2 Hidden Information, Screening

In this chapter we focus on the basic static adverse selection problem, with one principal facing one agent who has private information on her "type," that is, her preferences or her intrinsic productivity. This problem was first formally analyzed by Mirrlees (1971). We first explain how to solve such problems when the agent can be of only two types, a case that already allows us to obtain most of the key insights from adverse selection models. We do so by looking at the problem of nonlinear pricing by a monopolistic seller who faces a buyer with unknown valuation for his product.

We then move on to other applications, still in the case where the informed party can be of only two types: credit rationing, optimal income taxation, implicit labor contracts, and regulation. This is only a partial list of economic issues where adverse selection matters. Nevertheless, these are all important economic applications that have made a lasting impression on the economics profession. For each of them we underline both the economic insights and the specificities from the point of view of contract theory.

In the last part of the chapter we extend the analysis to more than two types, returning to monopoly pricing. We especially emphasize the continuum case, which is easier to handle. This extension allows us to stress which results from the two-type case are general and which ones are not. The methods we present will also be helpful in tackling multiagent contexts, in particular in Chapter 7.

2.1 The Simple Economics of Adverse Selection

Adverse selection naturally arises in the following context, analyzed first by Mussa and Rosen (1978), and subsequently by Maskin and Riley (1984a): Consider a transaction between a buyer and a seller, where the seller does not know perfectly how much the buyer is willing to pay for a good. Suppose, in addition, that the seller sets the terms of the contract. The buyer's preferences are represented by the utility function

$$u(q, T, \theta) = \int_0^q P(x, \theta)dx - T$$

where q is the number of units purchased, T is the total amount paid to the seller, and $P(x, \theta)$ is the inverse demand curve of a buyer with preference characteristics θ. Throughout this section we shall consider the following special and convenient functional form for the buyer's preferences:

$$u(q, T, \theta) = \theta v(q) - T$$

where $v(0) = 0$, $v'(q) > 0$, and $v''(q) < 0$ for all q. The characteristics θ are private information to the buyer. The seller knows only the distribution of θ, $F(\theta)$.

Assuming that the seller's unit production costs are given by $c > 0$, his profit from selling q units against a sum of money T is given by

$$\pi = T - cq$$

The question of interest here is, What is the best, that is, the profit-maximizing, pair (T, q) that the seller will be able to induce the buyer to choose? The answer to this question will depend on the information the seller has on the buyer's preferences. We treat in this section the case where there are only two types of buyers: $\theta \in \{\theta_L, \theta_H\}$, with $\theta_H > \theta_L$. The consumer is of type θ_L with probability $\beta \in [0, 1]$ and of type θ_H with probability $(1 - \beta)$. The probability β can also be interpreted as the proportion of consumers of type θ_L.

2.1.1 First-Best Outcome: Perfect Price Discrimination

To begin with, suppose that the seller is perfectly informed about the buyer's characteristics. The seller can then treat each type of buyer separately and offer her a type-specific contract, that is, (T_i, q_i) for type θ_i $(i = H, L)$. The seller will try to maximize his profits subject to inducing the buyer to accept the proposed contract. Assume the buyer obtains a payoff of $\bar{u}$ if she does not take the seller's offer. In this case, the seller will solve

$$\max_{T_i, q_i} T_i - cq_i$$

subject to

$$\theta_i v(q_i) - T_i \geq \bar{u}$$

We can call this constraint the participation, or individual-rationality, constraint of the buyer. The solution to this problem will be the contract $(\bar{q}_i, \bar{T}_i)$ such that

$$\theta_i v'(\bar{q}_i) = c$$

and

$$\theta_i v(\bar{q}_i) = \bar{T}_i + \bar{u}$$

Intuitively, without adverse selection, the seller finds it optimal to maximize total surplus by having the buyer select a quantity such that marginal utility

equals marginal cost, and then setting the payment so as to appropriate the full surplus and leave no rent to the buyer above $\bar{u}$. Note that in a market context, $\bar{u}$ could be endogenized, but here we shall treat it as exogenous and normalize it to 0.

Without adverse selection, the total profit of the seller is thus

$$\beta(T_L - cq_L) + (1 - \beta)(T_H - cq_H)$$

and the optimal contract maximizes this profit subject to the participation constraints for the two types of buyer. Note that it can be implemented by *type-specific two-part tariffs,* where the buyer is allowed to buy as much as she wants of the good at unit price c provided she pays a type-specific fixed fee equal to $\theta_i v(\bar{q}_i) - c\bar{q}_i$.

The idea that, without adverse selection, the optimal contract will maximize total surplus while participation constraints will determine the way in which it is shared is a very general one.[1] This ceases to be true in the presence of adverse selection.

2.1.2 Adverse Selection, Linear Pricing, and Simple Two-Part Tariffs

If the seller cannot observe the type of the buyer anymore, he has to offer the same contract to everybody. The contract set is potentially large, since it consists of the set of functions $T(q)$. We first look at two simple contracts of this kind.

2.1.2.1 Linear Pricing

The simplest contract consists in *traditional linear pricing,* which is a situation where the seller's contract specifies only a price P. Given this contract the buyer chooses q to maximize

$$\theta_i v(q) - Pq, \quad \text{where } i = L, H$$

From the first-order conditions

$$\theta_i v'(q) = P$$

we can derive the demand functions of each type:[2]

1. This idea also requires that surplus be freely transferable across individuals, which will not be the case if some individuals face financial resource constraints.

2. The assumed concavity of $v(.)$ ensures that there is a unique solution to the first-order conditions.

$q_i = D_i(P)$

The buyer's net surplus can now be written as follows:

$$S_i(P) = \theta_i v[D_i(P)] - PD_i(P)$$

Let

$$D(P) \equiv \beta D_L(P) + (1-\beta)D_H(P)$$

$$S(P) \equiv \beta S_L(P) + (1-\beta)S_H(P)$$

With linear pricing the seller's problem is the familiar monopoly pricing problem, where the seller chooses P to solve

$$\max_P (P-c)D(P)$$

and the monopoly price is given by

$$P_m = c - \frac{D(P)}{D'(P)}$$

In this solution we have both positive rents for the buyers [$S(P) > 0$] and inefficiently low consumption, that is, $\theta_i v'(q) = P > c$, since the seller can make profits only by setting a price in excess of marginal cost and $D'(\cdot) < 0$. Note that, depending on the values of β, θ_L, and θ_H, it may be optimal for the seller to serve only the θ_H buyers. We shall, however, proceed under the assumption that it is in the interest of the seller to serve both markets.

Can the seller do better by moving away from linear pricing? He will be able to do so only if buyers cannot make arbitrage profits by trading in a secondary market: if arbitrage is costless, only linear pricing is possible, because buyers would buy at the minimum average price and then resell in the secondary market if they do not want to consume everything they bought.

2.1.2.2 Single Two-Part Tariff

In this subsection we shall work with the interpretation that there is only one buyer and that β is a probability measure. Under this interpretation there are no arbitrage opportunities open to the buyer. Therefore, a single two-part tariff (Z, P), where P is the unit price and Z the fixed fee, will improve upon linear pricing for the seller. Note first that for any given price P, the minimum fixed fee the seller will set is given by $Z = S_L(P)$. (This is

the maximum fee a buyer of type θ_L is willing to pay.) A type-θ_H buyer will always decide to purchase a positive quantity of q when charged a two-part tariff $T(q) = S_L(P) + Pq$, since $\theta_H > \theta_L$. If the seller decides to serve both types of customers and therefore sets $Z = S_L(P)$, he also chooses P to maximize

$$\max_P S_L(P) + (P-c)D(P)$$

The solution for P under this arrangement is given by

$$S_L'(P) + D(P) + (P-c)D'(P) = 0$$

which implies

$$P = c - \frac{D(P) + S_L'(P)}{D'(P)}$$

Now, by the envelope theorem, $S_L'(P) = -D_L(P)$, so that $D(P) + S_L'(P)$ is strictly positive; in addition, $D'(P) < 0$, so that $P > c$. Thus, if the seller decides to serve both types of customers [and therefore sets $Z = S_L(P)$], the first-best outcome cannot be achieved and underconsumption remains relative to the first-best outcome.[3] Another conclusion to be drawn from this simple analysis is that an optimal single two-part tariff contract is always preferred by the seller to an optimal linear pricing contract [since the seller can always raise his profits by setting $Z = S_L(P_m)$]. We can also observe the following: If P_m, P_d, and P_c, respectively, denote the monopoly price, the marginal price in an optimal single two-part tariff, and the (first-best efficient) competitive price, then $P_m > P_d > P_c = c$. To see this point, note that a small reduction in price from P_m has a second-order (negative) effect on monopoly profits $(P_m - c)D(P_m)$, by definition of P_m. But it has a first-order (positive) effect on consumer surplus, which increases by an amount proportional to the reduction in price. The first-order (positive) effect dominates, and, therefore, the seller is better off lowering the price from P_m when he can extract the buyer's surplus with the fixed fee $Z = S_L(P_m)$. Similarly, a small increase in price from P_c has a first-order (positive) effect on $(P_c - c)D(P_c)$, but a second-order (negative) effect on $S(P_c)$, by definition of P_c.

3. If he decides to set an even higher fixed fee and to price the type-θ_L buyer out of the market, he does not achieve the first-best outcome either; either way, the first-best solution cannot be attained under a single two-part tariff contract.

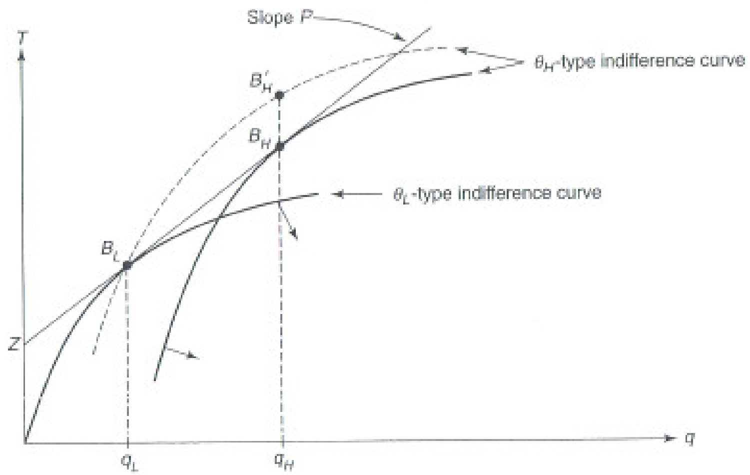

Figure 2.1
Two-Part Tariff Solution

An important feature of the optimal single two-part tariff solution is that the θ_H-type buyer strictly prefers the allocation B_H to B_L, as illustrated in Figure 2.1. As the figure also shows, by setting up more general contracts $C = [q, T(q)]$, the seller can do strictly better by offering the same allocation to the θ_L-type buyer, but offering, for example, some other allocation $B'_H \neq B_H$ to the θ_H-type buyer. Notice that at B'_H the seller gets a higher transfer $T(q)$ for the same consumption (this is particular to the example). Also, the θ_H buyer is indifferent between B_L and B'_H. These observations naturally raise the question of the form of optimal nonlinear pricing contract.

.1.3 Second-Best Outcome: Optimal Nonlinear Pricing

In this subsection we show that the seller can generally do better by offering more general nonlinear prices than a single two-part tariff. In the process we outline the basic methodology of solving for optimal contracts when the buyer's type is unknown. Since the seller does not observe the type of the buyer, he is forced to offer her a set of choices independent of her type. Without loss of generality, this set can be described as

$[q, T(q)]$; that is, the buyer faces a schedule from which she will pick the outcome that maximizes her payoff. The problem of the seller is therefore to solve

$$\max_{T(q)} \beta[T(q_L) - cq_L] + (1-\beta)[T(q_H) - cq_H]$$

subject to

$$q_i = \arg\max_q \theta_i v(q) - T(q) \quad \text{for } i = L, H$$

and

$$\theta_i v(q_i) - T(q_i) \geq 0 \quad \text{for } i = L, H$$

The first two constraints are the incentive-compatibility (IC) constraints, while the last two are participation or individual-rationality constraints (IR). This problem looks nontrivial to solve, since it involves optimization over a schedule $T(q)$ under constraints that themselves involve optimization problems. Such adverse selection problems can, however, be easily solved step-by-step as follows:

Step 1: Apply the revelation principle.

From Chapter 1 we can recall that without loss of generality we can restrict each schedule $T(q)$ to the pair of optimal choices made by the two types of buyers $\{[T(q_L), q_L]$ and $[T(q_H), q_H]\}$; this restriction also simplifies greatly the incentive constraints. Specifically, if we define $T(q_i) = T_i$ for $i = L, H$, then the problem can be rewritten as

$$\max_{T_i, q_i} \beta(T_L - cq_L) + (1-\beta)(T_H - cq_H)$$

subject to

$$\theta_H v(q_H) - T_H \geq \theta_H v(q_L) - T_L \qquad (ICH)$$

$$\theta_L v(q_L) - T_L \geq \theta_L v(q_H) - T_H \qquad (ICL)$$

$$\theta_H v(q_H) - T_H \geq 0 \qquad (IRH)$$

$$\theta_L v(q_L) - T_L \geq 0 \qquad (IRL)$$

The seller thus faces four constraints, two incentive constraints [(IC_i) means that the type-θ_i buyer should prefer her own allocation to the allocation of the other type of buyer] and two participation constraints

[(IR_i) means that the allocation that buyer of type θ_i chooses gives her a nonnegative payoff]. Step 1 has thus already greatly simplified the problem. We can now try to eliminate some of these constraints.

Step 2: Observe that the participation constraint of the "high" type will not bind at the optimum.

Indeed (IRH) will be satisfied automatically because of (IRL) and (ICH):

$$\theta_H v(q_H) - T_H \geq \theta_H v(q_L) - T_L \geq \theta_L v(q_L) - T_L \geq 0$$

where the inequality in the middle comes from the fact that $\theta_H > \theta_L$.

Step 3: Solve the relaxed problem without the incentive constraint that is satisfied at the first-best optimum.

The strategy now is to relax the problem by deleting one incentive constraint, solve the relaxed problem, and then check that it does satisfy this omitted incentive constraint. In order to choose which constraint to omit, consider the first-best problem. It involves efficient consumption and zero rents for both types of buyers, that is, $\theta_i v'(\bar{q}_i) = c$ and $\theta_i v(\bar{q}_i) = \bar{T}_i$. This outcome is not incentive compatible, because the θ_H buyer will prefer to choose $(\bar{q}_L, \bar{T}_L)$ rather than her own first-best allocation: while this inefficiently restricts her consumption, it allows her to enjoy a strictly positive surplus equal to $(\theta_H - \theta_L)\bar{q}_L$, rather than zero rents. Instead, type θ_L will not find it attractive to raise her consumption to the level $\bar{q}_H$: doing so would involve paying an amount $\bar{T}_H$ that exhausts the surplus of type θ_H and would therefore imply a negative payoff for type θ_L, who has a lower valuation for this consumption. In step 3, we thus choose to omit constraint (ICL). Note that the fact that only one incentive constraint will bind at the optimum is driven by the *Spence-Mirrlees single-crossing condition,* which can be written as

$$\frac{\partial}{\partial \theta}\left[-\frac{\partial u/\partial q}{\partial u/\partial T}\right] > 0$$

This condition means that the marginal utility of consumption (relative to that of money, which is here constant) rises with θ. Consequently, optimal consumption will have to rise with θ.

Step 4: Observe that the two remaining constraints of the relaxed problem will bind at the optimum.

Remember that we now look at the problem

$$\max_{T_i, q_i} \beta(T_L - cq_L) + (1-\beta)(T_H - cq_H)$$

subject to

$$\theta_H v(q_H) - T_H \geq \theta_H v(q_L) - T_L \qquad (ICH)$$

$$\theta_L v(q_L) - T_L \geq 0 \qquad (IRL)$$

In this problem, constraint (ICH) will bind at the optimum; otherwise, the seller can raise T_H until it does bind: this step leaves constraint (IRL) unaffected while improving the maximand. And constraint (IRL) will also bind; otherwise, the seller can raise T_L until it does bind: this step in fact relaxes constraint (ICH) while improving the maximand [it is here that having omitted (ICL) matters, since a rise in T_L could be problematic for this constraint].

Step 5: Eliminate T_L and T_H from the maximand using the two binding constraints, perform the unconstrained optimization, and then check that (ICL) is indeed satisfied.

Substituting for the values of T_L and T_H in the seller's objective function, we obtain the following unconstrained optimization problem:

$$\max_{q_L, q_H} \beta[\theta_L v(q_L) - cq_L] + (1-\beta)[\theta_H v(q_H) - cq_H - (\theta_H - \theta_L)v(q_L)]$$

The first term in brackets is the full surplus generated by the purchases of type θ_L, which the seller appropriates entirely because that type is left with zero rents. Instead, the second term in brackets is the full surplus generated by the purchases of type θ_H *minus her informational rent* $(\theta_H - \theta_L)v(q_L)$, which comes from the fact that she can "mimic" the behavior of the other type. This informational rent increases with q_L.

The following first-order conditions characterize the unique interior solution (q_L^*, q_H^*) to the relaxed program, if this solution exists:[4]

$$\theta_H v'(q_H^*) = c$$

$$\theta_L v'(q_L^*) = \frac{c}{1 - \left(\frac{1-\beta}{\beta}\frac{\theta_H - \theta_L}{\theta_L}\right)} > c$$

4. If the denominator of the second expression is not positive, then the optimal solution involves $q_L^* = 0$, while the other consumption remains determined by the first-order condition.

This interior solution implies $q_L^* < q_H^*$. One can then immediately verify that the omitted constraints are satisfied at the optimum $(q_i^*, T_i^*; i = L, H)$ given that (ICH) binds. Indeed,

$$\theta_H v(q_H^*) - T_H^* = \theta_H v(q_L^*) - T_L^* \qquad (ICH)$$

together with $\theta_L < \theta_H$ and $q_L^* < q_H^*$, implies

$$\theta_L v(q_H^*) - T_H^* \leq \theta_L v(q_L^*) - T_L^* \qquad (ICL)$$

We have therefore characterized the actual optimum. Two basic economic conclusions emerge from this analysis:

1. The second-best optimal consumption for type θ_H is the same as the first-best optimal consumption $(\bar{q}_H)$, but that of type θ_L is lower. Thus only the consumption of one of the two types is distorted in the second-best solution.

2. The type-θ_L buyer obtains a surplus of zero, while the other type obtains a strictly positive "informational" rent.

These two conclusions are closely related to each other: the consumption distortion for type θ_L is the result of the seller's attempt to reduce the informational rent of type θ_H. Since a buyer of type θ_H is more eager to consume, the seller can reduce that type's incentive to mimic type θ_L by cutting down on the consumption offered to type θ_L. By reducing type θ_H's incentives to mimic type θ_L, the seller can reduce the informational rent of (or, equivalently, charge a higher price to) type θ_H. Looking at the first-order conditions for q_i^* indicates that the size of the distortion, $\bar{q}_L - q_L^*$, is increasing in the potential size of the informational rent of type θ_H—as measured by the difference $(\theta_H - \theta_L)$—and decreasing in β. For β and $(\theta_H - \theta_L)$ large enough the denominator becomes negative. In that case the seller hits the constraint $q_L \geq 0$.

As the latter part of the chapter will show, what will remain true with more than two types is the inefficiently low consumption relative to the first best (except for the highest type: we will keep "efficiency at the top") and the fact that the buyer will enjoy positive informational rents (except for the lowest type). Before doing this extension, let us turn to other applications.

2.2 Applications

2.2.1 Credit Rationing

Adverse selection arises naturally in financial markets. Indeed, a lender usually knows less about the risk-return characteristics of a project than the borrower. In other words, a lender is in the same position as a buyer of a secondhand car:[5] it does not know perfectly the "quality" of the project it invests in. Because of this informational asymmetry, inefficiencies in the allocation of investment funds to projects may arise. As in the case of secondhand cars, these inefficiencies may take the form that "good quality" projects remain "unsold" or are denied credit. This type of inefficiency is generally referred to as "credit rationing." There is now an extensive literature on credit rationing. The main early contributions are Jaffee and Modigliani (1969), Jaffee and Russell (1976), Stiglitz and Weiss (1981), Bester (1985), and De Meza and Webb (1987). We shall illustrate the main ideas with a simple example where borrowers can be of two different types.

Consider a population of risk-neutral borrowers who each own a project that requires an initial outlay of $I = 1$ and yields a random return X, where $X \in \{R, 0\}$. Let $p \in [0, 1]$ denote the probability that $X = R$. Borrowers have no wealth and must obtain investment funds from an outside source. A borrower can be of two different types $i = s, r$, where s stands for "safe" and r for "risky." The borrower of type i has a project with return characteristics (p_i, R_i). We shall make the following assumptions:

A1: $p_i R_i = m$, with $m > 1$

A2: $p_s > p_r$ and $R_s < R_r$

Thus both types of borrowers have projects with the same expected return, but the risk characteristics of the projects differ. In general, project types may differ both in risk and return characteristics. It turns out that the early literature mainly emphasizes differences in risk characteristics.

A bank can offer to finance the initial outlay in exchange for a future repayment. Assume for simplicity that there is a single bank and excess

5. Akerlof (1970), in a pioneering contribution, has analyzed the role of adverse selection in markets by focusing in particular on used cars. See Chapter 13 on the role of adverse selection in markets more generally.