

Chapter 5

Producer Theory

Markets have two sides: consumers and producers. Up until now we have been studying the consumer side of the market. We now begin our study of the producer side of the market.

The basic unit of activity on the production side of the market is the firm. The task of the firm is take commodities and turn them into other commodities. The objective of the firm (in the neoclassical model) is to maximize profits. That is, the firm chooses the production plan from among all feasible plans that maximizes the profit earned on that plan. In the neoclassical (competitive) production model, the firm is assumed to be one firm among many others. Because of this (as in the consumer model), prices are exogenous in the neoclassical production model. Firms are unable to affect the prices of either their inputs or their outputs. Situations where the firm is able to affect the price of its output will be studied later under the headings of monopoly and oligopoly.

Our study of production will be divided into three parts: First, we will consider production from a purely technological point of view, characterizing the firm's set of feasible production plans in terms of its production set Y . Second, we will assume that the firm produces a single output using multiple inputs, and we will study its profit maximization and cost minimization problems using a production function to characterize its production possibilities. Finally, we will consider a special class of production models, where the firm's production function exhibits constant returns to scale.

5.1 Production Sets

Consider an economy with L commodities. The task of the firm is to change inputs into outputs. For example, if there are three commodities, and the firm uses 2 units of commodity one and 3 units of commodity two to produce 7 units of commodity three, we can write this production plan as $y = (-2, -3, 7)$, where, by convention, negative components mean that that commodity is an input and positive components mean that that commodity is an output. If the prices of the three commodities are $p = (1, 2, 2)$, then a firm that chooses this production plan earns profit of $\pi = p \cdot y = (1, 2, 2) \cdot (-2, -3, 7) = 6$.

Usually, we will let $y = (y_1, \dots, y_L)$ stand for a single production plan, and $Y \subset R^L$ stand for the set of all feasible production plans. The shape of Y is going to be driven by the way in which different inputs can be substituted for each other in the production process.

A typical production set (for the case of two commodities) is shown in MWG Figure 5.B.1. The set of points below the curved line represents all feasible production plans. Notice that in this situation, either commodity 1 can be used to produce commodity 2 ($y_1 < 0, y_2 > 0$), commodity 2 can be used to produce commodity 1 ($y_1 > 0, y_2 < 0$), nothing can be done ($y_1 = y_2 = 0$) or both commodities can be used without producing an output, ($y_1 < 0, y_2 < 0$). Of course, the last situation is wasteful – if it has the option of doing nothing, then no profit-maximizing firm would ever choose to use inputs and incur cost without producing any output. While this is true, it is useful for certain technical reasons to allow for this possibility.

Generally speaking, it will not be profit maximizing for the firm to be wasteful. What is meant by wasteful? Consider a point y inside Y in Figure 5.B.1. If y is not on the northeast frontier of Y then it is wasteful. Why? Because if this is the case the firm can either produce more output using the same amount of input or the same output using less input. Either way, the firm would earn higher profit. Because of this it is useful to have a mathematical representation for the frontier of Y . The tool we have for this is called the **transformation function**, $F(y)$, and we call the northeast frontier of the production set the **production frontier**. The transformation function is such that

$$\begin{aligned} F(y) &= 0 \text{ if } y \text{ is on the frontier} \\ &< 0 \text{ if } y \text{ is in the interior of } Y \\ &> 0 \text{ if } y \text{ is outside of } Y. \end{aligned}$$

Thus the transformation function implicitly defines the frontier of Y . Thus if $F(y) < 0$, y represents

some sort of waste, although $F(\cdot)$ tells us neither the form of the waste nor the magnitude.

The transformation function can be used to investigate how various inputs can be substituted for each other in the production process. For example, consider a production plan \bar{y} such that $F(\bar{y}) = 0$. The slope of the transformation frontier with respect to commodities i and j is given by:

$$\frac{\partial y_i}{\partial y_j} = -\frac{F_j(\bar{y})}{F_i(\bar{y})}.$$

The absolute value of the right-hand side of this expression, $\frac{F_j(\bar{y})}{F_i(\bar{y})}$, is known as the **marginal rate of transformation of good j for good i at \bar{y}** (MRT_{ji}).

$$MRT_{ji} = \frac{F_j(\bar{y})}{F_i(\bar{y})}$$

It tells how much you must increase the (net) usage of factor j if you decrease the net usage of factor i in order to remain on the transformation frontier. It is important to note that factor usage can be either positive or negative in this model. In either case, increasing factor usage means moving to the right on the number line. Thus if you are using -5 units of an input, going to -4 units of that input is an increase, as far as the MRT is concerned.

For example, suppose we are currently at $y = (-2, 7)$, that $F(-2, 7) = 0$, and that we are interested in MRT_{12} , the marginal rate of transformation of good 1 for good 2. $MRT_{12} = \frac{F_1(-2, 7)}{F_2(-2, 7)}$. Now, if the net usage of good 1 increases, say from -2 to -1 , then we move out of the production set, and $F(-1, 7) > 0$. Hence $F_1(-2, 7) > 0$. If we increase commodity 3 a small amount, say to 8, we also move out of the production set, and $F(-2, 8) > 0$. So, $MRT_{12} > 0$.

The slope of the transformation frontier asks how much the net usage of factor 2 must be changed if the net usage of factor 1 is increased. Thus it is a negative number. This is why the slope of the transformation frontier is negative when comparing an input and an output, but the MRT is positive.

5.1.1 Properties of Production Sets

There are a number of properties that can be attributed to production sets. Some of these will be assumed for all production sets, and some will only apply to certain production sets.

Properties of All Production Sets

Here, I will list properties that we assume all production sets satisfy.

1. **Y is nonempty.** (If Y is empty, then we have nothing to talk about).
2. **Y is closed.** A set is closed if it contains its boundary. We need Y to be closed for technical reasons. Namely, if a set does not contain its boundary, then if you try to maximize a function (such as profit) subject to the constraint that the production plan be in Y , it may be that there is no optimal plan – the firm will try to be as close to the boundary as possible, but no matter how close it is, it could always be a little closer.
3. **No free lunch.** This means that you cannot produce output without using any inputs. In other words, any feasible production plan y must have at least one negative component. Beside violating the laws of physics, if there were a “free lunch,” then the firm could make infinite profit just by replicating the free lunch point over and over, which makes the firm’s profit maximization problem impossible to solve.
4. **Free disposal.** This means that the firm can always throw away inputs if it wants. The meaning of this is that for any point in Y , points that use less of all components are also in Y . Thus if $y \in Y$, any point below and to the left is also in Y (in the two dimensional model). The idea is that you can throw away as much as you want, and while you have to buy the commodities you are throwing away, you don’t have to pay anybody to dispose of it for you. So, if there are two commodities, grapes and wine, and you can make 10 cases of wine from 1 ton of grapes, then it is also feasible for you to make 10 cases of wine from 2 tons of grapes (by just throwing one of ton of grapes away) or 5 cases of wine from 1 ton of grapes (by just throwing 5 cases of wine away at the end), or 5 cases of wine from 2 tons of grapes (by throwing 1 ton of grapes and 5 cases of wine away at the end). The upshot is that the production set is unbounded as you move down and to the left (in the standard diagram). Again, you should think of this as mostly a technical assumption.¹

Properties of Some Production Sets

The following properties may or may not hold for a particular production set. Usually, if the production set has one of these properties, it will be easier to choose the profit-maximizing bundle.

¹Basically, we are going to want to look for the tangency between the firm’s profit function and Y in solving the firm’s profit maximization problem. If Y is bounded below, i.e., free disposal doesn’t hold, then we may find a tangency below Y , which will not be profit maximizing. Thus assuming free disposal has something to do with second-order conditions. We want to make sure that the point that satisfies the first-order conditions is really a maximum.

1. **Irreversibility.** Irreversibility says that the production process cannot be undone. That is, if $y \in Y$ and $y \neq 0$, then $-y \notin Y$. Actually, the laws of physics imply that all production processes are irreversible. You may be able to turn gold bars into jewelry and then jewelry back into gold bars, but in either case you use energy. So, this process is not really reversible. The reason why I call this a property of some production sets is that, even though it is true of all real technologies, we often do not need to invoke irreversibility in order to get the results we are after. And, since we don't like to make assumptions we don't need, in many cases it won't be stated. On the other hand, you should beware of results that hinge on the reversibility of a technology, for the physics reasons I mentioned earlier.
2. **Possibility of inaction.** This property says that $0 \in Y$. That is, the firm can choose to do nothing. Of course, if it does so, it earns zero profit. This is good because it allows us to only consider positive profit production plans in the firm's optimization problem. Situations where $0 \notin Y$ arise when the firm has a **fixed factor of production**. For example, if the firm is obligated to pay rent on its factory, then it cannot do nothing. The cost of an unavoidable fixed factor of production is sometimes called a **sunk cost**. A production set with a fixed factor is illustrated in MWG Figure 5.B.3a. As you may remember, however, whether a cost item is fixed or not depends on the relevant time frame. Put another way, if the firm waits long enough, its lease will expire and it will no longer have to pay its rent. Thus while inaction is not a possibility in the short run, it is a possibility in the long run, provided that the long run is sufficiently long.

Global Returns to Scale Properties

The following properties refer to the entire production set Y . However, it is important to point out that many production sets will exhibit none of these. But, they are useful for talking about parts of production sets as well, and the idea of returns to scale in this abstract setting is a little different than the one you may be used to. So, it is worth working through them.

1. **Nonincreasing returns to scale.** Y exhibits nonincreasing returns to scale if any feasible production plan $y \in Y$ can be scaled down: $ay \in Y$ for $a \in [0, 1]$. What does that mean? A technology that exhibits increasing returns to scale is one that becomes more productive (on average) as the size of the output grows. Thus if you want to rule out increasing returns to scale, you want to rule out situations that require the firm to become more productive at

higher levels of production. The way to do this is to require that any feasible production plan y can be scaled down to ay , for $a \in [0, 1]$. If this holds, then the feasibility of y does not depend on the fact that it involves a large scale of production and the firm gets more efficient at large scale.

2. **Nondecreasing returns to scale.** Y exhibits nondecreasing returns to scale if any feasible production plan $y \in Y$ can be scaled up: $ay \in Y$ for $a \geq 1$. Decreasing returns to scale is a situation where the firm grows less productive at higher levels of output. Thus if we want to rule out the case of decreasing returns to scale, we must rule out the case where y is feasible, but if that same production plan were scaled larger to ay , it would no longer be feasible because the firm is less productive at the higher scale. Thus we require that $ay \in Y$ for $a \geq 1$.

- Note that if a firm has fixed costs, it may exhibit nondecreasing returns to scale but cannot exhibit nonincreasing returns to scale. See MWG Figure 5.B.6.

3. **Constant returns to scale.** Y exhibits constant returns to scale if it exhibits both non-increasing returns to scale and non-decreasing returns to scale at all \bar{y} . That is, for all $a \geq 0$, if $y \in Y$, then $ay \in Y$. Constant returns to scale means that the firm's productivity is independent of the level of production. Thus it means that any feasible production plan can either be scaled upward or downward.

- Constant returns to scale implies the possibility of inaction.

4. **Convexity.** If Y exhibits nonincreasing returns to scale, then Y is convex.

Note that the list of returns to scale properties is by no means exhaustive. In fact, most real technologies exhibit none of these. The "typical" technology that we think of is one that at first exhibits increasing returns to scale, and then exhibits decreasing returns to scale. This would be the case, for example, for a manufacturing firm whose factory size is fixed. At first, as output increases, its average productivity increases as it spreads the factory cost over more output. However, eventually the firm's output becomes larger than the factory is designed for. At this point, the firm's average productivity falls as the workers become crowded, machines become overworked, etc. A typical example of this type of technology is illustrated in MWG Figure 5.F.2.

So, while the previous definitions were global, we can also think of local versions. A technology exhibits nonincreasing returns at a point on the transformation frontier if the transformation frontier is locally concave there, and it exhibits nondecreasing returns at a point if the transformation function is locally convex there. Also note, decreasing returns to scale means that returns are nonincreasing and not constant – thus the transformation frontier is locally strictly concave. The opposite goes for increasing returns – the transformation frontier is locally strictly convex.

5.1.2 Profit Maximization with Production Sets

As we said earlier, the firm's objective is to maximize profit. Using the production plan approach we outlined earlier, the profit earned on production plan y is $p \cdot y$. Hence the firm's profit maximization problem (PMP) is given by:

$$\begin{aligned} & \max_y p \cdot y \\ \text{s.t. } & y \in Y. \end{aligned}$$

Since $Y = \{y | F(y) \leq 0\}$, this problem can be rewritten as:

$$\begin{aligned} & \max_y p \cdot y \\ \text{s.t. } & F(y) \leq 0. \end{aligned}$$

If $F()$ is differentiable, this problem can be solved using standard Lagrangian techniques. The graphical solution to this problem is depicted in Figure 5.1. The Lagrangian is:

$$L = p \cdot y - \lambda F(y)$$

which implies first-order conditions:

$$\begin{aligned} p_i &= \lambda F_i(y^*), \text{ for } i = 1, \dots, L \\ F(y^*) &\leq 0. \end{aligned}$$

As before, we can solve the first-order conditions for goods i and j in terms of λ and set them equal, yielding:

$$\begin{aligned} \frac{p_i}{F_i(y^*)} &= \lambda = \frac{p_j}{F_j(y^*)} \\ \frac{F_i(y^*)}{F_j(y^*)} &= \frac{p_i}{p_j}. \end{aligned} \tag{5.1}$$




Figure 5.1: The Profit Maximization Problem

Also, since we established that points strictly inside Y are wasteful, and wasteful behavior is not profit maximizing, we know that the constraint will bind. Hence $F(y^*) = 0$.

Condition 5.1 is a tangency condition, similar to the one we derived in the consumer's problem. It says that at the profit maximizing production plan y^* , the marginal rate of transformation between any two commodities is equal to the ratio of their prices.

Figure 5.1 presents the PMP. The profit isoquant is a downward sloping line with slope $-\frac{p_1}{p_2}$. We move the profit line up and to the right until we find the point of tangency. This is the point that satisfies the optimality condition 5.1. Note that unless Y is convex, the first-order conditions will not be sufficient for a maximum. Generally, second-order conditions will need to be checked. This can be seen in Figure 5.1, since there is a point of tangency that is not a maximum.

The solution to the profit maximization problem, $y(p)$, is called the firm's net supply function. The value function for the profit maximization problem, $\pi(p) = p \cdot y(p)$, is called the profit function.²

One thing to worry about is whether this solution exists. When Y is strictly convex and the production frontier becomes sufficiently flat (i.e., the firm experiences strongly decreasing returns) at some high level of production, it is not too hard to show that a solution exists.

Problems can arise when the firm's production set is not convex, i.e., its technology exhibits nondecreasing returns to scale. In this case, the production frontier is convex.³ If you go through the graphical analysis we did earlier, you'll discover that if the firm maximizes profit, it either produces nothing at all (if the output price is small relative to the input price), or else it can

²There is a direct correspondence to consumer theory. $y(p)$ is like $x(p, w)$, and $\pi(p)$ is like $v(p, w)$.

³Again, remember the distinction between a convex set and a convex function.

always increase profit by moving further out on its production frontier. Thus its production becomes infinite, as do its profits. This problem also arises when the firm's production function exhibits increasing returns to scale only over a region of its transformation frontier. It is fairly straightforward to show that the firm's profit maximizing point will never be on a part of the production frontier exhibiting increasing returns (formally, because the second order conditions will not hold). As a result, we will focus most of our attention on technologies that exhibit nonincreasing returns.

5.1.3 Properties of the Net Supply and Profit Functions



Just like the demand functions and indirect utility functions from consumer theory, the net supply function and profit function also have many properties worth knowing.

Let Y be the firm's production set, and suppose $y(p)$ solves the firm's profit maximization problem. Let $\pi(p) = p \cdot y(p)$ be the associated profit function. The net supply function $y(p)$ has the following properties:

1. **$y(p)$ is homogeneous of degree zero in p .** The reason for this is the same as always. The optimality conditions for the PMP involve a tangency condition, and tangencies are not affected by re-scaling all prices by the same amount.
2. **If Y is convex, $y(p)$ is a convex set. If Y is strictly convex, then $y(p)$ is a single point.** The reason for this is the same as in the UMP. But, notice that while it made sense to think of utility as being quasiconcave, it is not always reasonable to think of Y as being convex. Recall that a convex set Y corresponds to non-increasing returns. But, many real productive processes exhibit increasing returns, at least over some range. So, the convexity assumption rules out more important cases in profit maximization than quasiconcavity did in the consumer problem.

The profit function $\pi(p)$ has the following properties:

1. **$\pi(p)$ is homogenous of degree 1 in p .** Again, the reason for this is the same as the reason why $h(p, u)$ is homogenous of degree zero but $e(p, u)$ is homogeneous of degree 1. $\pi(p) = p \cdot y(p)$. $\pi(ap) = ap \cdot y(ap) = ap \cdot y(p) = a\pi(p)$. Since scaling all prices by $a > 0$ does not change relative prices, the optimal production plan does not change. However, multiplying all prices by a multiplies profit by a as well. Thus profit is homogenous of degree 1.

Figure 5.2: Convexity of $\pi(p)$ Figure 5.3: Convexity of $\pi(p)$

2. $\pi(p)$ is **convex**. The reason for this is the same as the reason why $e(p, u)$ is concave. That is, if the price of an output increases (or an input decreases) and the firm does not change its production plan, profit will increase linearly. However, since the firm will want to re-optimize at the new prices, it can actually do better. Profit will increase at a greater than linear rate. This is just what it means for a function to be convex. Figures 5.2 and 5.3 illustrate this for the cases where the price of an output increases and the price of an input increases, respectively. Note the difference is that profit is increasing in the price of an output but decreasing in the price of an input. To put it another way, consider two price

vectors p and p' , and $p^a = ap + (1 - a)p'$.

$$\begin{aligned} p^a \cdot y(p^a) &= ap \cdot y(p^a) + (1 - a)p' \cdot y(p^a) \\ &\leq ap \cdot y(p) + (1 - a)p' \cdot y(p') \\ &= a\pi(p) + (1 - a)\pi(p'). \end{aligned}$$

This provides a formal proof of convexity that mirrors the proof that $e(p, u)$ is concave.

The next two properties relate the net supply functions and the profit function. Note the similarity to the relationship between $h(p, u)$ and $e(p, u)$.⁴

1. **Hotelling's Lemma:** If $y(\bar{p})$ is single-valued at \bar{p} , then $\frac{\partial \pi(\bar{p})}{\partial p_i} = y_i(p)$. This follows directly from the envelope theorem. The Lagrangian of the PMP is

$$L = p \cdot y - \lambda F(y).$$

The envelope theorem says that

$$\frac{d\pi(\bar{p})}{dp_i} = \frac{\partial L}{\partial p_i} \Big|_{y=y(\bar{p})} = y_i(\bar{p}).$$

This is the direct analog to $\frac{\partial e(p, u)}{\partial p_i} = h_i(p, u)$. In other words, the increase in profit due to an increase in p_i is simply equal to the net usage of commodity i . Indirect effects (the effect of rearranging the production plan in response to the price change) can be ignored. If commodity i is an output, $y_i(\bar{p}) > 0$, and increasing p_i increases profit. On the other hand, if commodity i is an input, $y_i(\bar{p}) < 0$, and increasing p_i decreases profit.

To see the derivation without the Envelope Theorem, follow the method we used several times in consumer theory.

$$\begin{aligned} \frac{d}{dp_i} (\pi(p)) &= \frac{d}{dp_i} (p \cdot y(p)) = y_i + \sum p_j \frac{\partial y_j}{\partial p_i} \\ &= y_i + \lambda \sum F_j \frac{\partial y_j}{\partial p_i} \\ &= y_i + \lambda(0) = y_i. \end{aligned}$$

⁴The PMP is actually more similar to the EMP than the UMP because both the PMP and EMP do not have to worry about wealth effects, which was the chief complicating factor in the UMP. In fact, rewriting the EMP as a maximization problem shows that

$$\begin{aligned} &\max_x -p \cdot x \\ \text{s.t.} \quad &: u(x) \geq u \end{aligned}$$

is directly analogous to the PMP when there are distinct inputs and outputs.

The second line comes from substituting in the first-order condition $p_j = \lambda F_j$. The third line comes from noting that if the constraint binds, i.e., $F(y(p)) \equiv 0$, then $\sum_j F_j \frac{\partial y_j}{\partial p_i} = 0$ (you can see this by differentiating $F(y(p)) \equiv 0$ with respect to p_i).

2. **If $y()$ is single-valued and differentiable at \bar{p} , then $\frac{\partial y_i(\bar{p})}{\partial p_j} = \frac{\partial y_j(\bar{p})}{\partial p_i}$.** Again, the explanation is the same as in the EMP. Because each of the derivatives in the previous equality are equal to $\frac{\partial^2 \pi(\bar{p})}{\partial p_i \partial p_j}$, Young's theorem implies the result.
3. **If $y()$ is single-valued and differentiable at \bar{p} , the second-derivative matrix of $\pi(p)$, with typical term $\frac{\partial^2 \pi(\bar{p})}{\partial p_i \partial p_j}$, is a symmetric, positive semi-definite matrix.**

That $\frac{\partial^2 \pi(\bar{p})}{\partial p_i \partial p_j}$ is a symmetric, positive semi-definite (p.s.d.) matrix follows from the convexity of $\pi(p)$ in prices. Just as in the case of n.s.d. matrices, p.s.d. matrices have nice properties as well. One of them is that the diagonal elements are non-negative: $\frac{\partial y_i}{\partial p_i} \geq 0$. Hence if the price of an output increases, production of that output increases, and if the price of an input increases, utilization of that input decreases (since for an input, $y_i < 0$, and hence when it increases it becomes closer to zero, meaning that its magnitude decreases. Thus it becomes "less negative," meaning less of the input is used). This is a statement of what is commonly known as the **Law of Supply**. Note that there is no need for a "compensated law of supply" because there is nothing like a wealth effect in the PMP.

5.1.4 A Note on Recoverability

The profit function $\pi(p)$ gives the firm's maximum profit given the prices of inputs and outputs. At first glance, one would think that $\pi(p)$ contains less information about the firm than its technology set Y , since $\pi(p)$ contains only information about optimal behavior. However, a remarkable result in the theory of the firm is that if Y is convex, Y and $\pi(p)$ contain the exact same information. Thus $\pi(p)$ contains a complete description of the productive possibilities open to the firm. I'll briefly sketch the argument.

First, it is easy to show that $\pi(p)$ can be generated from Y , since the very definition of $\pi(p)$ is that it solves the PMP. Thus solving the PMP for any p gives you the profit function. The difficult direction is to show that if you know $\pi(p)$, you can recover the production set Y , provided that it is convex. The method, depicted in Figure 5.4, is as follows:

1. Choose a positive price vector $p \gg 0$, and find the set $\{y | p \cdot y \leq \pi(p)\}$. This gives the set




Figure 5.4: Recovering the Technology Set

of points that earn less profit than the optimal production plan, $y(p)$.

2. Since $y(p)$ is the optimal point in Y , we know that $Y \subset \{y | p \cdot y \leq \pi(p)\}$, and that any point not in $\{y | p \cdot y \leq \pi(p)\}$ cannot be in Y . So, eliminate all points not in $\{y | p \cdot y \leq \pi(p)\}$. If output is on the vertical axis and input on the horizontal axis, all points above and to the right of the price line are eliminated.
3. If we repeat steps 1-2 for all possible positive price vectors, we can eliminate all points that cannot be in Y for any price vector. And, if Y is convex, every point on the transformation frontier is optimal for some price vector. Thus by repeating this process we can trace out the entire transformation frontier, effectively recovering the set Y as

$$Y = \{y | p \cdot y \leq \pi(p) \text{ for all } p \gg 0\} \text{ whenever } Y \text{ is convex.}$$

The importance of this result is that $\pi(p)$ is analytically much easier to work with than Y . But, we should be concerned that by working with $\pi(p)$ we miss some important features of the firm's technology. However, since Y can be recovered in this way, we know that there is no loss of generality in working with the profit function and net supply functions instead of the full production set.

5.2 Production with a Single Output

An important special case of production models is where the firm produces a single output using a number of inputs. In this case, we can make use of a production function (which you probably

remember from intermediate micro). In order to distinguish between outputs and inputs, we will denote the (single) output by q and the inputs by z .⁵ In contrast to the production-plan approach examined earlier, inputs will be non-negative vectors when we use this approach, $z \in R_+^{L-1}$. When there is only one output, the firm's production set can be characterized by a production function $f(z)$, where $f(z)$ gives the quantity of output produced when input vector z is employed by the firm. That is, the relationship between output q and inputs z is given by $q = f(z)$. The firm's production set, Y , can be written as:

$$Y = \{(-z, q) \mid q - f(z) \leq 0 \text{ and } z \geq 0\}.$$

The analog to the marginal rate of transformation in this model is the **marginal rate of technical substitution**. The marginal rate of technical substitution of input j for input i when output is q is given by:

$$MRTS_{ji} = \frac{f_j(q)}{f_i(q)}.$$

$MRTS_{ji}$ gives the amount by which input j should be decreased in order to keep output constant following an increase in input i . Note that defined this way, the MRTS between two inputs will be positive, even though the slope of the production function isoquant is negative.

5.2.1 Profit Maximization with a Single Output

Let $p > 0$ be the price of the firm's output and $w = (w_1, \dots, w_{L-1}) \geq 0$ be the prices of the $L - 1$ inputs, z .⁶ Thus $w \cdot z$ is the cost of using input vector z . The firm's profit maximization problem can be written as

$$\begin{aligned} & \max_{z \geq 0} pq - w \cdot z \\ \text{s.t.} & \quad f(z) \geq q. \end{aligned}$$

Since $p > 0$, the constraint will always bind. Hence the firm's problem can be written in terms of the unconstrained maximization problem:

$$\max_{z \geq 0} pf(z) - w \cdot z.$$

⁵For simplicity, we will assume that none of the output will be used to produce the output. For example, electricity is used to produce electricity. We could easily generalize the model to take account of this possibility.

⁶We should be careful not to confuse w the input price vector with w the consumer's wealth. This is unfortunate, but there isn't much we can do about it.

Since this is an unconstrained problem, we don't need to set up a Lagrangian. However, we do need to be concerned with "corner solutions" since the firm may not use all inputs in production, especially if some are very close substitutes. The Kuhn-Tucker first-order conditions are given by:

$$pf_i(z^*) - w_i \leq 0, \text{ with equality if } z_i^* > 0, \text{ for } \forall i$$

As you may recall, $f_i(z^*)$ is the marginal product of input z_i , the amount by which output increases if you increase input z_i by a small amount. $pf_i(z^*)$ is then the amount by which revenue increases if you increase z_i by a small amount, which is sometimes called the marginal revenue product. Of course, w_i is the amount by which cost increases when you increase z_i by a small amount. So, the condition says that at the optimum it must be that the increase in revenue due to increasing z_i by a small amount is less than the increase in cost. If $z_i^* > 0$, then the increase in revenue must exactly equal the increase in cost. If $z_i^* = 0$, then it may be that $pf_i(z^*) < w_i$, meaning that the increase in revenue due to using even a small amount of input i is not greater than its cost.

We can also rearrange the optimality conditions above to get a familiar tangency condition. Assume that z_i^* and z_j^* are strictly positive. Then the above condition can be rewritten as:

$$\frac{f_i(z^*)}{f_j(z^*)} = \frac{w_i}{w_j},$$

which says that the marginal rate of technical substitution between any two inputs must equal the ratio of their prices. This is just a restatement in this version of the model that the marginal rate of transformation between any two commodities must equal the ratio of their prices, a condition we found in the production-plan version of the PMP.

For the case of a single input, this condition becomes $pf'(z^*) = w$, or $f'(z^*) = \frac{w}{p}$. Since inputs are positive, $f(z)$ increases in z on (z_1, q) space. Profit isoquant slopes upward at $\frac{w}{p}$, since $pq - wz = k$, $q = \frac{w}{p}z + \frac{k}{p}$. Thus, finding $f'(z^*) = \frac{w}{p}$ means finding the point of tangency between the profit isoquant and the production function. Graphically, we move the profit line up and to the left until tangency is found. Note that this is the same thing that is going on in the multiple-input case: we are finding the tangency between profit and the production function.

We can denote the solution to this problem as $z(w, p)$, which is called the **factor demand function**, since it says how much of the inputs are used at prices p and w . Note that $z(w, p)$ is an $L - 1$ dimensional vector. If we plug $z(w, p)$ into the production function, we get $q(w, p) = f(z(w, p))$, which is known as the (output) **supply function**. And, if we plug $z(w, p)$ into the objective function, we get another version of the firm's profit function:

$$\pi(w, p) = pf(z(w, p)) - w \cdot z = pq(w, p) - w \cdot z(w, p).$$

Note that the firm's net supply function $y(p)$ that was examined in the previous section is equivalent to:

$$y(w, p) = (-z_1(w, p), \dots, -z_{L-1}(w, p), q(w, p)).$$

Thus you should think of the model we are working with now as a special case of the one from the previous section, where now we have designated that commodity L is the output and commodities 1 through $L - 1$ are the inputs, whereas previously we had allowed for any of the commodities to be either inputs or outputs.

Because of this connection between the single-output model and the production-plan model, the supply function $q(w, p)$ and factor demand functions $z(w, p)$ will have similar properties to those enumerated in the previous section. A good test of your understanding would be if you could reproduce the results regarding the properties of the net supply function $y(p)$ and profit function $\pi(p)$ for the single-output model.

5.3 Cost Minimization

Consider the model of profit maximization when there is a single output, as in the previous section. At price p , let $z(w, p)$ be the firm's factor demand function and $q(w, p) = f(z(w, p))$ be the firm's supply function. At a particular price vector, (w^*, p^*) , let $q^* = q(w^*, p^*)$. An interesting implication of the profit maximization problem is that $z^* = z(w^*, p^*)$ solves the following problem:

$$\begin{aligned} \min_z \quad & w^* \cdot z \\ \text{s.t.} \quad & f(z) \geq q^*. \end{aligned}$$

That is, z^* is the input bundle that produces q^* at minimum cost when prices are (w^*, p^*) . If a firm maximizes profit by producing q^* using z^* , then z^* is also the input bundle that produces q^* at minimum cost. The reason for this is clear. Suppose there is another bundle $z' \neq z^*$ such that $w^* \cdot z' < w^* \cdot z^*$ and $f(z') \geq q^*$. But, if this is so, then:

$$p^* f(z') - w^* \cdot z' \geq p^* q^* - w^* \cdot z^* = p^* f(z^*) - w^* \cdot z^*,$$

which contradicts the assumption that z^* was the profit maximizing input bundle. If there was another bundle that produced q^* at lower cost than z^* does, the firm should have used it in the

first place. In other words, a necessary condition for the firm to be profit maximizing is that the input bundle it chooses minimize the cost of producing that level of output.

The fact that cost minimization is necessary for profit maximization points toward another way to attack the firm's problem. First, for any level of output, q , find the input bundle that minimizes the cost of producing that level of output at the current level of prices. That is, solve the following problem:

$$\begin{aligned} \min w \cdot z \\ \text{s.t.} \quad : \quad f(z) \geq q. \end{aligned}$$

This is known as the Cost Minimization Problem (CMP).⁷ The Lagrangian for this problem is given by:

$$L = w \cdot z - \lambda(f(z) - q)$$

which implies optimality conditions (for an interior solution):

$$w_i - \lambda f_i(z^*) = 0 \quad \forall i$$

Taking the optimality conditions for goods i and j together implies:

$$\frac{f_i(z^*)}{f_j(z^*)} = \frac{w_i}{w_j},$$

or that the firm chooses z^* so as to set the marginal rate of technical substitution between any two inputs equal to the ratio of their prices (this is directly analogous to $\frac{u_i}{u_j} = \frac{p_i}{p_j}$ in the UMP). The solution to this problem, known as the **conditional factor demand function**, is denoted $z(w, q)$. The function $z(w, q)$ is conditional because it is conditioned on the level of output. Note that the output price does not play a role in determining $z(w, q)$. The value function for this optimization problem is the **cost function**, which is denoted by $c(w, q)$ and defined as

$$c(w, q) = w \cdot z(w, q).$$

The optimized value of the Lagrange multiplier, $\lambda(w, q)$, is the shadow value of relaxing the constraint. Thus $\lambda(w, q)$ is the marginal cost savings of decreasing q , or the marginal cost of increasing q a small amount, $\frac{\partial c}{\partial q} = \lambda(w, q)$. This can be proven in the same way that $\frac{\partial v}{\partial w} = \lambda$ was proven or by application of the envelope theorem.

⁷Note the direct connection with the EMP here. If you think of the consumer as a firm that produces utility using commodities as inputs, then the CMP is just a version of the EMP, where the commodities are z with prices w , the utility function is $f(z)$, and the target utility level is q .

The conditional factor demand function gives the least costly way of producing output q when input prices are w , and the cost function gives the cost of producing that level of output, assuming that the firm minimizes cost. Thus we can rewrite the firm's profit maximization problem in terms of the cost function:

$$\max_q pq - c(w, q).$$

In other words, the firm's problem is to choose the level of output that maximizes profit, given that the firm will choose the input bundle that minimizes the cost of producing that level of output at the current input prices, w . Solving this problem will yield the same input usage and output production as if the PMP had been solved in its original form.

5.3.1 Properties of the Conditional Factor Demand and Cost Functions

As always, there are a number of properties of the cost function. Once again, they are similar to the properties derived in the consumer's version of this problem, the EMP. We start with properties of $z(w, q)$:

1. $z(w, q)$ is **homogeneous of degree zero in w** . The usual reason applies. Since the optimality condition determining $z(w, q)$ is a tangency condition, and the slope of the profit isoquant and the marginal rate of substitution between any pair of quantities is unaffected by scaling w to aw , for $a > 0$, the cost-minimizing point is also unaffected by this scaling (although the cost of that point will be).
2. If $\{z \geq 0 | f(z) \geq q\}$ is **convex**, then $z(w, q)$ is a **convex set**. If $\{z \geq 0 | f(z) \geq q\}$ is strictly convex, then $z(w, q)$ is single-valued. Again, the usual argument applies.

Next, there are properties of the cost function $c(w, q)$.

1. $c(w, q)$ is **homogeneous of degree one in w** . This follows from the fact that $z(w, q)$ is homogeneous of degree zero in w . $c(aw, q) = aw \cdot z(aw, q) = aw \cdot z(w, q) = ac(w, q)$.
2. $c(w, q)$ is **non-decreasing in q** . This one is straightforward. Suppose that $c(w, q') > c(w, q)$ but $q' < q$. In this case, $q = f(z(w, q)) > q'$, so $z(w, q)$ is feasible in the CMP for w and q' , and $c(w, q) < c(w, q')$, which contradicts the assumption that q' was optimal. Basically, this property just means that if you want to produce more output, you have to use more inputs, and as long as inputs are not free, this means that you will incur higher cost.

3. $c(w, q)$ is a **concave function of w** . The argument is exactly the same as the argument for why the expenditure function is concave. Let $w^a = aw + (1 - a)w'$.

$$\begin{aligned} c(w^a, q) &= aw \cdot z(w^a, q) + (1 - a)w' \cdot z(w^a, q) \\ &\geq aw \cdot z(w, q) + (1 - a)w' \cdot z(w', q) \\ &= ac(w, q) + (1 - a)c(w', q). \end{aligned}$$

As usual, going from the first to second line is the crucial step, and the inequality follows from the fact that $z(w^a, q)$ produces output q but is not the cost minimizing way to do so at either w or w' .

The next properties relate $z(w, q)$ and $c(w, q)$:

1. **Shepard's Lemma.** If $z(\bar{w}, q)$ is single valued at \bar{w} , then $\frac{\partial c(\bar{w}, q)}{\partial w_i} = z_i(\bar{w}, q)$. Again, proof is by application of the envelope theorem or by using the first-order conditions. I'll leave this one as an exercise for you. Hint: Think of $c(w, q)$ as $e(p, u)$, and $z(w, q)$ as $h(p, u)$.
2. **If $z(w, q)$ is differentiable, then $\frac{\partial^2 c(w, q)}{\partial w_i \partial w_j} = \frac{\partial z_j(w, q)}{\partial w_i} = \frac{\partial z_i(w, q)}{\partial w_j}$.** Again, the proof is the same.
3. **If $z(w, q)$ is differentiable, then the matrix of second derivatives $\left[\frac{\partial^2 c(w, q)}{\partial w_i \partial w_j} \right]$ is a symmetric (see number 2 above) negative semi-definite (since $c(w, q)$ is concave) matrix.**

5.3.2 Return to Recoverability

Previously we showed that the profit function contains all of the same information as the firm's production set Y . Surprisingly, the firm's cost function also contains all of the same information as Y under certain conditions (namely, that the production function is quasiconcave). In order to show this, we have to show that Y can be recovered from $c(w, q)$, and that $c(w, q)$ can be derived from Y . Since the cost function is generated by solving the cost minimization problem, $c(w, q)$ can clearly be derived from Y . Thus all we have to show is that Y can be recovered from $c(w, q)$. The method is similar to the method used to recover Y from $\pi(p)$.

Begin by fixing an output quantity, q . Then, choose an input price vector w and find the set

$$\{z \geq 0 \mid w \cdot z \geq c(w, q)\}.$$




Figure 5.5: Recovering the Production Function

Since $c(w, q)$ minimizes the cost of producing q over Y , the entire set Y must lay inside of this set, and no point outside of this set can be in Y . Repetition of this process for different w will eliminate more points. In the end, the points that are in the set

$$\{z \geq 0 \mid w \cdot z \geq c(w, q) \text{ for all } w\}$$

will be the set of input vectors that produce output at least q . That is, they will be the upper level set of the production function. By repeating this process for each q , all of the upper level sets can be recovered, which is the same as recovering the technology set (or production function).

Figure 5.5 shows how these sets trace out the level set of the production function.

To summarize, when $f(\cdot)$ is quasiconcave (i.e. $\{z \geq 0 \mid f(z) \geq q\}$ is convex for all q), then

$$Y = \{(-z, q) \mid w \cdot z \geq c(w, q) \text{ for all } w \gg 0\}.$$

5.4 Why Do You Keep Doing This to Me?

In our study of producer theory, we've seen a number of different approaches to the profit maximization problem. First, we looked at the production-plan (or net-output) model, where the firm does not have distinct inputs and outputs. Then we looked at the PMP where the firm has a single output and multiple inputs. Then we looked at the CMP in the single-output case. But, then I showed you that the firm's production set Y , its profit function $\pi(p)$, and its cost function $c(w, q)$ all contain the same information (assuming that the proper convexity conditions hold).

Why do I keep making you learn different approaches to the same problem that all yield the same result? The answer is that each approach is most useful in different situations. You want

to pick and choose the proper approach, depending on the problem you are working with.

The production-plan approach is useful for proving very general propositions, and necessary to deal with situations where you don't know which commodities will be inputs and which will be outputs. This approach is widely used in the study of general equilibrium, where the output of one industry is the input of another.

The CMP approach actually turns out to be very useful. We often have better data on the cost a firm incurs than its production function. But, thanks to the recoverability results, we know that the cost information contains everything we want to know about the firm. Thus the cost-function approach can be very useful from an empirical standpoint.

In addition, recall that the PMP is not very useful for technologies with increasing returns to scale. This is not so with the CMP. The CMP is perfectly well-defined as long as the production function is quasiconcave, even if it exhibits increasing returns. Hence the CMP may allow us to say things about firms when the PMP does not.

The CMP approach is useful for another reason. When the output price p is fixed, the firm's profit maximization problem in terms of the cost function is given by:

$$\max_q pq - c(w, q).$$

However, this approach can also be used when the price the firm charges is not fixed. In particular, suppose the price the firm can charge depends on the quantity it sells according to some function $P(q)$. The profit maximization problem can be written as

$$\max_q P(q)q - c(w, q).$$

Thus our study of the cost function can also be used when the firm is not competitive, as in the cases of monopoly and oligopoly. The other approaches to the PMP we have studied do not translate well to environments where the firm is not a price taker.

5.5 The Geometry of Cost Functions

MWG Section 5.D relates the work we have done on cost functions to the diagrammatic approach you may have studied in your previous micro classes. Although you should look at this section, I'm not going to write everything. You should re-familiarize yourself with the concepts of total cost, fixed cost, variable cost, marginal cost, average cost, etc. They will come in handy in the future (although they are really intermediate micro topics).

5.6 Aggregation of Supply

MWG summarizes this topic by saying, “If firms maximize profits taking prices as given, then the production side of the economy aggregates beautifully.” To this let me add that the whole problem with aggregation on the consumer side was with wealth effects. Since there is no budget constraint here, there are no wealth effects, and thus there is no problem aggregating.

To aggregate supply, just add up the individual supply functions. Let there be m producers, and let $y^m(p)$ be the net supply function for the m^{th} firm. Aggregate supply can be written as:

$$y^T(p) = \sum_{m=1}^M y^m(p).$$

In fact, we can easily think of the aggregate production function as having been produced by an aggregate producer. Define the aggregate production set $Y^T = Y^1 + \dots + Y^M$, where Y^m is the production set of the m^{th} consumer. Thus Y^T represents the opportunities that are available if all sets are used together. Now, consider y^T in Y^T . This is an aggregate production plan. It can be divided into parts, y^{T1}, \dots, y^{TM} , where $y^{Tm} \in Y^m$ and $\sum_m y^{Tm} = y^T$. That is, it can be divided into parts, each of which lies in the production set of some firm. Fix a price vector p^* and consider the profit maximizing aggregate supply vector $y^T(p^*)$. The question we want to ask is this: Can $y^T(p^*)$ be divided into parts $y^{Tm}(p^*)$ such that $y^{Tm}(p^*) \in Y^m$ and $\sum_m y^{Tm}(p^*) = y^T(p^*)$? In other words, is it the case that the profit maximizing production plan for the aggregate production set is the same as would be generated by allowing each of the firms to maximize profit separately and then adding up the profit-maximizing production plans for each firm? The answer is yes.

For all strictly positive prices, $p \gg 0$, $y^T(p) = \sum y^m(p)$, and $\pi^T(p) = \sum_m \pi^m(p)$. Thus the aggregate profit obtained by each firm maximizing profit is the same as is attained by choosing the profit-maximizing production plan from the aggregate production set.

The fact that production aggregates so nicely is a big help because aggregation tends to convexify the production set. For example, see MWG 5.E.2. Even if each individual firm’s production possibilities are non-convex, as in panel a), when you aggregate production possibilities and look at the average production set, it becomes almost convex. Since non-convexities are a big problem for economists, it is good to know that when you aggregate supply, the problems tend to go away.

5.7 A First Crack at the Welfare Theorems

When economists talk about the fundamental theorems of welfare economics, they are really not talking about welfare at all. What they are really talking about is efficiency. This is because, traditionally, economists have not been willing to say what increases society's "welfare," since doing so involves making judgments about how the benefits of various people should be compared. These kind of judgments, which deal with equity or fairness, it is thought, are not the subject of pure economics, but of public policy and philosophy.

However, while economists do not like to make equity judgments, we are willing to make the following statement: Wasteful things are bad. That is, if you are doing something in a wasteful manner, it would be better if you didn't. That (unfortunately) is what we mean by welfare.

Actually, I'm giving economists too hard a time. The reason why we focus on wastefulness (or efficiency) concerns is that while there will never be universal agreement on what is fair, there can be universal agreement on what is wasteful. And as long as we believe that wastefulness is bad, then anything that is fair can be made more fair by eliminating its wastefulness, at least in principle.

Now let's put this idea in practice. We will consider the question of whether profit maximizing firms will choose production plans that are not wasteful. We will call a production plan $y \in Y$ **efficient** if there is no vector $y' \in Y$ such that $y'_i \geq y_i$ for all i , and $y' \neq y$. In other words, a production plan is efficient if there is no other feasible production plan that could either: 1) produce the same output using fewer inputs (i.e. $y'_j = y_j$ and $y'_k > y_k$, where j denotes the output goods and k denotes at least one of the input goods; $y'_k > y_k$ when y' uses less of input k , since y'_k is less negative than y_k); or 2) produce greater output using the same amount of inputs (i.e. $y'_j > y_j$ and $y'_k = y_k$, where j denotes at least one of the output goods and k denotes all the input goods). Thus, production plans that are not efficient are wasteful.

Efficient production plans are located on the firm's transformation frontier. But, note that there may be points on the transformation frontier that are not efficient, as is the case in Figure 5.6, where the transformation frontier has a horizontal segment. Points on the interior of the horizontal part produce the same output as the point furthest to the right on the flat part, but use more inputs. Hence they are not efficient (i.e., they are wasteful).

The first result, which is a version of the **first fundamental theorem of welfare economics**, says that if a production plan is profit maximizing, then it is efficient. Formally, if there exists




Figure 5.6: The Transformation Frontier

$p \gg 0$ such that $y(p) = y$, then y is efficient. The reason for this is straightforward. Since the profit maximizing point is on the transformation frontier, the only case we have to worry about is when it lies on a flat part. But, of all of the points on a flat part, the one that maximizes profits is the one that produces the given output using the fewest inputs, which means that it is also an efficient point.

The second result deals with the converse question: Is it the case that any efficient point is chosen as the profit maximizing point for some price vector? The answer to this case is, “Almost.” If the production set is convex, then any efficient production plan $y \in Y$ is the profit-maximizing production plan for some price vector $p \geq 0$. This result is known as the **second fundamental theorem of welfare economics**. The reason why convexity is needed is illustrated in MWG Figure 5.F.2. Basically, a point on a non-convex part of the transformation frontier can never be chosen as profit-maximizing. However, any point that is on a convex part of the transformation frontier can be chosen as profit-maximizing for the appropriate price vector. Also, notice that $p \geq 0$, not $p \gg 0$. That is, some points, such as y in Figure 5.6, can only be chosen when some of the prices are zero.

5.8 Constant Returns Technologies

As we saw earlier, technologies exhibiting increasing returns are a strange and special case. If the technology exhibits increasing returns everywhere, then the firm will either choose to produce nothing or the firm’s profits will grow without bound as it increases its scale of operations. Further, if the firm exhibits increasing returns over a range of output, the firm will never choose to produce

an output on the interior of this range. It will either choose the smallest or largest production plan on the segment with increasing returns.

Because of the difficulties involved with increasing returns technologies, we tend to focus on nonincreasing returns technologies. But, nonincreasing returns technologies can be divided into two groups: technologies that exhibit decreasing returns and technologies that exhibit constant returns.

In the single-output case, decreasing returns technologies are characterized by strictly concave production functions. For example, a firm that uses machinery and labor to produce output may have a production function of the Cobb-Douglas form:

$$q = f(z_m, z_l) = bz_m^{\frac{1}{3}}z_l^{\frac{1}{3}},$$

where $b > 0$ is some positive real number (for the moment). This production function exhibits decreasing returns to scale:

$$f(az_m, az_l) = b(az_m)^{\frac{1}{3}}(az_l)^{\frac{1}{3}} = a^{\frac{2}{3}}bz_mz_l = a^{\frac{2}{3}}f(z_m, z_l).$$

We can tell a story accounting for decreasing returns in this case. Suppose the firm has a factory, and in it are some machines and some workers. At low levels of output, there is plenty of room for the machines and workers. However, as output increases, additional machines and workers are hired, and the factory begins to get crowded. Soon, the workers and machines begin to interfere with each other, making them less productive. Eventually, the place gets so crowded that nobody can get any work done. This is why the firm experiences decreasing returns to scale.

While the preceding story is realistic, notice that it depends on the fact that the size of the factory is fixed. To put it another way, if the firm's plant just kept growing as it hired more workers and machines, they would never interfere with each other. Because of this, there would be no decreasing returns.

The point of the argument I just gave is to point out that decreasing returns can usually be traced back to some fixed input into the productive process that has not been recognized. In this case, it was the fact that the firm's factory was fixed. To illustrate, suppose that the b from the previous production function really had to do with the fixed level of the plant. In particular, suppose $b = (z_b)^{\frac{1}{3}}$, where z_b is the current size of the firm's plant (building). Taking this into account, the firm's production function is really:

$$f(x_m, x_l, x_b) = x_b^{\frac{1}{3}}x_m^{\frac{1}{3}}x_l^{\frac{1}{3}},$$

and this production function exhibits constant returns to scale.

To paraphrase MWG on this point (bottom of p. 134), the production set Y represents technological possibilities, not limits on resources. If a firm's current production plan can be replicated, i.e., build an identical plant and fill it with identical machines and workers, then it should exhibit constant returns to scale. Of course, it may not actually be possible to replicate everything in practice, but it should be possible in theory. Because of this, it has been argued that decreasing returns must reflect the underlying scarcity of an input into the productive process. Frequently this is managerial know-how, special locations, or something else. However, if this factor could be varied, the technology would exhibit constant returns. Because of this, many people believe that constant returns technologies are the most important sub-category of convex technologies. And, because of that, we'll spend a little time looking at some of the peculiar features of constant returns production functions.

Suppose that the firm produces output q from inputs labor (L) and capital (K) according to $q = f(K, L)$. If $f(\cdot)$ exhibits constant returns,

$$q = f(tK, tL) \equiv tf(K, L), \text{ for any } t > 0.$$

Since this holds for any t , it also holds for $t = \frac{1}{q}$. Making this substitution:

$$1 = f\left(\frac{K}{q}, \frac{L}{q}\right).$$

If we let $\frac{K}{q} = k$ and $\frac{L}{q} = l$, then we can write the previous statement as:

$$f(k, l) = 1,$$

where k is the amount of capital used to produce one unit of output and l is the amount of labor used to produce one unit of output.

To see why this is important, consider the following cost minimization problem:

$$\begin{aligned} & \min_{K, L} rK + sL \\ \text{s.t. } & f(K, L) = q, \end{aligned}$$

where r is the price of capital and s is the price of labor. We can rewrite this problem as:

$$\begin{aligned} & \min_{K, L} q \frac{rK + sL}{q} \\ \text{s.t. } & \frac{1}{q} f(K, L) = 1, \end{aligned}$$

but $\frac{1}{q}f(K, L) = f\left(\frac{K}{q}, \frac{L}{q}\right) = f(k, l)$. Thus the CMP becomes:

$$\begin{aligned} \min q(rk + sl) \\ \text{s.t. } f(k, l) = 1. \end{aligned}$$

Thus if we want to solve the firm's CMP, we can first find the cost-minimizing way to produce one unit of output at the current input prices, and then replicate this production plan q times.

In other words, we can learn everything we want to know about the firm's production function by studying a single isoquant – in this case the unit isoquant.

As another special case of a constant returns technology, we can think of the situation where the firm only has a finite number of alternative production plans. For example, suppose that it can either produce 1 unit of output using 5 units of capital and 3 unit of labor or 2 units of capital and 7 units of labor. Frequently we will call the different production plans in this environment “activities.” Thus activity 1 is (5, 3) and activity 2 is (2, 7).

If the firm has the two activities above, then it can produce one unit of output using either inputs (5, 3) or inputs (2, 7). But, we can also allow the firm to mix the two activities. Thus it can also produce one unit of output by producing $0 \leq a \leq 1$ units of output using activity 1 and $1 - a$ units of output using activity 2. Hence for a firm with these two activities available, the firm's unit isoquant consists of the curve:

$$\begin{aligned} (k, 3) \text{ when } k > 5 \\ a(5, 3) + (1 - a)(2, 7) \text{ for } a \in [0, 1] \\ (2, l) \text{ for } l > 7. \end{aligned}$$

Figure 5.7 depicts this unit isoquant. Note that the vertical and horizontal segments follow from our assumption of free disposal. Also note that if we had more activities, we would begin to trace out something that looks like the nicely differentiable isoquants we have been dealing with all along. Thus differentiable isoquants can be thought of as the limit of having many activities and the “convexification” process we did above.

As a final point, let me make a connection to something you might have seen in your macro classes. Go back to:

$$q = f(tK, tL) \equiv tf(K, L), \text{ for any } t > 0.$$

Instead of dividing by q , as we did above, we can divide by the available labor, L . If we do this,




Figure 5.7: The Unit Isoquant

we get:

$$\frac{q}{L} = f\left(\frac{K}{L}, 1\right).$$

Here we have a version of the production function that gives per capita output as a function of capital. This is something that is useful in macro models. I just wanted to show you that this formula, which you may have seen, is grounded in the production theory we have been studying.

5.9 Household Production Models

Now that we have studied both consumer and producer theories, we can begin to think about how these parts fit together. Chapter 7 explores this topic at the level of the whole economy, but here we start by looking at the individual household as the unit of both production and consumption. In particular, we consider consumers who are able to produce one of the commodities using (at least in part) their own labor. Consider, for example, a consumer who owns a small farm and has preferences over leisure and consumption of the farm's output, which we will call "food." Labor for the farm can either be provided by the farmer or by hiring labor from the market, and food produced by the farm can either be consumed internally by the farmer or sold on the market. Thus it appears that the farmer's decision about how much leisure and food to consume is intertwined with his decision about how much food to produce on the farm and how much labor to hire. However, we will show that if the market for labor is complete, then the farmer's production and consumption decisions can be separated. The farmer can maximize utility by first choosing the amount of total labor that maximizes profit from the farm and then deciding how much of the labor he will provide himself.

5.9.1 Agricultural Household Models with Complete Markets

Models of situations such as we've been talking about are known as Agricultural Household Models (AHM). The version I'll give you here is a simplification of the model in Bardhan and Udry, *Development Microeconomics*, Chapter 2.

Consider a farm owner who can either work on the farm, work off the farm, or not work (leisure). Similarly, he can use his farm land for his own farming, rent it to others to use, or not use it. In addition, the farmer can buy additional labor from the market or rent additional land from the market. Total output produced by the farm depends on the total land used and total labor used on the farm.

We assume that the farmer faces a complete set of markets and that the buying and selling prices for food, labor, and land are the same. Let p be the price of output, w be the price of labor and r be the rental price of land.⁸

The farmer has utility defined over consumption of food and leisure:

$$u(c, l)$$

where c is consumption and l is leisure per person.

The total size of the output is given by

$$F(L, A)$$

where L is the total amount of labor employed on the farm and A is the total amount of land that is cultivated.

The farmer owns a farm with A_E acres. These acres can either be used on the farm or rented to others on the market. Let A_U be the number of acres used on the farm and A_S be the number of acres rented to (sold to) others. Since the farmer gets no utility from having idle land, it is straightforward to show that all acres will either be used internally or rented to the market. Similarly, the farmer has initial endowment of L_E of labor. Let L_U and L_S be the number of units of labor that are used on the farm and sold to the market, respectively. Thus we have two resource constraints:

$$\begin{aligned} A_E &= A_U + A_S \\ L_E &= L_U + L_S + l. \end{aligned}$$

⁸The natural units are for w to be the price per day of labor and r to be the price per season of an acre of land.

These constraints say that all land is either used internally or sold to the market, and all labor is either used internally, sold to the market, or consumed as leisure.

Let A_B and L_B , respectively, be the number of units of land and labor bought from the market. Total labor used on the farm is then given by:

$$L = L_U + L_B$$

and total land used on the farm is given by:

$$A = A_U + A_B.$$

The household's utility maximization problem can be written as:

$$\begin{aligned} & \max u(c, l) \\ & \text{s.t.} \quad : \\ & wL_B + rA_B \leq p(F(L, A) - c) + wL_S + rA_S \\ & L = L_U + L_B \\ & A = A_U + A_B \\ & l = L_E - L_U - L_S \\ & 0 = A_E - A_U - A_S. \end{aligned}$$

The first constraint is a budget constraint. The left-hand side consists of expenditure on labor (wL_B) and land (rA_B) purchased from the market. The right-hand side consists of net revenue from selling the crop to the market (the price is p , and the amount sold is equal to total production $F(L, A)$ less the amount consumption internally, c), plus revenue from selling labor to the market wL_S and revenue from selling land to the market rA_S . The remaining constraints are the definitions of land and labor used on the farm and the resource constraints on available time and land.

Through simplification, the constraints can be rewritten as:

$$\begin{aligned} w(L_B - L_S) + r(A_B - A_S) & \leq p(F(L, A) - c) \\ L - L_U & = L_B \\ l - L_E + L_U & = -L_S \\ A - A_U & = A_B \\ A_U - A_E & = -A_S. \end{aligned}$$

or

$$w(L + l - L_E) + r(A - A_E) \leq p(F(L, A) - c)$$

or

$$wl + pc \leq pF(L, A) - wL + wL_E - rA + rA_E$$

or

$$\begin{aligned} wl + pc &\leq \Pi + wL_E + rA_E \\ \Pi &= pF(L, A) - wL - rA. \end{aligned}$$

The last version says that the total expenditure on commodities, $wl + pc$, must be less than the profit earned by running the farm, Π , plus the value of the farmer's initial endowment, $wL_E + rA_E$. But, notice that L and A appear only on the right-hand side of the budget constraint. Because of this, we can solve the farmer's problem in two stages. First, solve

$$\max_{L, A} pF(L, A) - wL - rA$$

for the optimal total labor and land to be used on the farm. Second, solve the farmer's utility maximization problem:

$$\begin{aligned} &\max_{l, c} u(c, l) \\ \text{s.t.} \quad &: \quad wl + pc \leq \Pi^* + wL_E + rA_E \end{aligned}$$

where

$$\Pi^* = \max_{L, A} pF(L, A) - wL - rA$$

is the maximum profit that can be generated on the farm at prices p , w , and r .

What does this mean? Well, the variables l and c have to do with the farmer's consumption decisions. The variables L and A have to do with his production decisions. The essence of this result says that if you want to solve the farmer's overall utility maximization problem, you can separate his production and consumption decisions. First, choose the production variables that maximize the profit produced on the farm.⁹ Then, choose the consumption bundle that maximizes

⁹There is some degeneracy in the solution to this problem. To see why, suppose the farmer owns 100 acres of land and chooses to cultivate 80 acres himself. Based on the math of the problem, the farmer is indifferent between using 80 acres of his own land and renting 20 on the outside market and, for example, renting all 100 acres of his own land on the market and renting 80 different acres from the market. For simplicity, we'll just assume that the farmer first uses his own land and rents any remaining land to the market (or, if $A > A_E$, that he uses all of his own land and rents the remainder from the market).

utility subject to the “ordinary” budget constraint, where wealth is the sum of maximized profit and endowment wealth. To put it another way, the farmer has two separate decisions to make: how much land and labor to use on the farm, and how much leisure and food to consume. These two decisions can be made separately. In particular, the farmer can decide how much labor to use on the farm without deciding how much of **his own labor** he should use on the farm. Similarly, the farmer can decide how much food to produce without deciding how much food he, himself, is going to consume.

The result stated in the previous paragraph is known as the separation property of the AHP: When markets are complete, the production and consumption decisions of the household are separate from each other. The farmer chooses total labor and land in order to maximize profit, and then chooses how much labor and land to consume as if in an ordinary UMP, where wealth is the sum of maximized profit from production and endowment wealth.

The separation property is an implication of utility maximization and complete markets, and it arises endogenously as an implication of the model. Completeness of markets plays a critical role in the result. For our purposes, a market is complete if the farmer can buy or sell as much of the commodity as he wants at a particular price that is the same regardless of whether the farmer buys or sells. A market can fail to be complete if either the price at which someone buys an item is different than the price at which he sells the item, or if there are limits on the quantity that can be bought or sold of an item.¹⁰

A graphical illustration of the separation property may help. For illustration (because three dimensional graphs are hard to draw), we will assume that there is no market for land. The farmer has $A_E = A^*$ acres of land, and uses them all on the farm. The first stage of the optimization problem is to choose the amount of labor to use on the farm. This is done by maximizing the profit earned on the farm, $pc - wL$, subject to the constraint that $c = F(L, A^*)$. This profit maximization problem is illustrated in Figure 5.8. Thus the first stage in the farmer’s decision results in L^* units of labor being used on the farm. When this amount of labor is used, the farm’s profit is given by the level of the solid black isoprofit line at L^* . The equation of this line is given by $c = F(L^*, A^*) + \frac{w}{p}(L_U + L_S)$.¹¹ If the farmer chooses to work more than L^* by selling some

¹⁰One of the chief “market failures” in developing economies is incompleteness of labor markets. It is often impossible to buy labor regardless of the price offered because it is costly for the people who have jobs to find the people who are willing to work.

¹¹This isoprofit line makes intuitive sense if you think about it as a budget line. Consumption equals the total product of the farm plus the value of labor provided by the farmer, normalized to the units of consumption (w/p).




Figure 5.8: Profit Maximization

of his labor on the market, he moves to the right along the solid isoprofit line and increases the amount of wealth he has for consumption. If he chooses to work less than L^* , he moves to the left along the solid isoprofit line and decreases the wealth he has available for consumption.

In the second stage the farmer chooses how much labor he will provide. This is found by maximizing utility subject to the constraint that total spending satisfy the budget constraint $c \leq F(L^*, A^*) + \frac{w}{p}(L_U + L_S)$ (which we know will bind). The solution is depicted in Figure 5.9. As illustrated in the figure, $L^* < L_E$. Under the natural assumption that if the total labor used on the farm is less than the farmer's endowment of labor, he will buy no labor from the market, the farmer's choice of labor used on the farm (L_U^*), labor sold to the market (L_S^*), and leisure (l^*) are as shown in the diagram.

You can think of the parts of Figure 5.9 as evolving as follows. First, the farmer chooses how much labor, L^* , to use on the farm. This quantity is chosen in order to maximize profit without regard to how much of L^* will be provided by the farmer, himself, and how much will be hired from the market. Next, the farmer chooses how much of his own leisure to provide as labor, given that, in addition to any labor income he might earn, he will also have the profit from the farm to use to purchase the consumption good. Thus in the second stage the farmer chooses labor according to the budget set defined by the farm's maximized isoprofit line.

Generally speaking, the separation property will not hold if markets are sufficiently incomplete. Notice that in the model we presented, the separation result would continue to hold if either the market for labor or land were incomplete (i.e., labor or land could not be traded), but not both. However, if multiple markets were incomplete, the separation result would fail.




Figure 5.9: Separation Property

Suppose there is no market for land (that is, $A = A_E$) and there is an upper bound on how much labor the household can sell to the market (that is, $L_S \leq M$). The farmer's problem is given by:

$$\begin{aligned}
 & \max u(c, l) \\
 & s.t. : \\
 & wL_B \leq p(F(L, A_E) - c) + wL_S \\
 & l = L_E - L_S - L_U \\
 & L = L_B + L_U \\
 & L_S \leq M.
 \end{aligned}$$

This version of the model is the same as the complete market version except that we have eliminated the market for land and added the constraint that the amount of labor sold on the market must be less than M .

The optimal solution to this problem takes one of two forms. If, in the perfect markets version of the problem $L_S^* < M$, then the solution in the incomplete markets case is the same as in the complete markets case. On the other hand, if $L^* > M$ in the complete markets case, then the farmer is limited by the incomplete market and must choose $L_S^* = M$. Because of this, separation will fail to hold. Intuitively, this is because if separation holds the farmer can first decide how

much labor to use on the farm and then decide how much of his own labor to use. So, suppose the farmer would choose to operate a small farm (i.e., one that uses a small amount of labor). If markets are complete, he can then sell his surplus labor. But, if markets are not complete, the farmer will not be able to sell the surplus labor. Because of this, he will consume some of the excess as leisure, but he will also use some of it as additional labor on his own farm. Thus the optimal amount of labor to use on the farm will be affected by the amount of labor that the farmer can sell. Separation does not hold.

Formally, when the constraint on use of the labor market binds, the problem becomes:

$$\begin{aligned} \max u(c, l) \\ \text{s.t.} : pc = pF(L_E - M - l, A_E) + wM \end{aligned}$$

and the optimum is found at the point of tangency between the production function F and the utility function $u(c, l)$. This will generally not be the same point as the farmer would select if markets were complete.

Figure 5.10 illustrates the farmer's problem when the market for labor is incomplete. Suppose that in the absence of labor market imperfections, the farmer would like to sell a lot of labor: $L_S^* > M$. In this case, the constraint will bind, and leisure is given by: $l = L_E - L_U - M$. Thus the consumer can no longer independently choose L_U and l . Setting one determines the other.

So, consider three alternative values of L_U : L_U^1 , L_U^2 , and L_U^3 . L_U^1 is the value of L_U that maximizes profit from the farm. However, notice that because of the credit market imperfection, the farmer is only able to move along the profit line to the point labeled (1), where he works L_U^1 hours at home, sells M hours of labor, and consumes $l = L_E - M - L_U^1$ hours as leisure. If the farmer increases labor usage to L_U^2 , he changes the farm's profit, and so as he sells labor he moves along the new profit line to the point labeled (2), again by selling the maximum number of labor hours. Notice that point 2 involves more consumption but less leisure than point 1. Similarly, if the farmer works L_U^3 units of labor at home, then he can move along the new profit line to point (3). And, point 3 offers more consumption and less leisure than either point 2 or point 3.

Whether the farmer prefers point 1, 2, or 3 will depend on the shape of his isoquants, as illustrated in Figure 5.11. If they are relatively flat, then the consumer will prefer more consumption and less leisure: point 3 will be the best. However, if they are steep, then the farmer will prefer more leisure and less consumption - point 1 will be the best. Finally, for more intermediate preferences, point 2 may be the best. In any case, since the optimal point depends on the shape of the farmer's




Figure 5.10: Imperfect Labor Markets

utility isoquants and the shape of the production function, he will no longer be able to separate his production and consumption decisions.




Figure 5.11: AHP Without Separation

Chapter 6

Choice Under Uncertainty

Up until now, we have been concerned with choice under certainty. A consumer chooses which commodity bundle to consume. A producer chooses how much output to produce using which mix of inputs. In either case, there is no uncertainty about the outcome of the choice.

We now turn to considering choice under uncertainty, where the objects of choice are not certainties, but distributions over outcomes. For example, suppose that you have a choice between two alternatives. Under alternative A, you roll a six-sided die. If the die comes up 1, 2, or 3, you get \$1000. If it comes up 4, 5, or 6, you lose \$300. Under alternative B, you choose a card from a standard 52 card deck. If the card you choose is black, you pay me \$200. If it is a heart, you get a free trip to Bermuda. If it is a diamond, you have to shovel the snow off of my driveway all winter.

If I were to ask you whether you preferred alternative A or alternative B, you could probably tell me. Indeed, if I were to write down any two random situations, call them L_1 and L_2 , you could probably tell me which one you prefer. And, there is even the possibility that your preferences would be complete, transitive (i.e., rational), and continuous. If this is true then I can come up with a utility function representing your preferences over random situations, call it $U(L)$, such that L_1 is strictly preferred to L_2 if and only if $U(L_1) > U(L_2)$. Thus, without too much effort, we can extend our standard utility theory to utility under uncertainty. All we need is for the consumer to have well defined preferences over uncertain alternatives.

Now, recall that I said that much of what we do from a modeling perspective is add structure to people's preferences in order to be able to say more about how they behave. In this situation, what we would like to be able to do is say that a person's preferences over uncertain alternatives

should be able to be expressed in terms of the utility the person would assign to the outcome if it were certain to occur, and the probability of that outcome occurring. For example, suppose we are considering two different uncertain alternatives, each of which offers a different distribution over three outcomes: I buy you a trip to Bermuda, you pay me \$500, or you paint my house. The probability of each outcome under alternatives A and B are given in the following table:

	Bermuda	-\$500	Paint my house
A	.3	.4	.3
B	.2	.7	.1

What we would like to be able to do is express your utility for these two alternatives in terms of the utility you assign to each individual outcome and the probability that they occur. For example, suppose you assign value u_B to the trip to Bermuda, u_m to paying me the money, and u_p to painting my house. It would be very nice if we could express your utility for each alternative by multiplying each of these numbers by the probability of the outcome occurring, and summing. That is:

$$U(A) = 0.3u_B + 0.4u_m + 0.3u_p$$

$$U(B) = 0.2u_B + 0.7u_m + 0.1u_p.$$

Note that if this were the case, we could express the utility of any distribution over these outcomes in the same way. If the probabilities of Bermuda, paying me the money, and painting my house are p_B, p_m , and p_p , respectively, then the expected utility of the alternative is

$$p_B u_B + p_m u_m + p_p u_p.$$

This would be very useful, since it would allow us to base our study of choice under uncertainty on a study of choice over certain outcomes, extended in a simple way.

However, while the preceding equation, known as an **expected utility form**, is useful, it is not necessarily the case that a consumer with rational preferences over uncertain alternatives will be such that those alternatives can be represented in this form. Thus the question we turn to first is what additional structure we have to place on preferences in order to ensure that a person's preferences can be represented by a utility function that takes the expected utility form. After identifying these conditions, we will go on to show how utility functions of the expected utility form can be used to study behavior under uncertainty, and draw testable implications about people's behavior that are not implied by the standard approach.

6.1 Lotteries

In our study of consumer theory, the object of choice was a commodity bundle, x . In producer theory, the object of choice was a net input vector, y . In studying choice under uncertainty, the basic object of choice will be a **lottery**. A lottery is a probability distribution over a set of possible outcomes.

Suppose that there are N possible outcomes, denoted by a_1, \dots, a_N . Let $A = \{a_1, \dots, a_N\}$ denote the set of all possible outcomes. A **simple lottery** consists of an assignment of a probability to each outcome. Thus a **simple lottery** is a vector $L = (p_1, \dots, p_N)$ such that $p_n \geq 0$ for $n = 1, \dots, N$, and $\sum_n p_n = 1$.

A **compound lottery** is a lottery whose prizes are other lotteries. For example, suppose that I ask you to flip a coin. If it comes up heads, you roll a die, and I pay you the number of dollars that it shows. If the die comes up tails, you draw a random number between 1 and 10 and I pay you that amount of dollars. The set of outcomes here is $A = (1, 2, 3, 4, 5, 6, 7, 8, 9, 10)$. The coin flip is then a lottery whose prizes are the lotteries $(\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, 0, 0, 0, 0)$ and $(\frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10})$. Thus the coin flip represents a compound lottery. Notice that since the coin comes up heads or tails with probability $\frac{1}{2}$ each, the compound lottery can be reduced to a simple lottery:

$$\begin{aligned} & \frac{1}{2} \left(\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, 0, 0, 0, 0 \right) + \frac{1}{2} \left(\frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10} \right) \\ &= \left(\frac{2}{15}, \frac{2}{15}, \frac{2}{15}, \frac{2}{15}, \frac{2}{15}, \frac{2}{15}, \frac{1}{20}, \frac{1}{20}, \frac{1}{20}, \frac{1}{20} \right), \end{aligned}$$

where the final vector gives the probability of each outcome **before the coin is flipped**.

Generally, a compound lottery can be represented as follows. Suppose that there are K lotteries, denoted by L_1, \dots, L_K . Let Q be a lottery whose prizes are the lotteries L_1, \dots, L_K . That is, suppose that lottery Q awards the prize L_k with probability q_k . So, you can think of a compound lottery as a two-stage lottery. In the first stage, which lottery L_k you play in the second stage is determined. In the second stage, you play lottery L_k .

So, we call Q a compound lottery. It assigns probability q_k to L_k , where $q_k \geq 0$ and $\sum_k q_k = 1$. If p_n^k is the probability that L_k assigns to outcome n , this compound lottery can be reduced to a simple lottery where

$$p_n = \sum_k q_k p_n^k$$

is the probability of outcome n occurring. That is, p_n gives the probability of outcome n being the final outcome of the compound lottery before any of the randomizations have occurred.

If L and L' are lotteries, a compound lottery over these two lotteries can be represented as $aL + (1 - a)L'$, where $0 \leq a \leq 1$ is the probability of lottery L occurring.

6.1.1 Preferences Over Lotteries

We begin by building up a theory of rational preferences over lotteries. Once we do that, we'll know that there is a utility function that represents those preferences (under certain conditions). We'll then go on to ask whether those preferences can be represented by a utility function of the expected utility form.

Let \tilde{L} be the set of all possible lotteries. Thus \tilde{L} is like X from consumer theory, the set of all possible consumption bundles. We want our consumer to have rational preferences over lotteries. So, suppose that the relation \succsim represents strict preferences over lotteries, and suppose that these preferences are rational, i.e., complete and transitive.

We will also assume that the consumer's preferences are **consequentialist**. Basically, this means that consumers care only about the distribution over final outcomes, not whether this distribution comes about as a result of a simple lottery, or a compound lottery. In other words, the consumer is indifferent between any two compound lotteries that can be reduced to the same simple lottery. This property is often called **reduction of compound lotteries**.¹ Because of the reduction property, we can confine our attention to the set of all simple lotteries, and from now on we will let \tilde{L} be the set of all simple lotteries.

The other requirement we needed for preferences to be representable by a utility function was continuity. In consumer theory, continuity meant that you could not pass from the lower level set of a consumer's utility function to the upper level set without passing through the indifference set. Something similar is involved with continuity here, but what we are interested in is continuity in probabilities.

Suppose that $L \succ L'$. The essence of continuity is that adding a sufficiently small probability of some other lottery, L'' , to L should not reverse this preferences. That is:

$$\text{if } L \succ L', \text{ then } L \succ (1 - a)L' + aL'' \text{ for some } a > 0.$$

¹Another way to think about the reduction property is that we're assuming there is no process-oriented utility. Consumers do not enjoy the process of the gamble, only the outcome, eliminating the "fun of the gamble" in settings like casinos.

Formally, \succsim on \tilde{L} is continuous if for any L , L' , and L'' , and any $a \in (0, 1)$, the sets:

$$\{a \in (0, 1) \mid L \succsim (1 - a)L' + aL''\}$$

and

$$\{a \in (0, 1) \mid (1 - a)L' + aL'' \succsim L\}$$

are closed.

Continuity is mostly a technical assumption needed to get the existence of a utility function. But, notice that its validity is less compelling than it was in the certainty case. For example, suppose L is a trip to Bermuda and L' is \$3000, and that you prefer L' to L . Now suppose we introduce L'' , which is violent, painful, death. If preferences are continuous, then the trip to Bermuda should also be less preferred than \$3000 with probability $1 - a$ and violent painful death with probability a , provided that a is sufficiently small. For many people, there is no probability $a > 0$ such that this would be the case, even though when $a = 0$, L' is preferred to L .

If the consumer has rational, continuous preferences over \tilde{L} , we know that there is a utility function $U()$ such that $U()$ represents those preferences. In order to get a utility function of the expected utility form that represents those preferences, the consumer's preferences must also satisfy the **independence axiom**.

The preferences relation \succsim on \tilde{L} satisfies the **independence axiom** if for all L , L' , and L'' and $a \in (0, 1)$, $L \succsim L'$ if and only if $aL + (1 - a)L'' \succsim aL' + (1 - a)L''$.

The essence of the independence axiom is that outcomes that occur with the same probability under two alternatives should not affect your preferences over those alternatives. For example, suppose that I offer you the choice between the following two alternatives:

$$\begin{aligned} L &: \$5 \text{ with probability } \frac{1}{5}, 0 \text{ with probability } \frac{4}{5} \\ L' &: \$12 \text{ with probability } \frac{1}{10}, 0 \text{ with probability } \frac{9}{10}. \end{aligned}$$

Suppose you prefer L to L' . Now consider the following alternative. I flip a coin. If it comes up heads, I offer you the choice between L and L' . If it comes up tails, you get nothing. What the independence axiom says is that if I ask you to choose either L or L' before I flip the coin, your preference should be the same as it was when I didn't flip the coin. That is, if you prefer L to L' , you should also prefer $\frac{1}{2}L + \frac{1}{2}0 \succsim \frac{1}{2}L' + \frac{1}{2}0$, where 0 is the lottery that gives you 0 with probability 1.

The previous example illustrates why the independence axiom is frequently called “independence of irrelevant alternatives.” The irrelevant alternative is the event that occurs regardless of your choice. Thus, the independence axiom says that alternatives that occur regardless of what you choose should not affect your preferences.

Although the independence axiom seems straightforward, it is actually quite controversial. To illustrate, consider the following example, known as the Allais Paradox.

Consider a lottery with three possible outcomes: \$2.5 million, \$0.5 million, and \$0. Now, consider the following two lotteries (the numbers in the table are the probabilities of the outcome in the column occurring under the lottery in the row):

	\$2.5M	\$0.5M	\$0
L_1	0	1	0
L'_1	.1	.89	.01

That is, L_1 offer \$500,000 for sure. L'_1 offers \$2.5M with probability 0.1, \$500,000 with probability 0.89, and 0 with probability 0.01. Now, before going on, decide whether you would choose L_1 or L'_1 .

Next, consider the following two alternative lotteries over the same prizes.

	\$2.5M	\$0.5M	0
L_2	0	0.11	0.89
L'_2	.1	0	.9

It is not unusual for people to prefer L_1 to L'_1 , but L'_2 to L_2 . However, such behavior is a violation of the independence axiom. To see why, define lotteries $L_A = (0, 1, 0)$, $L_B = (\frac{10}{11}, 0, \frac{1}{11})$, and $L_C = (0, 0, 1)$. Notice that

$$\begin{aligned} L_1 &= 0.89L_A + 0.11L_A \\ L'_1 &= 0.89L_A + 0.11L_B. \end{aligned}$$

Thus preferences over L_1 should be preferred to L'_1 if and only if L_A is preferred to L_B .

Similarly, consider that L_2 and L'_2 can be written as:

$$\begin{aligned} L_2 &= 0.11L_A + 0.89L_C \\ L'_2 &= 0.11L_B + 0.89L_C. \end{aligned}$$

Thus if this person satisfies the independence axiom, L_2 should be preferred to L'_2 whenever L_A is preferred to L_B , which is the same as in the L_1 vs. L'_1 case above. Hence if L_1 is preferred to L'_1 , then L_2 should also be preferred to L'_2 .

Usually, about half of the people prefer L_1 to L'_1 but L'_2 to L_2 . Does this mean that they are irrational? Not really. What it means is that they do not satisfy the independence axiom. Whether or not such preferences are irrational has been a subject of debate in economics. Some people think yes. Others think no. Some people would argue that if your preferences don't satisfy the independence axiom, it's only because you don't understand the problem. And, once the nature of your failure has been explained to you, you will agree that your behavior should satisfy the independence axiom and that you must have been mistaken or crazy when it didn't. Others think this is complete nonsense. Basically, the independence axiom is a source of great controversy in economics. This is especially true because the independence axiom leads to a great number of paradoxes like the Allais paradox mentioned earlier.

In the end, the usefulness of the expected utility framework that we are about to develop usually justifies its use, even though it is not perfect. A lot of the research that is currently going on is trying to determine how you can have an expected utility theory without the independence axiom.

6.1.2 The Expected Utility Theorem

We now return to the question of when there is a utility function of the expected utility form that represents the consumer's preferences. Recall the definition:

Definition 9 *A utility function $U(L)$ has an expected utility form if there are real numbers u_1, \dots, u_N such that for every simple lottery $L = (p_1, \dots, p_N)$,*

$$U(L) = \sum_n p_n u_n.$$

The reduction property and the independence axiom combine to show that **utility function $U(L)$ has the expected utility form if and only if it is linear**, meaning it satisfies the property:

$$U\left(\sum_{k=1}^K t_k L_k\right) = \sum_{k=1}^K t_k U(L_k) \quad (6.1)$$

for any K lotteries. To see why, note that we need to show this in "two directions" - first, that the expected utility form implies linearity; then, that linearity implies the expected utility form.

1. Suppose that $U(L)$ has the expected utility form. Consider the compound lottery $\sum_{k=1}^K t_k L_k$.

$$U\left(\sum_{k=1}^K t_k L_k\right) = \sum_n u_n \left(\sum_{k=1}^K t_k p_n^k\right) = \sum_{k=1}^K t_k \left(\sum_n u_n p_n^k\right) = \sum_{k=1}^K t_k U(L_k).$$

So, it is linear.

2. Suppose that $U(L)$ is linear. Let L^n be the lottery that awards outcome a_n with probability 1. Then

$$U(L) = U\left(\sum_n p_n L^n\right) = \sum_n p_n U(L^n) = \sum_n p_n u_n.$$

So, it has the expected utility form.

Thus proving that a utility function has the expected utility form is equivalent to proving it is linear. We will use this fact momentarily.

The **expected utility theorem** says that if a consumer's preferences over simple lotteries are rational, continuous, and exhibit the reduction and independence properties, then there is a utility function of the expected utility form that represents those preferences. The argument is by construction. To make things simple, suppose that there is a best prize, a_B , and a worst prize, a_W , among the prizes. Let L^B be the "degenerate lottery" that puts probability 1 on a_B occurring, and L^W be the degenerate lottery that puts probability 1 on a_W . Now, consider some lottery, L , such that $L^B \succ L \succ L^W$. By continuity, there exists some number, a_L , such that

$$a_L L^B + (1 - a_L) L^W \sim L.$$

We will define the consumer's utility function as $U(L) = a_L$ as defined above, and note that $U(L^B) = 1$ and $U(L^W) = 0$. Thus the utility assigned to a lottery is equal to the probability put on the best prize in a lottery between the best and worst prizes such that the consumer is indifferent between L and $a_L L^B + (1 - a_L) L^W$.

In order to show that $U(L)$ takes the expected utility form, we must show that:

$$U(tL + (1 - t)L') = ta_L + (1 - t)a_{L'}.$$

If this is so, then $U(\cdot)$ is linear, and thus we know that it can be represented in the expected utility form. Now, $L \sim a_L L^B + (1 - a_L) L^W$, and $L' \sim a_{L'} L^B + (1 - a_{L'}) L^W$. Thus:

$$\begin{aligned} & U(tL + (1 - t)L') \\ &= U\left(t(a_L L^B + (1 - a_L) L^W) + (1 - t)(a_{L'} L^B + (1 - a_{L'}) L^W)\right) \end{aligned}$$

by the independence property. By the reduction property:

$$= U((ta_L + (1-t)a_{L'})L^B + (1 - (ta_L + (1-t)a_{L'}))L^W)$$

and by the definition of the utility function:

$$= ta_L + (1-t)a_{L'}.$$

This proves that $U(\cdot)$ is linear, and so we know that $U(\cdot)$ can be written in the expected utility form.²

I'm not going to write out the complete proof. However, I am going to write out the expected utility theorem.

Expected Utility Theorem: *Suppose that rational preference relation \succsim is continuous and satisfies the reduction and independence axioms on the space of simple lotteries \tilde{L} . Then \succsim admits a utility function $U(L)$ of the expected utility form. That is, there are numbers u_1, \dots, u_N such that $U(L) = \sum_{n=1}^N p_n^L u_n$, and for any two lotteries,*

$$L \succsim L' \text{ if and only if } U(L) \geq U(L').$$

Note the following about the expected utility theorem:

1. The expected utility theorem says that under these conditions, there is one utility function, call it $U(L)$, of the expected utility form that represents these preferences.
2. However, there may be other utility functions that also represent these preferences.
3. In fact, any monotone transformation of $U(\cdot)$ will also represent these preferences. That is, let $V(\cdot)$ be a monotone transformation, then $V(U(L)) = V(\sum_n p_n^L u_n)$ also represents these preferences.
4. However, it is not necessarily the case that $V(U(L))$ can be written in the expected utility form. For example, $v = e^u$ is a monotone transformation, but there is no way to write $V(U(L)) = \exp(\sum_n p_n^L u_n)$ in the expected utility form.
5. But, there are some types of transformations that can be applied to $U(\cdot)$ such that $V(U(\cdot))$ also has the expected utility form. It can be shown that the transformed utility function also has the expected utility form if and only if $V(\cdot)$ is linear. This is summarized as follows:

²This is not a formal proof, but it captures the general idea. There are some technical details that must be addressed in the formal proof, but you can read about these in your favorite micro text.

The Expected Utility Form is preserved only by positive linear transformations. If $U(\cdot)$ and $V(\cdot)$ are utility functions representing \succsim , and $U(\cdot)$ has the expected utility form, then $V(\cdot)$ also has the expected utility form if and only if there are numbers $a > 0$ and b such that:

$$U(L) = aV(L) + b.$$

In other words, the expected utility property is preserved by positive linear (affine) transformations, but any other transformation of $U(\cdot)$ does not preserve this property.

MWG calls the utility function of the expected utility form a **von-Neumann-Morgenstern (vNM)** utility function, and I'll adopt this as well. That said, it is important that we do not confuse the vNM utility function, $U(L)$, with the numbers u_1, \dots, u_N associated with it.³

An important consequence of the fact that the expected utility form is preserved only by positive linear transformations is that a vNM utility function imposes cardinal significance on utility. To see why, consider the utility associated with four prizes, u_1, u_2, u_3 , and u_4 , and suppose that

$$u_1 - u_2 > u_3 - u_4.$$

Suppose we apply a positive linear transformation to these numbers:

$$v_n = au_n + b.$$

Then

$$\begin{aligned} v_1 - v_2 &= au_1 + b - (au_2 + b) = a(u_1 - u_2) \\ &> a(u_3 - u_4) = au_3 + b - (au_4 + b) = v_3 - v_4. \end{aligned}$$

Thus $v_1 - v_2 > v_3 - v_4$ if and only if $u_1 - u_2 > u_3 - u_4$. And, since any utility function of the expected utility form that represents the same preferences will exhibit this property, differences in utility numbers are meaningful. The numbers assigned by the vNM utility function have cardinal significance. This will become important when we turn to our study of utility for money and risk aversion, which we do next.

6.1.3 Constructing a vNM utility function

Let $A = \{a_1, \dots, a_N\}$ denote the set of prizes. Suppose that a_N is the best prize and a_1 is the worst prize. We are going to show how to construct numbers u_1, \dots, u_N that make up a utility function with the expected utility form that represents preferences over these prizes.

³Later, when we allow for a continuum of prizes (such as monetary outcomes), the numbers u_1, \dots, u_N become the function $u(x)$, and we'll call the lowercase $u(x)$ function the Bernoulli utility function.

First, we are going to arbitrarily choose $u_N = 1$ and $u_1 = 0$. Why? Because we can. Remember, the point here is to construct a utility function that has the expected utility form. We could just as easily do it for arbitrary specifications of u_N and u_1 , but this is notationally a bit simpler.

Now, for each prize a_i , define u_i to be the probability such that the decision maker is indifferent between prize a_i for sure and a lottery that offers a_N with probability u_i and a_1 with probability $1 - u_i$. Let's refer to the lottery that offers prize a_N with probability u_i and prize a_1 with probability $(1 - u_i)$ lottery S_i . So, u_i 's are between 0 and 1. Notice that if we specify the numbers in this way it makes sense that $u_N = 1$ and $u_1 = 0$, since the decision maker should be indifferent between the best prize for sure and a lottery that offers the best prize with probability $u_1 = 1$, etc.

This gives us a way of defining numbers u_i . Now, we want to argue that this way of defining the u_i 's, combined with consequentialism (reduction of compound lotteries) and Independence of Irrelevant alternatives yields a utility function that looks like $U(L) = \sum p_i u_i$.

So, consider lottery $L = (p_1, \dots, p_N)$. This lottery offers prize a_i with probability p_i . But, we know that the decision maker is indifferent between a_i for sure and a lottery that offers prize a_N with probability u_i and price a_1 with probability $1 - u_i$. Thus, using IIA, we know that the decision maker is indifferent between lottery L and a compound lotter in which, with probability p_i , the decision maker faces another lottery: u_N with probability u_i and u_1 with probability $1 - u_i$. This lottery is depicted as L' in the following diagram. Note that L' only has two distinct prizes: a_N and a_1 . By reduction of compound lotteries, we can combine the total probability of each outcome, making an equivalent simple lottery, L'' . The utility for lottery L'' is $(\sum p_i u_i) u_N + (1 - \sum p_i u_i) u_1$. Since $u_N = 1$ and $u_1 = 0$, this gives that $U(L) = U(L') = U(L'') = \sum p_i u_i$, which is what we wanted to show. Defining utility in this way gives us a representation with the expected utility form.

6.2 Utility for Money and Risk Aversion

The theory of choice under uncertainty is most frequently applied to lotteries over monetary outcomes. The easiest way to treat monetary outcomes here is to let x be a continuous variable representing the amount of money received. With a finite number of outcomes, assign a number u_n to each of the N outcomes. We could also do this with the continuous variable, x , just by letting u_x be the number assigned to the lottery that assigns utility x with probability 1. In this




Figure 6.1:

case, there would be one value of u_x for each real number, x . But, this is just what it means to be a function. So, we'll let the function $u(x)$ play the role that u_n did in the finite outcome case. Thus $u(x)$ represents the utility associated with the lottery that awards the consumer x dollars for sure.

Since there is a continuum of outcomes, we need to use a more general probability structure as well. With a discrete number of outcomes, we represented a lottery in terms of a vector (p_1, \dots, p_N) , where p_n represents the probability of outcome n . When there is a continuum of outcomes, we will represent a lottery as a distribution over the outcomes. One concept that you are probably familiar with is using a probability density function $f(x)$. When we had a finite number of outcomes, we denoted the probability of any particular outcome by p_n . The analogue to this when there are a continuous number of outcomes is to use a probability density function (pdf). The pdf is defined such that:

$$\Pr(a \leq x \leq b) = \int_a^b f(x) dx.$$

Recall that when a distribution can be represented by a pdf, it has no atoms (discrete values of x with strictly positive probability of occurring). Thus the probability of any particular value of x being drawn is zero. The expected utility of a distribution $f()$ is given by:

$$U(f) = \int_{-\infty}^{+\infty} u(x) f(x) dx,$$

which is just the continuous version of $U(L) = \sum_n p_n u_n$. In order to keep things straight, we will call $u(x)$ the Bernoulli utility function, while we will continue to refer to $U(f)$ as the vNM utility function.

It will also be convenient to write a lottery in terms of its cumulative distribution function (cdf) rather than its pdf. The cdf of a random variable is given by:

$$F(b) = \int_{-\infty}^b f(x) dx.$$

When we use the cdf to represent the lottery, we'll write the expected utility of F as:

$$\int_{-\infty}^{+\infty} u(x) dF(x).$$

Mathematically, the latter formulation lets us deal with cases where the distribution has atoms, but we aren't going to worry too much about the distinction between the two.

The Bernoulli utility function provides a convenient way to think about a decision maker's attitude toward risk. For example, consider a gamble that offers \$100 with probability $\frac{1}{2}$ and 0

with probability $\frac{1}{2}$. Now, if I were to offer you the choice between this lottery and c dollars for sure, how small would c have to be before you are willing to accept the gamble?

The expected value of the gamble is $\frac{1}{2}100 + \frac{1}{2}0 = 50$. However, if offered the choice between 50 for sure and the lottery above, most people would choose the sure thing. It is not until c is somewhat lower than 50 that many people find themselves willing to accept the lottery. For me, I think the smallest c for which I am willing to accept the gamble is 40. The fact that $40 < 50$ captures the idea that I am **risk averse**. My expected utility from the lottery is less than the utility I would receive from getting the expected value of the gamble for sure. The minimum amount c such that I would accept the gamble instead of the sure thing is known as the **certainty equivalent** of the gamble, since it equals the certain amount of money that offers the same utility as the lottery. The difference between the expected value of the lottery, 50, and my certainty equivalent, 40, is known as my **risk premium**, since I would in effect be willing to pay somebody 10 to take the risk away from me (i.e. replace the gamble with its expected value).

Formally, let's define the certainty equivalent. Let $c(F, u)$ be the certainty equivalent for a person with Bernoulli utility function u facing lottery F , defined according to:

$$u(c(F, u)) = \int u(x) dF(x)$$

Although generally speaking people are risk averse, this is a behavioral postulate rather than an assumption or implication of our model. But, people need not be risk averse. In fact, we can divide utility functions into four classes:

1. Risk averse. A decision maker is risk averse if the expected utility of any lottery, F , is not more than the utility of the getting the expected value of the lottery for sure. That is, if:

$$\int u(x) dF(x) \leq u\left(\int x dF(x)\right) \text{ for all } F.$$

- (a) If the previous inequality is strict, we call the decision maker strictly risk averse.
- (b) Note also that since $u(c(F, u)) = \int u(x) dF(x)$ and $u(\cdot)$ is strictly increasing, an equivalent definition of risk aversion is that the certainty equivalent $c(F, u)$ is no larger than the expected value of the lottery, $\int x dF(x)$ for any lottery F .

2. Risk loving. A decision maker is risk loving if the expected utility of any lottery is not less than the utility of getting the expected value of the lottery for sure:

$$\int u(x) dF(x) \geq u\left(\int x dF(x)\right).$$

- (a) Strictly risk loving is when the previous inequality is strict.
- (b) An equivalent definition is that $c(F, u) \geq \int x dF(x)$ for all F .
3. Risk neutral. A decision maker is risk neutral if the expected utility of any lottery is the same as the utility of getting the expected value of the lottery for sure:

$$\int u(x) dF(x) = u\left(\int x dF(x)\right).$$

- (a) An equivalent definition is $c(F, u) = \int x dF(x)$.
4. None of the above. Many utility functions will not fit into any of the cases above. They'll be risk averse, risk loving, or risk neutral depending on the lottery involved.

Although many utility functions will fit into the “none of the above” category, risk aversion is by far the most natural way to think about actual people behaving, with the limiting case of risk neutrality. So, most of our attention will be focused on the cases of risk neutrality and risk aversion. Risk loving behavior does arise, but generally speaking people are risk averse, and so we start our study there.

Consider again the definition of risk aversion:

$$\int u(x) dF(x) \leq u\left(\int x dF(x)\right).$$

It turns out that this inequality is a version of Jensen's Inequality, which says that $h(\cdot)$ is a concave function if and only if

$$\int h(x) dF(x) \leq h\left(\int x dF(x)\right).$$

for all distributions $F(\cdot)$. Thus, risk aversion on the part of the decision maker is equivalent to having a concave Bernoulli utility function. Strict risk aversion is equivalent to having a strictly concave Bernoulli utility function.

Similarly, (strict) risk loving is equivalent to having a (strictly) convex Bernoulli utility function, and risk neutrality is equivalent to having a linear Bernoulli utility function.

The utility functions of risk averse and risk neutral decision makers are illustrated in MWG Figure 6.C.2. (panels a and b). In panel a, a risk averse consumer is diagrammed. Notice that with a strictly concave utility function, the expected utility of the lottery that offers 3 or 1 with equal probability is $\frac{1}{2}u(1) + \frac{1}{2}u(3) < u\left(\frac{1}{2}1 + \frac{1}{2}3\right)$. On the other hand, in panel b, $\frac{1}{2}u(1) + \frac{1}{2}u(3) = u\left(\frac{1}{2}1 + \frac{1}{2}3\right)$; the consumer is indifferent between the gamble and the sure thing. Thus

the manifestation of risk aversion in panel a is in the fact that the dotted line between $(1, u(1))$ and $(3, u(3))$ lies everywhere below the utility function.

To see if you understand, draw the diagram for a risk-loving decision maker, and convince yourself that $\frac{1}{2}u(1) + \frac{1}{2}u(3) > u(\frac{1}{2}1 + \frac{1}{2}3)$.

6.2.1 Measuring Risk Aversion: Coefficients of Absolute and Relative Risk Aversion

As we said, risk aversion is equivalent to concavity of the utility function. Thus one would expect that one utility function is “more risk averse” than another if it is “more concave.” While this is true, it turns out that measuring the risk aversion is more complicated than you might think (isn’t everything in this course?). Actually, it is only slightly more complicated.

You might be tempted to think that a good measure of risk aversion is that Bernoulli utility function $u_1()$ is more risk averse than Bernoulli utility function $u_2()$ if $|u_1''()| > |u_2''()|$ for all x . However, there is a problem with this measure, in that it is not invariant to positive linear transformations of the utility function. To see why, consider utility function $u_1(x)$, and apply the linear transformation $u_2() = au_1() + b$, where $a > 1$. We know that such a transformation leaves the consumer’s attitudes toward risk unchanged. However, $u_2''() = au_1''() > u_1''()$. Thus if we use the second derivative of the Bernoulli utility function as our measure of risk aversion, we find that it is possible for a utility function to be more risk averse than another, even though it represents the exact same preferences. Clearly, then, this is not a good measure of risk aversion.

The way around the problem identified in the previous paragraph is to normalize the second derivative of the utility function by the first derivative. Using $u_2()$ from the previous paragraph, we then get that:

$$\frac{u_2''()}{u_2'()} = \frac{au_1''()}{au_1'()} = \frac{u_1''()}{u_1'()}.$$

Thus this measure of risk aversion is invariant to linear transformations of the utility function. And, it’s almost the measure we will use. Because $u'' < 0$ for a concave function, we’ll multiply by -1 so that the risk aversion number is non-negative for a risk-averse consumer. This gives us the following definition:

Definition 10 *Given a twice-differentiable Bernoulli utility function $u()$, the Arrow-Pratt measure of absolute risk aversion is given by:*

$$r_A(x) = -\frac{u''(x)}{u'(x)}.$$

Note the following about the Arrow-Pratt (AP) measure:

1. $r_A(x)$ is positive for a risk-averse decision maker, 0 for a risk-neutral decision maker, and negative for a risk-loving decision maker.
2. $r_A(x)$ is a function of x , where x can be thought of as the consumer's current level of wealth. Thus we can admit the situation where the consumer is risk averse, risk loving, or risk neutral for different levels of initial wealth.
3. We can also think about how the decision maker's risk aversion changes with her wealth. How do you think this should go? Do you become more or less likely to accept a gamble that offers 100 with probability $\frac{1}{2}$ and -50 with probability $\frac{1}{2}$ as your wealth increases? Hopefully, you answered more. This means that you become less risk averse as wealth increases, and this is how we usually think of people, as having non-increasing absolute risk aversion.
4. The AP measure is called a measure of absolute risk aversion because it says how you feel about lotteries that are defined over absolute numbers of dollars. A gamble that offers to increase or decrease your wealth by a certain percentage is a relative lottery, since its prizes are defined relative to your current level of wealth. We also have a measure of **relative risk aversion**,

$$r_R(x) = -\frac{xu''(x)}{u'(x)}.$$

But, we'll come back to that later.

6.2.2 Comparing Risk Aversion

Frequently it is useful to know when one utility function is more risk averse than another. For example, risk aversion is important in the study of insurance, and a natural question to ask is how a person's desire for insurance changes as he becomes more risk averse. Fortunately, we already have the machinery in place for our comparisons. We say that utility function $u_2()$ is at least as risk averse as $u_1()$ if any of the following hold (in fact, they are all equivalent):

1. $c(F, u_2) \leq c(F, u_1)$ for all F .
2. $r_A(x, u_2) \geq r_A(x, u_1)$
3. $u_2()$ can be derived from $u_1()$ by applying an increasing, concave transformation, i.e., $u_2() = g(u_1(x))$, where $g()$ is increasing and concave. Note, this is what I meant when I said being

more risk averse is like being more concave. However, as you can see, this is not the most useful of the definitions we have come up with.

- Starting from any initial wealth position, x_0 , any lottery F that would be at least as good as x_0 for certain to a person with utility function $u_2()$ would also be acceptable to a person with utility function $u_1()$. That is,

$$\text{if } u_2(x_0) \leq \int u_2(x) dF(x), \text{ then } u_1(x_0) \leq \int u_1(x) dF(x).$$

Note that in MWG, they give definitions for “more risk averse” rather than “at least as risk averse.” Usually, you can go from what I say is “at least as risk averse” to something that is “more risk averse” by simply making the inequality strict for some value. That is, $u_2()$ is more risk averse than $u_1()$ if:

- $c(F, u_2) \leq c(F, u_1)$ for all F , with strict inequality for some F .
- $r_A(x, u_2) \geq r_A(x, u_1)$ for all x , with strict inequality for some x .
- $u_2()$ can be derived from $u_1()$ by applying an increasing, **strictly** concave transformation, i.e., $u_2() = g(u_1(x))$, where $g()$ is increasing and concave.
- Starting from any initial wealth position, x_0 , any lottery F that would be at least as good as x_0 for certain to a person with utility function $u_2()$ would be strictly preferred to x_0 for certain by a person with utility function $u_1()$. That is,

$$\text{if } u_2(x_0) \leq \int u_2(x) dF(x), \text{ then } u_1(x_0) < \int u_1(x) dF(x).$$

As usual, which definition is most useful will depend on the circumstances you are in. Practically, speaking, I think that number 3 is the least likely to come up, although it is useful in certain technical proofs. Note that it need not be the case that any two utility functions $u_2()$ and $u_1()$ are such that one is necessarily at least as risk averse as the other. In fact, the usual case is that you won't be able to rank them. However, most often we will be interested in finding out what happens to a particular person who becomes more risk averse, rather than actually comparing the risk aversion of two people.

In addition to knowing what happens when a person becomes more risk averse, we are also frequently interested in what happens to a person's risk aversion when her wealth changes. As I

mentioned earlier, the natural assumption to make (since it corresponds with how people actually seem to behave) is that people becomes less risk averse as they become wealthier. In terms of the measures we have for risk aversion, we say that a person exhibits non-increasing absolute risk aversion whenever $r_A(x)$ is non-increasing in x . In MWG Proposition 6.C.3, there are some alternate definitions. Figuring them out would be a useful exercise. Of particular interest, I think, is part iii), which says that having non-increasing (or decreasing) absolute risk aversion means that as your wealth increases, the amount you are willing to pay to get rid of a risk decreases. What does this say about insurance? Basically, it means that the wealthy will be willing to pay less for insurance and will receive less benefit from being insured. Formalizing this, let z be a random variable with distribution F and a mean of 0. Thus z is the prize of a lottery with distribution F . Let c_x (the certainty equivalent) be defined as:

$$u(c_x) = \int u(x + z) dF(z).$$

If the utility function exhibits decreasing absolute risk aversion, then $x - c_x$ (corresponding to the premium the person is willing to pay to get rid of the uncertainty) will be decreasing in x .

As before, it is natural to think of people as exhibiting nonincreasing relative risk aversion. That is, they are more likely to accept a proportional gamble as their initial wealth increases. Although the concept of relative risk aversion is useful in a variety of contexts, we will primarily be concerned with absolute risk aversion. One reason for this that many of the techniques we develop for studying absolute risk aversion translate readily to the case of relative risk aversion.

6.2.3 A Note on Comparing Distributions: Stochastic Dominance

We aren't going to spend a lot of time talking about comparing different distributions in terms of their risk and return because these are concepts that involve slightly more knowledge of probability and will most likely be developed in the course of any applications you see that use them. However, I will briefly mention them.

Suppose we are interested in knowing whether one distribution offers higher returns than another. There is some ambiguity as to what this means. Does it mean higher average monetary return (i.e., the mean of F), or does it mean higher expected utility? In fact, when a consumer is risk averse, a distribution with a higher mean may offer a lower expected utility if it is riskier. For example, a sufficiently risk averse consumer will prefer $x = 1.9$ for sure to a 50-50 lottery over 1 and 3. This is true even though the mean of the lottery, 2, is higher than the mean of the sure

thing, 1.9. Thus if we are concerned with figuring out which of two lotteries offers higher utility than another, simply comparing the means is not enough.

It turns out that the right concept to use when comparing the expected utility of two distributions is called **first-order stochastic dominance (FOSD)**. Consider two distribution functions, $F()$ and $G()$. We say that $F()$ first-order stochastically dominates $G()$ if $F(x) \leq G(x)$ for all x . That is, $F()$ FOSD $G()$ if the graph of $F()$ lies everywhere below the graph of $G()$. What is the meaning of this? Recall that $F(y)$ gives the probability that the lottery offers a prize that is less than or equal to y . Thus if $F(x) \leq G(x)$ for all x , this means that for any prize, y , the probability that $G()$'s prize is less than or equal to y is greater than the probability that $F()$'s prize is less than or equal to y . And, if it is the case that $F()$ FOSD $G()$, it can be shown that any consumer with a strictly increasing utility function $u()$ will prefer $F()$ to $G()$. That is, as long as you prefer more money to less, you will prefer lottery $F()$ to $G()$.

Now, it's important to point out that most of the time you will not be able to rank distributions in terms of FOSD. It will need not be the case that either $F()$ FOSD $G()$ or $G()$ FOSD $F()$. In particular, the example from two paragraphs ago (1.9 for sure vs. 1 or 3 with equal probability) cannot be ranked. As in the case of ranking the risk aversion of two utility functions, the primary use of this concept is in figuring out (in theory) how a decision maker would react when the distribution of prizes "gets higher." FOSD is what we use to capture the idea of "gets higher." And, knowing an initial distribution $F()$, FOSD gives us a good guide to what it means for the distribution to get higher: The new distribution function must lay everywhere below the old one.

So, FOSD helps us formalize the idea of a distribution "getting higher." In many circumstances it is also useful to have a concept of "getting riskier." The concept we use for "getting riskier" is called **second-order stochastic dominance (SOSD)**. One way to understand SOSD is in terms of mean preserving spreads. Let X be a random variable with distribution function $F()$. Now, for each value of x , add a new random variable z_x , where z_x has mean zero. Thus z_x can be thought of as a noise term, where the distribution of the noise depends on x but always has a mean of zero. Now consider the random variable $y = x + z_x$. Y will have the same mean as X , but it will be riskier because of all of the noise terms we have added in. And, we say that for any Y than has the same mean as X and can be derived from X by adding noise, Y is riskier than X . Thus, we say that X second-order stochastically dominates Y .

Let me make two comments at this point. First, as usual, it won't be the case that for any two random variables (or distributions), one must SOSD the other. In most cases, neither will SOSD.

Second, if you do have two distributions with the same mean, and one, say $F()$, SOSD the other, $G()$, then you can say that any risk averse decision maker will prefer $F()$ to $G()$. Intuitively, this is because $G()$ is just a noisy and therefore riskier version of $F()$, and risk-averse decision makers dislike risk.

OK. At this point let me apologize. Clearly I haven't said enough for you to understand FOSD and SOSD completely. But, I think that what we have done at this point is a good compromise. If you ever need to use these concepts, you'll know where to look. But, not everybody will have to use them, and using them properly involves knowing the terminology of probability theory, which not all of you know. So, at this point I think it's best just to put the definitions out there and leave you to learn more about them in the future if you ever have to.

6.3 Some Applications

6.3.1 Insurance

Consider a simple model of risk and insurance. A consumer has initial wealth w . With probability π , the consumer suffers damage of D . With probability $1 - \pi$, no damage occurs, and the consumer's wealth remains w . Thus, in the absence of insurance, the consumer's final wealth is $w - D$ with probability π , and w with probability $1 - \pi$.

Now, suppose that we allow the consumer to purchase insurance against the damage. Each unit of insurance costs q , and pays 1 dollar in the event of a loss. Let a be the number of units of insurance that the consumer purchases. In this case, the consumer's final wealth is $w - D + a - qa$ with probability π and $w - qa$ with probability $1 - \pi$. Thus the benefit of insurance is that it repays a dollars of the loss when a loss occurs. The cost is that the consumer must give up qa dollars regardless of whether the loss occurs. Insurance amounts to transferring wealth from the state where no loss occurs to the state where a loss occurs.

The consumer's utility maximization problem can then be written as:

$$\max_a \pi u(w - D + (1 - q)a) + (1 - \pi) u(w - qa).$$

The first-order condition for an interior solution to this problem is:

$$\pi(1 - q)u'(w - D + (1 - q)a^*) - (1 - \pi)qu'(w - a^*q) = 0$$

Let's assume for the moment that the insurance is fairly priced. That means that the cost to the consumer of 1 dollar of insurance is just the expected cost of providing that coverage; in insurance

jargon, this is called “actuarially fair” coverage. If the insurer must pay 1 dollar with probability π , then the fair price of insurance is $\pi * 1 = \pi$. Thus for the moment, let $q = \pi$, and the first-order condition becomes:

$$u'(w - D + (1 - \pi)a^*) = u'(w - \pi a^*).$$

Now, if the consumer is strictly risk averse, then $u'()$ is strictly decreasing, which means that in order for the previous equation to hold, it must be that:

$$w - D + (1 - \pi)a^* = w - \pi a^*.$$

This means that the consumer should equalize wealth in the two states of the world. Solving further,

$$D = a^*.$$

Thus, a utility-maximizing consumer will purchase insurance that covers the full extent of the risk - “full insurance” - if it is priced fairly.

What happens if the insurance is priced unfairly? That is, if $q > \pi$? In this case, the first-order condition becomes

$$\begin{aligned} \pi(1 - q)u'(w - D + (1 - q)a^*) - (1 - \pi)qu'(w - qa^*) &= 0 \text{ if } D > a^* > 0, \\ \pi(1 - q)u'(w - D) - (1 - \pi)qu'(w) &\leq 0 \text{ if } a^* = 0, \\ \pi(1 - q)u'(w - qD) - (1 - \pi)qu'(w - qD) &\geq 0 \text{ if } a^* = D. \end{aligned}$$

Now, if we consider the case where $a^* = D$, we derive the optimality condition for purchasing full insurance. Then,

$$u'(w - qD)(\pi(1 - q) - (1 - \pi)q) = u'(w - qD)(\pi - q) \geq 0.$$

Thus if the consumer is able to choose how much insurance she wants, she will never choose to fully insure when the price of insurance is actuarially unfair (since the above condition only holds if $\pi \geq q$, but by definition, unfair pricing means $q > \pi$).

There is another way to see that if insurance is priced fairly the consumer will want to fully insure. If it is actuarially fair, the price of full insurance is πD . Thus if the consumer purchases full insurance, her final wealth is $w - \pi D$ with probability 1, and her expected utility is $u(w - \pi D)$. If she purchases no insurance, her expected utility is:

$$\pi u(w - D) + (1 - \pi)u(w),$$

which, by risk aversion, is less than the utility of the expected outcome, $\pi(w - D) + (1 - \pi)w = w - \pi D$. Thus

$$u(w - \pi D) > \pi u(w - D) + (1 - \pi)u(w).$$

So, any risk-averse consumer, if offered the chance to buy full insurance at the actuarially fair rate, would choose to do so.

What is the largest amount of money that the consumer is willing to pay for full insurance, if the only other option is to remain without any insurance? This is found by looking at the consumer's certainty equivalent. Recall that the certainty equivalent of this risk, call it ce , solves the equation:

$$u(ce) = \pi u(w - D) + (1 - \pi)u(w).$$

Thus ce represents the smallest sure amount of wealth that the consumer would prefer to the lottery. From this, we can compute the maximum price she would be willing to pay as:

$$ce = w - p_{\max}.$$

The last two results in this example may be a bit confusing, so let me summarize. First, if the consumer is able to choose how much insurance she wants, and insurance is priced fairly, she will choose full insurance. But, if the consumer is able to choose how much insurance she wants and insurance is priced unfairly, she will choose to purchase less than full insurance. However, if the consumer is given the choice only between full insurance or no insurance, she will be willing to pay up to $p_{\max} = w - ce$ for insurance.

6.3.2 Investing in a Risky Asset: The Portfolio Problem

Suppose the consumer has utility function $u(\cdot)$ and initial wealth w . The consumer must decide how much of her wealth to invest in a riskless asset and how much to invest in a risky asset that pays 0 dollars with probability π and r dollars with probability $1 - \pi$. Let x be the number of dollars she invests in the risky asset. Note, for future reference, that the expected value of a dollar invested in the risky asset is $r(1 - \pi)$. The riskless asset yields no interest or dividend - its worth is simply \$1 per unit. The consumer's optimization problem is:

$$\max_x \pi u(w - x) + (1 - \pi)u(w + (r - 1)x).$$

The first-order condition for this problem is:

$$\begin{aligned} & \leq 0 && \text{if } x^* = 0 \\ -\pi u'(w-x) + (1-\pi)(r-1)u'(w+(r-1)x) & = 0 && \text{if } 0 < x^* < w \\ & \geq 0 && \text{if } x^* = w. \end{aligned}$$

The question we want to ask is when will it be the case that the consumer does not invest in the risky asset. That is, when will $x^* = 0$? Substituting x^* into the first-order condition yields:

$$\begin{aligned} -\pi u'(w) + (1-\pi)(r-1)u'(w) & \leq 0 \\ u'(w)(-\pi + (1-\pi)(r-1)) & \leq 0. \end{aligned}$$

But, since $u'(w) > 0$,⁴ for this condition to hold, it must be that:

$$-\pi + (1-\pi)(r-1) \leq 0$$

or

$$(1-\pi)r \leq 1.$$

Thus the only time it is optimal for the consumer not to invest in the risky asset at all is when $(1-\pi)r \leq 1$. But, note that $(1-\pi)r$ is just the expected return on the risky asset and 1 is the return on the safe asset. Hence only when the expected return on the risky asset is less than the return on the safe asset will the consumer choose not to invest at all in the risky asset. Put another way, whenever the expected return on the risky asset is greater than the expected return on the safe asset (i.e., it is actuarially favorable), the consumer will choose to invest at least some of her wealth in the risky asset. In fact, you can also show that if the consumer chooses $x^* > 0$, then $(1-\pi)r > 1$.

Assuming that the consumer chooses to invest $0 < x^* < w$ in the risky asset, which implies that $(1-\pi)r > 1$, we can ask what happens to investment in the risky asset when the consumer's wealth increases. Let $x(w)$ solve the following identity:

$$-\pi u'(w-x(w)) + (1-\pi)(r-1)u'(w+(r-1)x(w)) = 0$$

Differentiate with respect to w :

$$-\pi u''(w-x(w))(1-x'(w)) + (1-\pi)(r-1)u''(w+(r-1)x(w))(1+(r-1)x'(w)) = 0$$

⁴ $u'() > 0$ by assumption, since more money is better than less money (i.e. the marginal utility of money is always positive).

and to keep things simple, let $a = w - x(w)$ and $b = w + (r - 1)x(w)$. Solve for $x'(w)$

$$x'(w) = \frac{\pi u''(a) - (1 - \pi)(r - 1)u''(b)}{\pi u''(a) + (1 - \pi)(r - 1)^2 u''(b)}.$$

By concavity of $u(\cdot)$, the denominator is negative. Hence the sign of $x'(w)$ is opposite the sign of the numerator. The numerator will be negative whenever:

$$\begin{aligned} \pi u''(a) &< (1 - \pi)(r - 1)u''(b) \\ \frac{u''(a)}{u''(b)} &> \frac{(1 - \pi)(r - 1)}{\pi} \\ \frac{u''(a)}{u''(b)} \frac{u'(b)}{u'(a)} &> \frac{(1 - \pi)(r - 1)}{\pi} \frac{u'(b)}{u'(a)} = 1 \end{aligned}$$

where the inequality flip between the first and second lines follows from $u'' < 0$, and the equality in the last line follows from the first-order condition, $-\pi u'(a) + (1 - \pi)(r - 1)u'(b) = 0$. Thus this inequality holds whenever

$$\frac{u''(a)}{u'(a)} < \frac{u''(b)}{u'(b)},$$

(with another inequality flip because $\frac{u''}{u'} < 0$) or, multiplying by -1 (and flipping the inequality again),

$$r_A(w - x(w)) > r_A(w - x(w) + rx(w)).$$

So, the consumer having decreasing absolute risk aversion is sufficient for the numerator to be negative, and thus for $x'(w) > 0$. Thus the outcome of all of this is that whenever the consumer has decreasing absolute risk aversion, an increase in wealth will lead her to purchase more of the risky asset. Or, risky assets are normal goods for decision makers with decreasing absolute risk aversion. This means that the consumer will purchase more of the risky asset in absolute terms (total dollars spent), but not necessarily relative terms (percent of total wealth).

The algebra is a bit involved here, but I think it is a good illustration of the kinds of conclusions that can be drawn using expected utility theory. Notice that the conclusion is phrased, “If the decision maker exhibits decreasing absolute risk aversion...”, which is a behavioral postulate, not an implication of the theory of choice under uncertainty. Nevertheless, we believe that people do exhibit decreasing absolute risk aversion, and so we are willing to make this assumption.

6.4 Ex Ante vs. Ex Post Risk Management

Consider the following simple model, which develops a “separation” result in the choice under uncertainty model. Suppose, as before, the consumer has initial wealth w , and there are two states

of the world. With probability π , the consumer loses D dollars and has final wealth $w - D$. We call this the loss state. With probability $(1 - \pi)$, no loss occurs, and the consumer has final wealth w . Suppose there are perfect insurance markets, meaning that insurance is priced fairly: insuring against a loss of \$1 that occurs with probability π costs π (we showed this result earlier). Thus, by paying π dollars in both states, the consumer increases wealth in the loss state by 1 dollar. The net change is therefore that by giving up π dollars of consumption in the no-loss state, the consumer can gain $1 - \pi$ dollars of consumption in the loss state. Notice, however, that any such transfer (for any level of insurance, $0 \leq a \leq D$) keeps expected wealth constant:

$$\begin{aligned} & \pi(w - D + (1 - \pi)a) + (1 - \pi)(w - \pi a) \\ = & w - \pi D \end{aligned}$$

So, another way to think of the consumer's problem is choosing how to allocate consumption between the loss and no-loss states while keeping expected expenditure constant at $w - \pi D$. That is, another way to write the consumer's problem is:

$$\begin{aligned} & \max_{c_L, c_N} \pi u(c_L) + (1 - \pi) u(c_N) \\ \text{s.t.} \quad & \pi c_L + (1 - \pi) c_N \leq w - \pi D. \end{aligned}$$

This is where the separation result comes in. Notice that expected initial wealth only enters into this consumer's problem on the right-hand side (just as present value of lifetime income only entered into the right-hand side of the intertemporal budget constraint in the Fisher theorem, and profit from the farm only entered into the right-hand side of the farm's budget constraint in the agricultural household separation result).

So, suppose the consumer needs to choose between different vectors of state dependent income. That is, suppose $w = (w_L, w_N)$ and $w' = (w'_L, w'_N)$. How should the consumer choose between the two? The answer is that, if insurance is fairly priced, the consumer should choose the income vector with the highest expected value (i.e., w if $\pi w_L + (1 - \pi) w_N > \pi w'_L + (1 - \pi) w'_N$, w' if the inequality is reversed, either one if expected wealth is equal in the two states).

What happens if insurance markets are not perfect? Let's think about the situation where there are no insurance possibilities at all. In this case, the only way the consumer can reduce risk is to choose projects (i.e., state-contingent wealth vectors) that are, themselves, less risky. The result is that the consumer may choose a low-risk project that offers lower expected return, even though, if the consumer could access insurance markets, she would choose the risky project and reduce risk through insurance.

There are a number of applications of this idea. Consider, for example, a farmer who must decide whether to produce a crop that may not come in but earns a very high profit if it does, or a safe crop that always earns a modest return. It may be that if the farmer can access insurance markets, the optimal thing for him to do is to take the risky project and then reduce his risk by buying insurance. However, if insurance is not available, the farmer may choose the safer crop. The result is that the farmer consistently chooses the lower-value crop and earns lower profit in the long run than if insurance were available. A similar story can be told where the risky project is a business venture. In the absence of insurance, small businesses may choose safe strategies that ensure that the company stays in business but do significantly worse over the long run.

The distinction raised in the previous paragraph is one of *ex ante* vs. *ex post* risk reduction. *Ex post* risk reduction means buying insurance against risk. *Ex ante* risk reduction means engaging in activities up front that reduce risk (perhaps through choosing safer projects). In a very general way, if there are not good mechanisms for *ex post* risk reduction (either private or government insurance), then individuals will engage in excessive *ex ante* risk reduction. The result is that individuals are worse off than if there were good insurance markets, and society as a whole generates less output.

Chapter 7

Competitive Markets and Partial Equilibrium Analysis

Up until now we have concentrated our efforts on two major topics - consumer theory, which led to the theory of demand, and producer theory, which led to the theory of supply. Next, we will put these two parts together into a market. Specifically, we will begin with competitive markets. The key feature of a competitive market is that producers and consumers are considered **price takers**. That is, individual actors can buy or sell as much of the output as they want at the market price, but no one can take any unilateral action to affect the price. If this is the case, then the actors take prices as exogenous when making their decisions, which was a key feature in our analysis of consumer and producer behavior. Later, when we study monopoly and oligopoly, we will relax the assumption that firms cannot affect prices.

Our main goal here will be to determine how supply and demand interact to determine the way the market allocates society's resources. In particular, we will be concerned with:

1. When does the market allocate resources efficiently?
2. When, if the government wants to implement a specific allocation, can the allocation can be implemented using the market (possibly by rearranging people's initial endowments)?
3. Why does the market sometimes fail to allocate resources efficiently, and what can be done in such cases?

The third question will be the subject of the next chapter, on externalities and public goods. For now, we focus on the first and second questions, which bring us to the first and second fundamental

theorems of welfare economics, respectively.

7.1 Competitive Equilibrium

The basic idea in the analysis of competitive equilibrium is the “law of supply and demand.” Utility maximization by individual consumers determines individual demand. Summing over individual consumers determines aggregate demand, and the aggregate demand curve slopes downward. Profit maximization by individual firms determines individual supply, and summing over firms determines aggregate supply, which slopes upward. Adam Smith’s invisible hand acts to bring the market to the point where the two curves cross, i.e. supply equals demand. This point is known as a competitive equilibrium, and it tells us how much of the output will be produced and the price that will be charged for it.

Notation

We are going to be dealing with many consumers, many producers, and many commodities. To make things clear, I’ll denote which consumer or producer we are talking about with a superscript. For example, u^i is the utility function of consumer i , x^i is the commodity bundle chosen by consumer i , and y^j is the production plan chosen by firm j . For vectors x^i and y^j , I’ll denote the l^{th} component with a subscript. Hence $x^i = (x_1^i, \dots, x_L^i)$, and $y^j = (x_1^j, \dots, x_L^j)$. So x_L^j refers to consumer j ’s consumption of good L . This differs from MWG, which uses double subscripts. But, I think that this is clearer.

7.1.1 Allocations and Pareto Optimality

Our formal analysis of competitive markets begins with defining an allocation, and determining what we mean when we say that an allocation is efficient.

Consider an economy consisting of:

1. I consumers each with utility function u^i
2. J firms each maximizing its profit
3. L commodities.

Initially, there are $w_l \geq 0$ units of commodity l available. This societal endowment can either be consumed or used to produce other commodities. Because of this, it is most convenient to use the production plan/ net output vector approach to producer theory.¹

Each firm has production set Y_j and chooses production plan $y^j \in Y_j$ in order to maximize profit. Let y_l^j be the quantity of commodity l produced by firm j . Thus if each firm produces net-input vector y^j , the total amount of good l available for consumption in the economy is given by

$$w_l + \sum_j y_l^j.$$

The possible outcomes in this economy are called **allocations**. An allocation is a consumption vector $x^i \in X^i$ for each consumer i , and a production vector $y^j \in Y_j$. An allocation is **feasible** if

$$\sum_i x_l^i \leq w_l + \sum_j y_l^j.$$

for every l . That is, if total consumption of each commodity is no larger than the total amount of that commodity available.

Again, one of the things we will be most interested in is efficiency. In the context of producer theory, we considered productive efficiency, the question of whether firms choose production plans that are not wasteful. Currently, we are interested not only in productive efficiency but in consumption efficiency as well. That is, we are concerned that, given the availability of commodities in the economy, the commodities are allocated to consumers in such a way that no other arrangement could make everybody better off. The concept of “making everybody better off” is formalized by **Pareto optimality**.

Pareto Optimality

When an economist talks about efficiency, we refer to situations where no one can be made better off without making some one else worse off. This is the notion of **Pareto optimality**.

Formally, a feasible allocation (x, y) is **Pareto optimal** if there is no other feasible allocation (x', y') such that $u_i(x'_i) \geq u_i(x_i)$ for all i , with strict inequality for at least one i . Thus a Pareto optimal allocation is efficient in the sense that there is no other way to reorganize society’s productive facilities in order to make somebody better off without harming somebody else. Notice that we don’t care about producers in this definition of Pareto optimality. This is okay, because

¹Recall that in this approach, inputs enter into the production plan as negative elements.




Figure 7.1: Utility Possibility Frontier

all commodities will in the end find their way into the hands of consumers. A profit-maximizing firm will never buy inputs it doesn't use or produce output it doesn't sell, and firms are owned by consumers, so profit eventually becomes consumer wealth. Thus, looking at the utility of consumers fully captures the notion of efficiency.

If you draw the utility possibility frontier in two dimensions, as in Figure 7.1, Pareto optimal points are ones that lay on the northeast frontier. Note that Pareto optimality doesn't say anything about equity. An allocation that gives one person everything and the other nothing may be Pareto optimal. However, it is not at all equitable. Much of the job in policy making is in striking a balance between equity and efficiency – to put it another way, choosing the equitable point from among the efficient points.

7.1.2 Competitive Equilibria

We now turn to investigating competitive equilibria with the goal of determining whether or not the allocations determined by the market will be Pareto optimal. Again, we are concerned with competitive markets. Thus buyers and sellers are price takers in the L commodities. Further, we make the assumption that the firms in the market are owned by the consumers. Thus all profits from operation of the firms are redistributed back to the consumers. Consumers can then use this wealth to increase their consumption. In this way we “close” the model - it's entirely self contained.

Although our formal analysis will be of a partial equilibrium system, where we study only one or two markets, we will define an competitive equilibrium over all L commodities. In a competitive

economy, a market exists for each of the L goods, and all consumer and producers act as price takers. As usual, we'll let the vector of the L commodity prices be given by p , and suppose consumer i has endowment w_i^i of good i . We'll denote a consumer's entire endowment vector by w^i , and the total endowment of the good is given by $\sum_i w_i^i = w_l$.

We formalize the fact that consumers own the firms by letting θ_j^i ($0 \leq \theta_j^i \leq 1$) be the share of firm j that is owned by consumer i . Thus if firm j chooses production plan y^j , the profit earned by firm j is $\pi_j = p \cdot y^j$, and consumer i 's share of this profit is given by $\theta_j^i (p \cdot y^j)$. Consequently, consumer i 's total wealth is given by $p \cdot w_i + \sum \theta_j^i \pi_j$. Note that this means that all wealth is either in the form of endowment or firm share; there is no longer any exogenous wealth w . Of course, this depends on firms' decisions, but part of the idea of the equilibrium is that production, consumption, and prices will all be simultaneously determined. We now turn to the formal description of a competitive equilibrium.

There are three requirements for a competitive equilibrium, corresponding to the requirements that producers optimize, consumers optimize, and that "markets clear" at the equilibrium prices. An equilibrium will then consist of a production plan y^{j*} for each firm, a consumption vector x^{i*} for each consumer, and a price vector p^* .

Actually, the producer and consumer parts are just what we have been studying for the first half of the course. The market clearing condition says that at the equilibrium price, it must be that the aggregate supply of each commodity equals the aggregate demand for that commodity, when producers and consumers optimize. Formally, these requirements are:

1. Profit Maximization: For each firm, $y^j(p)$ solves

$$\max p y^j \text{ subject to } y^j \in Y_j.$$

2. Utility Maximization: For each consumer, $x^i(p)$ solves

$$\begin{aligned} & \max u^i(x^i) \text{ subject to} \\ & p \cdot x^i \leq p \cdot w^i + \sum \theta_j^i (p \cdot y^j(p)), \end{aligned}$$

where θ_j^i is consumer i 's ownership share in firm j . Note: this is just the normal UMP with the addition of the idea that the consumer has ultimate claim on the profit of the firm.²

²There is something a little strange here. Note that we won't know the firm's profit until after the price vector is determined. But, if we don't know the firm's profit, we can't derive consumers' demand functions, and so we can't

3. Market Clearing. For each good, p^* is such that

$$\sum_{i=1}^I x_l^i(p^*) = w_l + \sum_{j=1}^J y_l^j(p^*)$$

Of course, we must keep in mind that x^* and y^* will be a function of p . Thus operationally, the requirements for an equilibrium can be written as:

1. For each consumer, $x^{i*}(p, w^i, \theta^i)$ solves the UMP. Add up the individual demand curves to get aggregate demand, $D(p)$, as a function of prices.
2. For each firm, $y^{j*}(p)$ solves the PMP. Add up the individual supply curves to get aggregate supply, $S(p)$, as a function of prices.
3. Find the price where $D(p^*) = S(p^*)$.

The last step is the one that you are familiar with from intermediate micro. The first two steps are what we have developed so far in this course. Note that for consumers we will generally need to worry about aggregation issues. However, if consumer preferences take the Gorman form, things will aggregate nicely.

Since $x^i(p)$ and $y^j(p)$ are the demand and supply curves, and we know that these functions are homogeneous of degree zero in prices, we know that **if p^* induces a competitive equilibrium, αp^* also induces a competitive equilibrium for any $\alpha > 0$** . This allows us to normalize the prices without loss of generality, and we will usually do so by setting the price of good 1 equal to 1.

Although we will soon be working with only one or two markets, so far we have been thinking about an economy with L markets. It can be shown (MWG Lemma 10.B.1) that if you know that $L - 1$ of the market clear at price p^* , then the L^{th} market must clear as well, provided that consumers satisfy Walras Law and $p^* \gg 0$. That is, if

$$\sum_i x_l^i(p^*) = w_l + \sum_j y_l^j(p^*) \text{ for } \forall l \neq k,$$

then

$$\sum_i x_k^i(p^*) = w_k + \sum_j y_k^j(p^*) .$$

solve the UMP! Actually, this isn't really a problem. The difficulty arises from trying to put a dynamic interpretation on a static model. Really, what we are after is the price which, if it were to come about, would lead to equilibrium behavior. No agent would have any incentive to change what he/she/it is doing. The neoclassical equilibrium model doesn't say anything about how such an equilibrium comes about. Only that if it does, it is stable.

This lemma is a direct consequence of the idea that total wealth must be preserved in the economy. The nice thing about it is that when you are only studying two markets, as we do in the partial equilibrium approach, you know that if one market clears, the other must clear as well. Hence the study of two markets really reduces to the study of one market.

7.2 Partial Equilibrium Analysis

7.2.1 Set-Up of the Quasilinear Model

We now turn away from the general model to a simple case, known as Partial Equilibrium. It is ‘partial’ because we focus on a small part of the total economy, often on a two commodity world. We laid the groundwork for this type of approach in our discussion of consumer theory. If we are interested in studying a particular market, say the market for apples, we can make the assumption that the prices of all other commodities move in tandem. This justifies, through use of the composite commodity theorem, treating consumers as if they have preferences over apples and “everything else.” Hence we have justified a two-commodity model for this situation. Next, since each consumer’s expenditure on apples is likely to be only a small part of her total wealth, it is reasonable to think of there being no wealth effects on consumers’ demand for apples. And, recall, that quasilinear preferences correspond to the case where there are no wealth effects in the non-numeraire good. So, basically what we’ll do in our partial equilibrium approach (and what is implicitly underlying the approach you took in intermediate micro) is assume that there are two goods: a composite commodity (the numeraire) whose price is set equal to 1, and the good of interest. We’ll call the numeraire m (for “money”) and the good we are interested in x .

Now, we can set up the following simple model. Let x_i and m_i be consumer i ’s consumption of the commodity of interest and the numeraire commodity, respectively.³ Assume that each consumer has quasilinear utility of the form:

$$u^i(m_i, x_i) = m_i + \phi_i(x_i).$$

Further, we normalize $u^i(0, 0) = \phi_i(0) = 0$, and assume that $\phi_i' > 0$ and $\phi_i'' < 0$ for all $x_i \geq 0$. That is, we assume that the consumer’s utility is increasing in the consumption of x and that her marginal utility of consumption is decreasing.

³This is a change in notation from the set-up at the beginning of the chapter. Now, the subscripts refer to the consumer / firm, rather than the commodity.

Since we already set the price of m equal to 1, we only need to worry about the price of x . Denote it by p .

There are J firms in the economy. Each firm can transform m into x according to cost function $c_j(q_j)$, where q_j is the quantity of x that firm j produces, and $c_j(q_j)$ is the number of units of the numeraire commodity needed to produce q_j units of x . Thus, letting z_j denote firm j 's use of good m as an input, its technology set is therefore

$$Y_j = \{(-z_j, q_j) \mid q_j \geq 0 \text{ and } z_j \geq c_j(q_j)\}.$$

That is, you have to spend enough of good m to produce q_j units of x . We will assume that $c_j(q_j)$ is strictly increasing and convex for all j .

In order to solve the model, we also need to specify consumers' initial endowments. We assume there is no initial endowment of x , but that consumer i has endowment of m equal to $w_{mi} > 0$ and the total endowment is $\sum_i w_{mi} = w_m$.

7.2.2 Analysis of the Quasilinear Model

That completes the set-up of the model. The next step is to analyze it. Recall that in order to find an equilibrium, we need to derive the firms' supply functions, the consumers' demand functions, and find the market-clearing price.

1. **Profit maximization.** Given the equilibrium price p^* , firm j 's equilibrium output q_j^* must maximize

$$\max_{q_j} pq_j - c_j(q_j)$$

which has the necessary and sufficient first-order condition

$$p^* \leq c'_j(q_j^*)$$

with equality if $q_j^* > 0$.

2. **Utility Maximization:** Consumer i 's equilibrium consumption vector (x_i^*, m_i^*) must maximize

$$\begin{aligned} & \max m_i + \phi_i(x_i) \\ \text{s.t. } & m_i + p^*x_i \leq w_{mi} + \sum_j \theta_j^i (p^*q_j^* - c_j(q_j^*)) \end{aligned}$$

- We know that the budget constraint must hold with equality. We can substitute it into the objective function for m_i , which yields:

$$\max \phi_i(x_i) - p^* x_i + \left[w_{mi} + \sum \theta_j^i (p^* q_j^* - c_j(q_j^*)) \right].$$

The first-order condition is:

$$\phi_i'(x_i^*) \leq p^*$$

which holds with equality if $x_i^* > 0$.

3. Market clearing. Remember that lemma that said if one of the markets clears we know that the other one must clear as well? We will use that to formulate a plan of attack here. Basically, we will find a price vector such that aggregate demand for x equals aggregate supply of q , $\sum_i x_i(p^*) = \sum_j q_j(p^*)$, i.e. the market for the consumption commodity clears. Then, we'll use the budget equation to compute the equilibrium level of m_i for each consumer (since the lemma tells us that the market for the numeraire must clear as well). To begin, assume an interior solution to the UMP and PMP for each consumer and firm. Then p^* , q_j^* , and x_i^* must solve the system of equations

$$\begin{aligned} p^* &= c_j'(q_j(p)) \text{ for all } j \\ p^* &= \phi_i'(x_i(p)) \text{ for all } i \\ \sum_i x_i(p^*) &= \sum_j q_j(p^*) \end{aligned}$$

Notice that the first j equations determine each firm's supply function. We can then add them to get aggregate supply, the RHS of the third equation. The next i equations determine each consumer's demand function. We can add them up to get aggregate demand, which is the LHS of the third equation. The third equation is thus the requirement that at the equilibrium price supply equals demand.

Notice that the equilibrium conditions involve neither the initial endowments of the consumers nor their ownership shares. Thus the equilibrium allocation of x and the price of x are independent of the initial conditions. This follows directly from the assumption of quasilinear utility. However, since equilibrium allocations of the numeraire are found by using each consumer's budget constraint, the equilibrium allocations of the numeraire will depend on initial endowments and ownership shares.

From a graphical point of view, the partial equilibrium is as follows:

1. For each consumer, derive their Walrasian demand for the consumption good, $x_i(p)$. Add across consumers to derive the aggregate demand, $x(p) = \sum_i x_i(p)$. Since each demand curve is downward sloping, the aggregate demand curve will be downward sloping. Graphically, this addition is done by adding the demand curves “horizontally” (as in MWG Figure 10.C.1). Since individual demand curves are defined by the relation:

$$p = \phi'_i(q),$$

The price at which each individual’s demand curve intersects the vertical axis is $\phi'_i(0)$, and gives that individual’s marginal willingness to pay for the first unit of output. The intercept for the aggregate demand curve is therefore $\max_i \phi'_i(0)$. Hence if different consumers have different $\phi_i()$ functions, not all demand curves will have the same intercept, and the demand curve will become flatter as price decreases.

2. For each firm, derive the supply curve for the consumption good, $y_j(p)$. Add across firms to derive the aggregate supply, $y(p) = \sum_j y_j(p)$. For each firm, the supply curve is given by:

$$p = c'_j(q_j).$$

Thus each firm’s supply curve is the inverse of its marginal cost curve. Since we have assumed that $c''_j() \geq 0$, the supply curve will be upward sloping or flat. Again, addition is done by adding the supply curves horizontally, as in MWG Figure 10.C.2. The intercept of the aggregate supply curve will be the smallest $c'_j(0)$. If firms’ cost functions are strictly convex, aggregate supply will be upward sloping.

3. Find the price where supply equals demand: find p^* such that $x(p^*) = y(p^*)$. Since the market clears for good l , it must also clear for the numeraire. The equilibrium point will be at the price and quantity where the supply and demand curves cross.

At this point, we can talk a bit about the dynamics of how an equilibrium might come about. This is the story that is frequently told in intermediate micro courses, and I should point out that it is just a story. There is nothing in the model which justifies this approach since we have said nothing at all about how markets will behave if they are out of equilibrium. Nevertheless, I’ll tell the story.

Suppose p^* is the equilibrium price of x , but that currently the price is equal to $p^+ > p^*$. At this price, aggregate demand is less than aggregate supply: $D(p^+) < S(p^+)$. Because of this, there

is a “glut” on the market. There are more units of x available for sale than people willing to buy (think about cars sitting on a car lot at the end of the model year). Hence (so the story goes) there will be downward pressure on the price as suppliers lower their price in order to induce people to buy. As the price declines, supply decreases and demand increases until we reach equilibrium at $p = p^*$. Similarly, if initially the price p^- is such that $p^- < p^*$, then $D(p^-) > S(p^-)$. There is excess demand (think about the hot toy of the holiday season). The excess demand bids up the price as people fight to get one of the scarce units of x , and as the price rises, supply increases and demand decreases until equilibrium is reached, once again, at $p = p^*$.

Thus we have the “invisible hand” of the market working to bring the market into equilibrium. However, let me emphasize once again that stories such as these are not part of our model.

7.2.3 A Bit on Social Cost and Benefit

The firm’s supply function is $y_j(p)$ and satisfies $p = c'_j(y_j(p))$. Thus at any particular price, firms choose their quantities so that the marginal cost of producing an additional unit of production is exactly equal to the price. Similarly, the consumer’s demand function is $x_i(p)$ such that $p = \phi'_i(x_i(p))$. Thus at any price, consumers choose quantities so that the marginal benefit of consuming an additional unit of x is exactly equal to its price.⁴ When both firms and consumers do this, we get that, at equilibrium, the marginal cost of producing an additional unit of x is exactly equal to the marginal utility of consuming an additional unit of x . This is true both individually and in the aggregate. Thus at the equilibrium price, all units where the marginal social cost is less than or equal to the marginal social benefit are produced and consumed, and no other units are. Thus the market acts to produce an efficient allocation. We’ll see more about this in a little while, but I wanted to suggest where we are going before we take a moment to talk about a few other things.

7.2.4 Comparative Statics

As usual, one of the things we will be interested in determining in the partial equilibrium model is how the endogenous parameters of the model vary with changes in the environment. For example, suppose that a consumer’s utility function depends on a vector of exogenous parameters, $\phi_i(x_i, \alpha)$,

⁴This is true as long as we assume that for any level of output, consumers with the highest willingness to pay (i.e., marginal benefit) are the ones that are given the units of output to consume, which is a reasonable assumption in many circumstances.




Figure 7.2: Partial Equilibrium with a Tax

and a firm's cost function depends on another vector of parameters (possibly overlapping), $c_j(q_j, \beta)$. Note that α and β will include at least the prices of the commodities, but may include other things such as tax rates, taste parameters, etc. We can ask how the equilibrium prices and quantities change with a change in α or β .

One of the most studied situations of this type is the impact of a tax on good x . Suppose, for example, that the government collects a tax of t on each unit of output purchased. Let the price consumers pay be p_c and the price producers receive be p_f . Note that consumers only care about the price they have to pay to acquire the good, and producers only care about the price they receive when they sell a unit of the good. In particular, neither side of the market cares directly about the tax. The two prices are related by the tax rate:

$$p_c = p_f + t.$$

The equilibrium in this market is the point where supply equals demand, i.e., prices p_c and p_f such that

$$x(p_c) = q(p_f), \text{ and}$$

$$p_c = p_f + t.$$

Or,

$$x(p_f + t) = q(p_f).$$

The equilibrium is depicted in Figure 7.2:

Q^t is the quantity sold after the tax is implemented. At this quantity, the difference between the price paid by consumers and the price received by firms is exactly equal to the tax. Note that at Q^t , the marginal social cost of an additional unit of output is less than the marginal social benefit. Hence society could be made better off if additional units of output were produced and sold, all the way up to the point where Q^* units of output are produced. The loss suffered by society due to the fact that these units are not produced and consumed is called the deadweight loss (DWL) of taxation.

One question we may be interested in is how the price paid by consumers changes when the size of the tax increases. Let $p^*(t)$ be the equilibrium price received by firms. Thus consumers pay $p^*(t) + t$. The following identity holds for any tax rate t :

$$x(p^*(t) + t) \equiv q(p^*(t)).$$

Totally differentiating with respect to t yields:

$$\begin{aligned} x'(p^* + t)(p'(t) + 1) &= q'(p^*)p'(t) \\ p'(t) &= \frac{-x'(p^* + t)}{x'(p^* + t) - q'(p^*)} \end{aligned}$$

The numerator is positive by definition. Since $q'(p^*)$ is positive, the absolute value of the denominator is larger than the absolute value of the numerator, hence $-1 < p'(t) < 0$. This implies that as the tax rate increases, the price received by firms decreases, but by less than the full amount of the tax. As a consequence, the price paid by consumers must also increase, but by less than the increase in the tax. Further, the total quantity must decrease as well.

Consider the formula we just derived:

$$p'(t) = \frac{-x'(p^* + t)}{x'(p^* + t) - q'(p^*)}.$$

Evaluate at $t = 0$, and rewrite this in terms of elasticities:

$$\begin{aligned} \frac{dp}{dt} &= \frac{-\frac{dx}{dp} \frac{p}{x}}{\frac{dx}{dp} \frac{p}{x} - \frac{dq}{dp} \frac{p}{q}} \\ &= -\frac{|\varepsilon_d|}{|\varepsilon_d| + \varepsilon_s} \end{aligned}$$

where ε_d is the elasticity of demand and ε_s is the elasticity of supply. This says that the proportion of a small tax that is passed onto producers in the form of lower prices is given by $\frac{|\varepsilon_d|}{|\varepsilon_d| + \varepsilon_s}$. The proportion passed onto consumers in the form of higher prices is given by $\frac{\varepsilon_s}{|\varepsilon_d| + \varepsilon_s}$.

Now, suppose the government is considering two different tax programs: one that taxes a commodity with relatively inelastic demand, and one that taxes a commodity with relatively elastic demand. Which will result in the larger deadweight loss? The answer is, all else being equal, the commodity with the more elastic demand will have a larger deadweight loss. Why? The more elastic demand is, the flatter the demand curve will be. This means that a more elastic demand curve will respond to a tax with a relatively larger decrease in quantity. And, since this quantity distortion is the source of the deadweight loss, the more elastic demand curve will result in the larger deadweight loss.

Does this mean that we should only tax inelastic things, since this will make society better off? Not really. The main reason is that even though taxing inelastic things may be better for society as a whole from an efficiency standpoint, it may have undesirable redistributive effects. For example, we could tax cigarette smokers and force them to pay for road construction and schools. Since cigarette demand is relatively inelastic, this would result in a relatively small deadweight loss. However, is it really fair to force smokers to pay for roads and schools, even though they don't necessarily use the roads and schools any more intensely than other people? Probably not. We recently had a related issue in Massachusetts. The state wanted to increase tolls on the turnpike in order to pay for construction at the airport. Turnpike usage is relatively inelastic, but is it fair to make turnpike users pay for airport construction, even though turnpike users are no more likely to be going to the airport than other drivers? Issues of balancing efficiency and equity concerns such as this arise often in policy decisions.

7.3 The Fundamental Welfare Theorems

Recall that we ended our discussion of production by talking about efficiency, and showed that any profit-maximizing production plan is efficient (i.e. the same output cannot be produced using fewer inputs), and that any efficient production plan (under certain circumstances) is the profit maximizing production plan for some price vector. We now turn to ask the same questions of markets. That is, when are the allocations made by markets "efficient," and is every efficient allocation the market allocation for some initial conditions? Again, the reason we ask these questions has to do with decentralization. When can we decentralize the decisions we make in our society? Do we know that profit-maximizing firms and utility maximizing agents will arrive at a Pareto optimal allocation through the market? If we have a particular Pareto optimal allocation

in mind, can we rely on the market to get us there, provided we start at the right place (i.e., initial endowment for consumers)?

The proper concept of efficiency here is Pareto optimality. Recall that an allocation is Pareto optimal if there is no other feasible allocation that makes all agents at least as well off and some agent better off. We will study Pareto optimal allocations in the context of the quasilinear partial equilibrium model we introduced earlier. This greatly simplifies the analysis, since when preferences are quasilinear, the frontier of the utility possibility set is linear. That is, all points that are Pareto efficient involve the same consumption of the non-numeraire good by the consumers, and differ only in the distribution of the numeraire among the consumers.

To illustrate this point, suppose there are two consumers, and fix the consumption and production levels at \bar{x} and \bar{q} respectively. This will leave $w_m - \sum_j c_j(\bar{q}_j)$ of the numeraire to be distributed among the consumers. Since the numeraire can be traded one-for-one among consumers, the utility possibility set for this x^* and q^* is the set

$$\left\{ (u_1, u_2) \mid u_1 + u_2 \leq \phi_1(\bar{x}_1) + \phi_2(\bar{x}_2) + w_m - \sum_j c_j(\bar{q}_j) \right\}$$

This is the utility possibility set for any particular allocation of the consumption good, (\bar{x}, \bar{q}) , if we allow the remaining numeraire to be distributed among consumers in any possible way. The utility possibility set for the efficient allocation is the set generated by the x_1^* and x_2^* that maximize the right hand side of this expression: That is, Pareto optimal allocations satisfy:

$$\begin{aligned} x_1^*, x_2^*, q_j^* &\in \arg \max \phi_1(x_1) + \phi_2(x_2) + w_m - \sum_j c_j(q_j) \\ \text{subject to} & : x_1 + x_2 = \sum q_j \end{aligned}$$

We will call the x 's and q 's generated by such a procedure the optimal production and consumption levels of good x . If the firms have strictly concave production functions and $\phi_i(\cdot)$ is strictly concave, then there will be a unique (x, q) that maximizes the above expression.

We can rewrite the above problem in the multiple consumer case as:

$$\begin{aligned} \max_{x, q} & \left(\sum_i \phi_i(x_i) \right) + w_m - \left(\sum_j c_j(q_j) \right) \\ \text{subject to} & : \sum_{i=1}^I x_i - \sum_{j=1}^J q_j = 0 \end{aligned}$$

The top line just says to maximize the sum of the consumers' utilities. The constraint is that the total consumption of x is the same as the total production. Letting the Lagrange multiplier be μ , the first-order conditions for this problem are:

$$\begin{aligned}\phi'_i(x_i^*) &\leq \mu \text{ with equality if } x_i^* > 0 \\ c'_j(q_j^*) &\geq \mu \text{ with equality if } q_j^* > 0 \\ \sum_{i=1}^I x_i^* &= \sum_{j=1}^J q_j^*\end{aligned}$$

Note that these are exactly the conditions as the conditions defining the competitive equilibrium except that p^* has been replaced by μ . In other words, we know that the allocation produced by the competitive market satisfies these conditions, and that $\mu = p^*$. Thus the competitive market allocation is Pareto optimal, and the market clearing price p^* is the shadow value of the constraint: the additional social benefit generated by consuming one more unit of output or producing one less unit of output. Hence this is just another expression of the fact that at p^* the marginal social benefit of additional output equals the marginal social cost.

The preceding argument establishes **the first fundamental theorem of welfare economics** in the partial equilibrium case. *If the price p^* and the allocation (x^*, q^*) constitute a competitive equilibrium, then this allocation is Pareto optimal.*

The first theorem is just a formal expression of Adam Smith's invisible hand – the market acts to allocate commodities in a Pareto optimal manner. Since $p^* = \mu$, which is the shadow price of additional units of x , each firm acting in order to maximize its own profits chooses the output that equates the marginal cost of its production to the marginal social benefit, and each consumer, in choosing the quantity to consume in order to maximize utility, is also setting marginal benefit equal to the marginal social cost.

Note that while this is a special case, the first welfare theorem will hold quite generally whenever there are complete markets, no matter how many commodities there are. It will fail, however, when there are commodities (things that affect utility) that have no markets (as in the externalities problem we'll look at soon).

As we did with production, we can also look at this problem “backward.” Can any Pareto optimal allocation be generated as the outcome of a competitive market, for some suitable initial endowment vector? The answer to this question is yes.

To see why, recall that when all $\phi_i()$'s are strictly concave and all $c_j()$'s are strictly convex, there is a unique allocation of the consumption commodity x that maximizes the sum of the consumers'

utilities. The set of Pareto optimal allocations is derived by allocating the consumption commodity in this manner and varying the amount of the numeraire commodity given to each of the consumers. Thus the set of Pareto optimal allocations is a line with normal vector $(1,1,1,\dots,1)$ (see, for example, Figure 10.D.1 in MWG), since one unit of utility can be transferred from one consumer to another by transferring a unit of the numeraire.

Thus any Pareto optimal allocation can be generated by letting the market work and then appropriately transferring the numeraire. But, recall that firms' production decisions and consumers' consumption decisions do not depend on the initial endowment of the numeraire. Because of this, we could also perform these transfers before the market works. This allows us to implement any point along the Pareto frontier.

To see why, let (x^*, q^*) be the Pareto optimal allocation of the consumption commodity, and suppose we want to implement the point where each consumer gets (x_i^*, m_i^*) after the market works, where $\sum_i m_i^* = w_m - \sum_j c_j(q_j^*)$. If we want consumer i to have m_i^* units of the numeraire after the transfer, we need him to have m_i^0 before the transfer, where

$$m_i^0 + \sum_j \theta_j^i (p^* q_j^* - c_j(q_j^*)) = m_i^* + p^* x_i^*.$$

Hence if people have wealth m_i^0 before the market starts to work, allocation (x^*, q^*, m^*) will result.

This yields the second fundamental theorem of welfare economics. Let u_i^* be the utility in a Pareto optimal allocation for some initial endowment vector. There exists a set of transfers T_i (the amount of the numeraire given to consumer i) such that $\sum_i T_i = 0$ and the allocation generated by the competitive market yields the utility vector u^* .

The transfers are given by $T_i = m_i^0 - w_{mi}$.

7.3.1 Welfare Analysis and Partial Equilibrium

Recall in our discussion of consumer theory we said that equivalent variation is the proper measure of the impact of a policy change on consumers, and that EV is given by the area to the left of the Hicksian demand curve between the initial and final prices. However, since there are no wealth effects for the consumption good here, we know that the Hicksian and Walrasian demand curves are the same. So, the area to the left of the Walrasian demand curve is a proper measure of consumer welfare. Further, since utility is quasilinear, it makes sense to look at aggregate demand, and there is a normative representative consumer whose preferences are captured by the aggregate demand curve. Hence the area to the left of the Walrasian demand curve is a good measure of changes in

social welfare.

It is worthwhile to derive a measure of aggregate social surplus here, even though we already did it during our study of aggregation. Recall the Pareto optimality problem:⁵

$$\begin{aligned} & \max_{x,q} \left(\sum_i \phi_i(x_i) \right) + w_m - \left(\sum_j c_j(q_j) \right) \\ \text{subject to} \quad & : \quad \sum_{i=1}^I x_i - \sum_{j=1}^J q_j = 0. \end{aligned}$$

We said that the solution to this problem determines the allocation that maximizes consumers' welfare (and therefore societal welfare).

Consider the objective function:

$$\left(\sum_i \phi_i(x_i) \right) - \left(\sum_j c_j(q_j) \right) + w_m.$$

The last term is the initial aggregate endowment of the numeraire good, which is just a constant in the objective function. The first two terms represent the difference between aggregate utility from consumption and aggregate cost of production. It is this difference that we are maximizing in the Pareto optimality problem.

Consider a single unit of production. The difference between the utility derived from that production and the cost of production is the societal benefit from that good (since all profits return to consumers). If we add this surplus across all consumers, that gets us

$$\left(\sum_i \phi_i(x_i) \right) - \left(\sum_j c_j(q_j) \right)$$

which is the surplus generated by all units of production. This term, called Marshallian Aggregate Surplus (MAS), is the measure of social benefit that we want to use, since it tells us how much better off society is made when $\sum_i x_i = \sum_j q_j$ units of the non-numeraire good are produced and sold.

To better understand, note that we can break the surplus down into four parts:

1. (a) Some of the surplus comes from consumption, $\sum_i \phi_i(x_i)$,
 - (b) Some surplus is lost due to paying price p for the good.
- $\sum_i \phi_i(x_i) - px_i$ is the aggregate consumer surplus

⁵This corresponds to the utilitarian social welfare problem.

2. (a) Some surplus is gained by firms in the form of revenue, $\sum_j pq_j$
- (b) Some surplus is lost by firms in the form of production cost, $\sum_j c_j(q_j)$
- $\sum_j pq_j - c_j(q_j)$ is the aggregate producer surplus, which is then redistributed to consumers in the form of dividends, $\theta_{ij}(pq_j - c_j(q_j))$.

So, consumers receive part of the benefit through consumption of the non-numeraire good,

$$\sum_i \phi_i(x_i) - px_i$$

and part of the benefit through consumption of the dividends, which are measured in units of the numeraire:

$$\sum_j pq_j - c_j(q_j).$$

Aggregate surplus is found by adding these two together, and noting that $\sum_i px_i = \sum_j pq_j$

In our quasilinear model, the set of utility vectors that can be achieved by a feasible allocation is given by:

$$\left\{ (u_1, \dots, u_I) \mid \sum_i u_i \leq w_m + \left(\sum_i \phi_i(x_i) \right) - \left(\sum_j c_j(q_j) \right) \right\}$$

Suppose that the government (or you, or anybody) has a view of society that says the total welfare in society is given by:

$$W(u_1, \dots, u_I)$$

Thus this function gives a level of welfare associated with any utility vector – it allows us to compare any two distributions of utility in terms of their overall social welfare. The problem of the social planner would be to choose u in order to maximize $W(u)$ subject to the constraint that u lies in the utility possibility set. Clearly, then, the optimized level of W will be higher when the utility possibility set is larger, which occurs when the MAS is larger. Thus the total societal welfare achievable is increasing in the MAS. See MWG Figure 10.E.1.

In other words, if you want to maximize social welfare, you should first choose the production and consumption vectors that maximize MAS, and then redistribute the numeraire in order to maximize the welfare function. This gives us another separation result: If utility can be perfectly transferred between consumers (as in the quasilinear model), then social welfare is maximized by first choosing production and consumption plans that maximize MAS, and then choosing transfers such that $u_i = \phi_i(x_i) + m_i + t_i$ maximizes $W(u)$.

Now, how does MAS change when the quantity produced and consumed changes? Let $S(x, q)$ be the MAS, formally defined as follows:

$$S(x, q) = \left(\sum_i \phi_i(x_i) \right) - \left(\sum_j c_j(q_j) \right).$$

Consider a differential increase in consumption and production: $(dx_1, dx_2, \dots, dx_I, dq_1, \dots, dq_J)$ satisfying $\sum_i dx_i = \sum_j dq_j$. Note that under such a change, we increase total production and total consumption by the same amount.

The differential in S is given by

$$dS = \left(\sum_i \phi'_i(x_i) dx_i \right) - \left(\sum_j c'_j(q_j) dq_j \right).$$

Since consumers maximize utility, $\phi'_i(x_i) = p(x)$ for all i , and since producers maximize profit, $c'_j(q_j) = c'(q)$ for all j . Thus:

$$dS = \left(p(x) \sum_i dx_i \right) - \left(c'(q) \sum_j dq_j \right)$$

which by definition of our changes (and market clearing) implies

$$dS = (p(x) - c'(q)) dx$$

And, integrating this from 0 to \bar{x} yields that MAS equals

$$S(x) = \int_0^x (p(s) - c'(s)) ds.$$

Thus the total surplus is the area between the supply and demand curves between 0 and the quantity sold, \bar{x} .

Example: Welfare Effects of a Tax

We return to the idea of a commodity tax that we first considered in the context of consumer theory. Suppose there is a government that attempts to maximize the welfare of its citizens. The government keeps a balanced budget, and tax revenues are returned to consumers in the form of a lump sum transfer.

What are the welfare effects of this tax? Define $x_1^*(t), \dots, x_I^*(t)$ and $q_1^*(t), \dots, q_J^*(t)$ and $p^*(t)$, respectively, to be the consumptions, productions, and price paid by consumers when the per-unit

tax is t . Define $x^*(t) = \sum x_i^*(t)$ and $q^*(t) = \sum q_j^*(t)$ to be aggregate consumption and production, respectively. Letting $S^*(t) = \sum x_i^*(t) - \sum q_j^*(t)$ be the level of MAS (also equal to the area below the demand curve and above the supply curve) at tax rate t , the change in MAS when a tax of t is imposed is given by:

$$\begin{aligned} S^*(x^*(t)) - S^*(x^*(0)) &= x^*(t) - q^*(t) - (x^*(0) - q^*(0)) \\ &= \int_{x^*(0)}^{x^*(t)} p(s) - c'(s) ds, \end{aligned}$$

by the definition of MAS developed earlier. Thus the change in MAS is given by the change in the area between the aggregate demand and supply curves, between the equilibrium quantity when there is no tax and the equilibrium quantity after the tax is imposed. Note that this is just the area we called deadweight loss earlier.

7.4 Entry and Long-Run Competitive Equilibrium

Up until now, we have considered the competitive equilibrium holding the supply side of the market fixed. In particular, we have assumed:

1. Firms are unable to vary their fixed factors of production (plant size is fixed)
2. Firms are unable to enter or exit the market – the number of firms stay fixed

These assumptions are appropriate in the short run. However, if we want to examine the behavior of the market in the long run, we must explicitly allow for firms to change their fixed factors, including entering or exiting the industry.

In the long run, a perfectly competitive market is characterized by:

1. Firms and consumers are price takers
2. Free entry and exit

The free entry and exit condition doesn't mean that firms can enter at no cost. Rather, it means that there is no impediment to them incurring the cost and entering the market. For example, there are no laws against entry, there are no proprietary technologies or scarce resources, etc. Thus firms have the freedom to enter or exit, but this is not to say that they can do it for free. In our discussion up until now, we have considered only the price-taking requirement. In order to think about competitive equilibrium in the long run, we add the free entry condition.




Figure 7.3: Short-Run Equilibrium

7.4.1 Long-Run Competitive Equilibrium

Consider our analysis of perfect equilibrium in the short run.⁶ The short-run equilibrium price and quantity are found where supply equals demand for a given number of firms. However, note that if there happen to be a small number of firms in an industry, it may be that the firms are making large profits.

Figure 7.3 links the market equilibrium with the individual firm. On the left, aggregate supply and demand combine to determine the equilibrium price, P^* . At this price, the individual firm's behavior is depicted in the right panel. The firm chooses to produce q_i^* units of output, and earns $q_i^*(P^* - AC(q_i^*))$ profit. Total profit by the firm is shaded in the right-hand panel.

Now, if you are a business person, you see an industry where firms are making large profits, and you can enter it if you want, what do you do? You enter. When you enter, what happens to the supply curve? It shifts out to the right. And, as the supply curve shifts right, the equilibrium price decreases, which decreases the profit of the firms already in the industry.

How many firms enter? As long as a firm can earn a positive profit by entering the industry, it will choose to enter. Thus firms will continue to enter until they drive profit to zero. This situation is shown in Figure 7.4.

Note that profit equals zero when the equilibrium price, denoted by P' in the diagram, is such

⁶Implicit in the arguments for this section is the idea that all firms, both those that are in the market and those who could potentially enter the market, have the same technology. The conclusions can be adapted to the case where technologies are heterogeneous without changing the results too much.




Figure 7.4: Long-Run Equilibrium

that $P' = \min AC$.

Similarly, if there are initially too many firms in the industry, the firms in the industry will earn negative profits. This will drive firms to exit the industry. As they exit the industry, the price will increase, and the size of the firms' losses will decrease. When does exit stop? When the firms that are in the industry earn zero profit. You should draw the diagram for this case.

The dynamic story I've been telling suggests the following requirements for a long-run competitive equilibrium. First, we keep the requirements for a short-run equilibrium:

1. (a) Firms maximize profits
- (b) Consumers maximize utility
- (c) Market clearing: price adjusts so that supply equals demand

Second, we add the additional requirement that the equilibrium number of firms is found where in the short-run equilibrium firms make zero profit.⁷

Formally, if there are a very large number of identical firms that could potentially enter this market, this means that the equilibrium consists of q^*, p^*, J^* that satisfy:

1. q^* maximizes $p^*q - c(q)$ for each firm.
2. x_i^* maximizes $u_i(x_i) - px_i$ for all i

⁷There is something of an integer problem here. That is, it may be if there are J^* firms all firms earn a positive profit, but if there are $J^* + 1$ firms, all firms earn a negative profit. In this case, we will say that the equilibrium is the largest number of firms such that firms do not earn negative profits.

3. $x(p^*) = J^*q^*$: market clearing
4. $p^*q^* - c(q^*) = 0$: free entry

The last requirement is one way to think of the free entry condition. Entry continues until all firms make zero profit. Another way to think of it is as follows: Entry will continue until the point that with J^* firms in the industry, all firms make non-negative profits. With $J^* + 1$ firms in the industry, all firms make negative profits. Thus it need not be the case that all firms make zero profits. But, it must be the case that if one more firm enters, all firms will make negative profits. Thus J^* is the maximum number of firms that can be supported by this market.

What will the long-run aggregate supply correspondence look like? Since all firms are the same, the long-run aggregate supply correspondence as a function of p will look as follows:

$$\begin{aligned} Q(p) &= \infty \text{ if } \pi(p) > 0 \\ &= Jq \text{ for some integer } J \geq 0 \text{ and } q \in q(p) \text{ if } \pi(p) = 0 \\ &= 0 \text{ if } \pi(p) < 0 \end{aligned}$$

That is, if the price is such that profits are positive, an infinite number of firms will enter the industry – driving the quantity supplied to infinity. If price is such that profits are zero then some integer number of firms will enter the market and produce q according to its supply function $q(p)$. When profits are negative at a specific price, firms will supply nothing. Generally, however, we won't worry about the integer problems here. We'll assume that in the long run, the supply curve is horizontal at the level of minimum average cost.

Example: Consider the case of constant returns to scale technology with no fixed cost:

$$c(q) = cq.$$

- In this case, firms will supply an infinite amount when $p > c$, any positive amount when $p = c$, and zero when $p < c$.
- The long-run equilibrium will be where demand and long-run supply cross, which is at $p = c$. However, we don't know how many firms there will be, since the firms could split up the quantity they want to produce in any way they want. It is a long-run equilibrium with any number of firms!

Example: Firms have strictly increasing, strictly convex cost

- In this case, firms make a positive profit whenever $p > c'(0)$. Hence entry will continue, driving price down until it reaches $c'(0)$, and all firms make zero profit.
- This makes sense – convex cost is the same as decreasing returns to scale. In this case, the most efficient firms are those that are producing no output at all. So, as price decreases, firms are driven to produce more efficiently, and that involves producing lower quantities.

Example: Firms Have U-Shaped Cost Curves

Now suppose that the firm has a cost function that is first concave, and then convex, giving a U-shaped average cost curve. In the long run, price will be driven down to the point where $p = \min AC$. In that case, each firm will produce the quantity that minimizes AC . Thus the long-run aggregate supply function is given by:

$$\begin{aligned} Q(p) &= \infty \text{ if } p > \min AC \\ &= J\bar{q} \text{ if } p = \min AC, \text{ where } \bar{q} = q(\min AC) \\ &= 0 \text{ if } p < \min AC \end{aligned}$$

Note that it is possible that when $p = \min AC$, there is no J such that $J \cdot q(\min AC) = x(\min AC)$. This is known as the integer problem. There are several reasons to think that the integer problem is not such a horrible thing:

- If firms are small relative to the market, then there will be a J such that $J \cdot q(\min AC)$ is close to $x(\min AC)$
- Long-run equilibrium is a theoretical construct – we never really get there, but we're always moving toward there.
- If the price is slightly above $\min AC$, then the supply curve will be upward sloping. Thus we are really looking for the largest J such that $p > \min AC$. These firms will make a slight profit.

7.4.2 Final Comments on Partial Equilibrium

The approach to partial equilibrium we have adopted has been based on a quasilinear model. How crucial is this for the results? Well, the quasilinear utility is not critical for the determination of

the competitive equilibrium – you would still find it in the same way. However, it is critical for the welfare results. Without quasilinear utility, the area under the Walrasian demand curves doesn't mean anything – so we will need a different welfare measure. Further, with wealth effects, welfare will depend on the distribution of the numeraire, not just the consumption good, which means that we will have to do additional work.

Chapter 8

Externalities and Public Goods

8.1 What is an Externality?

We just showed that competitive markets result in Pareto optimal allocations – that is the market acts to make sure that those who value goods the most receive them, and those that can produce goods at the least cost produce them, and there is no way that everybody in society could be made better off. This gave us the first and second welfare theorems – the market allocates commodities efficiently, and any efficient allocation can be derived by a market with suitable ex ante transfers of wealth. Now we will take a look at one important circumstance where the welfare theorems do not hold.

When we talked about commodities in the past, they were always what are called “private goods.” That is, they were such that they were consumed by only one person, and that person’s consumption of the good had no effect on other people’s utility. But, this is not true of all goods. Think, for example, of a local bakery that produces bread. Earlier, we said that each person purchases the quantity of bread where the marginal benefit of consuming an additional loaf is just equal to the price of a loaf, and each firm produces bread up to the point where the marginal cost of producing the loaf is just equal to its price. In equilibrium, then, the marginal benefit of eating an additional loaf of bread is just equal to the marginal cost of producing an additional loaf. But, think about the following. People who walk by the bakery get the benefit from the pleasant smell of baking bread, and this is not incorporated into the price of bread. Thus at the equilibrium, the marginal social benefit of another loaf of bread is equal to the benefit people get from eating the bread as well as the benefit people get from the pleasant smell of baking bread. However, since bread purchasers do not take into account the benefit provided to people who do not purchase

bread, at the equilibrium price the total marginal benefit of additional bread will be greater than the marginal cost. From a social perspective, too little bread is produced.

We can also consider the case of a negative externality. One of the standard examples in this situation is the case of pollution. Suppose that a factory produces and sells tires. In the course of the production, smoke is produced, and everybody that lives in the neighborhood of the factory suffers because of it. The price consumers are willing to pay for tires is given by the benefit derived from using the tires. Hence at the market equilibrium, the marginal cost of producing a tire is equal to the marginal benefit of using the tire, but the market does not incorporate the additional cost of pollution imposed on those who live near the factory. Thus from the social point of view, too many tires will be produced by the market.

Another way to think about (some types of) goods with external costs or benefits is as public goods. A public good is a good that can be consumed by more than one consumer. Public goods can be classified based on whether people can be excluded from using them, and whether their consumption is rivalrous or not. For example, a non-excludable, non-rivalrous public good is national defense.¹ Having an army provides benefits to all residents of a country. It is non-excludable, since you cannot exclude a person from being protected by the army, and it is non-rivalrous, since one person consuming national defense does not diminish the effectiveness of national defense for other people.² Pollution is a non-rivalrous public good (or public bad), since consumption of polluted air by one person does not diminish the “ability” of other people to consume it. A bridge is also a non-rivalrous public good (up to certain capacity concerns), but it may be excludable if you only allow certain people to use it. Another example is premium cable television. One person having HBO does not diminish the ability of others to have it, but people can be excluded from having it by scrambling their signal.

Examples of externalities and public goods tend to overlap. It is hard to say what is an externality and what is a public good. This is as you would expect, since the two categories are really just different ways of talking about goods with non-private aspects. It turns out that a useful way to think about different examples is in terms of whether they are rivalrous or non-rivalrous, and whether they are excludable or not. Based on this, we can create a 2-by-2 matrix describing

¹Goods of this type are often called “pure public goods.”

²This is true in the case of national missile defense, which protects all people equally. However, in a nation where the military must either protect the northern region or the southern region, the army may be a rivalrous public good.

goods.³

	Non-rivalrous	Rivalrous
Non-Excludable	(Pure) Public Goods	Common-Pool Resources
Excludable	Club Goods	Private Goods

Private goods are goods where consumption by one person prevents consumption by another (an extreme form of rivalrous consumption), and one person has the right to prevent the other from consuming the object. When consumption is non-rivalrous but excludable, as in the case of a bridge, such goods are sometimes called **club goods**. Because club goods are excludable, inefficiencies due to external effects can often be addressed by charging people for access to the club goods, such as charging a toll for a bridge or a membership fee for a club. Pure public goods are goods such as national defense, where consumption is non-rivalrous and non-excludable. **Common-pool resources** are goods such as national fisheries or forests, where consumption is rivalrous but it is difficult to exclude people from consuming them. Both pure public goods and common-pool resources are situations where the market will fail to allocate resources efficiently. After considering a simple, bilateral externality, we will go on to study pure public goods and common pool resources in greater detail.

8.2 Bilateral Externalities

We begin with the following definition. An **externality** is present whenever the well-being of a consumer or the production possibilities of a firm are directly affected by the actions of another agent in the economy (and this interaction is not mediated by the price mechanism). An important feature of this definition is the word “directly.” This is because we want to differentiate between a true externality, and what is called a **pecuniary externality**. For example, return to the example of the bakery we considered earlier. We can think of three kinds of external effects. First, there is the fact, as we discussed earlier, that consumers walking down the street may get utility from the smell of baking bread. This is true regardless of whether the people participate in any market. Second, if the smells of the bread are pleasant enough, the bakery may be able to charge more for the bread it sells, and, the fact that the price of bread increases may have harmful effects on people who buy the bread because they must pay more for the bread. We call this type of effect a pecuniary externality, since it works through the price mechanism. Effects such as this are not

³Based on Ostrom, *Rules, Games, and Common Pool Resources*, University of Michigan Press, 1994.

really externalities, and will not have the distortionary effects we will find with true externalities.⁴ Third, there is the fact that being next to a bakery may increase rents in the area around it. While this is a situation where the bakery has effects outside of the bread market, this effect is captured by the rent paid by other stores in the area. Whether this is an externality or not depends on the particular situation. For example, if you own an apartment building next to the bakery before it opens and are able to increase rents after it begins to produce bread, then you have realized an external benefit from the bakery (since the bakery has increased the value of your property). On the other hand, if you purchase the building next to the bakery once it is already opened, then you will pay a higher price for the building, but this is the fair price for a building next to a bakery. Thus this situation is really more of a pecuniary externality than a true externality.

We will use the following example for our externality model. There are two consumers, $i = 1, 2$. There are L traded goods in the economy with price vector p , and the actions taken by these two consumers do not affect the prices of these goods. That is, the consumers are price takers. Further, consumer i has initial wealth w_i .

Each consumer has preferences over both the commodities he consumes and over some action h that is taken by consumer 1. That is,

$$u_i(x_1^i, \dots, x_L^i, h).$$

Activity h is something that has no direct monetary cost for person 1. For example, it could be playing loud music. Loud music itself has no cost. In order to play it, the consumer must purchase electricity, but electricity can be captured as one of the components of x^i .

From the point of view of consumer 2, h represents an external effect of consumer 1's action. In the model, we assume that

$$\frac{\partial u_2}{\partial h} \neq 0.$$

Thus the externality in this model lies in the fact that h affects consumer 2's utility, but it is not priced by the market. For example, h is the quantity of loud music played by person 1.

Let $v_i(p, w_i, h)$ be consumer i 's indirect utility function:

$$\begin{aligned} v_i(w_i, h) &= \max_{x_i} u_i(x_i, h) \\ \text{s.t. } p \cdot x^i &\leq w_i. \end{aligned}$$

⁴The key to being a true externality is that the external effect will usually be on parties that are not participants in the market we are studying, in this case the market for bread.

We will also make the additional assumption that preferences are quasilinear with respect to some numeraire commodity. If this were not so, then the optimal level of the externality would depend on the consumer's level of wealth, significantly complicating the analysis.

When preferences are quasilinear, the consumer's indirect utility function takes the form:

$$v_i(w_i, h) = \bar{\phi}_i(h) + w_i.^5$$

Since we are going to be concerned with the behavior of utility with respect to h but not p , we will suppress the price argument in the utility function. That is, let $\phi_i(h) = \bar{\phi}_i(p, h)$, when we hold the price p constant. We will assume that utility is concave in h : $\phi_i''(h) < 0$.

Now, we want to derive the competitive equilibrium outcome, and show that it is not Pareto optimal. How will consumer 1 choose h ? The function v_1 gives the highest utility the consumer can achieve for any level of h . Thus in order to maximize utility, the consumer should choose h in order to maximize v_1 . Thus the consumer will choose h in order to satisfy the following necessary and sufficient condition (assuming an interior solution):

$$\phi_1'(h^*) = 0.$$

Even though consumer 2's utility depends on h , it cannot affect the choice of h . Herein lies the problem.

What is the socially optimal level of h ? The socially optimal level of h will maximize the sum of the consumers' utilities (we can add utilities because of the quasilinear form) :

$$\max_h \phi_1(h) + \phi_2(h).$$

The first-order condition for an interior maximum is:

$$\phi_1'(h^{**}) + \phi_2'(h^{**}) = 0,$$

where h^{**} is the Pareto optimal amount of h .

The social optimum requires that the sum of the two consumers' marginal utilities for h is zero (for an interior solution). On the other hand, the level of the externality that is actually chosen depends only on person 1's utility. Thus the level of the externality will not generally be the socially optimal one. In the case where the externality is bad for consumer 2 (loud music), the level of $h^* > h^{**}$. That is, too much h is produced. In the case where the externality is good for consumer 2 (baking bread smell or yard beautification), too little will be provided, $h^* < h^{**}$. These situations are illustrated in Figures 8.1 and 8.2.




Figure 8.1: Negative Externality: $h^{**} < h^*$




Figure 8.2: Positive Externality: $h^{**} > h^*$

Note that the social optimum is not for the externality to be eliminated entirely. Rather, the social optimum is where the sum of the marginal benefit of the two consumers equals zero. In the case where there is a negative externality, this is where the marginal benefit to person 1 equals the marginal cost to person 2. In the case of a positive externality, this is where the sum of the marginal benefit to the two people is equal to zero.

The fact that the optimal level of a negative externality is greater than zero is true even in the case where the externality is pollution, endangered species preservation, etc. Of course, this still leaves open for discussion the question of how to value the harm of pollution or the benefit of saving wildlife. Generally, those who produce the externality (i.e., polluters) think that the optimal level of the externality is larger than those who are victims of it.

8.2.1 Traditional Solutions to the Externality Problem

There are two traditional approaches to solving the externality problem: quotas and taxes. Quotas impose a maximum (or minimum) amount of the externality good that can be produced. Taxes impose a cost of producing the externality good on the producer. Positive taxes will tend to decrease production of the externality, while negative taxes (subsidies) will tend to increase production of the externality.


Let's begin by considering a quota. Suppose that activity h generates a negative external effect, so that the privately chosen quantity h^* is greater than the socially optimal quantity h^{**} . In this case, the government can simply pass a quota, prohibiting production in excess of h^{**} . In the case of a positive externality, the government can require consumer 1 to produce at least h^{**} units of the externality (although this is less often seen in practice).

While the quota solution is simple to state, it is less simple to implement since it requires the government to enforce the quota. This involves monitoring the producer, which can be difficult and costly. One thing that would be nice would be if there were some adjustment we could make to the market so that it worked properly. One way to do this, known as Pigouvian Taxation, is to impose a tax on the production of the externality good, h .

Suppose consumer 1 were charged a tax of t_h per unit of h produced. His optimization problem would then be

$$\max \phi_1(h) - t_h h$$

⁵The argument for why is contained in footnote 3 on p. 353 in MWG.

Figure 8.3: Implementing h^{**} Using a Tax on h

with first-order condition

$$\phi'_1(h^t) = t_h$$

Thus setting $t_h = -\phi'_2(h^{**})$ (which is positive) will lead consumer 1 to choose $h^t = h^{**}$, implementing the social optimum. See Figure 8.3.

Note that the proper tax is equal to the marginal externality at the optimal level of h . By forcing consumer 1 to pay this, he is required to internalize the externality. That is, he must pay the marginal cost imposed on consumer 2 when the externality is set at its optimal level, h^{**} . When the tax rate is set in this way, consumer 1 chooses the Pareto optimal level of the externality.

In the case of a positive externality, the tax needed to implement the Pareto optimal level of the externality is negative. Consumer 1 is subsidized in the amount of the marginal external effect at the optimal level of the externality activity. And, when he internalizes the benefit imposed on the other consumer, he chooses the (larger) optimal level of h .

Another equivalent approach would be for the government to pay consumer 1 to reduce production of the externality. In this case, the consumer's objective function is:

$$\phi_1(h) + s_h(h^* - h) = \phi_1(h) - s_h h + s_h h^*$$

By setting $s_h = -\phi'_2(h^{**})$, the socially optimal level of h is implemented.

Note that it is key to tax the externality producing activity directly. If you want to reduce pollution from cars, you have to tax pollution, not cars. Taxes on cars will not restore optimality of pollution (since it does not affect the marginal propensity to pollute) and will distort people's car purchasing decisions. Similarly, if you want a tractor factory to reduce its pollution, you need

to tax pollution, not tractors. Taxing tractors will generally lead the firm to reduce output, but it won't necessarily lead it to reduce pollution (what if the increased costs lead it to adopt a more polluting technology?).

Note that taxes and quotas will restore optimality, but this result depends on the government knowing exactly what the correct level of the externality-producing activity is. In addition, it will require detailed knowledge of the preferences of the consumers.

8.2.2 Bargaining and Enforceable Property Rights: Coase's Theorem

A different approach to the externality problem relies on the parties to negotiate a solution to the problem themselves. As we shall see, the success of such a system depends on making sure that property rights are clearly assigned. Does consumer 1 have the right to produce h ? If so, how much? Can consumer 2 prevent consumer 1 from producing h ? If so, how much? The surprising result (known as Coase's Theorem) is that as long as property rights are clearly assigned, the two parties will negotiate in such a way that the optimal level of the externality-producing activity is implemented.

Suppose, for example, that we give consumer 2 the right to an externality-free environment. That is, consumer 2 has the right to prohibit consumer 1 from undertaking activity h . But, this right is contractible. Consumer 2 can sell consumer 1 the right to undertake h_2 units of activity h in exchange for some transfer, T_2 . The two consumers will bargain both over the size of the transfer T_2 and over the number of units of the externality good produced, h_2 .⁶

In order to determine the outcome of the bargaining, we first need to specify the bargaining mechanism. That is, who does what when, what are the other consumer's possible responses, and what happens following each response.⁷

Suppose bargaining mechanism is as follows:

1. Consumer 2 offers consumer 1 a take-it-or-leave-it contract specifying a payment T_2 and an activity level h_2 .
2. If consumer 1 accepts the offer, that outcome is implemented. If consumer 1 does not accept the offer, consumer 1 cannot produce any of the externality good, i.e., $h = 0$.

⁶The subscript 2 is used here because we will compare this with the case where 1 has the right to produce as much of the externality as it wants, and we'll denote the outcome with the subscript 1 in that case.

⁷Those of you familiar with game theory will recognize that what we are really doing here is setting up a game. If you don't know any game theory, revisit this after you see some, and it will be much clearer.

To analyze this, begin by considering which offers (h, T) will be accepted by consumer 1. Since in the absence of agreement, consumer 1 must produce $h = 0$, consumer 1 will accept (h_2, T_2) if and only if it offers higher utility than $h = 0$. That is, 1 accepts if and only if:⁸

$$\phi_1(h) - T \geq \phi_1(0).$$

Given this constraint on the set of acceptable offers, consumer 2 will choose (h_2, T_2) in order to solve the following problem.

$$\begin{aligned} & \max_{h, T} \phi_2(h) + T \\ \text{subject to} & : \phi_1(h) - T \geq \phi_1(0). \end{aligned}$$

Since consumer 2 prefers higher T , the constraint will bind at the optimum. Thus the problem becomes:

$$\max_h \phi_1(h) + \phi_2(h) - \phi_1(0).$$

The first-order condition for this problem is given by:

$$\phi_1'(h_2) + \phi_2'(h_2) = 0.$$

But, this is the same condition that defines the socially optimal level of h . Thus consumer 2 chooses $h_2 = h^{**}$, and, using the constraint, $T_2 = \phi_1(h^{**}) - \phi_1(0)$. And, the offer (h_2, T_2) is accepted by consumer 1. Thus this bargaining process implements the social optimum.

Now, we can ask the same question in the case where consumer 1 has the right to produce as much of the externality as she wants. We maintain the same bargaining mechanism. Consumer 2 makes consumer 1 a take-it-or-leave-it offer (h_1, T_1) , where the subscript indicates that consumer 1 has the property right in this situation. However, now, in the event that 1 rejects the offer, consumer 1 can choose to produce as much of the externality as she wants, which means that she will choose to produce h^* . Thus the only change between this situation and the previous example is what happens in the event that no agreement is reached.

In this case, consumer 2's problem is:

$$\begin{aligned} & \max_{h, T} \phi_2(h) + T \\ \text{subject to} & : \phi_1(h) - T \geq \phi_1(h^*) \end{aligned}$$

⁸In the language of game theory, this is called an incentive compatibility constraint.

Again, we know that the constraint will bind, and so consumer 2 chooses h_1 and T_1 in order to maximize

$$\max \phi_1(h) + \phi_2(h) - \phi_1(h^*)$$

which is also maximized at $h_1 = h^{**}$, since the first-order condition is the same. The only difference is in the transfer. Here $T_1 = \phi_1(h^{**}) - \phi_1(h^*)$.

While both property-rights allocations implement h^* , they have different distributional consequences. The transfer is larger in the case where consumer 2 has the property rights than when consumer 1 has the property rights. The reason for this is that consumer 2 is in a better bargaining position when the non-bargaining outcome is that consumer 1 is forced to produce 0 units of the externality good. However, note that in the quasilinear framework, redistribution of the numeraire commodity has no effect on social welfare.

The fact that regardless of how the property rights are allocated, bargaining leads to a Pareto optimal allocation is an example of the Coase Theorem: If trade of the externality can occur, then bargaining will lead to an efficient outcome no matter how property rights are allocated (as long as they are clearly allocated). Note that well-defined, enforceable property rights are essential for bargaining to work. If there is a dispute over who has the right to pollute (or not pollute), then bargaining may not lead to efficiency. An additional requirement for efficiency is that the bargaining process itself is costless.

Note that the government doesn't need to know about individual consumers here – it only needs to define property rights. However, it is critical that it do so clearly. Thus the Coase Theorem provides an argument in favor of having clear laws and well-developed courts.

8.2.3 Externalities and Missing Markets

The externality problem is frequently called a “missing market” problem. To see why, suppose now that there were a market for activity h . That is, suppose consumer 2 had the right to prevent all activity h , but could sell the right to undertake 1 unit of h for a price of p_h . In this case, in deciding how many rights to sell, player 2 will maximize

$$\phi_2(h) + p_h h$$

This has the first-order condition for an interior solution

$$\phi_2'(h) = -p_h,$$

which implicitly defines a supply function: $h_2(p_h)$.

In deciding how many rights to purchase, consumer 1 maximizes

$$\phi_1(h) - p_h h$$

This has the first-order condition for an interior solution:

$$\phi_1'(h) = p_h,$$

which implicitly defines a demand function, $h_1(p_h)$.

The market-clearing condition says that $h_1(p_h) = h_2(p_h)$, or that:

$$\phi_1'(h^m) = -\phi_2'(h^m)$$

at the equilibrium, h^m . But, note that this is the defining equation for the optimal level of the externality, h^{**} . Thus if we can create the missing market, that market will implement the Pareto optimal level of the externality.

This result depends on the assumption of price taking, which is unreasonable in this case. But, in most real markets with externalities, this is not an unreasonable assumption, since (as in the case of air pollution), there are many producers and consumers. This is the basic approach that is used in the case of tradeable pollution permits. The government creates a market for the right to pollute, and, once the missing market has been created, the market will work in such a way that it implements the socially optimal level of the externality good.

8.3 Public Goods and Pure Public Goods

Previously we looked at a simple model of an externality where there were only two consumers. We can also think of externalities in situations where there are many consumers. In situations such as these, it is useful to think of the externality-producing activity as a public good. Public goods are goods that are consumed by more than one consumer. As we described earlier, public goods can take a number of forms. Basically, the most useful way to classify them (I have found) is based on whether the consumption of the good is rivalrous (i.e., whether consumption by one person affects consumption by another person) and whether consumption is excludable (i.e., whether a person can be prevented from consuming the public good). We begin our study with pure public goods. A **pure public good** is a non-rivalrous, non-excludable public good. Consumption of the good

by one person does not affect its consumption by others, and it is difficult (impossible) to exclude a person from consuming it. The prototypical example is national defense.

Many goods are public goods, but are not pure public goods because their consumption is either rivalrous or excludable. Consider, for example, public grazing land or an open-access fishery. The more people who use this resource, the less benefit people get from using it. Resources like this are **common-pool resources**. We'll look at an example of a common-pool resource in Section 8.4.

A public good can also differ from a pure public good if its consumption is excludable. For example, you can exclude people from using a bridge or a park. Excludable, non-rivalrous public goods are called **club goods**. Excludability will play an important role in whether you can get people to pay for a public good or not. For example, how do you expect to get people to voluntarily pay for a pure public good like national defense when they cannot be excluded from consuming it? If there is no threat of being excluded, people will be tempted to free ride off of the contributions of others. On the other hand, in the case of a club good such as a park, the fact that people who do not contribute will be excluded from consuming the public good can be used to induce everybody, not just those who value the club good the most, to contribute.

Finally, not all public goods need to be "good." You can also have a public bad: pollution, poor quality roads, overgrazing on public land, etc. However, it will frequently be possible to redefine a public bad, such as pollution, in terms of a public good, pollution abatement or clean air. Thus the models we use will work equally well for public goods and public bads.

8.3.1 Pure Public Goods

Consider the following simple model of a pure public good. As usual, there are I consumers, and L commodities. Preferences are quasilinear with respect to some numeraire commodity, w . Let x denote the quantity of the public good. In this case, indirect utility takes the form

$$v_i(p, x, w) = \bar{\phi}_i(p, x) + w.$$

As in the case of the bilateral externality, we will not be interested in prices, and so we will let $\phi_i(x) = \bar{\phi}_i(p, x)$. Assume that ϕ_i is twice differentiable and concave at all $x \geq 0$. In the case of a public good, $\phi'_i > 0$, in the case of a public bad, $\phi'_i < 0$.

Assume that the cost of supplying q units of the public good is $c(q)$, where $c(q)$ is strictly increasing, convex, and twice differentiable. In the case of a public good whose production is

costly, $\phi'_i > 0$ and $c' > 0$. In the case of a public bad whose prevention is costly (such as garbage on the front lawn or pollution), $\phi'_i < 0$ and $c' < 0$.

In this model, a Pareto optimal allocation must maximize the aggregate surplus and therefore must solve

$$\max_{q \geq 0} \sum_{i=1}^I \phi_i(q) - c(q).$$

This yields the necessary and sufficient first-order condition, where q^0 is the Pareto optimal quantity,

$$\sum_{i=1}^I \phi'_i(q^0) - c'(q^0) = 0$$

for an interior solution.⁹ Thus the total marginal utility due to increasing the public good is equal to the marginal cost of increasing it.

Private provision of a public good

Now suppose that the public good is provided by private purchases by consumers. That is, the public good is something like national defense, and we ask people to pay for it by saying to them, “Give us some money, and we’ll use it to purchase national defense.”¹⁰ So, each consumer chooses how much of the public good x_i to purchase. We treat the supply side as consisting of profit-maximizing firms with aggregate cost function $c()$.

At a competitive equilibrium:

1. Consumers maximize utility:¹¹

$$\max \phi_i \left(x_i + \sum_{j \neq i} x_j^* \right) - px_i.$$

For an interior solution, the first-order condition is:

$$\phi'_i(x_i^* + x_{-i}^*) = p.$$

⁹As usual, we assume an interior solution here, but generally you would want to look at the Kuhn-Tucker conditions and determine endogenously whether the solution is interior or not.

¹⁰To see how well this works, think about public television.

¹¹In our treatment of the consumer’s problem, we model the consumer as assuming all other consumers choose x_j^* . Consumer i then chooses the level of x_i that maximizes his utility, given the choices of the other consumers. While this seems somewhat strange at first, note that in equilibrium, consumer i ’s beliefs will be confirmed. The other consumers will really choose to purchase x_j^* units of the public good. For those of you who know some game theory, what we’re doing here is finding a Nash equilibrium.

2. Firms maximize profit:

$$\max pq - c(q).$$

For an interior solution, the first-order condition is

$$p = c'(q^*).$$

3. Market clearing: at a competitive equilibrium the price adjusts so that,

$$x^* = \sum_i x_i^* = q^*.$$

Putting conditions 1 and 3 together, we know that

$$\phi'_i(x^*) = c'(x^*)$$

for any i that purchases a positive amount of the external good. Further, for all i that do not purchase the good, $\phi_i(x^*) > 0$. Without loss of generality, suppose consumers 1 through K do not contribute and consumers $K + 1$ through I do contribute. This implies that

$$\sum_i \phi'_i(x^*) = \sum_{i=1}^K \phi'_i(x^*) + \sum_{i=K+1}^I \phi'_i(x^*) \quad (8.1)$$

$$= \sum_{i=1}^K \phi'_i(x^*) + (I - (K + 1))c'(x^*) > c'(x^*). \quad (8.2)$$

whenever a positive amount of the public good is provided.

Now, compare this with the condition defining the Pareto optimal quantity of the public good:

$$\sum_{i=1}^I \phi'_i(q^0) = c'(q^0).$$

Thus, when people make voluntary contributions, the market will provide too little of the public good: $q^0 > q^*$.

The fact that the market provides too little of the public good can be understood in terms of externalities. Purchase of one unit of the public good by one consumer provides an external benefit on all other consumers. More formally, provision of one unit of a public good by consumer i involves a private cost of p^* , a private benefit, $\phi'_i(x^*)$, and a public benefit, $\sum_{j \neq i} \phi'_j(x^*)$. When purchasing units of the public good, individuals weigh the private benefit against the private cost. However, society as a whole is interested in weighing the total benefit against the cost (since $p^* = c'(q)$, the private cost is also the public cost). The fact that individuals do not consider the public benefit




Figure 8.4: The Free-Rider Problem

results in underprovision of the public good. This is frequently called the **free-rider problem**, depicted in Figure 8.4.

For a striking example of the free-rider problem, consider the case where consumers' marginal utilities are increasing in their index:

$$\phi'_1(x) < \phi'_2(x) < \dots < \phi'_I(x) \text{ for all } x.$$

In this case, condition

$$\phi'_i(x_i^* + x_{-i}^*) = p^*.$$

can hold for at most one consumer. Call the consumer who purchases the public good consumer j^* . All of the other consumers must choose $x_i^* = 0$. In the case where $x_i^* = 0$, the first-order condition is: $\phi'_i(x_i^*) \leq p = \phi'_j(x_{j^*}^*)$. This implies that j^* must be consumer I , since $\phi'_I(x) > \phi'_i(x)$ for all $i \neq I$.

The previous example is a particularly stark example of the free-rider problem. The only person who pays for the public good is the person who values it the most (on the margin). All others contribute nothing toward the public good. Real-world examples of something like this include contributions to public television. For another example, think about whether you've ever shared an apartment with a person who either is much neater or much sloppier than you are. Who does all of the cleaning in this case?

8.3.2 Remedies for the Free-Rider Problem

As in the case of bilateral externalities, there are also a number of remedies for the free rider problem in public goods environments. Some remedies include government intervention in the market for the public good. For example, the government may mandate the amount of the public good that consumers must purchase. The government may pass a law requiring inoculations or imposing limits on pollution. For other public goods, such as roads and bridges or national defense, the government may simply take over provision of the public good, taking the decision out of the hands of individual consumers entirely.

The government may also engage in price-based interventions. For example, the government could tax or subsidize the provision of public goods in such a way that private incentives are brought into line with public incentives. Suppose there are I consumers, each with benefit function $\phi_i(x)$. Using the Pigouvian taxation example from the bilateral externality case, we can implement the optimal consumption x^0 by setting the per unit subsidy to each consumer equal to

$$s_i = \sum_{j \neq i} \phi'_j(x^0).$$

This is because (assuming that the other consumers choose x_j^0) the consumer maximizes

$$\phi_i(x_i + x_{-i}^0) + s_i x_i - p x_i.$$

The necessary and sufficient first-order condition for this problem is:

$$\phi'_i(x_i + x_{-i}^0) + s_i = p.$$

Substituting in the above subsidy and combining with the market-clearing condition,

$$\phi'_i(x_i + x_{-i}^0) + \sum_{j \neq i} \phi'_j(x^0) = p^* = c'(x)$$

which is satisfied when $x = x^0$. Thus the optimum is implemented.¹²

While the subsidies described above will implement the Pareto optimal level of the public good, it might be very expensive for the government to do so, since the subsidies can be quite large, and each person is paid the marginal value to all other consumers.

¹²This is an example of the Groves-Clark mechanism. The Groves-Clark mechanism is basically a class of mechanisms in which the external effect of each consumer's decision is added to the private effect, in order to bring individual preferences in line with social preferences.

Lindahl Equilibrium

As in the bilateral externality case, both the quantity- and price-based government interventions require the government to have detailed knowledge of the preferences of consumers and firms. We now present a market-based solution to the problem, known as the Lindahl equilibrium. The idea behind the Lindahl equilibrium is that the public good is unbundled into I private goods, where each good is “person i ’s enjoyment of the public good,” each with its own price (known as the Lindahl price). The equilibrium in this market is known as the Lindahl equilibrium, and it turns out that the Lindahl equilibrium implements the Pareto optimal allocation of the public good.

Suppose that for each i , there is a market for “person i ’s enjoyment of the public good.” Denote the price of this personalized good as p_i . Given an equilibrium price, the consumer chooses the total amount of the good to maximize

$$\phi_i(x) - p_i x_i,$$

which has necessary and sufficient first-order condition

$$\phi_i'(x) = p_i \text{ for each } i. \quad (8.3)$$

Now, consider the producer side of the market. When the firm produces a single unit of the public good, it produces one unit of the personalized public good for every person. That is, one unit of national defense is a bundle of one unit of defense for person 1, one unit for person 2, etc. Hence for each unit of the public good that the firm produces and sells, it earns $\sum_i p_i$ dollars. Hence the firm’s problem is written as

$$\max_q \left(\sum_i p_i q \right) - c(q),$$

which has first-order condition

$$\sum_i p_i = c'(q).$$

Combining this with the consumer’s optimality condition, Equation 8.3, yields

$$\sum_i \phi_i'(q^{**}) = c'(q^{**}),$$

which is the defining equation for the efficient level of the public good. The corresponding prices are $p_i^{**} = \phi_i'(q^{**})$. Thus the Lindahl equilibrium results in the efficient level of the public good being provided.

The Lindahl equilibrium illustrates that the right kind of market can implement the Pareto optimal allocation, even in the public good case. However, Lindahl equilibrium may not be realistic. In particular, the Lindahl equilibrium depends on consumers behaving as price takers, even when they are the only buyers of a particular good. Still, it may be reasonable to think that the consumer has no power to force the producer of the public good to lower its price, especially if the producer is the government. Second, and more troubling, is the idea that in order for the Lindahl equilibrium to work, the consumer has to believe that if they do not purchase any of the public good, they will not be able to consume any of it. Of course, since one of the defining features of a public good is that it is non-excludable, it is unlikely that consumers will believe this.

8.3.3 Club Goods

However, while Lindahl equilibria may not be reasonable for pure (non-excludable) public goods, they are reasonable if the good is excludable, which we earlier called **club goods**. The Lindahl price p_i^{**} can then be thought of as the price of a membership in the “club,” i.e. the right for access to the club good. In this case, the market will result in efficient provision of the club good (although you still have to worry about the price-taker assumption).

8.4 Common-Pool Resources

A common-pool resource (CPR) is a good where consumption is rivalrous and non-excludable. Some of the prototypical examples of CPRs are local fishing grounds, common grazing land, or irrigation systems. In such situations, individuals will tend to overuse the CPR since they will choose the level of usage at which the individual marginal utility is zero, but the Pareto optimum is the level at which the total marginal benefit is zero, which is generally a lower level of consumption than the market equilibrium.

Consider the following example of a common-pool resource. The total number of fish caught in a non-excludable local fishery is given by $f(k)$, where k is the total number of fishing boats that work the fishing ground. Assume that $f' > 0$ and $f'' < 0$, and that $f(0) = 0$. That is, we assume total fish production is an increasing concave function of the number of boats working the fishing ground. Also, note that as a consequence of the concavity of $f()$, $\frac{f(k)}{k} > f'(k)$. That is, the number of fish caught per boat is always larger than the marginal product of adding another boat. This follows from the observation that average product is decreasing for a concave production

function. Let $AP(k) = \frac{f(k)}{k}$. Then $AP'(k) = \frac{1}{k}(f'(k) - AP(k)) < 0$. Hence $AP(k) > f'(k)$.

Fishing boats are produced at a cost $c(k)$, where k is the total number of boats, and $c(\cdot)$ is a strictly increasing, strictly convex function. The price of fish is normalized to 1.

The Pareto efficient number of boats is found by solving

$$\max_k f(k) - c(k),$$

which implies the first-order condition for the optimal number of boats k^0 :

$$f'(k^0) = c'(k^0).$$

Let k_i be the number of boats that fisher i employs, and assume that there are I total fishers. Hence $k = \sum_i k_i$. If p is the market-clearing price of a fishing boat, the boat producers solve

$$\max_k pk - c(k),$$

which has optimality condition

$$c'(k) = p.$$

Under the assumption that each fishing boat catches the same number of fish, each fisher solves the problem

$$\max_{k_i} \frac{k_i}{k_i + k_{-i}} f(k) - pk_i,$$

where $k_{-i} = \sum_{j \neq i} k_j$. The optimality condition for this problem is

$$f'(k^*) \frac{k_i^*}{k^*} + \frac{f(k^*)}{k^*} \left(\frac{k_{-i}^*}{k^*} \right) = p.$$

Market clearing then implies that:

$$f'(k^*) \frac{k_i^*}{k^*} + \frac{f(k^*)}{k^*} \left(\frac{k_{-i}^*}{k^*} \right) = c'(k^*).$$

Since all of our producers are identical, the optimum will involve $k_i^* = k_j^*$ for all i, j . That is, all fishers will choose the same number of boats.¹³ If there are n total fishers, we can rewrite this condition as:

$$f'(k^*) \frac{1}{n} + \frac{f(k^*)}{k^*} \left(\frac{n-1}{n} \right) = c'(k^*).$$

¹³In fact, since $f(\cdot)$ is strictly concave, we could show this if we wanted to.

Thus the left-hand side is a convex combination of the marginal product, $f'(k)$, and the average product, $\frac{f(k)}{k}$. And, since we know that $\frac{f(k)}{k} > f'(k)$, $f'(k^*) \frac{1}{n} + \frac{f(k^*)}{k} \left(\frac{n-1}{n}\right) > f'(k^*)$. Finally, since $c'(k)$ is increasing in k , this implies that $k^* > k^0$. **The market overuses the fishery.**

What causes people to overuse the CPR? Consider once again the fisher's optimality condition:

$$f'(k^*) \frac{k_i^*}{k^*} + \frac{f(k^*)}{k^*} \left(\frac{k_{-i}^*}{k^*}\right) = p.$$

The terms on the left-hand side correspond to two different phenomena. First, if fisher i buys another boat, that increases the total catch, and fisher i gets $\frac{k_i^*}{k^*}$ of that increase. This is the term $f'(k^*) \frac{k_i^*}{k^*}$. Second, because fisher i now has one more boat, he has a greater proportion of the total boats, and so he gains from the fact that a greater proportion of the total catch is given to him. This second effect, which corresponds to the term $\frac{f(k^*)}{k^*} \left(\frac{k_{-i}^*}{k^*}\right)$, can be thought of as a "market stealing" effect.

Note that both of these effects can lead the market to act inefficiently. Since the fisher only gets $\frac{k_i^*}{k^*}$ of the increase in the number of fish due to adding another boat, this will tend to make the market choose too few boats. However, since adding another boat also results in stealing part of the catch from the other fishers, and this is profitable, this tends to make fishers buy too many boats. When $f()$ is concave, we know that the latter effect dominates the former, and there will be too many boats.

The previous phenomenon, that the market will tend to overuse common-pool resources, is known as the **tragedy of the commons**. We won't have time to formally go through all of the solutions to this problem, but they include the same sorts of tools that we have already seen. For example, placing a quota on the number of boats each fisher can own or the number of fish that each boat can catch would help to solve the problem. In addition, putting a tax on boats or fish would also help to solve the problem. The appropriate boat tax would be

$$t^* = \frac{k_{-i}^*}{k^*} \left(\frac{f(k^*)}{k^*} - f'(k^*) \right).$$

In this case, the fisher's problem becomes:

$$\max_{k_i} \frac{k_i}{k_i + k_{-i}} f(k) - pk_i - t^* k_i.$$

The first-order condition is:

$$f'(k) \frac{k_i}{k} + \frac{f(k)}{k} \left(\frac{k_{-i}}{k}\right) - t^* = p.$$

Substituting in the value for t^*

$$f'(k) \frac{k_i}{k} + \frac{f(k)}{k} \left(\frac{k_{-i}}{k} \right) - \frac{k_{-i}^*}{k^*} \left(\frac{f(k^*)}{k^*} - f'(k^*) \right) = p.$$

When combined with the market-clearing condition, this becomes:

$$f'(k) \frac{k_i}{k} + \frac{f(k)}{k} \left(\frac{k_{-i}}{k} \right) - \frac{k_{-i}^*}{k^*} \left(\frac{f(k^*)}{k^*} - f'(k^*) \right) = c'(k),$$

which is solved at $k_i = \frac{k^*}{n}$ for all i .

Privatization (or nationalization) is another solution. If the whole CPR is held by one person, then they will choose $f'(k) = p$, and (assuming price-taking) $f'(k^*) = c'(k^*)$. The owner of the CPR and the consumers can then bargain over usage. However, privatization can have problems of its own, including managerial difficulties and adverse distributional consequences.

In addition to the tools I've already mentioned, it is important to note that while the tragedy of the commons is a real problem, people have been solving it for hundreds of years, at least in part. Frequently, when the people who have access to a CPR (such as common grazing land, an irrigation system, or an open-access fishery) are part of a close community, such as the residents of a village, they find informal ways to cooperate with each other. Since the villagers know each other, they can find informal ways to punish people who overuse the resources.

Chapter 9

Monopoly

As you will recall from intermediate micro, monopoly is the situation where there is a single seller of a good. Because of this, it has the power to set both the price and quantity of the good that will be sold. We begin our study of monopoly by considering the price that the monopolist should charge.¹

9.1 Simple Monopoly Pricing

The object of the firm is to maximize profit. However, the price that the monopolist charges affects the quantity it sells. The relationship between the quantity sold and the price charged is governed by the (aggregate) demand curve $q(p)$. Note, in order to focus on the relationship between q and p , we suppress the wealth arguments in the aggregate demand function.

We can thus state the monopolist's problem as follows:

$$\max_p pq(p) - c(q(p)).$$

Note, however, that there is a one-to-one correspondence between the price charged and the quantity the monopolist sells. Thus we can rewrite the problem in terms of quantity sold instead of the price charged. Let $p(q)$ be the inverse demand function. That is, $p(q(p)) = p$. The firm's profit maximization problem can then be written as

$$\max_q p(q)q - c(q).$$

It turns out that it is usually easier to look at the problem in terms of setting quantity and letting price be determined by the market. For this reason, we will use the quantity-setting approach.

¹References: Tirole, Chapter 1; MWG, Chapter 12; Bulow, "Durable-Goods Monopolists," *JPE* 90(2) 314-332.




Figure 9.1: The Monopolist's Marginal Revenue

In order for the solution to be unique, we need the objective function to be strictly concave (i.e. $\frac{d^2\pi}{dq^2} < 0$). The second derivative of profit with respect to q is given by

$$\frac{d^2}{dq^2} (p(q)q - c(q)) = p''(q)q + 2p'(q) - c''(q).$$

If cost is strictly convex, $c''(q) > 0$, and since demand slopes downward, $p'(q) < 0$. Hence the second and third terms are negative. Because of this, we don't need inverse demand to be concave. However, it can't be "too convex." Generally speaking, we'll just assume that the objective function is concave without making additional assumptions on $p(\cdot)$. Actually, to make sure the maximizing quantity is finite, we need to assume that eventually costs get large enough relative to demand. This will always be satisfied if, for example, the demand and marginal cost curves cross.

The objective function is maximized by looking at the first derivative. At the optimal quantity, q^* ,

$$p'(q^*)q^* + p(q^*) = c'(q^*)$$

On the left-hand side of the expression is the marginal revenue of increasing output a little bit. This has two parts - the additional revenue due to selling one more unit, $p(q^*)$ (area B in Figure 9.1), and the decrease in revenue due to the fact that the firm receives a lower price on all units it sells (area A in Figure 9.1). Hence the monopolist's optimal quantity is where marginal revenue is equal to marginal cost, and price is defined by the demand curve $p(q^*)$.² See Figure 9.2 for the graphical depiction of the optimum.

²This is also true for the competitive firm. However, since competitive firms are price takers, their marginal revenue is equal to price.




Figure 9.2: Monopolist's Optimal Price and Quantity




Figure 9.3: The Monopolist Cannot Make a Profit

If the monopolist's profit is maximized at $q = 0$, it must be that $p(0) \leq c'(0)$. This corresponds to the case where the cost of producing even the first unit is more than consumers are willing to pay. We will generally assume that $p(0) > c'(0)$ to focus on the interesting case where the monopolist wants to produce a positive output. However, even if we assume that $p(0) > c'(0)$, the monopolist may not want to choose a positive output. It may be that shutting down is still preferable to producing a positive output. That is, the monopolist may have fixed costs that are so large that it would rather exit the industry. Such a situation is illustrated in Figure 9.3. Thus we interpret the condition that $p(0) > c'(0)$ (along with the appropriate second-order conditions) as saying that if the monopolist does not exit the industry, it will produce a positive output.

If there is to be a maximum at a positive level of output, it must be that the first derivative equals zero, or:

$$p'(q^*)q^* + p(q^*) = c'(q^*)$$

Note that one way to rewrite the left side is:

$$p'(q^*)q^* + p(q^*) = p(q^*) \left(\frac{dp}{dq} \frac{q^*}{p^*} + 1 \right) = p(q^*) \left(1 - \frac{1}{|\varepsilon_p^*|} \right)$$

where ε_p^* is the price elasticity of demand evaluated at (q^*, p^*) . Now we can rewrite the monopolist's first-order condition as:

$$\frac{p(q^*) - c'(q^*)}{p(q^*)} = p'(q^*) \frac{-q^*}{p(q^*)} = \frac{1}{|\varepsilon_p^*|}.$$

The left-most quantity in this expression, $(p - mc)/p$ is the “markup” of price over marginal cost, expressed as a fraction of the price. This quantity, called the Lerner index, is frequently used to measure the degree of market power in an industry.

Note that at the quantity where $MR = MC$, $p > MC$ (since $p'(q)q$ is negative). Thus the monopolist charges more than marginal cost.³ The social optimum would be for the monopolist to sell output as long as consumers are willing to pay more for the last unit produced than it costs to produce it. That is, produce up until the point where $p = MC$. But, the monopolist cuts back on production because it cares about profit, not social optimality. And, it is willing to reduce q in order to increase the amount it makes per unit. This results in what is known as the **deadweight loss of monopoly**, and it is equal to the area between the demand curve and the marginal cost curve, and to the right of the optimal quantity, as in Figure 9.4. This area represents social surplus that could be generated but is not in the monopoly outcome.

9.2 Non-Simple Pricing

The fact that the monopolist sells less than the societally optimal amount of the output arises from the requirement that the monopolist must sell all goods at the same price. Thus if it wants to earn higher profits on any particular item, it must raise the price on all items, which lowers the quantity sold. If the monopolist could raise the price on some items but not others, it could earn higher profits and still sell the efficient quantity. We now consider two examples of more complicated pricing mechanisms along these lines, non-linear pricing and two-part tariffs.

³Also, from these two equations we can show that the monopolist will always choose a quantity such that price is elastic, i.e. $|\varepsilon_p^*| > 1$.




Figure 9.4: Deadweight Loss of Monopoly

9.2.1 Non-Linear Pricing

Consider the case where the monopolist charges a price scheme where each unit is sold for a different price.

For the moment, we assume that all consumers are identical and that consumers are not able to resell items once they buy them. In this case, the monopolist solves its profit-maximization problem by designing a scheme that maximizes the profit earned on any one consumer and then applying this scheme to all consumers.

In the context of the quasilinear model, we know that the height of a consumer's demand function at a particular quantity represents his marginal utility for that unit of output. In other words, the height of the demand curve represents the maximum a consumer will pay for that unit of output.

With this in mind, if the monopolist is going to charge a different price for each unit of output, how should it set that price? Obviously, it wants to set the price of each unit equal to the consumer's maximum willingness to pay for that unit, i.e. equal to the height of the demand curve at that q . If the monopolist employs a declining price scheme, where $p(q) = D^{-1}(q)$, up to the point where demand and marginal cost cross, the monopolist can actually extract all of the social surplus. Note: if we are worried that output can only be sold in whole units, then the price of unit q should be given by:

$$p_q = \int_{q-1}^q D^{-1}(q) dq.$$




Figure 9.5: Non-Linear Pricing

To illustrate non-linear pricing, consider a consumer who has demand curve $P = 100 - Q$, and suppose the monopolist's marginal cost is equal to 10. At a price of 90, the consumer demands 10 units of output. While the consumer would not purchase any more output at a price of 90, it would purchase more output if the price were lower. So, suppose the monopolist sells the first 10 units of output at a price of 90, and the second 10 units at a price of 80. Similarly, suppose the monopolist sells units 21-30 at a price of 70, 31-40 at a price of 60, 41-50 at a price of 50, etc., all the way up to units 71-80, which are sold at a price of 20. This yields Figure 9.5.

The monopolist's producer surplus is equal to the shaded region. Thus by decreasing the price as the number of units purchased increases, the monopolist can appropriate much of the consumer surplus. In fact, as the number of "steps" in the pricing scheme increases, the producer surplus approaches the entire triangle bounded by the demand curve and marginal cost. Thus if the monopolist were able to charge the consumer $\int_{q-1}^q D^{-1}(q) dq$ for the block of output consisting of unit q , it would appropriate the entire consumer surplus.

In practice, declining block pricing is found most often in utility pricing, where a large buyer may be charged a high price for the first X units of output, a lower price for the next X units, and so on.

Recall that we said that a monopolist charging a single price can make a positive profit only if the demand curve is above the firm's average total cost (ATC) curve at some point. If the firm can engage in non-linear pricing, this is no longer the case. Since the firm's revenue is the entire shaded region in the above diagram, it may be possible for the firm to earn a positive profit even if the demand curve is entirely below the ATC curve. For example, if ATC in the previous example




Figure 9.6: Profit Even With Demand Below ATC

is as in Figure 9.6, then the monopolist can make a profit whenever the area of the shaded region is larger than $TC(90) = ATC(90) \times 90$.

9.2.2 Two-Part Tariffs

Another form of non-simple pricing that is frequently employed is the two-part tariff. A two-part tariff consists of a fixed fee and a price per unit consumed. For example, an amusement park may charge an admission fee and a price for each ride, or a country club may charge a membership fee and a fee for each round of golf the member plays. The question we want to address is how, in this context, the fixed fee (which we'll call F) and the use fee (which we'll call p) should be set.

As in the non-linear pricing example, we continue to assume that all consumers are identical and that the monopolist produces output at constant marginal cost. For simplicity, assume that $MC = 0$, but note that a positive marginal cost is easily incorporated into the model. Further, we continue to think of the consumer as having quasilinear utility for output q and a numeraire good. In this case, we know that the consumer's inverse demand curve, $p(q)$, gives the consumer's marginal utility from consuming unit q , and that the consumer's net benefit from consuming q units is given by the consumer surplus at this quantity, $CS(q)$.

The firm's profit from two-part tariff (F, p) is given by $F + p * q(p)$. We break our analysis into two steps. First, for any p , how should F be set? Second, which p should be chosen? So, fix a p . How should F be set. Since the monopolist's profit is increasing in F , the monopolist wants to set F as large as possible while still inducing the consumer to participate. That is, the consumer's net benefit must be at least zero.




Figure 9.7: The Two-Part Tariff

At price p , the consumer chooses to consume $q(p)$ units of output, and earns surplus $CS(q(p))$. Net surplus is therefore $CS(q(p)) - F$, and we want this to be non-negative: $CS(q(p)) - F \geq 0$. As argued above, the firm wants to set F as large as possible. Therefore, for any p , setting $F(p) = CS(q(p))$ maximizes profit. This makes sense, since the firm wants to set the membership fee in order to extract all of the consumer's surplus from using its product.

Next, how should p be chosen? Given $F(p)$ as defined above, the monopolist's profit from the two-part tariff is given by: $CS(q(p)) + p * q(p)$. Note that this is the sum of consumer and producer surplus (see Figure 9.7). As is easily seen from the diagram, this expression is maximized by setting the use fee equal to zero (marginal cost) and the fixed fee equal to $CS(q(0))$. Thus the optimal two-part tariff involves $p = 0$ and $F = CS(0)$.⁴

We can understand the optimal two-part tariff in another way. When the monopolist chooses F optimally, it claims the entire surplus created from the market. Thus the monopolist has an incentive to maximize the surplus created in this market. By the first welfare theorem, we know that total surplus is maximized by the perfectly competitive outcome. That is, when p is such that $p^* = MC$. Thus, in order to maximize total surplus, the firm should set $p = MC$. If it does so, "producer surplus" is zero. However, the firm sets $F = CS(q(p^*))$, and extracts the entire social surplus in the form of the fixed fee.

⁴More generally, if the firm has positive marginal cost, then p should be set such that $p = MC$ at $q(p)$, and F should be set equal to $CS(q(p))$.

Interestingly, there seems to be a trend toward this sort of two-part pricing in recent years. For example, at Walt Disney World, the standard admission package charges for entry into one of the several theme parks but charges nothing for rides once inside the park. This makes sense, since the marginal cost of an additional ride is very low.⁵

9.3 Price Discrimination

Previously we considered the case where the monopolist faces one or many identical consumers, and we investigated various pricing schemes the monopolist might pursue. However, typically, the monopolist faces another problem, which is that some consumers will have a high willingness to pay for the product and others will have a low willingness to pay. For example, business travellers will be willing to pay more for airline tickets than leisure travellers, and working-age adults may be more willing to pay for a movie ticket than senior citizens. We refer to the monopolist's attempts to charge different prices to these different groups of people as **price discrimination**. Frequently, three types of price discrimination are identified, although the distinctions are, at least to some extent, arbitrary. They are called first-degree, second-degree, and third-degree price discrimination.⁶

9.3.1 First-Degree Price Discrimination

First-degree price discrimination - also called "perfect price discrimination" - refers to the situation where the monopolist is able to sell each unit of output to the person that values it the most for their maximum willingness to pay. For example, think of 100 consumers, each of whom have demand for one unit of a good. Let $p(q)$ be the q^{th} consumer's willingness to pay for the good (note, this is just the inverse demand curve). A perfectly discriminating monopolist is able to charge each consumer $p(q)$ for the good: Each consumer is charged her maximum willingness-to-pay. In this case, the monopolist's marginal revenue is equal to $p(q)$, and so the monopolist sets quantity q so that $p(q) = c'(q)$, which is exactly the condition for Pareto optimality. Figure 9.8 depicts the price-discrimination optimum, which is also the Pareto optimal quantity.

We can also think of first-degree price discrimination where each individual has a downward

⁵This contrasts with the scheme Disney used when I was younger, which involved both an admission fee and positive prices for each ride. In fact, they even charged higher prices for the most popular rides.

⁶There are many different definitions of the various kinds of price discrimination. The ones I use are based on Varian's *Intermediate Microeconomics*.




Figure 9.8: First-Degree Price Discrimination is Efficient

sloping demand for the good. In this case, the monopolist will sell a quantity such that $p^* = p(q) = c'(q)$. However, each individual will be sold quantity q_i^* , where $p^* = p_i(q_i^*)$, and $p_i(q_i)$ is individual i 's inverse demand curve. The monopolist will charge the consumer the total consumer surplus associated with q_i^* units of output,

$$t_i^* = \int_0^{q_i^*} p_i(s) ds.$$

Thus the monopolist's selling scheme is given by a quantity and total charge (q_i, t_i) for each consumer. This scenario is depicted in Figure 9.9: In the left panel of the diagram, we see the aggregate demand curve, and the price-discriminating monopolist sets total quantity equal to where marginal cost equals aggregate demand (unlike the non-discriminating monopolist, who sets quantity where marginal cost equals marginal revenue). The panel on the right shows that the demand side of the market consists of two different types of consumers (D_1 and D_2), and the monopolist maximizes profits by dividing the quantity between the two buyers such that the marginal willingness-to-pay (WTP) is equal across the two buyers and equals marginal cost: $WTP_1 = WTP_2 = MC$. And the price is set for each consumer as specified above for t_i^* , to capture the full surplus.

In order for first-degree price discrimination to be possible, let alone successful, the monopolist must be able to identify the consumer's willingness to pay (or demand curve) and charge a different price to each consumer. There are two problems with this. First, it is difficult to identify a consumer's willingness to pay (they're not just going to tell you); and second, it is often impractical or illegal to tailor your pricing to each individual. Because of this, first-degree price discrimination




Figure 9.9: First-Degree Price Discrimination

is perhaps best thought of as an extreme example of the maximum (but rarely attainable) profit the monopolist can achieve.

9.3.2 Second-Degree Price Discrimination

Second-degree price discrimination refers to the case where the monopolist cannot perfectly identify the consumers. For example, consider a golf course. Some people are going to use a golf course every week and are willing to pay a lot for each use, while others are only going to use the course once or twice a season and place a relatively low value on a round of golf. The owner of the golf course would like to charge a high price to the high-valued users, and a low price to the low-valued users. What's the problem with this? On any given day, people will have an incentive to act like low-valued users and pay the lower price. So, in order to be able to charge the high-valued people a high price and the low-valued people a low price, the monopolist needs to design a pricing scheme such that the high-valued people do not want to pretend to be low-valued people. This practice goes by various names, including second-degree price discrimination, non-linear pricing, and screening.

Let's think about a simple example.⁷ Suppose the monopolist has zero marginal cost, and the demand curves for the high- and low-valued consumers are given by $D_H(p)$ and $D_L(p)$. Assume for the moment that there are equal numbers of high- and low-valued buyers. This scenario is depicted in Figure 9.10.

⁷This example is based on the analysis in Varian's *Intermediate Microeconomics*.




Figure 9.10: Second-Degree Price Discrimination

If the monopolist could perfectly price discriminate, it would charge the low type a total price equal to the area of region B and sell the low type Q_L units, and charge the high type total price $A + B$ for Q_H units. Notice that under perfect price discrimination, consumers earn zero surplus. The consumer surplus from consuming the object is exactly offset by the total cost of consuming the good.

Now consider what would happen if the firm cannot identify the low and high types and charge them different prices. One thing it could do is simply offer (Q_L, B) and $(Q_H, A + B)$ and let the consumers choose which offer they want. Clearly, all of the low-valued buyers would choose (Q_L, B) (why?). But, what about the high-valued buyers? If they choose $(Q_H, A + B)$, their net surplus is zero. But, if they choose (Q_L, B) , they earn a positive net surplus equal to the area between D_H and D_L to the left of Q_L (again, why?). Thus, given the opportunity, all of the high-valued buyers will self-select and choose the bundle intended for the low-valued buyers.

The monopolist can get around this self-selection problem by changing the bundles. Specifically, if the bundles intended for the high-valued and low-valued consumers offered the same net surplus to the high-valued consumers, the high-valued consumers would choose the bundle intended for them.⁸ If the monopolist lowered the price charged to the high types, they would earn more surplus from taking the bundle intended for them, and this may induce them to actually select that bundle. For example, consider Figure 9.11.

⁸Here we make the common assumption that if an agent is indifferent between two actions, he chooses the one the economist wants. This is pretty innocuous, since we could always change the offer slightly to make the agent strictly prefer one of the options. Technically, the reason for this assumption is to avoid an “open set” problem: The set of bundles that the high type strictly prefers to a particular bundle is open, and thus there may not be a profit-maximizing bundle for the monopolist, if we don’t assume that indifference yields the desired outcome.




Figure 9.11: Monopolistic Screening

The optimal **first-degree** pricing scheme is to offer $\beta_1 = (Q_3, E + F + G)$ to the low type and $\beta_2 = (Q_4, A + B + C + D + E + F + G)$ to the high type. However, high prefers β_1 , which offers him surplus $A + B + C$ to β_2 . In order to get the high-valued type to accept a bundle offering Q_4 , the monopolist can charge no more than $E + F + G + D$. Let's call this bundle $\beta_3 = (Q_4, E + F + G + D)$. If the monopolist has zero production costs, selling β_3 to the highs and β_1 to the lows will be better than selling β_1 to all consumers.

The key point in the previous paragraph is that the monopolist must design the bundles so that the buyers self-select the proper bundle. If you want the highs to accept β_3 , it had better be that there is no other bundle that offers them higher surplus.

The monopolist can do even better than offering β_1 and β_3 . Suppose the monopolist offers $\beta_4 = (Q_2, E + F)$ to low and $\beta_5 = (Q_4, C + D + E + F + G)$ to high. Since high earns surplus $A + B$ from either bundle, the self-selection constraint is satisfied. The monopolist earns profit $C + D + 2E + 2F + G$, as compared to $2E + 2F + 2G + D$ from offering β_1 and β_3 . So, as long as area C is larger than area G , the monopolist earns more profit from offering β_4 and β_5 than from offering β_1 and β_3 .

What is the optimal menu of bundles? In the previous paragraph, G is the revenue given up by making less from the lows, while C is the revenue gained by making more from the highs. The optimal bundle will be where these two effects just offset each other. In the diagram, they are shown by $\beta_6 = (Q_1, E)$ and $\beta_7 = (Q_4, B + C + D + E + F + G)$.

Figure 9.12 gives a clean version of the diagrams we have been looking at above. The optimal second-degree pricing scheme is illustrated. The monopolist should offer menu (Q_L^*, A) and




Figure 9.12: Optimal Second-Degree Pricing Scheme

$(Q_H^*, A + B + C)$. Under this scheme, the low-valued consumers choose (Q_L^*, A) and earn no surplus. Notice that Q_L^* is less than the Pareto optimal quantity for the low-valued consumers. The high-valued consumers are indifferent between the two bundles, and so choose the one we want them to choose: $(Q_H^*, A + B + C)$. Notice that under this scheme the high-valued consumers are offered the Pareto optimal quantity, Q_H^* , but earn a positive surplus. Overall, as the quantity offered to the low-valued consumers decreases, the monopolist gains on the margin the area marked “gain” in the diagram and loses the area market “loss.” The optimal level of Q_L^* is where these lengths are just equal.

Once we have derived the optimal second-degree pricing scheme, there is still one thing we have to check. The monopolist could always decide it is not worth it to separate the two types of consumers. Rather it could just offer Q_H^* at price $A + B + C + D$ in Figure 9.12 and sell only to the high-valued consumers. Hence, after deriving the optimal self-selection scheme, we must still make sure that the profit under this scheme, $2A + C + B$, is greater than the profit from selling only to the high types, $A + B + C + D$. Clearly, this depends on the relative size of A (the profit earned by selling to the low-types under the optimal self-selection mechanism) and D (the rent that must be given to the high type to get them to buy when offer (Q_L^*, A) is also on the menu).⁹

The previous graphical analysis is different in style than you are used to. I show it to you for a couple of reasons. First, it is a nice illustration of a type of problem that we will see over and over again, called the monopolistic screening (or hidden-information principal-agent) problem.¹⁰ One

⁹If there were different numbers of high- and low- type consumers, we would have to weight the size of the gains and losses by the size of their respective markets.

¹⁰Since you haven’t seen these other problems yet, this paragraph may not make sense to you. That’s okay. We’ll

party (call her the principal) offers a menu of contracts to another party (call him the agent) about which she does not know some information. The contracts are designed in such a way that the agents self-select themselves, revealing that information to the principal. The menu of contracts is distorted away from what would be offered if the principal knew the agent's private information. Further, it is done in a specific way. The high type is offered the full-information quantity, but a lower price (phenomena like these are often referred to as “no distortion at the top” results), while the low type is offered less than the full-information quantity at a lower price. Distorting the quantity offered to the low type allows the principal to extract more profit from the high type, but in the end the high type earns a positive surplus. This is known as the informational rent – the extra payoff high gets due to the fact that the principal does not know his type.¹¹ As I said, this is a theme we will return to over and over again.

9.3.3 Third-Degree Price Discrimination

Third-degree price discrimination refers to a situation where the monopolist sells to different buyers at different prices based on some observable characteristic of the buyers. For example, senior citizens may be sold movie tickets at one price, while adults pay another price, and children pay a third price. Self-selection is not a problem here, since the characteristic defining the groups is observable and verifiable, at least in principle. For example, you can always check a driver's license to see if someone is eligible for the senior citizen price or not.

The basics of third-degree price discrimination are simple. In fact, you really don't need to know any economics to figure it out. Let $p_1(q_1)$ and $p_2(q_2)$ be the aggregate demand curves for the two groups. Let $c(q_1 + q_2)$ be the firm's (strictly convex) cost function, based on the total quantity produced. The monopolist's problem is:

$$\max_{q_1, q_2} p_1(q_1)q_1 + p_2(q_2)q_2 - c(q_1 + q_2).$$

come back to this point later.

¹¹This is only one type of principal-agent problem. Often, when people refer to “the” principal-agent problem, they are referring to a situation where the agent makes an unobservable effort choice and the principal must design an incentive contract that induces him to choose the correct effort level. To be precise, such models should be referred to as “hidden-action principal-agent” models, to differentiate them from “hidden-information principal-agent” problems, such as the monopolistic screening or second-degree price discrimination problems. See MWG Chapter 14 for a discussion of the different types of principal-agent models.

The first-order conditions for an interior maximum are:

$$\begin{aligned} p_1'(q_1^*) q_1^* + p_1(q_1^*) &= c'(q_1^* + q_2^*) \\ p_2'(q_2^*) q_2^* + p_2(q_2^*) &= c'(q_1^* + q_2^*). \end{aligned}$$

In general, you would need to check second-order conditions, but let's just assume they hold. The first-order conditions just say the monopolist should set marginal revenue in each market equal to total marginal cost. This makes sense. Consider the last unit of output. If the marginal revenue in market 1 is greater than the marginal revenue in market 2, you should sell it in market 1, and vice versa. Hence any optimal selling scheme must set MR equal in the two markets. And, we already know that MR should equal MC at the optimum.

So, third-degree price discrimination is basically simple. But, it can have interesting implications. For example, suppose there are two types of demand for Harvard football tickets. Alumni have demand $p_a(q_a) = 100 - q_a$, while students have demand $p_s(q_s) = 20 - 0.1q_s$. Suppose the marginal cost of an additional ticket is zero. How many tickets of each type should Harvard sell? Set MR = MC in each market:

$$\begin{aligned} 20 - 0.2q_s &= 0 \\ 100 - 2q_a &= 0, \end{aligned}$$

So, $q_a^* = 50$ and $q_s^* = 100$ as well. Alumni tickets are sold at $p_a = 50$, while student tickets are sold at price $p_s = 10$. Total profit is $50 * 50 + 100 * 10 = 2500 + 1000 = 3500$.

Now, suppose that the stadium capacity is 151 seats (for the sake of argument), and that all seats must be sold. Should the remaining seat be sold to alumni or students? Think of your answer before going on.

If the additional seat is sold to alumni, $q_a = 51$, $p_a = 49$, and total profit on alumni sales is 2499, yielding total profit 3499. On the other hand, if the additional seat is sold to students, $q_s = 101$, $p_s = 20 - 0.1(101) = 9.9$, total profit on student sales is 999.9, and so total profit is 3499.9. Hence it is better to sell the extra ticket to the students, even though the price paid by the alumni is higher. Why? The answer has to do with marginal revenue. We know that in either case selling another unit of output decreases total profit (why?). By selling an additional alumni ticket, you have to lower the price more than when selling another student ticket. Hence it is better to sell the ticket to a student, not because more is made on the additional ticket, but because less is lost due to lowering price on the tickets sold to all other buyers in that market.

9.4 Natural Monopoly and Ramsey Pricing

From the point of view of efficiency, monopolies are a bad thing because they impose a deadweight loss.¹² The government takes a number of steps to prevent monopolies. For example, patent law grants a firm exclusive rights to an invention for a number of years in exchange for making the design of the item available, and allowing other firms to license the use of the technology after a period of time. Other ways the government opposes monopolies is through anti-trust legislation, under which the government may break up a firm deemed to exercise too much monopoly power (such as AT&T) or prevent mergers between competitors that would create monopolies..

However, there are some monopolies that the government chooses not to break up. Instead the government allows the monopoly to operate (and even sanctions its operation) but regulates the prices it can charge. Why would the government do such a thing? The government allows monopolies to exist when they are so-called natural monopolies. A **natural monopoly** is an industry where there are high fixed costs and relatively small variable costs. This implies that AC will be decreasing in output. Because of this, it makes sense to have only one firm.¹³ The best examples of natural monopolies are utilities such as the electric, gas, water, and (formerly) telephone companies. Take the water company: In order to provide households with water, you need to purify and filter the water and pass it through a network of pipes leading from the filtration plant to the consumer's house. The filtration plants are expensive to build, and the network of pipes is expensive to install and maintain. Further, they are inconvenient, since laying and maintaining pipe can disrupt traffic, businesses, etc. Think of how inefficient it would be if there were four or five different companies all trying to run pipes into a person's house!

Because of the inefficiency involved in having multiple providers in an industry that is a natural monopoly, the government will allow a single firm to be the monopoly provider of that product, but regulate the price that it is allowed to charge. What price does the government choose? Consider the diagram of a natural monopolist in Figure 9.13.

¹²Of course, there are also critical distribution issues with monopolies, and these issues can persist even in situations in which the DWL of monopoly has been eliminated. For example, with perfect price discrimination, there is no DWL - the efficient quantity is sold - but the distribution of surplus (all to the monopolist, zero to consumers) is clearly a cause for societal concern.

¹³The technical definition of a natural monopoly is that the industry cost function is subadditive. That is, $c(q_1 + q_2) \leq c(q_1) + c(q_2)$. Hence it is always cheaper to produce $q_1 + q_2$ units of output using a single firm than using two (or more) firms. This is a slightly weaker definition of a natural monopoly than decreasing average cost, especially in multiple dimensions.




Figure 9.13: Ramsey Pricing

The monopolist has constant marginal cost, $c'(q) = MC$, and positive fixed cost. Note that this implies that $AC > MC$ for all q , the defining feature of a natural monopolist.

If the monopoly is not regulated, it will charge price p^M . The Pareto optimal quantity to sell is where inverse demand equals marginal cost, labeled Q^E in the diagram. The corresponding price is $P^E = MC$. However, since this price is below AC , the monopolist will not be able to cover its costs if the government forces it to charge P^E . In order for the monopolist to cover its production costs, it must be allowed to charge a higher price. However, as it increases price above MC , the quantity drops below Q^E , and there is a corresponding deadweight loss. Thus the government's task is to strike a balance between allowing the monopolist to cover its costs and keeping prices (and deadweight loss) low. The price that does this is the smallest price at which the monopolist is able to cover its cost. This price is labeled P^R in the diagram. The R stands for Ramsey, and the practice of finding the prices that balance deadweight loss and allow the monopolist to cover its costs is called Ramsey pricing. Finding the Ramsey price is easy in this example, but when the monopolist produces a large number of products (such as electricity at different times of the day and year, and electricity for different kinds of customers), the Ramsey pricing problem becomes much harder.¹⁴ Ramsey pricing (and other related pricing practices) is one of the major topics in the economics of regulation (or at least it was for a long time).¹⁵

¹⁴Ramsey pricing is really a topic for monopolists that produce multiple outputs. The Ramsey prices are then the prices that maximize consumer surplus subject to the constraint that the firm break even overall.

¹⁵Another type of regulatory mechanism is the regulatory-constraint mechanism, where the monopolist is allowed

In recent years, technology has progressed to the point where many of the industries that were traditionally thought to be natural monopolies are being deregulated. Most of these have been “network” industries such as electric or phone utilities. Technological advances have made it possible for a number of providers to use the same network. For example, a firm can generate electricity and put it on the “power grid” where it can then be sold, or competing local exchange carriers (phone companies) can provide phone service by purchasing access to Ameritech’s network at prespecified rates.¹⁶ Because it is now possible for multiple firms to use the same network, these industries are being deregulated, at least in part. Generally, the “network” remains a regulated monopoly, with specified rates and terms for allowing access to the network. Provision of services such as electricity generation or connecting phone calls is opened up to competition.

9.4.1 Regulation and Incentives

When a monopolist is regulated, the government chooses the output price so that the monopolist just covers its costs.¹⁷ Because the monopolist knows that it will cover its costs, it will not have an incentive to keep its costs low. As a result, it may incur costs that it would not incur if it were subjected to market discipline. For example, it could purchase fancy office equipment, and decorate the corporate headquarters. This phenomenon is sometimes known as gold plating. Gold-plating is something regulators look out for when they are determining the cost base for the regulated firm.

An interesting phenomenon occurs when a firm operates both in a regulated and an unregulated industry. For example, consider local phone providers, which are regulated monopolists on local service (the “loop” from the switchbox to your house) but one of many competitors on “local toll” calls (calls over a certain distance – like 15 miles). Since the firm knows it will cover its costs on the local service, it may try to classify some of the costs of operating its competitive service as costs of local service, thereby gaining a competitive advantage in the local toll market.

The extent to which the problems I mentioned here are real problems depend on the industry. However, they are things that regulators worry about. Whenever a regulated monopoly goes before a rate commission to ask for a rate increase, the regulators ask whether the costs are appropriate to do whatever it wants, subject to a constraint such as its return on assets can be no larger than a pre-specified number. If you are interested in such things, see Berg and Tschirhart, *Natural Monopoly Regulation* or Spulber, *Regulation and Markets*.

¹⁶Ameritech is the local phone company in the Chicago area, where I started writing these notes. In the Boston area (where I’m adding this note), the relevant company is Verizon.

¹⁷Actual regulation usually allows the firm to cover its costs and earn a specified rate of return on its assets.

and whether they should be allocated to the regulated or competitive sector of the monopolist's business. The question of how regulated firms respond to regulatory mechanisms, i.e. regulatory incentives, is another major subject in the economics of monopolies.

9.5 Further Topics in Monopoly Pricing

9.5.1 Multi-Product Monopoly

When a monopolist sells more than one product, it must take into account that the price it charges for one of its products may affect the demand for its other products.¹⁸ This is true in a wide variety of contexts, but let's start with a simple example, that of a monopolist that sells goods that are perfect complements. For example, think of a firm that sells vacation packages that consist of a plane trip and a hotel stay. Consumers care only about the total cost of the vacation. The higher the price of a hotel room, the less people will be willing to pay for the airline ticket, and vice versa.

Demand for vacations is given by:

$$q(p_V) = 100 - p_V,$$

where p_V is the price of a vacation, $p_V = p_A + p_H$, and p_A and p_H are the prices of airline travel and hotel travel respectively. Each airline trip costs the firm c_H , and each hotel stay costs c_H .

To begin, consider the case where the firm realizes that consumers care only about the price of a vacation, and so it chooses p_V in order to maximize profit.

$$\max_{p_V} (100 - p_V)(p_V - (c_H + c_A)),$$

since the cost of a vacation is $c_H + c_A$. The first-order condition for this problem is:

$$\begin{aligned} \frac{d}{dp_V} ((100 - p_V)(p_V - (c_H + c_A))) &= 0 \\ p_V^* &= \frac{100 + c_H + c_A}{2} = 50 + \frac{c_H + c_A}{2}. \end{aligned}$$

Thus if the firm is interested in maximizing total profit, it should set p_V^* as above, and divide the cost among the plane ticket and hotel any way it wants. In fact, if you've ever bought a tour, you

¹⁸See Tirole, *Theory of Industrial Organization*, starting at p. 70.

know that you usually don't get separate prices for the various components. Profit is given by:

$$\begin{aligned} & \left(100 - \frac{100 + c_H + c_A}{2}\right) \left(\frac{100 + c_H + c_A}{2} - c_H - c_A\right) \\ &= \frac{1}{4}(100 - (c_H + c_A))^2. \end{aligned}$$

Now consider the case where the price of airline seats and hotel stays are set by separate divisions, each of which cares only about its own profits. In this case, the hotel division takes the price of the airline division as given and chooses p_H in order to maximize its profit:

$$(100 - p_H - p_A)(p_H - c_H)$$

which implies that the optimal choice of p_H responds to the airline price p_A according to the "reaction curve" or "best-response function" :

$$p_H = \frac{100 + c_H - p_A}{2}.$$

Similarly, the airline division sets p_A in order to maximize its profit, yielding reaction curve:

$$p_A = \frac{100 + c_A - p_H}{2}.$$

At equilibrium, the optimal prices set by the two divisions must satisfy the two equation system:¹⁹

$$\begin{aligned} p_A &= \frac{100 + c_A - p_H}{2} \\ p_H &= \frac{100 + c_H - p_A}{2} \end{aligned}$$

which has solution:

$$\begin{aligned} p_A^* &= \frac{100 + 2c_A - c_H}{3} \\ p_H^* &= \frac{100 - c_A + 2c_H}{3} \end{aligned}$$

The total price of a tour is thus:

$$p_H^* + p_A^* = \frac{200 + c_A + c_H}{3} \simeq 66.7 + \frac{c_A + c_H}{3}.$$

Unless $c_H + c_A > 100$ (in which case the cost of production is greater than consumers' maximum willingness to pay), the total price of a tour when the hotel and airline prices are set separately is greater than the total price when they are set jointly, $p_V^* = 50 + \frac{c_H + c_A}{2}$.

¹⁹This kind of best-response equilibrium is known in game theory as a Nash equilibrium - more on that in the game theory chapters.

When prices are set separately, the values of p_H^* and p_A^* above imply quantity:

$$\begin{aligned} & 100 - \left(\frac{200}{3} + \frac{1}{3}c_A + \frac{1}{3}c_H \right) \\ &= \frac{100}{3} - \frac{1}{3}c_A - \frac{1}{3}c_H \end{aligned}$$

and total profit:

$$\begin{aligned} & \left(\frac{100}{3} - \frac{1}{3}c_A - \frac{1}{3}c_H \right) \left(\frac{200}{3} + \frac{1}{3}c_A + \frac{1}{3}c_H - c_H - c_A \right) \\ & \frac{2}{9} (100 - c_H - c_A)^2 \end{aligned}$$

Finally, since $\frac{2}{9} < \frac{1}{4}$, the firm earns higher profit when it sets both prices jointly than when the prices are set independently by separate divisions.

The previous example shows that when the firm's divisions set prices separately, they set the total price too high relative to the prices that maximize joint profits. The firm as a whole would be better off lowering prices – the increased demand would more than make up for the decrease in price.

What is going on here? Begin with the case where the firm is charging p_V^* , and suppose for the sake of simplicity that the hotel price and airline price are equal. At p_V^* , the marginal revenue to the entire firm is equal to its marginal cost. That is,

$$MR_{Firm} = 100 - 2p_V^* = c_H + c_A.$$

Now think about the incentives for the hotel manager in this situation. If the hotel increases its price by a small amount, beginning from $\frac{p_V^*}{2}$, its marginal revenue is:

$$MR_{Hotel} = 100 - 2p_H^* - p_A^* = 100 - 2\frac{p_V^*}{2} - \frac{p_V^*}{2} > 100 - 2p_V^* = MR_{Firm}.$$

Thus the hotel's marginal gain in revenue due to raising its price is greater than the marginal gain in revenue to the entire firm. Why? Raising the price reduces the quantity demanded, but the increase in price all goes to the hotel, while the decrease in quantity is split between the hotel and airline. But, the hotel manager doesn't care about this latter effect on the airline. The same logic holds for the airline manager's choice of the airline price. Because of this, each division will charge a price that is too high.

What about if the goods were substitutes instead of complements? Think about a car company pricing its product line. If its lines are priced separately, each division head has an incentive to

lower the price and steal some business from the other divisions. Because all division heads have this incentive, the prices for the cars are lower when prices are set separately than they would be if the firm set all prices centrally. This is sometimes known as cannibalization. The firm must worry that by lowering the price on one line it is really just stealing business from the other product lines.

We will return to these types of issues when we study oligopoly. But, for now, we just want to motivate the idea that even the monopolist has to worry about issues of strategy.²⁰

9.5.2 Intertemporal Pricing

Consider the following scenario. A monopolist produces a single good that is sold in two consecutive periods, 1 and 2. Using the quantity-based approach again, let $p_1(q_1)$ be the inverse demand curve in the first period and $p_2(q_2, q_1)$ be the inverse demand in the second period. Note that this formulation indicates that the first-period demand does not depend on the second-period quantity. This implies that first-period consumers do not plan ahead in their purchase decisions.

Let $c_1(q_1)$ and $c_2(q_2)$ be the cost functions in the two periods, with a discount factor $\delta = \frac{1}{1+r}$. The monopolist maximizes:

$$p_1(q_1) \cdot q_1 - c_1(q_1) + \delta (p_2(q_2, q_1) \cdot q_2 - c_2(q_2))$$

The first-order conditions are:

$$\begin{aligned} p_1'(q_1) \cdot q_1 + p_1(q_1) + \delta \left(\frac{\partial p_2(q_2, q_1)}{\partial q_1} \cdot q_2 \right) &= c_1'(q_1) \\ \frac{\partial p_2(q_2, q_1)}{\partial q_2} \cdot q_2 + p_2(q_2, q_1) &= c_2'(q_2) \end{aligned}$$

Thus the monopolist sets marginal revenue equal to marginal cost in the second period. But, what about the first period? In this period, the monopolist must take into account the effect of the quantity sold in the first period on the quantity it will be able to sell in the second period. There are two possible ways this effect could go:

- Goodwill: $\frac{\partial p_2(q_2, q_1)}{\partial q_1} > 0$. Selling more in the first period generates “goodwill” that increases demand in the second period, perhaps through reputation effects or good word-of-mouth. In this case, the monopolist will produce more in the first period than it would if there were only one period in the model (or if demand in the two periods were independent).

²⁰Strategic issues are the main subject of the rest of these notes, starting with game theory, economists’ tool for studying strategic interactions.

- Fixed Pie: $\frac{\partial p_2(q_2, q_1)}{\partial q_1} < 0$. Selling more in the first period means that demand is lower in the second period. This would be true if the market for the product is fixed. In this case the additional sales in period 1 are cannibalized from period 2, leaving nobody to buy in period 2. In this case, the monopolist will want to sell less than the amount it would in the case where demand was independent in order to keep demand up in the second period.

Note - the stories told here are not quite rigorous enough. In Industrial Organization (IO) economics, people work on models to account for these things explicitly. Our object here is to illustrate that the firm's strategy in the first period will depend on what the firm is going to do in the second period.

9.5.3 Durable Goods Monopoly

Another way in which a monopolist can compete against itself is if it produces a product that is durable. For example, think about a company that produces refrigerators. Substitute products include not only refrigerators produced today by other firms, but refrigerators produced yesterday as well. Put another way, the refrigerator is competing not only against other refrigerator makers, it is also competing against past and future versions of its own refrigerator.

Consider the following simple model of a durable goods monopoly. A monopolist sells durable goods that last forever. Each unit of the good costs c dollars. There are N consumers numbered $1, 2, \dots, N$. Consumer n values the durable good at n dollars. That is, she will pay up to n dollars, but no more. Each day the monopolist quotes a price, p_t , and agrees to sell the good to whomever wants to purchase it at that price. How should the monopolist choose the path of prices? Assuming no discounting, the monopolist is willing to wait, so it should charge price N on day 1, $N - 1$ on day 2, etc. Each day, it skims off the highest value customers that remain. If there is a positive discount rate, then the optimal price path will balance the extra revenue gained by skimming against the cost of putting off the revenue of the lower valued customers into the future. Either way, the monopolist is able to garner almost all the social surplus as profit. Can you think of examples of this type of behavior? What about hard cover vs. soft cover books?

But, there is the problem with this model. If customers know that the price will be lower tomorrow, they will wait to buy. Of course, how willing they are to wait will depend on the length of the period. The longer they have to wait for the price to fall, the more likely they are to buy today. We can turn the previous result on its head by asking what will happen as the length of

the period gets very short. In this case, consumers will know that by waiting a very short time, they can get the product at an even lower price. Because of this they will tend to wait to buy. But, knowing that consumers will wait to buy, the monopolist will have an incentive to lower the price even faster. The limit of this argument is that the monopolist is driven to charge $p = MC$ immediately, when the period is very short.

So, we have seen that the durable goods monopolist will make zero profit when it has the opportunity to lower its price as fast as it wants to. This phenomenon is known as the Coase Conjecture, so called because Coase believed it but didn't prove it. It has since been proven. Notice that the monopolist's flexibility to charge different prices over time actually hurts the firm. It would be better off if it could commit to charging the declining price schedule we mentioned earlier. In fact, the monopolist would be better off if it could commit never to lower prices and just charging the monopoly price forever. In this case, the firm wouldn't sell any units in the second period, but it would also face no pressures to lower the price in the first period.

How can a monopolist commit to never lowering its prices? The market gives us many examples:

- Print the price on the package
- Some products never go on sale by reputation, such as Tumi luggage
- Third-party commitment, such as using a retailer who is contractually prohibited from putting items on sale
- "Destroy" the factory after producing in the first period, by limiting production runs (such as "limited edition" collectibles or the "retirement" of Beanie Babies)
- Money back guarantee – if the price is lowered in the future, the monopolist will refund the difference to all purchasers.
- Planned obsolescence – if the goods aren't that durable, there isn't a problem. In the absence of an ability to commit to keeping prices high, firms should produce goods that aren't particularly durable.
- Leasing goods instead of selling them

This same phenomenon happens even in a two-period model. The monopolist would like to charge a high price in both period. But after the first period has passed it has an incentive to

lower the price in order to increase sales. But the customers, knowing that the monopolist will lower prices in the second period, will not buy in the first period. Thus, the monopolist will have to charge a lower price in the first period as well. The monopolist would be better off if it could somehow commit to keeping prices high in both periods in this model.

Consider the following scenario. Instead of selling the durable good, the monopolist will lease it for a year. In this case, units of the good rented in period 2 are no longer substitutes for units of the good rented in period 1. In fact, somebody who likes the good will rent a unit in each year.

What should the monopolist do in this case? The monopolist should charge the monopoly rental price (high price) in each period to rent the good. Thus if the monopolist can commit to rental instead of selling, it can earn higher profits. Renting is similar to the other tools listed above, in that it is an example of the firm's intentionally restricting its own flexibility. Thus the monopolist, by eliminating its opportunity to cut prices on sales later, is able to do better. In this context, choice is not always a good thing.

Formal Model of Renting vs. Selling²¹ There are two periods, and goods produced in period 1 may be used in period 2 as well (i.e. it is durable). After period 2 the good becomes obsolete. Assume that the cost of production is zero to make things simple, and the monopolist and consumers have discount factor δ . Demand in each period is given by $q(p) = 1 - p$. The monopolist can either lease the good for each period or sell it for both periods. If the monopolist decides to sell the good, then consumers who purchased in the first period can resell it during the second period.

Suppose the monopolist decides to lease. The optimal price for the monopolist to charge in each period is $1/2$. This yields quantity $1/2$ in the first period, which can be leased again in period 2. No additional quantity is produced in period 2. Thus discounted profits are given by

$$\pi_{lease} = \frac{1}{4}(1 + \delta)$$

Suppose the monopolist decides to sell. In this case, the quantity offered for sale in period 1 is reoffered by the resale market in period 2. Thus the residual demand in period 2 is given by $p_2 = 1 - q_2 - q_1$. Thus in period 2 the monopolist chooses q_2 to solve:

$$\max q_2(1 - q_1 - q_2)$$

which implies that the monopolist sells $q_2 = \frac{1-q_1}{2}$ and earns profit $\left(\frac{1-q_1}{2}\right)^2$.

²¹From Tirole, pp. 81-84.

Now, what should the monopolist do in period 1? The price that consumers are willing to pay is given by the willingness to pay in period 1 plus the discounted price in period 2, since the consumer can always lease the object (or resell it) in period 2 for the market price. Thus consumers are willing to pay

$$(1 - q_1) + \delta p_2^a$$

where p_2^a is their belief about what the price will be in period 2.

We suppose that consumers correctly anticipate the price in the second period.²² That is, $p_2^a = p_2 = \frac{1-q_1}{2}$. Thus as a function of q_1 the maximum price the monopolist can charge is given by:

$$p_1 = (1 - q_1) + \delta \frac{1 - q_1}{2} = (1 - q_1) \left(1 + \frac{\delta}{2}\right)$$

Thus in the first period the monopolist chooses q_1 to maximize:

$$\pi_{sales} = \pi_1 + \delta \pi_2 = (1 - q_1) \left(1 + \frac{\delta}{2}\right) q_1 + \delta \left(\frac{1 - q_1}{2}\right)^2$$

The first-order condition yields:

$$\begin{aligned} \frac{d}{dq_1} \left((1 - q_1) \left(1 + \frac{\delta}{2}\right) q_1 + \delta \left(\frac{1 - q_1}{2}\right)^2 \right) &= 0 \\ (1 - q_1) \left(1 + \frac{\delta}{2}\right) - q_1 \left(1 + \frac{\delta}{2}\right) - \delta \left(\frac{1 - q_1}{2}\right) &= 0 \end{aligned}$$

$$\begin{aligned} q_1^* &= \frac{2}{4 + \delta} < \frac{1}{2} \\ q_2^* &= \frac{1 - q_1^*}{2} = \frac{1 - \frac{2}{4 + \delta}}{2} = \frac{1}{2} \cdot \frac{2 + \delta}{4 + \delta} \\ q_1^* + q_2^* &= \frac{2}{4 + \delta} + \frac{1}{2} \cdot \frac{2 + \delta}{4 + \delta} = \frac{1}{2} \cdot \frac{6 + \delta}{4 + \delta} > \frac{1}{2} \end{aligned}$$

In terms of prices,

$$\begin{aligned} p_2^* &= 1 - (q_1^* + q_2^*) \\ p_2^* &= \left(1 - \frac{1}{2} \cdot \frac{6 + \delta}{4 + \delta}\right) = \frac{1}{2} \cdot \frac{2 + \delta}{4 + \delta} < \frac{1}{2} \\ p_1^* &= (1 - q_1^*) \left(1 + \frac{\delta}{2}\right) = \left(1 - \frac{2}{4 + \delta}\right) \left(1 + \frac{\delta}{2}\right) = \frac{1}{2} \frac{(2 + \delta)^2}{4 + \delta} < \frac{1 + \delta}{2} \end{aligned}$$

²²In game theory terms, this is a Perfect Bayesian equilibrium.

Note that $\frac{1+\delta}{2}$ is the monopoly price if the second period is ignored ($\delta = 0$). Going back to the expression for total profit:

$$\begin{aligned}
 \pi_{sales} &= (1 - q_1) \left(1 + \frac{\delta}{2}\right) q_1 + \delta \left(\frac{1 - q_1}{2}\right)^2 \\
 &= \frac{1}{2} \frac{(2 + \delta)^2}{4 + \delta} \cdot \frac{2}{4 + \delta} + \delta \left(\frac{1}{2} \cdot \frac{2 + \delta}{4 + \delta}\right)^2 \\
 &= \left(\frac{2 + \delta}{4 + \delta}\right)^2 + \delta \left(\frac{1}{2} \cdot \frac{2 + \delta}{4 + \delta}\right)^2 \\
 &= \left(1 + \frac{\delta}{4}\right) \left(\frac{2 + \delta}{4 + \delta}\right)^2
 \end{aligned}$$

Comparing the sales profit, $(1 + \frac{\delta}{4}) \left(\frac{2+\delta}{4+\delta}\right)^2$, to the leasing profit, $\frac{1}{4}(1 + \delta)$, we can see that $\pi_{lease} = \pi_{sales}$ only when $\delta = 0$. It turns out that for all $\delta > 0$ (meaning the second period matters, at least somewhat), then $\pi_{lease} > \pi_{sales}$. Leasing is more profitable. Why? When selling the good, the monopolist cannot resist the temptation to lower the price in the second period and sell more; thus, it cannot sell as much in the first period ($q_1^* < \frac{1}{2}$) as would be optimal, and this reduces overall profits. The monopolist would be better off if it could commit to leasing rather than selling.