CHAPTER 9

AUCTIONS AND MECHANISM DESIGN

In most real-world markets, sellers do not have perfect knowledge of market demand as we have maintained so far in this book. Instead, sellers typically have only *statistical* information about market demand. Only the buyers themselves know precisely how much of the good they are willing to buy at a particular price. In this chapter, we will revisit the monopoly problem under this more typical circumstance.

Perhaps the simplest situation in which the above elements are present occurs when a single object is put up for auction. There, the seller is typically unaware of the buyers' values but may nevertheless have some information about the distribution of values across buyers. In such a setting, there are a number of standard auction forms that the seller might use to sell the good—first-price, second-price, Dutch, English. Do each of these standard auctions raise the same revenue for the seller? If not, which is best? Is there a nonstandard yet even better selling mechanism for the seller? To answer these and other questions, we will introduce and employ some of the tools from the theory of *mechanism design*.

Mechanism design is a general theory about how and when the design of appropriate institutions can achieve particular goals. This theory is especially germane when the designer requires information possessed only by others to achieve her goal. The subtlety in designing a successful mechanism lies in ensuring that the mechanism gives those who possess the needed information the incentive to reveal it to the designer. Although we will not explore the general theory of mechanism design here, this chapter provides an introduction to the topic by employing its techniques to study the design of revenue-maximizing auctions.

9.1 THE FOUR STANDARD AUCTIONS

Consider a seller with a single object for sale who wishes to sell the object to one of N buyers for the highest possible price. How should the seller go about achieving this goal? One possible answer is to hold an auction. Many distinct auctions have been put to use at one time or another, but we will focus on the following four standard auctions.[1]

[1] We shall assume throughout and unless otherwise noted that in all auctions ties in bids are broken at random: Each tied bidder is equally likely to be deemed the winner.

- **First-Price, Sealed-Bid:** Each bidder submits a sealed bid to the seller. The high bidder wins and pays his bid for the good.
- **Second-Price, Sealed-Bid:** Each bidder submits a sealed bid to the seller. The high bidder wins and pays the second-highest bid for the good.
- **Dutch Auction:** The seller begins with a very high price and begins to reduce it. The first bidder to raise her hand wins the object at the current price.
- **English Auction:** The seller begins with very low price (perhaps zero) and begins to increase it. Each bidder signals when he wishes to drop out of the auction. Once a bidder has dropped out, he cannot resume bidding later. When only one bidder remains, he is the winner and pays the current price.

Can we decide even among these four which is best for the seller? To get a handle on this problem, we must begin with a model.

9.2 The Independent Private Values Model

A single risk-neutral seller wishes to sell an indivisible object to one of N risk-neutral buyers. The seller values the object at zero dollars.[2] Buyer i's value for the object, v_i, is drawn from the interval [0, 1] according to the distribution function $F_i(v_i)$ with density function $f_i(v_i)$.[3] We shall assume that the buyers' values are mutually independent. Each buyer knows his own value but not the values of the other buyers. However, the density functions, $f_1, \ldots, f_N$, are public information and so known by the seller and all buyers. In particular, while the seller is unaware of the buyers' exact values, he knows the distribution from which each value is drawn. If buyer i's value is v_i, then if he wins the object and pays p, his payoff (i.e., von Neumann-Morgenstern utility) is $v_i - p$, whereas his payoff is $-p$ if he must pay p but does not win the object.[4]

This is known as the "independent, private values" model. **Independent** refers to the fact that each buyer's *private information* (in this case, each buyer's value) is independent of every other bidder's private information. **Private value** refers to the fact that once a buyer employs his own private information to assess the value of the object, this assessment would be unaffected were he subsequently to learn any other buyer's private information, i.e., each buyer's *private* information is sufficient for determining his *value*.[5]

Throughout this chapter, we will assume that the setting in which our monopolist finds himself is well-represented by the independent private values model. We can now begin to think about how the seller's profits vary with different auction formats. Note that with the

[2]This amounts to assuming that the object has already been produced and that the seller's use value for it is zero.

[3]Recall that $F_i(v_i)$ denotes the probability that i's value is less than or equal to v_i, and that $f_i(v_i) = F_i'(v_i)$. The latter relation can be equivalently expressed as $F_i(v_i) = \int_0^{v_i} f_i(x)\,dx$. Consequently, we will sometimes refer to f_i and sometimes refer to F_i since each one determines the other.

[4]Although such an outcome is not possible in any one of the four auctions above, there are other auctions (i.e., all-pay auctions) in which payments must be made whether or not one wins the object.

[5]There are more general models in which buyers with private information would potentially obtain yet *additional* information about the value of the object were they to learn *another buyer's* private information, but we shall not consider such models here.

production decision behind him and his own value equal to zero, profit-maximization is equivalent to revenue-maximization.

Before we can determine the seller's revenues in each of the four standard auctions, we must understand the bidding behavior of the buyers across the different auction formats. Let's start with the first-price auction.

9.2.1 BIDDING BEHAVIOR IN A FIRST-PRICE, SEALED-BID AUCTION

To understand bidding behavior in a first-price auction, we shall, for simplicity, assume that the buyers are ex-ante symmetric. That is, we shall suppose that for all bidders $i = 1, \ldots, N$, $f_i(v) = f(v)$ for all $v \in [0, 1]$.

Clearly, the main difficulty in determining the seller's revenue is in determining how the buyers, let's agree to call them *bidders* now, will bid. But note that if you are one of the bidders, then because you'd prefer to win the good at a lower price rather than a higher one, you will want to bid low when the others are bidding low and you will want to bid higher when the others bid higher. Of course, you do not know the bids that the others submit because of the sealed-bid rule. Yet, *your optimal bid will depend on how the others bid.* Thus, the bidders are in a strategic setting in which the optimal action (bid) of each bidder depends on the actions of others. Consequently, to determine the behavior of the bidders, we shall employ the game-theoretic tools developed in Chapter 7.

Let's consider the problem of how to bid from the point of view of bidder i. Suppose that bidder i's value is v_i. Given this value, bidder i must submit a sealed-bid, b_i. Because b_i will in general depend on i's value, let's write $b_i(v_i)$ to denote bidder i's bid when his value is v_i. Now, because bidder i must be prepared to submit a bid $b_i(v_i)$ for each of his potential values $v_i \in [0, 1]$, we may view bidder i's **strategy** as a *bidding function* b_i: $[0, 1] \to \mathbb{R}_+$, mapping each of his values into a (possibly different) nonnegative bid.

Before we discuss payoffs, it will be helpful to focus our attention on a natural class of bidding strategies. It seems very natural to expect that bidders with higher values will place higher bids. So, let us restrict attention to *strictly increasing* bidding functions. Next, because the bidders are ex-ante symmetric, it is also natural to suppose that bidders with the same value will submit the same bid. With this in mind, we shall focus on finding a strictly increasing bidding function, $\hat{b} : [0, 1] \to \mathbb{R}_+$, that is optimal for each bidder to employ, given that all other bidders employ this bidding function as well. That is, we wish to find a symmetric Nash equilibrium in strictly increasing bidding functions.

Now, let's suppose that we find a symmetric Nash equilibrium given by the strictly increasing bidding function $\hat{b}(\cdot)$. By definition it must be payoff-maximizing for a bidder, say i, with value v to bid $\hat{b}(v)$ given that the other bidders employ the same bidding function $\hat{b}(\cdot)$. Because of this, we can usefully employ what may at first appear to be a rather mysterious exercise.

The mysterious but useful exercise is this: Imagine that bidder i cannot attend the auction and that he sends a friend to bid for him. The friend knows the equilibrium bidding function $\hat{b}(\cdot)$, but he does not know bidder i's value. Now, if bidder i's value is v, bidder i would like his friend to submit the bid $\hat{b}(v)$ on his behalf. His friend can do this for him once bidder i calls him and tells him his value. Clearly, bidder i has no incentive to lie to his friend about his value. That is, among all the values $r \in [0, 1]$ that bidder i with value

v can report to his friend, his payoff is maximized by reporting his true value, v, to his friend. This is because reporting the value r results in his friend submitting the bid $\hat{b}(r)$ on his behalf. But if bidder i were there himself he would submit the bid $\hat{b}(v)$.

Let's calculate bidder i's expected payoff from reporting an arbitrary value, r, to his friend when his value is v, given that all other bidders employ the bidding function $\hat{b}(\cdot)$. To calculate this expected payoff, it is necessary to notice just two things. First, bidder i will win only when the bid submitted for him is highest. That is, when $\hat{b}(r) > \hat{b}(v_j)$ for all bidders $j \neq i$. Because $\hat{b}(\cdot)$ is strictly increasing this occurs precisely when r exceeds the values of all $N-1$ other bidders. Letting F denote the distribution function associated with f, the probability that this occurs is $(F(r))^{N-1}$ which we'll denote $F^{N-1}(r)$. Second, bidder i pays only when he wins and he then pays his bid, $\hat{b}(r)$. Consequently, bidder i's expected payoff from reporting the value r to his friend when his value is v, given that all other bidders employ the bidding function $\hat{b}(\cdot)$, can be written

$$u(r, v) = F^{N-1}(r)(v - \hat{b}(r)). \tag{9.1}$$

Now, as we have already remarked, because $\hat{b}(\cdot)$ is an equilibrium, bidder i's expected payoff maximizing bid when his value is v must be $\hat{b}(v)$. Consequently, (9.1) must be maximized when $r = v$, i.e., when bidder i reports his true value, v, to his friend. So, we may differentiate the right-hand side with respect to r and set the derivative equal to zero when $r = v$. Differentiating yields

$$\frac{dF^{N-1}(r)(v - \hat{b}(r))}{dr} = (N-1)F^{N-2}(r)f(r)(v - \hat{b}(r)) - F^{N-1}(r)\hat{b}'(r). \tag{9.2}$$

Setting this equal to zero when $r = v$ and rearranging yields

$$(N-1)F^{N-2}(v)f(v)\hat{b}(v) + F^{N-1}(v)\hat{b}'(v) = (N-1)vf(v)F^{N-2}(v). \tag{9.3}$$

Looking closely at the left-hand side of (9.3), we see that it is just the derivative of the product $F^{N-1}(v)\hat{b}(v)$ with respect to v. With this observation, we can rewrite (9.3) as

$$\frac{dF^{N-1}(v)\hat{b}(v)}{dv} = (N-1)vf(v)F^{N-2}(v). \tag{9.4}$$

Now, because (9.4) must hold for every v, it must be the case that

$$F^{N-1}(v)\hat{b}(v) = (N-1)\int_0^v xf(x)F^{N-2}(x)\,dx + \text{constant}.$$

Noting that a bidder with value zero must bid zero, we conclude that the constant above must be zero. Hence, it must be the case that

$$\hat{b}(v) = \frac{N-1}{F^{N-1}(v)}\int_0^v xf(x)F^{N-2}(x)\,dx,$$

which can be written more succinctly as

$$\hat{b}(v) = \frac{1}{F^{N-1}(v)} \int_0^v x dF^{N-1}(x). \tag{9.5}$$

There are two things to notice about the bidding function in (9.5). First, as we had assumed, it is strictly increasing in v (see Exercise 9.1). Second, it has been uniquely determined. Hence, in conclusion, we have proven the following.

THEOREM 9.1 ***First-Price Auction Symmetric Equilibrium***

If N bidders have independent private values drawn from the common distribution, F, then bidding

$$\hat{b}(v) = \frac{1}{F^{N-1}(v)} \int_0^v x dF^{N-1}(x)$$

whenever one's value is v constitutes a symmetric Nash equilibrium of a first-price, sealed-bid auction. Moreover, this is the only symmetric Nash equilibrium.[6]

EXAMPLE 9.1 Suppose that each bidder's value is uniformly distributed on $[0, 1]$. Then $F(v) = v$ and $f(v) = 1$. Consequently, if there are N bidders, then each employs the bidding function

$$\begin{aligned}
\hat{b}(v) &= \frac{1}{v^{N-1}} \int_0^v x dx^{N-1} \\
&= \frac{1}{v^{N-1}} \int_0^v x(N-1)x^{N-2}\, dx \\
&= \frac{N-1}{v^{N-1}} \int_0^v x^{N-1}\, dx \\
&= \frac{N-1}{v^{N-1}} \frac{1}{N} v^N \\
&= v - \frac{v}{N}.
\end{aligned}$$

So, each bidder *shades* his bid, by bidding less than his value. Note that as the number of bidders increases, the bidders bid more aggressively. □

Because $F^{N-1}(\cdot)$ is the distribution function of the highest value among a bidder's $N-1$ competitors, the bidding strategy displayed in Theorem 9.1 says that each bidder bids the expectation of the second highest bidder's value conditional on his own value being

[6]Strictly speaking, we have not shown that this is an equilibrium. We have shown that *if* a symmetric equilibrium exists, then this must be it. You are asked to show that this is indeed an equilibrium in an exercise. You might also wonder about the existence of asymmetric equilibria. It can be shown that there are none, although we shall not do so here.

highest. But, because the bidders use the *same strictly increasing* bidding function, having the highest value is equivalent to having the highest bid and so equivalent to winning the auction. So, we may say that—

In the unique symmetric equilibrium of a first-price, sealed-bid auction, each bidder bids the expectation of the second-highest bidder's value conditional on winning the auction.

The idea that one ought to bid *conditional on winning* is very intuitive in a first-price auction because of the feature that one's bid matters only when one wins the auction. Because this feature is present in other auctions as well, this idea should be considered one of the basic insights of our strategic analysis.

Having analyzed the first-price auction, it is an easy matter to describe behavior in a Dutch auction.

9.2.2 BIDDING BEHAVIOR IN A DUTCH AUCTION

In a Dutch auction, each bidder has a single decision to make, namely, "At what price should I raise my hand to signal that I am willing to buy the good at that price?" Moreover, the bidder who chooses the highest price wins the auction and pays this price. Consequently, by replacing the word "price" by "bid" in the previous sentence we see that this auction is equivalent to a first-price auction! So, we can immediately conclude the following.

THEOREM 9.2 ***Dutch Auction Symmetric Equilibrium***

If N bidders have independent private values drawn from the common distribution, F, then raising one's hand when the price reaches

$$\frac{1}{F^{N-1}(v)}\int_0^v x\,dF^{N-1}(x)$$

whenever one's value is v constitutes a symmetric Nash equilibrium of a Dutch auction. Moreover, this is the only symmetric Nash equilibrium.

Clearly then, the first-price and Dutch auctions raise exactly the same revenue for the seller, ex-post (i.e., for every realization of bidder values $v_1, \ldots, v_N$).

We now turn to the second-price, sealed-bid auction.

9.2.3 BIDDING BEHAVIOR IN A SECOND-PRICE, SEALED-BID AUCTION

One might wonder why we would bother considering a second-price auction at all. Isn't it obvious that a first-price auction must yield higher revenue for the seller? After all, in a first-price auction the seller receives the *highest* bid, whereas in a second-price auction she receives only the *second-highest* bid.

While this might sound convincing, it neglects a crucial point: *The bidders will bid differently in the two auctions.* In a first-price auction, a bidder has an incentive to raise her bid to increase her chances of winning the auction, yet she has an incentive to reduce her bid to lower the price she pays when she does win. In a second-price auction, the second effect is absent because when a bidder wins, the amount she pays is independent of her bid. So, we should expect bidders to bid *more aggressively* in a second-price auction than they would in a first-price auction. Therefore, there is a chance that a second-price auction will generate higher expected revenues for the seller than will a first-price auction. When we recognize that bidding behavior changes with the change in the auction format, the question of which auction raises more revenue is not quite so obvious, is it?

Happily, analyzing bidding behavior in a second-price, sealed-bid auction is remarkably straightforward. Unlike our analysis of the first-price auction, we need not restrict attention to the case involving symmetric bidders. That is, we shall allow the density functions $f_1, \ldots, f_N$, from which the bidders' values are independently drawn, to differ.[7]

Consider bidder i with value v_i, and let B denote the highest bid submitted by the other bidders. Of course, B is unknown to bidder i because the bids are sealed. Now, if bidder i were to win the auction, his bid would be highest and B would then be the second-highest bid. Consequently, bidder i would have to pay B for the object. In effect, then, the price that bidder i must pay for the object is the highest bid, B, submitted by the other bidders.

Now, because bidder i's value is v_i, he would strictly want to win the auction when his value exceeds the price he would have to pay, i.e., when $v_i > B$; and he would strictly want to lose when $v_i < B$. When $v_i = B$ he is indifferent between winning and losing. Can bidder i bid in a manner that guarantees that he will win when $v_i > B$ and that he will lose when $v_i < B$, even though he does not know B? The answer is yes. He can guarantee precisely this simply by bidding his value, v_i!

By bidding v_i, bidder i is the high bidder, and so wins, when $v_i > B$, and he is not the high bidder, and so loses, when $v_i < B$. Consequently, bidding his value is a payoff-maximizing bid for bidder i *regardless of the bids submitted by the other bidders* (recall that B was the highest bid among any arbitrary bids submitted by the others). Moreover, because bidding below one's value runs the risk of losing the auction when one would have strictly preferred winning it, and bidding above one's value runs the risk of winning the auction for a price above one's value, bidding one's value is a weakly dominant bidding strategy. So, we can state the following.

THEOREM 9.3 ***Second-Price Auction Equilibrium***

If N bidders have independent private values, then bidding one's value is the unique weakly dominant bidding strategy for each bidder in a second-price, sealed-bid auction.

This brings us to the English auction.

9.2.4 BIDDING BEHAVIOR IN AN ENGLISH AUCTION

In contrast to the auctions we have considered so far, in an English auction there are potentially many decisions a bidder has to make. For example, when the price is very low,

[7]In fact, even the independence assumption can be dropped. (See Exercise 9.3.)

he must decide at which price he would drop out when no one has yet dropped out. But, if some other bidder drops out first, he must then decide at which price to drop out *given* the remaining active bidders, and so on. Despite this, there is a close connection between the English and second-price auctions.

In an English auction, as in a second-price auction, it turns out to be a dominant strategy for a bidder to drop out when the price reaches his value, regardless of which bidders remain active. The reason is rather straightforward. A bidder i with value v_i who, given the history of play and the current price $p < v_i$, considers dropping out can do no worse by planning to remain active a little longer and until the price reaches his value, v_i. By doing so, the worst that can happen is that he ends up dropping out when the price does indeed reach his value. His payoff would then be zero, just as it would be if he were to drop out now at price p. However, it might happen, were he to remain active, that all other bidders would drop out before the price reaches v_i. In this case, bidder i would be strictly better off by having remained active since he then wins the object at a price strictly less than his value v_i, obtaining a positive payoff. So, we have the following.

THEOREM 9.4 ***English Auction Equilibrium***

If N bidders have independent private values, then dropping out when the price reaches one's value is the unique weakly dominant bidding strategy for each bidder in an English auction.[8]

Given this result, it is easy to see that the bidder with the highest value will win in an English auction. But what price will he pay for the object? That, of course, depends on the price at which his *last remaining competitor* drops out of the auction. But his last remaining competitor will be the bidder with the *second-highest value*, and he will, like all bidders, drop out when the price reaches his value. Consequently, the bidder with highest value wins and pays a price equal to the second-highest value. Hence, we see that the outcome of the English auction is identical to that of the second-price auction. In particular, the English and second-price auctions earn exactly the same revenue for the seller, ex-post.

9.2.5 REVENUE COMPARISONS

Because the first-price and Dutch auctions raise the same ex-post revenue and the second-price and English auctions raise the same ex-post revenue, it remains only to compare the revenues generated by the first- and second-price auctions. Clearly, these auctions need not raise the same revenue ex-post. For example, when the highest value is quite high and the second-highest is quite low, running a first-price auction will yield more revenue than a second-price auction. On the other hand, when the first- and second-highest values are close together, a second-price auction will yield higher revenues than will a first-price auction.

Of course, when the seller must decide which of the two auction forms to employ, he does not know the bidders' values. However, knowing how the bidders bid as functions

[8] As in the second-price auction case, this weak dominance result does not rely on the independence of the bidder's values. It holds even if the values are correlated. However, it is important that the values are *private*.

of their values, and knowing the distribution of bidder values, the seller can calculate the *expected revenue* associated with each auction. Thus, the question is, which auction yields the highest expected revenue, a first- or a second-price auction? Because our analysis of the first-price auction involved symmetric bidders, we must assume symmetry here to compare the expected revenue generated by a first-price versus a second-price auction. So, in what follows, $f(\cdot)$ will denote the common density of each bidder's value and $F(\cdot)$ will denote the associated distribution function.

Let's begin by considering the expected revenue, R_{FPA}, generated by a first-price auction (FPA). Because the highest bid wins a first-price auction and because the bidder with the highest value submits the highest bid, if v is the highest value among the N bidder values, then the seller's revenue is $\hat{b}(v)$. So, if the highest value is distributed according to the density $g(v)$, the seller's expected revenue can be written

$$R_{FPA} = \int_0^1 \hat{b}(v)g(v)\,dv.$$

Because the density, g, of the maximum of N independent random variables with common density f and distribution F is NfF^{N-1},[9] we have

$$R_{FPA} = N\int_0^1 \hat{b}(v)f(v)F^{N-1}(v)\,dv. \tag{9.6}$$

We have seen that in a second-price auction, because each bidder bids his value, the seller receives as price the second-highest value among the N bidder values. So, if $h(v)$ is the density of the second-highest value, the seller's expected revenue, R_{SPA}, in a second-price auction can be written

$$R_{SPA} = \int_0^1 vh(v)\,dv.$$

Because the density, h, of the second-highest of N independent random variables with common density f and distribution function F is $N(N-1)F^{N-2}f(1-F)$,[10] we have

$$R_{SPA} = N(N-1)\int_0^1 vF^{N-2}(v)f(v)(1-F(v))\,dv. \tag{9.7}$$

[9] To see this, note that the highest value is less than or equal to v if and only if all N values are, and that this occurs with probability $F^N(v)$. Hence, the distribution function of the highest value is F^N. Because the density function is the derivative of the distribution function the result follows.

[10] One way to see this is to treat probability density like probability. Then the probability (density) that some particular bidder's value is v is $f(v)$ and the probability that exactly one of the remaining $N-1$ other bidders' values is above this is $(N-1)F^{N-2}(v)(1-F(v))$. Consequently, the probability that this particular bidder's value is v and it is second-highest is $(N-1)f(v)F^{N-2}(v)(1-F(v))$. Because there are N bidders, the probability (i.e., density) that the second-highest value is v is then $N(N-1)f(v)F^{N-2}(v)(1-F(v))$.

We shall now compare the two. From (9.6) and (9.5) we have

$$
\begin{aligned}
R_{FPA} &= N \int_0^1 \left[\frac{1}{F^{N-1}(v)} \int_0^v x dF^{N-1}(x) \right] f(v) F^{N-1}(v)\, dv \\
&= N(N-1) \int_0^1 \left[\int_0^v x F^{N-2}(x) f(x)\, dx \right] f(v)\, dv \\
&= N(N-1) \int_0^1 \int_0^v [x F^{N-2}(x) f(x) f(v)]\, dx dv \\
&= N(N-1) \int_0^1 \int_x^1 [x F^{N-2}(x) f(x) f(v)]\, dv dx \\
&= N(N-1) \int_0^1 x F^{N-2}(x) f(x) (1 - F(x))\, dx \\
&= R_{SPA},
\end{aligned}
$$

where the fourth equality follows from interchanging the order of integration (i.e., from $dxdv$ to $dvdx$), and the final equality follows from (9.7).

EXAMPLE 9.2 Consider the case in which each bidder's value is uniform on [0, 1] so that $F(v) = v$ and $f(v) = 1$. The expected revenue generated in a first-price auction is

$$
\begin{aligned}
R_{FPA} &= N \int_0^1 \hat{b}(v) f(v) F^{N-1}(v)\, dv \\
&= N \int_0^1 \left[v - \frac{v}{N} \right] v^{N-1} dv \\
&= (N-1) \int_0^1 v^N dv \\
&= \frac{N-1}{N+1}.
\end{aligned}
$$

On the other hand, the expected revenue generated in a second-price auction is

$$
\begin{aligned}
R_{SPA} &= N(N-1) \int_0^1 v F^{N-2}(v) f(v) (1 - F(v))\, dv \\
&= N(N-1) \int_0^1 v^{N-1} (1 - v)\, dv \\
&= N(N-1) \left[\frac{1}{N} - \frac{1}{N+1} \right] \\
&= \frac{N-1}{N+1}.
\end{aligned}
$$

□

Remarkably, the first- and second-price auctions raise the *same* expected revenue, regardless of the common distribution of bidder values! So, we may state the following:

If N bidders have independent private values drawn from the common distribution, F, then all four standard auction forms (first-price, second-price, Dutch, and English) raise the same expected revenue for the seller.

This *revenue equivalence* result may go some way toward explaining why we see all four auction forms in practice. Were it the case that one of them raised more revenue than the others on average, then we would expect that one to be used rather than any of the others. But what is it that accounts for the coincidence of expected revenue in these auctions? Our next objective is to gain some insight into why this is so.

9.3 The Revenue Equivalence Theorem

To explain the equivalence of revenue in the four standard auction forms, we must first find a way to fit all of these auctions into a single framework. With this in mind, we now define the notion of a *direct selling mechanism.*[11]

A **direct selling mechanism** is a collection of N *probability assignment functions*, $p_1(v_1, \ldots, v_N), \ldots, p_N(v_1, \ldots, v_N)$, and N *cost functions* $c_1(v_1, \ldots, v_N), \ldots, c_N(v_1, \ldots, v_N)$. For each i and every vector of values $(v_1, \ldots, v_N)$, $p_i(v_1, \ldots, v_N) \in [0, 1]$ denotes the probability that bidder i receives the object and $c_i(v_1, \ldots, v_N) \in \mathbb{R}$ denotes the payment that bidder i must make to the seller.[12] Consequently, the sum of the probabilities, $p_1(v_1, \ldots, v_N) + \cdots + p_N(v_1, \ldots, v_N)$, must never exceed unity. On the other hand, we allow this sum to fall short of unity because we want to allow the seller to keep the object.[13]

A direct selling mechanism works as follows. Because the seller does not know the bidders' values, he asks them to report them to him simultaneously. He then takes those reports, $r_1, \ldots, r_N$, which need not be truthful, and assigns the object to one of the bidders according to the probabilities $p_i(r_1, \ldots, r_N)$, $i = 1, \ldots, N$, keeping the object with the residual probability, and secures the payment $c_i(r_1, \ldots, r_N)$ from each bidder $i = 1, \ldots, N$. It is assumed that the entire direct selling mechanism—the probability assignment functions and the cost functions—are public information, and that the seller must carry out the terms of the mechanism given the vector of reported values.

Clearly, the seller's revenue will depend on the reports submitted by the bidders. Will they be induced to report truthfully? If not, how will they behave? These are very good questions, but let's put them aside for the time being. Instead, let us consider a different question: How are the four standard auctions related to direct selling mechanisms?

What we will show is that each of the four standard auctions can be equivalently viewed as an **incentive-compatible** direct selling mechanism. That is, a direct selling mechanism *in which it is an equilibrium for the bidders to report their values truthfully*. These mechanisms

[11]Our presentation is based upon Myerson (1981).

[12]Note, first, that a bidder's cost may be negative and, second, that a bidder's cost may be positive even when that bidder does not receive the object (i.e., when that bidder's probability of receiving the object is zero).

[13]This is more generality than we need at the moment because the seller never keeps the object in any of the four standard auctions. However, this will be helpful a little later.

will prove to be central. Indeed, understanding incentive-compatible direct selling mechanisms will not only be the key to understanding the connection among the four standard auctions, but it will be central to our understanding revenue-maximizing auctions as well.

Consider a first-price auction with symmetric bidders. We'd like to construct an "equivalent" direct selling mechanism in which truth-telling is an equilibrium. To do this, we shall employ the first-price auction equilibrium bidding function $\hat{b}(\cdot)$. The idea behind our construction is simple. Instead of the bidders submitting bids computed by plugging their values into the equilibrium bidding function, the bidders will be asked to submit their values and the seller will then compute their equilibrium bids for them. Recall that because $\hat{b}(\cdot)$ is strictly increasing, a bidder wins the object in a first-price auction if and only if he has the highest value.

Consider, then, the following direct selling mechanism, where $\hat{b}(\cdot)$ is the equilibrium bidding function for the first-price auction given in (9.5):

$$p_i(v_1, \ldots, v_N) = \begin{cases} 1, & \text{if } v_i > v_j \text{ for all } j \neq i \\ 0, & \text{otherwise,} \end{cases}$$
$$\text{and} \tag{9.8}$$
$$c_i(v_1, \ldots, v_N) = \begin{cases} \hat{b}(v_i), & \text{if } v_i > v_j \text{ for all } j \neq i \\ 0, & \text{otherwise.} \end{cases}$$

Look closely at this mechanism. Note that the bidder with the highest reported value, v, receives the object and he pays $\hat{b}(v)$ for it, just as he would have in a first-price auction equilibrium. So, if the bidders report their values truthfully, then the bidder with the highest value, v, wins the object and makes the payment $\hat{b}(v)$ to the seller. Consequently, if this mechanism is incentive-compatible, the seller will earn exactly the same ex-post revenue as he would with a first-price auction.

To demonstrate that this mechanism is incentive-compatible we need to show that truth-telling is a Nash equilibrium. So, let us suppose that all other bidders report their values truthfully and that the remaining bidder has value v. We must show that this bidder can do no better than to report his value truthfully to the seller. So, suppose that this bidder considers reporting value r. He then wins the object and makes a payment of $\hat{b}(r)$ if and only if $r > v_j$ for all other bidders j. Because the other $N-1$ bidders' values are independently distributed according to F, this event occurs with probability $F^{N-1}(r)$. Consequently, this bidder's expected payoff from reporting value r when his true value is v is

$$F^{N-1}(r)(v - \hat{b}(r)).$$

But this is exactly the payoff in (9.1), which we already know is maximized when $r = v$. Hence, the direct selling mechanism (9.8) is indeed incentive-compatible.

Let's reconsider what we have accomplished here. Beginning with the equilibrium of a first-price auction, we have constructed an incentive-compatible direct selling mechanism whose truth-telling equilibrium results in the same ex-post assignment of the object to bidders and the same ex-post payments by them. In particular, it results in the same ex-post revenue for the seller. Moreover, this method of constructing a direct mechanism is quite general. Indeed, beginning with the equilibrium of any of the four standard auctions, we can similarly construct an incentive-compatible direct selling mechanism that yields the same ex-post assignment of the object to bidders and the same ex-post payments by them. (You are asked to do this in an exercise.)

In effect, we have shown that each of the four standard auctions is equivalent to some incentive-compatible direct selling mechanism. Because of this, we can now gain insight into the former by studying the latter.

9.3.1 INCENTIVE-COMPATIBLE DIRECT SELLING MECHANISMS

Consider an incentive-compatible direct selling mechanism with probability assignment functions $p_i(\cdot)$ and cost functions $c_i(\cdot)$, $i = 1, \ldots, N$. By incentive compatibility, each bidder must find it optimal to report his true value given that all other bidders do so. Let us consider the implications of this.

Suppose that bidder i's value is v_i and he considers reporting value r_i. If all other bidders report their values truthfully, then bidder i's expected payoff is

$$u_i(r_i, v_i) = \int_0^1 \cdots \int_0^1 (p_i(r_i, v_{-i})v_i - c_i(r_i, v_{-i}))f_{-i}(v_{-i})dv_{-i},$$

where $f_{-i}(v_{-i}) = f(v_1)\ldots f(v_{i-1})f(v_{i+1})\ldots f(v_N)$ and $dv_{-i} = dv_1 \ldots dv_{i-1}dv_{i+1}\ldots dv_N$.

For every $r_i \in [0, 1]$, let

$$\bar{p}_i(r_i) = \int_0^1 \cdots \int_0^1 p_i(r_i, v_{-i})f_{-i}(v_{-i})dv_{-i}$$

and

$$\bar{c}_i(r_i) = \int_0^1 \cdots \int_0^1 c_i(r_i, v_{-i})f_{-i}(v_{-i})dv_{-i}.$$

Therefore, $\bar{p}_i(r_i)$ is the probability that i receives the object when he reports r_i and $\bar{c}_i(r_i)$ is i's expected payment when he reports r_i, with both of these being conditional on all others reporting truthfully. Consequently, bidder i's expected payoff when his value is v_i and he reports it to be r_i can be written as

$$u_i(r_i, v_i) = \bar{p}_i(r_i)v_i - \bar{c}_i(r_i), \tag{9.9}$$

when all other bidders report their values truthfully.

So, the mechanism is incentive-compatible if and only if for every v_i, $u_i(r_i, v_i)$ is maximized in r_i at $r_i = v_i$; i.e., $u_i(v_i, v_i) \geq u_i(r_i, v_i)$ for all $r_i \in [0, 1]$.

The following result is very useful. It completely characterizes incentive-compatible direct selling mechanisms.

THEOREM 9.5 ***Incentive-Compatible Direct Selling Mechanisms***

A direct selling mechanism $(p_i(\cdot), c_i(\cdot))_{i=1}^N$ is incentive-compatible if and only if for every bidder i

(i) $\bar{p}_i(v_i)$ *is nondecreasing in* v_i *and,*
(ii) $\bar{c}_i(v_i) = \bar{c}_i(0) + \bar{p}_i(v_i)v_i - \int_0^{v_i} \bar{p}_i(x)\,dx$, *for every* $v_i \in [0, 1]$.

Proof: Suppose the mechanism is incentive-compatible. We must show that (i) and (ii) hold.

To see that (i) holds, note that by incentive compatibility, for all $r_i, v_i \in [0, 1]$,

$$\bar{p}_i(r_i)v_i - \bar{c}_i(r_i) = u_i(r_i, v_i) \le u_i(v_i, v_i) = \bar{p}_i(v_i)v_i - \bar{c}_i(v_i).$$

Adding and subtracting $\bar{p}_i(v_i)r_i$ to the right-hand side, this implies

$$\bar{p}_i(r_i)v_i - \bar{c}_i(r_i) \le [\bar{p}_i(v_i)r_i - \bar{c}_i(v_i)] + \bar{p}_i(v_i)(v_i - r_i).$$

But a careful look at the term in square brackets reveals that it is $u_i(v_i, r_i)$, bidder i's expected payoff from reporting v_i when her true value is r_i. By incentive compatibility, this must be no greater than $u_i(r_i, r_i)$, her payoff when she reports her true value, r_i. Consequently,

$$\begin{aligned}\bar{p}_i(r_i)v_i - \bar{c}_i(r_i) &\le [\bar{p}_i(v_i)r_i - \bar{c}_i(v_i)] + \bar{p}_i(v_i)(v_i - r_i)\\ &\le u_i(r_i, r_i) + \bar{p}_i(v_i)(v_i - r_i)\\ &= [\bar{p}_i(r_i)r_i - \bar{c}_i(r_i)] + \bar{p}_i(v_i)(v_i - r_i).\end{aligned}$$

That is,

$$\bar{p}_i(r_i)v_i - \bar{c}_i(r_i) \le [\bar{p}_i(r_i)r_i - \bar{c}_i(r_i)] + \bar{p}_i(v_i)(v_i - r_i),$$

which, when rewritten, becomes

$$(\bar{p}_i(v_i) - \bar{p}_i(r_i))(v_i - r_i) \ge 0.$$

So, when $v_i > r_i$, it must be the case that $\bar{p}_i(v_i) \ge \bar{p}_i(r_i)$. We conclude that $\bar{p}_i(\cdot)$ is nondecreasing. Hence, (i) holds.

To see that (ii) holds, note that because bidder i's expected payoff must be maximized when he reports truthfully, the derivative of $u_i(r_i, v_i)$ with respect to r_i must be zero when $r_i = v_i$.[14] Computing this derivative yields

$$\frac{\partial u_i(r_i, v_i)}{\partial r_i} = \bar{p}_i'(r_i)v_i - \bar{c}_i'(r_i),$$

[14]We are ignoring two points here. The first is whether $u_i(r_i, v_i)$ is in fact differentiable in r_i. Although it need not be everywhere differentiable, incentive compatibility implies that it must be differentiable almost everywhere and that the analysis we shall conduct can be made perfectly rigorous. We will not pursue these details here. The second point we ignore is the first-order condition at the two noninterior values $v_i = 0$ or 1. Strictly speaking, the derivatives at these boundary points need not be zero. But there is no harm in this because these two values each occur with probability zero.

and setting this to zero when $r_i = v_i$ yields

$$\bar{c}_i'(v_i) = \bar{p}_i'(v_i)v_i. \tag{9.10}$$

Because v_i was arbitrary, (9.10) must hold for every $v_i \in [0, 1]$. Consequently,

$$\begin{aligned}\bar{c}_i(v_i) - \bar{c}_i(0) &= \int_0^{v_i} \bar{c}_i'(x)\,dx \\ &= \int_0^{v_i} \bar{p}_i'(x)x\,dx \\ &= \bar{p}_i(v_i)v_i - \int_0^{v_i} \bar{p}_i(x)\,dx,\end{aligned}$$

where the first equality follows from the fundamental theorem of calculus, the second from (9.10), and the third from integration by parts. Consequently, for every bidder i and every $v_i \in [0, 1]$,

$$\bar{c}_i(v_i) = \bar{c}_i(0) + \bar{p}_i(v_i)v_i - \int_0^{v_i} \bar{p}_i(x)\,dx, \tag{9.11}$$

proving (ii).

We must now show the converse. So, suppose that (i) and (ii) hold. We must show that $u_i(r_i, v_i)$ is maximized in r_i when $r_i = v_i$. To see this, note that substituting (ii) into (9.9) yields

$$u_i(r_i, v_i) = \bar{p}_i(r_i)v_i - \left[\bar{c}_i(0) + \bar{p}_i(r_i)r_i - \int_0^{r_i} \bar{p}_i(x)\,dx\right]. \tag{9.12}$$

This can be rewritten as

$$u_i(r_i, v_i) = -\bar{c}_i(0) + \int_0^{v_i} \bar{p}_i(x)\,dx - \left\{\int_{r_i}^{v_i} (\bar{p}_i(x) - \bar{p}_i(r_i))\,dx\right\},$$

where this expression is valid whether $r_i \leq v_i$ or $r_i \geq v_i$.[15] Because by (i) $\bar{p}_i(\cdot)$ is nondecreasing, the integral in curly brackets is nonnegative for all r_i and v_i. Consequently,

$$u_i(r_i, v_i) \leq -\bar{c}_i(0) + \int_0^{v_i} \bar{p}_i(x)\,dx. \tag{9.13}$$

But, by (9.12), the right-hand-side of (9.13) is equal to $u_i(v_i, v_i)$. Consequently,

$$u_i(r_i, v_i) \leq u_i(v_i, v_i),$$

so that $u_i(r_i, v_i)$ is indeed maximized in r_i when $r_i = v_i$. ∎

[15] Recall the convention in mathematics that when $a < b$, $\int_b^a f(x)\,dx = -\int_a^b f(x)\,dx$.

Part (ii) of Theorem 9.5 says that if a direct mechanism is incentive-compatible there must be a connection between the probability assignment functions and the cost functions. In particular, it says that once the probability assignment function has been chosen and once a bidder's expected cost conditional on having value zero is chosen, the remainder of the expected cost function is chosen as well. To put it differently, under incentive compatibility a bidder's expected payment conditional on his value is completely determined by his expected payment when his value is zero and his probability assignment function. This observation is essential for understanding the following result.

THEOREM 9.6 ***Revenue Equivalence***

If two incentive-compatible direct selling mechanisms have the same probability assignment functions and every bidder with value zero is indifferent between the two mechanisms, then the two mechanisms generate the same expected revenue for the seller.

Proof: The seller's expected revenue is

$$\begin{aligned}
R &= \int_0^1 \cdots \int_0^1 \sum_{i=1}^N c_i(v_1, \ldots, v_N) f(v_1) \ldots f(v_N)\, dv_1 \ldots dv_N \\
&= \sum_{i=1}^N \int_0^1 \cdots \int_0^1 c_i(v_1, \ldots, v_N) f(v_1) \ldots f(v_N)\, dv_1 \ldots dv_N \\
&= \sum_{i=1}^N \int_0^1 \left[\int_0^1 \cdots \int_0^1 c_i(v_i, v_{-i}) f_{-i}(v_{-i}) dv_{-i} \right] f_i(v_i)\, dv_i \\
&= \sum_{i=1}^N \int_0^1 \bar{c}_i(v_i) f_i(v_i)\, dv_i \\
&= \sum_{i=1}^N \int_0^1 \left[\bar{c}_i(0) + \bar{p}_i(v_i) v_i - \int_0^{v_i} \bar{p}_i(x)\, dx \right] f_i(v_i)\, dv_i \\
&= \sum_{i=1}^N \int_0^1 \left[\bar{p}_i(v_i) v_i - \int_0^{v_i} \bar{p}_i(x)\, dx \right] f_i(v_i)\, dv_i + \sum_{i=1}^N \bar{c}_i(0),
\end{aligned}$$

where the fourth equality follows from the definition of $\bar{c}_i(v_i)$ and the fifth equality follows from (9.11).

Consequently, the seller's expected revenue depends only on the probability assignment functions and the amount bidders expect to pay when their values are zero. Because a bidder's expected payoff when his value is zero is completely determined by his expected payment when his value is zero, the desired result follows. ∎

The revenue equivalence theorem provides an explanation for the apparently coincidental equality of expected revenue among the four standard auctions. We now see that this follows because, with symmetric bidders, each of the four standard auctions has the same probability assignment function (i.e., the object is assigned to the bidder with the highest value), and in each of the four standard auctions a bidder with value zero receives expected utility equal to zero.

The revenue equivalence theorem is very general and allows us to add additional auctions to the list of those yielding the same expected revenue as the four standard ones. For example, a first-price, all-pay auction, in which the highest among all sealed bids wins but *every* bidder pays an amount equal to his bid, also yields the same expected revenue under bidder symmetry as the four standard auctions. You are asked to explore this and other auctions in the exercises.

9.3.2 EFFICIENCY

Before closing this section, we briefly turn our attention to the allocative properties of the four standard auctions. As we have already noted several times, each of these auctions allocates the object to the bidder who values it most. That is, each of these auctions is efficient. In the case of the Dutch and the first-price auctions, this result relies on bidder symmetry. Without symmetry, different bidders in a first-price auction, say, will employ different strictly increasing bidding functions. Consequently, if one bidder employs a lower bidding function than another, then the one may have a higher value yet be outbid by the other.

9.4 REVENUE-MAXIMIZATION: AN APPLICATION OF MECHANISM DESIGN

By now we understand very well the four standard auctions, their equilibria, their expected revenue, and the relation among them. But do these auctions, each generating the same expected revenue (under bidder symmetry), maximize the seller's expected revenue? Or is there a better selling mechanism for the seller? If there is a better selling mechanism what form does it take? Do the bidders submit sealed bids? Do they bid sequentially? What about a combination of the two? Is an auction the best selling mechanism?

Apparently, finding a revenue-maximizing selling mechanism is likely to be a difficult task. Given the freedom to choose any selling procedure, where do we start? The key observation is to recall how we were able to construct an incentive-compatible direct selling mechanism from the equilibrium of a first-price auction, and how the outcome of the first-price auction was exactly replicated in the direct mechanism's truth-telling equilibrium. As it turns out, the same type of construction can be applied to any selling procedure. That is, given an arbitrary selling procedure and a Nash equilibrium in which each bidder employs a strategy mapping his value into payoff-maximizing behavior under that selling procedure, we can construct an equivalent incentive-compatible direct selling mechanism. The requisite probability assignment and cost functions map each vector of values to the probabilities and costs that each bidder would experience according to the equilibrium strategies in the original selling procedure. So constructed, this direct selling mechanism is incentive-compatible and yields the same (probabilistic) assignment of the object and the same expected costs to each bidder as well as the same expected revenue to the seller.

Consequently, if some selling procedure yields the seller expected revenue equal to R, then so too does some incentive-compatible direct selling mechanism. But this means that *no selling mechanism among all conceivable selling mechanisms yields more revenue* *for the seller than the revenue-maximizing, incentive-compatible direct selling mechanism.* We can, therefore, restrict our search for a revenue-maximizing selling procedure to the (manageable) set of incentive-compatible direct selling mechanisms. In this way, we have simplified our problem considerably while losing nothing.

9.4.1 INDIVIDUAL RATIONALITY

There is one additional restriction we must now consider. Because participation by the bidders is entirely voluntary, no bidder's expected payoff can be negative given his value. Otherwise, whenever he has that value, he will simply not participate in the selling mechanism. Thus, we must restrict attention to incentive-compatible direct selling mechanisms that are **individually rational**, i.e., that yield each bidder, regardless of his value, a nonnegative expected payoff in the truth-telling equilibrium.

Now, in an incentive-compatible mechanism bidder i with value v_i will receive expected payoff $u_i(v_i, v_i)$ in the truth-telling equilibrium. So, an incentive-compatible direct selling mechanism is individually rational if this payoff is always nonnegative, i.e., if

$$u_i(v_i, v_i) = \bar{p}_i(v_i)v_i - \bar{c}_i(v_i) \geq 0 \text{ for all } v_i \in [0, 1].$$

However, by incentive compatibility, (ii) of Theorem 9.5 tells us that

$$\bar{c}_i(v_i) = \bar{c}_i(0) + \bar{p}_i(v_i)v_i - \int_0^{v_i} \bar{p}_i(x)\,dx, \text{ for every } v_i \in [0, 1].$$

Consequently, an incentive-compatible direct selling mechanism is individually rational if and only if

$$u_i(v_i, v_i) = \bar{p}_i(v_i)v_i - \bar{c}_i(v_i) = -\bar{c}_i(0) + \int_0^{v_i} \bar{p}_i(x)\,dx \geq 0 \text{ for every } v_i \in [0, 1],$$

which clearly holds if and only if

$$\bar{c}_i(0) \leq 0. \tag{9.14}$$

Consequently, an incentive-compatible direct selling mechanism is individually rational if and only if each bidder's expected cost when his value is zero is nonpositive.

9.4.2 AN OPTIMAL SELLING MECHANISM

We have now reduced the task of finding the optimal selling mechanism to maximizing the seller's expected revenue among all individually rational, incentive-compatible direct selling mechanisms, $p_i(\cdot)$ and $c_i(\cdot)$, $i = 1, \ldots, N$. Because Theorem 9.5 characterizes all incentive-compatible selling mechanisms, and because an incentive-compatible direct selling mechanism is individually rational if and only if $\bar{c}_i(0) \leq 0$, our task has been reduced to solving the following problem: Choose a direct selling mechanism $p_i(\cdot)$, $c_i(\cdot)$,

$i = 1, \ldots, N$, to maximize

$$R = \sum_{i=1}^{N} \int_0^1 \left[\bar{p}_i(v_i)v_i - \int_0^{v_i} \bar{p}_i(x)\,dx \right] f_i(v_i)\,dv_i + \sum_{i=1}^{N} \bar{c}_i(0)$$

subject to

(i) $\bar{p}_i(v_i)$ is nondecreasing in v_i,
(ii) $\bar{c}_i(v_i) = \bar{c}_i(0) + \bar{p}_i(v_i)v_i - \int_0^{v_i} \bar{p}_i(x)\,dx$, for every $v_i \in [0, 1]$,
(iii) $\bar{c}_i(0) \le 0$,

where the expression for the seller's expected revenue follows from incentive compatibility precisely as in the proof of Theorem 9.6.

It will be helpful to rearrange the expression for the seller's expected revenue.

$$\begin{aligned} R &= \sum_{i=1}^{N} \int_0^1 \left[\bar{p}_i(v_i)v_i - \int_0^{v_i} \bar{p}_i(x)\,dx \right] f_i(v_i)\,dv_i + \sum_{i=1}^{N} \bar{c}_i(0) \\ &= \sum_{i=1}^{N} \left[\int_0^1 \bar{p}_i(v_i)v_i f_i(v_i)\,dv_i - \int_0^1 \int_0^{v_i} \bar{p}_i(x) f_i(v_i)\,dx dv_i \right] + \sum_{i=1}^{N} \bar{c}_i(0). \end{aligned}$$

By interchanging the order of integration in the iterated integral (i.e., from $dxdv_i$ to $dv_i dx$), we obtain

$$\begin{aligned} R &= \sum_{i=1}^{N} \left[\int_0^1 \bar{p}_i(v_i)v_i f_i(v_i)\,dv_i - \int_0^1 \int_x^1 \bar{p}_i(x) f_i(v_i)\,dv_i dx \right] + \sum_{i=1}^{N} \bar{c}_i(0) \\ &= \sum_{i=1}^{N} \left[\int_0^1 \bar{p}_i(v_i)v_i f_i(v_i)\,dv_i - \int_0^1 \bar{p}_i(x)(1 - F_i(x))\,dx \right] + \sum_{i=1}^{N} \bar{c}_i(0). \end{aligned}$$

By replacing the dummy variable of integration, x, by v_i, this can be written equivalently as

$$\begin{aligned} R &= \sum_{i=1}^{N} \left[\int_0^1 \bar{p}_i(v_i)v_i f_i(v_i)\,dv_i - \int_0^1 \bar{p}_i(v_i)(1 - F_i(v_i))\,dv_i \right] + \sum_{i=1}^{N} \bar{c}_i(0) \\ &= \sum_{i=1}^{N} \int_0^1 \bar{p}_i(v_i) \left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right] f_i(v_i)\,dv_i + \sum_{i=1}^{N} \bar{c}_i(0). \end{aligned}$$

Finally, recalling that

$$\bar{p}_i(r_i) = \int_0^1 \cdots \int_0^1 p_i(r_i, v_{-i}) f_{-i}(v_{-i})\,dv_{-i}$$

we may write

$$R = \sum_{i=1}^{N} \int_0^1 \cdots \int_0^1 p_i(v_1, \ldots, v_N) \left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right] f_1(v_1) \ldots f_N(v_N)\, dv_1 \ldots dv_N + \sum_{i=1}^{N} \bar{c}_i(0),$$

or

$$R = \int_0^1 \cdots \int_0^1 \left\{ \sum_{i=1}^{N} p_i(v_1, \ldots, v_N) \left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right] \right\} f_1(v_1) \ldots f_N(v_N)\, dv_1 \ldots dv_N + \sum_{i=1}^{N} \bar{c}_i(0). \tag{9.15}$$

So, our problem is to maximize (9.15) subject to the constraints (i)–(iii) above. For the moment, let's concentrate on the first term in (9.15), namely

$$\int_0^1 \cdots \int_0^1 \left\{ \sum_{i=1}^{N} p_i(v_1, \ldots, v_N) \left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right] \right\} f_1(v_1) \ldots f_N(v_N)\, dv_1 \ldots dv_N. \tag{9.16}$$

Clearly, (9.16) would be maximized if the term in curly brackets were maximized for each vector of values $v_1, \ldots, v_N$. Now, because the $p_i(v_1, \ldots, v_N)$ are nonnegative and sum to one or less, the $N + 1$ numbers $p_1(v_1, \ldots, v_N), \ldots, p_N(v_1, \ldots, v_N)$, $1 - \sum_{i=1}^{N} p_i(v_1, \ldots, v_N)$ are nonnegative and sum to one. So, the sum above in curly brackets, which can be rewritten as

$$\sum_{i=1}^{N} p_i(v_1, \ldots, v_N) \left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right] + \left(1 - \sum_{i=1}^{N} p_i(v_1, \ldots, v_N) \right) \cdot 0,$$

is just a weighted average of the $N + 1$ numbers

$$\left[v_1 - \frac{1 - F_1(v_1)}{f_1(v_1)} \right], \ldots, \left[v_N - \frac{1 - F_N(v_N)}{f_N(v_N)} \right], \ 0.$$

But then the sum in curly brackets can be no larger than the largest of these bracketed terms if one of them is positive, and no larger than zero if all of them are negative. Suppose now that no two of the bracketed terms are equal to one another. Then, if we define

$$p_i^*(v_1, \ldots, v_N) = \begin{cases} 1, & \text{if } v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} > \max\left(0, v_j - \frac{1 - Fj(v_j)}{f_j(v_j)}\right) \text{ for all } j \neq i, \\ 0, & \text{otherwise,} \end{cases} \tag{9.17}$$

it must be the case that

$$\sum_{i=1}^{N} p_i(v_1, \ldots, v_N)\left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right] \le \sum_{i=1}^{N} p_i^*(v_1, \ldots, v_N)\left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right].$$

Therefore, if the bracketed terms are distinct with probability one, we will have

$$\begin{aligned} R &= \int_0^1 \cdots \int_0^1 \left\{\sum_{i=1}^{N} p_i(v_1, \ldots, v_N)\left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right]\right\} f_1(v_1) \ldots f_N(v_N)\, dv_1 \ldots dv_N \\ &\quad + \sum_{i=1}^{N} \bar{c}_i(0) \\ &\le \int_0^1 \cdots \int_0^1 \left\{\sum_{i=1}^{N} p_i^*(v_1, \ldots, v_N)\left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right]\right\} f_1(v_1) \ldots f_N(v_N)\, dv_1 \ldots dv_N \\ &\quad + \sum_{i=1}^{N} \bar{c}_i(0), \end{aligned}$$

for all incentive-compatible direct selling mechanisms $p_i(\cdot), c_i(\cdot)$. For the moment, then, let's assume that the bracketed terms are distinct with probability one. We will introduce an assumption on the bidders' distributions that guarantees this shortly.[16]

Because constraint (iii) implies that each $\bar{c}_i(0) \le 0$, we can also say that for all incentive-compatible direct selling mechanisms $p_i(\cdot), c_i(\cdot)$, the seller's revenue can be no larger than the following upper bound:

$$R \le \int_0^1 \cdots \int_0^1 \left\{\sum_{i=1}^{N} p_i^*(v_1, \ldots, v_N)\left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right]\right\} f_1(v_1) \ldots f_N(v_N)\, dv_1 \ldots dv_N. \tag{9.18}$$

We will now construct an incentive-compatible direct selling mechanism that *achieves* this upper bound. Consequently, this mechanism will maximize the seller's revenue, and so will be optimal for the seller.

To construct this optimal mechanism, let the probability assignment functions be the $p_i^*(v_1, \ldots, v_N)$, $i = 1, \ldots, N$, in (9.17). To complete the mechanism, we must define cost functions $c_i^*(v_1, \ldots, v_N)$, $i = 1, \ldots, N$. But constraint (ii) requires that for each v_i, bidder i's expected cost and probability of receiving the object, $\bar{c}_i^*(v_i)$ and $\bar{p}_i^*(v_i)$, be related as follows

$$\bar{c}_i^*(v_i) = \bar{c}_i^*(0) + \bar{p}_i^*(v_i)v_i - \int_0^{v_i} \bar{p}_i^*(x)\, dx.$$

[16]The assumption is given in (9.22).

Now, because the $\bar{c}_i^*$ and $\bar{p}_i^*$ are averages of the c_i^* and p_i^*, this required relationship between *averages* will hold if it holds for *each and every* vector of values $v_1, \ldots, v_N$. That is, (ii) is guaranteed to hold if we define the c_i^* as follows: For every $v_1, \ldots, v_N$,

$$c_i^*(v_1, \ldots, v_N) = c_i^*(0, v_{-i}) + p_i^*(v_1, \ldots, v_N)v_i - \int_0^{v_i} p_i^*(x, v_{-i})\,dx. \tag{9.19}$$

To complete the definition of the cost functions and to satisfy constraint (iii), we shall set $c_i^*(0, v_{-i}, \ldots, v_{-i}) = 0$ for all i and all $v_2, \ldots, v_n$. So, our candidate for a revenue-maximizing, incentive-compatible direct selling mechanism is as follows: For every $i = 1, \ldots, N$ and every $v_1, \ldots, v_N$

$$p_i^*(v_1, \ldots, v_N) = \begin{cases} 1, & \text{if } v_i - \frac{1-F_i(v_i)}{f_i(v_i)} > \max\left(0, v_j - \frac{1-F_j(v_j)}{f_j(v_j)}\right) \text{ for all } j \neq i, \\ 0, & \text{otherwise;} \end{cases} \tag{9.20}$$

and

$$c_i^*(v_1, \ldots, v_N) = p_i^*(v_1, \ldots, v_N)v_i - \int_0^{v_i} p_i^*(x, v_{-i})\,dx. \tag{9.21}$$

By construction, this mechanism satisfies constraints (ii) and (iii), and it achieves the upper bound for revenues in (9.18). To see this, simply substitute the p_i^* into (9.15) and recall that by construction $\bar{c}_i^*(0) = 0$ for every i. The result is that the seller's revenues are

$$R = \int_0^1 \cdots \int_0^1 \left\{ \sum_{i=1}^N p_i^*(v_1, \ldots, v_N)\left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right]\right\} f_1(v_1) \ldots f_N(v_N)\,dv_1 \ldots dv_N,$$

their maximum possible value.

So, if we can show that our mechanism's probability assignment functions defined in (9.20) satisfy constraint (i), then this mechanism will indeed be the solution we are seeking.

Unfortunately, the p_i^* as defined in (9.20) need not satisfy (i). To ensure that they do, we need to restrict the distributions of the bidders' values. Consider, then, the following assumption: For every $i = 1, \ldots, N$

$$v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \text{ is } \textit{strictly increasing} \text{ in } v_i. \tag{9.22}$$

This assumption is satisfied for a number of distributions, including the uniform distribution. Moreover, you are asked to show in an exercise that it holds whenever each F_i is any convex function, not merely that of the uniform distribution.[17] Note that in addition to ensuring that (i) holds, this assumption also guarantees that the numbers

[17]When this assumption fails, the mechanism we have constructed here is not optimal. One can nevertheless construct the optimal mechanism, but we shall not do so here. Thus, the additional assumption we are making here is only for simplicity's sake.

$v_1 - (1 - F_1(v_1))/f_1(v_1), \ldots, v_N - (1 - F_N(v_N))/f_N(v_N)$ are distinct with probability one, a requirement that we earlier employed but had left unjustified until now.

Let us now see why (9.22) implies that (i) is satisfied. Consider some bidder i and some fixed vector of values, v_{-i}, for the other bidders. Now, suppose that $\bar{v}_i > \underline{v}_i$ and that $p_i^*(\underline{v}_i, v_{-i}) = 1$. Then, by the definition of p_i^*, it must be the case that $\underline{v}_i - (1 - F_i(\underline{v}_i))/f_i(\underline{v}_i)$ is positive and strictly greater than $v_j - (1 - Fj(v_j))/f_j(v_j)$ for all $j \neq i$. Consequently, because $v_i - (1 - F_i(v_i))/f_i(v_i)$ is strictly increasing it must also be the case that $\bar{v}_i - (1 - F_i(\underline{v}_i))/f_i(\underline{v}_i)$ is both positive and strictly greater than $v_j - (1 - Fj(v_j))/f_j(v_j)$ for all $j \neq i$, which means that $p_i^*(\bar{v}_i, v_{-i}) = 1$. Thus, we have shown that if $p_i^*(v_i, v_{-i}) = 1$, then $p_i^*(v_i', v_{-i}) = 1$ for all $v_i' > v_i$. But because p_i^* takes on either the value 0 or 1, $p_i^*(v_i, v_{-i})$ is nondecreasing in v_i for every v_{-i}. This in turn implies that $\bar{p}_i^*(v_i)$ is nondecreasing in v_i, so that constraint (i) is indeed satisfied.

In the end then, our hard work has paid off handsomely. We can now state the following.

THEOREM 9.7 ***An Optimal Selling Mechanism***

If N bidders have independent private values with bidder i's value drawn from the continuous positive density f_i satisfying (9.22), then the direct selling mechanism defined in (9.20) and (9.21) yields the seller the largest possible expected revenue.

9.4.3 A CLOSER LOOK AT THE OPTIMAL SELLING MECHANISM

Let's see if we can simplify the description of the optimal selling mechanism by studying its details. There are two parts to the mechanism, the manner in which it allocates the object—the p_i^*—and the manner in which it determines payments—the c_i^*.

The allocation portion of the optimal mechanism is straightforward. Given the reported values $v_1, \ldots, v_N$, the object is given to the bidder i whose $v_i - (1 - F_i(v_i))/f_i(v_i)$ is strictly highest and positive. Otherwise, the seller keeps the object. But it is worth a little effort to try to interpret this allocation scheme.

What we shall argue is that $v_i - (1 - F_i(v_i))/f_i(v_i)$ represents the marginal revenue, $MR_i(v_i)$, that the seller obtains from increasing the probability that the object is assigned to bidder i when his value is v_i. To see this without too much notation we shall provide an intuitive argument. Consider the effect of increasing the probability that the object is awarded bidder i when his value is v_i. This enables the seller to increase the cost to v_i so as to leave his utility unchanged. Because the density of v_i is $f_i(v_i)$, the seller's revenue increases at the rate $v_i f_i(v_i)$ as a result of this change. On the other hand, incentive compatibility forces a connection between the probability that the good is assigned to bidder i with value v_i and the cost assessed to all higher values $v_i' > v_i$. Indeed, according to constraint (ii), increasing the probability that lower values receive the object reduces one-for-one the cost that all higher values can be assessed. Because there is a mass of $1 - F_i(v_i)$ values above v_i, this total reduction in revenue is $1 - F_i(v_i)$. So, altogether the seller's revenues increase by $v_i f_i(v_i) - (1 - F_i(v_i))$. But this is the total effect due to the density $f_i(v_i)$ of values equal to v_i. Consequently, the marginal revenue associated with each v_i is $MR_i(v_i) = v_i - (1 - F_i(v_i))/f_i(v_i)$.

The allocation rule now makes perfect sense. If $MR_i(v_i) > MR_j(v_j)$. The seller can increase revenue by reducing the probability that the object is assigned to bidder j and increasing the probability that it is assigned to bidder i. Clearly then, the seller maximizes her revenue by assigning all probability (i.e., probability one) to the bidder with the highest $MR_i(v_i)$, so long as it is positive. If all the marginal revenues are negative, the seller does best by reducing all of the bidders' probabilities to zero, i.e., the seller keeps the object.

The payment portion of the mechanism is a little less transparent. To get a clearer picture of what is going on, suppose that when the (truthfully) reported values are $v_1, \ldots, v_N$, bidder i does not receive the object, i.e., that $p_i^*(v_i, v_{-i}) = 0$. What must bidder i pay according to the mechanism? The answer, according to (9.21), is

$$\begin{aligned} c_i^*(v_i, v_{-i}) &= p_i^*(v_i, v_{-i})v_i - \int_0^{v_i} p_i^*(x, v_{-i})\,dx \\ &= 0 \cdot v_i - \int_0^{v_i} p_i^*(x, v_{-i})\,dx. \end{aligned}$$

But recall that, by virtue of assumption (9.22), $p_i^*(\cdot, v_{-i})$ is nondecreasing. Consequently, because $p_i^*(v_i, v_{-i}) = 0$, it must be the case that $p_i^*(x, v_{-i}) = 0$ for every $x \leq v_i$. Hence the integral above must be zero so that

$$c_i^*(v_i, v_{-i}) = 0.$$

So, we have shown that according to the optimal mechanism, if bidder i does not receive the object, he pays nothing.

Suppose now that bidder i does receive the object, i.e., that $p_i^*(v_i, v_{-i}) = 1$. According to (9.21), he then pays

$$\begin{aligned} c_i^*(v_i, v_{-i}) &= p_i^*(v_i, v_{-i})v_i - \int_0^{v_i} p_i^*(x, v_{-i})\,dx \\ &= v_i - \int_0^{v_i} p_i^*(x, v_{-i})\,dx. \end{aligned}$$

Now, because p_i^* takes on the value 0 or 1, is nondecreasing and continuous from the left in i's value, and $p_i^*(v_i, v_{-i}) = 1$, there must be a largest value for bidder i, $r_i^* < v_i$, such that $p_i^*(r_i^*, v_{-i}) = 0$. Note that r_i^* will generally depend on v_{-i} so it would be more explicit to write $r_i^*(v_{-i})$. Note then that by the very definition of $r_i^*(v_{-i})$, $p_i^*(x, v_{-i})$ is equal to 1 for every $x > r_i^*(v_{-i})$, and is equal to 0 for every $x \leq r_i^*(v_{-i})$. But this means that

$$\begin{aligned} c_i^*(v_i, v_{-i}) &= v_i - \int_{r_i^*(v_{-i})}^{v_i} 1\,dx \\ &= v_i - (v_i - r_i^*(v_{-i})) \\ &= r_i^*(v_{-i}). \end{aligned}$$

So, when bidder i wins the object, *he pays a price, $r_i^*(v_{-i})$, that is independent of his own* *reported value*. Moreover, the price he pays is the maximum value he could have reported, given the others' reported values, without receiving the object.

Putting all of this together, we may rephrase the revenue-maximizing selling mechanism defined by (9.20) and (9.21) in the following manner.

THEOREM 9.8 ***The Optimal Selling Mechanism Simplified***

If N bidders have independent private values with bidder i's value drawn from the continuous positive density f_i and each $v_i - (1 - F_i(v_i))/f_i(v_i)$ is strictly increasing, then the following direct selling mechanism yields the seller the largest possible expected revenue:

For each reported vector of values, $v_1, \ldots, v_N$, the seller assigns the object to the bidder i whose $v_i - (1 - F_i(v_i))/f_i(v_i)$ is strictly largest and positive. If there is no such bidder, the seller keeps the object and no payments are made. If there is such a bidder i, then only this bidder makes a payment to the seller in the amount r_i^*, where $r_i^* - (1 - F_i(r_i^*))/f_i(r_i^*) = 0$ or $\max_{j \neq i} v_j - (1 - F_j(v_j)/f_j(v_j)$, whichever is largest. Bidder i's payment, r_i^*, is, therefore, the largest value he could have reported, given the others' reported values, without receiving the object.

As we know, this mechanism is incentive-compatible. That is, truth-telling is a Nash equilibrium. But, in fact, the incentive to tell the truth in this mechanism is much stronger than this. In this mechanism it is, in fact, a *dominant strategy* for each bidder to report his value truthfully to the seller; even if the other bidders do not report their values truthfully, bidder i can do no better than to report his value truthfully to the seller. You are asked to show this in one of the exercises.

One drawback of this mechanism is that to implement it, the seller must know the distributions, F_i, from which the bidders' values are drawn. This is in contrast to the standard auctions that the seller can implement without any bidder information whatsoever. Yet there is a connection between this optimal mechanism and the four standard auctions that we now explore.

9.4.4 EFFICIENCY, SYMMETRY, AND COMPARISON TO THE FOUR STANDARD AUCTIONS

In the optimal selling mechanism, the object is not always allocated efficiently. Sometimes the bidder with the highest value does not receive the object. In fact, there are *two ways* that inefficiency can occur in the optimal selling mechanism. First, the outcome can be inefficient because the seller sometimes keeps the object, even though his value for it is zero and all bidders have positive values. This occurs when every bidder i's value v_i is such that $v_i - (1 - F_i(v_i))/f_i(v_i) \leq 0$. Second, even when the seller does assign the object to one of the bidders, it might not be assigned to the bidder with the highest value. To see this, consider the case of two bidders, 1 and 2. If the bidders are asymmetric, then for some $v \in [0, 1]$, $v - (1 - F_1(v))/f_1(v) \neq v - (1 - F_2(v))/f_2(v)$. Indeed, let us suppose that for this particular value, v, $v - (1 - F_1(v))/f_1(v) > v - (1 - F_2(v))/f_2(v) > 0$. Consequently, when both bidders' values are v, bidder 1 will receive the object. But, by continuity, even if bidder 1's value falls slightly to $v' < v$, so long as v' is close enough to v, the inequality

$v' - (1 - F_1(v'))/f_1(v') > v - (1 - F_2(v))/f_2(v) > 0$ will continue to hold. Hence, bidder 1 will receive the object even though his value is strictly below that of bidder 2.

The presence of inefficiencies is not surprising. After all, the seller is a monopolist seeking maximal profits. In Chapter 4, we saw that a monopolist will restrict output below the efficient level so as to command a higher price. The same effect is present here. But, because there is only one unit of an indivisible object for sale, the seller here restricts supply by sometimes keeping the object, depending on the vector of reports. But this accounts for only the first kind of inefficiency. The second kind of inefficiency that arises here did not occur in our brief look at monopoly in Chapter 4. The reason is that there we assumed that the monopolist was unable to distinguish one consumer from another. Consequently, the monopolist had to charge all consumers the same price. Here, however, we are assuming that the seller *can* distinguish bidder i from bidder j and that the seller knows that i's distribution of values is F_i and that j's is F_j. This additional knowledge allows the monopolist to discriminate between the bidders, which leads to higher profits.

Let's now eliminate this second source of inefficiency by supposing that bidders are symmetric. Because the four standard auctions all yield the same expected revenue for the seller under symmetry, this will also allow us to compare the standard auctions with the optimal selling mechanism.

How does symmetry affect the optimal selling mechanism? If the bidders are symmetric, then $f_i = f$ and $F_i = F$ for every bidder i. Consequently, the optimal selling mechanism is as follows: If the vector of reported values is $(v_1, \ldots, v_N)$, the bidder i with the highest positive $v_i - (1 - F(v_i))/f(v_i)$ receives the object and pays the seller r_i^*, the largest value he could have reported, given the other bidder's reported values, without winning the object. If there is no such bidder i, the seller keeps the object and no payments are made.

But let's think about this for a moment. Because we are assuming that $v - (1 - F(v))/f(v)$ is strictly increasing in v, the object is actually awarded to the bidder i with the strictly highest value v_i, so long as $v_i - (1 - F_i(v_i))/f_i(v_i) > 0$—that is, so long as $v_i > \rho^* \in [0, 1]$, where

$$\rho^* - \frac{1 - F(\rho^*)}{f(\rho^*)} = 0. \tag{9.23}$$

(You are asked to show in an exercise that a unique such ρ^* is guaranteed to exist.)

Now, how large can bidder i's reported value be before he is awarded the object? Well, he does not get the object unless his reported value is strictly highest and strictly above ρ^*. So, the largest his report can be without receiving the object is the largest of the other bidders' values or ρ^*, whichever is larger. Consequently, when bidder i does receive the object he pays either ρ^* or the largest value reported by the other bidders, whichever is larger.

Altogether then, the optimal selling mechanism is as follows: The bidder whose reported value is strictly highest and strictly above ρ^* receives the object and pays the larger of ρ^* and the largest reported value of the other bidders.

Remarkably, this optimal direct selling mechanism can be mimicked by running a second-price auction with reserve price ρ^*. That is, an auction in which the bidder with the highest bid strictly above the reserve price wins and pays the second-highest bid or the reserve price, whichever is larger. If no bids are above the reserve price, the seller keeps the object and no payments are made. This is optimal because, just as in a standard second-price auction, it is a dominant strategy to bid one's value in a second-price auction with a reserve price.

This is worth highlighting.

THEOREM 9.9 ***An Optimal Auction Under Symmetry***

If N bidders have independent private values, each drawn from the same continuous positive density f, where $v-(1-F(v))/f(v)$ is strictly increasing, then a second price auction with reserve price ρ^* satisfying $\rho^*-(1-F(\rho^*))/f(\rho^*)=0$, maximizes the seller's expected revenue.

You might wonder about the other three standard auctions. Will adding an appropriate reserve price render these auctions optimal for the seller too? The answer is yes, and this is left for you to explore in the exercises.

So, we have now come full circle. The four standard auctions—first-price, second-price, Dutch, and English—all yield the same revenue under symmetry. Moreover, by supplementing each by an appropriate reserve price, the seller maximizes his expected revenue. Is it any wonder then that these auctions are in such widespread use? We will leave you with that thought.

9.5 Exercises

9.1 Show that the bidding strategy in (9.5) is strictly increasing.

9.2 Show in two ways that the symmetric equilibrium bidding strategy of a first price auction with N symmetric bidders each with values distributed according to F, can be written as

$$\hat{b}(v)=v-\int_0^v\left(\frac{F(x)}{F(v)}\right)^{N-1}dx.$$

For the first way, use our solution from the text and apply integration by parts. For the second way, use the fact that $F^{N-1}(r)(v-\hat{b}(r))$ is maximized in r when $r=v$ and then apply the envelope theorem to conclude that $d(F^{N-1}(v)(v-\hat{b}(v))/dv=F^{N-1}(v)$; now integrate both sides from 0 to v.

9.3 This exercise will guide you through the proof that the bidding function in (9.5) is in fact a symmetric equilibrium of the first-price auction.

(a) Recall from (9.2) that

$$\frac{du(r,v)}{dr}=(N-1)F^{N-2}(r)f(r)(v-\hat{b}(r))-F^{N-1}(r)\hat{b}'(r).$$

Using (9.3), show that

$$\begin{aligned}\frac{du(r,v)}{dr}&=(N-1)F^{N-2}(r)f(r)(v-\hat{b}(r))-(N-1)F^{N-2}(r)f(r)(r-\hat{b}(r))\\&=(N-1)F^{N-2}(r)f(r)(v-r).\end{aligned}$$

(b) Use the result in part (a) to conclude that $du(r,v)/dr$ is positive when $r<v$ and negative when $r>v$, so that $u(r,v)$ is maximized when $r=v$.

9.4 Throughout this chapter we have assumed that both the seller and all bidders are risk neutral. In this question, we shall explore the consequences of risk aversion on the part of bidders.

There are N bidders participating in a first-price auction. Each bidder's value is independently drawn from $[0,1]$ according to the distribution function F, having continuous and strictly positive density f. If a bidder's value is v and he wins the object with a bid of $b < v$, then his von Neumann–Morgenstern utility is $(v-b)^{\frac{1}{\alpha}}$, where $\alpha \geq 1$ is fixed and common to all bidders. Consequently, the bidders are risk averse when $\alpha > 1$ and risk neutral when $\alpha = 1$. (Do you see why?) Given the risk aversion parameter α, let $\hat{b}_\alpha(v)$ denote the (symmetric) equilibrium bid of a bidder when his value is v. The following parts will guide you toward finding $\hat{b}_\alpha(v)$ and uncovering some of its implications.

(a) Let $u(r, v)$ denote a bidder's expected utility from bidding $\hat{b}_\alpha(r)$ when his value is v, given that all other bidders employ $\hat{b}_\alpha(\cdot)$. Show that

$$u(r, v) = F^{N-1}(r)(v - \hat{b}_\alpha(r))^{\frac{1}{\alpha}}.$$

Why must $u(r, v)$ be maximized in r when $r = v$?

(b) Use part (a) to argue that

$$[u(r, v)]^\alpha = [F^\alpha(r)]^{N-1}(v - \hat{b}_\alpha(r))$$

must be maximized in r when $r = v$.

(c) Use part (b) to argue that a first-price auction with the $N-1$ risk averse bidders above whose values are each independently distributed according to $F(v)$, is equivalent to a first-price auction with $N-1$ *risk neutral* bidders whose values are each independently distributed according to $F^\alpha(v)$. Use our solution for the risk neutral case (see Exercise 9.2 above) to conclude that

$$\hat{b}_\alpha(v) = v - \int_0^v \left(\frac{F(x)}{F(v)}\right)^{\alpha(N-1)} dx.$$

(d) Prove that $\hat{b}_\alpha(v)$ is strictly increasing in $\alpha \geq 1$. Does this make sense? Conclude that as bidders become more risk averse, the seller's revenue from a first-price auction increases.

(e) Use part (d) and the revenue equivalence result for the standard auctions in the risk neutral case to argue that when bidders are risk averse as above, a first-price auction raises *more* revenue for the seller than a second-price auction. Hence, these two standard auctions no longer generate the same revenue when bidders are risk averse.

(f) What happens to the seller's revenue as the bidders become infinitely risk averse (i.e., as $\alpha \to \infty$)?

9.5 In a private values model, argue that it is a weakly dominant strategy for a bidder to bid her value in a second-price auction even if the joint distribution of the bidders' values exhibits correlation.

9.6 Use the equilibria of the second-price, Dutch, and English auctions to construct incentive-compatible direct selling mechanisms for each of them in which the ex-post assignment of the object to bidders as well as their ex-post payments to the seller are unchanged.

9.7 In a first-price, all-pay auction, the bidders simultaneously submit sealed bids. The highest bid wins the object and *every* bidder pays the seller the amount of his bid. Consider the independent private values model with symmetric bidders whose values are each distributed according to the distribution function F, with density f.

(a) Find the unique symmetric equilibrium bidding function. Interpret.

(b) Do bidders bid higher or lower than in a first-price auction?

(c) Find an expression for the seller's expected revenue.

(d) Both with and without using the revenue equivalence theorem, show that the seller's expected revenue is the same as in a first-price auction.

9.8 Suppose there are just two bidders. In a second-price, all-pay auction, the two bidders simultaneously submit sealed bids. The highest bid wins the object and both bidders pay the *second-highest* bid.

(a) Find the unique symmetric equilibrium bidding function. Interpret.

(b) Do bidders bid higher or lower than in a first-price, all-pay auction?

(c) Find an expression for the seller's expected revenue.

(d) Both with and without using the revenue equivalence theorem, show that the seller's expected revenue is the same as in a first-price auction.

9.9 Consider the following variant of a first-price auction. Sealed bids are collected. The highest bidder pays his bid but receives the object only if the outcome of the toss of a fair coin is heads. If the outcome is tails, the seller keeps the object and the high bidder's bid. Assume bidder symmetry.

(a) Find the unique symmetric equilibrium bidding function. Interpret.

(b) Do bidders bid higher or lower than in a first-price auction?

(c) Find an expression for the seller's expected revenue.

(d) Both with and without using the revenue equivalence theorem, show that the seller's expected revenue is exactly half that of a standard first-price auction.

9.10 Suppose all bidders' values are uniform on $[0, 1]$. Construct a revenue-maximizing auction. What is the reserve price?

9.11 Consider again the case of uniformly distributed values on $[0, 1]$. Is a first-price auction with the same reserve price as in the preceding question optimal for the seller? Prove your claim using the revenue equivalence theorem.

9.12 Suppose the bidders' values are *i.i.d.*, each according to a uniform distribution on $[1, 2]$. Construct a revenue-maximizing auction for the seller.

9.13 Suppose there are N bidders with independent private values where bidder i's value is uniform on $[a_i, b_i]$. Show that the following is a revenue-maximizing, incentive-compatible direct selling mechanism. Each bidder reports his value. Given the reported values $v_1, \ldots, v_N$, bidder i wins the object if v_i is strictly larger than the $N - 1$ numbers of the form $b_i/2 + \max(0, v_j - b_j/2)$ for $j \neq i$. Bidder i then pays the seller an amount equal to the largest of these $N - 1$ numbers. All other bidders pay nothing.

9.14 A drawback of the direct mechanism approach is that the seller must know the distribution of the bidders' values to compute the optimal auction. The following exercise provides an optimal auction that is distribution-free for the case of two asymmetric bidders, 1 and 2, with independent private values. Bidder i's strictly positive and continuous density of values on $[0, 1]$ is f_i with distribution F_i. Assume throughout that $v_i - (1 - F_i(v_i))/f_i(v_i)$ is strictly increasing for $i = 1, 2$.

The auction is as follows. In the first stage, the bidders each simultaneously submit a sealed bid. Before the second stage begins, the bids are publicly revealed. In the second stage, the bidders must simultaneously declare whether they are willing to purchase the object at the other bidder's revealed sealed bid. If one of them says "yes" and the other "no," then the "yes" transaction is carried out. If they both say "yes" or both say "no," then the seller keeps the object and no payments are made. Note that the seller can run this auction without knowing the bidders' value distributions.

(a) Consider the following strategies for the bidders: In the first stage, when her value is v_i, bidder $i \neq j$ submits the sealed bid $b_i^*(v_i) = b_i$, where b_i solves

$$b_i - \frac{1 - F_j(b_i)}{f_j(b_i)} = \max\left(0, v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right).$$

(Although such a b_i need not always exist, it will always exist if the functions $v_1 - (1 - F_1(v_1))/f_1(v_1)$ and $v_2 - (1 - F_2(v_2))/f_2(v_2)$ have the same range. So, assume this is the case.)

In the second stage each bidder says "yes" if and only if her value is above the other bidder's first-stage bid.

Show that these strategies constitute an equilibrium of this auction. (Also, note that while the seller need not know the distribution of values, each bidder needs to know the distribution of the other bidder's values to carry out her strategy. Hence, this auction shifts the informational burden from the seller to the bidders.)

(b) (i) Show that in this equilibrium the seller's expected revenues are maximized.
(ii) Is the outcome always efficient?

(c) (i) Show that it is also an equilibrium for each bidder to bid his value and then to say "yes" if and only if his value is above the other's bid.
(ii) Is the outcome always efficient in this equilibrium?

(d) Show that the seller's revenues are *not* maximal in this second equilibrium.

(e) Unfortunately, this auction possesses *many* equilibria. Choose any two strictly increasing functions $g_i : [0, 1] \to \mathbb{R}_2$ $i = 1, 2$, with a common range. Suppose in the first stage that bidder $i \neq j$ with value v_i bids $\bar{b}_i(v_i) = b_i$, where b_i solves $g_j(b_i) = g_i(v_i)$ and says "yes" in the second stage if and only if his value is strictly above the other bidder's bid. Show that this is an equilibrium of this auction. Also, show that the outcome is always efficient if and only if $g_i = g_j$.

9.15 Show that condition (9.22) is satisfied when each F_i is a convex function. Is convexity of F_i necessary?

9.16 Consider the independent private values model with N possibly asymmetric bidders. Suppose we restrict attention to *efficient* individually rational, incentive-compatible direct selling mechanisms; i.e., those that always assign the object to the bidder who values it most.

(a) What are the probability assignment functions?

(b) What then are the cost functions?

(c) What cost functions among these maximize the seller's revenue?

(d) Conclude that among *efficient* individually rational, incentive-compatible direct selling mechanisms, a second-price auction maximizes the seller's expected revenue. (What about the other three standard auction forms?)

9.17 Call a direct selling mechanism $p_i(\cdot)$, $c_i(\cdot)$, $i = 1, \ldots, N$ *deterministic* if the p_i take on only the values 0 or 1.

(a) Assuming independent private values, show that for every incentive-compatible deterministic direct selling mechanism whose probability assignment functions, $p_i(v_i, v_{-i})$, are nondecreasing in v_i for every v_{-i}, there is another incentive-compatible direct selling mechanism with the *same* probability assignment functions (and, hence, deterministic as well) whose cost functions have the property that a bidder pays only when he receives the object and when he does win, the amount that he pays is independent of his reported value. Moreover, show that the new mechanism can be chosen so that the seller's expected revenue is the same as that in the old.

(b) How does this result apply to a first-price auction with symmetric bidders, wherein a bidder's payment depends on his bid?

(c) How does this result apply to an all-pay, first-price auction with symmetric bidders wherein bidders pay whether or not they win the auction?

9.18 Show that it is a weakly dominant strategy for each bidder to report his value truthfully in the optimal direct mechanism we derived in this chapter.

9.19 Under the assumption that each bidder's density, f_i, is continuous and strictly positive and that each $v_i - (1 - F_i(v_i))/f_i(v_i)$ is strictly increasing,

(a) Show that the optimal selling mechanism entails the seller keeping the object with strictly positive probability.

(b) Show that there is precisely one $\rho^* \in [0, 1]$ satisfying $\rho^* - (1 - F(\rho^*))/f(\rho^*) = 0$.

9.20 Show that when the bidders are symmetric, the first-price, Dutch, and English auctions all are optimal for the seller once an appropriate reserve price is chosen. Indeed, show that the optimal reserve price is the same for all four of the standard auctions.