
Bayesian Mechanism Design

Bayesian Mechanism Design

Jason D. Hartline

Northwestern University

Evanston, 60208

USA

hartline@eecs.northwestern.edu

now

the essence of **know**ledge

Boston – Delft

Foundations and Trends[®] in Theoretical Computer Science

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
USA
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is J. D. Hartline, Bayesian Mechanism Design, Foundation and Trends[®] in Theoretical Computer Science, vol 8, no 3, pp 143–263, 2012

ISBN: 978-1-60198-670-2

© 2013 J. D. Hartline

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc. for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1-781-871-0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

**Foundations and Trends[®] in
Theoretical Computer Science**
Volume 8 Issue 3, 2012
Editorial Board

Editor-in-Chief:

Madhu Sudan

*Microsoft Research New England
One Memorial Drive
Cambridge, Massachusetts 02142
USA*

Editors

Bernard Chazelle (Princeton)
Oded Goldreich (Weizmann Inst.)
Shafi Goldwasser (MIT and Weizmann Inst.)
Sanjeev Khanna (University of Pennsylvania)
Jon Kleinberg (Cornell University)
László Lovász (Microsoft Research)
Christos Papadimitriou (UC. Berkeley)
Prabhakar Raghavan (Yahoo! Research)
Peter Shor (MIT)
Madhu Sudan (Microsoft Research)
Éva Tardos (Cornell University)
Avi Wigderson (IAS)

Editorial Scope

Foundations and Trends[®] in Theoretical Computer Science

will publish survey and tutorial articles in the following topics:

- Algorithmic game theory
- Computational algebra
- Computational aspects of combinatorics and graph theory
- Computational aspects of communication
- Computational biology
- Computational complexity
- Computational geometry
- Computational learning
- Computational Models and Complexity
- Computational Number Theory
- Cryptography and information security
- Data structures
- Database theory
- Design and analysis of algorithms
- Distributed computing
- Information retrieval
- Operations Research
- Parallel algorithms
- Quantum Computation
- Randomness in Computation

Information for Librarians

Foundations and Trends[®] in Theoretical Computer Science, 2012, Volume 8, 4 issues. ISSN paper version 1551-305X. ISSN online version 1551-3068. Also available as a combined paper and online subscription.

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.

Contents

1	Introduction	1
1.1	Topics Covered	4
1.2	Topics Omitted	8
2	Equilibrium	11
2.1	Equilibrium	13
2.2	Independent, Single-dimensional, and Linear Utilities	16
2.3	Equilibrium Characterization	17
2.4	Solving for Equilibrium	19
3	Optimal Mechanisms	21
3.1	Single-dimensional Environments	22
3.2	Single-agent Revenue Maximization	25
3.3	Multi-agent Revenue Maximization	33
3.4	Regular Distributions and Ironing	35
4	Approximation Mechanisms	41
4.1	Reserve Pricing	42
4.2	Posted Pricing	45
4.3	Prior-independent Approximation	54

5 Multi-dimensional and Non-linear Preferences	61
5.1 Single-agent Optimization	62
5.2 Multi-agent Optimization with Revenue Linearity	67
5.3 Multi-agent Optimization without Revenue Linearity	71
5.4 Multi-agent Optimization with Multi-dimensional Externalities	80
6 Approximation for Multi-dimensional and Non-linear Preferences	85
6.1 Single-agent Approximation	85
6.2 Multi-agent Approximation in Service Constrained Environments	90
6.3 Multi-agent Approximation with Multi-dimensional Externalities	96
7 Computation and Approximation Algorithms	101
7.1 Single-dimensional Monotonization via Ironing	103
7.2 Multi-dimensional Monotonization via Matching	107
A Mathematical Reference	111
A.1 Submodular Set Functions	111
A.2 Matroid Set Systems	111
A.3 Convex Optimization	113
Acknowledgments	117
References	119

1

Introduction

The Bayesian¹ 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.

2 Introduction

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. Truth-telling 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 prefers 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

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

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 truth-telling 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 $\mathbf{t} = (t_1, \dots, t_n)$.

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 $\mathbf{s} = (s_1, \dots, s_n)$ 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 truth-telling.

Definition 2.1. A *dominant strategy equilibrium* (DSE) is a strategy profile $\mathbf{s}$ such that for all i , t_i , and $\mathbf{b}_{-i}$ (where $\mathbf{b}_{-i}$ 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 $\mathbf{t}$ are drawn at random from a *prior distribution* $\mathbf{F}$ (a joint probability distribution over type profiles) and this prior distribution is *common knowledge*.

The distribution $\mathbf{F}$ over $\mathbf{t}$ 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 $\mathbf{F}_{-i}|_{t_i}$. Of course, when the distribution of types is independent, i.e., $\mathbf{F}$ is the *product distribution* $F_1 \times \cdots \times F_n$, then $\mathbf{F}_{-i}|_{t_i} = \mathbf{F}_{-i}$.

A strategy profile is a mapping for each agent of her private type to an action in the game. Notice that a prior $\mathbf{F}$ and strategies $\mathbf{s}$ 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 $\mathbf{F}$ is a strategy profile $\mathbf{s}$ such that for all i and t_i , $s_i(t_i)$ is a best response when other agents play $\mathbf{s}_{-i}(\mathbf{t}_{-i})$ for $\mathbf{t}_{-i} \sim \mathbf{F}_{-i}|_{t_i}$.

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 $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{p} = (p_1, \dots, p_n)$ 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 $\mathbf{v} = (v_1, \dots, v_n)$ to an allocation $\mathbf{x}(\mathbf{v})$ and payments $\mathbf{p}(\mathbf{v})$. 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 $x_i(z) = \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[x_i(z, \mathbf{v}_{-i})]$ and $p_i(z) = \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[p_i(z, \mathbf{v}_{-i})]$. 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 , $\mathbf{v}$, and z ; that agent i has higher utility for her strategy at v_i than following her strategy for value z , i.e.,

$$v_i \cdot x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i \cdot x_i(z, \mathbf{v}_{-i}) - p_i(z, \mathbf{v}_{-i}) \quad (\text{DSE})$$

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

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 truth-telling 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):

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

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 $x_i(z, \mathbf{v}_{-i})$ and $p_i(z, \mathbf{v}_{-i})$. 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,

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

because of the dominant strategy truth-telling 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,

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

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

¹The calculation proceeds as follows with $v_{(i)}$ denoting the i th 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 $\mathbf{E}[v_{(2)} \mid \text{case (iii) occurs}] = 2/3$. 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 $\mathbf{x} = (x_1, \dots, x_n)$ and prices $\mathbf{p} = (p_1, \dots, p_n)$. 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 $\mathbf{x}$ is feasible then $\mathbf{x}_{-i} = (x_1, \dots, x_{i-1}, 0, x_{i+1}, \dots, x_n)$ 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., $\max_{\mathbf{x} \in \mathcal{X}} \sum_i v_i x_i$ for weights $\mathbf{v} = (v_1, \dots, v_n)$. 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 $\mathbf{v}$ and outcome $(\mathbf{x}, \mathbf{p})$ 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 $(\mathbf{x}, \mathbf{p})$ 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 $\text{OPT}(\mathbf{v}) = \max_{\mathbf{x} \in \mathcal{X}} \sum_i v_i x_i$, the payment of an agent i is:

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

Notice that when $x_i = 1$, $\text{OPT}(\mathbf{v}) - v_i$ 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 $\mathbf{v}$ drawn from distribution $\mathbf{F}$. 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., $V(q) = F^{-1}(1 - q)$.




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.




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).

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

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:

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

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