Opportunity cost—The cost foregone by not choosing the next best alternative. For a firm, it is the profit foregone by not producing in the next most profitable industry.

Own-price elasticity—The effect of an increase in the price of a good on the quantity of it demanded.

Production possibilities frontier—Shows the maximum amount of one good that can be produced, given the amount produced of the other good. On the frontier, resources and technology are used to their maximum efficiency.

Rate of product transformation—The rate at which the resources used in the production of one good can be shifted to the production of another good while maintaining the greatest output efficiency possible. It is equal to the slope of the production possibilities frontier.

Second-order condition—Necessary to ensure that we are obtaining a maximum value and not a minimum value.

REFERENCE

Howarth, R. 2005. The present value criterion and environmental taxation: The suboptimality of first-best decision rules. *Land Economics* 81, 321–336.

5

MARKET FAILURE AND THE SECOND FUNDAMENTAL THEOREM OF WELFARE ECONOMICS

What is it we mean by "market failure"? Typically, at least in allocation theory, we mean the failure of a more or less idealized system of price-market institutions to sustain "desirable" activities or stop "undesirable" activities. The desirability of an activity, in turn, is evaluated relative to the solution values of some explicit or implied maximum welfare problem.

—*Francis Bator, "The Anatomy of Market Failure,"* Quarterly Journal of Economics *72(3) (1958), 351*

The model of perfect competition presented in the last chapter is the heart and soul of neoclassical welfare economics. It is a mathematical representation of an economic system that exactly duplicates the operation of a voluntary, frictionless, barter economy. As we saw in Chapter 1, the starting point for the model is the individual consumer whose behavior is described by the rational actor model. The consumer chooses from among the array of market goods based solely on the prices of those goods, and producers supply these goods using the most efficient combinations of inputs as indicated by their prices (see Figure 5.1). Two critical assumptions drive the demand for market goods in this model. The first is the assumption of rational choice on the part of consumers and producers, and the second is faith in the ability of prices to correctly capture all the relevant information about market goods. If the price signals are "wrong," then so too will be the collection of market goods chosen by the rational consumer.

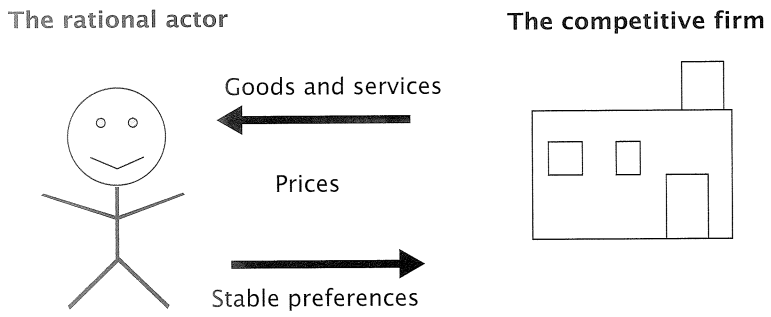

More is always preferred to less
Consistency in choices
Preferences are stable
Responds rationality to price signals
Every preference is equally valid
No interaction

Perfect information
Mobility of factors
No barriers to entry
Uses best technology
Constant returns to scale
No interaction

Figure 5.1. Assumptions of the Walrasian model

Much of the criticism of neoclassical economics is centered on the First Fundamental Theorem, namely, a competitive economy will ensure the most efficient allocation of society's resources in consumption and production. It is easy to see that real economies bear little resemblance to the competitive ideal. But this is recognized by neoclassical economics. Sometimes prices are distorted and the market fails to achieve the socially optimal result. Instances of **market failure** include **externalities, public goods,** and **market power.** In all these cases, prices are distorted and there is a legitimate role for the government to intervene and correct the market failures.

EXTERNALITIES

An externality occurs when the action of one economic agent (consumer or firm) affects another agent, and this effect is not included in market prices. When this occurs, the private market price is not equal to the **social price** of the good.

In Figure 5.2, a factory producing some good (perhaps widgets) is emitting air pollution that affects nearby residents. This air pollution might cause health

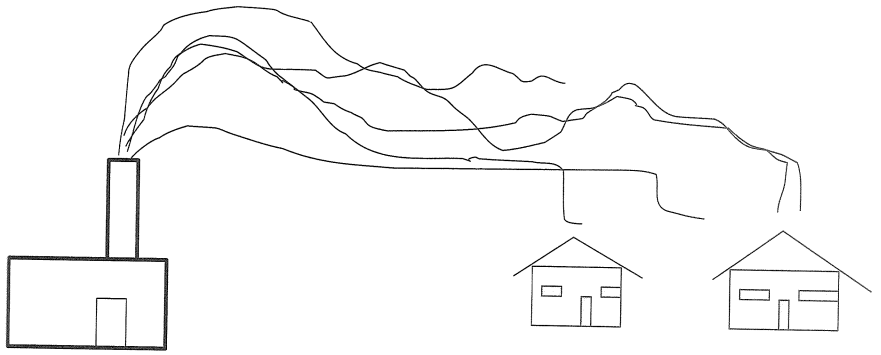

Figure 5.2. An externality in production

damages, extra cleaning costs, or reduced visibility. If these damages are not included in production costs, the "true" (social) cost of producing a widget would be higher than the cost to the (private) producer and thus higher than the market price. This is the case even if all the assumptions of Walrasian theory hold, including perfect competition, the rational actor model, and so on. Figure 5.3 shows the effect of a negative externality on the price and quantity of a good.

If the external costs of production were not accounted for, the (private) price would be set at P_P and the quantity produced would be Q_P. Including the external effects of production would increase the price to P_S (the social price) and decrease the quantity demanded to Q_S. When externalities are present, consumers respond rationally to prices, and prices are capable of containing all relevant information, but the wrong prices are sent to the market. Once these prices are corrected, a socially efficient allocation of resources can be obtained.

Two classic solutions to externalities are (1) the use of taxes and subsidies and (2) completely assigning property rights. Arthur Pigou, writing in the 1920s, argued for a tax on negative externalities such as pollution and subsidies for positive externalities such as beekeeping (which has a positive effect on the production of many crops). Ronald Coase argued that the socially optimal amount of an externality would automatically be arrived at if property rights were assigned to either the polluter or pollutee. In the first case the pollutee would pay the polluter not to pollute, and in the second case the

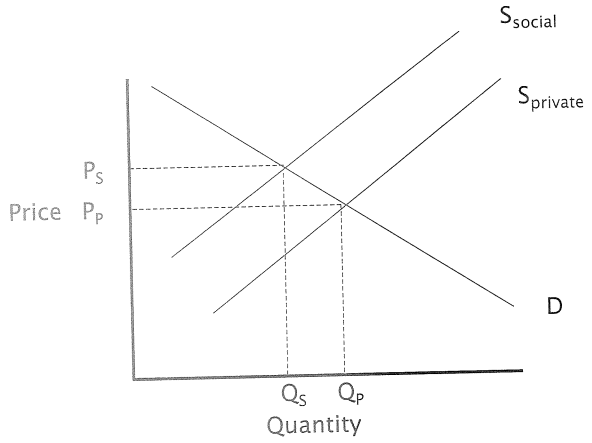

Figure 5.3. Price and quantity of a good associated with a negative externality

polluter would pay for the right to pollute. No matter to which party property rights are assigned, with perfect information and no transactions costs, the situation would be reached where the marginal cost of pollution would exactly equal the marginal benefit from the polluting activity.

PUBLIC GOODS

Public goods have two characteristics that preclude assigning them a proper price so they can be efficiently allocated.

1. Public goods are non-exclusive. Once a public good is provided, anyone can use it whether or not that person paid for its use.
2. Public goods are non-rival. Once a public good is provided, one person's use does not affect another person's use of that good.

An example of a pure public good is public radio. Once a public radio station is operating, anyone can listen to that station whether or not that person sends in a donation to the station. Also, any number of people can listen to the station without affecting any other person's use. The marginal cost of adding one more person is zero.

Many goods have one of these characteristics and not the other. For example, cable TV is exclusive but not rival. People can be excluded from using it if they do not pay a cable fee, but up to a point more customers can be added without affecting the supply of cable TV services. Other goods are rival but not exclusive. The fishery is a classic example of a good of this sort. If a fishing ground is not regulated by some government or private cooperative, anyone can go there and fish. But, because the number of fish is limited in any given area, if one person catches a fish no one else can catch that fish. The private incentive is to catch the fish before someone else does. This is called the *tragedy of the commons.*

The Tragedy of the Commons

In the case of a finite, open-access resource there is an incentive for each individual to use the resource as quickly as possible before someone else uses it first. One solution to the tragedy is to assign property rights, either to private individuals or to a public entity.

We saw in Chapter 4 that a private good is efficiently provided by the market if its price is equal to its marginal cost. Price is determined by the intersection of demand and supply in the market for that good. In the case of a public good, efficient allocation results when the marginal social value of the good is equal to the marginal cost of providing it. The social benefit of a public good is determined by vertically summing the demand curves for the consumers in the market for that good, as shown in Figure 5.4.

Figure 5.4 shows the usual textbook solution to the public goods problem. The efficient provision of the public good is determined by the point where the public good demand curve (vertical sum of D1 + D2) equals the marginal cost of producing the good. So the efficient price would be P_{pg} and the optimal amount provided would be Q_{pg}. This solution is a little disingenuous because for a pure public good, once it exists, the marginal cost of providing it to one more user is zero. In that case, the price would be zero and the optimal amount provided would be Q_{ppg}.

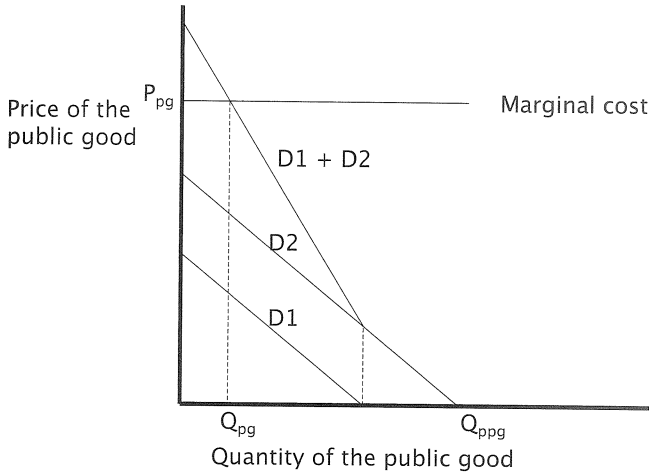

Figure 5.4. The optimal provision of a public good

MARKET POWER

A fundamental characteristic of perfect competition is that firms are price takers; each firm is so small relative to the market that its actions have no effect on price. The price a firm can receive for a good is dictated by the forces of supply and demand in the market for that good. In economic jargon, the firm faces a horizontal, perfectly elastic demand curve. For a monopolist, however, the firm controls the entire market, and so it faces a downward sloping demand curve.

As shown in Figure 5.5, the monopolist faces a downward sloping demand curve, and a marginal cost curve lying under the demand curve. In the case of a linear demand curve, the demand curve (in inverse form) is $P=a+2Q$ (where a is a constant) and the slope would be $dP/dQ=2$. Total revenue is the quantity sold times price or $PQ=(a+2Q)Q=aQ+2Q^2$, and marginal revenue is slope of the TR curve $d(PQ)/dQ$ or $(a+4Q)$, so the intercepts of the demand and marginal revenue curves are the same but the slope of the marginal revenue curve ($dMR/dQ=4$) is twice as great. As in the case of perfect competition, the firm will produce where $MC=MR$, and in this case the output will be Q_M. The price is taken from the demand curve and the quantity Q_M is

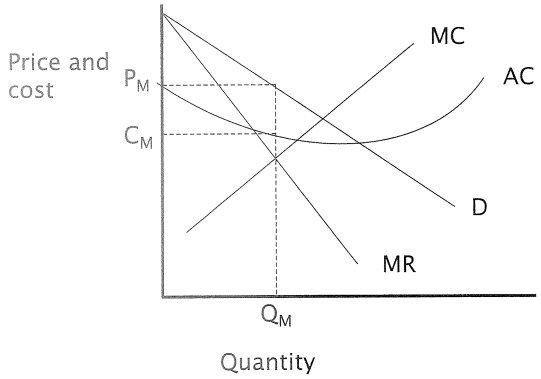

Figure 5.5. Price, output, and average cost under monopoly

associated with a price of P_M. At that quantity the average unit cost of production is C_M, which is less than the price, so the firm makes a profit of $(P_M - C_M)$ Q_M. Unlike the competitive firm, a firm with market power can earn an economic profit, that is, a rate of profit greater than the average rate prevailing in the economy. The monopolistic firm produces at a point other than the minimum on the long-run average cost curve and can control price by changing the quantity sold.

The efficiency of monopoly compared with perfect competition can be seen in Figure 5.6. If the market for the good being produced were competitive, the amount produced would be Qc, determined by the point where the marginal cost curve intersects with the demand curve (the demand curve would represent the sum of the marginal revenue curves for all the producers). The price would be Pc (price = marginal cost = marginal revenue).

In the case of a monopolistic firm, output would be determined where MC = MR, the quantity produced would be Qm, and the price would be Pm. A monopolist produces less and charges more than a firm operating under conditions of perfect competition. In the Walrasian general equilibrium framework, this represents a loss to society because less of this good is being produced than is socially optimal.

The effect of this can be seen using the concept of **consumer surplus**. Under competitive conditions, the price of the good in Figure 5.6 would be Pc.

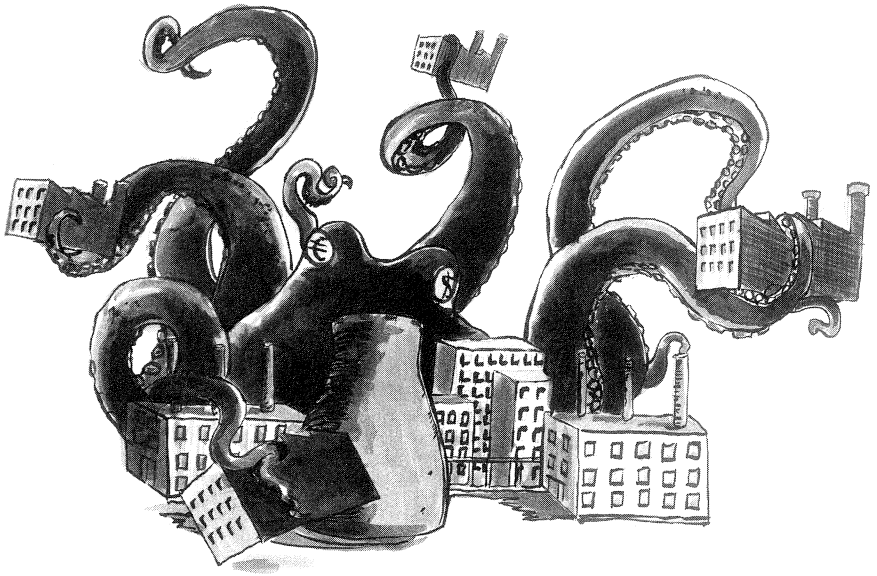

Looking at the demand curve, if the price were higher than Pc, would there still be some demand for the good? The answer is clearly yes. Some people are willing to pay a higher price than Pc for the good. So they are in a sense receiving a "bonus," or surplus, compared with what they would be willing to pay. Total consumer surplus under competitive conditions in Figure 5.6 is the area of the triangle acPc.

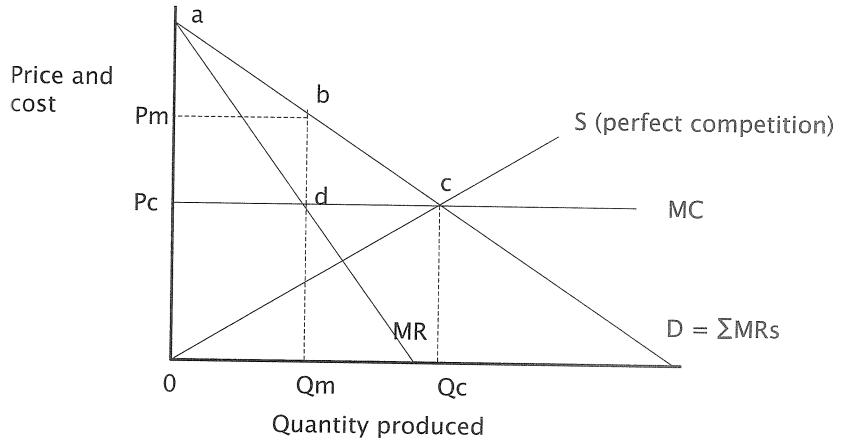

Figure 5.6. The loss of social welfare under monopoly

Looking at Figure 5.6 from the point of view of the producer, we can see that some sellers would be willing to sell the good at a price less than the competitive price of Pc. They too are getting a bonus, and this is called **producer surplus**. In Figure 5.6, under competitive conditions, producer surplus would be the area under the competitive supply curve cOPc.

Suppose this industry were characterized by monopoly power? In that case, the firm would set the price at Pm and consumer surplus would be reduced to the area abPm. Compared with the competitive situation, this represents a loss of consumer surplus equivalent to the area (acPc−abPm) or PcPmbc. But part of this loss goes to the producer as producer surplus (the area PcPmbd). This does not represent a loss to society; it is simply a transfer from consumers to producers. But the area bcd totally disappears. This is called a **deadweight loss** and represents the cost to society of the monopolization of this industry. The rationale for this is that society would prefer that more scarce resources be used to make this good but instead they are going to their next best use (whatever that is).

PRICES AND MARKET FAILURE

In all three classic cases of market failure, the economy fails to achieve general Pareto efficiency because the wrong prices are sent to consumers. When externalities are present, prices do not reflect the "external" effects of production. Because of the free-riding problem, the prices of public goods underestimate their true value to society. Monopolists can control the prices they charge in order to influence the quantity sold. All of these instances of market failure are ubiquitous in real economies. This fact is recognized by most economists. Most economists are not free-market fundamentalists, and they recognize a legitimate role for government intervention to correct market failures. And within the Walrasian system, there is a theoretically tractable way of doing this.

We saw in Chapters 1 (consumption) and 2 (production) that any particular Pareto-efficient outcome depends on the initial distribution of society's resources (goods or productive inputs). So if market failure results in an unsatisfactory outcome from society's point of view, all we have to do is change the initial distribution of resources (a lump-sum transfer) in order to move

to the socially desirable outcome. This is perhaps the most important idea for Walrasian economic policy. It is called the Second Fundamental Theorem of Welfare Economics, and it provides the intellectual rationale for policies promoting efficient markets.

The Second Fundamental Theorem of Welfare Economics

Assuming that consumers and producers are rational, self-regarding price takers, almost any Pareto-efficient outcome can be supported by lump-sum transfers.

Broadly interpreted, the Second Fundamental Theorem says that if market imperfections exist, it is possible to correct them through enlightened intervention. It is usually interpreted more narrowly to mean that if an imperfection exists, it is better to address the problem through income transfers rather than by adjusting prices. That is, the best policy is to adjust initial endowments and let the market take over.

For completeness we should also mention the Third Fundamental Theorem of Welfare Economics, the Arrow impossibility theorem.

The Third Fundamental Theorem of Welfare Economics

This is simply the Arrow impossibility theorem. There is no social welfare function that satisfies Arrow's criteria of universality, non-dictatorship, Pareto efficiency, and independence.

The power of neoclassical welfare economics comes directly from the First and Second Welfare Theorems. These theorems are nothing short of the intellectual foundation for arguments for the efficacy of free-market economies in achieving the highest possible social welfare. Of course, if these theorems are shown to be invalid it does not automatically throw out the case for market allocation of society's scarce resources. But it does call into question market efficiency as a universal first principle of resource allocation. If these theorems

do not hold, the case for free markets must be based on specific evidence and the specific conditions of particular situations. This is, in fact, the current trend in economic research and policy. There is a growing recognition of the theoretical intractabilities and the behavioral limitations of Walrasian welfare economics. This is the subject of Part Two of this book.

APPENDIX

Income and Substitution Effects

A fundamental principle of economics is the **law of demand**. All other things being equal, the price of something and the quantity demanded of it are inversely related, or in mathematical notation, $dQ/dP<0$. There are two reasons for this relationship. First, if the price of a good increases, then the consumer's **real income** (income adjusted for inflation) will decline; and if the good is a **normal good** ($dQ/dI>0$), then fewer units of that good will be demanded. This is called the **income effect**. Second, an increase in the price of one good (X in this case) relative to another ($[Px/Py]\uparrow$) will cause consumers to substitute good X for good Y. This is the **substitution effect**.

Accounting for the loss of real income affects the demand curve. The ordinary, or **Marshallian demand curve**, named for Alfred Marshall, shows both the income and substitution effects. The **Hicksian demand curve**, named from John Hicks, is *income compensated*, that is, the consumer is compensated for increases or decreases in real income as prices change. In Figure 5.7 notice that the Hicksian demand curve is steeper, less elastic (less price-responsive) compared with the ordinary market demand curve. This is because part of the response to a price change has been taken away.

It is straightforward to show how these effects can be separated using graphs or mathematics. Figure 5.8 is a graphical representation of the income and substitution effects. Suppose there is a price increase in good Y from P_Y' to P_Y. This shifts the budget line inward because at the higher price the consumer can buy less of good Y with his fixed budget M. The optimal amount of good Y purchased at price P_Y' is Y3, where the indifference curve I2 is just tangent to the original budget line at point c. The optimal amount of good Y purchased at price P_Y is Y1, where indifference curve I1 is just tangent to the new budget line at point a.

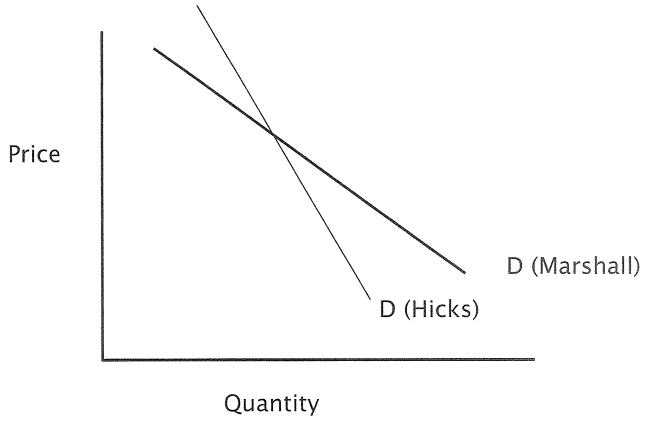

Figure 5.7. Demand according to Hicks and Marshall

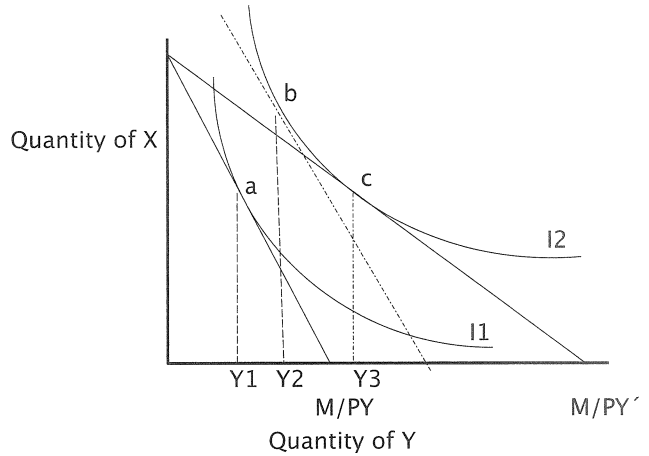

Figure 5.8. Income and substitution effects for a price increase in a normal good

We can separate the income and substitution effects by giving the consumer enough income so that the consumer can be back on the original indifference curve I2. This is shown by the dotted budget line that is just tangent to the indifference curve I2 at point b. So we have compensated the consumer for a loss of real income due to the price increase in good Y by giving the consumer

enough money to move back to the original indifference curve. When the price of Y went up, the total effect was to reduce the amount of Y purchased from Y3 to Y1 (the movement from c to a in Figure 5.8). The movement from Y3 to Y2 (c to b along the original indifference curve) is the substitution effect. The movement from Y2 to Y1 (b to a) is the income effect.

The total effect of a price change is always negative (the law of demand), $P_Y\uparrow Q_Y\downarrow$. The substitution effect is also always negative, $(P_Y/P_X)\uparrow Q_Y\downarrow$. As the price of good Y increases relative to the price of good X, a consumer will substitute good X for good Y. In the case of a normal good the income effect is also negative, $P\uparrow M\downarrow Q_Y\downarrow$. For an **inferior good,** however, the income effect is positive, $P\uparrow M\downarrow Q_Y\uparrow$, partially (but not entirely) offsetting the negative substitution effect.

Mathematically, this result can be derived using a variety of approaches. An interesting interpretation of the income and substitution effects is given by the *Slutsky equation*, named for the early twentieth-century mathematician Eugen Slutsky. Stated in words, the Slutsky equation shows that the slope of the ordinary (Marshallian) demand curve for some good X is equal to the slope of the income-compensated (Hicksian) demand curve minus the optimal amount of good X purchased times the slope of the **Engel curve**. The Engel curve shows the change in the amount of a good (X) purchased resulting from a change in income (dX/dM) keeping prices constant. In elasticity form the Slutsky equation is:

$$(5.1) \qquad E_{11} = E_{11\,(\text{utility held constant})} - \alpha_1 E_{1M}$$

The own-price elasticity (E_{11}) is equal to the substitution effect (the income-compensated own-price elasticity) minus the income elasticity (E_{1M}) times the budget share of the good ($\alpha_1 = P_1Q_1/M$).

The relationships in equation (5.1) implies that the Hicksian and Marshallian measures should be about the same either if the budget share of the good is small or if demand for the good is relatively insensitive to income changes. This is an important point because if the term $\alpha_1 E_{1M}$ is near zero, we can use the market demand curve (and market data directly) for estimates of welfare changes. The three measures of consumer surplus should be about the same. How much

people are willing to pay for a gain (WTP) should be about the same as how much they would have to be compensated to accept a similar loss (WTA).

Consumer Surplus According to Hicks and Marshall

The two different kinds of demand curves, ordinary and income-compensated, yield different measures of consumer surplus. Figure 5.9 shows the effect on consumer surplus of a price change for a good. The curve labeled D_M is the Marshallian (ordinary or market) demand curve. Consumer surplus is the area A+B under this demand curve corresponding to the two different prices. In this case the effect on consumer surplus of a price increase, −(A+B), or price decrease+(A+B) is the same.

In the case of the Hicksian demand curve we get different measures of consumer surplus depending on the direction of the price change. This is because, at a higher price (P1 in Figure 5.9), the consumer has a lower real income, and at a lower price (P2), the consumer has a higher real income. The relevant Hicksian demand curve for price P1 is H1 (lower real income), and

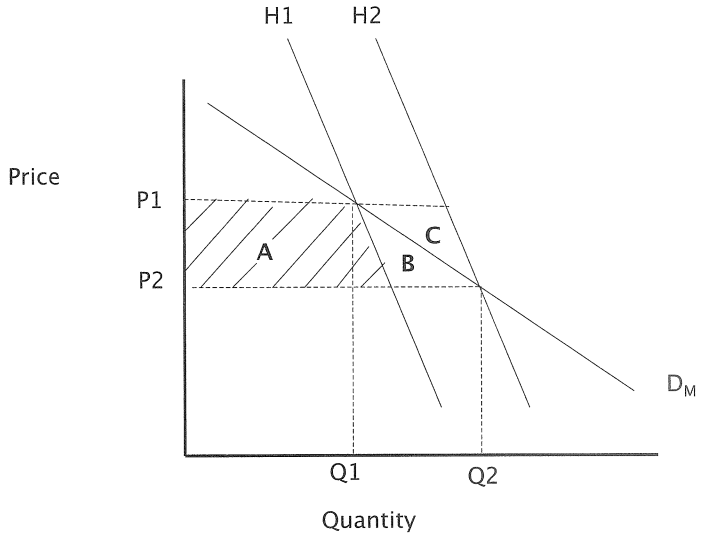

Figure 5.9. Consumer surplus according to Hicks and Marshall

the Hicksian demand curve for price P2 is H2 (higher real income). The gain in consumer surplus resulting from a price move from P1 to P2 is the area under H1, or A. This is called the **compensating variation**. The loss in consumer surplus resulting from a price increase from P2 to P1 is the area under the Hicksian demand curve H2, or A+B+C. This is called the **equivalent variation**.

A = compensating variation

A+B = consumer surplus

A+B+C = equivalent variation

Referring to equation (5.1), the differences between these three measures disappear as the expression $(-\alpha_1 E_{1M})$ becomes smaller.

This discussion of consumer surplus may seem esoteric but it is critically important to **cost-benefit analysis**, one of the basic tools of economics. Traditional cost-benefit analysis relies on Marshallian measures of consumer surplus and they assume that WTA and WTP are similar. As Chapter 8 shows, empirical studies reveal very large differences between the two measures. People are loss averse; they are much less willing to accept losses than they are willing to pay for equivalent gains. This has important consequences for cost-benefit analysis and public policy.

GLOSSARY

Compensating variation—A measure of the welfare gain resulting from a decrease in the price of a good or service. It is equal to the loss in income that would exactly offset the welfare gain from the fall in price.

Consumer surplus—The difference between what consumers actually pay for a good and the extra amount they would be willing to pay.

Cost-benefit analysis—Calculating the benefits and costs of public policies based on identifying potential Pareto improvements. Also called benefit-cost analysis.

Deadweight loss—The net loss of consumer surplus plus producer surplus.

Engel curve—The change in the amount of a good (X) purchased resulting from a change in income (dX/dM), keeping prices constant.

Equivalent variation—A measure of the welfare loss resulting from an increase in the price of a good or service. It is equal to the gain in income that would exactly offset the welfare loss from the price increase.

Externalities—An effect from production or consumption that is not taken account of by market prices. An externality may be either positive or negative.

Hicksian demand curve—A demand curve that includes only changes in relative prices of the good. Also called an income-compensated demand curve.

Income effect—The effect on consumption of a gain or loss in real income.

Inferior good—A good for which an increase in income decreases the demand for it.

Law of demand—The quantity demanded of a good varies inversely with changes in its price.

Market failure—The failure of the market to send "correct" price signals to consumers and firms. Instances include externalities, public goods, and market power.

Market power—The ability of a firm to affect prices and output in a particular market.

Marshallian demand curve—A demand curve that includes the effects of both changes in real income and changes in price. Also called an ordinary demand curve.

Normal good—A good for which an increase in income increases the demand for it.

Producer surplus—The difference between the price sellers actually get for a product and the (lower) price they would be willing to accept.

Public goods—Goods with the characteristics of being non-rival and non-exclusive.

Social price—The price of a good that includes all positive and negative external effects of producing or consuming it.

Substitution effect—The change in the demand for a good arising solely from a change in relative prices, holding utility (usually measured by real income) constant.

PART TWO

MODERN WELFARE ECONOMICS

INTRODUCTION AND OVERVIEW

Part One of this book presents the core framework of Walrasian welfare economics. Its two key ideas are (1) the system was formalized as a theory of exchange in a pure barter economy—prices are added on more or less as an afterthought, and (2) the key assumption that holds the system together is the independent rational actor. Part One is a presentation of standard economic theory, a theory than has been around more or less in its present form for fifty years or more. There is no need for extensive footnotes and references in Part One because the material is presented in dozens of microeconomic textbooks, although not as a logical progression from barter to prices to the fundamental theorems. The material presented in Part One is consistent with other contemporary microeconomic texts, although the behavioral aspects of microeconomic theory are given more emphasis to prepare for the discussion in Part Two.