8
Multi-dimensional and Non-linear Preferences

Chapters 2–7 focused on environments where the 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 single-agent problems (like those in Section 3.4 on page 79) may be analytically or computationally intractable. Second, we should not expect the revenue-linearity condition of Definition 3.16 on page 83 to hold when agents have non-linear preferences (in fact, it also does not generally hold for linear but 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 the order given.

8.1 Optimal Single-agent Mechanisms

A general agent has a type t drawn from an abstract type space $\mathcal{T}$ according to a distribution F. A mechanism can produce an outcome w from a general space of outcomes $\mathcal{W}$. Outcomes can be complicated objects but we will project them down onto our familiar notation for allocation and payment as follows. When the outcome includes a payment, it is denoted by p. In general an outcome can include multiple alternatives by which an agent is allocated or not allocated, $x = 1$ denotes the former and $x = 0$ denotes the latter. Both x and p are encoded by w, which may also encode other aspects of the outcome.

Recall any mechanism for a single agent, by the taxation principle, can be represented by a menu, i.e., a set of outcomes, where the agent picks her favorite outcome from the menu. Simple mechanisms can best be described by the set of outcomes they allow. For individually rational mechanisms it is without loss to assume that the outcome $\emptyset$, which does not allocate and requires no payment, is available in $\mathcal{M}$ and that all types obtain zero utility for this outcome. When listing the outcomes of the mechanism we will omit $\emptyset$. When describing more complex mechanisms it will be convenient to index the outcomes by the types that prefer them, e.g., as $w(t)$ for $t \in \mathcal{T}$; thus,

$$\mathcal{M} = \{w(t) : t \in \mathcal{T}\}.$$

Indexing as such, incentive compatibility and individual rationality can be expressed as follows, respectively.

$$\begin{aligned} u(t, w(t)) &\geq u(t, w(s)), && \forall t, s \in \mathcal{T}, \\ u(t, w(t)) &\geq 0, && \forall t \in \mathcal{T}. \end{aligned}$$

The subsequent developments of this section and chapter will be illustrated two representative examples, (a) a single-dimensional agent with a public budget (a non-linear preference), and (b) a (multi-dimensional)

Chapter 8: Topics Covered.

- unit-demand and public budget preferences,
- revenue linearity (revisited from Section 3.4.4 on page 83),
- interim feasibility,
- implementation by stochastic weighted optimization, and
- multi-dimensional virtual values via amortized analysis.

unit-demand agent with linear utility given by her value for the alternative obtained minus her payment.

Definition 8.1. A *public budget* agent has a (single-dimensional) value t for service and a public budget B. Her utility is linear for outcomes with required payment that is within her budget, and infinitely negative for outcomes with payments that exceed her budget. For outcome $w = (x, p)$, where x denotes the probability she is allocated and p is her required payment, her utility is $u(t, w) = t\,x - p$ when $p \leq B$ (and $-\infty$, otherwise).

Definition 8.2. A *unit-demand* agent desires one of m alternatives. Her type $t = (\{t\}_1, \ldots, \{t\}_m)$ is m-dimensional where $\{t\}_j$ is her value for alternative j. Her utility is linear; an outcome w is given by a payment and a probability measure over the m alternatives and nothing. For outcome $w = (\{x\}_1, \ldots, \{x\}_m, p)$, where $\{x\}_j$ denotes the probability she obtains alternative j and p is her required payment, her utility is $u(t, w) = \sum_j \{t\}_j \{x\}_j - p$. The probability she receives any allocation is $x = \sum_j \{x\}_j$.

The single-agent problems we consider are (i) the unconstrained single-agent problem, (ii) the ex ante constrained single-agent problem, and (iii) the interim constrained single-agent problem. Similar to Section 3.4, we are looking to understand the single-agent mechanisms that correspond to $R(1)$, $R(\hat{q})$, and $\mathbf{Rev}[\hat{y}]$.

Definition 8.3. A *single-agent problem* is specified by the type space, outcome space, and distribution over types as $(\mathcal{T}, \mathcal{W}, F)$, and a feasibility constraint. The feasibility constraints of single-agent problems are:

(i) *Unconstrained*: any mechanism is feasible. The optimal unconstrained mechanism's revenue is denoted $R(1)$. The *monopoly quantile* $\hat{q}^\star$ is defined to be its ex ante sale probability.

(ii) *(Weak) ex-ante constrained*: for ex ante constraint $\hat{q}$, a mechanism is feasible if its ex ante allocation probability at most $\hat{q}$, i.e., $\mathbf{E}_{t \sim F}[x(t)] \leq \hat{q}$. The optimal $\hat{q}$ ex ante mechanism's revenue is denoted $R(\hat{q})$.

(iii) *(Weak) interim constrained*: for interim constraint $\hat{y}$ (a monotone non-increasing function from $[0, 1]$ to $[0, 1]$), a mechanism $\mathcal{M}$ is feasible if, for any subspace of types $S \subset \mathcal{T}$ with measure $\hat{q} = \mathbf{Pr}[t \in S]$ under distribution F, the probability that a type in the subset is allocated under $\mathcal{M}$ is at most that of a quantile $q \in [0, \hat{q}]$ under constraint $\hat{y}$. In other words, the allocation constraint $\hat{y}$ is satisfied if for type distribution F, quantile $q \sim U[0,1]$, and subspace of types S with $\mathbf{Pr}_{t \sim F}[t \in S] = \hat{q}$,

$$\mathbf{E}[x(t) \mid t \in S] \leq \mathbf{E}[\hat{y}(q) \mid q \leq \hat{q}] \, .$$

The optimal $\hat{y}$ interim mechanism's revenue is denoted $\mathbf{Rev}[\hat{y}]$.

The optimal revenues from these convex maximization problems satisfy the standard concavity properties. Thus,

- the revenue curve $R(\cdot)$ is concave, and
- the interim optimal revenue is concave, i.e., $\mathbf{Rev}[\hat{y}] \geq \mathbf{Rev}[\hat{y}^{\dagger}] + \mathbf{Rev}[\hat{y}^{\ddagger}]$ for $\hat{y} = \hat{y}^{\dagger} + \hat{y}^{\ddagger}$.

See the Technical Note on page 255 for further discussion.

8.1.1 Public Budget Preferences

In this section we will describe optimal mechanisms for the three single-agent problems and agents with a public budget. Formal derivations of these optimal mechanisms are deferred to Section 8.6.

Example 8.1. The exemplary *uniform public-budget agent* has (single-dimensional) private type t uniform on type space $\mathcal{T} = [0,1]$ and public budget $B = 1/4$.

Technical Note. In Section 3.4 the ex ante constraints of $R(\cdot)$ and $\mathbf{Rev}[\cdot]$ were required to hold with equality, i.e., $\mathbf{E}[x(t)] = \hat{q}$ and $\mathbf{E}[x(t)] = \mathbf{E}[\hat{y}(q)]$, respectively. For multi-dimensional and non-linear preferences, single-agent mechanisms can be ill-behaved when required to serve types that the optimal unconstrained mechanism would reject. The "weak" definitions of Definition 8.3 avoid the resulting technicalities and are without loss for downward-closed environments. These weakened definitions of the single-agent problems, relative to those of Chapter 3, satisfy the following additional properties.

- The revenue curve $R(\cdot)$ is monotonically non-decreasing on $\hat{q} \in [0, \hat{q}^{\star}]$ and constant on $\hat{q} \in [\hat{q}^{\star}, 1]$.
- The unconstrained optimal revenue is given by $R(1) = R(\hat{q}^{\star})$.

The results of this chapter will be restricted to downward closed environments.

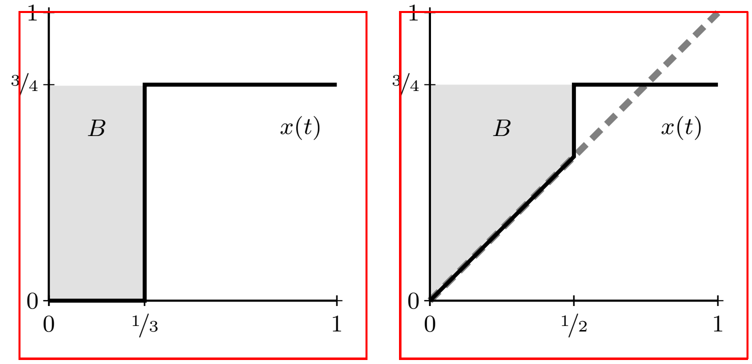

(a) Lottery allocation rule. (b) All-pay allocation rule.

Figure 8.1. Depicted are the allocation rules for the $^3/_4$-lottery and the two-agent all-pay auction (for public budget $B = {}^1/_4$). In each case, the ex ante allocation probability is $\hat{q} = {}^1/_2$ and the highest type $t = 1$ pays her full budget B. In subfigure (a), the $^2/_3$ measure of highest types buy the $^3/_4$-lottery at price $B = {}^1/_4$ for a total revenue of $^1/_6$. In subfigure (b), the allocation rule of the two-agent all-pay auction without the budget constraint is also depicted (think, gray, dashed line).

We begin by describing a few mechanisms for an agent with uniformly distributed private value and public budget $B = {}^1/_4$ (Example 8.1). Recall any mechanism for a single agent, by the taxation principle, can be represented by a menu, where the agent picks her favorite outcome from the menu. An outcome is a pair $w = (x, p)$, and individual rationality requires that the outcome $\emptyset = (0, 0)$ is implicitly in the menu of any mechanism. The following are mechanisms that sell with ex ante probability $\hat{q} = {}^1/_2$ and do not exceed the agent's budget.

- A $^3/_4$ *lottery at price* $^1/_4$ is $\mathcal{M} = \{({}^3/_4, {}^1/_4)\}$. The utility of an agent with type t for this lottery pricing is $^3/_4\, t - {}^1/_4$ and, thus, types $t \in [{}^1/_3, 1]$ buy. The ex ante sale probability is $^3/_4 \cdot {}^2/_3 = \hat{q} = {}^1/_2$ as desired. Its allocation rule is the following:

$$x(t) = \begin{cases} 0 & \text{if } t \leq {}^1/_3, \text{ and} \\ {}^3/_4 & \text{otherwise.} \end{cases}$$

- In a *two-agent all-pay auction*, types $t \in [{}^1/_2, 1]$ bid the budget $B = {}^1/_4$, remaining types bid $^1/_2\, t^2$ (the usual all-pay equilibrium, cf. Section 2.8 on page 38). The agents are symmetric; thus, each wins with ex ante probability $\hat{q} = 1/2$. Each agent's allocation rule is the following:

$$x(t) = \begin{cases} t & \text{if } t \leq 1/2 \text{, and} \\ 3/4 & \text{otherwise.} \end{cases}$$

These allocation rules are depicted in Figure 8.1. In each the payment of an agent with the highest type $t = 1$ is exactly her budget $B = 1/4$. To understand where the second allocation rule comes from, consider running an all-pay auction for two agents with types uniformly distributed on $[0, 1]$. Absent a budget constraint, an agent with the highest type $t = 1$ would win with probability one and pay $1/2$. This exceeds the budget and this agent would prefer to lower her bid relative to the non-budgeted equilibrium. In fact, types in $[3/4, 1]$ prefer to lower their bids to B, this causes types in $(1/2, 3/4]$ to prefer to raise their bids to B, and leaves type $t = \frac{1}{2}$ indifferent between bidding in the unconstrained all-pay equilibrium and bidding B (Exercise 8.1). The given allocation rule results. Naturally, there are many other possible mechanisms that satisfy the budget constraint and have ex ante sale probability of $\hat{q} = 1/2$; the optimal one, however, is the $3/4$-lottery.

Section 8.6 derives a characterization of optimal mechanisms for the three single-agent optimization problems for an agent with a public budget and uniformly distributed type. For the example of budget $B = 1/4$ and uniformly distributed types, these optimal mechanisms are as follows.

- The unconstrained optimal mechanism posts price $B = 1/4$, sells to the $3/4$ measure of types $t \in [1/4, 1]$, and has expected revenue $3/16$. Its ex ante sale probability is $\hat{q}^\star = 1 - B = 3/4$.
- The ex ante optimal mechanism for $\hat{q} \leq \hat{q}^\star$ is the $\hat{q} + B$ lottery at price B, i.e., $\mathcal{M}^{\hat{q}} = \{(\hat{q} + B, B)\}$. The top $\hat{q}/\hat{q}+B$ measure of types choose to buy this lottery. See Figure 8.1(a) for the special case of $\hat{q} = 1/2$ and $B = 1/4$ where $R(1/2) = 1/6$. For $\hat{q} \geq \hat{q}^\star$ the optimal mechanism with sale probability at most $\hat{q}$ is the optimal unconstrained mechanism, above. The revenue curve is given by

$$R(\hat{q}) = \begin{cases} \hat{q}\, B/\hat{q}+B & \text{if } \hat{q} \leq \hat{q}^\star, \text{ and} \\ 1 - B & \text{otherwise.} \end{cases}$$

- The interim optimal mechanism for allocation constraint $\hat{y}$ is given by types $0 \leq \hat{t}^\ddagger \leq \hat{t}^\dagger \leq 1$ and has allocation rule given by reserve pricing the weak types $t \in [0, \hat{t}^\ddagger)$, ironing the strong types $t \in (\hat{t}^\dagger, 1]$, and allocating maximally to intermediate types $t \in [\hat{t}^\ddagger, \hat{t}^\dagger)$. The reserve

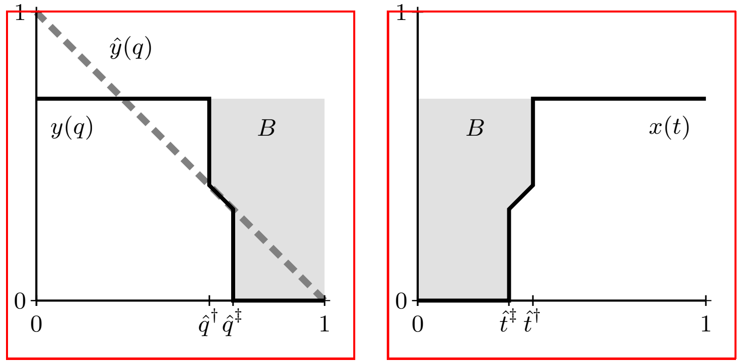

(a) Quantile space: $y(q)$. (b) Type space: $x(q)$.

Figure 8.2. For an agent with uniform type on $[0, 1]$ and public budget $B = 1/4$, depicted is the allocation rule for the optimal mechanism that satisfies the interim allocation constraint $\hat{y}(q) = 1 - q$. The allocation rule in quantile space is depicted in subfigure (a); the allocation rule in type space is depicted in subfigure (b). For the uniform distribution, these are mirror images of each other.

price is set to optimize revenue; the ironed interval is set to meet the budget constraint with equality. The example of $\hat{y}(q) = 1-q$ is given in Figure 8.2. It is most natural to describe the resulting allocation rule in quantile space. Recall that the quantile of a type is the measure of stronger types; for the uniform distribution the quantile of t is $q = 1-t$. Thus, the allocation rule of this mechanism is $x(t) = y(1-t)$ with $\hat{q}^\dagger = 1 - \hat{t}^\dagger$, $\hat{q}^\ddagger = 1 - \hat{t}^\ddagger$, and

$$y(q) = \begin{cases} 1/\hat{q}^\dagger \int_0^{\hat{q}^\dagger} \hat{y}(z)\,\mathrm{d}z & \text{if } q \in [0, \hat{q}^\dagger), \\ \hat{y}(q) & \text{if } q \in [\hat{q}^\dagger, \hat{q}^\ddagger), \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

While in these single-dimensional public budget problems, there is a natural ordering on types by strength, it is not the case that optimal mechanisms always break ties in the same way. In particular the interval of ironing, i.e. $[0, \hat{q}^\dagger]$, in the interim mechanism design problem depends on the allocation constraint $\hat{y}$. Unlike the case of single-dimensional linear preferences described in Chapter 3, there is no fixed virtual value function for which optimization of virtual surplus gives the optimal mechanism. Instead, the appropriate virtual value function will

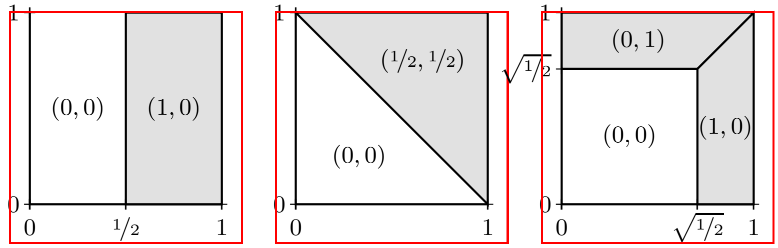

(a) Pricing item 1. (b) Uniform lottery. (c) Uniform pricing.

Figure 8.3. Depicted are the outcomes $x = (\{x\}_1, \{x\}_2)$ for the example mechanisms given in the text for the two-alternative uniform unit-demand agent (Example 8.2). The regions of allocation (gray) and non-allocation (white) have area $\hat{q} = 1 - \hat{q} = 1/2$.

depend on the environment and level of competition from other agents. This approach is further described subsequently in Section 8.6.

8.1.2 Unit-demand Preferences

In this section we will describe optimal mechanisms for the three single-agent problems and (multi-dimensional) agents with unit-demand. Formal derivations of these optimal mechanisms are deferred to Section 8.7.

Example 8.2. The exemplary *two-alternative uniform unit-demand agent* has type $t \in \mathcal{T} = [0, 1]^2$ uniformly distributed, i.e., $t \sim U[0, 1]^2$; equivalently her value for each alternative is i.i.d. and uniform on $[0, 1]$. Recall, $t = (\{t\}_1, \{t\}_2)$.

We begin by describing a few mechanisms for the two-alternative uniform unit-demand agent of Example 8.2. Recall any mechanism for a single agent, by the taxation principle, can be represented by a menu where the agent picks her favorite outcome from the menu. An outcome is a triple $w = (\{x\}_1, \{x\}_2, p)$ and individual rationality requires that the outcome $\emptyset = (0, 0, 0)$ is implicitly in the menu of any mechanism. The following are mechanisms that sell with ex ante probability $\hat{q} = 1/2$:

- Sell only item 1 for price $1/2$: $\mathcal{M} = \{(1, 0, 1/2)\}$.
- Sell the *uniform lottery* for price $1/2$: $\mathcal{M} = \{(1/2, 1/2, 1/2)\}$.
- Sell either item at a *uniform price* $\sqrt{1/2}$: $\mathcal{M} = \{(1, 0, \sqrt{1/2}), (0, 1, \sqrt{1/2})\}$.

The first two mechanisms obtain revenue $1/2$ when they sell and, thus, obtain an expected revenue of $1/4$. The final mechanism obtains revenue

$\sqrt{1/2} > 1/2$ when it sells. Of these three mechanisms, the latter has the highest revenue. Figure 8.3 depicts the outcomes of these mechanisms. Of course, these are just three of an infinite number of mechanisms that sell with ex ante probability $\hat{q} = 1/2$. As we will describe below, the ex ante optimal mechanism for $\hat{q} = 1/2$ is in fact the uniform pricing of both alternatives at $\sqrt{1/2}$.

Section 8.7 derives a characterization of optimal mechanisms for the three single-agent optimization problems (Definition 8.3) for a unit-demand agent with a uniformly distributed type. For $m = 2$ alternatives, these optimal mechanisms are as follows.

- The unconstrained optimal mechanism is the uniform pricing at $\sqrt{1/3}$, i.e., $\mathcal{M}^{\hat{q}} = \{(1, 0, \sqrt{1/3}), (0, 1, \sqrt{1/3})\}$. This mechanism serves with ex ante probability $\hat{q}^\star = 2/3$ and has revenue $R(1) = \sqrt{4/27} \approx 0.38$.
- The ex ante optimal mechanism for $\hat{q} \leq 2/3$ is the uniform pricing at $\sqrt{1 - \hat{q}}$. For $\hat{q} > 2/3$ it is the $2/3$ ex ante optimal mechanism, i.e., the optimal unconstrained mechanism, above. The revenue curve is given by

$$R(\hat{q}) = \begin{cases} \hat{q}\sqrt{1-\hat{q}} & \text{if } \hat{q} \leq 2/3, \text{ and} \\ \sqrt{4/27} \approx 0.38 & \text{otherwise.} \end{cases}$$

- The interim optimal mechanism for allocation constraint $\hat{y}$ is has allocation rule (cf. Section 3.4.2 on page 81 and see Figure 8.4).

$$y(q) = \begin{cases} \hat{y}(q) & \text{if } q \leq 2/3, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

  Figure 8.4(c) is the example of $\hat{y}(q) = 1 - q$; the menu of the optimal mechanism for this interim constraint is

$$\begin{aligned} \mathcal{M} = &\left\{ \left(x, 0, \textstyle\int_0^x \sqrt{\max(z, 1/3)}\, \mathrm{d}z\right) \,:\, x \in [1/3, 1] \right\} \\ &\cup \left\{ \left(0, x, \textstyle\int_0^x \sqrt{\max(z, 1/3)}\, \mathrm{d}z\right) \,:\, x \in [1/3, 1] \right\}. \end{aligned}$$

Unlike the example of a single dimensional agent with a public budget (Example 8.1), there is no implicit ordering on types by strength. Which is stronger type $t^\dagger = (.8, .1)$ or type $t^\ddagger = (.6, .6)$? In fact, which of these types is stronger generally depends on the mechanism. The uniform pricing of $\sqrt{1/2} \approx .71$ (Figure 8.3(c)) serves type $t^\dagger$ and rejects type $t^\ddagger$ while the uniform lottery at price $1/2$ (Figure 8.3(b)) serves type $t^\ddagger$ and rejects type $t^\dagger$. If we consider optimal mechanisms for any strictly decreasing interim constraint $\hat{y}$, however, it is clear that types are ranked as stronger based on their maximum coordinate $\max(\{t\}_1, \{t\}_2)$. Moreover,

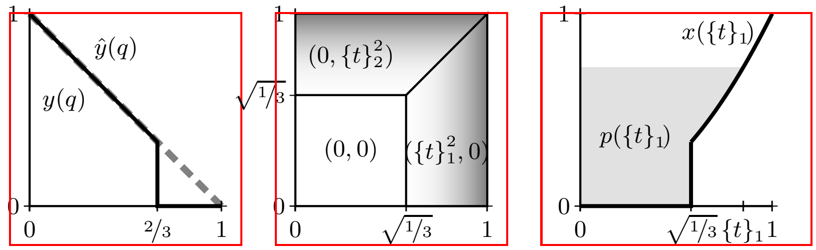

(a) Allocation rules. (b) Allocation. (c) Payment, $\{t\}_1 \geq \{t\}_2$.

Figure 8.4. For a unit-demand agent with uniformly distributed type for $m = 2$ alternatives (Example 8.2), the allocation constraint $\hat{y}(q) = 1 - q$, its optimal allocation rule $y(\cdot)$, the allocation in the two-dimensional type space $\mathcal{T} = [0,1]^2$, and the payment as a function of $\{t\}_1 > \{t\}_2$ are depicted. In subfigure (b), the degree of shading represents the probability with which a type $t = (\{t\}_1, \{t\}_2)$ receives her preferred alternative. When alternative 1 is preferred (i.e., $\{t\}_1 > \{t\}_2$) and at least the reservation value (i.e., $\{t\}_1 \geq \sqrt{1/3}$), the probability of receiving alternative 1 is $\{t\}_1^2$. For a type $t = (\{t\}_1, \{t\}_2)$ with $\{t\}_1 > \{t\}_2$, subfigure (c) depicts the allocation rule $x(\{t\}_1)$ as a function of the agent's value for alternative 1. The agent's payment for her preferred outcome with such a type is calculated as the area above this allocation rule $x(\cdot)$, by integrating $x^{-1}(\cdot)$ vertically, as $p(\{t\}_1) = \int_0^{\{t\}_1^2} \sqrt{\max(z, 1/3)}\, \mathrm{d}z$ for $\{t\}_1 \geq \sqrt{1/3}$.

the optimal mechanism for an interim constraint can be viewed as a convex combination of optimal mechanisms for ex ante constraints. Though these simplifying properties of optimal single-agent mechanisms do not hold in general for unit-demand agents, in Section 8.7 we will describe sufficient conditions, beyond the uniform distribution, under which they extend. As we already observed in the preceding study of public budgets these properties do not hold for non-linear preferences.

Important differences in families of single-agent problems have been exhibited above in our study unit-demand and public budget preferences under the uniform distribution. In particular, unit-demand preferences drawn from the uniform distribution (Example 8.2) behave similarly to the single-dimensional linear preferences of Chapter 3, whereas public budget preferences do not (Example 8.1). In the subsequent sections as we describe optimal multi-agent mechanisms both for families of preferences that behave similarly to the uniform unit-demand example and families of preferences that behave similarly to the uniform public budget example. This latter class of mechanisms will be completely general.

8.2 Service Constrained Environments

We will consider environments, like the single-dimensional environments of Section 3.1 on page 54, 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 an 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 environments that exhibit multi-dimensional externalities is deferred to Section 8.5.

Definition 8.4. 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. Subsets of the n agents N that can be simultaneously served are given by $\mathcal{X} \subset 2^N$.

In the subsequent sections we will reduce the problem of multi-agent mechanism design to a collection of single-agent mechanism design problems. These sections will not further address the details of how to solve these single-agent problems, instead they will focus on how the multi-agent mechanism is constructed from the single-agent components. We begin in Section 8.3 where the simplifying assumption of revenue linearity enables optimal multi-agent mechanisms to be constructed from the single-agent ex ante optimal mechanisms. In Section 8.4 we consider the more general case where revenue linearity does not hold. In this case we describe how to construct multi-agent mechanisms from the solution to the single-agent interim optimal mechanism design problems.

The definition of service constrained environments, above, corresponds to the general feasibility environments of Section 3.1. The framework can be easily extended to incorporate service costs that are a function of the set of agents served (see Exercise 8.8).

8.3 The Ex Ante Reduction

In this section we construct optimal multi-agent mechanisms for agents whose single-agent problems behave similarly to the single-dimensional linear agents of Chapter 3. We will use, as a running example of such agents, the uniform unit-demand agent (Example 8.2 on page 259). Our approach follows and extends that of Section 3.4 on page 79. In this approach the multi-agent mechanism design problem is reduced to the single-agent ex ante optimal mechanism design problem. This single-agent optimization gives rise to a revenue curve $R(\cdot)$. The revenue linearity property (Definition 3.16), specifically that $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[\hat{y}^{\dagger}] + \mathbf{Rev}[\hat{y}^{\ddagger}]$ for $\hat{y} = \hat{y}^{\dagger} + \hat{y}^{\ddagger}$, implies that any interim optimal mechanism can be expressed in terms of marginal revenue $\mathbf{Rev}[\hat{y}] = \mathbf{E}_q\big[R'(q)\,\hat{y}(q)\big]$ (Proposition 3.17).[1] The optimal mechanism is, thus, the one that maximizes marginal revenue (cf. Theorem 3.20).

Recall from Definition 8.3, the ex ante optimal mechanism design problem is given an upper bound $\hat{q}$ on the ex ante probability of serving the agent, over randomization in the agent's type and the mechanism, and determines the outcome rule $w^{\hat{q}}$ which maps types to outcomes. Encoded in an outcome $w^{\hat{q}}(t)$ for type t is a probability of service, denoted $x^{\hat{q}}(t)$, and a payment, denoted $p^{\hat{q}}(t)$. This mechanism can alternatively be thought of as the menu $\mathcal{M}^{\hat{q}} = \{w^{\hat{q}}(t) \,:\, t \in \mathcal{T}\}$. The revenue of the ex ante optimal mechanism for every $\hat{q} \in [0,1]$ defines the revenue curve $R(\hat{q})$.

We will use the uniform unit-demand agent of Example 8.2, which is revenue linear, to illustrate this construction and then give the formal definition, derivation, and proof of correctness. For this example, recall that the $\hat{q}$ ex ante optimal mechanism posts the uniform price $\sqrt{1-\hat{q}}$ for each of the two alternatives (see Figure 8.3(c)).

[1] For review: View allocation constraint $\hat{y}$, a monotone non-increasing function from $[0,1]$ to $[0,1]$, as a convex combination of reverse step functions each of which steps from 1 to 0 at some $\hat{q}$. In this convex combination $\hat{q}$ is drawn with cumulative distribution function $G^{\hat{y}}(q) = 1 - \hat{y}(q)$ and density function $g^{\hat{y}}(q) = -\hat{y}'(q)$. The optimal revenue for each $\hat{q}$, of the $\hat{q}$ ex ante optimal mechanism, defines $R(\hat{q})$, the revenue of the convex combination is thus $\mathbf{E}[(-\hat{y}'(q)\,R(q)]$. Integration by parts with $R(1) = R(0) = 0$ gives the marginal revenue $\mathbf{MargRev}[\hat{y}] = \mathbf{E}[R'(q)\,\hat{y}(q)]$.

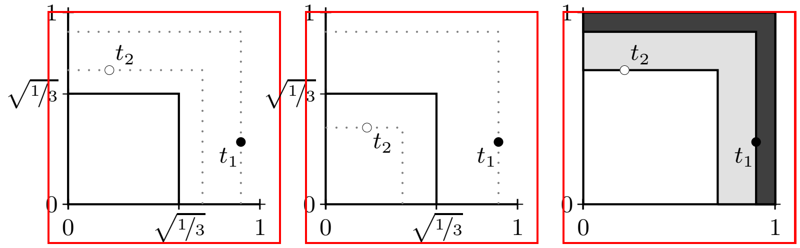

(a) Agent 2 above reserve. (b) Agent 2 below reserve. (c) Mapping types to quantiles.

Figure 8.5. The optimal auction for the red-blue car environment of Example 8.3 is illustrated. In subfigure (a), the value agent 2 has for her preferred color $\{t_2\}_2$ exceeds the reserve price of $\sqrt{1/3}$, agent 1 wins her preferred alternative 1 and pays $\{t_2\}_2$. In subfigure (b), the value agent 2 has for her preferred item is below the reserve, agent 1 wins alternative 1 and pays the reserve price. In subfigure (c), the types stronger than agent 1 (dark gray region) and agent 2 (light and dark gray regions) are depicted. As the distribution on types is uniform, the quantile q_1 of agent 1 with type t_1 and q_2 of agent 2 with t_2 can be calculated as the areas of these regions, respectively.

8.3.1 Example: Uniform Unit-demand Preferences

This section illustrates the general construction of optimal mechanisms for revenue-linear agents. The first example considers two (identically distributed, i.e., symmetric) uniform unit-demand agents (Example 8.2) in a single-item environment, and the second considers (non-identically distributed, i.e., asymmetric) agents.

Example 8.3. There are two agents, the seller has one car that he can paint red or blue on its sale. The agents' types, i.e., values for each color, are independently, identically, and uniformly distributed on $[0, 1]$. The second-price auction for each agent's preferred color with a reserve of $\sqrt{1/3}$ is revenue optimal.

To explain Example 8.3, we will follow the construction of optimal mechanisms for single-dimensional linear agents as described in Chapter 3. In particular, we map types to quantiles by their relative strength. We calculate marginal revenue of a type t with quantile q as $R'(q) = \frac{\mathrm{d}}{\mathrm{d}q} R(q)$. We allocate the car to the type with the highest positive marginal revenue. To determine the payment and what color to paint the car, we look at the weakest quantile, given the quantile(s) of the other agent(s), at which the winner still wins and allocate according to ex ante mechanism for this critical quantile. These four steps are described in detail below; the mechanism is illustrated in Figure 8.5.

Observe that agent 2 may win the car and this imposes an interim constraint on agent 1. As we have observed previously, the optimal single-agent mechanism for agent 1 orders her types by her value for her preferred alternative. This ordering on types allows types to be mapped to quantiles. Recall, the quantile of a type designates its strength relative to the distribution of types and is defined as the measure of stronger types. For the example type $t = (\{t\}_1, \{t\}_2)$ is weaker than all types s with $\max(\{s\}_1, \{s\}_2) > \max(\{t\}_1, \{t\}_2)$ and stronger than all types s with $\max(\{s\}_1, \{s\}_2) < \max(\{t\}_1, \{t\}_2)$. As the distribution F is uniform on $[0,1]^2$, the quantile of a type t is $q = 1 - [\max(\{t\}_1, \{t\}_2)]^2$; see Figure 8.5(c).

The revenue curves $R(\cdot)$ is defined from the solution to the ex ante optimal mechanism for each $\hat{q} \in [0,1]$. For the two-alternative uniformly distributed types of the example and $\hat{q} \leq 2/3$, $\hat{q}$ ex ante optimal mechanism posts price $\sqrt{1-\hat{q}}$ and obtains revenue $R(\hat{q}) = \hat{q}\sqrt{1-\hat{q}}$. For two symmetric agents, as in our example, the details of the revenue curve before its maximum, $\hat{q}^\star = 2/3$ for the example, are irrelevant as long as it is strictly concave (unlike the asymmetric example given subsequently). The agent with the stronger quantile wins, as long as that quantile is at least the quantile reserve, which is given by the unconstrained optimal mechanism and is $\hat{q}^\star = 2/3$.

Agent 1 will win the auction whenever her quantile is less than agent 2's quantile and the quantile reserve. Agent 1's critical quantile is thus $\hat{q}_1 = \min(q_2, \hat{q}^\star)$. Fixing agent 2's type and quantile, agent 1 faces the $\hat{q}_1$ ex ante optimal mechanism. For the example, this mechanism is given by the menu $\{(1, 0, \sqrt{1-\hat{q}_1}), (0, 1, \sqrt{1-\hat{q}_1})\}$, i.e., it is a uniform pricing of $\sqrt{1-\hat{q}_1}$. When this critical quantile comes from agent 2, the uniform price is exactly the value agent 2 has for her preferred alternative. When this critical quantile comes from the reserve, then the uniform price is from the optimal unconstrained mechanism, i.e., it is $\sqrt{1/3}$. Thus, agent 1 is offered a uniform price that is the higher of the reserve and agent 2's value for her preferred alternative. The optimal mechanism is the second-price auction for the agents' preferred alternative with a uniform reserve of $\sqrt{1/3}$.

The next example environment shows that the approach taken above can treat asymmetric agent preferences and that the dimensionality of the preferences need not be the same. Due to the asymmetry, however,

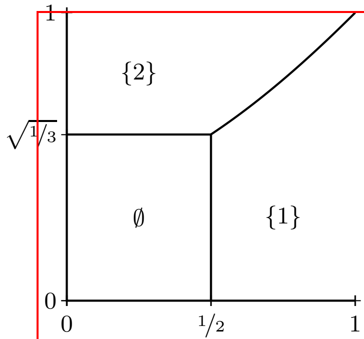

Figure 8.6. The allocation of the optimal auction for the red-blue-green car environment of Example 8.4 is depicted. The horizontal axis is the value of agent 1 for the green car t_1; the vertical axis is the value of agent 2 for her preferred color of red and blue $\{t_2\}_{\max} = \max_j \{t_2\}_j$.

the resulting optimal mechanism is more complex and will depend more finely on the details of the agents' revenue curves.

Example 8.4. There are two agents. A seller has one car that he can paint red, blue, or green on its sale. Agent 1 is single-dimensional and has uniformly distributed value for a green car (and no value for a red or blue car); agent 2 has independent and uniformly distributed values for red or blue cars (and no value for a green car). The winner of the optimal mechanism is depicted in Figure 8.6, the winner always receives her preferred color car.

The optimal mechanism of the red-blue-green car example (Example 8.4) is constructed as follows. We map agent 1's single-dimensional type to a virtual value via the standard transformation from value to quantile to marginal revenue as per the construction of Section 3.4. With type t_1, agent 1's virtual value is $\phi_1(t_1) = 2t_1 - 1$. We map agent 2's multi-dimensional type to a (single-dimensional) virtual value similarly:

(i) $q_2 = 1 - [\max(\{t_2\}_1, \{t_2\}_2)]^2$.
(ii) $R_2'(q) = \frac{\mathrm{d}}{\mathrm{d}q}[q\sqrt{1-q}] = \frac{1-3/2\,q}{\sqrt{1-q}} = [1 - 3/2\,q]\,[1-q]^{-1/2}$.
(iii) $\phi_2(t_2) = 3/2\,\{t_2\}_{\max} - 1/2\,\{t_2\}_{\max}^{-1}$ where $\{t_2\}_{\max} = \max_j \{t_2\}_j$.

The winner is the agent i with the highest virtual value. She receives her preferred alternative at a price of $\phi_i^{-1}\big(\max\{\phi_{3-i}(t_{3-i}), 0\}\big)$.

The remainder of this section will formalize the construction of the marginal revenue mechanism for revenue-linear agents in general and prove its optimality among all mechanisms.

8.3.2 Orderability

Fundamental to the examples above is the identification of an ordering on types from which types can be mapped to quantiles (and then to marginal revenues). In this section we show that the existence of such an ordering is a consequence of the revenue linearity property (restating Definition 3.16):

$$\mathbf{Rev}[\hat{y}^\dagger + \hat{y}^\ddagger] = \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger].$$

The uniform unit-demand agent (Example 8.2) is revenue linear. This observation follows from the fact that the interim optimal mechanism for constraint $\hat{y}$ is a convex combination of ex ante optimal mechanisms. Revenue linearity immediately implies that the surplus of marginal revenue (Definition 3.15) is equal to the optimal revenue (Proposition 3.17):

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

However, it does not immediately suggest how to implement surplus of marginal revenue maximization as the mapping from type space to quantile space is not explicit in a multi-dimensional type space (as it is in a single-dimensional type space).

Definition 8.5. A single-agent problem given by a type space, outcome space, and distribution over types is *orderable* if there is an equivalence relation on types and an ordering over equivalence classes, such that for every allocation constraint $\hat{y}$, an 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 agent (Example 8.1) is not orderable by the above definition. Though, the agent's value for service gives a natural ordering on types, the optimal mechanism for $\hat{y}$ irons the strongest quantiles so as to meet the budget constraint with equality and this ironed interval depend on the allocation constraint $\hat{y}$ (see Figure 8.2 on page 258).

Theorem 8.1. *For any single-agent problem, revenue linearity implies orderability.*

The main intuition behind Theorem 8.1 comes from the observation that revenue linearity implies that the allocation rule y that is obtained from optimization subject to interim constraint $\hat{y}$ 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 of quantile space where the marginal revenue is constant). The proof illustrates how revenue linearity enables interim optimal mechanisms to be understood in terms of ex ante optimal mechanisms.

Recall the definition of the cumulative allocation rule as $Y(\hat{q}) = \int_0^{\hat{q}} y(q)\,\mathrm{d}q$ and, by integration by parts, we can express the marginal revenue of any allocation rule y as (recall that $Y(0) = 0$):

$$\begin{aligned}\mathbf{MargRev}[y] &= \Big[R'(q)\,Y(q)\Big]_0^1 - \int_0^1 R''(q)\,Y(q)\,\mathrm{d}q \\ &= R'(1)\,Y(1) - \mathbf{E}\big[R''(q)\,Y(q)\big]\,. \end{aligned} \tag{8.1}$$

We will prove Theorem 8.1 by combining the following two lemmas. The first lemma shows that for any quantile $\hat{q}$ where marginal revenue is strictly decreasing, i.e., where $R''(\hat{q}) < 0$, the interim optimal allocation rule y for interim constraint $\hat{y}$ allocates to the maximum extent possible, i.e., the interim constraint of Definition 8.3 is tight at $\hat{q}$, i.e., $Y(\hat{q}) = \hat{Y}(\hat{q})$.

Lemma 8.2. *For a revenue-linear agent, allocation rule y that is optimal for allocation constraint $\hat{y}$, and any ex ante probability $\hat{q}$ with $R''(\hat{q}) \neq 0$, the cumulative allocation rule and constraint satisfy $Y(\hat{q}) = \hat{Y}(\hat{q})$.*

Proof. If we optimize revenue for allocation constraint $\hat{y}$ and obtain a mechanism with allocation rule y, then $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[y]$. Revenue linearity implies that optimal revenues are equal to marginal revenues, i.e., $\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}]$ and $\mathbf{Rev}[y] = \mathbf{MargRev}[y]$, respectively. Writing the difference between these marginal revenues and employing equation (8.1), we have:

$$\begin{aligned}0 &= \mathbf{MargRev}[\hat{y}] - \mathbf{MargRev}[y] \\ &= R'(1)\,[\hat{Y}(1) - Y(1)] + \mathbf{E}[[-R''(q)]\,[\hat{Y}(q) - Y(q)]]\,. \end{aligned} \tag{8.2}$$

By interim feasibility of y for $\hat{y}$, $[\hat{Y}(q) - Y(q)] \geq 0$. By concavity of revenue curves $[-R''(q)] \geq 0$. By monotonicity of revenue curves $R'(1) \geq 0$. Thus, every term in equation (8.2) is non-negative. The only way it can be identically zero if $[-R''(q)] > 0$ implies that $[\hat{Y}(q) - Y(q)] = 0$ as required by the lemma. (Also observe, though unnecessary for the lemma, that $R'(1) > 0$ implies that $[\hat{Y}(1) - Y(1)] = 0$.) □

An immediate corollary of Lemma 8.2 is that the ex ante optimal mechanism for any $\hat{q}$ with $R''(\hat{q}) < 0$ deterministically serves or rejects each type, i.e., the allocation rule in type space is $x^{\hat{q}}(t) \in \{0, 1\}$. Contrast this corollary to the uniform public-budget example where the optimal mechanism for ex ante constraint $\hat{q} = 1/2$ is pricing the $3/4$ lottery while $R(\hat{q}) = \hat{q}\,B/\hat{q}+B$ is strictly convex at $\hat{q} = 1/2$ (Section 8.1.1).

Corollary 8.3. *For a revenue-linear agent and any ex ante probability $\hat{q}$ with $R''(\hat{q}) \neq 0$, the $\hat{q}$ ex ante optimal mechanism deterministically serves or rejects each type $t \in \mathcal{T}$.*

Proof. Notice that an ex ante constraint of $\hat{q}$ is equivalent to the interim constraint given by the reverse-step function that steps from 1 to 0 at $\hat{q}$, denoted $\hat{y}^{\hat{q}}$. By Lemma 8.2, the optimal allocation rule for this constraint is the reverse-step function itself. A mechanisms whose allocation rule is a reverse-step function deterministically allocates or rejects each type. □

For $\hat{q}$ where the revenue curve is strictly concave, Corollary 8.3 implies the type space $\mathcal{T}$ is partitioned into types that are served and those that are rejected. Denote the allocation rule of the $\hat{q}$ ex ante optimal mechanism in type space by $x^{\hat{q}}(\cdot)$ and the subset of types it serves by

$$S^{\hat{q}} = \{t \in \mathcal{T} \,:\, x^{\hat{q}}(t) = 1\}. \tag{8.3}$$

The following lemma shows that the sets of types served by the ex ante optimal mechanisms are nested.

Lemma 8.4. *For a revenue-linear agent and any ex ante probabilities $\hat{q}^{\dagger} < \hat{q}^{\ddagger}$ with $R''(\hat{q}^{\dagger}) \neq 0$ and $R''(\hat{q}^{\ddagger}) \neq 0$, then $S^{\hat{q}^{\dagger}} \subset S^{\hat{q}^{\ddagger}}$.*

Proof. This proof is illustrated in Figure 8.7. Let $\hat{y}^{\dagger}$ and $\hat{y}^{\ddagger}$ denote the interim allocation constraint corresponding to the ex ante constraints $\hat{q}^{\dagger}$ and $\hat{q}^{\ddagger}$, respectively. Consider the interim constraint $\hat{y} = 1/2\,\hat{y}^{\dagger} + 1/2\,\hat{y}^{\ddagger}$. The constraint $\hat{y}$ is a reverse stair function that steps from 1 to $1/2$ at $\hat{q}^{\dagger}$ and from $1/2$ to 0 at $\hat{q}^{\ddagger}$. Suppose for a contradiction that $S^{\hat{q}^{\dagger}}$ contains a measurable (with respect to the distribution F) set of types that is not also contained in $S^{\hat{q}^{\ddagger}}$. By revenue linearity the optimal mechanism for $\hat{y}$ is the convex combination of the optimal mechanisms for $\hat{y}^{\dagger}$ and $\hat{y}^{\ddagger}$; denote

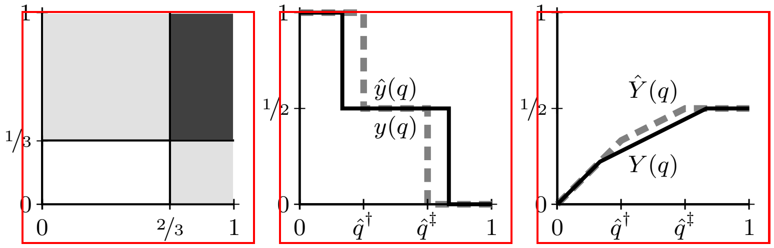

(a) Type-allocation rule. (b) (Quantile) allocation rules. (c) Cumulative allocation rules.

Figure 8.7. For a two-alternative unit-demand agent, the proof of Lemma 8.4 is illustrated. In this hypothetical situation the $\hat{q}^{\dagger}$ (resp. $\hat{q}^{\ddagger}$) optimal mechanism posts a price for selling alternative 1 (resp. 2) only. In subfigure (a) the two-dimensional type space $\mathcal{T} = [0, 1]$ and the allocation of the convex combination of the $\hat{q}^{\dagger}$ and $\hat{q}^{\ddagger}$ optimal mechanism are depicted. The types $t \in S^{\hat{q}^{\dagger}} \cap S^{\hat{q}^{\ddagger}}$ (dark gray region) have allocation probability $x(t) = 1$, types t in the symmetric difference of $S^{\hat{q}^{\dagger}}$ and $S^{\hat{q}^{\ddagger}}$ (light gray region) have allocation probability $x(t) = 1/2$, and the remaining types have allocation probability $x(t) = 0$. In subfigure (b) the allocation constraint $\hat{y}$ and allocation rule y are depicted. In subfigure (c) the cumulative allocation constraint $\hat{Y}$ and cumulative allocation rule Y are depicted. Allocation constraints are depicted with thick, dashed, gray lines and allocation rules are depicted with thin, solid, black lines.

its allocation rules in quantile and type space by y and x, respectively. The following table tallies the measure and service probability of each relevant subset of types (with ρ defined by the probability of the second line):

S	$\mathbf{Pr}[t \in S]$	$x(t)$
$S^{\hat{q}^{\dagger}} \cap S^{\hat{q}^{\ddagger}}$	$\hat{q}^{\dagger} - \rho$	1
$S^{\hat{q}^{\dagger}} \setminus S^{\hat{q}^{\ddagger}}$	ρ	$1/2$
$S^{\hat{q}^{\ddagger}} \setminus S^{\hat{q}^{\dagger}}$	$\hat{q}^{\ddagger} - \hat{q}^{\dagger} + \rho$	$1/2$
$\mathcal{T} \setminus S^{\hat{q}^{\dagger}} \setminus S^{\hat{q}^{\ddagger}}$	$1 - \hat{q}^{\ddagger} + \rho$	0

The allocation rule y is a reverse-stair function that steps from 1 to $1/2$ at $\hat{q}^{\dagger} - \rho$ and from $1/2$ to 0 at $\hat{q}^{\ddagger} + \rho$. Inspection reveals, for a contradiction, that Lemma 8.2 is violated at $\hat{q}^{\dagger}$ and $\hat{q}^{\ddagger}$. For example, $\hat{Y}(\hat{q}^{\dagger}) = \hat{q}^{\dagger} > Y(\hat{q}^{\dagger}) = \hat{q}^{\dagger} - 1/2\,\rho$. □

Proof of Theorem 8.1. Define type subspace $S^{\hat{q}}$ as in equation (8.3). Define the marginal revenue of a type t as $\inf\{R'(\hat{q}) : S^{\hat{q}} \ni t\}$. The equivalence classes of Definition 8.5 are sets of types with the same marginal revenue; types with marginal revenue zero are in the equivalence class $\perp$. By Lemma 8.4, these definitions are well defined.

Consider the $\hat{q}$ optimal mechanism. If $R''(\hat{q}) \neq 0$ then it is optimal to serve types $S^{\hat{q}}$. Greedy by marginal revenue (as defined above) serves these types. If $R''(\hat{q}) = 0$, then let $(\hat{q}^\dagger, \hat{q}^\ddagger)$ be the interval on which $R'(\hat{q})$ is constant. An optimal mechanism randomizes between the $\hat{q}^\dagger$ ex ante optimal and $\hat{q}^\ddagger$ ex ante optimal mechanisms so that the total sale probability is $\hat{q}$. By Lemma 8.4, types in $S^{\hat{q}^\dagger}$, with marginal revenue strictly greater than $R'(\hat{q})$, are served with certainty and types in $S^{\hat{q}^\ddagger} \setminus S^{\hat{q}^\dagger}$, with marginal revenue equal to $R'(\hat{q})$ are served with probability $\hat{q}-\hat{q}^\dagger/\hat{q}^\ddagger-\hat{q}^\dagger$. One way to achieve these service probabilities is to randomly order order the types by marginal revenue with ties broken randomly and to greedily serve the first $\hat{q}$ measure of types. Thus, all ex ante optimal mechanisms order types by marginal revenue and serve them greedily.

By revenue linearity the optimal mechanism for allocation constraint $\hat{y}$ is a convex combination of ex ante optimal mechanisms. As these ex ante optimal mechanisms all order the types greedily by marginal revenue with ties broken randomly, so does the optimal mechanism for $\hat{y}$. □

Theorem 8.1 says that while there is not an inherent ordering on type space that is respected by all mechanisms, there is one that, for all interim allocation constraints, is consistent with an optimal mechanism for the constraint.

Definition 8.6. For an orderable agent (Definition 8.5) and an implicit arbitrary total order on types that is consistent with the partial order on types, the *quantile q of a type t* is the probability that a random type $s \in \mathcal{T}$ from the distribution F precedes type t in the total order.

8.3.3 The Marginal Revenue Mechanism

We now define the marginal revenue mechanism for orderable agents.

Definition 8.7. The *marginal revenue mechanism for orderable agents* works as follows:

(i) Map the profile of agents' types $\boldsymbol{t}$ to a profile of quantiles $\boldsymbol{q}$ via Definition 8.6.
(ii) Calculate the profile of marginal revenues for the profile of quantiles.

(iii) Calculate a feasible allocation to optimize the surplus of marginal revenue, i.e., $\boldsymbol{x} = \operatorname{argmax}_{\boldsymbol{x}^\dagger} \sum_i x_i^\dagger R_i'(q_i) - c(\boldsymbol{x}^\dagger)$. For each agent i, calculate the supremum quantile $\hat{q}_i$ she could possess for which she is would be allocated in the above calculation of $\boldsymbol{x}$.

(iv) Offer each agent i the $\hat{q}_i$ optimal single-agent mechanism.

Theorem 8.5. *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. Consider any agent i. Analogously to the proof of Theorem 3.5, monotonicity of marginal revenue curves and Lemma 3.1 implies that, for every profile of reports of the other agents, there is a critical quantile $\hat{q}_i$ for agent i. The $\hat{q}_i$ ex ante mechanism is dominant strategy incentive compatible. Thus, the composition is dominant strategy incentive compatible.

The critical quantile $\hat{q}_i$ for agent i is at the boundary of service and non-service thus $R_i''(\hat{q}_i) \neq 0$ which implies that the $\hat{q}_i$ optimal mechanism, by Corollary 8.3, deterministically serves or rejects the agent. The set of agents served are exactly those served by $\boldsymbol{x}$ which is feasible.

The mechanism is revenue optimal because (a) its revenue is equal to its surplus of marginal revenue, (b) its surplus of marginal revenue is pointwise at least that of any other feasible mechanism, and (c), by revenue linearity, the revenue of any mechanism is at most its marginal revenue. □

While it is often assumed that the optimality of the marginal revenue mechanism is special to single-dimensional linear agents (as in Chapter 3), we have seen here that the condition required is revenue linearity not single dimensionality. It is instructive to contrast the simple optimal mechanism for revenue-linear agents to the complex optimal mechanism for non-revenue-linear agents that is derived in the next section.

8.4 The Interim Reduction

Without the revenue linearity property, which was assumed in the preceding section, single-agent interim optimal mechanisms can not be described solely in terms of the single-agent ex ante optimal mechanisms. For this reason, more sophisticated single-agent mechanisms are needed to enable the optimization of general multi-agent mechanisms.

In this section we characterize optimal multi-agent mechanisms for service constrained environments (Definition 8.4) in terms of the solution to single-agent interim optimal mechanism design problems (and without the simplifying revenue-linearity property). It is useful to contrast the complexity of the optimal mechanism for general preferences with that of the optimal mechanism with the revenue linearity assumption of Section 8.3.

The (quantile) allocation rule of a mechanism can be determined as follows. Recall that a single-agent mechanism $\mathcal{M}$ is given by an outcome rule $w(\cdot)$ which maps types to outcomes. The mechanism can alternatively be thought of as the menu $\{w(t) : t \in \mathcal{T}\}$. Encoded in an outcome $w(t)$ for type t is a probability of service denoted $x(t)$ and a payment denoted $p(t)$. For finite type spaces where $f(\cdot)$ denotes the probability mass function of the distribution F, the (quantile) allocation rule $y(\cdot)$ of the mechanism can be found by making a rectangle for each type t with height $x(t)$ and width $f(t)$ and sorting these rectangles in decreasing order of height. Equivalently and generally for continuous type spaces, consider the function defined as the measure of types that are served with at least a given service probability, the (quantile) allocation rule is the inverse of this function, i.e., $y(q) = \sup\{x^\dagger : \mathbf{Pr}_{t\sim F}[x(t) \geq x^\dagger] \geq q\}$.

Recall from Definition 8.3 that the single-agent interim optimal mechanism for allocation constraint $\hat{y}$ has allocation rule y that is no stronger than $\hat{y}$, i.e., the cumulative allocation rules satisfy $Y(\hat{q}) \leq \hat{Y}(\hat{q})$. Moreover, it optimizes revenue, denoted $\mathbf{Rev}[\hat{y}]$, over all such mechanism.

8.4.1 Symmetric Single-item Environments

We will illustrate the approach of this section with an example environment where symmetry enables the optimal mechanism to be easily identified.

Example 8.5. There are two agents competing for a single item each with private value t independently, identically, and uniformly distributed on $[0, 1]$ and a public budget of $B = 1/4$ (as in Example 8.1). The revenue-optimal mechanism fixes $(\hat{t}^\ddagger, \hat{t}^\dagger) \approx (0.32, 0.40)$, rejects agents with values less than $\hat{t}^\ddagger$, and allocates to the item to the remaining agent i for which $\min(t_i, \hat{t}^\dagger)$ is highest, breaking ties randomly. Each agent makes deterministic payment according to the interim allocation rule $x(\cdot)$. In particular, types $t_i \geq \hat{t}^\dagger$ pay B. See Figure 8.2 or Figure 8.8.

The main observation that enables the identification of an optimal mechanism in symmetric environments, i.e., with identically distributed agent types and symmetric feasibility constraint, is that convexity of the mechanism design problem, i.e., that convex combinations of mechanisms are valid mechanisms, implies that there is always an optimal mechanisms that is symmetric. The search for the optimal mechanism is then facilitated by symmetry.

Proposition 8.6. *In any symmetric environment there is an optimal mechanism that is symmetric.*

Proof. Consider any optimal incentive compatible mechanism that is asymmetric. Symmetry of the environment implies that permuting the identities of the agents gives a (potentially distinct) incentive compatible mechanism that is also optimal. The convex combination of, i.e., randomization over, incentive compatible mechanisms is incentive compatible. In particular, the convex combination of mechanisms for the uniform distribution over all permutations of the identities of agents is optimal, incentive compatible, and symmetric. □

We will separate the process of finding the (symmetric) optimal mechanism into two parts. The first part will be to identify a symmetric profile of allocation constraints $\hat{\boldsymbol{y}} = (\hat{y}, \ldots, \hat{y})$ that is feasible for a mechanism. The second part will be to find the single-agent mechanism, with allocation rule y, that is optimal for the identified constraint $\hat{y}$. The optimal multi-agent mechanism is then found by optimizing over the first part and combining with the second part. In the subsequent discussion, this approach is illustrated for the two-agent public-budget environment of Example 8.5.

The following theorem resolves the first part by identifying a single allocation constraint that is feasible and stronger than all other symmetric interim feasible allocation rules. In particular, the optimization over solutions to the first part is given by this allocation constraint.

Theorem 8.7. *Let $\hat{\boldsymbol{y}} = (\hat{y}, \ldots, \hat{y})$ be the n-agent allocation constraints induced by the k strongest-agents-win mechanism and $\boldsymbol{y} = (y, \ldots, y)$ the allocation rules induced by any symmetric k-unit mechanism for n i.i.d. agents, then y is feasible for $\hat{y}$.*

Proof. We prove the $n = 2$ agent $k = 1$ unit special case; the general result is left for Exercise 8.2. The two agent strongest-agent-wins mechanism induces allocation constraint $\hat{y}(q) = 1 - q$.

We argue, as follows, that $\hat{y}$ is the strongest symmetric allocation rule.

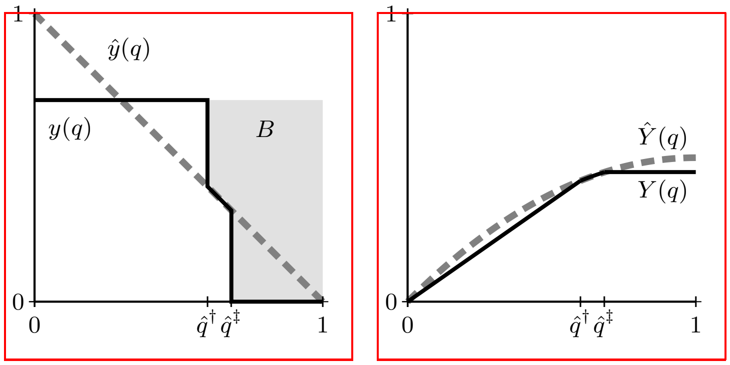

(a) $\hat{y}$ and y. (b) $\hat{Y}$ and Y.

Figure 8.8. Depicted in (a) are the allocation constraint (thick, gray, dashed line) and allocation rule (thin, black, solid line) for the optimal mechanism of the two-agent public budget environment of Example 8.5. Similarly, in (b) are the cumulative allocation rule and constraint.

For one of the agents, the probability she has quantile stronger than $\hat{q}$ is $\hat{q}$ while the probability she has quantile stronger than $\hat{q}$ and is served by the mechanism is $\hat{Y}(\hat{q}) = \int_0^{\hat{q}} \hat{y}(q)\,\mathrm{d}q = \hat{q} - 1/2\,\hat{q}^2$. For both agents, the probability at least one has quantile stronger than $\hat{q}$ is

$$1 - \mathbf{Pr}[q \geq \hat{q}]^2 = 1 - (1-\hat{q})^2 = 2\,\hat{q} - \hat{q}^2. \tag{8.4}$$

Meanwhile and by linearity of expectation, the expected number of agents with stronger quantile than $\hat{q}$ that are served is

$$2\hat{Y}(\hat{q}) = 2\,\hat{q} - \hat{q}^2. \tag{8.5}$$

Importantly the quantity of (8.4) is equal to the quantity of (8.5).

Consider the allocation rule y induced by any symmetric mechanism. For a contradiction, assume than y is stronger than $\hat{y}$ at $\hat{q}$, i.e., with $Y(\hat{q}) > \hat{Y}(\hat{q})$. This allocation rule cannot result from any ex post feasible mechanism as the analogous quantity to (8.5) for y would exceed quantity (8.4). Because there is only one unit available, it is impossible for the expected number of units allocated to agents with quantile in $[0, \hat{q}]$ to be greater than the probability that there is at least one such agent. □

For the second part, recall the interim optimal mechanism for the uniform public-budget agent (Example 8.1) and allocation constraint $\hat{y}(q) = 1 - q$ described in Section 8.1.1. It it has allocation rule $y(\cdot)$ that is equal to $\hat{y}(\cdot)$ except that it irons quantiles on interval $[0, \hat{q}^{\dagger} \approx 0.60]$ and rejects those below quantile reserve $\hat{q}^{\ddagger} \approx 0.68$. See Figure 8.8. For the uniform public-budget agent, these quantiles map back to type space via $t = 1 - q$ as $\hat{t}^{\dagger} = 0.40$ and $\hat{t}^{\ddagger} = 0.32$.

These two parts can be easily combined to observe that the mechanism described in Example 8.5 is revenue optimal. Shortly, in Section 8.4.4 we will want to generalize this construction, so it is instructive to see how we might construct the optimal mechanism knowing nothing about $\hat{y}$ except that there is an ex post mechanism, i.e., a mapping from quantile profiles to an ex post allocation that is feasible, that induces the interim allocation $\hat{y}$, and that we want to iron the strongest $[0, \hat{q}^{\dagger}]$ quantiles and reject the weakest $(\hat{q}^{\ddagger}, 1]$ quantiles. Denote the ex post allocation for quantile profile $\boldsymbol{q}$ by $\hat{\boldsymbol{y}}^{EP}(\boldsymbol{q})$; e.g., the two-agent strongest-quantile-wins ex post allocation is

$$\hat{y}_i^{EP}(\boldsymbol{q}) = \begin{cases} 1 & \text{if } q_i < q_{3-i}, \\ 0 & \text{otherwise.} \end{cases}$$

The following steps suffice to convert this ex post allocation $\hat{\boldsymbol{y}}^{EP}(\boldsymbol{q})$ that implements $\hat{\boldsymbol{y}} = (\hat{y}, \hat{y})$ to an ex post allocation rule $\boldsymbol{y}^{EP}(\boldsymbol{q})$ that implements the desired $\boldsymbol{y} = (y, y)$.

(i) Calculate $\boldsymbol{q}^{\dagger}$ by ironing on $[0, \hat{q}^{\dagger}]$ as

$$q_i^{\dagger} = \begin{cases} U[0, \hat{q}^{\dagger}] & \text{if } q_i \in [0, \hat{q}^{\dagger}], \\ q_i & \text{otherwise.} \end{cases}$$

(ii) Calculate $\boldsymbol{y}$ with quantile reserve $\hat{q}^{\ddagger}$ as

$$y_i = \begin{cases} \hat{y}_i^{EP}(\boldsymbol{q}^{\dagger}) & \text{if } q_i \in [0, \hat{q}^{\ddagger}], \\ 0 & \text{otherwise.} \end{cases}$$

These operations are easy to interpret on the cumulative allocation rule (see Figure 8.8). We are given the cumulative allocation constraint $\hat{Y}$ and we wish to implement cumulative allocation rule Y that satisfies $Y(q) \leq \hat{Y}(q)$ for all $q \in [0, 1]$; both the cumulative constraint and rule are concave. Ironing by resampling quantiles on an interval replaces the original curve with a line segment. A quantile reserve replaces the original curve with a horizontal line from the quantile reserve and over weaker quantiles. Combinations of these operations can produce any such Y from any such $\hat{Y}$.

The following proposition summarizes and generalizes the discussion of optimal mechanism for symmetric single-item environments. In the next section, these methods are further generalized to asymmetric environments.

Proposition 8.8. *For symmetric n-agent single-item environments, the optimal mechanism has expected revenue $n\,\mathbf{Rev}[\hat{y}]$ with allocation constraint $\hat{y}(q) = (1-q)^{n-1}$. For uniform public budget preferences, there exists quantiles $0 \leq \hat{q}^{\dagger} \leq \hat{q}^{\ddagger} \leq 1$ such that the optimal mechanism irons the strongest $[0, \hat{q}^{\dagger}]$ and reserve prices the weakest $(\hat{q}^{\ddagger}, 1]$ quantiles.*

8.4.2 Interim Feasibility

Consider any multi-agent mechanism $\boldsymbol{\mathcal{M}}$. 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 at the onset of Section 8.4; 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 (in particular, $\boldsymbol{\mathcal{M}}$) that induces it.[2] Note that interim feasibility is unrelated to incentives, any function $\boldsymbol{\mathcal{M}}$ that maps profiles of types $\boldsymbol{t}$ to feasible allocations induces interim feasible allocation rules.

Definition 8.8. A profile of allocation rules $\boldsymbol{y}$ is *interim feasible* if it is induced by some ex post feasible mechanism $\boldsymbol{\mathcal{M}}$ and type distribution $\boldsymbol{F}$.

Example 8.6. Consider selling a single item to one of two agents each with one of two interim allocation rules:

$$y^{\dagger}(q) = 1/2, \qquad\qquad y^{\ddagger}(q) = \begin{cases} 1 & \text{if } q \in [0, 1/2], \\ 0 & \text{otherwise.} \end{cases} \tag{8.6}$$

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

[2] Such a profile of interim allocation rules is sometimes also called the *reduced form* of the mechanism.

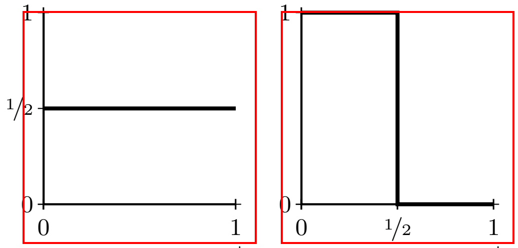

(a) Allocation rule $y^\dagger$. (b) Allocation rule $y^\ddagger$.

Figure 8.9. The allocation rules $y^\dagger$ and $y^\ddagger$ demonstrating interim feasibility. For single-item environments $(y^\dagger, y^\dagger)$ and $(y^\dagger, y^\ddagger)$ are feasible while $(y^\ddagger, y^\ddagger)$ is infeasible.

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 *serial-dictator* mechanism that serves agent 1 if she has a high type (i.e., $q_1 \in [0, 1/2]$) and agent 2 otherwise. The *serve-high-types* profile of interim allocation rules $\boldsymbol{y} = (y^\ddagger, y^\ddagger)$, on the other hand, 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 as there is only one item.

Our goal is to maximize expected revenue over (Bayesian incentive compatible and interim individually rational) mechanisms subject to ex post feasibility. Decomposing this goal into optimization of single-agent revenue subject to interim feasibility, we obtain the following program.

$$\max_{\hat{\boldsymbol{y}}} \ \sum\nolimits_i \mathbf{Rev}[\hat{y}_i] \qquad (8.7)$$
$$\text{s.t. “}\hat{\boldsymbol{y}} \text{ is interim feasible.”}$$

Recall that the optimal revenue for a single agent as solved by $\mathbf{Rev}[\cdot]$ is a convex optimization problem and thus $\mathbf{Rev}[\cdot]$ is concave, i.e., $\mathbf{Rev}[\hat{y}^\dagger + \hat{y}^\ddagger] \geq \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger]$. Observe that while the constraint of interim feasibility on $\hat{\boldsymbol{y}}$ is somewhat opaque at this point, it is nonetheless a convex constraint. Simply, the convex combination of two interim feasible mechanisms is interim feasible. The ex post mechanism that implements the convex combination is exactly the convex combination of the ex post mechanisms that implement the two original mechanisms. It will be the task of the remainder of this section to further elucidate the constraint imposed by interim feasibility.

The following proposition shows that the revenue optimal mechanism can be found by optimizing expected revenue over profiles of allocation constraints subject to interim feasibility.

Proposition 8.9. *The optimal multi-agent revenue is given by optimizing single-agent revenue subject to interim feasibility, i.e., solving program* (8.7)*.*

Proof. We will argue the two directions of this proof separately. First note that the optimal revenue from the program (8.7) is at least the revenue of the optimal mechanism. To see this, observe that any mechanism induces a profile of interim allocation rules $\boldsymbol{y}$. The ex post feasibility of this mechanism implies that this profile of allocation rules is interim feasible. The revenue from each agent i in this mechanism is at most the revenue of the interim optimal mechanism subject to allocation rule y_i as a constraint, i.e., at most $\mathbf{Rev}[y_i]$. Thus, the program upper bounds the optimal revenue.

For the other direction we will construct, from any ex post mechanism $\boldsymbol{\mathcal{M}}$ that induces the profile $\hat{\boldsymbol{y}}$ of interim allocation rules that attains the maximum of the program and each agent i's $\hat{y}_i$-optimal mechanism $\mathcal{M}_i$, a Bayesian incentive compatible mechanism with revenue equal to the revenue of the program, i.e., $\sum_i \mathbf{Rev}[\hat{y}_i]$. The remainder of this proof is deferred to Section 8.4.4 where the construction is generalized by Definition 8.13 and shown to be correct by Theorem 8.18. ☐

Optimization of mathematical program (8.7) in asymmetric environments 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., $Y_1(\hat{q}_1), \ldots, Y_n(\hat{q}_n)$. The expected number of these agents served is thus $\sum_i Y_i(\hat{q}_i)$. Of course the probability that one or more agents agent i with quantile bounded by $\hat{q}_i$ are realized is $1 - \prod_i (1 - \hat{q}_i)$.[3] Given the single-item ex post feasibility constraint that allows only one such agent to be served at once, the expected number served must be at least the probability that at least one is realized. In fact this necessary condition is also sufficient, as we will see by the max-flow-min-cut

[3] The probability that one or more such agents show up is one minus the probability that none show up.

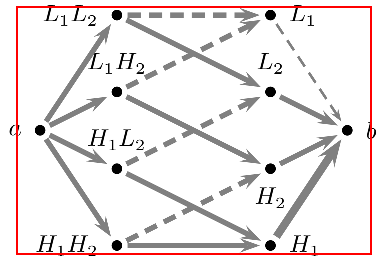

(a) Feasible

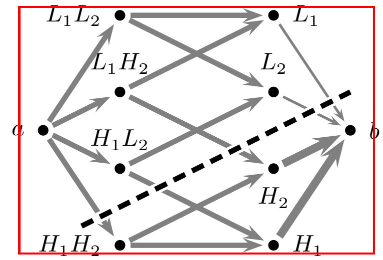

(b) Infeasible

Figure 8.10. The flow constructions are illustrated for (a) the feasibility of the serial-dictator mechanism and (b) the infeasibility of the serve-high-types profile of interim allocation rules from Example 8.6. In both flow graphs the edges depicted in heavy, medium, and light weight correspond to capacities 1/2, 1/4, and zero, respectively. To translate between the quantile-space allocation rules of Example 8.6 and the type-space allocation rules in the proof of Theorem 8.10, type H will correspond to strong quantiles $[0, 1/2]$ and type L corresponds to weak quantiles $(1/2, 1]$. Subfigure (a) depicts the flow graph that corresponds to the profile of interim allocation rules $\boldsymbol{y} = (y^{\ddagger}, y^{\dagger})$ and flow (solid gray edges) that corresponds to the ex post allocation rule of the serial-dictator mechanism. Recall that this serial-dictator mechanism allocates to agent 1 if she has a high type H_1 and to agent 2 otherwise. This ex post allocation rule can be determined by inspecting the out-going flow from vertices corresponding to type profiles, i.e., the left-side column. Subfigure (b) depicts the flow graph that corresponds to the profile of interim allocation rules $\boldsymbol{y} = (y^{\ddagger}, y^{\ddagger})$ which require that an agent is served if an only if she has a high type. This profile is infeasible, which can be seen as the minimum a–b cut, depicted with the black dashed line, has cost 3/4 (it cuts three edges with capacity 1/4 each and two edges with capacity zero) while the total capacity of edges incident on sink b is one. Inequality (8.9) is violated for subsets of types $\boldsymbol{S}^{\star}$ with $S_i^{\star} = \{H_i\}$ for each i, i.e., corresponding to the vertices in the right-side column that are on the sink b side of the cut.

style argument of the proof below. The following theorem is often called Border's Theorem in recognition of Kim Border's pioneering study of interim feasibility.

Theorem 8.10. *For single-item environments, a profile of allocation rules* $\boldsymbol{y}$ *(with cumulative allocation profile* $\boldsymbol{Y}$*) is interim feasible if and only if,*

$$\sum\nolimits_i Y_i(\hat{q}_i) \leq 1 - \prod\nolimits_i (1 - \hat{q}_i), \qquad \forall \hat{\boldsymbol{q}} \in [0,1]^n. \tag{8.8}$$

Proof. This proof is most instructive to see in type space. For finite type spaces, the inequality (8.8) of the theorem is equivalent to the following (see Exercise 8.7). For any subsets of agents' types $S_i \subset \mathcal{T}_i$ for all i,

$$\sum_i \sum_{t_i \in S_i} x_i(t_i)\, f_i(t_i) \leq 1 - \prod_i (1 - f_i(S_i)) \tag{8.9}$$

with $f_i(t_i)$ denoting the probability that agent i has type t_i and $f_i(S_i) = \sum_{t_i \in S_i} f_i(t_i)$ denoting the probability that i has type $t_i \in S_i$. Notice that the left-hand side of the equation is the expected number of items allocated to agents i with types $t_i \in S_i$. Notice that the right-hand side of the equation is simply the probability that one or more agents i have realized type $t_i \in S_i$. As describe above, the necessity of the condition for interim feasibility is straightforward.

The following argument shows sufficiency, specifically, that if a profile of allocation rules is infeasible that there exists subsets of agents' types $S_1^\star \subset \mathcal{T}_1, \ldots, S_n^\star \subset \mathcal{T}_n$ for which inequality (8.9) is violated. The approach of this proof is (a) to show that interim feasibility is equivalent to whether a specific cut in a network flow graph is the minimum cut, and (b) to use the minimum cut corresponding to interim infeasible allocation rules to identify the subsets of agents' types that violate inequality (8.9).

Consider the following network flow problem, equivalently a weighted directed graph where the weights are referred to as capacities; see Figure 8.10. This graph is defined on the following vertices (left to right in Figure 8.10):

- a source vertex a,
- a vertex $\boldsymbol{t}$ for each type profile $\boldsymbol{t} \in \boldsymbol{\mathcal{T}}$,
- a vertex t_i for each type $t_i \in \mathcal{T}_i$ of each agent i, and
- a sink vertex b.

Directed weighted edges connect these vertices as follows (left to right in Figure 8.10):

- source a is connected to each vertex $\boldsymbol{t}$ with capacity $f(\boldsymbol{t}) = \prod_i f_i(t_i)$, i.e., the probability that type profile $\boldsymbol{t}$ is realized;
- each vertex $\boldsymbol{t}$ is connected to vertex t_i for each i with capacity $f(\boldsymbol{t})$; and
- each vertex t_i is connected to sink b with capacity $x_i(t_i)\, f_i(t_i)$, i.e., the probability that agent i has type t_i and is allocated by the interim allocation rule x_i.

A profile of interim allocation rules $\boldsymbol{x}$ is feasible if any only if there is a flow in the flow graph constructed above that saturates all edges incident on the sink b; see Figure 8.10(a). For the "only if" direction, consider any ex post feasible mechanism that induces interim allocation rules $\boldsymbol{x}$ and construct a flow as follows. Flow from source a to vertex $\boldsymbol{t}$ represents the probability that type profile $\boldsymbol{t}$ is realized. The flow from vertex $\boldsymbol{t}$ to vertices t_i for each i represents the probability that $\boldsymbol{t}$ is realized and agent i is served by the ex post mechanism. Since the total flow into vertex $\boldsymbol{t}$ is $f(\boldsymbol{t})$ the cumulative flow out can be at most $f(\boldsymbol{t})$ which satisfies the ex post feasibility constraint that at most one of the agents is served. The flow on the edge from vertex $\boldsymbol{t}$ to vertex t_i is $x_i(\boldsymbol{t})\, f(\boldsymbol{t})$ as follows. Vertex t_i aggregates flow from each type profile $\boldsymbol{t}$ containing t_i and thus the flow that can go from vertex t_i to sink b is the cumulative probability that t_i is realized and is served, i.e., $x_i(t_i)\, f_i(t_i)$. Thus, the edges incident on sink b are saturated. For the "if" direction, given any flow that saturates all the edges incident on the sink b, an ex post mechanisms can be inferred. The ex post allocation on type profile $\boldsymbol{t}$ picks an agent with probability equal to the flow from $\boldsymbol{t}$ to t_i divided by $f(\boldsymbol{t})$.

Non-existence of a flow that saturates the edges incident on sink b implies that the profile of allocation rules $\boldsymbol{x}$ is infeasible. We will now show that this non-existence of a flow will enable us to identify subsets of types $\boldsymbol{S}^\star = (S_1^\star, \ldots, S_n^\star)$ that violate inequality (8.9) and thus the inequality is sufficient for the interim feasibility of $\boldsymbol{x}$; see Figure 8.10(b). An a–b cut in a directed graph is partitioning of the vertices into two sets $\{a\} \cup A$ and $\{b\} \cup B$. The capacity of the cut is the sum of the capacities of edges that cross from $\{a\} \cup A$ to $\{b\} \cup B$. The proof will show that the inequality (8.9) is satisfied for interim allocation rules $\boldsymbol{x}$ only if $B = \emptyset$ is minimum capacity a–b cut, i.e., the capacity of the minimum cut is equal to the sum of the capacities of edges from vertices $t_i \in \mathcal{T}_i$ and all i to vertex b, specifically, $\sum_i \sum_{t_i \in \mathcal{T}_i} x_i(t_i)\, f_i(t_i)$.

Observe that the value of the maximum a–b flow in the graph is upper bounded by the capacity of any a–b cut in the graph. Simply, there is no way to get more flow across this cut than the total capacity of the cut. More precisely, the well known max-flow min-cut theorem states that the value of the maximum a–b flow in a flow graph is equal to the capacity of its minimum a–b cut. We now show that there is a flow that saturates the edges incident on sink b, equivalently, that $B = \emptyset$ is a minimum cut, if and only if there is no profile of subsets of type space $(S_1, \ldots, S_n)$ for which inequality (8.9) is violated.

We will calculate the difference between the capacity of the $B = \emptyset$ cut, i.e., the capacity edges incident on sink b, and the minimum cut.

When this difference is strictly positive we will identify a violation of inequality (8.9). Denote by $(A^\star, B^\star)$ the minimum capacity a–b cut. The subsets of each agent's type space that are candidates for violation of inequality (8.9) are $S_i^\star = B^\star \cap \mathcal{T}_i$. The difference between the capacities of these two cuts is calculated as follows.

- For edges incident on sink b: The capacity of the cut (A, B) (with $B = \emptyset$) is equal to the sum of the capacities of edges incident on sink b. Subtracting from this the sum of capacities of edges crossing cut $(A^\star, B^\star)$, the difference is the sum of capacities of edges that are not cut by $(A^\star, B^\star)$. These uncut edges are the ones from vertices in $B^\star$ to sink b which correspond to types $t_i \in S_i^\star$ for all i. The total contribution from these edges to the difference is thus, $\sum_i \sum_{t_i \in S_i^\star} x_i(t_i)\, f_i(t_i)$, i.e., the left-hand side of inequality (8.9).
- For edges incident on vertices $\boldsymbol{t} \in \boldsymbol{\mathcal{T}}$: Vertices corresponding to type profiles, e.g., $\boldsymbol{t}$, are either in $B^\star$, in which case we have cut the edge from source a and must subtract $f(\boldsymbol{t})$, or in $A^\star$, in which case we have cut edges with capacity $f(\boldsymbol{t})$ for each i with $t_i \in S_i^\star$ (i.e., with $t_i \in B^\star$) and must subtract $f(\boldsymbol{t}) \cdot |\{i : t_i \in S_i\}|$. Since $(A^\star, B^\star)$ is a minimum cut, we must have chosen the smaller of these two quantities, i.e., $f(\boldsymbol{t}) \cdot \min(1, |\{i : t_i \in S_i\}|)$. Summing this quantity to be subtracted over all type profiles $\boldsymbol{t}$ equates to the probability that one or more types $t_i \in S_i^\star$ are realized, i.e., the right-hand side of inequality (8.9) of $1 - \prod_i (1 - f_i(S_i^\star))$.

Combine these two contributions to the difference, and observe that when the difference is strictly positive then inequality (8.9) is violated for the subsets of types $S_1^\star, \ldots, S_n^\star$. □

This characterization of interim feasibility extends naturally to matroid environments where the right-hand side becomes the *expected rank*, with short-hand notation $\mathrm{rank}(\hat{\boldsymbol{q}})$ representing $\mathbf{E}_S[\mathrm{rank}(S)]$ where each i is in S independently with probability $\hat{q}_i$ (cf. Section 4.3).

Theorem 8.11. *For matroid environments, a profile of allocation rules $\boldsymbol{y}$ is interim feasible if and only if,*

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

One way this characterization of interim feasibility is helpful is as follows. It can be shown that the matroid rank function $\mathrm{rank}(\cdot)$ is submodular. This submodularity implies that the interim feasibility constraint has a polymatroidal structure which, in turn, implies that the vertices corresponding to the feasible region can be implemented by greedily ordering types and serving each type to the maximum extent possible. Instead of introducing this polymatroidal theory of optimization, we will give an alternative first-principles proof of this result in the next section (see Corollary 8.15).

For feasibility constraints beyond matroid, we will not get a succinct formula like inequality (8.8) in Theorem 8.10 that characterizes interim feasibility. Nonetheless, in the next section we will describe a simple family of ex post mechanisms from which any profile of interim feasible allocation rules can be derived.

8.4.3 Interim Feasibility by Stochastic Weighted Optimization

In this section we show that, for any service constrained environment, any interim feasible profile of allocation rules is implementable as a stochastic weighted optimization. This characterization is derived by observing that:

(i) there is an isomorphism between profiles of interim allocation rules to points in a high dimensional Euclidean space,

(ii) the set of interim feasible points by this isomorphism is convex and, specifically, a polytope,

(iii) any point in the interior of this polytope can be given as a convex combination of points on the exterior of the set, specifically, vertices of the polytope, and

(iv) these vertices can be implemented by weighted optimization, i.e., a mapping of each type of each agent in the profile to a weight and selection of the ex post feasible set of agents with highest cumulative weight.

We relax two constraints from the Bayesian mechanism design problem. Relaxing incentive compatibility, allocation rules need not be monotone and we will, thus, work in type space rather than quantile space.[4] Relaxing the independence of the distribution of types across agents, we

[4] The approach we take is similar to that of the characterization of interim feasibility for single-item and matroid environments (Theorem 8.10 and Theorem 8.11) which was also described in type space. At the end of this section we will describe how to modify the construction for quantile space and when this approach is helpful.

draw type profile $\boldsymbol{t}$ from joint distribution $\boldsymbol{F}$ and denote by $f_i(t)$ the marginal probability that agent i has type t, i.e., $\mathbf{Pr}_{\boldsymbol{t}\sim\boldsymbol{F}}[t = t_i]$.[5]

For Step (i), consider a finite space of type profiles $\boldsymbol{\mathcal{T}} = \mathcal{T}_1 \times \cdots \times \mathcal{T}_n$ with size $\ell = \sum_i |\mathcal{T}_i|$ and map profiles of interim allocation rules $\boldsymbol{x}$, with $x_i : \mathcal{T}_i \to [0,1]$, to points $\boldsymbol{z} \in [0,1]^\ell$, in a high dimensional Euclidean space. In this mapping the vector $\boldsymbol{z}$ will be indexed by agent-type pairs it as z_{it}.

Definition 8.9. For joint type space $\boldsymbol{\mathcal{T}}$ with size $\ell = \sum_i |\mathcal{T}_i|$ and joint distribution $\boldsymbol{F}$, the *flattened ex post allocation rule* $\boldsymbol{z}^{EP} : \boldsymbol{\mathcal{T}} \to [0,1]^\ell$ and *flattened interim allocation* $\boldsymbol{z} \in [0,1]^\ell$ are induced by ex post and interim allocation rules $\boldsymbol{x}^{EP}$ and $\boldsymbol{x}$ and indexed by it for agent i and type $t \in \mathcal{T}_i$ as:

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

Notice that, by the above definition, the flattened interim allocation is in fact specifying the ex ante probability that each type of each agent is served. The normalization by the density function in the definition of the flattened interim allocation serves a similar purpose in the geometry of interim feasibility as the mapping of types to quantiles. Definition 8.9 is useful as it immediately gives the following propositions; the second of which concludes Step (ii).

Proposition 8.12. *The flattened interim allocation is the expectation of the flattened ex post allocation rule:*

$$\boldsymbol{z} = \mathbf{E}_{\boldsymbol{t}\sim\boldsymbol{F}}\big[\boldsymbol{z}^{EP}(\boldsymbol{t})\big]. \tag{8.10}$$

Proposition 8.13. *For joint type space $\boldsymbol{\mathcal{T}}$ with size $\ell = \sum_i |\mathcal{T}_i|$ and distribution $\boldsymbol{F}$, the space $\mathcal{Z} \subset [0,1]^\ell$ of feasible flattened interim allocations $\boldsymbol{z}$ is convex.*

Proof. Randomized mechanisms are feasible, and flattened interim allocations are linear with respect to convex combinations. ☐

The flattened ex post allocation rule is a redundant representation of the ex post allocation rule, it specifies allocation probabilities for all types an agent might possess. For a given type profile, all of these probability must be zero except for the ones that correspond to types in the

[5] The relaxation to (possibly) correlated distributions over type profiles will allow the results of this section to generalize beyond service constrained environments as in Section 8.5.

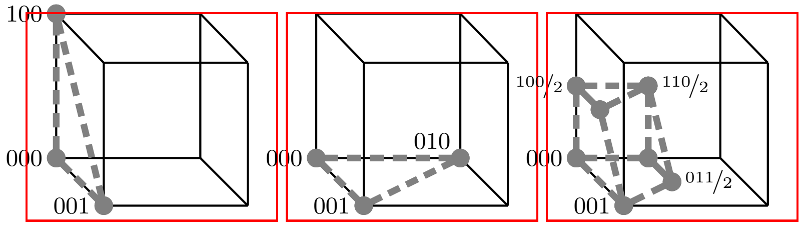

(a) $\mathcal{Z}(H,M)$ (b) $\mathcal{Z}(L,M)$ (c) $\mathcal{Z}$

Figure 8.11. The polytopes that define the feasible flattened allocations of Example 8.7 are depicted as convex subsets of the unit cube. Selected vertices of these polytopes are labeled in short-hand as bit vectors and correspond to $\boldsymbol{z} = (z_{1H}, z_{1L}, z_{2M})$. The vertical axis is z_{1H}; the horizontal axis is z_{1L}; and the outward axis is z_{2M}. The ex post feasible flattened allocations for type profiles $\boldsymbol{t} = (H, M)$ and $\boldsymbol{t} = (L, M)$ are depicted in figures (a) and (b), respectively. The interim feasible flattened allocations when t_1 is uniform on $\{L, H\}$ and $t_2 = M$ (deterministically) are depicted in figure (c).

type profile. Specifically, consider an agent-type pair it with $t \in \mathcal{T}_i$ and type profile $\boldsymbol{t}$, if $t_i \neq t$ in type profile $\boldsymbol{t}$ then $z_{it}^{EP}(\boldsymbol{t}) = 0$ as a mechanism cannot serve a type that "does not show up." The types served must also satisfy the feasibility constraint of the service constrained environment.

Denote the feasibility constraint imposed by the service constrained environment by $\mathcal{X}$. Randomized mechanisms are allowed, thus $\mathcal{X}$ is convex. For example in a single-item environment deterministic outcomes correspond to allocating to a single agent i or not allocating, the convex closure of these outcomes gives $\mathcal{X} = \{\boldsymbol{x} \in [0,1]^n : \sum_i x_i \leq 1\}$. Ex post feasibility for flattened allocation rules $z^{EP}(\boldsymbol{t})$ is the projection of the service constrained feasibility constraint onto the indices $\{it_i\}_{i\in[n]}$ of the types in the given type profile $\boldsymbol{t}$. Specifically, for any profile $\boldsymbol{t}$, $\boldsymbol{z}^{EP}(t) \in \mathcal{Z}(t)$ is ex post feasible, where

$$\mathcal{Z}(t) = \left\{\boldsymbol{z} \in [0,1]^\ell : \bigtimes_i z_{it_i} \in \mathcal{X} \wedge z_{\{it\,:\,t\neq t_i\}} = \mathbf{0}\right\}.$$

This projection is depicted in Figure 8.11 (for Example 8.7, below).

Example 8.7. Consider two agents and a single-item environment. Agent 1 has type t_1 uniformly drawn from $\mathcal{T}_1 = \{L, H\}$; agent 2 has type deterministically $t_2 = M$ (i.e., with $\mathcal{T}_2 = \{M\}$). A flattened allocation is $\boldsymbol{z} = (z_{1H}, z_{1L}, z_{2M})$. Ex post feasible allocations for type profile $\boldsymbol{t} = (H, M)$ are convex combinations of $\{(0,0,0), (1,0,0), (0,0,1)\}$; ex post feasible allocations for type profile $\boldsymbol{t} = (L, M)$ are convex combinations of $\{(0,0,0),(0,1,0),(0,0,1)\}$. Interim feasible flattened allocations are convex combinations of $\{(0,0,0),(1/2,0,0),(0,1/2,0),(0,0,1),(1/2,1/2,0),(1/2,0,1/2),(0,1/2,1/2)\}$. The vertices of the interim feasible polytope are given by the ordinal mechanisms, i.e., where there is an ordered subset of types and after the types are realized the first one in order receives the item (or none if no type in the order is realized); see Corollary 8.15. The ordered subsets that correspond to the vertices above are $\{\emptyset,(1H),(1L),(1M),(1H,1L),(1H,2M),(1L,2M)\}$. See Figure 8.11.

We have seen that interim feasibility is convex (Proposition 8.13); as previously observed, the single-agent optimal revenues given by $\mathbf{Rev}[\cdot]$ are concave. Suppose that, instead of the concave objective given by the sum of the single agent revenues, the objective was linear and given by weights $\boldsymbol{w} \in \mathbb{R}^{\ell}$ indexed in the flattened space by agent-type pair it. The optimization of the expected *surplus of weights*, i.e., $\sum_{it} z_{it}\, w_{it}$, subject to interim feasibility, is achieved by optimizing the surplus of weights pointwise for each profile of types $\boldsymbol{t}$ subject to ex post feasibility. Moreover, for any such weights (which we view as a direction in the flattened space) the corresponding interim feasible allocation vector, denoted $\boldsymbol{z}^{\boldsymbol{w}}$, is given by Proposition 8.12.

A *vertex* of a convex subset $\mathcal{Z}$ of ℓ-dimensional Euclidean space is a point that is uniquely optimal for some direction. Vertices can be specified equivalently as the direction, e.g., $\boldsymbol{w}$, or the point, e.g., $\boldsymbol{z}^{\boldsymbol{w}}$. Any other point in $\mathcal{Z}$ can be represented as a convex combination of $\ell+1$ vertices; therefore, we can implement any interim feasible allocation rule by sampling a direction from a distribution over $\ell+1$ vectors of weights, and then for the type profile realized, optimizing the weights given by the direction subject to ex post feasibility.

A weights $\boldsymbol{w}$ in the space of flattened allocation rules correspond, in the original space of allocation rules, to a profile of functions that map each type to a weight. The following theorem summarizes the construction above in the original space of allocation rules.

Definition 8.10. A *stochastic weighted optimizer* is given by a joint distribution over profiles of weight functions $\boldsymbol{w}$ with $w_i : \mathcal{T}_i \to \mathbb{R}$ as follows for type profile $\boldsymbol{t}$:

(i) Draw weight functions $\boldsymbol{w}$ from the distribution.
(ii) Output allocation $\boldsymbol{x} = \operatorname{argmax}_{\boldsymbol{x}^\dagger \in \mathcal{X}} \sum_i w_i(t_i)\, x_i^\dagger$.

Theorem 8.14. *For any joint distribution on type profiles and service* *constrained environment, any interim feasible allocation profile can be ex post implemented by a stochastic weighted optimization.*

In the special case that the service constrained environment is ordinal, e.g., multi-unit environments, matroid environments, and position environments, the surplus of weights is optimized by the greedy algorithm (See Section 4.6 on page 129). The greedy algorithm, by definition, considers only the order of weights of each type and not magnitudes of the weights. The following corollary refines Theorem 8.14 for ordinal environments.

Definition 8.11. A *stochastic ordered-subset algorithm* is given by a joint distribution over ordered subsets of joint type space $\bigcup_i \mathcal{T}_i$ as follows for type profile $\boldsymbol{t}$:

(i) Draw an ordered subset from the distribution.
(ii) Output allocation $\boldsymbol{x}$ obtained by the greedy algorithm on agents ordered by the rank of their types in the ordering; agents whose types are not present in the subset are discarded.

Corollary 8.15. *For any joint distribution on type profiles and service constrained environment that is given by the independent sets of a matroid, any interim feasible allocation profile can be ex post implemented by a stochastic ordered-subset algorithm.*

This characterization of interim feasible allocation rules as a convex subset of high-dimensional Euclidean space is central to the design of computationally-efficient revenue-optimal mechanisms. The main challenges to be resolved is in quickly finding the distribution over weights Theorem 8.14. Discussion of the computational issues involved are deferred to Section 10.5.

Though this section approached the characterization of interim feasibility in type space, it can be equivalently characterized in quantile space as well. For example, discretize quantile space into intervals and apply the construction in this section with each discrete interval for each agent as a type. This approach is advantageous when the agents' type spaces are very large, or high dimensional as quantiles are always single dimensional. Again, further discussion is deferred to Section 10.5.

8.4.4 Combining Ex post Feasibility and Bayesian Incentive Compatibility

This section formalizes the general constructions for ex post implementation of interim mechanisms. Let $\hat{\boldsymbol{\mathcal{M}}}$ denote an ex post feasible mechanism (that is not necessarily incentive compatible or possessing revenue guarantees). Its ex post allocation rule maps quantile profiles to distributions over ex post feasible allocations via $\hat{\boldsymbol{y}}^{EP} : [0,1]^n \to \Delta(\mathcal{X})$. Recall from Section 2.4 that the induced interim allocation rule $\hat{y}_i : [0,1] \to [0,1]$ for agent i is defined as $\hat{y}_i(q_i) = \mathbf{E}_{\boldsymbol{q}}[\hat{y}_i^{EP}(\boldsymbol{q}) \mid q_i]$. Let $\boldsymbol{\mathcal{M}}$ denote a Bayesian incentive compatible mechanism (potentially with good revenue properties, but that is not necessarily ex post feasible). Its interim allocation rules are denoted by $\boldsymbol{y}$ with $y_i : [0,1] \to [0,1]$ for each agent i. We can compose these two mechanisms to obtain a mechanism with the ex post feasibility of $\hat{\boldsymbol{\mathcal{M}}}$ and the Bayesian incentive compatibility (and revenue properties) of $\boldsymbol{\mathcal{M}}$ if and only if the allocation rules $\boldsymbol{y}$ are feasible for allocation constraints $\hat{\boldsymbol{y}}$.

This construction can be instantiated with the revenue-optimal mechanisms of the previous section. In such an instantiation, $\hat{\boldsymbol{\mathcal{M}}}$ is the ex post mechanism that induces the interim allocation rules $\hat{\boldsymbol{y}}$ that optimize $\sum_i \mathbf{Rev}[\hat{y}_i]$ subject to interim feasibility, and $\boldsymbol{\mathcal{M}}$ is the profile of interim mechanism that optimize revenue subject to the allocation constraints $\hat{\boldsymbol{y}}$, i.e., with $\mathcal{M}_i$ as the $\hat{y}_i$ interim optimal mechanism.

First, given any single-agent allocation constraint $\hat{y} : [0,1] \to [0,1]$ and single-agent mechanism $\mathcal{M}$ with allocation rule y that satisfies constraint $\hat{y}$, we given an ex post implementation of y from $\hat{y}$. Second, given any ex post implementation $\hat{\boldsymbol{y}}^{EP}$ that induces interim constraints $\hat{\boldsymbol{y}}$ and a profile of single-agent mechanisms $(\mathcal{M}_1, \ldots, \mathcal{M}_n)$, where y_i satisfies $\hat{y}_i$ for each agent i, we give an ex post implementation of a mechanism with allocation rules $\boldsymbol{y}$.

Recall that a single-agent mechanism $\mathcal{M}$ is equivalently a menu of outcomes $\{w(t) : t \in \mathcal{T}\}$. A deterministic outcome is either a service outcome or a non-service outcome. Outcomes are closed under convex combination, i.e., they may be randomized. The (type) allocation rule $x : \mathcal{T} \to [0,1]$ gives a probability of service for each type $t \in \mathcal{T}$. For $x(t) \in [0,1]$, the outcome distribution $w(t)$ is a distribution over service and non-service outcomes. Denote by $w^x(t)$ the distribution of outcomes conditioned on the allocation $x \in \{0,1\}$.

Any single-agent mechanism $\mathcal{M}$ induces an ordering on types which in turn induces a mapping from types to quantiles. Correctness requires that the distribution of quantiles from this mapping be uniform on the $[0, 1]$ interval.

Definition 8.12. The *quantile mapping* for mechanism $\mathcal{M}$ with (type) allocation rule $x(\cdot)$ is $Q(\cdot)$ defined as follows. For any type t, calculate interval $[\hat{q}^\dagger, \hat{q}^\ddagger]$ as

$$\hat{q}^\dagger = \mathbf{Pr}_{t^\dagger \sim F}\left[x(t^\dagger) > x(t)\right], \qquad \hat{q}^\ddagger = \mathbf{Pr}_{t^\ddagger \sim F}\left[x(t) < x(t^\ddagger)\right].$$

The stochastic mapping from types to quantiles is Q defined as:

$$Q(t) \sim U[\hat{q}^\dagger, \hat{q}^\ddagger].$$

Lemma 8.16. *For types $t \in \mathcal{T}$ from distribution F, the quantile distribution of the quantile mapping $Q(t)$ (of Definition 8.12) is uniform on* $[0, 1]$.

From the above induced mapping from types to quantiles and a procedure for the allocation rule $y(\cdot)$ of mechanism $\mathcal{M}$, the mechanism can be implemented with $y(\cdot)$ as follows:

(i) Calculate the agent's quantile as $q = Q(t)$.
(ii) Calculate the agent's service as

$$x = \begin{cases} 1 & \text{w.p. } y(q) \\ 0 & \text{otherwise.} \end{cases}$$

(iii) Calculate the agent's outcome as $w = w^x(t)$.

To generalize this construction to procedures for allocation constraints $\hat{y}(\cdot)$ that allocation rule $y(\cdot)$ satisfies, we need to convert the procedure for $\hat{y}$ to a procedure for y.

Suppose we have an interim allocation constraint $\hat{y}$ and a mechanism $\mathcal{M}$ with allocation rule y that satisfies the constraint, i.e., $Y(q) \leq \hat{Y}(q)$ for all q. The following lemma shows that we can implement y from $\hat{y}$; thus, by the above construction, we can implement $\mathcal{M}$.

Lemma 8.17. *Any allocation rule y that satisfies allocation constraint $\hat{y}$ can be implemented by a* quantile reserve pricing $\hat{q}$ *and a stationary* quantile resampling transformation σ.

Proof. This proof will be by construction; see Figure 8.12. First, we will construct $\hat{y}^\dagger$ from $\hat{y}$ with a quantile reserve so as to equate the ex ante service probabilities $\hat{Y}^\dagger(1) = Y(1)$ while preserving feasibility of

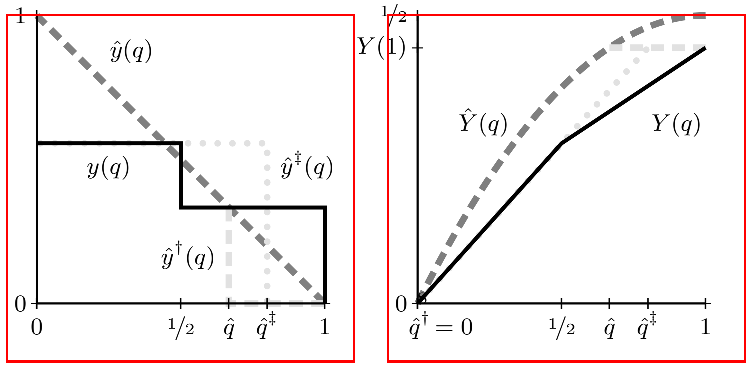

(a) Allocation Rules (b) Cumulative Allocation Rules

Figure 8.12. The reserve price and resampling transformation construction of Lemma 8.17 is depicted for a piecewise constant allocation rule y with $\ell = 2$ pieces. The allocation constraint $\hat{y}$ is thick, dashed, and dark gray, the allocation rule y is thin, solid, and black. The allocation constraint $\hat{y}^{\dagger}$ (think, dashed, and light gray) is constructed from $\hat{y}$ by reserve pricing at quantile $\hat{q}$. The allocation constraint $\hat{y}^{\ddagger}$ (think, dotted, and light gray) is constructed from $\hat{y}^{\dagger}$ by ironing on interval $[\hat{q}^{\dagger}, \hat{q}^{\ddagger}]$.

y for $\hat{y}^{\dagger}$, i.e., so $Y(q) \leq \hat{Y}^{\dagger}(q)$ for all q. Second, we will give a stationary quantile resampling transformation σ, i.e., with $\sigma(q)$ uniformly distributed on $[0, 1]$ if q is uniform on $[0, 1]$, that transforms $\hat{y}^{\dagger}$ to y, i.e., $y(q) = \mathbf{E}_{\sigma}[\hat{y}^{\dagger}(\sigma(q))]$.

The allocation constraint $\hat{y}^{\dagger}$ is obtained from allocation constraint $\hat{y}$ by quantile reserve pricing at $\hat{q} = \hat{Y}^{-1}(Y(1))$. Recall, quantile reserve pricing has the effect of replacing the cumulative allocation rule with a constant function after the quantile reserve; thus $\hat{Y}^{\dagger}(1) = Y(1)$.

Assume that y is piece-wise constant, equivalently that Y is piece-wise linear, with ℓ equal-width pieces. This assumption can be removed by considering y in the limit, as ℓ goes to infinity, of such a piece-wise constant allocation rule. The following inductive procedure gives a stationary resampling transformation σ that constructs y from $\hat{y}^{\dagger}$. Let $\hat{q}^{\dagger}$ be the lower end point of the first piece on which $\hat{Y}^{\dagger}(\cdot)$ and $Y(\cdot)$ are distinct, equivalently, where the upper end point q of the piece satisfies $\hat{Y}^{\dagger}(q) > Y(q)$. Note that the right slope of $Y(\cdot)$ at $\hat{q}^{\dagger}$ is equal to the right value of $y(\cdot)$ at $\hat{q}^{\dagger}$. Set $\hat{q}^{\ddagger} > \hat{q}^{\dagger}$ to the quantile at which the line through point $(\hat{q}^{\dagger}, Y(\hat{q}^{\dagger}))$ with slope $y(\hat{q}^{\dagger})$ next intersects the cumulative allocation constraint $\hat{Y}^{\dagger}(\cdot)$. Define $\sigma^{\dagger}$ to be the interval resampling transformation that irons on quantile interval $[\hat{q}^\dagger, \hat{q}^\ddagger]$, i.e.,

$$\sigma^\dagger(q) = \begin{cases} q^\dagger \sim U[\hat{q}^\dagger, \hat{q}^\ddagger] & \text{if } q \in [\hat{q}^\dagger, \hat{q}^\ddagger], \\ q & \text{otherwise.} \end{cases}$$

By the line-segment interpretation of ironing on the cumulative allocation rule, this resampling transformation gives an allocation constraint $\hat{y}^\ddagger(q) = \mathbf{E}_{\sigma^\dagger}[\hat{y}^\dagger(\sigma^\dagger(q))]$ with $\hat{Y}^\ddagger(q) = Y(q)$ for q in the piece.

By this construction $\hat{y}^\ddagger$ differs from y on (at least) one fewer piece than $\hat{y}^\dagger$. By induction we can construct a sequence of interval resampling transformations that, when composed, transform $\hat{y}^\dagger$ to y. The transformation σ in the statement of the lemma is this composition of interval resampling transformations. The stationarity of each interval resampling transformation implies that the transformation σ is stationary. □

Definition 8.13. For distribution $\boldsymbol{F}$ and mechanisms $\hat{\boldsymbol{\mathcal{M}}}$ and $\boldsymbol{\mathcal{M}}$ with allocation rules $\hat{\boldsymbol{y}}$ and $\boldsymbol{y}$ satisfying $\boldsymbol{y} \preceq \hat{\boldsymbol{y}}$, the *interim composite mechanism* is:

(i) For each agent i, map type to quantile from $q_i = Q_i(t_i)$ according to $\mathcal{M}_i$.
(ii) For each agent i, calculate the quantile reserve $\hat{q}_i$ and resampling transformation σ_i by which y_i can be constructed from $\hat{y}_i$. Set $q_i^\dagger = \sigma_i(q_i)$.
(iii) Run $\hat{\boldsymbol{\mathcal{M}}}$ on $\boldsymbol{q}^\dagger$ to get $\boldsymbol{x}^\dagger = \hat{\boldsymbol{y}}^{EP}(\boldsymbol{q}^\dagger)$. Set $\boldsymbol{x}$ to incorporate the reserves $\hat{\boldsymbol{q}}$ as

$$x_i = \begin{cases} x_i^\dagger & \text{if } q_i^\dagger \leq \hat{q}_i, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

(iv) For each agent i, select the outcome distribution of $\mathcal{M}_i$ conditioned on x_i, i.e., $w_i^{x_i}(t_i)$.

Theorem 8.18. *For any type distribution $\boldsymbol{F}$ and mechanisms $\hat{\boldsymbol{\mathcal{M}}}$ and $\boldsymbol{\mathcal{M}}$ with allocation rules $\hat{\boldsymbol{y}}$ and $\boldsymbol{y}$ satisfying $\boldsymbol{y} \preceq \hat{\boldsymbol{y}}$, the composite mechanism (Definition 8.13) induces a distribution over allocations that is in the downward closure of the distribution of allocations of $\hat{\boldsymbol{\mathcal{M}}}$ and the same interim mechanisms as $\boldsymbol{\mathcal{M}}$.*[6]

Proof. See Exercise 8.9. □

[6] One distribution of allocations is *in the downward closure of* a second distribution of allocations if there is a coupling of the distributions so that the set of agents served by the first is a subset of those served by the second.

Corollary 8.19. *For any downward-closed service constrained environment, type distribution $\boldsymbol{F}$, ex post feasible mechanism $\hat{\boldsymbol{\mathcal{M}}}$, and Bayesian incentive compatible $\boldsymbol{\mathcal{M}}$ with allocation rules $\hat{\boldsymbol{y}}$ and $\boldsymbol{y}$ satisfying $\boldsymbol{y} \preceq \hat{\boldsymbol{y}}$, the composite mechanism (Definition 8.13) is ex post feasible, Bayesian incentive compatible, and has the same expected revenue as $\boldsymbol{\mathcal{M}}$.*

8.5 Multi-dimensional Externalities

This section considers optimal mechanisms for agents with multi-dimensional preferences where the way an agent is served imposes a multi-dimensional externality on the other agents via the feasibility constraint of the environment. For example, 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 be so assigned.

Definition 8.14. In a *multi-service service-constrained environment* there are n agents N and m services M. The subset of agent-service pairs that can be simultaneously assigned is given by $\mathcal{X} \subset \{0,1\}^{N \times M}$.

The multi- to single-agent reduction that was described in the previous section separates the problem of producing an outcome that is ex post feasible from the problem of ensuring that the mechanism is incentive compatible for each agent. This section takes the same approach; it applies generally to environments where each agent's utility linearly separates across distinct services in which she is interested. For simplicity we will state all results for the special case of additive agents.

Definition 8.15. An *additive* agent desires subsets of m services. Her type $t = (\{t\}_1, \ldots, \{t\}_m)$ is m-dimensional where $\{t\}_j$ is her value for alternative j. Her utility is linear; an outcome w is given by a payment and a marginal probability for each of the m services. For outcome $w = (\{x\}_1, \ldots, \{x\}_m, p)$, where $\{x\}_j$ denotes the marginal probability with which she obtains alternative j and p is her required payment, her utility is $u(t, w) = \sum_j \{t\}_j \{x\}_j - p$.

Note that it is possible within an additive multi-service environment to model more complex preferences. As a first example, unit-demand preferences (Definition 8.2) can be incorporated into the model by modifying the feasibility constraints $\mathcal{X}$ of Definition 8.14 so that it is infeasible to serve an agent more than one unit. It is also possible to model any general utility function over bundles of services as a unit-demand utility function over the power set of services, i.e., $\{0, 1\}^M$. We will see shortly, however, that the complexity of the construction depends on the number m of services, and thus moving to the power-set representation comes at an exponential blowup in complexity.

The assumption that the agents are additive implies, as is stated in the definition, the expected utility of an agent is determined by the marginal probability by which she is allocated each service. This property does not hold for general multi-dimensional utility functions. For example, a problematic case is when the agent views the services as complementary, e.g., she has high value to receive two services together but low value to receive either of the services individually. A mechanism that serves her both services or neither service with equal probability has the same marginal probabilities of allocating each service as the mechanism that serves her one or the other with equal probability. The agent has a higher utility for the former outcome distribution than the latter; thus, marginal probabilities are insufficient for determining such an agent's utility. In fact, any non-linearity of utility renders marginal probabilities similarly insufficient.

It is possible to decompose this mechanism design problem into a collection of single-agent problems that can be combined into a multi-agent mechanism, as we did in the previous section. In such a decomposition the single-agent problem is specified by m allocation constraints, one for each service. The difference between mechanisms for additive multi-service service constrained environments and the (single-service) service constrained environments is that the incentive compatibility constraints of the agents bind across multiple services not a single service. We will not formalize this approach here, instead we show that for additive multi-service service constrained environments, the optimal mechanism is a stochastic weighted optimizer (cf. Definition 8.10).

Our objective is to optimize expected revenue subject to Bayesian incentive compatibility and ex post feasibility. As we did in previous sections, we will replace the ex post feasibility constraint with an equivalent interim feasibility constraint. Ex post feasibility of a multi-service service constrained environment is equivalent to ex post feasibility of the following representative environment which is service constrained as per Definition 8.4. The intuition behind this representative environment is that we replace each multi-service agent, i.e., who desires subsets of the m services, with m single-service agents.

Definition 8.16. The *representative environment* for an n-agent m-service multi-service service-constrained environment is given by $n^\dagger = nm$ agents $N^\dagger = N \times M$ and service-constrained feasibility constraint $\mathcal{X}^\dagger = \mathcal{X} \subset 2^{N^\dagger}$. The type profile $\boldsymbol{t}$ for the original environment is extended to representative environment by duplicating each agent i's type across her m representatives, i.e., $t_{ij} = t_i$ for all i and j.

Importantly, ex post feasibility of the representative environment and the original multi-service environment is the same. Consequently, our discussion of interim feasibility extends directly. Recall that our discussion of interim feasibility relaxed the requirement that the agent types be independently distributed. This relaxation is important as in the representative environment for the multi-service environment all the representatives ij for $j \in M$ have the same type t_i, i.e., they are perfectly correlated. Thus, the characterization of interim feasibility and ex post implementations (Theorem 8.14) of the previous section hold for the representative environment.

Corollary 8.20. *For any joint distribution on type profiles and multi-service service-constrained environment, any interim feasible allocation profile can be ex post implemented by a stochastic weighted optimizer with weights that correspond to each type-service pair of each agent, i.e.,* $\boldsymbol{w}$ *with* $w_i : \mathcal{T}_i \times M \to \mathbb{R}$ *for each* i.

Recall, that a weighted optimizer for the representative environment assigns a weight to each type t_{ij} of each representative ij (see Definition 8.10); in the original environment such an assignment of weights corresponds to a weight for each type t_i, agent i, and service j (though weights for services that are infeasible for agent i can be omitted).

Theorem 8.21. *For any additive multi-service service-constrained environment, there is a stochastic weighted optimizer, with weights that map each feasible type-service pair, that is Bayesian incentive compatible and revenue-optimal among all Bayesian incentive compatible mechanisms.*

Proof. Consider any optimal mechanism $\mathcal{M}^\star$. The optimal mechanism must produce interim feasible allocation rules (specifying the marginal probability by which an agent of a given type receives each service). By Corollary 8.20 any profile of interim feasible allocation rules can be implemented as a stochastic weighted optimizer. Consider the mechanism given by this stochastic weighted optimizer and the same payment rule of $\mathcal{M}^\star$.

By the definition of additive utility agents, each agent's utility for a randomized outcome depends only on the marginal probability that she receives each service (and expected payment). These marginal probabilities and expected payments are the same for both mechanisms; thus, incentive compatibility of $\mathcal{M}^\star$ implies incentive compatibility of the stochastic weighted optimizer. Since both mechanisms have the same payment rule, they both have the same revenue; the optimality of $\mathcal{M}^\star$ implies the optimality of the stochastic weighted optimizer. □

8.6 Public Budget Preferences

In the sections below we will prove the optimality of the single-agent mechanisms described in Section 8.1 for agents with public budgets. In particular, we will show that for a large family of well-behaved distributions the revenue-optimal single-agent mechanism will have an all-pay payment rule and will reserve-price the low valued agents and iron the top valued agents.

Recall the public budget preference where the agent has a single dimensional value t drawn from distribution F and public budget B. The agent's utility for allocation x and payment p is $tx - p$ when $p \leq B$ and negative infinity of $p > B$. We will assume that the distribution F is continuous and supported on types $\mathcal{T} = [0, h]$.

We begin by observing that, for an agent with a public budget, the optimal mechanism, satisfying the usual Bayesian incentive compatibility (BIC) and interim individual rationality (IIR) constraints, is an all-pay mechanism. In other words, the agent makes a bid and pays this bid always, though she may only win some of time. All-pay mechanisms may seem unnatural as they are not ex post individually rational, i.e., an agent will sometimes have negative utility. Notice, however, that in most economic interactions there are upsides and downsides that strategic agents must trade off; ex post individual rationality is the exception rather than the rule. Moreover, as non-all-pay mechanisms will generally be suboptimal, ex post individual rationality comes at a loss in performance, in this case revenue, relative to the optimal all-pay format.

Proposition 8.22. *The revenue-optimal Bayesian-incentive-compatible and interim-individually-rational mechanism for single-dimensional agents with public budgets is an all-pay mechanism.*

Proof. An agent with public budget is quasi-linear except for her budget constraint. Therefore, unless the budget constraint is violated, revenue equivalence of Section 2.7 on page 37 implies two mechanisms with the same allocation rule in equilibrium have the payment rule (in the interim stage of the mechanism).

Consider any mechanism where, in equilibrium, the agent's budget constraint is not violated. Recall that the payment rule $p(t)$ is defined as the expected payment of the agent. For a given valuation t the payment of an agent is a random variable, potentially a function of randomization in the mechanism and randomization in the types of other agents. By the definition of expectation, the maximum payment in the support of the distribution of payments is at least the expected payment. As the budget is not violated for this maximum payment, it is not violated for the expected payment, i.e., $p(t) \leq B$. Thus, the all-pay mechanism that requires deterministic payment $p(t)$ does not violate the budget either. Therefore, it is incentive compatible and obtains the same revenue. □

We now proceed to characterize the optimal single-agent all-pay mechanism subject to an interim feasibility constraint $\hat{y}(\cdot)$. This optimization problem is similar to that for the single-dimensional linear agent that was previously solved in Section 3.3; however, the solution to the optimization must additionally satisfy the (all-pay) budget constraint that $p(t) \leq B$ for all $t \in \mathcal{T}$. As payments are non-decreasing in the agent's type, the budget constraint for all types $t \in \mathcal{T} = [0, h]$ is implied by the budget constraint for the highest type h. In other words, the revenue-optimization problem has only the additional constraint $p(h) \leq B$.

Our approach will be to write a mathematical program for the revenue maximization problem where the budget appears as a constraint. We will then use Lagrangian relaxation to move the budget constraint into the objective.[7] We will proceed by optimizing this Lagrangian objective in

[7] For maximization problems, Lagrangian relaxation of a constraint (a) rewrites it as a quantity that is at least zero and (b) and moves the terms of the constraint, scaled by a Lagrangian parameter λ, to the objective scales. Thus, satisfying the constraint is consistent with the objective, i.e., this term is larger when the constraint is satisfied. The Lagrangian parameter λ allows the emphasis of the Lagrangian objective to traded off between the original objective and satisfaction of the constraint. At $\lambda = 0$ the constraint is ignored and is only satisfied if it was not binding in the first place. At $\lambda = \infty$, the objective is ignored and the constraint is satisfied with slack (if it is satisfiable by any assignment of the variables of the program). The original program with the constraint is optimized by finding the λ where the constraint is met with equality; at such a point the Lagrangian program trades off emphasis on the objective and the constraint perfectly. Notice that when the constraint is met with equality, the contribution to the objective is zero and the objective of the Lagrangian program is the optimal value of the original program.

the same manner as our revenue optimization for single-dimensional linear agents, cf. Section 3.3.4. We will rewrite the objective in terms of the allocation rule and Lagrangian revenue curves. For a given Lagrangian parameter λ, and these revenue curves, we will be able to identify the optimal allocation rule for any allocation constraint, cf. Section 3.4.5. Finally, we choose the Lagrangian parameter so that the budget is met with equality.

We begin by writing a mathematical program for the interim revenue maximization problem and using Lagrangian relaxation to move the budget constraint into the objective. The other constraints of the problem will not play a roll in most of our discussion so we will not write them formally. The original and relaxed programs are as follows; recall the value of the highest type is denoted h.

$$\begin{aligned} \sup_{(x,p)} \ & \mathbf{E}_{t\sim F}[p(t)] \\ \text{s.t. } & (x,p) \text{ are BIC, IIR,} \\ & \text{and feasible for } \hat{y}; \\ & p(h) \leq B. \end{aligned} \qquad \begin{aligned} \sup_{(x,p)} \ & \mathbf{E}_{t\sim F}[p(t)] + \lambda\, B - \lambda\, p(h) \\ \text{s.t. } & (x,p) \text{ are BIC, IIR,} \\ & \text{and feasible for } \hat{y}. \end{aligned}$$

We will fix the Lagrangian parameter λ and characterize the optimizer of the Lagrangian objective. Notice that this Lagrangian objective is linear and therefore can be treated with the methods of Chapter 3. With such a characterization, the Lagrangian parameter λ can be chosen to be zero if the budget is not binding or so that the budget constraint is met with equality if it is binding.

With Lagrangian λ fixed, the $\lambda\, B$ term in the objective is a constant and does not affect optimization. The optimization is to find allocation and payment rules (x,p) to maximize $\mathbf{E}[p(t)] - \lambda\, p(h)$. Our approach to this optimization problem will mirror our approach to revenue optimization without budgets, cf. Section 3.3.4. We will define Lagrangian revenue curves, we will write the Lagrangian objective of any allocation rule in terms of these revenue curves, and then we will directly interpret the form of the Lagrangian optimizer.

Recall that price-posting revenue curves are defined by considering the ex ante constraint $\hat{q}$ and the mechanism $(x^{\hat{q}}, p^{\hat{q}})$ that posts the price $V(\hat{q})$ that is accepted with probability $\hat{q}$. Consider the Lagrangian objective $\mathbf{E}_t[p^{\hat{q}}(t)] - \lambda\, p^{\hat{q}}(h)$ for the mechanism that posts price $V(\hat{q})$. The revenue from such a price is $\mathbf{E}_t[p^{\hat{q}}(t)] = P(\hat{q}) = \hat{q}\, V(\hat{q})$ where, recall, $P(\hat{q})$ denotes the price-posting revenue curve of the single-dimensional linear utility agent.

For $\hat{q} > 0$ (strictly positive), the price $V(\hat{q})$ is strictly less than the value of the highest type $V(0) = h$, so $p^{\hat{q}}(h) = V(\hat{q})$. Thus, the Lagrangian objective for $\hat{q} \in (0, 1]$ is $P_\lambda(\hat{q}) = P(\hat{q}) - \lambda V(\hat{q})$. For $\hat{q} = 0$ the highest type is indifferent between buying and not buying. This indifference does not matter as highest type (quantile $q = 0$) is realized with measure zero (i.e., never) and so this type cannot affect the optimization. It will be technically convenient and consistent with the highest type "not mattering" to assume, with respect to the budget constraint, that indifference of the highest type to the price h is resolved in favor of rejecting the price. For posting price $V(\hat{q} = 0) = h$, the expected revenue is $\mathbf{E}[p^0(t)] = 0$ and, by this indifference-resolution assumption, the payment of the highest type is $p^0(h) = 0$; therefore, the expected Lagrangian objective from posting price $V(0)$ is identically zero. Having identified the expected revenue and payment of the highest type for every $\hat{q}$ price posting we have identified the Lagrangian price-posting revenue curve; see Figure 8.13(a). Notice that this revenue curve is discontinuous at $\hat{q} = 0$ (unless $\lambda = 0$, i.e., when the budget constraint is not binding).

Proposition 8.23. *The Lagrangian price-posting revenue curve for an agent with public budget satisfies*

$$P_\lambda(\hat{q}) = \begin{cases} 0 & \text{if } \hat{q} = 0, \text{ and} \\ P(\hat{q}) - \lambda V(\hat{q}) & \text{otherwise.} \end{cases}$$

Notice that on $q \in (0, 1]$ this Lagrangian price-posting revenue curve is the difference between the original revenue curve and the scaled value function. If the original revenue curve is concave and the value function $V(q) = F^{-1}(1 - q)$ is convex (equivalently, the cumulative distribution function $F(\cdot)$ is convex; equivalently, the density function $f(\cdot)$ is monotone non-decreasing), then this Lagrangian price-posting revenue curve is concave (on $q \in (0, 1]$).

Definition 8.17. A single-dimensional public budget agent is *regular* if for all Lagrangian parameters $\lambda \geq 0$ the Lagrangian price-posting revenue curve is concave on interval $(0, 1]$. The value distribution F of such a regular public-budget agent is *public-budget regular.*

Proposition 8.24. *A single-dimensional public budget agent is regular*

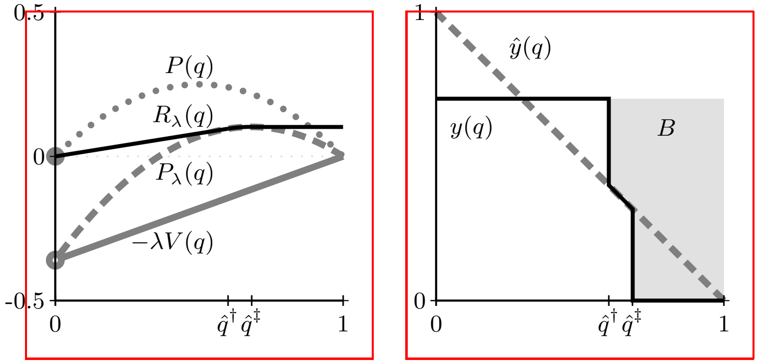

(a) Lagrangian Revenue Curves (b) Allocation Rules

Figure 8.13. Depicted in (a) are the Lagrangian revenue curves corresponding to an agent with type distributed uniformly on $[0, 1]$ and Lagrangian parameter $\lambda > 0$. Note that the Lagrangian price-posting revenue curve $P_\lambda(\cdot)$ is discontinuous at $q = 0$ with $P_\lambda(0) = 0$ and $\lim_{q \to 0} P_\lambda(q) = -\lambda\, V(0)$. The Lagrangian price-posting revenue curve (thick, gray, dashed line) on $q \in (0, 1]$ is the sum of the single-dimensional linear price-posting revenue curve $P(q)$ (thick, gray, dotted line) and the relaxed budget constraint $-\lambda\, V(q)$ (thick, gray, solid line). Depicted in (b) are the allocation rules corresponding to optimization of the Lagrangian objective for two i.i.d. agents with budget $B = 1/4$. The allocation rule y (thin, black, solid line) is derived from the allocation constraint $\hat{y}$ (thick, gray, dashed line) by averaging on $[0, \hat{q}^\dagger)$ and zeroing on $(\hat{q}^\star, 1]$. The Lagrangian parameter was selected to meet the budget constraint with equality for this example.

if (a) her type distribution F is regular (for single-dimensional linear agents) and (b) its cumulative distribution function $F(\cdot)$ is convex.

Due to the discontinuity at $q = 0$, public-budget regularity does not imply that the entire Lagrangian price-posting revenue curve is concave. Recall that when a price-posting revenue curves is not concave, as in the case of an irregular distribution with a single-dimensional linear agent, optimization subject to incentive compatibility (i.e., monotonicity of the allocation rule) is simplified by ironing (cf. Section 3.3.5 on page 75). With respect to the Lagrangian price-posting revenue curve, ironing is equivalent to taking the concave hull, i.e., the smallest concave upper bound. Geometrically, it is easy to see that this ironing replaces the revenue curve with a line segment from the origin to the point where it is tangent to the original curve. This point is uniquely identified by $\hat{q}^\dagger$ satisfying $P_\lambda(\hat{q}^\dagger) = \hat{q}^\dagger P'_\lambda(\hat{q}^\dagger)$. The resulting revenue curve is continuous and concave (Proposition 8.25). The *Lagrangian revenue curve* denotes, for every ex ante constraint $\hat{q}$, the optimal Lagrangian objective value from a mechanism with ex ante sale probability $\hat{q}$.

Proposition 8.25. *The Lagrangian revenue curve $R_\lambda(\cdot)$ for an agent with public budget and value drawn from a public-budget regular distribution satisfies*

$$R_\lambda(q) = \begin{cases} q\,P'_\lambda(\hat{q}^\dagger) & \text{if } q \in [0, \hat{q}^\dagger], \\ P_\lambda(q) & \text{if } q \in [\hat{q}^\dagger, \hat{q}^\ddagger], \text{ and} \\ P_\lambda(\hat{q}^\ddagger) & \text{if } q \in [\hat{q}^\ddagger, 1], \end{cases}$$

with $\hat{q}^\dagger$ and $\hat{q}^\ddagger$ set to satisfy $P_\lambda(\hat{q}^\dagger) = \hat{q}^\dagger P'_\lambda(\hat{q}^\dagger)$ and $P'_\lambda(\hat{q}^\ddagger) = 0$, respectively.

The revenue curves $P_\lambda(\cdot)$ and $R_\lambda(\cdot)$ that correspond to price posting and ex ante optimization can be extended to describe the Lagrangian objective, respectively, for any allocation rule $y(\cdot)$ and for optimization with respect to constraint $\hat{y}(\cdot)$. These extensions follow from reinterpreting an allocation rule or constraint as a distribution over ex ante constraints. For example, the a mechanism with allocation rule $y(\cdot)$ can be obtained by drawing a random quantile $\hat{q}$ from distribution G^y with cumulative distribution function $G^y(z) = 1 - y(z)$ and offering the agent the price that corresponds to this quantile i.e., $V(\hat{q})$ (and then applying the revenue equivalence to convert this mechanism to its all-pay equivalent). Thus, the expected Lagrangian objective for allocation rule $y(\cdot)$ is $\mathbf{E}_{\hat{q}\sim G^y}[P_\lambda(\hat{q})] = P_\lambda(1)\,y(1) + \mathbf{E}_{q\sim U[0,1]}\big[P_\lambda(q)\,[-y'(q)]\big]$. Similarly, the optimal Lagrangian objective for allocation constraint $\hat{y}$ is $\mathbf{Rev}[\hat{y}] = \mathbf{E}\big[R_\lambda(q)\,[-\hat{y}'(q)]\big]$.

The usual integration by parts approach, with the fact that the Lagrangian revenue curve satisfies $R_\lambda(0) = 0$, implies that the optimal Lagrangian objective can be rewritten in terms of the Lagrangian marginal revenue curve $R'_\lambda(\cdot)$ as $\mathbf{Rev}[\hat{y}] = \mathbf{E}\big[R'_\lambda(q)\,\hat{y}(q)\big]$. Monotonicity of this marginal revenue curve, the theory of Lagrangian relaxation, and Corollary 3.6 on page 65 gives the following theorem.

Theorem 8.26. *For a public-budget agent, the revenue-optimal mechanism is given by optimizing the surplus of Lagrangian marginal revenue with Lagrangian parameter $\lambda > 0$ when the budget constraint met with equality for the highest type, or with $\lambda = 0$ when the budget constraint is not binding.*

As in the linear utility case, the optimal mechanism for the Lagrangian objective can be interpreted from the Lagrangian revenue curves. In particular, we get the optimal allocation rule y subject to constraint $\hat{y}$ from ironing on the intervals where the Lagrangian price-posting revenue curve is ironed, and reserve pricing at its peak (cf. Section 3.4.5 on page 85). Thus, it is optimal to iron the high-valued types, i.e., quantiles $q \in [0, \hat{q}^\dagger)$ with $\hat{q}^\dagger$ as described above, and reserve price the low-valued types, i.e., quantiles $q \in (\hat{q}^\star, 1]$ with $\hat{q}^\star$ defined as the maximizer of $R_\lambda(\cdot)$. The remaining types, which correspond to quantiles $q \in [\hat{q}^\star, \hat{q}^\dagger]$, are served greedily according to the allocation constraint $\hat{y}(\cdot)$. The resulting allocation rule $y(\cdot)$ can be interpreted as averaging $\hat{y}(\cdot)$ on $[0, \hat{q}^\dagger)$ and setting it to zero on $(\hat{q}^\star, 1]$; see Figure 8.13(b).

Corollary 8.27. *For a regular public-budget agent and interim allocation constraint $\hat{y}(\cdot)$, the optimal single-agent mechanisms allocates as by $\hat{y}(\cdot)$ except that types with quantiles in $[0, \hat{q}^\dagger)$ are ironed, and types with quantiles in $(\hat{q}^\star, 1]$ are reserve priced.*

8.7 Unit-demand Preferences

In the sections below we will prove the optimality of the single-agent mechanisms described in Section 8.1 for agents with unit-demand preferences. More generally, we will show that for a large family of well-behaved and alternative-symmetric distributions, the optimal mechanism is given by a uniform reserve price (i.e., the same across all alternative) and sells the agent her favorite alternative.

This section gives a generalization to multi-dimensional preferences of the framework of virtual values (cf. Section 3.3.1). Recall that virtual values are an amortization of revenue in the sense that they can be evaluated pointwise, but equate to revenue in expectation. The pointwise optimization of virtual surplus, then, gives a revenue-optimal incentive compatible mechanism. The approach is to (a) cover type space with paths, (b) solve the problem restricted to the path, and then (c) find sufficient conditions on the distribution of over types that implies that the optimal mechanisms on the path are consistent. This approach will generally fail unless the right paths are identified.

We will use this approach to solve the single-agent problems corresponding to an unit-demand agent with uniformly distributed types on the unit square (Example 8.2). For such an agent the optimal mechanism projects the multi-dimensional agent type onto a single dimension that corresponds to the agent's value for her favorite alternative. In addition to proving this result, we will give sufficient conditions on the distribution, beyond uniform, under which this projection continues to be optimal.

This single-dimensional projection result gives insight on the role of second-degree price discrimination, i.e., whether a seller can make more money with a differentiated product. For example, a seller might introduce a high-quality and low-quality product to segment the market between high-valued consumers (to buy the high-quality product at a premium) and low-valued consumers (to buy the low-quality product at a discount). Intuitively, this approach can be profitable if high-valued consumers are more sensitive to quality than low-valued consumers. This section develops a proof of the inverse, that if high-valued consumers are less sensitive to quality than low-valued consumers, then there is no benefit to quality-based second-degree price discrimination. For example, movie tickets are predominantly sold with a uniform price. Such a mechanism is suggested by the results of this section under the assumption that film buffs tend not to have a higher willingness to pay than the general public.

We begin the section with a simple warmup exercise that single-agent problems for the two-alternative uniform unit-demand agent of Example 8.2. The approach is to solve the mechanism design problem independently on rays from the origin and relies solely on the single-dimensional theory of mechanism design from Chapter 3. To solve more complex multi-dimensional problems we generalize the characterization of incentive compatible mechanisms to multi-dimensional agents. We then solve the multi-dimensional mechanism design problem for more general families of paths. For the right choice of paths the approach of the warmup can be generalized. To identify the right paths, we develop a multi-dimensional framework of virtual values.

8.7.1 Warmup: The Uniform Distribution

As a warmup, consider selling one of two alternatives to a unit-demand agent with type drawn from the uniform distribution over type space $\mathcal{T} = [0,1]^2$ (Example 8.2). We claimed without proof in Section 8.1.2 that the optimal (unconstrained) single-agent mechanism is to post a uniform price of $\sqrt{1/3}$. A simple argument for this result is as follows.

First, restrict the problem to the alternative-1 preferred subspace of types, i.e., where $\{t\}_1 > \{t\}_2$ (the solution for the other part will be symmetric). A uniform pricing always sells the agent her favorite alternative, so with this restriction, the uniform pricing sells alternative 1 only. The conditional distribution on $\{t\}_1$ is the distribution of the maximum of two i.i.d. uniform random variables and has cumulative distribution function $F_{\max}(z) = \mathbf{Pr}[\{t\}_1 \leq z \wedge \{t\}_2 \leq z] = z^2$, density function $f_{\max}(z) = 2z$, single-dimensional virtual value $\phi_{\max}(z) = z - {}^{1-z^2}/_{2z}$, and monopoly price $\hat{v}^\star_{\max} = \sqrt{1/3}$.

Now consider restricting the type space to paths that coincide with rays from the origin. Parameterizing such a path by its slope θ, a type on this path can be expressed in terms of $\{t\}_1$ as $t = (\{t\}_1, \theta\,\{t\}_1)$. Notice that all types $t \in \mathcal{T}_\theta = \{(v, \theta\, v) : v \in [0, 1]\}$ have the same value for receiving alternative 1 with probability θ or alternative 2 with certainty. Thus, restricting the type space to the path $\mathcal{T}_\theta$, the problem of selling the agent alternative 1 or 2 is equivalent to that of selling the agent alternative 1 with probability one or alternative 1 with probability θ. Recall from Section 3.3 that the optimal single-dimensional mechanism, which is allowed to probabilistically allocate, is always deterministic. It sells to the agent with probability one if she has a non-negative virtual value and with probability zero otherwise. In other words, it posts the monopoly price. Thus, the optimal mechanism for $\mathcal{T}_\theta$ posts a price for alternative 1.

This restriction on type space to a path can be equivalently viewed as giving the mechanism designer extra knowledge, specifically, the knowledge of θ. With this extra knowledge, the conditional distribution on $\{t\}_1$ is $F_{\max}$, thus, the designer with this knowledge would post a price of $\hat{v}^\star_{\max} = \sqrt{1/3}$ for alternative 1 (and not sell alternative 2). This solution is independent of θ, and the designer can do as well without knowledge of θ as with it. Thus, there is no loss with respect to the optimal mechanism from relaxing the incentive constraints between types that are not on the same path. The optimal mechanism for a unit-demand agent with types uniformly drawn from the full type space $\mathcal{T} = [0, 1]^2$ is the uniform price of $\sqrt{1/3}$.

In the remainder of the section this approach is generalized to a richer family of distributions. In particular, the same uniform-pricing result holds for any distribution where the conditional distributions of $\theta = {}^{\{t\}_2}/_{\{t\}_1}$ with respect to $\{t\}_1$ is ordered according to $\{t\}_1$ by first-order stochastic dominance. In other words, $\mathbf{Pr}[{}^{\{t\}_2}/_{\{t\}_1} \leq \theta \mid \{t\}_1]$, for all fixed θ, is monotone in $\{t\}_1$.

8.7.2 Multi-dimensional Characterization of Incentive Compatibility

Chapter 2 characterized incentive compatible mechanisms for single-dimensional linear agents in Theorem 2.2 and Corollary 2.12. These results concluded that a mechanisms with allocation and payment rules (x, p) is incentive compatible if and only if

- the allocation rule $x(\cdot)$ is monotone non-decreasing, and
- the payment rule satisfies the payment identity:

$$p(v) = v\, x(v) - \int_0^v x(z)\, \mathrm{d}z.$$

Recall that the first term in the payment identity is the surplus and the second term is, thus, the agent's utility. We can reinterpret this characterization in terms of utility as follows. The utility function $u(\cdot)$ corresponds to an incentive compatible mechanism with allocation rule x if and only if

- it is convex, and
- related to the allocation rule by the *utility derivative identity*:

$$x(v) = \tfrac{\mathrm{d}}{\mathrm{d}v} u(v).$$

Moreover, under our usual interpretation of the allocation rule $x(v)$ as denoting the probability that the agent with value v is served, the utility derivative identity combined with $x(v) \in [0, 1]$ imply that the utility function is non-decreasing and has derivative at most one.

The multi-dimensional characterization of incentive compatibility generalizes this reinterpretation.

Theorem 8.28. *For an agent with linear utility, allocation rule and utility functions (x, u) correspond to an incentive compatible mechanism if and only if*

(i) *(convexity) $u(\cdot)$ is convex, and*
(ii) *(utility gradient identity) $x(t) = \nabla u(t)$.*[8]

[8] Technically, the gradient ∇u of a convex function is only guaranteed to exist almost everywhere (and not everywhere). For types t where the gradient does not exist, the allocation $x(t)$ can be any *subgradient*, i.e., the gradient of any plane through point $(t, u(t))$ that lower bounds the utility function $u(\cdot)$; convexity of the utility function guarantees that such a plane exists.

Proof. Incentive compatibility is equivalent to the following inequality holding for all pairs of types $(t, t^\dagger)$:

$$u(t) \geq u(t^\dagger) + (t - t^\dagger) \cdot x(t^\dagger). \tag{8.11}$$

The right-hand side of equation (8.11) is the utility that t obtains for the outcome of $t^\dagger$. The only difference between the utility of t for an outcome and the utility of $t^\dagger$ for an outcome is the surplus from the allocation. Thus, t's utility for the outcome of $t^\dagger$ is equal to the utility $t^\dagger$ for this outcome plus the difference in surplus for t and $t^\dagger$ for the outcome.

Like the proof of Theorem 2.2, this proof is broken into three parts.

(i) The allocation rule and utility function (x, u) correspond to an incentive compatible mechanisms if convexity and the utility gradient identity hold.

Consider any type $t^\dagger$ and the plane orthogonal to the surface of $u(\cdot)$ at $t^\dagger$. By convexity, this plane is lower-bounds the utility at any other type t. In other words,

$$u(t) \geq u(t^\dagger) + (t - t^\dagger) \cdot \nabla u(t^\dagger).$$

In this equation, the right-hand side is the point on this plane at t. By the gradient utility identity, we can substitute the $x(t^\dagger)$ for $\nabla u(t^\dagger)$ in the right-hand side to obtain the defining inequality (8.11) of incentive compatibility.

(ii) The allocation rule and utility function (x, u) correspond to an incentive compatible mechanisms only if convexity holds.

Consider types $t^\dagger$, $t^\ddagger$, and convex combination $t = \gamma\, t^\dagger + (1 - \gamma)\, t^\ddagger$. Incentive compatibility requires equation (8.11) hold for all pairs of types. Restating the equation for type pairs $(t^\dagger, t)$, $(t^\ddagger, t)$ we have:

$$u(t^\dagger) \geq u(t) + (t^\dagger - t) \cdot x(t).$$
$$u(t^\ddagger) \geq u(t) + (t^\ddagger - t) \cdot x(t).$$

A convex combination of these equations gives:

$$\begin{aligned} \gamma\, u(t) + (1 + \gamma)\, u(t^\ddagger) &\geq u(t) + (\gamma\, t^\dagger + (1 - \gamma)\, t^\ddagger - t) \cdot x(t) \\ &= u(t) \end{aligned} \tag{8.12}$$

The final equation above comes from the definition of t as the convex combination of $t^\dagger$ and $t^\ddagger$. Inequality (8.12) implies convexity of utility as desired.

(iii) The allocation rule and utility function (x, u) correspond to an incentive compatible mechanisms only if the utility gradient equality holds.

Let e_j be the unit vector corresponding to allocation of alternative j, i.e., $\{e_j\}_j = 1$ and $\{e_j\}_{j^\dagger} = 0$ for $j \neq j^\dagger$. For small constant ϵ, apply equation (8.11) to type pairs $(t + \epsilon e_j, t)$ and $(t - \epsilon e_j, t)$ to conclude:

$$u(t + \epsilon\, e_j) - u(t) \geq \epsilon\, \{x(t)\}_j, \text{ and}$$
$$u(t - \epsilon\, e_j) - u(t) \geq -\epsilon\, \{x(t)\}_j.$$

Combine these equations to obtain upper and lower bounds on $\{x(t)\}_j$ as:

$$1/\epsilon\, [u(t + \epsilon\, e_j) - u(t)] \geq \{x(t)\}_j \geq 1/\epsilon\, [u(t - \epsilon\, e_j) - u(t)]$$

Assuming the partial derivative of $u(\cdot)$ with respect to $\{t\}_j$ is defined at t, the limit as ϵ goes to zero is defined and both the upper and lower bound, above, are equal to the partial derivative of u with respect to $\{t\}_j$ at t which is the jth coordinate of the gradient $\{\nabla u(t)\}_j$. If the partial derivative is not defined, then the same limit argument implies that $x(t)$ is a subgradient of the utility function at type t. □

The subsequent developments of this section will rely heavily on Theorem 8.28.

8.7.3 Optimal Mechanisms for Paths

A seller who segments the market by offering a differentiated product line is engaging in what is called *second-degree price discrimination.* One way to offer a differentiated product is to offer lotteries for the same product. For example, a seller could offer (1) the good at a high price or (2) the same good with probability $1/2$ (and nothing otherwise) at a low price. Our analysis of single-dimensional agents of Chapter 3 has shown pricing these lotteries is never beneficial.

In the example above, a buyer who has a value v for the allocation of (1) and will have value $v/2$ for the allocation of (2). Viewing these allocations as two alternatives, the buyer's type can be mapped into the two dimensional space corresponding to her value for each alternative. The buyer's type space is degenerate and lies on the line with slope $1/2$. Of course, the optimal mechanism when the buyer's value is drawn from a distribution is to post the monopoly price for the distribution for alternative 1 (and never sells alternative 2). In this section, this monopoly-pricing result is generalized to type spaces given by more general families of paths.

Definition 8.18. A *path-based agent* is specified by a path $C : [0,1] \to [0,1]$ with $C(v) \leq v$ (thus, $C(0) = 0$) and distribution $F_{\max}$ where the agent's type is given by $t^v = (v, C(v))$ with v drawn from $F_{\max}$. The type space is $\mathcal{T} = \{(v, C(v)) : v \in [0,1]\}$.

Now we follow the same approach as Section 3.3 and convert the problem of optimizing revenue in expectation to the problem of optimizing a virtual surplus pointwise. The approach is the following. Expected profit is equal to expected surplus minus expected utility. We will use integration by parts on the path to write the expected utility as the integral of the gradient of the utility. By Theorem 8.28, the gradient of utility is equal to the allocation. The two terms for expected surplus and utility can then be combined to give a virtual surplus.

Lemma 8.29. *For a path-based agent with path C and distribution $F_{\max}$, and any incentive compatible mechanism with allocation rule x, the agent's expected utility is*

$$\mathbf{E}[u(t^v)] = u(t^0) + \mathbf{E}\big[x(t^v) \cdot (1, C'(v))\,(1 - F_{\max}(v))\,{}^1\!/\!{}_{f_{\max}(v)}\big]\,.$$

Proof. An explanation of the following calculus is given below.

$$\begin{aligned}
\mathbf{E}[u(t^v)] &= \int_0^1 u(t^v)\, f_{\max}(v)\, \mathrm{d}v \\
&= -\int_0^1 u(t^v)\, \tfrac{\mathrm{d}}{\mathrm{d}v}[1 - F_{\max}(v)]\, \mathrm{d}v \\
&= -\Big[u(t^v)\,[1 - F_{\max}(v)]\Big]_0^1 + \int_0^1 \tfrac{\mathrm{d}}{\mathrm{d}v}[u(t^v)]\,[1 - F_{\max}(v)]\, \mathrm{d}v \\
&= u(t^0) + \int_0^1 \nabla u(t^v) \cdot (1, \tfrac{\mathrm{d}}{\mathrm{d}v} C(v))\,[1 - F_{\max}(v)]\, \mathrm{d}v \\
&= u(t^0) + \mathbf{E}\big[x(t^v) \cdot (1, C'(v))\,(1 - F_{\max}(v))\,{}^1\!/\!{}_{f_{\max}(v)}\big]\,.
\end{aligned}$$

The first line is from the definition of expectation. The second line is from the definition of the density function $f_{\max}$ as the derivative of the cumulative distribution function $F_{\max}$. The third line is by integration by parts. The first part of the third line simplifies by substituting $1 - F_{\max}(0) = 1$ and $1 - F_{\max}(1) = 0$; the second part of the third line simplifies by taking the derivative of the utility; and the fourth line results. The final line is from the definition of expectation. □

Theorem 8.30. *For a path-based agent with path C and distribution $F_{\max}$, and any incentive compatible mechanism (x, p), the expected revenue is*

$$\mathbf{E}[p(t^v)] = p(t^0) + \mathbf{E}\big[x(t^v) \cdot \big[t^v - (1, C'(v))\,[1 - F_{\max}(v)]\,{}^1\!/\!{}_{f_{\max}(v)}\big]\big]\,.$$

Proof. Expected revenue is equal to expected surplus minus expected utility, i.e.

$$\mathbf{E}[p(t^v)] = \mathbf{E}[x(t^v) \cdot t^v] - \mathbf{E}[u(t^v)]\,.$$

Lemma 8.29 allows the expected utility to be rewritten in terms of the allocation rule and the utility of type $t^0 = (0, 0)$. The agent with type t^0 has no surplus, so her only utility is from the negation of her payment, i.e., $p(t^0) = -u(t^0)$. Combining these two equations, we have the theorem. □

Theorem 8.30 shows that the vector field $\phi(t^v) = t^v - (1, C'(v))\,[1 - F_{\max}(v)]\,{}^1\!/\!{}_{f_{\max}(v)}$ gives a virtual value function for revenue. The remaining question for revenue maximization is to choose the allocation rule to optimize virtual surplus with respect to ϕ subject to incentive compatibility. As in Section 3.3, we relax the incentive compatibility constraint, and choose allocation x to optimize the virtual surplus $\phi(t^v) \cdot x$ pointwise for each $t^v \in \mathcal{T}$. We then check for conditions on the environment, in this case the path C and distribution $F_{\max}$, that imply that the resulting allocation rule is incentive compatible.

Our goal in this section is not to identify the optimal incentive compatible mechanism. Instead we are looking for sufficient conditions on the distribution to imply that the optimal mechanism posts a price for alternative 1 only. Notice that the first coordinate of the virtual value function is exactly the single-dimensional virtual value corresponding to distribution $F_{\max}$, i.e., $\{\phi(t^v)\}_1 = v - \frac{1-F_{\max}(v)}{f_{\max}(v)}$. If the distribution is regular (i.e., this function is monotone non-decreasing; Definition 3.4) then selling alternative 1 to maximize virtual surplus will post the monopoly price $\hat{v}^\star_{\max}$ that solves $\hat{v}^\star_{\max} - \frac{1-F_{\max}(\hat{v}^\star_{\max})}{f_{\max}(\hat{v}^\star_{\max})} = 0$. Pointwise optimization of $\phi(t^v) \cdot x$ serves the agent the alternative with the highest positive virtual value. Thus, for virtual surplus maximization to be equivalent to posting a price for alternative 1 only, it better be that when $\{\phi\}_1 > 0$ that $\{\phi\}_1 \geq \{\phi\}_2$ and when $\{\phi\}_1 \leq 0$ that $\{\phi\}_2 \leq 0$.

Definition 8.19. A path $C(\cdot)$ is *ratio monotone* if the ratio ${}^{C(v)}\!/\!{}_v$ is monotone non-decreasing in v; see Figure 8.14.

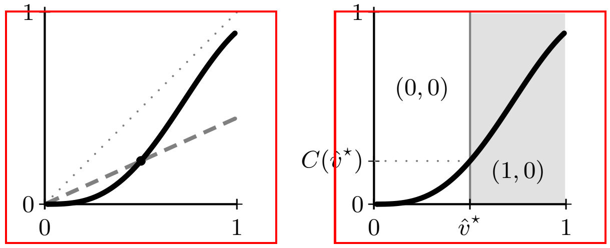

(a) Ratio montone path (b) Optimal mechanism

Figure 8.14. In subfigure (a), a ratio-monotone path (solid, thick, black line) is depicted. Ratio monotonicity of the path implies that the slope of the path at a type t is greater than that of the line through the type and the origin (depicted for type $t^{1/2}(1/2, C(1/2))$; gray, dashed line). Subfigure (b) depicts the optimal mechanism, when value v is drawn uniformly from $[0, 1]$. This mechanism post price $\hat{v}^\star = 1/2$ for alternative 1.

Theorem 8.31. *For a path-based agent with ratio-monotone path C and regular distribution $F_{\max}$, the optimal mechanism is to post a price for alternative 1.*

Proof. For the allocation that optimizes virtual surplus $\phi(t^v) \cdot x$ pointwise to never sell alternative 2, it better be that when $\{\phi\}_1 > 0$ that $\{\phi\}_1 \geq \{\phi\}_2$ and when $\{\phi\}_1 \leq 0$ that $\{\phi\}_2 \leq 0$. A sufficient condition is, for all v,

$$\tfrac{C(v)}{v} \{\phi(t^v)\}_1 \geq \{\phi(t^v)\}_2. \tag{8.13}$$

Ratio-monotonicity is equivalent to the property that rays from the origin only cross the path from above to below, i.e., at the point of intersection, the slope of the ray is at most the slope of the path, i.e.,

$$C'(v) \geq C(v)/v.$$

The sufficient condition of equation (8.13) can be derived from ratio monotonicity as follows,

$$\begin{aligned} \tfrac{C(v)}{v} \{\phi(t^v)\}_1 &= \tfrac{C(v)}{v} \left[v - \tfrac{1-F_{\max}(v)}{f_{\max}(v)}\right] \\ &\geq C(v) - C'(v)\, \tfrac{1-F_{\max}(v)}{f_{\max}(v)} \\ &= \{\phi(t^v)\}_2. \end{aligned}$$

Thus, pointwise virtual surplus maximization never sells alternative 2. For virtual surplus maximization to additionally correspond to posting a price alternative 1, regularity of the distribution $F_{\max}$, as assumed in the statement of the theorem, is sufficient. □

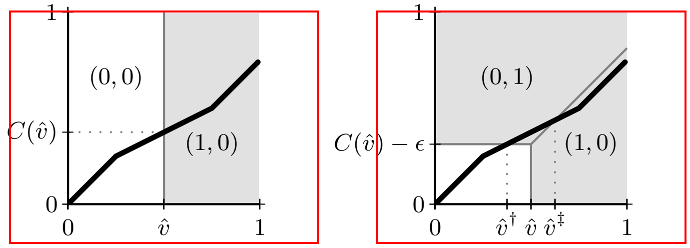

(a) Pricing alternative 1 (b) Discounting alternative 2

Figure 8.15. Depicted is a path (thick, solid, black line) that is not ratio monotone at $\hat{v}$. In subfigure (a), the allocation from posting price $\hat{v}$ for alternative 1 is depicted by the shaded regions. In subfigure (b), the allocations from posting prices $\hat{v}$ for alternative 1 and $C(\hat{v}) - \epsilon$ for alternative 2 are depicted by the shaded region. On the path, types with value $\{t\}_1$ for alternative 1 in interval $[\hat{v}^\dagger, \hat{v}^\ddagger]$ will buy alternative 2. Relative to simply posting a price for alternative 1, offering alternative 2 at a discount adds revenue from types with $\{t\}_1 \in [\hat{v}^\dagger, \hat{v}]$ and loses revenue from types with $\{t\}_1 \in [\hat{v}, \hat{v}^\ddagger]$. With constant density $f_{\max}$, types that add revenue have measure $(\hat{v} - \hat{v}^\dagger)\, f_{\max} = \epsilon f_{\max}/C'(\hat{v})$; types that lose revenue have measure $(\hat{v}^\ddagger - \hat{v})\, f_{\max} = \epsilon f_{\max}/1-C'(\hat{v})$.

One way to view the result above is that, fixing a ratio-monotone path C, regularity of the distribution $F_{\max}$ of value for alternative 1 implies that the optimal mechanism is to post a price for alternative 1. We now show that regularity implies that price posting is optimal only if the path is ratio monotone.

Theorem 8.32. *For any non-ratio-monotone path C there exists a regular distribution $F_{\max}$ such that posting a price for alternative 1 is not optimal for the path-based agent defined by C and $F_{\max}$.*

Proof. This proof is by counter example. Suppose that the path C is not monotone at some value $\hat{v} \in (0, 1)$, i.e.,

$$C'(\hat{v}) < C(\hat{v})/\hat{v}. \tag{8.14}$$

Consider the uniform distribution on $[0, 2\hat{v}]$ truncated at 1 with a point-mass. By construction the monopoly price is $\hat{v}$ (thus, $\hat{v}$ is the optimal price to post for alternative 1) and the density function is constant at $f_{\max} = 1/2\hat{v}$.

For the remainder of the proof, assume that $C(\cdot)$ is locally linear at $\hat{v}$ (the general proof is deferred to Exercise 8.13). Consider adding the option of buying alternative 2 at price $C(\hat{v}) - \epsilon$. There is a gain from types who were not buying before who now buy and a loss from types who were buying alternative 1 before but now switch to the lower cost alternative 2; see Figure 8.15. These are:

$$\begin{aligned} \text{Gain}(\epsilon) &= (C(\hat{v}) - \epsilon)\, \tfrac{\epsilon f_{\max}}{C'(\hat{v})}. \\ \text{Loss}(\epsilon) &= (\hat{v} - C(\hat{v}) + \epsilon)\, \tfrac{\epsilon f_{\max}}{1 - C'(\hat{v})}. \end{aligned}$$

The first term in each expression above is the gain or loss from each type; the second term is the measure of such types.

To see that the gain is more than the loss in the limit as ϵ goes to zero, we can divide each by $\epsilon f_{\max}$ and take their limits.

$$\begin{aligned} \lim_{\epsilon \to 0} \left[{}^{\text{Gain}(\epsilon)}\!/\!{}_{\epsilon f_{\max}} \right] &= {}^{C(\hat{v})}\!/\!{}_{C'(\hat{v})} > \hat{v}. \\ \lim_{\epsilon \to 0} \left[{}^{\text{Loss}(\epsilon)}\!/\!{}_{\epsilon f_{\max}} \right] &= {}^{\hat{v} - C(\hat{v})}\!/\!{}_{1 - C'(\hat{v})} < \hat{v}. \end{aligned}$$

The inequalities of both lines follow from equation (8.14), by rearranging as ${}^{C(\hat{v})}\!/\!{}_{C'(\hat{v})} > \hat{v}$ for the first line and because it implies $\hat{v} - C(\hat{v}) < \hat{v}\,[1 - C'(\hat{v})]$ for the second line. Thus, the gain is strictly more than the loss and posting price $\hat{v}$ for alternative 1, which is optimal among such price postings, is not optimal among all mechanisms. □

8.7.4 Uniform Pricing for Ratio-monotone Distributions

This section duplicates the analysis from the warmup (Section 8.7.1) for more general distributions F on the alternative-1 preferred type space, i.e., $\mathcal{T} = \{t \in [0,1]^2 : \{t\}_2 \leq \{t\}_1\}$. In the previous analysis the mechanism design problem was decomposed into a collection of paths, solved on each path in the collection, and then it was argued that these solutions are consistent with a single mechanism. Critical in this analysis is the choice of the paths. In The results of this section are based on a natural guess at the right paths; a principled method for determining the right paths will be described in subsequent sections.

The following discussion gives a natural guess as to the right paths on which to decompose the multi-dimensional mechanism design problem. We would like to solve the mechanism design problem independently on these paths and for the optimal mechanism on each path to post the same price for alternative 1 (and never sell alternative 2). One sufficient condition to guarantee that the optimal mechanisms for selling alternative 1 on each path posts the same price is to require the distribution of the agent's value for alternative 1, conditioned on the path on which the agent's type lies, be the same for all paths.

Definition 8.20. The *equiquantile path* C_θ solves

$$\mathbf{Pr}_{t\sim F}[\{t\}_2 \leq C_\theta(v) \mid \{t\}_1 = v] = \theta.$$

The *equiquantile type subspace* is $\mathcal{T}_\theta = \{(v, C_\theta(v)) : v \in [0,1]\}$.

Lemma 8.33. *For any quantile θ, the conditional distribution of $\{t\}_1$ given $t \in \mathcal{T}_\theta$ for $t \sim F$ is equal to the unconditional distribution of $\{t\}_1$, i.e., for all $z \in [0,1]$,*

$$\mathbf{Pr}_{t\sim F}[\{t\}_1 \leq z \mid t \in \mathcal{T}_\theta] = \mathbf{Pr}_{t\sim F}[\{t\}_1 \leq z]\,.$$

Proof. By definition, given $\{t\}_1$ the quantile θ corresponding to the path type space $\mathcal{T}_\theta$ that contains type t is uniformly distributed. Therefore, θ is independent of $\{t\}_1$; equivalently, $\{t\}_1$ is independent of θ. Thus, the conditional distribution of $\{t\}_1$ given θ is the same as its unconditional distribution. □

We are now ready to complete the construction. The set of equiquantile paths $\{C_\theta : \theta \in [0,1]\}$ partition type space. Supposing the designer knew the path C_θ on which the type was drawn, then the designer would employ the optimal mechanism for that path. If the distribution of $\{t\}_1$ is regular and the equiquantile curves are ratio monotone, then by Theorem 8.31 the optimal mechanism for each path is to post the monopoly price $\hat{v}^\star_{\max}$ for the distribution $F_{\max}$ of $\{t\}_1$. This mechanism is the same regardless of the path; thus, posting price $\hat{v}^\star_{\max}$ for alternative 1 is revenue optimal. From this argument, we can conclude the following theorem.

Theorem 8.34. *For distribution F on the alternative-1 preferred type space $\mathcal{T} = \{t \in [0,1]^2 : \{t\}_2 \leq \{t\}_1\}$ satisfying (a) the distribution of $\{t\}_1$ is regular and (b) the equiquantile paths are ratio monotone, the revenue optimal mechanism is to post the monopoly price for alternative 1.*

This theorem can easily be extended to distributions on the unit square for which the distribution of the agent's value for her preferred alternative is independent of which alternative is preferred. Under this assumption, the problem can be independently solved under each conditioning, and the mechanism that posts a uniform price for each alternative is optimal.

This theorem can also be easily generalized to selling a single item in one of two configurations (which we will continue to refer to as alternatives) to several agents. Notice that if there was a cost c for serving the agent, then the optimal mechanism of Theorem 8.34 would still be a posted price for alternative 1. The price $\hat{v}$ however would be increased to solve $\hat{v} - \frac{1-F_{\max}(\hat{v})}{f_{\max}(\hat{v})} = c$. We conclude that the optimal mechanism for several multi-dimensional agents (that each satisfy the assumptions of the theorem) is simply the optimal mechanism that projects each agent into a single dimension according her value for her favorite alternative.

The next section will formalize the method of virtual values employed above to prove the optimality of posting the monopoly price for alternative 1. The section following will use this formulation to give a general method for solving for the appropriate paths on which to solve the mechanism design problem.

8.7.5 Multi-dimensional Virtual Values

In this section we generalize the virtual-value-based approach to mechanism design from the single-dimensional agents of Section 3.3.2 on page 64 to multi-dimensional agents. Assume that the seller has a cost function for producing a given allocation of alternatives x that is specified by $c(x)$. The definitions below are given for a single agent and the objective of profit, i.e., expected payment minus expected cost, but they could be equally well defined for multiple agents and any objective.

Definition 8.21. A *virtual value function* ϕ is a vector field that satisfies three properties:

(i) *Amortization of revenue*: For any incentive compatible mechanism (x, p), the agent's expected virtual surplus is an upper bound on expected revenue, i.e., $\mathbf{E}[\phi(t) \cdot x(t)] \geq \mathbf{E}[p(t)]$.
(ii) *Incentive compatibility*: A point-wise virtual surplus maximizer $x^\star(t) \in \operatorname{argmax}_x \phi(t) \cdot x - c(x)$ is incentive compatible, i.e., there exists a payment rule $p^\star$ such that mechanism $(x^\star, p^\star)$ is incentive compatible.
(iii) *Tightness*: For this point-wise virtual surplus maximizer $(x^\star, p^\star)$, the agent's expected virtual surplus is equal to the expected revenue, i.e., $\mathbf{E}\big[\phi(t) \cdot x^\star(t)\big] = \mathbf{E}\big[p^\star(t)\big]$.

Definition 8.21 makes a distinction between the agent's virtual surplus $\phi(t) \cdot x$ and the virtual surplus of the mechanism $\phi(t) \cdot x - c(x)$ which includes the seller's cost. A special case of interest is a *uniform cost* c where virtual surplus is $\phi(t) \cdot x - c \sum_j \{x\}_j$. This uniform cost could represent the *opportunity cost* the seller faces for serving this agent. For example, with two agents with virtual value functions and a single-item environment, the opportunity cost of serving one agent is the maximum of the virtual value of the other agent and zero.

Proposition 8.35. *For any mechanism design problem that admits a virtual value function, a virtual surplus maximizer is the optimal mechanism.*

Proof. Denote the incentive compatible virtual surplus maximizer (guaranteed to exist by Definition 8.21) by allocation and payment rules $(x^\star, p^\star)$; denote any other incentive compatible mechanism by allocation and payment rules (x, p); then,

$$\begin{aligned}\mathbf{E}_t\big[p^\star(t) - c(x^\star(t))\big] &= \mathbf{E}\big[\phi(t) \cdot x^\star(t) - c(x^\star(t))\big] \\ &\geq \mathbf{E}[\phi(t) \cdot x(t) - c(x(t))] \geq \mathbf{E}\big[p(t) - c(x^\star(t))\big] .\end{aligned}$$

The first equality is by tightness (the expected cost term $\mathbf{E}\big[c(x^\star(t))\big]$ is the same on both sides of the equality), the second inequality is by the fact that $x^\star$ is a virtual surplus maximizer, the third inequality is because ϕ is an amortization of revenue (again, the expected cost term $\mathbf{E}[c(x(t))]$ is the same on both sides of the inequality). □

In Section 8.7.3, we derived a virtual value function for types on a ratio-monotone path and with regularly distributed value for the preferred alternative. In Section 8.7.4, we guessed a set of paths, solved the mechanism design problem on each path, and argued that under some distributional assumptions these optimal mechanisms are consistent with one mechanism. In this section we will develop a general framework for deriving virtual value functions absent a good guess of the decomposition to paths. The main idea is to leave the paths as variables, that can then be solved for later. To do this we will employ a multi-dimensional integration by parts (which is defined with respect to any vector field; see the Mathematical Note on page 316), with the constraint that this vector field corresponds to paths.

Our goal is to use integration by parts to rewrite the expected utility $\mathbf{E}[u(t)] = \int_{t \in \mathcal{T}} u(t)\, f(t)\, \mathrm{d}t$ in terms of the gradient of the utility ∇u, which by Theorem 8.28 is equal to the allocation x, and a boundary integral. Thus, we need to identify a vector field α with divergence equal to the (negated) density. The term corresponding to the boundary integral, we would prefer to be zero, but for an upper bound it would be sufficient for it to be negative. Thus, we seek a vector field α to satisfy properties of the following definition.

Definition 8.22. For distribution F and type space $\mathcal{T}$, a vector field $\alpha : \mathbb{R}^m \to \mathbb{R}^m$ satisfies

- the *divergence density equality* if $\nabla \cdot \alpha(t) = -f(t)$ for all $t \in \mathcal{T}$, and

- *boundary influx* if $\alpha(t) \cdot \eta(t) \leq 0$ for all $t \in \partial\mathcal{T}$.

Theorem 8.36. *If $\alpha : \mathbb{R}^m \to \mathbb{R}^m$ is a vector field satisfying the divergence density equality and boundary influx on type space $\mathcal{T}$ then vector field $\phi(\mathcal{T}) = t - {}^{\alpha(t)}/_{f(t)}$ is an amortization of revenue. Moreover, the amortization ϕ is tight for incentive compatible mechanisms that have binding individual rationality constraint $u(t) = 0$ on all boundary types t with non-trivial flux, i.e., $\alpha(t) \cdot \eta(t) \neq 0$.*

Mathematical Note. *Multi-dimensional integration by parts* is defined for function $u : \mathbb{R}^m \to \mathbb{R}$ and vector field $\alpha : \mathbb{R}^m \to \mathbb{R}^m$ on region $\mathcal{T}$ with boundary $\partial\mathcal{T}$ as follows:

$$\int_{t \in \mathcal{T}} \nabla u(t) \cdot \alpha(t) \, \mathrm{d}t = \int_{t \in \partial\mathcal{T}} u(t) \, (\alpha(t) \cdot \eta(t)) \, \mathrm{d}t - \int_{t \in \mathcal{T}} u(t) \, (\nabla \cdot \alpha(t)) \, \mathrm{d}t.$$

In the formula above, $\nabla \cdot \alpha(t)$ is the divergence of α at point $t \in \mathcal{T}$ and is defined as $\nabla \cdot \alpha(t) = \sum_j \{\nabla \alpha(t)\}_j$; and $\eta(t)$ is a unit-length normal vector to the boundary at point $t \in \partial\mathcal{T}$.

The *divergence theorem* is the application of multi-dimensional integration by parts to the vector field α and the function $u(\cdot) = 1$ (which has trivial gradient $\nabla u(\cdot) = (0, \ldots, 0)$). Viewing the vector field as a flow, the divergence theorem shows that the divergence of a flow in a region $\mathcal{T}$ is equal to magnitude of the flux out of the region.

$$\int_{t \in \mathcal{T}} \nabla \cdot \alpha(t) \, \mathrm{d}t = \int_{t \in \partial\mathcal{T}} \alpha(t) \cdot \eta(t) \, \mathrm{d}t.$$

Proof. Rewrite expected utility as

$$\begin{aligned}
\mathbf{E}[u(t)] &= \int_{t\in\mathcal{T}} u(t)\, f(t)\,\mathrm{d}t \\
&= -\int_{t\in\mathcal{T}} u(t)\,(\nabla\cdot\alpha(t))\,\mathrm{d}t \\
&= -\int_{t\in\partial\mathcal{T}} u(t)\,(\alpha(t)\cdot\eta(t))\,\mathrm{d}t + \int_{t\in\mathcal{T}} \nabla u(t)\cdot\alpha(t)\,\mathrm{d}t \\
&\geq \int_{t\in\mathcal{T}} \nabla u(t)\cdot\alpha(t)\,\mathrm{d}t \\
&= \int_{t\in\mathcal{T}} x(t)\cdot\alpha(t)\,\mathrm{d}t \\
&= \mathbf{E}[x(t)\cdot {}^{\alpha(t)}/_{f(t)}]\,.
\end{aligned}$$

The first line is the definition of expectation, the second line applies the divergence density equality, the third line is integration by parts, the fourth line follows from individual rationality, i.e., $u(t) \geq 0$ for all $t \in \mathcal{T}$, and boundary influx (implying that the first term on the third line is non-negative), the fifth line is from Theorem 8.28, and the sixth line is the definition of expectation.

A type t with binding participation constraint has zero utility, i.e., $u(t) = 0$. If all boundary types with non-trivial boundary influx, i.e., with $\alpha(t)\cdot\eta(t) \neq 0$, have binding participation constraint $u(t) = 0$ then $u(t)\,(\alpha(t)\cdot\eta(t)) = 0$ at all boundary types $t \in \partial\mathcal{T}$. In this case, the first term on the third line is identically zero and the whole sequence of inequalities is tight.

Expected revenue is equal to the expected surplus less the agent's expected utility, i.e.,

$$\begin{aligned}
\mathbf{E}[p(t)] &= \mathbf{E}[t\cdot x(t)] - \mathbf{E}[u(t)] \\
&\leq \mathbf{E}[t\cdot x(t)] - \mathbf{E}[{}^{\alpha(t)}/_{f(t)}\cdot x(t)] \\
&= \mathbf{E}[\phi(t)\cdot x(t)]\,.
\end{aligned}$$

The second line is from the previous derivation, and the third line is from the definition of vector field ϕ. This sequence of inequalities is tight when the previous sequence of inequalities is tight. □

While Theorem 8.36 does not explicitly mention paths, paths are implicit in the choice of vector fields α that give tight amortizations. With two-dimensional type space oriented as for considering the optimality of pricing only alternative 1 with weaker types towards the left and

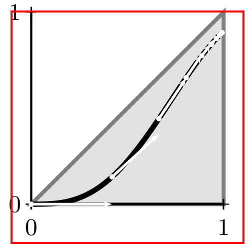

Figure 8.16. Depicted is the type space $\mathcal{T} = \{\boldsymbol{t} \in [0,1]^2 : \{t\}_2 \leq \{t\}_1\}$ (light gray region), its boundary $\partial\mathcal{T}$ (dark gray boarder), and a path (thick, black line) through type space. The vector field α (white arrows) that corresponds to this path is shown. The lengths of the arrows are proportional to the magnitude of vectors in the field (the first coordinate of which is the remaining probability density f left to distribute on the path). All such paths originate at type $(0,0)$ (the left boundary) and terminate on the right boundary, i.e., with types $\{t \in \mathcal{T} : \{t\}_1 = 1\}$.

stronger types towards the right, the boundary will have four regions; see Figure 8.16. Paths will originate on the left boundary with an influx of flow. On this path the direction of vector field α is the direction of the path; the magnitude of the first coordinate $\{\alpha(t)\}_1$ of this flow is cumulative the density on the remainder of the path. The path terminates on the right boundary where vector field α is the zero vector (there is no remaining density according to f). The top and bottom boundary regions are parallel to paths and thus the dot-product of vector field α with the normal to the boundary is zero. By these interpretation, the influx on the boundary is trivial (equal to zero) on all types except those on the left where the paths originate. If these types t on the left are chosen so that the mechanisms under consideration have binding individual rationality constraint, i.e., $u(t) = 0$, then vector field ϕ constructed in Theorem 8.36 is tight, as desired.

Definition 8.23. An amortization of revenue ϕ is *canonical* if is derived as $\phi(t) = t - {}^{\alpha(t)}/_{f(t)}$ from vector field α that satisfies the divergence density equality and boundary influx.

Notice that a vector field α that satisfies the divergence density equality and boundary influx will have divergence -1 on type space $\mathcal{T}$. Consequentially, by the divergence theorem, the outflux on the boundary must also be -1. For mechanism where this non-trivial outflux is concentrated on boundary types that have zero utility, then the amortization of revenue ϕ defined from α is tight.

Example 8.8. Consider a single-dimensional agent with type t uniformly distributed on type space $\mathcal{T} = [1, 2]$. The density function is $f(t) = 1$. There is only one path and it goes from type 1 to type 2. The (single-dimensional) vector field α at t is the remaining cumulative density on $[t, 2]$, i.e., $\alpha(t) = 2 - t$. Notice that there is boundary influx only at type $t = 1$. Any mechanism where type 1 has zero utility, e.g., posting a price at 1 or higher, will have a binding participation constraint for type 1. Thus, $u(1)\alpha(1) \cdot \eta(1) = 0$ where the (single-dimensional) normal vector at 1 is $\eta(1) = -1$. The resulting amortization of revenue is vector field $\phi(t) = t - {}^{\alpha(t)}/_{f(t)} = 2t - 2$. Notice that $\alpha(t)$ is equal to $1 - F(t)$ (for cumulative distribution function $F(t) = t - 1$), so this formula is identical to the single-dimensional virtual value derived in Chapter 3. Notice that for mechanisms that post prices less than one, say, at $1/2$, the amortization of revenue is not tight. The expected virtual surplus of this mechanism is one, while its revenue is $1/2$, the virtual surplus less the utility of the weakest type, i.e., type 1.

Unfortunately, except in single dimensional environments, canonical amortizations of revenue are not unique. Any covering of type space by paths will give a canonical amortization. Generally, at most one of these canonical amortizations can be a virtual value function. In the next section we will develop a systematic method for identifying a virtual value function, or equivalently, the right set of paths.

We conclude this section by observing that the existence of a virtual value function for the family of single-agent environments with uniform costs implies revenue linearity (Definition 3.16), i.e., $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger]$ for $\hat{y} = \hat{y}^\dagger + \hat{y}^\ddagger$. Essentially, virtual surplus is a linear objective. Thus, as described in Section 8.3, multi-agent service constrained mechanism design problems reduce to single-agent ex ante problems.

Theorem 8.37. *Consider a unit-demand agent (given by type space and distribution), if vector field ϕ is a virtual value function for the single-agent environment with any non-negative uniform cost c then the agent is revenue linear.*

Proof. Sort the types t in decreasing order of the virtual value of the alternative with the highest virtual value, i.e. $\max_j \{\phi(t)\}_j$. Let $\hat{q}^\star$ be the measure of types where this highest virtual value is non-negative. The $\hat{q}$ ex ante optimal mechanism serves the first $\min(\hat{q}, \hat{q}^\star)$ measure of types in this order. The $\hat{y}$ interim optimal mechanism serves the first

$\hat{q}^{\star}$ measure of types greedily by this order (and discards the remaining types). By the linearity of virtual surplus, the expected virtual surplus of the latter is the appropriate convex combination of the expected virtual surplus of the former. Since expected virtual surplus equals expected revenue, the agent is revenue linear. □

8.7.6 Reverse Solving for Virtual Values

In multi-dimensional environments, because there are multiple ways to cover type space by paths, there are multiple canonical amortizations of revenue. If we can find an amortization of revenue that is incentive compatible, i.e., for which its pointwise optimization gives an incentive compatible mechanism (Definition 8.21), then Proposition 8.35 implies that this mechanism is optimal. Except in edge cases, optimal mechanisms for a single agent are unique. Thus, we are searching among these canonical amortizations for the one, if any, that is incentive compatible. In Section 8.7.4, we guessed the right paths, in this section we give a principled approach for identifying them.

Consider an agent with the two-dimensional alternative-1 preferred type space, i.e., $\mathcal{T} = \{t \in [0,1]^2 : \{t\}_1 \geq \{t\}_2\}$. The goal of this setting is to describe sufficient conditions on the distribution F (as specified by density function f) so that posting a price for alternative 1 (only) is revenue optimal. As in Section 8.7.4 the solution to this problem will generalize to the full type space $[0,1]^2$; moreover, it will also generalize to $m \geq 2$ alternatives.

Definition 8.24. The *single-dimensional favorite-alternative projection* is given by value $v = \{t\}_1$, distribution function $F_{\max}$, density function $f_{\max}$, amortization of revenue $\phi_{\max}(v) = v - \frac{1-F_{\max}(v)}{f_{\max}(v)}$, and monopoly price $\hat{v}^{\star}_{\max}$ that solves $\phi_{\max}(\hat{v}^{\star}_{\max}) = 0$.

Proposition 8.38. *For non-negative uniform costs c, a vector field ϕ is a virtual value that proves the the optimality of the favorite-alternative single-dimensional projection, i.e., the mechanism that projects the agent's type to her value for alternative-1 and is optimal for this projection, if (a) the alternative-1 virtual value $\{\phi(\cdot)\}_1$ is a virtual value for the single-dimensional projection, and (b) the alternative 2 virtual value never maximizes virtual surplus, i.e., $\{\phi(t)\}_1 \geq 0$ implies $\{\phi(t)\}_2 \leq \{\phi(t)\}_1$ and $\{\phi(t)\}_1 \leq 0$ implies $\{\phi(t)\}_2 \leq 0$ for all types t.*

Proof. By property (b), virtual surplus maximization only serves alternative 1 (or nothing if $\{\phi(t)\}_1 < c$);[9] by property (a) and Proposition 8.35, virtual surplus maximization is optimal among all mechanisms that only sell alternative alternative 1. □

The goal of this section is to identify a vector field ϕ that satisfies the conditions of Proposition 8.38. The approach is to use property (a) of the proposition to reduce a degree of freedom in defining a canonical amortization, and then to identify conditions on the distribution that are sufficient to imply property (b). Specifically, set the first coordinate of the virtual value function, denoted $\{\phi(t)\}_1$, to the virtual value of the single-dimensional projection, denoted $\phi_{\max}(\{t\}_1)$. The definition of the canonical amortization ϕ for vector field α (Definition 8.23) gives $\{\alpha(t)\}_1$ from $\{\phi(t)\}_1$; the divergence density equality gives $\{\alpha(t)\}_2$ from $\{\alpha(t)\}_1$ (and identifies the right paths); and Definition 8.23, again, gives $\{\phi(t)\}_2$ from $\{\alpha(t)\}_2$. With amortization $\phi(\cdot)$ fully defined, sufficient conditions on the distribution to imply property (b) can be identified.

Definition 8.25. The *two-dimensional extension of the favorite-alternative projection* defines vector fields ϕ and α as follows:

(i) $\{\phi(t)\}_1 = \phi_{\max}(\{t\}_1)$,
(ii) $\{\alpha(t)\}_1 = \big[\{t\}_1 - \{\phi(t)\}_1\big]\, f(t) = \frac{1-F_{\max}(\{t\}_1)}{f_{\max}(\{t\}_1)}\, f(t)$,
(iii) $\{\alpha(t)\}_2 = -\int_0^{\{t\}_2} \big[f(\{t\}_1, z) + {}^{\mathrm{d}}/_{\mathrm{d}\{t\}_1}\{\alpha(\{t\}_1, z)\}_1\big]\, \mathrm{d}z$.
(iv) $\{\phi(t)\}_2 = \{t\}_2 - {}^{\{\alpha(t)\}_2}/_{f(t)}$.

Lemma 8.39. *The two-dimensional extension of the favorite-item projection defines vector field α that satisfies the divergence density equality and boundary inflow, and vector field ϕ is a canonical amortization that is tight for mechanisms for which individual rationality binds on type $(0,0)$.*

Proof. By Theorem 8.36, it suffices to show that vector field α satisfies the divergence density equality and trivial boundary influx at $t \in \partial\mathcal{T} \setminus \{(0,0)\}$.

The divergence density equality is satisfied by definition; to see this, differentiate both sides of the definition of $\{\alpha(t)\}_2$ with respect to $\{t\}_2$.

[9] Notice that the assumptions of the proposition are insufficient if the uniform service cost c is negative. With a negative service cost the agent may receive a non-trivial alternative even when her virtual value for this alternative is negative. This restriction limits the applicability of the proposition to single-agent problems where the ex ante constraint holds as an inequality, and to multi-agent service-constrained environments that are downward closed.

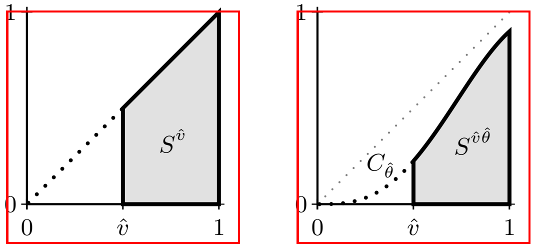

Figure 8.17. Depicted are the regions to which the divergence theory is applied in Lemma 8.39 and Lemma 8.40. The flux on the right and bottom boundaries are zero by definition. In subfigure (a), the divergence of the subspace of types $S^{\hat{v}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v}\}$ (gray region) and the outflux of the left boundary are both equal to $-(1 - F_{\max}(\hat{v}))$. In subfigure (b), the divergence of the the subspace of types $S^{\hat{v}\hat{\theta}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v} \wedge \{t\}_2 \leq C_{\hat{\theta}}(\{t\}_1)\}$ (gray region) and the outflux of the left boundary are both equal to $-\hat{\theta}\,(1 - F_{\max}(\hat{v}))$. For both, the remaining flux out the top boundary must be zero.

We now analyze the flux $\alpha(t) \cdot \eta(t)$ for types t on the right, bottom, and top boundaries.

- Right boundary, i.e., t with $\{t\}_1 = 1$: Normal $\eta(t) = (1, 0)$ and $\{\alpha(t)\}_1 = 0$ as $1 - F_{\max}(1) = 0$, so $\alpha(t) \cdot \eta(t) = 0$.
- Bottom boundary, i.e., t with $\{t\}_2 = 0$: Normal $\eta(t) = (0, -1)$ and $\{\alpha(t)\}_2 = 0$ as the integral from 0 to 0 of any function is zero, so $\alpha(t) \cdot \eta(t) = 0$.
- Top boundary, i.e., t with $\{t\}_1 = \{t\}_2$: Apply the divergence theorem to vector field α on the type subspace with value at least $\hat{v}$ for alternative 1, i.e., $S^{\hat{v}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v}\}$; see Figure 8.17. The divergence theorem requires that the divergence of α on this subspace is equal the outflux. By the divergence density equality, the divergence of subspace $S^{\hat{v}}$ is $-\int_{t \in S^{\hat{v}}} f(t)\,\mathrm{d}t = -(1 - F_{\max}(\hat{v}))$. The outflux on the left-boundary is $-\int_0^1 \{\alpha(\hat{v}, z)\}_1\,\mathrm{d}z = -\int_0^1 \frac{1 - F_{\max}(\hat{v})}{f_{\max}(\hat{v})} f(\hat{v}, z)\,\mathrm{d}z$. Note that the density of the favorite-alternative projection is $f_{\max}(\hat{v}) = \int_0^1 f(\hat{v}, z)\,\mathrm{d}z$ and, thus, this outflux is $-(1 - F_{\max}(\hat{v}))$. The remaining total outflux for the top boundary is zero. This equality holds for all values $\hat{v}$, thus the outflux at each type t on the top boundary is identically zero.

Though unnecessary for the proof, observe that the above calculation of the outflux on the left-boundary applied to full type space $\mathcal{T}$ (equal to the subspace $S^{\hat{v}}$ when $\hat{v} = 0$) implies that the outflux at type $(0, 0)$ in the direction of $(-1, 0)$ is -1. ☐

The remaining task is to identify sufficient conditions on the distribution F so that uniform pricing optimizes virtual surplus with respect to the amortization defined by the two-dimensional extension of the favorite-alternative projection. Recall that the monopoly price $\hat{v}^{\star}_{\max}$ for the distribution $F_{\max}$ is the optimal price to post for alternative 1 when the agent's value for the alternative is drawn from distribution $F_{\max}$. A sufficient condition on the virtual value function ϕ is that (a) for types t with $\{t\}_1 \geq \hat{v}^{\star}_{\max}$ that $\{\phi(t)\}_2 \leq \{\phi(t)\}_1$; and (b) for types t with $\{t\}_1 < \hat{v}^{\star}_{\max}$, that $\{\phi(t)\}_2 \leq 0$.

Our approach will be to show that the paths defined by vector field α, i.e., the direction of α, are the equiquantile paths (Definition 8.20) and that ratio-monotonicity of these paths (Definition 8.19) implies both conditions (a) and (b) when the distribution of the agents value $\{t\}_1$ for her favorite alternative is regular.

Lemma 8.40. *The vector field α of the two-dimensional extension of the favorite-item projection corresponds to the equiquantile paths.*

Proof. Consider subspace of types who value alternative 1 more than $\hat{v}$ but lie below the $\hat{\theta}$-equiquantile path $C_{\hat{\theta}}$, i.e., $S^{\hat{v}\hat{\theta}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v} \wedge \{t\}_2 \leq C_{\hat{\theta}}(\{t\}_1)\}$; see Figure 8.17. We will show that the outflux of α on the top boundary, namely the types on the path $C_{\hat{\theta}}$, is zero by applying the divergence theorem to subspace $S^{\hat{v}\hat{\theta}}$.

By the divergence density equality, the divergence of α on subspace $S^{\hat{v}\hat{\theta}}$ can be evaluated as:

$$\begin{aligned}\int_{t \in S^{\hat{v}\hat{\theta}}} \nabla \cdot \alpha(t)\, \mathrm{d}t &= -\int_{t \in S^{\hat{v}\hat{\theta}}} f(t)\, \mathrm{d}t \\ &= -\mathbf{Pr}\big[t \in S^{\hat{v}\hat{\theta}}\big] \\ &= -\hat{\theta}\, [1 - F_{\max}(\hat{v})].\end{aligned}$$

The first line is by the divergence density, and the second line is by the definition of probability. The third line follows because the probabilities that $\{t\}_1 \geq \hat{v}$ and $\{t\}_2 \leq C_{\hat{\theta}}(\{t\}_1)$ are independent (Lemma 8.33); equal to $[1 - F_{\max}(\hat{v})]$ and $\hat{\theta}$, respectively; and the probability of the intersection of two independent events is the product of their probabilities.

As before, the outflux of α on the right and bottom boundary is zero. The outflux of α on the left boundary can be calculated by integrat-

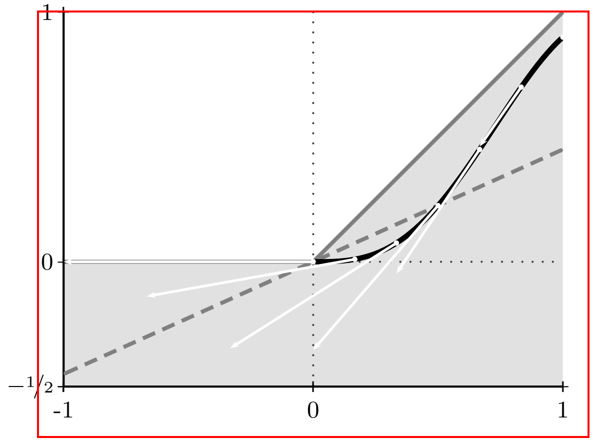

Figure 8.18. An equiquantile path C_θ (solid, thick, black line) is depicted. For the case where the distribution of $\{t\}_1$ is uniform on $[0,1]$ the virtual values for six types on this path are depicted. By the definition of canonical amortization, $\phi(t) = t - {}^{\alpha(t)}/_{f(t)}$; thus, the virtual values can be calculated by starting at the type t (white bullets) and moving backwards (white arrows) a magnitude of $|{}^{\alpha(t)}/_{f(t)}|$ in the direction of $-\alpha(t)$. The gray shaded region represents the allowable space $\{\phi \in [-1,1]^2 : \{\phi\}_2 \leq \{\phi\}_1 \vee \{\phi\}_2 \leq 0\}$ of virtual values for which selling only alternative 1 is optimal. Ratio monotonicity of the path implies that the slope of the path at a type t is greater than that of the line through the type and the origin (depicted for type $({}^1/_2, C_\theta({}^1/_2))$; gray, dashed line). Since the direction of α is tangent to the path, the virtual value for the type is below the line and, consequentially, within the allowable space.

ing the formula $\{\alpha(t)\}_1 = \frac{1-F_{\max}(\{t\}_1)}{f_{\max}(\{t\}_1)} f(t)$ with respect to $\{t\}_2$ when $\{t\}_2 = \hat{v}$. The first term in this integral is independent of $\{t\}_2$ and can be factored out. Integrating second-term, i.e., the density, on the left boundary up to quantile $\hat{\theta}$ gives $\hat{\theta}\, f_{\max}(\hat{v})$ as the total density of types with $\{t\}_1 = \hat{v}$ is $f_{\max}(\hat{v})$. The product of these quantities gives influx

$$\frac{1 - F_{\max}(\hat{v})}{f_{\max}(\hat{v})} \times \hat{\theta}\, f_{\max}(\hat{v}) = \hat{\theta}\,(1 - F_{\max}(\hat{v})),$$

and its negation is the outflux. Thus, the remaining outflux for the top boundary of type subspace $S^{\hat{v}\hat{\theta}}$ is zero. This equality holds for all values $\hat{v}$, thus the outflux at each type t on the top boundary is identically zero. In other words, α is tangent to the equiquantile paths. □

Theorem 8.41. *For distribution F on type space $\mathcal{T} = \{t \in [0,1]^2 : \{t\}_2 \leq \{t\}_1\}$ satisfying (c) the distribution of $\{t\}_1$ is regular and (d) the* *equiquantile paths are ratio monotone, the two-dimensional extension of the favorite-alternative projection ϕ satisfies properties (a) and (b) of Proposition 8.38, i.e., the alternative 1 coordinate of ϕ is a virtual value for the single dimensional projection and alternative 2 coordinate of ϕ never maximizes virtual surplus.*

Proof. Property (c), the regularity of the distribution of $\{t\}_1$, implies that the amortization of revenue for the single-dimensional projection is a virtual value function. Hence, property (a) holds.

We argue that property (d) implies property (b) as follows. Recall from Definition 8.23 that $\phi(t) = t - {}^{\alpha(t)}/_{f(t)}$. Consider this vector addition geometrically as the vector from the origin to t plus the vector back in the direction of $-\alpha(t)$ (the magnitude will not be important). Since the paths defined by vector field α are ratio monotone, this vector back lies below the line that connects the type t to the origin; see Figure 8.18. In other words,

$$\{\phi(t)\}_2 \leq \tfrac{\{t\}_2}{\{t\}_1} \{\phi(t)\}_1. \tag{8.15}$$

Thus, when $\{\phi\}_1 > 0$ then $\{\phi\}_2 \leq \{\phi\}_1$, and when $\{\phi\}_1 \leq 0$ then $\{\phi\}_2 \leq 0$. □

We wrap this section up by applying this virtual value theory to the two-alternative uniform unit-demand agent of Example 8.2 to solve the three single-agent problems of Section 8.1.

Example 8.9. For the two-alternative uniform unit-demand agent of Example 8.2 the two-dimensional extension of the single-dimensional projection satisfies the assumptions of Proposition 8.38. The virtual value function (on $t \in \{t \in \mathcal{T} : \{t\}_2 \leq \{t\}_1\}$):

$$\phi(t) = (1, {}^{\{t\}_2}/_{\{t\}_1}) \left[\{t\}_1 - {}^{1-F_{\max}(\{t\}_1)}/_{f_{\max}(\{t\}_1)}\right].$$

The three single-agent problems are solved by optimizing revenue under the single-dimensional favorite-alternative projection. The optimal unconstrained mechanisms posts a uniform price of $\sqrt{1/3}$, the $\hat{q}$ ex ante optimal mechanism for post a uniform price that sells with ex ante probability $\hat{q}$ (for $\hat{q} \leq \hat{q}^\star = 2/3$), the $\hat{y}$ interim optimal mechanism is the distribution over uniform prices that give allocation rule $y(q) = \hat{y}(q)$ for $q \leq \hat{q}^\star$ (and $y(q) = 0$, otherwise).

Exercises

8.1 Consider two agents with independent, identical, and uniformly distributed values on $[0,1]$ and budget $B = 1/4$. Solve for the equilibrium of the highest-bid-wins all-pay auction by identifying the critical type that is indifferent between following the budget unconstrained equilibrium of bidding $s(t) = t^2/2$ (see Section 2.8) and bidding the budget B assuming all higher types bid the budget. Is this equilibrium unique?

8.2 Prove Theorem 8.7: *Let* $\hat{\boldsymbol{y}} = (\hat{y}, \ldots, \hat{y})$ *be the* n*-agent allocation constraints induced by the* k *strongest-agents-win mechanism and* $\boldsymbol{y} = (y, \ldots, y)$ *the allocation rules induced by any symmetric* k*-unit mechanism for* n *i.i.d. agents, then* y *is feasible for* $\hat{y}$.

8.3 Let $\hat{\boldsymbol{y}} = (\hat{y}, \ldots, \hat{y})$ be the n-agent allocation constraints induced by the assortative matching of agents to n positions with weights $\boldsymbol{w} = (w_1, \ldots, w_n)$, i.e., stronger agents are matched to positions with larger weights. The position weights correspond to a stochastic probability of service. Let $\boldsymbol{y} = (y, \ldots, y)$ be the allocation rules induced by any symmetric mechanism for position weights $\boldsymbol{w}$ and n i.i.d. agents. Prove that y is feasible for $\hat{y}$. (Hint: Use Theorem 8.7.)

8.4 Consider a three-agent position environment with position weights 1, $1/2$, and 0 for the first, second, and third positions; respectively. Recall that in a position environment (Definition 7.18), an agent assigned to the kth position is served with probability given by the kth position weight. Consider three agents with values drawn independently, identically, and uniformly from the interval $[0,1]$ and each with a public budget constraint of $B = 1/4$ (Example 8.1). Derive the revenue optimal auction.

8.5 Consider the all-or-none set system that corresponds to a public project. Consider two agents with types drawn uniformly from type space $\mathcal{T} = \{L, H\}$ and characterize the class of symmetric interim feasible allocation rules. Specifically, what are the pairs of allocation probabilities $(x(L), x(H))$ that are induced by a symmetric ex post feasible mechanism?

8.6 Consider the all-or-none set system that corresponds to a public project. Consider two agents with quantiles drawn independently, identically, and uniformly from $[0,1]$ and characterize the class of interim feasible allocation rules $y : [0,1] \to [0,1]$ that are induced by a symmetric ex post feasible mechanism. (Recall, y is monotone non-increasing by definition.)

8.7 Prove that the quantile space and type space inequalities (8.8) and (8.9) that characterize interim feasibility are equivalent (Theorem 8.10).

8.8 Extend Theorem 8.14 from general feasibility environments, i.e., where only $\boldsymbol{x} \in \mathcal{X}$ are feasible, to general cost environments, i.e., where $\boldsymbol{x} \in \{0,1\}^n$ has service cost $c(\boldsymbol{x})$. Prove that if there exists a mechanism that induces the profile of allocation rules $\boldsymbol{y}$ with expected cost C, then there is a stochastic weighted optimizer that induces $\boldsymbol{y}$ with expected cost C.

8.9 Prove Theorem 8.18: *For any type distribution $\boldsymbol{F}$ and mechanisms $\hat{\boldsymbol{\mathcal{M}}}$ and $\boldsymbol{\mathcal{M}}$ with allocation rules $\hat{\boldsymbol{y}}$ and $\boldsymbol{y}$ satisfying $\boldsymbol{y} \preceq \hat{\boldsymbol{y}}$, the composite mechanism (Definition 8.13) induces a distribution over allocations that is in the downward closure of the distribution of allocations of $\hat{\boldsymbol{\mathcal{M}}}$ and the same interim mechanisms as $\boldsymbol{\mathcal{M}}$.*[10]

8.10 Use the analysis of the public budget agent from Section 8.6 with the uniform public-budget agent (Example 8.1; budget $B = 1/4$) to solve the single-agent problems as described in Section 8.1.1. In particular derive the unconstrained optimal mechanism, the $1/2$ ex ante optimal mechanism, and the $\hat{y}$ interim optimal mechanism for allocation constraint $\hat{y}(q) = 1 - q$. In your derivation of each of these mechanisms explicitly identify the correct Lagrangian parameter λ that gives the right Lagrangian revenue curves.

8.11 Consider the objective of optimizing welfare for an agent with a public budget. Adapt a version of Corollary 8.27 for the welfare objective. Specifically:

   (a) Identify the optimal auction for any interim allocation constraint $\hat{y}$ under sufficiently general assumptions on the distribution of types, e.g., regularity.
   (b) Clearly state the necessary assumptions on the distribution.
   (c) Identify the optimal auction for two agents with uniformly distributed types on $[0, 1]$ and public budget $B = 1/4$.

8.12 Consider a single-item, i.i.d., public-budget regular environment. Prove that the expected revenue of the highest-bid-wins all-pay auction (with no reserve or explicit intervals on which types are ironed) on $n + 1$ agents obtains at least optimal revenue for n agents. In other words, generalize Theorem 5.1 to public budgets.

[10] One distribution of allocations is *in the downward closure of* a second distribution of allocations if there is a coupling of the distributions so that the set of agents served by the first is a subset of those served by the second.

8.13 Extend the proof Theorem 8.32 given in the text to relax the assumption of locally linearity of the path $C(\cdot)$ at the price $\hat{v}$ where it is non-ratio-monotone. Hint: Instead of deriving limit equations for $\lim_{\epsilon \to 0}[{}^{\text{Gain}(\epsilon)}/{\epsilon f_{\max}}]$ give a general expression for Gain(ϵ), take its derivative, and evaluate at $\epsilon = 0$; likewise for Loss(ϵ).

Chapter Notes

Exact reductions from multi-agent to single-agent mechanism design problems for multi-dimensional and non-linear agents and the objective of revenue were considered by Alaei et al. (2012) and Alaei et al. (2013). The former considered the general case of non-revenue-linear agents; the latter defined revenue linearity as a property of interest and generalized the optimality of the marginal revenue mechanism of Myerson (1981) and Bulow and Roberts (1989).

For single-item environments the necessary and sufficient conditions for ex post implementation of an interim mechanism were developed by Maskin and Riley (1984), Matthews (1984), and Border (1991). For symmetric single-item environments, the latter gave a characterization of interim feasible mechanisms that is similar to the one presented here (characterizing them as stochastic weighted optimizers). These results were generalized to asymmetric single-item environments by Border (2007) and Mierendorff (2011) and to matroid environments by Alaei et al. (2012) and Che et al. (2013). In Cai et al. (2012a,b) and Alaei (2012) these results were generalized beyond matroids to show that any interim feasible allocation for a general feasibility environment could be implemented as a stochastic weighted optimizer optimization. Cai et al. (2012a,b) additionally address environments with multi-dimensional externalities and the results presented here for multi-service service constrained environments, such as n-agent m-item matching environments, are an adaptation of their results.

Optimization of revenue and welfare for single-dimensional agents with public budgets was considered by Laffont and Robert (1996) and Maskin (2000), respectively. The derivation in this text is a simplification of the one from Laffont and Robert (1996) that is enabled by the marginal revenue framework of Bulow and Roberts (1989) and can be found in Devanur et al. (2013). For the analogous, and more challenging, optimization problem where the budget of the agent is private see Pai and Vohra (2014) and Alaei et al. (2012).

The multi-dimensional characterization of Bayesian incentive compatibility is due to Rochet (1985). The canonical amortizations which underlie the theory of multi-dimensional virtual values were characterized by Rochet and Choné (1998) and further refined in Rochet and Stole, 2003. Armstrong (1996) developed the approach of integration by parts on rays from the origin which, as we described in the text, can be used to prove the optimality of uniform pricing for an agent with multi-dimensional type drawn from the uniform distribution. The methods given in the text for solving for optimal mechanisms on paths and for reverse-solving for the right paths are from Haghpanah and Hartline (2015).

Haghpanah and Hartline (2015) apply their framework for multi-dimensional virtual values to prove the optimality of uniform pricing, broadly. This result can be viewed as a "no haggling" result for substitutes. Single-dimensional no-haggling theorems come from Stokey (1979) and Riley and Zeckhauser (1983). The no-haggling characterization of Haghpanah and Hartline for multi-dimensional types on paths shows that these single-dimensional no-haggling result are on the boundary between haggling and no haggling. The part of this characterization that shows when haggling can be expected is a simplification of an example from Thanassoulis (2004).

A similar characterization of expected revenue (to the canonical amortizations presented in the text) was given by Daskalakis et al. (2015) where, instead of integration by parts to rewrite expected revenue in terms of the allocation rule, they use integration by parts to rewrite expected revenue in terms of expected utility. This alternative approach is also useful in identifying optimal mechanisms, e.g., see Giannakopoulos and Koutsoupias (2014).