Chapter 6

AUCTIONS

6.1 Introduction and Industry Overview

Auctions provide an alternative means of dynamically adjusting prices to match market conditions. An auction is simply a set of rules (called a *mechanism*) for specifying how information is revealed among customers and the firm, how goods are awarded to customers, and what payments are made from customers to the firm based on the revealed information. They differ from a dynamic posted-price mechanism in that typically customers are the ones who offer a price they are willing to pay—their *bid*—and the firm then decides which bids to accept. However, there are some auction formats that rather resemble posted pricing, in that the firm names a price, and customers simply indicate their willingness to buy at the offered price. As we show below, the prices in an auction depend both on the number of customers bidding and their valuations for items—and, not surprisingly, the more customers there are, or the more each customer values the items, the higher the prices generated by an auction. In this way, auction prices effectively "adapt" to market conditions, and hence they are often viewed as *price-discovery mechanisms.*

Auctions are important both practically and theoretically. On a practical level, auctions are encountered in many markets, including those for treasury bonds, livestock, used cars, electricity, foreign exchange, real estate, art and rare collectibles, fish, fresh flowers, industrial procurement, public-works contracts, and the sale of natural-resource rights (such as offshore oil and gas leases, logging rights, radio spectrum licenses, and so on). More recently, auctions have gained popularity with the success of e-commerce auction sites such as eBay.

From a revenue management perspective, in particular, auctions have some appealing features. First, they hold out the potential of achieving near-perfect, first-degree price discrimination, and although customers still retain some "information rent" that prevents a firm from capturing the entire consumer surplus, the revenue benefits over using a single price—or even second- and third-degree price discrimination—can be significant. Second, auctions have the potential to directly uncover these near-optimal prices without the need to estimate customers' demand functions or willingness to pay, though this statement again must be qualified somewhat as we explain shortly. Nevertheless, it is fair to say that most auction mechanisms generally require less information about customers than do alternative price-discrimination mechanisms.

On a theoretical level, auctions are important because they provide a rich framework for studying pricing mechanisms in settings where customers act strategically. Indeed, as we show in this chapter, auction theory can often be used to design *optimal mechanisms*—that is, mechanisms that maximize revenues among essentially all possible pricing mechanisms, under certain assumptions of course. In other cases, the theory provides convenient ways to compare the revenues produced by different pricing mechanisms. Also, the theory is based on a strategic (rational) consumer model, which adds to the realism of auction models relative to the (mainly) myopic models studied in Chapter 5.

We first look at some common examples of auctions in practice. Then, in Section 6.2, we describe the classical auction models and theory. Next, we look at dynamic auctions, both in the setting of selling a fixed capacity as in Chapter 2 and in a replenishment setting, where the firm orders and auctions over an infinite horizon as in the inventory-pricing problem of Section 5.3.2. Finally, we consider network auctions and discuss their relationship to the network RM problems of Chapter 3.

6.1.1 An Overview of Auctions in Practice

Auctions are used in a wide range of markets, including industrial, financial, and consumer markets. We briefly survey next each of these markets in turn.

6.1.1.1 Traditional Auction Houses

Traditional auction houses—the two largest being Christie's and Sotheby's—provide auctions for selling art, antiques, jewelry, wine, and other rare, high-value collectibles. Both have been doing so for a very long time indeed; Sotheby's was founded in 1744 and Christie's in 1766. Christie's is the market leader with sales of $2.3 billion in 2000. Both use a variation of an ascending, open-price (English) auction. (See Section 6.1.2 below for definitions of these auction types.) As of January 2000, Sotheby's started offering online auctions. These traditional auction houses generally limit themselves to high-value items, and their clientele are largely wealthy individuals and institutional collectors.

6.1.1.2 Financial-Market Auctions

Auctions have been used for many years in financial markets. Most government bonds and bills are sold at auctions, which are conducted at regular intervals to finance national debts. Investors (both institutional and individual) bid for the minimum interest rate they are willing to receive. The selling agency then sorts the bids and the bonds or bills are awarded to the lowest bidders until the desired amount of the issue is reached.

Auctions are also used by securities exchanges for trading stocks, bonds and foreign exchange. Typically, these are *double auctions* in which *bid* offers are made by customers and *ask* offers are made by sellers. A queue of bid and ask offers is maintained and trade takes place when the highest bid offer in queue exceeds the lowest ask offer in queue. (The rules for the price paid and how this matching takes place are usually specific to each exchange.)

6.1.1.3 Government Auctions

Governments use auctions for the sale of many public assets, including public lands, public industries (privatization sales), and natural-resource rights. A prominent example is the radio spectrum auctions for third-generation (3G) cellular phone service in Europe and the United States These spectrum auctions involved complex combinatorial features, in which communications companies bid for combinations of geographical areas to achieve coverage in a given market area. The sale prices produced by some of these auctions were staggering, and indeed the resulting debts incurred to finance these purchases have left many of the winning companies in a precarious financial position.

6.1.1.4 Industrial-Procurement Auctions

Auctions are also used in many industries for procurement of materials, services, and general subcontracting of production. Typically, this occurs through a request-for-quote (RFQ) process in which a buying firm details its requirements for a certain input, and selling firms submit price quotes to supply the input. Factors other than price, such as quality levels and delivery schedules, are typically important in the final selection as well.

Online versions of procurement auctions have also increased in the past several years. In the auto industry, the exchange Covisint was formed in early 2000 as a joint venture by Daimler-Chrysler, Ford Motor Company, and General Motors with technology provided by Commerce One and Oracle. The goal of the exchange is to facilitate integration and collaboration among suppliers and automakers, with the aim of lowering costs and facilitating more efficient business practices. The Covisint exchange supports a range of auction formats for procurement. FreeMarkets, which has been in operation since 1995, combines software products with market-making services that help facilitate real-time procurement auctions over the Internet. The company reports sales transactions of over $35 billion to date on their reverse-price auction systems and services. Many manufacturers also host their own private online auctions for procurement.

6.1.1.5 Consumer Online Auctions

Online consumer auctions have become popular, largely due to the success of eBay. eBay provides a platform for users to conduct auctions to buy and sell a wide range of items—a sort of Sotheby's OR Christie's for the common man.

An immense variety of items are sold on eBay—new, used, and collectibles, by both individuals and small businesses. It is by some measures, the most popular shopping site on the Internet as of this writing.[1] In 2001, eBay transacted more than $9.3 billion in gross merchandise sales. Most significantly, the company has proven that the Internet can be used to facilitate communication and trade among geographically dispersed individual buyers and sellers, allowing for the sort of real-time auction mechanisms that in the past required the physical presence of market participants.

Priceline.com provides a different online auction mechanism. It is based on what they term a "buyer-driven conditional purchase offer" mechanism,[2] in which customers declare what they are willing to pay for products and supplying firms accept or reject these offers. In return, consumers agree to varying degrees of flexibility in the brand and product features they receive for their offered price. This mechanism has proved quite popular as a channel for selling surplus airline seats and is gaining popularity for products such as discount phone service and home mortgages.

---

[1]For example, in 2000 eBay was the shopping site with the highest number of total user minutes according to Media Metrix.
[2]Priceline.com has been granted a United States patent for this invention.

Priceline is attractive to sellers in large part because the mechanism does not divulge the identity of the seller until after the purchase offer is accepted. (Customers bid on generic products and features, not specific brands.) This creates less of a pricing risk for a firm because it can discount without fear that its discounted prices will become widely known to other customers and to competitors. This feature produces *brand shielding* and such selling formats are often referred to as *opaque channels* in industry terminology. However, as we show in Section 6.3.3 below, under certain assumptions, this mechanism theoretically offers no benefit over list prices. (Priceline.com is discussed further in Chapter 10.)

6.1.2 Types of Auctions

There are a variety of mechanisms one can use to conduct an auction. For simplicity we focus first on the case of a firm auctioning a single indivisible good to a group of N customers. We then consider several variations of these simple, single-unit auctions.

6.1.2.1 Standard Auction Types

There are four common types of auctions for selling a single object:

- **Open ascending (English) auction** In an open ascending auction, the firm announces a progressively increasing sequence of prices. Customers indicate (say by raising their hand or showing a number) their willingness to buy an item at the announced price. The firm increases the price until only one customer is left willing to buy at the announced price. This is the mechanism commonly used to sell art and valuables at major auction houses such as Christie's and Sotheby's.
- **Open descending (Dutch) auction** In an open descending auction, the firm announces a progressively decreasing sequence of prices. The first customer to indicate willingness to buy at the announced price wins the item and pays the current price. The Aalsmeer and Naaldwijk flower markets in Holland have long used this type of auction, which explains the name.
- **Sealed-bid, first-price auction** In the sealed-bid, first-price auction, customers submit sealed bids to the firm. The customer submitting the highest bid wins the auction and pays the amount of his bid. This form of auction is used (in its minimization form) for awarding many government contracts.
- **Sealed-bid, second-price (Vickrey) auction** In the sealed-bid, second-price auction, customers again submit sealed bids, and the

customer submitting the highest bid wins the auction. However, the amount the winner pays is equal to the second-highest bid submitted. While this auction form has certain desirable theoretical properties, as shown by Vickrey [533], it is somewhat less common in practice.[3]

These basic auction types can be varied: for example, one may impose a reserve price or minimum bid increments. Moreover, there are other, less standard, auction types that are encountered in practice as well, such as the *uniform price auction* used in many financial markets. The above four types, however, are the most common.

An auction is called a *reverse auction* if customers are competing to sell to the auctioneer by submitting cost (or willingness to sell) bids rather than price (or willingness to buy) bids, such as in a procurement auction. Reverse auctions are essentially equivalent to regular auctions if we put a "minus sign" on the rewards (one involves maximization of price while the other involves minimization of cost), and hence we do not address them separately here.

6.1.2.2 Multiunit Auctions

Multiunit versions of the above auction types can also be defined in the natural way. For example, suppose the firm has C homogeneous items to sell and each customer wants only at most one item. Then in the C-unit open ascending auction, the firm announces increasing prices, and customers indicate their willingness to pay the offered prices. The price is increased until only C customers remain and each is awarded an item at the prevailing price. In an open descending auction, the price declines until a customer indicates willingness to pay the announced price. The customer is awarded a unit at that price, and the firm continues to decrease the price until a second customer is willing to pay the announced price, and so on until all C units have been awarded.

In a sealed-bid, first-price auction, the C highest bids are accepted, and each pays his bid; in a sealed-bid, second-price auction, the C highest bidders are awarded the item and each pays the $(C+1)^{\text{st}}$ highest bid.

Again, more complex multiunit auctions exist in practice. For example, customers may bid for multiple units. In a sealed-bid, first-price auction, this is accomplished by having customers submit a *demand schedule*—a list of quantities and prices they are willing to pay for each marginal unit they buy. The firm then awards items to the C highest marginal values, which may involve awarding multiple units to a sin-

[3] Though Lucking-Reily [351] points out that the Vickrey auction is more commonly used than most people realize.

gle customer. In this chapter, we only consider the simple, single-unit demand version of multiunit auctions.

6.1.2.3 Combinatorial Auctions

Another complexity in many procurement auctions is that a customer may require several products simultaneously. For example, to complete production of a product, a manufacturer may need both metal and plastic resin, or to provide cell phone service in a particular region, a communications company may need licenses in several contiguous regions. Such problems create dependencies, in which customers are willing to pay more for certain combinations of items than the sum of what they would be willing to pay for each item alone.

In such cases, one can construct auctions where the customers submit bids for various combinations of items rather than individual bids for each item alone, and the firm must then decide on which combinations to award based on these bids. Such problems may require solving complex, combinatorial optimization problems to simply determine the winners of the auction. Understanding the customers' behavior in the face of such complex auctions is quite difficult. We examine one such combinatorial auction in Section 6.5 below, in which customers bid for "products" that require a subset of "resources" and the firm has to allocate a finite supply of these resources to the customers based on their bids. This problem closely matches the network problems of Chapter 3.

6.2 Independent Private-Value Theory

In this section, we present the basic theory of auctions for the so-called *independent private-value* model, which is the most widely studied in the literature. In addition, we focus here on the revenue-generating properties of auctions and largely ignore welfare and allocative efficiency properties. Readers interested in these properties and other extensions of the basic theory are referred to survey papers by Klemperer [305], Matthews [366], McAfee [369], and Milgrom [381].

6.2.1 Independent Private-Value Model and Assumptions

Consider an auction in which we are selling one or more homogeneous objects to N potential customers. Each customer desires at most one of the objects. Customer i values an object at v_i. The valuations v_i are private information to the individual customers, but it is common knowledge that v_i's are i.i.d. with a distribution F. We assume that

F is strictly increasing with a continuous density function $f(\cdot)$ and has bounded support on the interval $[0, \overline{v}]$, so $F(0) = 0$ and $F(\overline{v}) = 1$.

Note that the assumption that customers have i.i.d. valuations and all know F is not equivalent to saying all customers are the same. Indeed, because customers valuations are draws from a distribution, some customers will have high valuations, and some will have low ones; F merely describes the distribution of valuations in the customer population. In addition, customers know their valuation; thus, a customer with a high (low) valuation will *know* that his valuation is higher (or lower) than average and will bid accordingly. The assumption of i.i.d. valuations and symmetry is more precisely a statement about the views the participants hold about the market. It is equivalent to saying that all customers and the firm have the *same belief* about the likely valuations of other customers and that there is no discernable difference among customers a *priori.*

6.2.2 An Informal Analysis of Sealed-Bid, First- and Second-Price Auctions

First, to build some initial intuition we start with a somewhat informal analysis of the sealed-bid, first- and second-price auctions. A formal equilibrium analysis is then provided in Section 6.2.3.

A key feature of auction models is that they assume customers are rational; that is, they bid so as to maximize their surplus (the value of the item minus the price they pay). Hence, for each mechanism we need to analyze customers' bids as a function of their valuations—called their *bidding strategy.* When formulating his bidding strategy, a rational customer will take into account the bidding strategies of the other customers. Our auction analysis therefore relies on the concept of an *equilibrium set of strategies;* that is, a set of strategies such that each customer has no incentive to change his strategy provided the other customers do not change their strategies (Nash equilibrium in game-theory terminology; see Appendix F).

We are also interested in the revenues produced by a given auction. These revenues depend on the strategic, equilibrium response of customers. So changes in the mechanism will lead to changes in the equilibrium bidding strategies of customers, which in turn will affect the revenues the firm generates. Thus, a "good" mechanism induces a more profitable equilibrium, and this makes the revenue analysis of auctions qualitatively different from the analyses we have seen in the previous chapters.

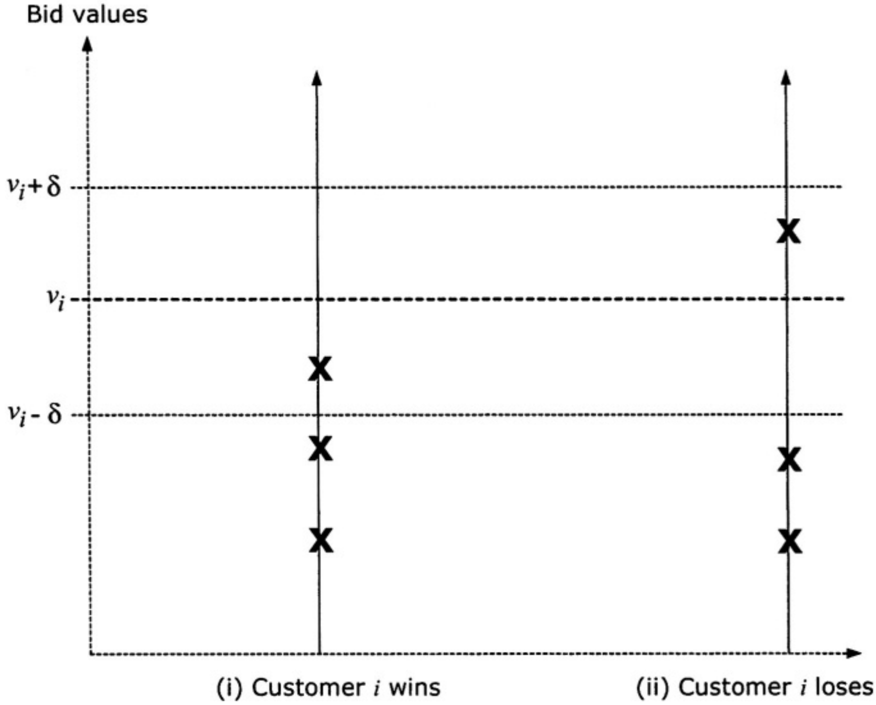

Figure 6.1. Perturbing the bid v_i in a second-price auction.

6.2.2.1 Equilibrium Strategies for a Second-Price Auction

Consider first a single-unit, sealed-bid, second-price auction with N customers. Recall that in this case each of the N customers submits a bid, the firm awards the item to the customer with the highest bid, and the winner pays the value of the second-highest bid. Let $\mathbf{b} = (b_1, \ldots, b_N)$ denote the vector of bids submitted by the N customers and let $b_{[i]}$ denote the i^{th} reverse-order statistic: that is, $b_{[1]} \geq b_{[2]} \cdots \geq b_{[N]}$. Hence, $b_{[2]}$ denotes the value of the second-highest bid (the winner's payment). A bidding strategy for customer i specifies the bid customer i will submit as a function of his valuation v_i and is denoted $b_i(v_i)$. A bidding strategy that is an equilibrium strategy is denoted $b_i^*(v_i)$ (to denote that it is an optimal response to the strategies of other customers).

How would a rational customer bid in this type of auction? The answer, it turns out, is surprisingly simple. Each customer i cannot do better than to simply bid his own valuation v_i; that is, the strategy of bidding $b_i^*(v_i) = v_i$ is optimal for all customers i.

To see this, note that the amount the winner pays in a second-price auction is not affected by his bid since he pays an amount equal to the second-highest bid. In other words, a customer's bid affects whether he wins but not how much he pays if he wins. Now suppose customer i bids v_i. Consider the two possible cases—customer i wins or customer i loses—and see whether customer i can do better by changing his bid v_i in either case. The situation is shown in Figure 6.1.

First, consider case (i) on the left of Figure 6.1, where customer i wins the auction by bidding v_i. In this case, customer i has a surplus of $v_i - b_{[2]} \geq 0$. Now if he increases his bid to $v_i + \delta$, it has no effect because he is still the highest bidder and still pays an amount equal to the second-highest bid. So customer i cannot do better by increasing his bid. If customer i is a winner and decreases his bid to $v_i - \delta$, there is again no change in his surplus as long as he remains the highest bidder. However, if—as shown on the left-hand side of Figure 6.1—he lowers his bid enough to become the second-highest bidder, then he is no longer the winner and his surplus is zero. Since his surplus was positive beforehand, this is not an improvement either. Thus, customer i cannot do better than bidding v_i in case (i).

Now consider case (ii) on the right of Figure 6.1, in which customer i bids v_i and loses. Customer i's surplus in this case is zero because he does not get the item and pays nothing. Note also in this case, the highest bid is strictly greater than v_i; that is, $b_{[1]} > v_i$. Now if he decreases his bid to $v_i - \delta$, he remains one of the losers, and his surplus is still zero. If he increases his bid to $v_i + \delta$, again there is no change unless he increases his bid enough to become the new highest bidder. But in this case, he must pay an amount equal to the previous high-bidder's bid, which is strictly greater than his own valuation v_i (else he would have been the high bidder originally). So his surplus in this case is negative, he is worse off. Hence, he cannot do better than bidding v_i in case (ii) either. Therefore, in both cases (i) and (ii), bidding v_i is an optimal decision for customer i.

Note that this strategy is optimal regardless of the bids placed by other customers. Indeed, our analysis did not make any assumptions on the strategies used by other customers; the strategy $b_i^*(v_i) = v_i$ is optimal for *any realization* of competing bids. Such a strategy is called a *dominant strategy,* and the set of such strategies $b_i^*(v_i) = v_i, i = 1, \ldots, N$ is called a *dominant-strategy equilibrium* for the auction. (See Appendix F.)

A dominant-strategy equilibrium is a robust equilibrium. It applies under very general conditions; essentially, we need assume only that customers have private valuations (the valuation v_i that customer i has for the item is not influenced by the valuations of other customers) and customers are rational so that they recognize the benefit of this strategy. We need little else beyond these two assumptions. For example, customers can have different distributions of valuations, have different information about the distributions, and may be risk-averse. None of these change the equilibrium under the second-price mechanism because of the strong dominance of the bidding strategy.

Under this equilibrium, the firm earns a revenue equal to the second reverse-order statistic of the distribution $F(v)$, a quantity that is not difficult to evaluate (at least numerically). The following example illustrates both the equilibrium and the revenue calculation.

Example 6.1 There are N customers with valuations uniformly distributed on [0,1], so $F(v) = v$ on this interval. Under the second-price auction, it is a dominant-strategy equilibrium for each customer to adopt the strategy $b^*(v_i) = v_i$.

The expected revenue earned by the firm is just $E[v_{[2]}]$—the value of the second highest bid. It is not hard to show for the uniform distribution that

$$E[v_{[2]}] = \frac{N-1}{N+1}.$$

Thus, the firm's average revenue increases with the number of bidders N.

6.2.2.2 Equilibrium Strategies for a First-Price Auction

Consider next a first-price auction, in which the highest bidder wins the item and pays his bid. We consider only symmetric bidding strategies in this case. That is, we assume each customer i uses the same bidding strategy $b(v)$ and therefore bids an amount $b_i = b(v_i)$. This is a reasonable assumption given that customer valuations are symmetric (have valuations independently drawn from the same distribution $F(v)$). We also assume that a customer's bid is increasing in his valuation (customers with higher valuations bid more), so the bid strategy $b(v)$ is increasing in v. This assumption is verified afterward. As before, an equilibrium strategy is denoted $b^*(v)$.

Again, we are looking for an equilibrium bid function such that if all customers are using the strategy $b^*(v)$, then no customer is able to improve his expected surplus by bidding anything other than $b^*(v)$. We then use the first-order conditions for this equilibrium to derive a differential equation for the bid function $b^*(v)$.

To begin, note that a customer i with valuation v_i will win the item if he is the highest bidder; that is, if $b^*(v_i) > b^*(v_j)$, for all $j \neq i$. So customer i's probability of winning is[4]

$$\begin{aligned} P(b^*(v_i)) &= \Pi_{j \neq i} P(b^*(v_i) > b^*(v_j)) \\ &= \Pi_{j \neq i} P(v_i > v_j) \end{aligned}$$

---

[4]Equation (6.1) contains a minor abuse of notation: the $P(\cdot)$ on the left-hand side represents a function of the bids (albeit a probability), while the right-hand side $P(\cdot)$ stands, as throughout the book, for probability of an event.

$$= F^{N-1}(v_i), \tag{6.1}$$

where the second inequality follows from the assumption that strategy $b^*(v)$ is strictly increasing in v. Since the argument is generic to any customer i, we henceforth drop the subscript and consider an arbitrary customer with valuation v.

Now suppose our customer could improve his expected surplus by adopting the strategy of a customer with valuation $\tilde{v}$ different from v. Specifically, the customer would bid $b^*(\tilde{v})$ and thus win with probability $P(\tilde{v})$ but would still value the item at v. In this case, his expected surplus would be

$$S(b^*(\tilde{v}), v) = P(\tilde{v})(v - b^*(\tilde{v})). \tag{6.2}$$

If the strategy $b^*(v)$ is truly an equilibrium, this surplus should be maximized at $\tilde{v} = v$ (otherwise, $b^*(v)$ would not be the customer's optimal bid). Therefore, applying the first-order optimality conditions, we can differentiate (6.2) with respect to $\tilde{v}$, set the result to zero at $\tilde{v} = v$, and obtain the following differential equation for $b^*(v)$:

$$b^{*\prime}(v) = \frac{P'(v)}{P(v)}(v - b^*(v)). \tag{6.3}$$

The solution to this differential equation is somewhat tedious to derive, but one can verify that it is[5]

$$b^*(v) = v - \frac{\int_0^v P(s)ds}{P(v)}, \tag{6.4}$$

where $P(v) = F^{N-1}(v)$ is the probability of winning given by (6.1).[6]

Note that the equilibrium strategy (6.4) is continuous and increasing in v and increasing in N (higher-valuation customers bid more; and the more customers there are, the higher a given customer bids). One can also show that it is the unique symmetric equilibrium for this problem (Riley and Samuelson [441]).

Note from (6.4) that $b^*(v) < v$, so customers in a first-price auction will bid strictly less than their valuation. Hence, unlike in the second-price auction, they *shade* their true valuations when bidding. This is to be expected because customers are required to pay what they bid, so they must shade their bids to make a positive surplus from winning.

---

[5] To check this, just use the fact that $b^{*\prime}(v) = [\int_0^v P(s)ds][P(v)]^{-2}P'(v)$, and substitute (6.4) into the right-hand side of (6.3).

[6] A boundary condition of $b^*(0) = 0$ is required as well; see Appendix 6.A.

Finally, the revenue to the firm is the expected value of the highest bidder's bid because the winner pays his valuation. So the firm's expected revenue is $E[b^*(v_{[1]})]$. Again, the mean of this order statistic is not difficult to compute numerically or by simulation.

To illustrate, consider again the example of uniformly distributed valuations:

Example 6.2 There are N customers with valuations uniformly distributed on [0,1], so $F(v) = v$ on this interval. In this case, $P(v) = v^{N-1}$ and the equilibrium bidding strategy is

$$\begin{aligned} b^*(v) &= v - \frac{\int_0^v P(s)ds}{P(v)} \\ &= v - \frac{\int_0^v s^{N-1}ds}{v^{N-1}} \\ &= v(1 - \frac{1}{N}). \end{aligned}$$

So each customer bids a fraction $1 - \frac{1}{N}$ of his valuation; hence, customers with higher valuations bid more, and the more customers N, the closer each bids to his actual valuation.

Since the highest bidder $v_{[1]}$ wins, the expected revenue to the firm is then $E[v_{[1]}](1 - \frac{1}{N})$. It is not hard to show for the uniform distribution that $E[v_{[1]}] = \frac{N}{N+1}$. Therefore, the firm's expected revenue is

$$E[v_{[1]}]\left(1 - \frac{1}{N}\right) = \frac{N}{N+1}\left(1 - \frac{1}{N}\right) = \frac{N-1}{N+1}.$$

Note that the expected revenue for the firm is the same in this example and in Example 6.2 for the second-price auction. In other words, the firm generates the same expected revenue regardless of which auction it runs. As we show below, this is not a coincidence; rather, it is a consequence of general conditions that guarantee that these two auctions are always revenue equivalent under the private-value model.

6.2.2.3 Strategic Equivalence of Open and Sealed-Bid Auctions

In the private-value model, the open descending (Dutch) auction is strategically equivalent to the sealed-bid, first-price auction in the sense that the equilibrium strategies for the two mechanisms are the same. That is, if $b^*(\cdot)$ is a symmetric equilibrium in a sealed-bid, first-price auction, then it is also a symmetric equilibrium in a open descending auction, and vice versa. This is true because in an open descending auction, each customer (knowing his valuation v) calculates his expected surplus at each price b, given that there are no other customers willing to buy at b. He then determines the value $b^*(v)$ at which this surplus is maximized and bids when the price drops to $b^*(v)$. But this is exactly the same calculation the customer must make when submitting a bid in a sealed-bid, first-price auction. Hence, the equilibrium strategies are the same.

Likewise, an open ascending auction can be shown to be strategically equivalent to a sealed-bid, second-price (Vickrey) auction under the independent private-value model. In an open ascending auction, it is always optimal for a customer to stay in the bidding as long as the announced price b is below his valuation v—and to drop out once the price exceeds v. But this is equivalent to the strategy of bidding $b^*(v) = v$ in a second-price auction, since in both cases if the customer wins, he ends up paying the valuation of the second-highest customer. And as we showed in Section 6.2.2.1, $b^*(v) = v$ is a dominant-strategy equilibrium in the Vickrey auction. Hence, the two auctions are strategically equivalent.

Because of this equivalence, we henceforth refer to these two cases as simply the *first-price* and *second-price* auctions—without specifying whether the mechanism is the open- or sealed-bid version.

6.2.3 Formal Game-Theoretic Analysis

We now formalize and generalize the analysis of bidding equilibria for a general auction mechanism. Formally, a *bidding strategy* for customer i is a function $b_i(v_i)$ that specifies the bid that customer i will submit conditional on his valuation v_i. We let $\mathbf{v} = (v_1, \ldots, v_N)$ denote the vector of valuations and $\mathbf{b}(\mathbf{v}) = (b_1(v_1), \ldots, b_N(v_N))$ denote the vector of bidding strategies used by the N customers. We let $\mathbf{v}_{-i} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_N)$; that is, the vector $\mathbf{v}$ without the i^{th} component. Similarly, let

$$\mathbf{b}_{-i}(\mathbf{v}_{-i}) = (b_1(v_1), \ldots, b_{i-1}(v_{i-1}), b_{i+1}(v_{i+1}), \ldots, b_N(v_N))$$

denote the bid strategies for all customers other than i.

An *auction mechanism* is specified by a pair of mappings $\tilde{\mathbf{y}} : \Re_+^N \to \{0,1\}^N$ that defines the allocations of the goods and $\tilde{\mathbf{p}} : \Re_+^N \to \Re_+^N$ that defines payments made by the customers (equivalently, revenue received by the firm) as a function of their bids. The firm chooses the auction mechanism before the auction is conducted and announces it to all customers, so the mechanism too is common knowledge.

Suppose customer i chooses a strategy $b_i(v_i)$. Then $\tilde{y}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))$ is the allocation of goods to customer i, which is equal to 1 if he is awarded a unit, and 0 otherwise. Given his bid $b_i(v_i)$, the probability that customer i is awarded a unit is given by

$$P_i(b_i(v_i)) = E_{\mathbf{v}_{-i}}[\tilde{y}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))].$$

Similarly, $\tilde{p}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))$ is the payment made by customer i given the bid vector $\mathbf{b}(\mathbf{v})$, and his expected payment is

$$R_i(b_i(v_i)) = E_{\mathbf{v}_{-i}}[\tilde{p}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))].$$

Note the expected payment is the expected revenue received by the firm. When the number of players N is random, each player computes his optimal action by conditioning both on the valuations of the other players and the total number of players in the game.

Customers are assumed to be rational and attempt to maximize their expected net utility (the value of the item less the price paid to the firm). Therefore, customer i chooses his strategy $b_i(v_i)$ to maximize his expected surplus

$$S_i(b_i(v_i), v_i) = v_i P_i(b_i(v_i)) - R_i(b_i(v_i)). \tag{6.5}$$

For example, in the case of the single-unit, first-price auction, the item is awarded to the highest bidder who pays the auctioneer the value of his bid; all other bidders pay nothing. Then if $b_i = b_{[1]}$ (i is the highest bidder and wins the item), then $\tilde{y}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})) = 1$, and $\tilde{p}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})) = b_i(v_i)$, and if $b_i < b_{[1]}$ (i is not the winning bidder), $\tilde{y}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})) = 0$, and $\tilde{p}_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})) = 0$. So the expected net utility is simply $P_i(b_i(v_i))(v_i - b_i(v_i))$.

We assume that customers choose their strategies without collusion. In this case, they play a noncooperative game of incomplete information. An appropriate solution concept in this context is that of the Bayesian equilibrium of Harsanyi [241], an extension of the ordinary Nash equilibrium [402]. Specifically, a vector of strategies $(b_1^*(\cdot), \ldots, b_N^*(\cdot))$ is an *equilibrium strategy* if, for all i, customer i's best response is to maintain his strategy $b_i^*(\cdot)$ provided all other customers maintain their strategies $\mathbf{b}_{-i}^*(\cdot)$. Formally,

$$S_i(b_i^*(v_i), v_i) \geq S_i(b, v_i) \quad \forall b \in [0, \bar{v}], \ \forall i = 1, \ldots, N.$$

In other words, no customer has an incentive to change his strategy if all other customers maintain their strategies. We further restrict ourselves and consider only *symmetric* equilibria; that is, strategies for which the equilibrium strategy $b_i^*(\cdot) = b^*(\cdot)$ is the same for all i. As mentioned, this assumption is reasonable given that customer valuations are assumed symmetric; however, it is a restriction nevertheless, and one cannot rule out the fact that asymmetric bidding strategies may exist. Henceforth, we let $b^*(\cdot)$ denote such a symmetric equilibrium strategy.

6.2.3.1 Direct-Revelation Mechanisms

The analysis of equilibrium bidding strategies is greatly simplified by considering what are called *direct-revelation mechanisms.* Essentially, a direct-revelation mechanism is one in which a customer's equilibrium strategy is to bid his true valuation v. For any mechanism that has an equilibrium it turns out, we can always find an equivalent direct-revelation mechanism.

To see this, note that if $b^*(v)$ is a symmetric equilibrium for some given auction mechanism, then the firm can always define an alternative mechanism (the direct-revelation mechanism) in which customers submit bids, the firm inserts these bids into the function $b^*(\cdot)$, and the resulting values are treated as bids under the rules of the original auction mechanism. The situation is illustrated in Figure 6.2. Since $b^*(\cdot)$ is an equilibrium strategy, it follows that under the direct-revelation mechanism it is an optimal strategy for every customer i to bid his valuation v_i, since otherwise it would contradict the fact that $b^*(\cdot)$ is an equilibrium strategy. Conversely, if there does not exist a direct-revelation mechanism defined by some $b^*(\cdot)$ in which bidding v is an equilibrium, then there cannot be any equilibrium bidding strategy under the original mechanism, otherwise the corresponding equilibrium $b^*(\cdot)$ would define such a direct-revelation mechanism.

In this way, we can reduce the equilibrium analysis of any mechanism to an analysis of the corresponding direct-revelation mechanism, in which case we can view the allocation and payments as being directly a function of the customers' valuations—denoted $y_i(v_i, \mathbf{v}_{-i})$ and $p_i(v_i, \mathbf{v}_{-i})$, respectively (because the optimal strategy is for customers to bid their valuations). This approach is illustrated in Figure 6.2.

Let $R_i(v)$ denote customer i's expected payment ($R_i(v_i) = E_{\mathbf{v}_{-i}}[p_i(v_i, \mathbf{v}_{-i})]$) under a direct-revelation mechanism. The equilibrium can be analyzed by noting that the expected surplus in the direct-revelation mechanism, defined by

$$S_i(v_i) = v_i P_i(v_i) - R_i(v_i),$$

must satisfy

$$S_i(v_i) \geq S_i(\tilde{v}) + P_i(\tilde{v})(v_i - \tilde{v}) \quad \forall \tilde{v} \in [0, \overline{v}]$$

for all customers i. In other words, for each customer i, revealing his true valuation v_i is no worse than pretending to have another valuation $\tilde{v}$. This condition is called the *incentive compatibility constraint* because it requires that it be in customer i's self-interest to truthfully reveal his valuation.

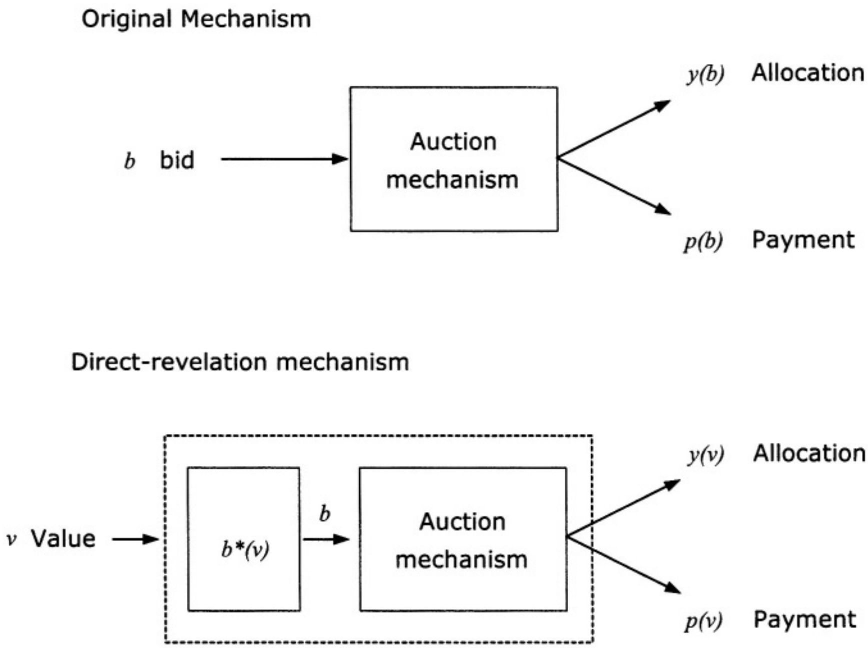

Figure 6.2. Illustration of the direct-revelation mechanism.

6.2.4 Revenue Equivalence

How much revenue is generated for the firm by a given auction mechanism? At first, answering this question would appear to be a hopeless task because each auction mechanism leads to different equilibrium bidding strategies and equilibrium payments. Finding and evaluating these various equilibria for a reasonable range of mechanisms (or ideally all possible mechanisms) is a daunting task. However, it turns out that the expected revenue generated by a private-value auction can be reduced to an analysis only of the resulting allocations $y_i(v_i, \mathbf{v}_{-i})$—without explicitly solving for the equilibrium bidding strategies. The only conditions required are that the functions $y_i(v_i, \mathbf{v}_{-i})$ are increasing in v_i (so that higher valuations lead to a higher probability of allocation)[7] and customers with valuation zero have zero expected surplus in equilibrium.[8] Specifically, we have:

[7]Verifying that the allocations are increasing in the valuations v_i may require analyzing monotonicity properties of the equilibrium strategy—that higher-valuation customers v bid higher in equilibrium.

[8]The requirement that customers with valuations zero have zero expected surplus is called a *participation constraint;* intuitively, it means we cannot force a customer to participate in an auction in which his expected surplus is negative.

THEOREM 6.1 *[Revenue Equivalence Theorem] Consider the private-value model, in which there are C items and N customers with i.i.d. valuation independently drawn from a continuously differentiable, strictly increasing distribution F on* $[0, \overline{v}]$. *Consider any mechanism in which (i) the allocation to customer* i, $y_i(v_i, \mathbf{v}_{-i})$, *is increasing in* v_i *for all* i, *and (ii) customers with valuations of zero have zero expected surplus. Then the expected revenue for the firm is given by*

$$E\left[\sum_{i=1}^{N} J(v_i) y_i(v_i, \mathbf{v}_{-i})\right], \tag{6.6}$$

where

$$J(v) = v - \frac{1 - F(v)}{f(v)}.$$

A proof of the this theorem is given in Appendix 6.A. Moreover, this revenue equivalence holds for the customers as well; a customer's expected payment is the same under all mechanisms satisfying the above conditions. However, the equivalence in both cases is only in expectation; the payments on a sample-path basis may be quite different under different mechanisms.

Note $J(v)$ is precisely the marginal revenue function (7.14) encountered in our analysis of revenue-maximizing prices in Chapter 7. This is not a coincidence; Bulow and Roberts [95] show the revenue function (6.6) can in fact be interpreted as a variant of third-degree monopoly price discrimination among the N customers (see Section 8.3.3.2).

In auction theory, $J(v_i)$ is sometimes referred to as customer i's *virtual value* because (6.6) implies that the firm can hope to collect only $J(v_i) < v_i$ from customer i (in expectation) and not his entire valuation v_i. The difference $v_i - J(v_i)$ is referred to as the *information rent* of customer i because it is the surplus that customer i retains due to his private information about his own valuation v_i.

As a result of Theorem 6.1, note that any two mechanisms that produce the same allocation for every realization of $v_1, \ldots, v_N$ (the same customers are awarded units under each mechanism) produce the same expected revenue for the firm. This is true despite the fact that the bidding strategies and payments may be very different under each mechanism. For this reason, Theorem 6.1 is referred to as the *revenue equivalence theorem.*

To illustrate, consider a standard first-price auction on a single unit. We saw in Section 6.2.2.2 that the equilibrium bidding strategy was strictly increasing in the value v and that the item is awarded to the highest bidder. In a second-price auction, the customer with the highest valuation also wins that auction. Thus, by Theorem 6.1, the expected revenue to the firm must be the same in each case. We illustrate this result with a continuation of our previous example:

Example 6.3 Suppose there are $N = 2$ customers with valuations uniformly distributed on [0,1]. Let $v_{\max} = \max\{v_1, v_2\}$ and $v_{\min} = \min\{v_1, v_2\}$. From Example 6.2 we know that in a first-price auction, the customers will bid $v(1 - \frac{1}{N}) = v/2$ in equilibrium. The highest bidder will win, and the firm's expected revenue is

$$\frac{1}{2}E[v_{\max}] = (1/2)(2/3) = 1/3.$$

Now consider a second-price auction of Example 6.1. The highest bidder wins but pays the price of the second-highest bid, and each customer bids his valuation in equilibrium. The firm's expected revenue is

$$E[v_{\min}] = 1/3.$$

Hence, the two expected revenues are equal. Moreover, note that since $F(v) = v$, $J(v) = 2v - 1$, and therefore since in both auctions $y_i(v_1, v_2) = 1$ if and only if $v_i = v_{\max}$, we have

$$\begin{aligned} E[J(v_1)y_1(v_1, v_2) + J(v_2)y_2(v_1, v_2)] &= E[2v_{\max} - 1] \\ &= 2(2/3) - 1 \\ &= 1/3 \end{aligned}$$

as well. Finally, for $N > 2$ customers, the customer with valuation $v_{\max}$ wins, and the same analysis shows that

$$\begin{aligned} E\left[\sum_{i=1}^{N} J(v_i)y_i(v)\right] &= E[2\max\{v_1, \ldots, v_N\} - 1] \\ &= 2\left(\frac{N}{N+1}\right) - 1 \\ &= \frac{N-1}{N+1}, \end{aligned}$$

which is precisely the expected revenue found in Examples 6.1 and 6.2.

Similarly, in a standard C-unit auction, one can show that both the first-price and second-price auctions award the goods to the customers with the C highest valuations. Thus, the allocation $y(v)$ is the same for each v, and hence the two mechanisms generate the same expected revenue for the firm.

6.2.5 Optimal Auction Design

The revenue expression (6.6) can be used to design an optimal mechanism by simply choosing the allocation rule $\mathbf{y}^*(\mathbf{v})$ that maximizes

$$\sum_{i=1}^{N} J(v_i)y_i(v_i, \mathbf{v}_{-i}) \tag{6.7}$$

subject to any constraints one might have on the allocation.

Toward this end, it is useful to make the same regularity Assumption 7.2 on the distribution function F that we impose in dynamic pricing problems: namely, that $J(v)$ is strictly increasing in v. Note that

$$J(v) = v - \frac{1}{\rho(v)},$$

where $\rho(v)$ is the hazard rate of the distribution F. The marginal revenue $J(v)$ satisfies this monotone condition as long as the hazard rate $\rho(v)$ is increasing—or not decreasing too quickly with v.[9]

To illustrate, consider designing an optimal C-unit auction using (6.6). Note that with C units to allocate and given a realization of $\mathbf{v}$, we want to maximize (6.7) subject to the constraint that

$$\sum_{i=1}^{N} y_i(v_i, \mathbf{v}_{-i}) \leq C,$$

and $y_i(\cdot) \in \{0, 1\}$ for all i. It is easy to see what the optimal allocation is by inspection. Indeed, define

$$v^* = \max\{v : \; J(v) = 0\} \tag{6.8}$$

(and by convention, $v^* = \infty$ if $J(v) < 0, \; \forall v$). Then since $J(\,.\,)$ is assumed to be increasing, it follows that it is never optimal to allocate a unit to a customer with valuation $v_i < v^*$ because awarding units to such customers results in a negative contribution to the sum (6.7). Among the remaining customers with $v_i \geq v^*$, we want to award units to those with the highest valuations v_i. Thus, the optimal allocation is to award units to the C highest-valuation customers v_i above v^*, and if there are less than C customers with $v_i \geq v^*$, to award units only to those customers and discard the remaining units.

How can we achieve such an allocation? One possibility is to introduce a *reserve price* into the standard first- or second-price C-unit auction mechanism. A reserve price is a lower bound on bids that the firm sets before the auction; only bids above the announced reserve price are considered.

To illustrate, consider a first-price auction; it is easy to see that if we set a reserve price of v^*, customers with valuations $v_i < v^*$ will not submit bids. One can show that the remaining customers with valuations

[9] More precisely, it is satisfied when $\rho'(v)/\rho^2(v) > -1$ for all $v \geq 0$. One can show that this condition is satisfied by many standard distributions, including the uniform, normal, logistic, exponential, and extreme value (double exponential) distributions [25].

$v_i \geq v^*$ will submit bids according to an increasing equilibrium strategy similar to (6.4). Indeed, the resulting symmetric equilibrium strategy is now

$$b^*(v) = v - \frac{\int_{v^*}^{v} F^{N-1}(s)ds}{F^{N-1}(v)}. \tag{6.9}$$

The C units are awarded to the C highest bids above the reserve price, and the resulting allocation is exactly the same as the optimal allocation. Hence, the first-price mechanism with reserve price v^* is optimal. A similar argument holds for the second-price auction, in which case one can show that it is optimal to post a reserve price of v^*, where winners pay the minimum of the $(C+1)^{\text{st}}$ highest bid above v^*, or v^* if there are fewer than $C + 1$ bids above v^*. We therefore have the following theorem:

THEOREM 6.2 *Under the private-value model, the standard C-unit first-price and second-price auctions with reserve price v^* (given by (6.8)) are optimal for the firm.*

Hence, with a properly chosen reserve price, the standard first-price and second-price auctions are revenue maximizing among all possible pricing mechanisms. This is a rather remarkable result; under the private-value model assumptions, a firm simply cannot do better than to sell using one of these two auction formats. Again, we illustrate this result by returning to our uniform-distribution example:

Example 6.4 Suppose there are N customers with valuations uniformly distributed on [0, 1]. Since $F(v) = v$ implies $J(v) = 2v - 1$, the optimal reserve price $v^* = 1/2$ since this satisfies $J(v^*) = 0$. From (6.9), the customers with valuations over 1/2 therefore bid

$$\begin{aligned} b^*(v) &= v - \frac{\int_{1/2}^{v} s^{N-1} ds}{v^{N-1}} \\ &= v(1 - \frac{1}{N}) + \frac{(1/2)^N}{N v^{N-1}}, \end{aligned}$$

which is strictly greater than the bid of $v(1-1/N)$ submitted by these same customers in the first-price auction without reserve prices. Also, since the item is allocated to the highest-value customer and the distribution of the highest valuation is $F^N(v) = v^N$, the firm's optimal expected revenue is

$$\begin{aligned} E[J(\max\{v_1, \ldots, v_N\})] &= \int_{1/2}^{1} (2v - 1) N v^{N-1} dv \\ &= \frac{N - 1 + (1/2)^N}{N+1}, \end{aligned}$$

which is again larger than the expected revenue of $(N - 1)/(N + 1)$ generated when no reserve prices are used but approaches the no-reserve-price revenue when N is large.

The fact that a reserve price primarily benefits the firm most when there are few customers is intuitive. In essence, the reserve price serves to create "extra competition" for customers—forcing them to bid higher in a first-price auction or pay more if they win in a second-price auction—than they otherwise would without a reserve price. However, with lots of competition from other customers, the need for the firm to introduce this extra incentive is less important, as the customers themselves create sufficient competition.

6.2.6 Relationship to List Pricing

How is the optimal auction mechanism related to a traditional list-price mechanism? There are several close connections worth examining.

First, note that a list-price mechanism qualifies as one of the possible allocation and payment mechanisms studied above for the C-unit auction. In particular, using a fixed list price p, customers indicate their willingness to pay p (just as in an ascending auction). If there are C or fewer customers willing to pay p, each receives a unit and pays the fixed amount p; if there are more than C customers willing to pay p, the C units are randomly rationed to these customers. This produces an allocation and payment rule just as in the standard auction types. In the list price case, it is easy to see that it is a dominant strategy for a customer to "bid" (indicate his willingness to pay p) if his valuation v is more than p. Thus, a dominant, symmetric equilibrium strategy exists in which all customers with valuations greater than p attempt to buy.

Given this observation, it is easy to compute the firm's expected revenue. Let $N(v)$ denote the number of customers with valuations greater than v. Then the expected revenue to the firm as a function of p is

$$r(p) = pE[\min\{N(p), C\}]. \tag{6.10}$$

Another way of deriving this revenue is to use the expression (6.6) for the firm's equilibrium revenue and to note that for the list-price mechanism

$$E\left[\sum_{i=1}^{N} J(v_i)y_i(v_i, v_{-i})\right] = E[\min\{N(p), C\}]E[J(v)|v > p],$$

since $E[\min\{N(p), C\}]$ units are allocated, and each customer i to which we allocate a unit ($y_i = 1$) has a valuation $v_i > p$, so $E[J(v)|v > p]$ is the expected value of the corresponding term $J(v_i)$. Using the fact that $J(v) = v - (1 - F(v))/f(v)$, it is then a simple exercise to show that $E[J(v)|v > p] = p$, which gives us the same expression as (6.10).

A direct optimization of (6.10) does not lead to a clean expression, but several special cases are simple and provide useful insight. We look at these next.

6.2.6.1 Capacity Is Unconstrained

The first case is where the number of customers $N \leq C$, so there are always fewer customers than there are units. In this case, $E[\min\{N(p), C\}] = E[N(p)] = N(1 - F(p))$ and the revenue for a list price of p is $r(p) = Np(1 - F(p))$. Differentiating and setting the result equal to zero, we find that the optimal price p^* satisfies

$$p^* f(p^*) - (1 - F(p^*)) = 0.$$

But since $f(p) > 0$ (F is strictly increasing), rearranging this is equivalent to $J(p^*) = 0$, which is the condition for the revenue-maximizing price and is also the condition for determining the optimal reserve price. Thus, the optimal price when $N \leq C$ is the same as the optimal reserve price—that is, $p^* = v^*$. Moreover, the revenue under this optimal price can be written

$$\begin{aligned} E[N(v^*)]E[J(v)|v > v^*] &= E\left[\sum_{i=1}^{N} J(v_i)\mathbf{1}(v_i > v^*)\right] \\ &= E\left[\max_{y_i \in \{0,1\}} \{\sum_{i=1}^{N} J(v_i) y_i\}\right], \end{aligned}$$

where $\mathbf{1}(v_i > v^*)$ is the indicator function of the event $v_i > v^*$. But the expression $E\left[\max_{y_i \in \{0,1\}} \{\sum_{i=1}^{N} J(v_i) y_i\}\right]$ above is simply the optimal auction revenue when $J(v)$ is strictly increasing and $N \leq C$. Thus we have

PROPOSITION 6.1 *If capacity is not constrained in a C-unit, private-value auction ($N \leq C$), then using a fixed list price p^* satisfying $J(p^*) = 0$ is an optimal mechanism for the firm.*

6.2.6.2 Large Capacity and Sales Volumes

Another case in which list pricing is provably good is when both the number of customers and the number of units for sale is large. Specifically, let θ be a positive integer, and consider a problem with θC units and θN customers for some $N > C > 0$. If we set a fixed price of p, then the number of customers willing to purchase at this price is denoted $N_\theta(p)$ with mean $N\theta(1 - F(p))$. Moreover, by the law of large numbers, as $\theta \to \infty$,

$$\frac{N_\theta(p)}{\theta} \to N(1 - F(p)) \quad (a.s.),$$

and the firm's revenue satisfies

$$p\frac{\min\{N_\theta(p), \theta C\}}{\theta} = p \min\{\frac{N_\theta(p)}{\theta}, C\} \tag{6.11}$$

$$\to \quad p \min\{N(1-F(p)), C\} \quad (a.s.).$$

Just as in the capacity-constrained pricing problem of Section 5.2.1.2, the asymptotically optimal price is given by

$$p^* = \max\{p^0, \bar{p}\},$$

where p^0 is the revenue-maximizing price, determined by $J(p^0) = 0$, and $\bar{p}$ is the run-out price, determined by equating the expected number of customers willing to pay $\bar{p}$ to the supply C, so $N(1 - F(\bar{p})) = C$. When $p^* = \bar{p}$, the expected revenue is $C\bar{p}$, and when $p^* = p^0$, the expected revenue is $Np^0(1 - F(p^0))$.

Similarly, one can analyze the scaled optimal auction revenue. Note that the scaled expected optimal auction revenue can be written

$$\frac{1}{\theta} E\left[\sum_{i=1}^{\min\{\theta C, N_\theta(v^*)\}} J(v_{[i]})\right],$$

where $v_{[i]}$ denotes the i^{th} largest valuations $(v_{[1]} \geq v_{[2]} \geq \cdots \geq v_{[N]})$.

First, consider the case $N(1 - F(v^*)) > C$. Then as $\theta \to \infty$, with probability one $N_\theta(v^*) > \theta C$. So the above becomes

$$\begin{aligned} \frac{1}{\theta} E\left[\sum_{i=1}^{\theta C} J(v_{[i]})\right] &\to N \int_{F^{-1}(1-\frac{C}{N})}^{\bar{v}} J(v) f(v) dv \\ &= N\left[\int_{\bar{p}}^{\bar{v}} v f(v) dv - \int_{\bar{p}}^{\bar{v}} (1 - F(v)) dv\right] \\ &= N\bar{p}(1 - F(\bar{p})) \\ &= C\bar{p}, \end{aligned}$$

which is exactly the asymptotic fixed-price revenue given by (6.11) when $p = \bar{p}$.

In the alternative case where $N(1 - F(v^*)) < C$, as $\theta \to \infty$, with probability one $N_\theta(v^*) < \theta C$, and similar reasoning shows that the

$$\frac{1}{\theta} E[\sum_{i=1}^{\theta C} J(v_{[i]})] \to Np^0(1 - F(p^0)),$$

which is again exactly the asymptotic fixed-price revenue given by (6.11) when $p = p^0$. These arguments can be formalized to show

PROPOSITION 6.2 *If the number of customers and the number of units for sale in the private-value auction model are, respectively, θN and*

θC for some integers θ, and $N > C > 0$. Then as $\theta \to \infty$, a list-price mechanism is asymptotically optimal, in the sense that the ratio of the optimal expected list price revenue to the optimal expected auction revenue tends to one.

As a result, in high-sales-volume settings, using a fixed price will be near optimal. This implies auction benefits are something of a "small-numbers" phenomenon, which is consistent with the auctions one encounters in practice.

6.2.6.3 Dynamic Pricing

Another close connection between auctions and list-price mechanisms is obtained by considering a dynamic pricing policy as a particular allocation and payment mechanism in a private-value auction model. Making this connection yields several important insights.

For example, consider the problem of selling a single unit to a population of N strategic consumers. As in the price-skimming model of Section 5.5.2, the private-value model considers the N customers to have i.i.d. private valuations v_i for the item. The Dutch-auction mechanism calls for the firm to continuously reduce the price over time until a customer decides to bid at the offered price. The customer then pays this offered price. However, this is precisely what happens in a (continuous-time) dynamic pricing policy as well, so a descending dynamic price can effectively achieve the Dutch-auction outcome. By simply adding an optimal reserve price v^*, below which we will not lower the price, such a dynamic pricing mechanism becomes optimal.

More generally, by the revenue equivalence theorem, *any* dynamic pricing policy that results in the C highest-valuation customers, with valuations in excess of v^*, receiving the units, will be revenue-maximizing for the firm and thus produce the same expected revenue as the optimal auction.

For example, Bulow and Klemperer [94] analyze the C-unit, private-value model under a dynamic pricing mechanism. In their mechanism, the firm uses a list price that is lowered continuously until one or more customers offers to buy at the current price. If the number of customers willing to buy at the current price is less than the remaining supply, these customers are awarded the items at this price, and the firm continues to lower the price. If the number of customers willing to pay the current price exceeds the remaining supply, the firm does not sell the items; it instead increases the price discontinuously and then tries again to lower the price. Since a customer's probability of getting an item is higher when he attempts to buy early (if he attempts to buy and fails, he can always try again later, so his probability of obtaining the item cannot decrease by attempting to buy early), it is not hard to show that customers with the highest valuations are the ones that attempt to buy first. Therefore, the firm allocates the items first to the customers with the highest valuation. As a result, Bulow and Klemperer [94] argue that by revenue equivalence, if the firm does not lower the price below v^*, this dynamic pricing mechanism is optimal. Similar arguments hold for many other dynamic pricing policies as well.

This shows there is a rather close connection between optimal auction theory and dynamic pricing theory with strategic consumers. Indeed, using the revenue equivalence theorem, the seemingly difficult task of analyzing the customer equilibrium produced by a dynamic pricing strategy is greatly simplified, and it shows in fact that a range of pricing mechanisms are optimal.

6.2.7 Departures from the Independent Private-Value Model

Many of our conclusions thus far depend to a greater or lesser extent on the assumptions of the independent, private-value model. What happens when these assumptions are relaxed? In this section, we look briefly at a few cases that are especially relevant for RM. Each has implications for the types of auctions that are optimal for the firm.

6.2.7.1 The Common-Value Model

The private-value model assumes that each customer's valuation is independent of the valuations of other customers. Thus, if a customer learns the value that another customer places on the item, it has no impact on his valuation. Such an assumption is reasonable if the item is going to be used for personal enjoyment or consumption. However, in other cases the item may have a common commercial value, may be resold at some future point in time, or may be of uncertain quality, so the valuations others have on the item could reveal useful information about the value of the item to a given customer.

A canonical example of such a setting is selling an offshore oil lease. The value of the lease to a customer is dependent on two key factors: the volume of oil it contains and the cost of extracting that oil. Typically, there is a high degree of uncertainty about both these factors. Because of differences in survey data or technological expertise, different customers may have independent information on the value of a given lease, and so on. As a result, knowing how another customer values the lease may change your assessment of its profitability.

A simple model of such a setting is the following: consider auctioning a single item that has a *common value* a, which is the same for all customers. However, the value a is uncertain. All customers have the same prior knowledge of a, embodied as a distribution over values of a. This distribution is common knowledge. A value of a is drawn from this distribution, and then each customer i receives a (noisy) signal t_i of the form

$$t_i = a + \xi_i,$$

where $\xi_i, i = 1, \ldots, N$ are i.i.d. random-noise terms with mean zero. The distribution of ξ_i is also common knowledge.

Note that given only the signal t_i, customer i's expected value for the item is

$$E[a|t_i] = E[t_i - \xi_i|t_i] = t_i,$$

with variance given by the variance of ξ_i. However, if one were to aggregate the customers' signals by averaging them, the estimate would be $a + \frac{1}{N}\sum_{i=1}^{N}\xi_i$, which provides a much better (lower variance) information on the value a than do the signals t_i alone.[10] More generally, customer i's estimate of a may be altered by information he receives about the signals of other customers. This sort of behavior significantly affects the auction outcomes.

For example, one phenomenon that arises in this setting is the so-called *winner's curse.* To illustrate the idea, consider a sealed-bid, second-price auction. Suppose customer i were to bid his expected valuation t_i for the item, as in the private-value case. The customer might (incorrectly) reason that bidding his own expected valuation t_i is a dominant strategy because bidding more than t_i increases his chance of winning only in cases where his expected surplus is negative, and bidding less than t_i decreases his chances of winning only in cases where his expected surplus is positive. The reasoning is false, however, because customer i's expected valuation conditional on winning the auction is less than his unconditional expected valuation t_i. Indeed,

$$E[a|t_i = \max\{t_1, \ldots, t_N\}] = t_i - E[\max\{\xi_1, \ldots, \xi_N\}] < t_i,$$

since $E[\max\{\xi_1, \ldots, \xi_N\}] > E[\xi_i] = 0$ (provided that ξ_i has nonzero probability of exceeding its mean zero).

Intuitively, winning should indicate to customer i that his noise term ξ_i is the largest and therefore his initial estimate t_i is upwardly biased. Therefore, if he were to bid his unconditional estimate t_i, winning the

[10]For example, the variance of the aggregate signal $\frac{1}{N}\sum_{i=1}^{N}\xi_i$ is a factor $1/\sqrt{N}$ smaller than the variance of ξ.

auction would indeed be bad news. It would indicate his expected surplus was negative; hence the *winner's curse*. To overcome this "curse," a rational customer must adjust his bid downward, considering the fact that it is the expected valuation of the item conditioned on having the highest signal that matters in determining his winnings.

The tendency of customers to reduce their bids to avoid the winner's curse changes the revenue equivalence of the basic auction types. In particular, while the sealed-bid auction conveys no information to customers, the Dutch (open descending price) and English (open ascending price) auctions provides them some information because they can observe how many customers are still willing—or not willing—to buy at the current price, when each drops out, and so on. The information about other customers' valuations tends to reduce the negative impact of the winner's curse.

For example, when an item has a common-value component, one can show the firm is better off using an English (open ascending price) auction than a sealed-bid, second-price auction—auctions that are strategically equivalent under the private-value model. Moreover, one can show that if the firm has its own signal (some private information) positively correlated with the item's value (like past price data of similar items or an appraisal), it benefits by sharing that information with the customers. This is because customers will tend to increase their estimate of the item's value as a result and bid more aggressively. Reserve prices also benefit the firm, but unlike in the private-value case, the optimal reserve price may vary with the type of auction and the number of customers.

6.2.7.2 Risk Aversion

Another factor affecting the results of the independent, private-value model is the assumption that both the firm and customer are risk-neutral. (See Appendix E for a discussion of risk preferences.) While the assumption of risk neutrality for a firm is often reasonable (for example, when the firm is a large, participating in many auctions over time), the assumption of risk neutrality for individual consumers is typically less realistic. However, it is easy to determine the relative performance of the standard auction types under risk aversion.

First, consider the case where the firm is risk-neutral and the customers are risk-averse. By revenue equivalence, note that a customer i's expected payment conditioned on his valuation v_i being the highest is the same under the first- and second-price auctions, since this expected payment is simply the expected revenue to the firm. However, the customer who wins a first-price auction pays a certain amount $b^*(v_i)$, while the same customer in a second-price auction will pay an uncertain amount with the same mean—namely, the valuation of the second-highest customer conditioned on the fact that v_i is the highest valuation. Thus, a risk-averse customer will prefer the first-price auction to the second-price auction. Given this preference, in the first-price auction risk-averse customers will tend to increase their bids above the risk-neutral equilibrium bid $b^*(v)$. (Bidding one's own valuation is still a dominant strategy under risk aversion in the second-price auction, so the bidding strategy in this case is not affected.) The higher resulting equilibrium bids in the first-price auction mean that the firm's expected revenue is higher as well, so the firm prefers this auction format.

Now consider the opposite case, where the firm is risk-averse and the customers are risk-neutral. By the same reasoning as above, the firm's revenue in the first-price auction conditioned on the winning value being v is certain while in the second-price auction it is uncertain. Therefore, unconditioning on v, the revenue in the second-price auction is more variable as well—also with the same mean as in the first-price auction. Thus, a risk-averse firm will also prefer the first-price auction.

The fact that the firm prefers the first-price auction in both cases (and is no worse in the second-price auction if all parties are risk-neutral) has been offered as one explanation for the relative popularity of first-price auctions over second-price auctions in practice.

6.2.7.3 Asymmetry Among Customers

Yet another departure from the private-value model is to relax the assumption of symmetry. The simplest case is to assume that there are two types of customers, types 1 and 2, with different valuations for the item drawn from different distributions, denoted $F_1(v)$ and $F_2(v)$, with corresponding marginal revenue (virtual value) functions $J_1(v)$ and $J_2(v)$. For example, type 1 customers may be experienced customers, and type 2 customers may be novice customers, or type 1 customers may be individuals while type 2 are industrial customers.

To see what can happen in this case, assume the first N_1 customers are of type 1 and the next N_2 are of type 2, and assume the marginal revenue functions are both increasing. The optimal allocation for the firm is obtained by maximizing

$$E\left[\sum_{i=1}^{N_1} J_1(v_i)y_i(v_i, \mathbf{v}_{-i}) + \sum_{i=N_1+1}^{N_1+N_2} J_2(v_i)y_i(v_i, \mathbf{v}_{-i})\right],$$

subject to the constraint that the total allocation is one

$$\sum_{i=1}^{N_1} y_i(v_i, \mathbf{v}_{-i}) + \sum_{i=N_1+1}^{N_1+N_2} y_i(v_i, \mathbf{v}_{-i}) = 1.$$

As before, it is optimal to allocate the item to the customer with the highest marginal value. However, note since $J_1(v)$ and $J_2(v)$ may differ, the customer with the highest marginal value is not necessarily the one with the highest valuation v.

This has important consequences for the optimal auction. For example, it means that it can be optimal for the firm to set different reserve prices for different types of customers, and the firm may systematically favor one class of customers over another in awarding the item. Indeed, one can show that in certain cases, it is optimal for the firm to favor the type of customers that tend to value the item less. The rationale for this is that by favoring these low-value types, the firm encourages the high-value types to bid even higher. The resulting higher equilibrium bids it receives from the high-value types more than compensates the firm for the loss he occasionally takes in favoring the low-value types.

In other words, it is optimal for the firm to *discriminate* among customers in the offering terms for the auction. This behavior is similar to the classical third-degree price-discrimination policy of offering different prices to different customer groups based on their different willingness to pay. (See Section 8.3.3.1.)

6.2.7.4 Collusion

The private-value model assumes the firm defines a game among the customers, intended to extract the highest prices possible from them. A key assumption in this game is that customers do not cooperate. Yet in practice, there is the possibility of *collusion* among customers, in which a coalition of customers (popularly called a *bidding ring*) cooperates and agrees to submit bids that are designed to reduce the price paid by the winner. Such collusion has been reported, for example, in the awarding of some government contracts.

There are several practical devices to reduce the likelihood of collusion among customers, most of which involve reducing the ability of customers to communicate among themselves. For example, one technique is to keep the identity of all customers secret, so customers cannot identify each other and form a bidding ring. (Though this may fail if the number of potential customers is so small that most customers know, *a priori,* the pool of likely participants—such as major suppliers in a procurement auction). Another technique is to reduce the amount of information relayed about bids to the minimum necessary to conduct the auction. For example, the firm might report only the highest current bid in an ascending auction, not the number of bids received, the time bids were received, or the history of bid values. This prevents customers from using such data to "signal" their intentions to each other.[11]

Because collusion can take so many forms, it is difficult to make general recommendations on the firm's "optimal response" to collusion. Nevertheless, to give some sense of the effect that it has consider a case where all N customers in the private-value model can collude perfectly. That is, they can get together and agree to submit bids, make payments, and allocate the item among themselves to maximize the surplus they receive as a group. In this case, the group of N customers effectively acts as a single "big customer" with valuation $\hat{v} = \max\{v_1, \ldots, v_N\}$, with distribution

$$P(\hat{v} \leq v) = F^N(v),$$

where $F(\cdot)$ is the distribution of the valuations for the N customers with density $f(\cdot)$. The marginal value for this distribution is

$$J_N(v) = v - \frac{1 - F^N(v)}{NF^{N-1}(v)f(v)}.$$

Of course, when faced with a single customer, the optimal auction is still the usual one: conduct a first or second-price auction with a reserve price set according to the marginal value of the single customer. So assuming $J_N(v)$ is increasing, the firm should set a reserve price v^* satisfying

$$J_N(v^*) = 0.$$

In this case, the bidding ring will be forced (yet willing) to pay v^* when its maximum valuation, $\hat{v}$, is at least this large. Also, one can show that this optimal reserve price is higher than the noncooperative optimal reserve price and that it increases with the number of customers N in the coalition. Thus, the possibility of collusion creates an incentive to use higher reserve prices than when customers do not collude. Indeed, the desire to thwart collusion is one of the main motivations for using reserve prices in practice.

[11]For example, at a keynote address to the Institute for Mathematics and its Applications (IMA) in December 2000, Robert Weber reported an instance in which bidders in a auction used the least significant digits in their bid amounts as a signaling mechanism. To overcome this, the auctioneer imposed larger minimum bid increments, thus reducing—or at least raising the cost of—this sort of signaling.

6.3 Optimal Dynamic Single-Resource Capacity Auctions

We next consider a dynamic auction problem that is in essence the auction equivalent of the single-resource problem of Chapter 2. In contrast to the traditional auction problem, in this case the firm receives bids from T groups of customers who are separated over time. In particular, in each period t, we assume that a new set of customers arrives and bids for the remaining capacity. The firm must determine winners in period t before observing the bids (or even the number of customers) in future periods. This dynamic feature parallels the traditional RM model, in which the firm must determine the capacity to sell in a given period before observing demand in future periods.

Such separation of customers over time is common in RM practice, a canonical example being the airline industry. Leisure travelers typically make travel plans months in advance of departure because they frequently must coordinate their vacation travel with other arrangements, like reserving resort accommodations, taking time off work or finding child care, and so on. In contrast, business travelers may not even know of their need to travel until a few days in advance of departure. As a result, if an airline were to conduct a single auction months in advance of departure, they would likely lose many business travelers; if they conducted a single auction a week before departure, they would likely lose many leisure travelers. This creates an incentive for them to conduct auctions at multiple points in time.

Other industries face similar situations, in which customers' needs are realized at different points of time (the need to buy a gift for a birthday, for example)—or are based on other contingent events (a new order to a manufacturer triggering a need for new supplies) that effectively separate customers in time. In such situations, a firm attempting to use a single auction at a single point in time would find itself eliminating many potential customers. By conducting multiple auctions over time, it can reach a larger pool of customers.

We next look at the optimal auction-design problem for this dynamic auction setting. We also compare the optimal auction to a traditional RM mechanism based on using dynamic list prices and capacity controls.

6.3.1 Formulation

A firm has an initial capacity of C units of a good that it wants to sell over a finite time horizon T. It does this by conducting a sequence of auctions indexed by $t = 1, \ldots, T$.

Customers are separated in time. In period t, N_t risk-neutral potential customers arrive. N_t is a nonnegative, discrete-valued random variable distributed according to a known p.m.f. $g(\cdot)$ with support $\{0,\ldots,M\}$ for some $M > 0$ and strictly positive first moment.

The assumptions parallel the private-value model: Each customer wishes to purchase at most one unit and has a reservation value v_i^t, $1 \leq i \leq N_t$. When the context is clear, we omit the time index and write v_i. Reservation values are private information, i.i.d. samples from a distribution $F(\cdot)$, which, as in the private-value model, is assumed strictly increasing with a continuous density function $f(\cdot)$ on the support $[0, \overline{v}]$, with $F(0) = 0$ and $F(\overline{v}) = 1$. To simplify notation and subsequent analysis, we assume that the distribution functions g and F do not depend on the time t but the extension to time-dependent distributions is straightforward.

The distributions F and g are assumed common knowledge to the firm and all potential customers (although this assumption can be relaxed for the second-price mechanism below). In addition, customer i knows his own (private) valuation v_i. Without loss of generality we assume that the unit salvage value for the firm at time $t = 0$ is $v_0 = 0$.

The firm's problem is to design an auction mechanism that maximizes its expected revenue. To do so, it must solve for an optimal allocation $y(v)$ in each period, given the values of N_t, v in each period and knowing only the probabilistic information (distributions) of these values in future periods.

Define the value function $V_t(x)$ as the maximum expected revenue obtainable from periods $t, t+1, \ldots, T$ given that there are x units in period t. Using (6.6) for the expected revenue in each period, the Bellman equation for $V_t(x)$ in terms of the allocation variables $y(v)$ can be written

$$V_t(x) = E_{N_t,v}\left[\max\left\{\sum_{i=1}^{N_t} J(v_i)\, y_i + V_{t+1}(x-k) : \; y_i \in \{0,1\}, \; k = \sum_{i=1}^{N_t} y_i, \; k \leq x\right\}\right], \tag{6.12}$$

where k is the total number of units awarded in period t. The boundary conditions are

$$V_{T+1}(x) = 0, \; x = 1, \ldots, C, \tag{6.13}$$

where C denotes the initial capacity. An allocation $y(\cdot)$ that achieves the maximum above given x, t and v will be an optimal dynamic allocation policy. (See Appendix D.)

6.3.2 Optimal Dynamic Allocations and Mechanisms

We first analyze the theoretical properties of the dynamic program (6.12)–(6.13). From this structure, one can show that variants of the classic first- and second-price auctions are optimal for this problem.

6.3.2.1 Optimal Allocations

As in the traditional single-resource RM model, the solution of the dynamic program (6.12)–(6.13) hinges on the monotonicity of the marginal values $\Delta V_t(x) = V_t(x) - V_t(x-1)$. Indeed, one can show the following [542]:

PROPOSITION 6.3 *$\Delta V_t(x)$ is decreasing in x for any fixed t and is decreasing in t for any fixed x.*

These are quite natural economic properties. At any point in time, the marginal benefit of each additional unit declines because the future number of customers is limited; therefore, the chance of selling the marginal unit—and the expected revenue if we sell it—decreases. Similarly, for any given remaining quantity x, the marginal benefit of an additional unit decreases with t because as time progresses, the number of future customers declines; therefore, the chance of selling the marginal unit—and the expected revenue if we sell it—goes down.

Proposition 6.3 simplifies the optimal allocation. To see this, note that since $J(\cdot)$ is assumed to be increasing, if the firm decides to award k units, it is optimal to allocate them to the highest $J(v_i)$'s (that is, to the highest v_i's). Therefore, define

$$R(k) \equiv \begin{cases} 0 & \text{if } k = 0 \\ \sum_{i=1}^{\min\{k,N_t\}} J(v_{[i]}) & \text{if } k > 0, \end{cases} \tag{6.14}$$

and note that

$$R(k) = \max\left\{ \sum_{i=1}^{N_t} J(v_i) y_i : y_i \in \{0,1\},\ \sum_i y_i = \min\{k, N_t\} \right\}.$$

Also, define $\Delta R(i) \equiv R(i) - R(i-1)$. Then the formulation (6.12) can be rewritten in terms of k as follows:

$$V_t(x) = E_{N_t,v}\left[\max_{0 \le k \le x} \left\{ \sum_{i=1}^{k} [\Delta R(i) - \Delta V_{t+1}(x-i+1)] \right\} \right] + V_{t+1}(x), \tag{6.15}$$

where the sum is defined to be 0 if $k = 0$. Let k^* be the optimal solution above (the optimal number of bids to accept) at time t in state x.

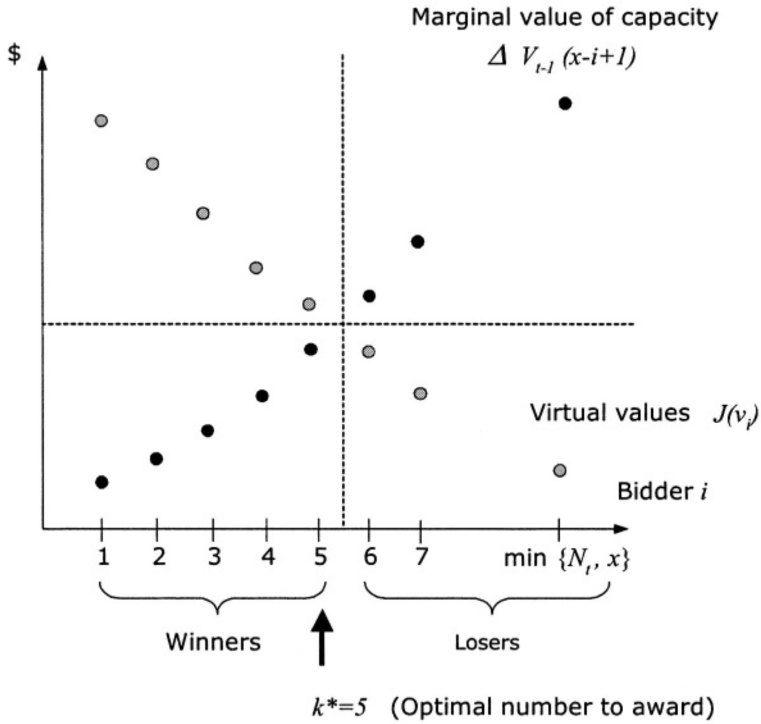

Figure 6.3. **Illustration of optimal allocations in the dynamic auction model.**

Let N_t denote a realization of the random variable N_t and v be a realization of customers' types. The following proposition characterizes the optimal allocation and follows from (6.15) and Proposition 6.3:

PROPOSITION 6.4 *For any realization* (N_t, v)*, the optimal number of units to allocate in state* (x, t) *is given by*

$$k^* = \max\{1 \leq k \leq \min\{x, N_t\} : \Delta R(k) > \Delta V_{t+1}(x - k + 1)\}$$

if $R(1) > \Delta V_{t+1}(x)$ *and by* $k^* = 0$ *otherwise. Moreover, it is optimal to award these* k^* *units to those customers with the highest valuations* v_i*.*

This shows how the firm should run the auction—provided it can infer the valuations v_i of the customers. In particular, note that $\Delta R(i) = J(v_{[i]})$ for $i = 1, \dots, \min\{x, N_t\}$, so the decision rule in Proposition 6.4 about the optimal number of bids to accept is simply based on sorting the values v_i and progressively awarding items to the highest-value customers until $J(v_{[i]})$ drops below the marginal opportunity cost $\Delta V_{t+1}(x - i+1)$. The situation is illustrated in Figure 6.3. Thus, given the customer valuations v_i and the value function $V_{t+1}(x)$, the optimal allocation rule is simple.

6.3.2.2 Optimal Mechanisms

We next demonstrate that appropriately modified versions of two standard procedures—the first- and the second-price auctions—achieve the optimal dynamic allocation.

Second-price auction In a straightforward application of the second-price mechanism in the dynamic auction setting, it is no longer optimal for customers to bid their valuation. The following informal reasoning shows why. Suppose it is optimal to bid truthfully under the second-price mechanism and let

$$\hat{v}_i \equiv J^{-1}(\Delta V_{t+1}(x - i + 1)), \quad i \geq 1. \tag{6.16}$$

The thresholds $\hat{v}_i$ are directly computable from the solution of (6.23) described in the previous section, which uses common knowledge information, and are in principle known to all customers and the firm. Following Theorem 6.4, the firm will accept bid $v_{[i]}$ as long as $v_{[i]} > \hat{v}_i$. Now suppose the firm decides to award k units. That means $v_{[i]} > \hat{v}_i$, $i = 1, \ldots, k$; and $v_{[i]} \leq \hat{v}_i$, $i = k+1, \ldots, N_t$. However, if the first loser, $v_{[k+1]}$, had bid $\hat{v}_{k+1} + \varepsilon$ instead (which in fact verifies $v_{[k+1]} < \hat{v}_{k+1} + \varepsilon$), the firm would include him among the winners and award $k+1$ units, and the customer would pay only $v_{[k+2]}$ and make a positive profit. Hence, customers have some incentive to bid above their own valuations (a pure second-price mechanism fails to elicit truthful bids).

However, the following modification to the second-price mechanism avoids this pitfall. In each period t, the firm first computes the thresholds $\hat{v}_i$ using the current capacity x. Given the vector of submitted bids $\mathbf{b}$, the firm will award k units, where

$$k = \max\{i \geq 1 : \; b_{[i]} > \hat{v}_i\}, \tag{6.17}$$

and $k = 0$ if $b_{[1]} \leq \hat{v}_1$; and all winners will pay

$$b^{(2^{\text{nd}})}_{[k+1]} = \max\{b_{[k+1]}, \hat{v}_k\}, \tag{6.18}$$

where $b_{[k+1]}$ is the $(k+1)^{\text{st}}$ highest bid and $\hat{v}_k$ is the threshold to award the k^{th} unit. Ties between bids are broken by randomization. For simplicity we refer to (6.17)-(6.18) as the *modified second-price* mechanism. One can then show the following result [542]:

PROPOSITION 6.5 *For the modified second-price auction with allocation and payments given by (6.17)-(6.18), a customer's dominant strategy* *is to bid his own valuation. Moreover, under this dominant-strategy equilibrium, the modified second-price mechanism is optimal.*

First-Price auction In a first-price auction, items are awarded to the highest bidders, and winners pay their bids. This type of mechanism may be more natural in many applications.

To establish that the first-price auction achieves the same expected revenue as the second-price mechanism described above, one needs to show that (1) items are again awarded according to the optimal allocation rule derived in the previous section and (2) customers with zero value have zero expected surplus. To do this, it suffices to show that there exists a symmetric equilibrium bidding strategy $b^*(\cdot)$ that is strictly increasing in the customer's valuation. In this case, the firm can use this bid function to invert a bid and infer the customer's valuation, which it can then use to correctly compute the number of items to award.

The main result for this case is the following [542]:

PROPOSITION 6.6 *Under the first-price auction, there exists a symmetric, strictly increasing, bidding strategy equilibrium $b^*(v_i)$. The strategy b^* depends on the current values of $\mathbf{t}$ and x as given by*

$$\hat{b}^*(v_i) = v_i - \frac{\int_{\hat{v}_1}^{v_i} P(v)\,\mathrm{d}v}{P(v_i)} \qquad \text{and} \quad b^*(v_i) \equiv \lim_{\varepsilon \downarrow 0} \hat{b}^*(v_i - \varepsilon), \tag{6.19}$$

where $P(v)$ is the probability that a customer with valuation v is among the winners,

$$P(v) = \begin{cases} 0 & \text{if } k^* = 0, \\ \sum_{n=1}^{M} \left\{ \sum_{k=0}^{k^*-1} \binom{n-1}{k} [1 - F(v)]^k \, [F(v)]^{n-1-k} \right\} g(n) & \text{if } k^* \geq 1, \end{cases} \tag{6.20}$$

$k^* = \max\{0 \leq i \leq \min\{x, N_t\} : v_i > \hat{v}_i\}$*, and by convention* $\hat{v}_0 < 0$. *Moreover, under this symmetric equilibrium, the first-price auction is optimal.*

Note that (6.19) shows—since winners are required to pay what they bid—that under a first-price mechanism customers shade their valuations to make some positive surplus. Since $b^*(\cdot)$ is strictly increasing, the units are sold to the players with the highest valuations. Moreover, once the firm observes bids $b_1^*, \ldots, b_{N_t}^*$, it can calculate the valuations $v_1, \ldots, v_{N_t}$ through the well-defined inverse bidding function $b^{*-1}(\cdot)$.

An important practical observation from this result is that the optimal first-price mechanism is not greedy, in the sense that it does not maximize the sum of *observable* revenue in the current period plus the expected revenue to go, because the firm compares the values of $J(v_i)$'s with the marginal value $\Delta V_t(\cdot)$ rather than with the bids themselves. As a result, the firm may (1) accept bids below the marginal value when $J(v_{[k]}) > \Delta V_{t+1}(x-k+1) \geq b^*(v_{[k]})$ and (2) reject bids that are above marginal value when $b^*(v_{[k]}) > \Delta V_{t+1}(x-k+1) \geq J(v_{[k]})$. Numerical experiments show that both cases may occur.

This behavior is somewhat counterintuitive because at first blush is seems that any bid that exceeds the marginal value of capacity ought to be worth accepting. However, such reasoning neglects the effect that the acceptance policy has on the bidding strategy of the customers. If the firm accepts all bids that are *ex post* profitable, then customers end up bidding *less* in equilibrium than they do when the firm follows the optimal acceptance strategy. The net result is to lower the firm's total revenue. In short, the firm has to occasionally refuse profitable bids to induce the customers to bid more aggressively—and in equilibrium it benefits by taking these short-run losses. This is simply an extension of the rationale for using reserve prices in a standard auction.

6.3.3 Comparisons with Traditional RM

We next compare the optimal auction mechanisms with a variation of a traditional quantity-based RM mechanism as in Chapter 2. The firm sets a list price at the beginning of each period and calculates a threshold on the number of units it is willing to award at the list price. Both the price and the capacity limit are optimized. We call this mechanism the *dynamic list price, capacity-controlled mechanism* (DLPCC). Note that unlike in a traditional RM mechanism, in DLPCC prices are set optimally rather than being given exogenously.

Customers who are interested in acquiring one unit at that list price submit *acceptances* (an offer to buy). If the number of acceptances exceeds the capacity limit set by the firm, the units are randomly rationed to the customers. It is easy to see in this case that a dominant strategy for customers is to submit an acceptance if and only if their valuations exceed the firm's reserve price.

6.3.3.1 Theoretical Comparisons

One can show that the DLPCC mechanism is, in fact, optimal in several cases. Indeed, we have [542]

PROPOSITION 6.7 *The DLPCC is optimal if the following cases:*
(i) There is at most one customer per period ($N_t \leq 1$ (w.p.1)).
(ii) There are more units to sell than there are potential customers

($\sum_t N_t < C$ (w.p.1)).
(iii) Asymptotically as the number of customers and units to sell grows large $(C, N_t \uparrow \infty)$.

That is, unless customers can be aggregated in time, the number of customers and objects is not too large, and there is some scarcity, there is no advantage to using a bidding mechanism over simple list pricing. These results are analogous to those in Section 6.2.6 for the single-period auction.

6.3.3.2 Numerical Comparisons

We next consider some numerical examples that illustrate the conditions under which an optimal pricing mechanism significantly outperforms DLPCC. In the examples that follow, the dynamic program associated with the optimal mechanism is solved using simulation, and customers' valuations are assumed to be uniformly distributed.

The first experiment shows how the revenue changes as the *concentration of customers,* defined as the number of customers per period, is varied. The firm starts with $C = 16$ units, and the total number of customers in all periods is constant at 64. The number of periods varies from 1 to 64, so that the number of customers per period varies. That is, the example runs from 64 customers in one period (high concentration of customers) to one customer in each of 64 consecutive periods (low concentration of customers).

The results are given in Table 6.1. Observe that the optimal revenue increases as the concentration of customers increases. This is intuitive, since as the firm observes more customers' valuations per period, it is making allocation decisions with reduced uncertainty about future bid values. Moreover, an increase in concentration increases direct bidding competition amongst customers. The gap reaches over 6% in the extreme case of a single period with 64 customers, which is significant. The second experiment compares the suboptimality gaps of the DLPCC mechanism under various levels of capacity and demand. The number of periods is kept constant at $T = 5$. The number of customers per period is fixed at $N_t = 10$, 30, 50 and 100; and for each of these, three choices of capacity–$C = 0.1\,T N_t$, $C = 0.3\,T N_t$ and $C = 0.5\,T N_t$—are used. Results are shown in Table 6.2. The gaps for DLPCC tend to decrease from left to right (which corresponds to increasing the capacity to demand ratio) and from top to bottom (which corresponds to increasing proportional number of customers per period and number of units in stock) in each table. Note that the gaps of 2% or more occur only in the

Table 6.1. Dynamic auction revenues for different concentrations of customers.

Customers	*Number*	*Optimal Revenue*		*DLPCC Revenue*	
Per Period	*of Periods*	*Mean*	*95% CI*	*Mean*	*Gap*
1	64	11.410	(11.390, 11.430)	11.412	0.16%
2	32	11.434	(11.420, 11.448)	11.401	0.41%
4	16	11.480	(11.466, 11.495)	11.383	0.98%
8	8	11.534	(11.511, 11.556)	11.348	1.79%
16	4	11.621	(11.602, 11.639)	11.292	2.99%
32	2	11.722	(11.704, 11.740)	11.201	4.59%
64	1	11.796	(11.780, 11.812)	11.060	6.36%

Table 6.2. DLPCC suboptimality gaps relative to a dynamic auction for different demand to capacity ratios.

N_t	$C = 0.1TN_t$	$C = 0.3TN_t$	$C = 0.5TN_t$
10	2.37%	2.32%	0.58%
30	1.77%	1.77%	0.38%
50	1.43%	1.43%	0.21%
100	1.06%	1.13%	0.14%

case where the number of customers is moderate (such as 10) and the capacity is constrained $(C = 0.1TN_t)$.

Other numerical experiments of [542] show that the relative benefit of the dynamic auction increase as the variance in the customer's reservation value v_i increases and as the variance in the number of customers N_t increases. Hence, variability in the demand environment appears to favor the use of a dynamic auction mechanism.

6.4 Optimal Dynamic Auctions with Replenishment

We next consider an infinite-horizon auction problem with replenishment, which is essentially the auction equivalent of the dynamic pricing and inventory problem of Section 5.3.2. A firm orders, stores, and then sells units of a homogeneous good over an infinite time horizon. The firm starts a period with an initial (integral) inventory x, and it reorders at a unit cost c at the end of the period. Replenishment orders arrive instantly, and we do not allow backlogging. In each period, a convex, strictly increasing holding cost of $h(x)$ is charged on the starting inventory level x.[12]

The firm sells its goods through a sequence of auctions indexed by $t \geq 1$. The problem is assumed to be stationary, so the statistics of demand are the same for all periods t. Private-value assumptions apply. In each period, N risk-neutral customers arrive. N is a nonnegative, discrete-valued random variable distributed according to a known probability mass function $g(\cdot)$ with support $[0, M]$ for some $M > 0$ and with a strictly positive first moment. Each customer requires one unit and has a private valuation v_i, $1 \leq i \leq N$, i.i.d. with a distribution $F(\cdot)$, which is strictly increasing with a continuous-density function $f(\cdot)$ on the support $[0, \overline{v}]$. We assume that the marginal value $J(\cdot)$ derived from $F(\cdot)$ is strictly increasing.

As in the single-resource capacity auction case, we use $\mathbf{v}$ both for the random vector of valuations (from the firm's perspective) and for a particular realization. The distribution functions g and F are constant through time t and are assumed common knowledge to the firm and all potential customers. We assume that both the number of customers N and their valuations $\mathbf{v}$ are independent from one period to the next. Thus, each period is an independent draw of N and $\mathbf{v}$.

The firm's problem is to design an auction mechanism and find a replenishment policy that maximizes its expected discounted profit. As before, we analyze this by first finding an optimal allocation and then finding mechanisms that achieve the optimal allocation.

6.4.1 Dynamic Programming Formulation

We analyze this problem using a dynamic programming formulation in terms of allocation variables $y(\mathbf{v})$. Define the value function $V(x)$ as the maximum expected discounted profit given an initial inventory $x = 0, 1, \ldots$, which satisfies the Bellman's equation:

$$V(x) = E\left[\max_{\substack{y \in \{0,1\}^N \\ q \in \mathbb{Z}_+}} \left\{ \sum_{i=1}^{N} J(v_i)\, y_i + \alpha\, V(x - k + q) - h(x - k + q) - c\,q : \sum_{i=1}^{N} y_i = k,\ k \leq x, \right\} \right] \tag{6.21}$$

[12] One can also analyze the case where holding cost is charged on the ending inventory level. The results are qualitatively the same as long as the holding cost $h(x)$ is linear.

where $0 < \alpha < 1$ is the discount factor, k is the total number of units awarded, and q is the replenishment order for the next period. Note from first principles that the state space can be bounded by M because at most M customers will arrive in any period, and since we can reorder at the end of every period, there is no need to stock more than M. Our objective is finding an optimal stationary policy consisting of an allocation $\mathbf{y}(\cdot)$ and a replenishment order $q(\cdot)$, that achieves $V(x)$.

Assuming $J(\cdot)$ is monotone increasing (Assumption 7.2), it again follows that if the firm allocates k units, it is optimal to allocate them to the highest $J(v_i)$'s (to the highest v_i's). So, as before, define

$$R(k) \equiv \begin{cases} 0 & \text{if } k = 0 \\ \sum_{i=1}^{\min\{k,N\}} J(v_{[i]}) & \text{if } k > 0, \end{cases} \tag{6.22}$$

and note that $R(k)$ is a random function that solves

$$R(k) = \max\{\sum_{i=1}^{N} J(v_i)y_i : 0 \le y_i \le 1, \ \sum_i y_i = \min\{k, N\}\}.$$

Therefore, we can rewrite (6.21) in terms of k as follows:

$$V(x) = E\left[\max_{\substack{0\le k\le x \\ q\in \mathcal{Z}_+}} \{R(k) + \alpha\, V(x-k+q) - h(x-k+q) - c\,q\}\right], \quad x = 0, 1, \ldots, M. \tag{6.23}$$

Note that above we are assuming that excess stock can be eliminated without cost (*free disposal*) when $N < k \le x$. This assumption is not essential for the analysis, but it helps to simplify the notation.

6.4.2 Optimal Auction and Replenishment Policy

We next characterize the optimal auction and replenishment policy for this problem. The first statement is presented in algorithmic form [527]:

PROPOSITION 6.8 *Consider the inventory-pricing problem described in (6.23). Define the optimal base-stock level by*

$$z^* = \max\{z \in \mathcal{Z}_+ : \alpha\, \Delta V(z) - \Delta h(z) - c > 0\}.$$

Then the optimal stationary policy is to allocate units to customers and replenish stock as follows:

STEP 1 (Allocate units):

FOR $k = 1, 2, \ldots, \min\{x, N\}$, *allocate the* k^{th} *unit if either:*
(i) $x - k \geq z^*$ *and* $J(v_{[k]}) > \alpha \Delta V(x - k + 1) - \Delta h(x - k + 1)$
(ii) $x - k < z^*$ *and* $J(v_{[k]}) > c$
ELSE GOTO STEP 2.

STEP 2 (Replenish stock):

IF $x - k < z^*$, *then order up to* z^*, *i.e.,* $q = z^* + k - x$; *ELSE order nothing* ($q = 0$).

The policy says that while the current inventory is above the optimal base-stock level z^* (case (i)), then we will award the k^{th} unit if the benefit from accepting the k^{th} bid (its virtual value $J(v_{[k]})$) exceeds the profit of keeping the k^{th} unit for the next period less the marginal holding cost for keeping it. The k^{th} unit is not replenished in this case. Once the inventory reaches the optimal level z^* (case (ii)), the firm awards a unit as long as the benefit from accepting a bid exceeds the cost of replacing the unit awarded; each such unit is replenished. This policy is illustrated in Figure 6.4.

An interesting result of this allocation policy is that when the inventory is less than the optimal base-stock level z^*, the firm can achieve the optimal allocation by simply running a standard first-price or second-price auction in each period with a fixed reserve price

$$\hat{v} \doteq J^{-1}(c). \tag{6.24}$$

Indeed, the following characterization of the optimal policy in this case [527]:

PROPOSITION 6.9 *Once the inventory reaches* z^* *units, the optimal policy in all subsequent periods is to (i) run a standard first- or second-price* z^**-unit auction with fixed reserve price* $\hat{v}$ *and then (ii) at the end of each period, order up to the optimal base-stock level* z^**.*

Since the problem is over an infinite horizon and the optimal policy calls only for ordering when the inventory drops below z^*, the firm eventually reaches a point where the above simple auction and replenishment policy are optimal for all remaining time. That is, z^* is the unique recurrent state in the resulting Markov chain that governs the evolution of the inventory over time under the optimal policy.

This result is significant on several levels. First, it shows that the classical first-price and second-price mechanisms remain optimal in the

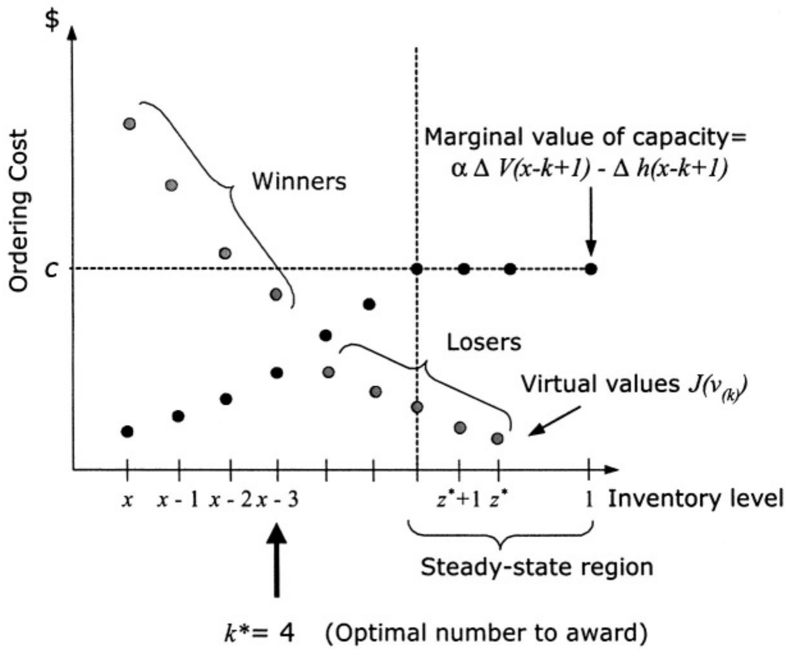

Figure 6.4. **Illustration of optimal allocations in the dynamic-auction model with replenishment.**

dynamic-inventory setting. These are both familiar auction mechanisms, which are easy for customers to understand and easy for firms to implement. The inventory-replenishment policy is also a familiar and simple base-stock policy. This combination makes the optimal policy quite practical. On a theoretical level, the result is as simple as one could hope for in this setting. Finally, it is convenient as well from a computational perspective because it reduces the optimal policy to a search over the single parameter z^*, as we show next.

6.4.3 Average-Profit Criterion

Consider maximizing the long-run average profit. One can show that the optimal policy for the α-discounted problem is in fact Blackwell optimal;[13] that is, it is simultaneously optimal for all discounted problems with discount factors $\alpha \in (\bar{\alpha}, 1)$, for some $0 < \bar{\alpha} < 1$. As a result, one can show (see [527]) that the optimal average-profit policy will again be to run a standard first-price or second-price auction in each period with

[13]See Bertsekas [57, Section 4.2, Definition 1.1]) for a formal definition of Blackwell optimality.

reserve price $\hat{v} = J^{-1}(c)$ and then order up to a fixed base-stock level z^* at the end of each period.

Indeed, because of this fact, one can develop a simple procedure for finding the optimal base-stock level z^*. Let

$$\Pi(z) \equiv E\left[\max_{0\le k\le \min\{z,N\}} \{R(k) - c\,k\}\right] - h(z) \qquad (6.25)$$

be the average profit when following a policy of reordering up to a fixed base-stock level z. We know that such a policy will be optimal for some z^*; therefore, we simply need to search for a value z that maximizes $\Pi(z)$. In fact, one can verify that the profit function $\Pi(z)$ is concave in z and that $z \le M$. $\Pi(z)$ can be evaluated by simulation and in special cases by closed-form expressions. Taking advantage of the concavity of $\Pi(z)$, a binary search over the range for z^* therefore gives an overall algorithm complexity of $O(N \log M)$. Henceforth, we denote the optimal objective value $\Pi^* \equiv \Pi(z^*)$.

6.4.4 Comparison with a List-Price Mechanism

We next consider how the optimal auction policy compares with a traditional, fixed-price policy. Specifically, we consider the base-stock, list price policy of Section 5.3.2, in which the firm sets a fixed list price p in each period and then replenishes by ordering up to a fixed base-stock level z. To be consistent, we assume we incur the holding cost at the beginning of the period, and we assume customers who are interested in acquiring one unit at the posted price submit acceptances. If the number of acceptances exceeds the current inventory of the firm, we randomly ration the units to the customers. It is easy to see that under this pricing mechanism, a dominant strategy for customers is to submit an acceptance if and only if their own valuations are higher than the list price.

We compare the profits earned under the optimal mechanism with those under the base-stock, list-price mechanism for an optimal choice of p and z. We give theoretical comparisons first, followed by a numerical comparison of the two policies.

6.4.4.1 Theoretical Comparisons

We restrict ourselves to the average-cost case, where the optimal profit is given by optimizing (6.25) over z, though similar results can be developed for the discounted case. One can show the following [527]:

PROPOSITION 6.10 *The base-stock, list price policy is optimal when*
(i) The number of customers is at most one—$N \le 1$ with probability

one.
(ii) The number of customers is large, $N \to \infty$, *and the holding cost is linear,* $h(z) = a + hz$.
(iii) The holding cost is zero, $h(x) = 0$.

Part (i) shows that if the firm is receiving isolated bids (as in some consumer online auctions, such as Priceline.com's mechanism), there is no inherent advantage to using auctions over list pricing. Some aggregation of customers is needed to gain a strict advantage through an auction mechanism. Intuitively, this is because one needs to generate some bidding competition among customers to realize a benefit from an auction. With at most one customer bidding, there is no competition. Part (ii) is analogous to the finite-horizon problem. As the number of customers in each period becomes large, the fraction with valuations above any given price p converges to a deterministic function of p, and hence the ratio of the auction and the fixed-price revenues tends to one. The intuitive reason for part (iii) is that with no holding cost, the firm will stock the maximum inventory M at the start of each period under both the optimal auction and list price policies. As a result, there is no rationing of product, and thus customers do not face any bidding competition. Without bidding competition, the auction produces the same profits as the base-stock, list price policy.

6.4.4.2 Numerical Comparisons

We next present the results of some numerical simulations from [527] with the average-profit criterion. The following base case is used as a starting point. The ordering cost is normalized at $c = 1$; customers' valuations are assumed uniform of width $\Delta = 0.5$ centered at c (that is, customers' valuations are centered at the cost, with Δ representing the dispersion in valuations); there are a constant $N = 50$ customers per period; and the holding cost is linear of the form $h(z) = c\,\delta\, z$, where $\delta = 1\%$ is the one-period interest rate.

The individual parameters of this base case are varied to see the effect on the absolute and relative performance of each policy. Along with expected profit, a *fill rate* is computed for each policy, defined as the expected number of customers who are awarded an item divided by the expected number who attempt to purchase (those with valuations above the reserve price in the auction or those with valuations above the fixed price in the base-stock, list price case).[14] The fill rate gives a measure of

[14] Formally, if $N(v)$ denotes the number of customers with valuations greater than v, then the fill rate is the ratio $E[\min\{N(\hat{v}), z^*\}]/E[N(\hat{v})]$ in the auction case and $E[\min\{N(p^*), z^*_{LP}\}]/E[N(p^*)]$ in the list price case, where p^* is the optimal list price.

the scarcity of inventory relative to demand and is a traditional service measure in inventory problems.

The first experiment shows how the profit is affected by the number of customers in each period. The number of customers N is assumed constant, but N is varied from 1 to 1,000. All other parameters are the same as in the base case. The results are summarized in Table 6.3. As

Table 6.3. **Dynamic auction and replenishment profits for different numbers of customers.**

N Customers	*Base-Stock, Auction*			*Base-Stock, List Price*			*Profit*
per Period	*Profit*	z^*	*Fill Rate*	*Profit*	z^*_{LP}	*Fill Rate*	*Gap*
1	0.02	1	100.00%	0.02	1	100.00%	0.00%
5	0.13	2	90.39%	0.12	3	98.74%	3.20%
10	0.27	4	95.93%	0.26	4	96.62%	2.50%
50	1.40	14	95.03%	1.38	16	98.88%	1.62%
100	2.83	26	94.90%	2.80	30	99.32%	1.31%
1000	28.72	242	95.86%	28.54	259	99.80%	0.62%

one would expect, the profits and inventory levels increase in both policies as the number of customers increases. Also, as shown theoretically in Proposition 6.10, the base-stock, list-price mechanism is optimal in the limiting case of just one customer per period. In the other extreme, as N gets large, again the base-stock, list price profit approaches the optimal auction profit, as predicted by the asymptotic result of Proposition 6.10. The biggest benefit from the auction occurs at a moderate value of five customers per period, where it achieves a 3.2% increase in profits over list pricing.

Note that the fill rate and inventory level are also higher in the base-stock, list price case. This suggests that the auction policy deliberately introduces some scarcity in the available goods to create more bidding competition among the customers.

The next experiment shows the effect of varying the interest rate δ—or equivalently varying the holding cost rate since $h(z) = c\delta z$ (with $c = 1$ in our case). Typically, this interest rate represents a cost of capital plus a rate of depreciation in the product's value over time. Table 6.4 shows the results. The small difference in the expected profits for the lowest interest rate confirms the result of Proposition 6.10—low holding cost leads to high inventory levels, which reduces the bidding competition and hence the benefit of the auction. As the interest rate rises, the auction performs relatively better, achieving a large 21.67% improvement over list pricing when the interest rate reaches 10%. This is simply the reverse effect: a high holding cost means the firm is unwilling to stock much

Table 6.4. Dynamic auction and replenishment profits for different holding costs.

Interest	*Base-Stock, Auction*			*Base-Stock, List Price*			*Profit*
Rate δ	*Profit*	z^*	*Fill Rate*	*Profit*	z^*_{LP}	*Fill Rate*	*Gap*
0.01%	1.560	21	99.97%	1.560	23	100.00%	0.01%
0.10%	1.543	18	99.58%	1.541	20	99.92%	0.14%
1.00%	1.404	14	95.03%	1.381	16	98.88%	1.62%
5.00%	0.932	10	77.36%	0.845	11	93.10%	9.37%
10.00%	0.502	7	55.77%	0.393	7	82.41%	21.67%

inventory. Since the number of customers per period is unchanged, the number of customers per unit of inventory increases; more competition among customers is created and hence the auction mechanism performs relatively better.

It is worth pointing out, however, that there are few practical situations where interest rates of over 1% per period are observed, especially if one is considering auctions that are held relatively frequently (such as weekly). Rates this high are observed for products such as personal computers, which become obsolete quickly, but for most goods, weekly rates of less than 1% are the norm. This suggests that either the product has to suffer rapid depreciation or selling events have to be relatively infrequent (such as monthly or semiannual periods, not weekly) for the firm to realize a significant benefit from using auctions over list pricing.

Finally, as in the finite-horizon problem without replenishment, numerical experiments show that variability in the valuations v_i or variability in the number of customers N increases the relative benefits of the auction policy.

6.5 Network Auctions

We next consider an auction mechanism for a network RM problem of the type studied in Chapter 3, which is based on Cooper and Menich [129]. Customers in this case bid for products (combinations of resources), and the firm awards resources based on these product-level bids. Such auctions are also relevant to procurement settings, where customers bid for a mix of inputs required to produce a given product. Customers desire the entire *bill of materials* (the complete set of resources) and a firm, with stockpiles of the various resources, must solicit bids and award the resources given a collection of *package* bids.

We next look at such an auction based on a network version of the Vickrey (second-price, sealed-bid) auction (a so-called Vickrey-Clarke-Groves mechanism [533, 120, 226]). We describe the basic mechanism and the resulting equilibrium bidding strategies and then explore the connections between this problem and traditional, network-capacity-control problems.

6.5.1 Problem Definition and Mechanism

The problem definition and notation are similar to those in Chapter 3, but are slightly modified to be consistent with the auction notation of this chapter. There are *N* customers, each with a private valuation v_j for one unit of product j, which requires one or more of m resources. We define $a_{ij} = 1$ if the product required by customer j uses resource i and $a_{ij} = 0$ otherwise. The incidence matrix is defined by $\mathbf{A} = [a_{ij}]$. The vector of current remaining capacities of the m resources is $\mathbf{x} = (x_1, \ldots, x_m)$. Because the mechanism is based on a generalized Vickrey auction, minimal assumptions about customers and their valuations are needed, as we show below. In fact, we require only that customers are rational and that their valuations are finite.

The mechanism is defined as follows. As in the classical auction setting, let $\mathbf{y} = (y_1, \ldots, y_N)$ denote the allocation vector, $\mathbf{p} = (p_1, \ldots, p_N)$ denote the payment vector, and $\mathbf{b}(\mathbf{v}) = (b_1(v_1), \ldots, b_N(v_N))$ denote the vector of bidding strategies. Customers submit a sealed bid $b_j(v_j)$ for their desired product j. The firm collects all *N* bids $\mathbf{b} = (b_1, \ldots, b_N)$ and then solves the following integer program:

$$
\begin{aligned}
z^*(\mathbf{b}) = \quad \max \quad & \mathbf{b}^\top \mathbf{y} \\
\text{s.t.} \quad & \mathbf{A}\mathbf{y} \leq \mathbf{x} \\
& y_j \in \{0, 1\}, \quad j = 1, \ldots, N.
\end{aligned}
\tag{6.26}
$$

Let $\mathbf{y}^*(\mathbf{b})$ denote an optimal solution to this integer program. The set of winning customers is denoted $\mathcal{W}(\mathbf{b}) = \{j : y_j^* = 1\}$.

It is important to note that the optimal value of this integer program is *not* the revenue earned by the firm; rather, it is solved simply as a means of determining winners and losers in the auction. The revenue to the firm will be determined by the vector of payments $\mathbf{p}(\mathbf{b})$ that are requested from the winning customers, which we look at next.

Note that the surplus of customer j is

$$S_j(\mathbf{b}, v_j) = v_j y_j^*(\mathbf{b}) - p_j(\mathbf{b}). \tag{6.27}$$

Let $\mathbf{e}_j$ denote the j^{th} unit vector and note that $\mathbf{b} - b_j\mathbf{e}_j$ is the vector of bids with the j^{th} component replaced by zero—that is, the vector of bids without the bid of customer j. Then the scheme calls for the winning customers to pay

$$p_j(\mathbf{b}) = y_j^*(\mathbf{b}) \left[b_j - (z^*(\mathbf{b}) - z^*(\mathbf{b} - b_j\mathbf{e}_j)) \right]. \tag{6.28}$$

Note that the term $z^*(\mathbf{b}) - z^*(\mathbf{b} - b_j\mathbf{e}_j)$ is simply the network benefit of having customer j's bid. And also clearly $z^*(\mathbf{b}) - z^*(\mathbf{b} - b_j\mathbf{e}_j) \leq b_j$, since when adding customer j's bid of b_j the optimal value of the problem (6.26) cannot increase by more than b_j. If $z^*(\mathbf{b}) - z^*(\mathbf{b} - b_j\mathbf{e}_j) < b_j$, it is because other winning bids were displaced to include the bid of customer j in the optimal solution. Hence, $b_j - (z^*(\mathbf{b}) - z^*(\mathbf{b} - b_j\mathbf{e}_j))$ represents the *displacement cost* produced by including customer j in the winning set, and hence in this scheme a customer pays his displacement cost. As we show below, this displacement-cost interpretation of the payment has a close connection to the bid-price values from the network problems studied in Chapter 3.

6.5.2 Equilibrium Analysis

We next analyze the equilibrium produced by this mechanism. An important relation is obtained by rewriting (6.28) as

$$p_j(\mathbf{b}) = z^*(\mathbf{b} - b_j\mathbf{e}_j) - \sum_{j' \neq j} b_{j'} y^*_{j'}(b). \tag{6.29}$$

This holds because

$$y^*_j(\mathbf{b})(-z^*(\mathbf{b}) + z^*(\mathbf{b} - b_j\mathbf{e}_j)) = -z^*(\mathbf{b}) + z^*(\mathbf{b} - b_j\mathbf{e}_j),$$

which is trivially true when $y^*_j(\mathbf{b}) = 1$; when $y^*_j(\mathbf{b}) = 0$, then it is true because in this case $z^*(\mathbf{b}) = z^*(\mathbf{b} - b_j\mathbf{e}_j))$.

Therefore, substituting (6.29) into (6.27), we find that the customer j's net utility can be written as

$$S_j(\mathbf{b}, v_j) = v_j y^*_j(\mathbf{b}) + \sum_{j' \neq j} b_{j'} y^*_{j'}(\mathbf{b}) - z^*(\mathbf{b} - b_j\mathbf{e}_j).$$

This shows that customer j's payoff does not depend on his bid b_j but only on whether his bid places him in the winning set $\mathcal{W}(\mathbf{b}) = \{j : y^*(j) = 1\}$. Thus, as in a second-price auction, one can show that if a customer bids less than his valuation, it reduces his chances of winning only in cases where he would have a positive net surplus, and bidding more than his valuation increases his chance of winning only in cases where his net surplus is negative. Indeed, one can prove

PROPOSITION 6.11 *Under a sealed-bid mechanism with allocations determined by an optimal solution $y^*(\mathbf{b})$ to (6.26) and payments determined by (6.28), then $b^*_j(v_j) = v_j$ is a dominant strategy for all customers, and hence $b^*(v) = v$ is a dominant-strategy equilibrium.*

As a result of this fact, we can assume $\mathbf{b} = \mathbf{v}$, and the equilibrium revenue collected by the firm is therefore given by

$$\begin{aligned}\sum_{j=1}^{N} p_j^*(\mathbf{v}) &= \sum_{j=1}^{N} y_j^*(\mathbf{v})v_j - z^*(\mathbf{v}) + z^*(\mathbf{v} - \mathbf{e}_j v_j) \\ &= \sum_{j=1}^{N} z^*(\mathbf{v} - \mathbf{e}_j v_j) - (N-1)z^*(\mathbf{v}).\end{aligned}$$

Note that this revenue is less than $z^*(\mathbf{v})$ since $p_j(\mathbf{v}) \le v_j$ and $p_j(\mathbf{v}) = 0$ when $y_j^*(v) = 0$. So the revenue collected by the firm is less than the optimal solution to the integer program (6.26) as claimed.

6.5.3 Relationship to Traditional Auctions

We next show that this mechanism is indeed the network generalization of the second-price mechanism in a traditional C-unit auction. To see this, note we can formulate the C-unit auction as an instance of the network auction with $m = 1$ and $x_1 = C$. In this case, solving the integer program (6.26) is trivial. We simply award the C items to the C customers with the highest bids b_j, which is the same as the classical second-price allocation. Also, note that the optimal value is

$$z^*(\mathbf{b}) = \sum_{j'=1}^{C} b_{[j']},$$

where $b_{[j]}$ denotes the j^{th} highest bid. As a result, by (6.28) each winner pays

$$\begin{aligned}p_j(\mathbf{b}) &= y_j^*(\mathbf{b})\left[b_j - z^*(\mathbf{b}) + z^*(\mathbf{b} - \mathbf{e}_j b_j)\right] \\ &= b_j - \sum_{j'=1}^{C} b_{[j']} + \left(\sum_{j'=1}^{C} b_{[j']} - b_j + b_{[C+1]}\right) \\ &= b_{[C+1]},\end{aligned}$$

which is just the usual second-price auction payment with no reserve price. Thus, the allocations and payments reduce to those of the classical C-unit second-price auction in the $m = 1$ case.

6.5.4 Relationship to Traditional Network RM

This network-auction mechanism also has an interesting connection to bid prices in traditional network RM. We proceed informally here to illustrate the ideas, but the connections can be made rigorous.

Consider the linear programming relaxation of (6.26), which is

$$\begin{aligned} z^*(\mathbf{b}) = \quad & \max \quad \mathbf{b}^\top \mathbf{y} \\ & \text{s.t.} \quad \mathbf{A}\mathbf{y} \leq \mathbf{x} \\ & \qquad\; 0 \leq y_j \leq 1, \quad j = 1, \ldots, N. \end{aligned} \tag{6.30}$$

Note that this is the exactly same form as the deterministic linear programming (DLP) model of Section 3.3.1, interpreting the demands for product j to be one for all j. Let $y^*(\mathbf{b})$ denote the optimal solution of (6.30). As in the DLP model, let $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_m)$ denote a vector of optimal dual variables for the capacity constraints $\mathbf{Ay} \leq \mathbf{x}$.

Note that if we remove customer j from the problem, then the reduction in revenue in this relaxed problem is zero if $y_j^*(\mathbf{b}) = 0$, while if $y_j^*(\mathbf{b}) = 1$, it is approximately given by[15]

$$z^*(\mathbf{b}) - z^*(\mathbf{b} - \mathbf{e}_j b_j) \approx b_j - \sum_{i \in A_j} \pi_i,$$

since removing customer j eliminates his bid b_j but frees up a unit of capacity on each resource i used by product j, and $\boldsymbol{\pi}$ gives the marginal benefit of this freed-up capacity. So the right-hand side above is approximately the net benefit of having customer j in the problem.

As a result, the amount a winning customer j pays from (6.28) is approximately

$$\begin{aligned} p_j(\mathbf{b}) &= b_j - (z^*(\mathbf{b}) - z^*(\mathbf{b} - \mathbf{e}_j b_j)) \\ &\approx \sum_{i \in A_j} \pi_i. \end{aligned}$$

Thus, roughly speaking, winning customers pay the bid prices of the set of resources required by the product they are bidding for.[16] Of course, the actual bidding mechanism uses an integer program rather than a linear program, but the connection is still close.

For example, if we allow continuous allocations in the auction (customers can receive a fractional quantity of the product they bid on and are willing to accept any quantity between 0 and 1), the two problems

[15]Here, we are ignoring the possibility that the dual is degenerate, and we are assuming the allowable decrease in the right-hand side of the constraints $\mathbf{Ay} \leq \mathbf{x}$ is at least one, so that $\mathbf{A}_j^\top \boldsymbol{\pi}$ measures the change in the optimal objective function when the capacity is reduced by the vector A_j.

[16]Despite the close connection between bid prices and the price paid by customers in this network Vickrey auction, the use of the term bid price is purely a coincidence; the two problems were not connected in the literature or in practice until quite recently.

coincide exactly. In this case, by linear programming duality, one can say that a customer j gets a positive allocation only if his bid b_j is at least as large as the sum of the bid prices, $\sum_{i \in A_j} \pi_i$. The payment of these customers with positive allocations is also given by the sum of the bid prices.

6.5.5 Revenue Maximization and Reserve Prices

While the network mechanism outlined above has a well-defined, dominant-strategy equilibrium, it is not revenue maximizing for the firm. To see this, suppose that $z^*(\mathbf{b}) - z^*(\mathbf{b} - \mathbf{e}_j b_j) = b_j$ for any customer j. This would occur, for example, if the resources required by the product requested by customer j are not capacity constrained, so that including j as one of the winners would not displace any other winners. In this case, the payment according to (6.28) is simply $p_j(b_j) = 0$. Therefore, any customers requesting unconstrained resources would win and pay nothing. However, clearly, the firm would increase its revenue by charging these customers something positive.

Just as in the classical auction, reserve prices can be used to increase the revenue in the network case. However, there is no theory showing how to construct optimal reserve prices in this case. Still, one can heuristically consider a scheme whereby the firm imposes reserve prices, denoted $\hat{\pi}_i$, on each resource i and requires each customer j to submit bids that exceed the sum of the reserve prices of the resources requested—that is, $b_j \geq \mathbf{A}_j^\top \hat{\boldsymbol{\pi}}$. It is still a dominant strategy for customer j to bid his valuation $b_j = v_j$, provided it exceed $\mathbf{A}_j^\top \hat{\boldsymbol{\pi}}$; otherwise, his dominant strategy is not to bid at all.

Numerical results show that one can increase revenue significantly by using such reserve prices. For example, Table 6.5 shows the simulated revenues for an example with two resources and three customer types. Type 1 customers require only resource 1, type 2 customers require only resource 2, and type 3 customers require both resources 1 and 2. The number of customers of each type is an independent Poisson random variable. Customers of types 1 and 2 have valuations with a mean of 100 and variance of 10; customers of type 3 have a valuation with mean 200 and variance of 20. There are 20 units of capacity for each of the two resources. Two demand scenarios are tested—a high-demand scenario in which the mean number of customers of each type is 15 and a low-demand scenario in which the mean number of customers of each type is 5. Symmetric reserve prices are used for each of the two resources, but they are varied.

Table 6.5 shows the effect of the reserve prices on the average revenues in the two scenarios. Note that in the high-demand scenario (Poisson-

15), the reserve price has a minimal effect on the average revenue for low reserve prices, with a maximum occurring at $70. However, revenues decrease significantly once the reserve price approaches $100, the mean valuation that customers have for each resource. In contrast, in the low-demand scenario (Poisson-5), the reserve price significantly increases revenues, achieving a maximum with a reserve price of $80. Again, revenues fall when we increase the reserve price beyond this point. This behavior is consistent with a traditional C-unit auction, where reserve prices affect only the revenue when there are fewer than C customers willing to bid above the reserve price.

Table 6.5. Network auction simulation results: average revenues as a function of reserve price.[a]

Resource Reserve Price	*Average Revenue (Poisson): 15 Customers per Product*	*Average Revenue (Poisson): 5 Customers per Product*
0	3652.7	3.6
10	3645.6	101.0
20	3701.3	799.0
30	3697.6	618.0
40	3680.6	798.0
50	3722.2	1018.7
60	3713.8	1209.7
70	3737.8	1408.8
80	3725.7	1566.9
90	3703.3	1468.5
100	2936.6	970.2
110	1019.7	374.2
120	46.7	16.2

[a] *Note:* Data as reported by Cooper and Menich [129].

6.6 Notes and Sources

The formal study of auctions stems from the seminal work of Vickrey [533], who derived the equilibrium strategies and the revenue equivalence of standard first and second-price auctions. The extensive two-volume collection edited by Klemperer [306] provides an excellent source for much of the literature on auction theory; see also Klemperer's [305] excellent survey article contained therein. Other survey articles on the private-value model are Matthews [366], McAfee [369], and Milgrom [381].

The analysis of optimal auction mechanisms for the private-value model, as described in Section 6.2.5, stems from the seminal paper of

Myerson [398], Maskin and Riley [364] extended Myerson's optimal auction analysis to multiunit auctions. The optimal discriminatory auction discussed in Section 6.2.7.3 is addressed in more detail in the survey of McAfee and McMillan [369] but was again originally due to Myerson [398].

The dynamic RM auction model in Section 6.3 is from Vulcano, van Ryzin, and Maglaras [542]. The infinite-horizon version with replenishment discussed in Section 6.4 is from van Ryzin and Vulcano [527].

The problem and results in Section 6.5 on network auctions are from Cooper and Menich [129]. For a an in-depth survey of other combinatorial auctions, see de Vries and Vohra [156].

APPENDIX 6.A: Proof of the Revenue-Equivalence Theorem

This surprisingly simple proof of the revenue-equivalence Theorem 6.1 is from Klemperer [305]:

Proof

Consider any symmetric equilibrium. Let $P(v)$ denote the probability that a customer winning under this equilibrium given his valuation is v (a type v customer), and let $S(v)$ denote the expected surplus of a customer with valuation v, defined by

$$S(v) = vP(v) - R(v),$$

where $R(v)$ is the expected payment. Since we are assuming an equilibrium, we must have that

$$S(v) \geq S(\tilde{v}) + (v - \tilde{v})P(\tilde{v}), \quad \forall \tilde{v}.$$

This follows because $P(\tilde{v})$ is the probability a customer wins if they were to follow the strategy of a customer with valuation $\tilde{v}$ instead of v. And if a customer with a valuation v wins by doing so, they value the item an amount $v - \tilde{v}$ different from a type $\tilde{v}$ customer. Hence, the right-hand side above is the expected surplus for a customer of valuation v if he follows the strategy of a type $\tilde{v}$ customer. However, as we are in equilibrium, a type v customer's surplus cannot be improved by deviating from the equilibrium strategy v.

Considering that a type v customer would not want to mimic a type $v + dv$ customer, we then have

$$S(v) \geq S(v + dv) + (-dv)P(v + dv),$$

and similarly since a type $v+dv$ customer would not want to mimic a type v customer,

$$S(v + dv) \geq S(v) + (dv)P(v).$$

Combining these two inequalities and we have that

$$P(v + dv) \geq \frac{S(v + dv) - S(v)}{dv} \geq P(v).$$

Since by assumption the allocation $y_i(v_i, \mathbf{v}_{-i})$ is increasing in v_i for all $\mathbf{v}_{-i}$, then $P(v)$ is increasing in v (since $P(v) = P(y_i(v, \mathbf{v}_{-i}) = 1)$), so the above inequalities are always feasible. Letting $dv \to 0$, shows

$$\frac{d}{dv}S(v) = P(v),$$

where upon integrating, we obtain

$$\begin{aligned} S(v) &= S(0) + \int_0^{\overline{v}} P(x)dx \\ &= \int_0^{\overline{v}} P(x)dx, \end{aligned} \tag{6.A.1}$$

where the last equality follows by the assumption 5(0) = 0 (i.e., customers with valuation zero have zero expected surplus).

Next, note that the expected payment, $R(v) = vP(v) - S(v)$, is equal to the expected revenue received by the firm. This means the firm's expected revenue from type v customer is

$$\int_0^{\overline{v}} (vP(v) - S(v))f(v)dv. \tag{6.A.2}$$

To evaluate this, note that by (6.A.1), we have

$$\begin{aligned} \int_0^{\overline{v}} S(v)f(v)dv &= \int_0^{\overline{v}} f(v)\left[\int_0^{v} P(x)dx\right]dv \\ &= \int_0^{\overline{v}} (1 - F(v))P(v)dv, \end{aligned} \tag{6.A.3}$$

where the last equality is obtained by integrating by parts, since

$$\begin{aligned} \int_0^{\overline{v}} f(v)\left[\int_0^{v} P(x)dx\right]dv &= \left[F(v)\int_0^{\overline{v}} P(x)dx\right]\Bigg|_0^{\overline{v}} \\ &\quad - \int_0^{\overline{v}} F(v)P(v)dv \\ &= \int_0^{\overline{v}} (1 - F(v))P(v)dv \end{aligned}$$

Substituting (6.A.3) into (6.A.2), the firm's expected revenue from each customer i is

$$\int_0^{\overline{v}} \left(v - \frac{1 - F(v)}{f(v)}\right) P(v)f(v)dv = E_{v_i}[J(v_i)P(v_i)],$$

where recall that $J(v) = v - \frac{1-F(v)}{f(v)}$ (which is always well defined since, by assumption, F is strictly increasing, so we always have that $f(v) > 0$). Summing over all N customers, the firm's total expected revenue is

$$\sum_{i=1}^{N} E_{v_i}[J(v_i)P(v_i)].$$

Finally, noting that the allocation variable $y_i(v_i, \mathbf{v}_{-i}) = 1$, if customer i is awarded an item and is zero otherwise, we have that $P(v_i) = E_{\mathbf{v}_{-i}}[y_i(v_i, \mathbf{v}_{-i})|v_i]$. Hence, the above expected revenue can be written

$$E_{\mathbf{v}}\left[\sum_{i=1}^{N} J(v_i) y_i(v_i, \mathbf{v}_{-i})\right].$$

QED