transformation to each agent independently and then running the algorithm on the transformed types (with payments as per Proposition 7.3) is Bayesian incentive compatible and has expected surplus at least that of the algorithm minus the sum of transportation costs of the agents.

Hartline et al. [59] give bounds on the transportation cost for several settings of interest. In particular, for single-dimensional preferences on and values on $[0,1]$, the loss from an agent in the reduction can be reduced to ϵ with a replica surrogate matching with size m that is quadratic in $1/\epsilon$. Unsurprisingly, the best general bounds obtained for multi-dimensional type spaces require a replica surrogate matching of size exponential in the dimension.

A

Mathematical Reference

We provide herein additional discussion of a few basic mathematical objects that play a prominent role in Bayesian mechanism design.

A.1 Submodular Set Functions

Given a subset S of a ground set N consider a real valued set function $g : 2^N \to \mathbb{R}$. Intuitively, *submodularity* corresponds to diminishing returns. Adding an element i to a set increases the value of the set function less than it would have for adding it to a subset.

Definition A.1. A set function g is *submodular* if for $S^\dagger \subset S^\ddagger$ and $i \notin S^\ddagger$,

$$g(S^\dagger \cup \{i\}) - g(S^\dagger) \geq g(S^\ddagger \cup \{i\}) - g(S^\ddagger).$$

A.2 Matroid Set Systems

Matroid set systems capture environments with inherent substitutability. A matroid is a set system that is given by a ground set N and a collection of feasible, a.k.a., *independent*, sets, $\mathcal{I} \subset 2^N$. The defining two properties of a matroid set system $(N,\mathcal{I})$ are,

downward closure: Subsets of independent sets are independent. I.e., if $S \in \mathcal{I}$ and $S^\dagger \subset S$ then $S^\dagger \in \mathcal{I}$.

augmentation: There is always an element in a big independent set that can augment a small independent set such that the augmented set is independent. I.e., if $S^\dagger, S^\ddagger \in \mathcal{I}$ and $\left|S^\dagger\right| < \left|S^\ddagger\right|$ then exists $i \in S^\ddagger \setminus S^\dagger$ such that $S^\dagger \cup \{i\} \in \mathcal{I}$.

The augmentation property implies, importantly, that all maximal independent sets are of the same cardinality. The *matroid rank function* is a set function defined as the largest cardinality of an independent subset of the given set; the rank of the matroid is the rank of the ground set; the matroid rank function is denoted by $\text{rank}(S)$ for $S \subset N$. If the elements of the ground set have weights, then we can similarly define the matroid weighted-rank function as the maximum cumulative weight of an independent subset. Importantly, the matroid rank and weighted-rank functions are submodular (Definition A.1). Therefore, the matroid structure imposes diminishing returns.

Theorem A.1. The matroid rank function is submodular; for any real valued weights, the matroid weighted-rank function is submodular.

The consequence of the matroid structure most important to auction theory is given by the following algorithmic characterization.

Theorem A.2. An set system $(N,\mathcal{I})$ is a matroid if and only if for any weights the greedy-by-weight algorithm selects the independent set with maximum cumulative weight.

There are a number of examples of matroid set systems that are relevant to auction theory.

uniform: The independent sets are those with cardinality at most a given number (e.g., k). The k-uniform matroid corresponds to the feasibility constraint of a k-unit environment.

transversal: The ground set is the vertices of one part of a bipartite graph; independent sets correspond to subsets that can be simultaneously matched to vertices in the other part. Transversal matroids correspond to single-dimensional matching environments (i.e., or-preferences over a set of items).

graphical: The ground set is the set of edges in a graph; independent sets correspond to acyclic subgraphs. For graphical matroids the (optimal) greedy-by-weight algorithm is known as Kruskal's algorithm [67].

A.3 Convex Optimization

Consider maximizing a concave real-valued function h defined on a convex feasible subset $\mathcal{X}$ of the n-dimensional real-valued vector field $\mathbb{R}^n$. Convexity of the feasible set $\mathcal{X}$ is defined by: For any two vectors $\boldsymbol{x}^\dagger$ and $\boldsymbol{x}^\ddagger$ in $\mathcal{X}$, the convex combination $\boldsymbol{x} = \alpha \cdot \boldsymbol{x}^\dagger + (1-\alpha) \cdot \boldsymbol{x}^\ddagger$ for any $\alpha \in [0,1]$ is in $\mathcal{X}$. Concavity of h is defined by: The function value on a convex combination of vectors is at least the convex combination of the function value on each vector, i.e., for any two vectors $\boldsymbol{x}^\dagger$ and $\boldsymbol{x}^\ddagger$ in $\mathcal{X}$ and convex combination $\boldsymbol{x} = \alpha \cdot \boldsymbol{x}^\dagger + (1-\alpha)\boldsymbol{x}^\ddagger$, then $h(\boldsymbol{x}) \geq \alpha \cdot h(\boldsymbol{x}^\dagger) + (1-\alpha) \cdot h(\boldsymbol{x}^\ddagger)$. The mathematical program for convex optimization is,

$$\begin{aligned} \max_{\boldsymbol{x}} \ & h(\boldsymbol{x}) \\ \text{s.t.} \ & \boldsymbol{x} \in \mathcal{X}. \end{aligned} \tag{COP}$$

A.3.1 Structure

For convex optimization, a local optimum is globally optimal, and if the objective function is strictly concave then the global optimum is unique.

If the feasibility constraint of $\mathcal{X}$ is given by a finite number of linear constraints, then it is referred to as a *polytope*. A point in the polytope is a vertex if it is uniquely optimal for some linear objective function, i.e., i.e., h given by weights $\boldsymbol{w} = (w_1, \ldots, w_n)$ with $h(\boldsymbol{x}) = \boldsymbol{w} \cdot \boldsymbol{x} = \sum_i w_i x_i$. Conversely, for any linear objective there is a vertex of the polytope that is optimal. A polytope has a finite set of vertices and is uniquely specified as the convex hull of its vertices. For a general concave objective function h, optima may be on the interior of $\mathcal{X}$. These interior points can be expressed as a convex combination vertices, moreover, by Carathéodory [26], $n+1$ vertices are sufficient.

Theorem A.3. (Carathéodory [26]) Vertices of an n-dimensional polytope correspond to maximal solutions of linear objective; any point in the polytope can be expressed as a convex combination of $n+1$ vertices.

Sometimes the feasibility space $\mathcal{X}$ of a convex optimization exhibits special structure. One such special structure is that of a *polymatroid*. A polytope $\mathcal{X}$ is a polymatroid if there is a submodular function g (see Section A.1) such that for all $\boldsymbol{x} \in \mathcal{X}$,[1]

$$\sum\nolimits_{i \in S} x_i \leq g(S) \qquad\qquad \forall S \subset \{1,\ldots,n\}, \tag{A.1}$$

and $\boldsymbol{x}$ has non-negative coordinates, i.e., $x_i \geq 0$ for all i.

Polymatroids have two important properties with respect to linear objectives. First, optimization of any two linear objectives $\boldsymbol{w}$ and $\boldsymbol{w}^\dagger$ (e.g., corresponding to objective functions of the form $h^{\boldsymbol{w}}(\boldsymbol{x}) = \sum_i w_i x_i$) for which the weights are in the same order, give the same optimal vertex (or vertices if there are multiple optima), i.e., $\boldsymbol{x}^{\boldsymbol{w}} = \boldsymbol{x}^{\boldsymbol{w}^\dagger}$. Moreover, the greedy algorithm is optimal. Start with the all-zeros vector (which is feasible), sort the coordinates by their weights, then in this order increase each coordinate from zero as much as possible (until feasibility prevents further increase).

Proposition A.4. Vertices of a polymatroid correspond to greedy maximization according to an ordering of coordinates of the vector space; any point in the polymatroid can be expressed as a convex combination of vertices.

---

[1] Often the notation $x(S)$ employed as a short-hand for $\sum_{i \in S} x_i$ and then the inequality is written as $x(S) \leq g(S)$.

A.3.2 Optimization

Consider the computational problem of optimizing a linear objective $\boldsymbol{w}$ over a convex space $\mathcal{X}$. Clearly, the way the space $\mathcal{X}$ is expressed is important. For instance, if it is given by a polynomial number of constraints (i.e., facets of the polytope) then it can be solved by linear programming. Suppose instead it is expressed as the convex hull of a polynomial number of points (i.e., the vertices of this polytope is a subset of these points). The number of facets (i.e., linear constraints) of the convex hull of k points is $O(k^{n/2})$; therefore, a direct reduction to linear programming is not tractable (see, e.g., Chazelle [33]). Suppose instead that the polytope is the polymatroid expressed via the weighted rank function of a graphical matroid as per (A.1) and this weighted rank function is given by the graph (which has polynomial size in the number of elements of the ground set, i.e., the edges of the graph). The number of potential facets (i.e., linear constraints) explicitly given by (A.1) is 2^n; again, a direct reduction to linear programming is not tractable (see, e.g., Grötschel et al. [52]). Of course, by Theorem A.2 optimization subject to such a polymatroid constraint is computationally tractable.

The problem of convex optimization (and linear optimization) reduce to a *separation problem*. The separation problem is, given a query point $\boldsymbol{x}$, either conclude that the point is feasible, i.e., $\boldsymbol{x} \in \mathcal{X}$, or find a *separating hyperplane*. A separating hyperplane for a query point $\boldsymbol{x}$ and feasible space $\mathcal{X}$ is given by a direction $\boldsymbol{w}$ (orthogonal to the plane) and a point $\boldsymbol{x}^{\boldsymbol{w}}$ such that the plane through $\boldsymbol{x}^{\boldsymbol{w}}$ in the direction of $\boldsymbol{w}$ (weakly) has the query point $\boldsymbol{x}$ on the $\boldsymbol{w}$ side of the hyperplane and the entire feasible space $\mathcal{X}$ (weakly) on the $-\boldsymbol{w}$ side.

A.3.3 Approximate Computation

Optimal solutions to convex optimization problems may have irrational coordinates. Writing down such a coordinate exactly generally requires infinite precision. To generically solve a convex optimization problem we must therefore give up on exact optimality. The reduction formally stated below shows that if we can solve the separation problem to a given degree of accuracy then we can solve the optimization problem to the same degree of accuracy. The reduction requires the separation oracle to be polynomial time in the logarithm of the degree of accuracy.

The *weak convex optimization problem* for concave objective h, convex feasibility constraint $\mathcal{X} \subset \mathbb{R}^n$, and error tolerance ϵ is the following. Find a point $\boldsymbol{x} \in \mathbb{R}^n$ that is approximately feasible and approximately optimal in time polynomial in n and $\log(1/\epsilon)$. Formally, $\boldsymbol{x}$ is approximately feasible if there exists an $\boldsymbol{x}^\dagger \in \mathcal{X}$ such that $\left|\boldsymbol{x} - \boldsymbol{x}^\dagger\right| \leq \epsilon$; and $\boldsymbol{x}$ is approximately optimal if all $\boldsymbol{x}^\ddagger \in \mathcal{X}$ satisfy $h(\boldsymbol{x}) \geq h(\boldsymbol{x}^\ddagger) - \epsilon$.

The *weak convex separation problem* for convex feasibility constraint $\mathcal{X} \subset \mathbb{R}^n$ and error tolerance ϵ is the following. Given any query point $\boldsymbol{x} \in \mathbb{R}^n$, if $\boldsymbol{x}$ is not approximately feasible find a unit-length direction (a.k.a., a linear objective) $\boldsymbol{w}$ for which $\boldsymbol{x}$ approximately exceeds all feasible points in time polynomial in n and $\log(1/\epsilon)$. Formally, $\boldsymbol{x}$ approximately exceeds all feasible points in direction $\boldsymbol{w}$ if all $\boldsymbol{x}^\ddagger \in \mathcal{X}$ satisfy $\boldsymbol{w} \cdot \boldsymbol{x} > \boldsymbol{w} \cdot \boldsymbol{x}^\ddagger - \epsilon$.

Theorem A.5 (Schrijver [86]). For concave objective h, convex feasibility constraint $\mathcal{X} \subset \mathbb{R}^n$, and error tolerance ϵ, the weak optimization problem reduces in polynomial time in n and $\log(1/\epsilon)$ to the weak separation problem.

Acknowledgments

This survey presents over a decade of recent work on algorithmic aspects of mechanism design in the context of the classical economic theory of Bayesian mechanism design. My perspective on these problems, solutions, and concepts; and thus their presentation herein; has been shaped by a number of collaborations. In chronological order, they are with colleagues Andrew Goldberg, Anna Karlin, Shuchi Chawla, Bobby Kleinberg, Tim Roughgarden, Qiqi Yan, Brendan Lucier, Saeed Alaei, and Nima Haghpanah. Additionally, Saeed Alaei, Hu Fu, Nima Haghpanah, Anna Karlin, Balasubramanian Sivan, Vasilis Syrganis, and Qiqi Yan provided useful comments and suggestions on parts of this survey.

References

[1] S. Agrawal, Y. Ding, A. Saberi, and Y. Ye, "Price of correlations in stochastic optimization," *Operations Research*, vol. 60, no. 1, pp. 150–162, 2012.

[2] S. Alaei, "Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers," in *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science*, pp. 512–521, IEEE, 2011.

[3] S. Alaei, "Mechanism design with general utilities," PhD thesis, University of Maryland, 2012.

[4] S. Alaei, H. Fu, N. Haghpanah, J. Hartline, and A. Malekian, "Bayesian optimal auctions via multi-to single-agent reduction," in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 17–17, 2012. Available at: http://arxiv.org/abs/1203.5099.

[5] S. Alaei, H. Fu, N. Haghpanah, J. Hartline, and A. Malekian, "The simple economics of approximately optimal auctions," arXiv preprint arXiv:1206.3541, 2012.

[6] M. Armstrong, "Price discrimination by a many-product firm," *Review of Economic Studies*, vol. 66, no. 1, pp. 151–168, 1999.

[7] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour, "Reducing mechanism design to algorithm design via machine learning," *Journal of Computer and System Sciences*, vol. 74, no. 8, pp. 1245–1270, 2008.

[8] S. Baliga and R. Vohra, "Market research and market design," *Advances in Theoretical Economics*, vol. 3, 2003.

[9] X. Bei and Z. Huang, "Bayesian incentive compatibility via fractional assignments," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 720–733, SIAM, 2011.

[10] D. Bergemann and S. Morris, "Robust mechanism design," *Econometrica*, vol. 73, pp. 1771–1813, 2005.

[11] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala, "Budget constrained auctions with heterogeneous items," in *Proceedings of the 42nd ACM symposium on Theory of computing*, pp. 379–388, ACM, 2010.

[12] K. Bhawalkar and T. Roughgarden, "Welfare guarantees for combinatorial auctions with item bidding," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 700–709, SIAM, 2011.

[13] S. Bikhchandani, "Auctions of heterogeneous objects," *Games and Economic Behavior*, vol. 26, no. 2, pp. 193–220, 1999.

[14] K. Border, "Implementation of reduced form auctions: A geometric approach," *Econometrica*, vol. 59, no. 4, pp. 1175–1187, 1991.

[15] K. Border, "Reduced form auctions revisited," *Economic Theory*, vol. 31, no. 1, pp. 167–181, 2007.

[16] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis.* Cambridge University Press, 1998.

[17] P. Briest, "Uniform budgets and the envy-free pricing problem," in *Automata, Languages and Programming*, pp. 808–819, Springer, 2008.

[18] P. Briest, S. Chawla, R. Kleinberg, and M. Weinberg, "Pricing randomized allocations," in *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 585–597, Society for Industrial and Applied Mathematics, 2010.

[19] J. Bulow and P. Klemperer, "Auctions versus negotiations," *American Economic Review*, vol. 86, pp. 180–194, 1996.

[20] J. Bulow and J. Roberts, "The simple economics of optimal auctions," *The Journal of Political Economy*, vol. 97, pp. 1060–1090, 1989.

[21] Y. Cai and C. Daskalakis, "Extreme-value theorems for optimal multidimensional pricing," in *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science*, pp. 522–531, IEEE, 2011.

[22] Y. Cai, C. Daskalakis, and M. Weinberg, "An algorithmic characterization of multi-dimensional mechanisms," in *Symposium on Theory of Computation*, pp. 459–478, 2012.

[23] Y. Cai, C. Daskalakis, and M. Weinberg, "Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization," in *Foundations of Computer Science*, pp. 130–139, 2012.

[24] B. Caillaud and J. Robert, "Implementation of the revenue-maximizing auction by an ignorant seller," *Review of Economic Design*, vol. 9, no. 2, pp. 127–143, 2005.

[25] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, and É. Tardos, "On the efficiency of equilibria in generalized second price auctions," arXiv preprint arXiv:1201.6429, 2012.

[26] C. Carathéodory, "Über den Variabilitätsbereich der Fourierschen Konstanten von positiven harmonischen Funktionen," *Rendiconti del Circolo Matematico di Palermo (1884–1940)*, vol. 32, no. 1, pp. 193–217, 1911.

[27] T. Chakraborty, E. Even-Dar, S. Guha, Y. Mansour, and S. Muthukrishnan, "Approximation schemes for sequential posted pricing in multi-unit auctions," *Internet and Network Economics*, pp. 158–169, 2010.

[28] S. Chawla and J. Hartline, "Auctions with unique equilbria," in *Proceedings of the ACM Conference on Electronic Commerce*, 2013.
[29] S. Chawla, J. Hartline, and R. Kleinberg, "Algorithmic pricing via virtual valuations," in *Proceedings of the ACM Conference on Electronic Commerce*, 2007.
[30] S. Chawla, J. Hartline, D. Malec, and B. Sivan, "Sequential posted pricing and multi-parameter mechanism design," in *Proceedings of the 41st ACM Symposium on Theory of Computing*, 2010.
[31] S. Chawla, N. Immorlica, and B. Lucier, "On the limits of black-box reductions in mechanism design," in *Proceedings of the 44th symposium on Theory of Computing*, pp. 435–448, ACM, 2012.
[32] S. Chawla, D. Malec, and B. Sivan, "The power of randomness in bayesian optimal mechanism design," in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 149–158, 2010.
[33] B. Chazelle, "An optimal convex hull algorithm in any fixed dimension," *Discrete & Computational Geometry*, vol. 10, no. 1, pp. 377–409, 1993.
[34] G. Christodoulou, A. Kovács, and M. Schapira, "Bayesian combinatorial auctions," *Automata, Languages and Programming*, pp. 820–832, 2008.
[35] E. H. Clarke, "Multipart Pricing of Public Goods," *Public Choice*, vol. 11, pp. 17–33, 1971.
[36] T. Cover and P. Hart, "Nearest neighbor pattern classification," *Information Theory, IEEE Transactions on*, vol. 13, no. 1, pp. 21–27, 1967.
[37] N. Devanur, B. Ha, and J. Hartline, "Prior-free auctions for budgeted agents," in *Proceedings of the ACM Conference on Electronic Commerce*, 2013.
[38] N. Devanur and J. Hartline, "Limited and online supply and the bayesian foundations of prior-free mechanism design," in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 41–50, 2009.
[39] N. Devanur, J. Hartline, A. Karlin, and T. Nguyen, "Prior-independent multi-parameter mechanism design," *Internet and Network Economics*, pp. 122–133, 2011.
[40] N. Devanur, J. Hartline, and Q. Yan, "Envy freedom and prior-free mechanism design," *CoRR*, abs/1212.3741, 2012.
[41] P. Dhangwatnotai, T. Roughgarden, and Q. Yan, "Revenue maximization with a single sample," in *ECOM10*, 2010.
[42] S. Dobzinski and J. Vondrák, "The computational complexity of truthfulness in combinatorial auctions," in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 405–422, ACM, 2012.
[43] S. Dughmi and T. Roughgarden, "Black-box randomized reductions in algorithmic mechanism design," in *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pp. 775–784, IEEE, 2010.
[44] C. Dwork, A. Goldberg, and M. Naor, "On memory-bound functions for fighting spam," *Advances in Cryptology-Crypto 2003*, pp. 426–444, 2003.
[45] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar, "Rank aggregation methods for the web," in *Proceedings of the 10th international conference on World Wide Web*, pp. 613–622, ACM, 2001.
[46] D. Fain and J. Pedersen, "Sponsored search: A brief history," *Bulletin of the American Society for Information Science and Technology*, vol. 32, no. 2, pp. 12–13, 2006.

[47] H. Fu, J. Hartline, and D. Hoy, "Prior-independent Auctions for Risk-averse Agents," in *Proceedings of the ACM Conference on Electronic Commerce*, 2013.

[48] A. Goldberg and J. Hartline, "Collusion-Resistant Mechanisms for Single-Parameter Agents," in *Proceedings of the 16th ACM Symposium on Discrete Algorithms*, 2005.

[49] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright, "Competitive Auctions," *Games and Economic Behavior*, vol. 55, pp. 242–269, 2006.

[50] A. Goldberg, J. Hartline, and A. Wright, "Competitive auctions and digital goods," in *Proceedings of the 12th ACM Symposium on Discrete Algorithms*, pp. 735–744, ACM/SIAM, 2001.

[51] R. Gomes and K. Sweeney, "Bayes-Nash equilibria of the generalized second-price auction," *Games and Economic Behavior*, 2012.

[52] M. Grötschel, L. Lovász, and A. Schrijver, "The ellipsoid method and its consequences in combinatorial optimization," *Combinatorica*, vol. 1, no. 2, pp. 169–197, 1981.

[53] T. Groves, "Incentives in Teams," *Econometrica*, vol. 41, pp. 617–631, 1973.

[54] F. Gul and E. Stacchetti, "Walrasian equilibrium with gross substitutes," *Journal of Economic Theory*, vol. 87, no. 1, pp. 95–124, 1999.

[55] Z. Gyongyi and H. Garcia-Molina, "Web spam taxonomy," in *First international workshop on adversarial information retrieval on the web (AIRWeb 2005)*, 2005.

[56] N. Haghpanah, "Optimal multi-dimensional pricings," manuscript, 2013.

[57] G. Hardy, J. Littlewood, and G. Pólya, "Some simple inequalities satisfied by convex functions," *Messenger of Math*, vol. 58, pp. 145–152, 1929.

[58] J. Hartline and A. Karlin, "Profit maximization in mechanism design," in *Algorithmic Game Theory*, ch. 13, (N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, eds.), pp. 331–362, Cambridge University Press, 2007.

[59] J. Hartline, R. Kleinberg, and A. Malekian, "Bayesian incentive compatibility via matchings," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 734–747, SIAM, 2011.

[60] J. Hartline and B. Lucier, "Bayesian Algorithmic Mechanism Design," in *Proceedings of the 41st ACM Symposium on Theory of Computing*, 2010.

[61] J. Hartline and T. Roughgarden, "Optimal Mechanism Design and Money Burning," in *Proceedings of the 39th ACM Symposium on Theory of Computing*, 2008.

[62] J. Hartline and T. Roughgarden, "Simple versus Optimal Mechanisms," in *Proceedings of the 10th ACM Conference on Electronic Commerce*, 2009.

[63] A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan, "Non-price equilibria in markets of discrete goods," in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 295–296, New York, NY, USA: ACM, 2011.

[64] D. Huffman, "A method for the construction of minimum-redundancy codes," *Proceedings of the IRE*, vol. 40, no. 9, pp. 1098–1101, 1952.

[65] M. Jackson and H. Sonnenschein, "Overcoming incentive constraints by linking decisions," *Econometrica*, vol. 75, no. 1, pp. 241–257, 2007.

[66] R. Kleinberg and M. Weinberg, "Matroid prophet inequalities," in *Proceedings of the 44th symposium on Theory of Computing*, pp. 123–136, ACM, 2012.

[67] J. Kruskal, "On the shortest spanning subtree of a graph and the traveling salesman problem," in *Proceedings of the American Mathematical Society*, vol. 7, pp. 48–50, 1956.

[68] J.-J. Laffont and J. Robert, "Optimal auction with financially constrained buyers," *Economics Letters*, vol. 52, no. 2, pp. 181–186, 1996.

[69] R. Lavi and C. Swamy, "Truthful and Near-Optimal Mechanism Design via Linear Programming," in *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, 2005.

[70] D. Lehmann, L. O'Callaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *Journal of the ACM (JACM)*, vol. 49, no. 5, pp. 577–602, 2002.

[71] B. Lucier and A. Borodin, "Price of anarchy for greedy auctions," in *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, pp. 537–553, 2010.

[72] B. Lucier, R. Paes Leme, and É. Tardos, "On revenue in the generalized second price auction," in *Proceedings of the 21st international conference on World Wide Web*, pp. 361–370, ACM, 2012.

[73] S. Matthews, "Selling to risk averse buyers with unobservable tastes," *Journal of Economic Theory*, vol. 30, no. 2, pp. 370–400, August 1983.

[74] R. Myerson, "Optimal Auction Design," *Mathematics of Operations Research*, vol. 6, pp. 58–73, 1981.

[75] N. Nisan and A. Ronen, "Algorithmic Mechanism Design," *Games and Economic Behavior*, vol. 35, pp. 166–196, 2001.

[76] J. Oxley, *Matroid Theory.* Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1992.

[77] R. Paes Leme, V. Syrgkanis, and E. Tardos, "The Dining Bidder Problem:a la russe eta la française," *SIGECOM Exchanges*, pp. 25–28, 2012.

[78] R. Paes Leme, V. Syrgkanis, and É. Tardos, "Sequential auctions and externalities," in *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 869–886, SIAM, 2012.

[79] R. Paes Leme and E. Tardos, "Pure and Bayes-Nash price of anarchy for generalized second price auction," in *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 735–744, IEEE Computer Society, 2010.

[80] M. Pai and R. Vohra, "Optimal auctions with financially constrained bidders," Discussion Papers, Northwestern University, August 2008.

[81] J. Rochet, "The taxation principle and multi-time Hamilton-Jacobi equations," *Journal of Mathematical Economics*, vol. 14, no. 2, pp. 113–128, 1985.

[82] J. Rochet, "A necessary and sufficient condition for rationalizability in a quasilinear context," *Journal of Mathematical Economics*, vol. 16, no. 2, pp. 191–200, 1987.

[83] T. Roughgarden, "The price of anarchy in games of incomplete information," in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 862–879, ACM, 2012.

[84] T. Roughgarden, I. Talgam-Cohen, and Q. Yan, "Supply-limiting mechanisms," in *Proceedings the ACM Conference on Electronic Commerce*, pp. 844–861, 2012.

[85] E. Samuel-Cahn, "Comparison of threshold stop rules and maximum for independent nonnegative random variables," *The Annals of Probability*, vol. 12, no. 4, pp. 1213–1216, 1984.

[86] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency (Algorithms and Combinatorics).* Springer, July 2003.

[87] I. Segal, "Optimal pricing mechanisms with unknown demand," *American Economic Review*, 2003.

[88] V. Syrgkanis and E. Tardos, "Bayesian sequential auctions," in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 929–944, 2012.

[89] V. Syrgkanis and E. Tardos, "Composable and efficient mechanisms," arXiv preprint arXiv:1211.1325, 2012.

[90] M. Talagrand, "Matching random samples in many dimensions," *Annals of Applied Probability*, vol. 2, no. 4, pp. 846–856, 1992.

[91] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, vol. 16, pp. 8–37, 1961.

[92] V. Vishnumurthy, S. Chandrakumar, and E. Sirer, "Karma: A secure economic framework for peer-to-peer resource sharing," in *Workshop on Economics of Peer-to-Peer Systems*, 2003.

[93] R. Wilson, *Nonlinear Pricing.* Oxford University Press, 1997.

[94] Q. Yan, "Mechanism design via correlation gap," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 710–719, SIAM, 2011.