Chapter 2
A Brief Review of Microeconomic Theory

This chapter reviews microeconomics theory. As the suggested guide for teachers in the Introduction to this manual has indicated, not every class using this text will find that they need to cover this material. Some instructors tell us that even if their students do not know microeconomic theory, they do not take time early in the semester to review this chapter. Rather, they tell the students to read through the chapter on their own.

No one should expect students who have never studied economics to emerge from their reading of Chapter 2 a full-blown economist. Perhaps the better message to give to students who have not heretofore studied economics is to read through the chapter solely to familiarize themselves with where in this chapter they can later find treatment of various economic topics. For example, a student may not want to labor over the material on game theory now. Rather, she may want to know where this treatment is in Chapter 2 so that when she meets game theory in Chapters 4 and 8 she can return to this point in Chapter 2 for a refresher.

Other teachers tell us that they spend up to three weeks of a fifteen-week semester reviewing microeconomic theory before they get to the legal material.

Still others take three or four lectures at the beginning of the semester or term to go over selected parts of Chapter 2. For example, in teaching law students (some of whom were economics majors as undergraduates, but many of whom have never studied economics) one might focus on the sections in Chapter 2 on market failure and decisionmaking under uncertainty.

Our experience has been that because a great deal of the economic analysis of law consists of demonstrating that legal commands and institutions can correct for such market imperfections as external costs and benefits, public goods, information asymmetries, and collective action and group coordination problems, students need to be thoroughly familiar with those concepts before launching into Chapter 4 and the subsequent material.

Ulen frequently uses the first week of the semester to stress only limited aspects of micro theory for the law students he teaches. For example, he discusses welfare economics much more extensively than is the case in the text, stressing the sources of market failure and their correctives, spending time on the Arrow Impossibility Theorem, and introducing some intriguing uses of economics to discuss important public policy issues, which we describe in more detail below and in the *PowerPoint* presentations that are available on the Instructor's Resource website.

■ Notions of Efficiency

Be certain to distinguish between Pareto and Kaldor-Hicks efficiency. Pareto efficiency requires consent—that is, the gainers from a reallocation must receive the explicit consent of the losers. Presumably, they will only be able to do so when their gains are greater than the total of all losses. (In Chapter 4 we refer to this difference as "cooperative surplus.") In contrast, Kaldor-Hicks efficiency is a species of cost-benefit analysis in which a change is deemed efficient if the total gains exceed the total losses but without the requirement that the gainers compensate the losers.

There are several reasons for taking care with this distinction between two very different definitions of efficiency. First, many economics students will not be aware of the difference. Modern microeconomic theory uses the notion of Pareto optimality almost exclusively, so that even those with some familiarity with microeconomics may not be aware of Kaldor-Hicks efficiency. Law and economics uses Pareto optimality as its efficiency criterion in relatively few instances, most having to do with contractual matters. Instead, the literature uses, implicitly or explicitly, Kaldor-Hicks efficiency as its central efficiency concept.

Second, the difference between Pareto and Kaldor-Hicks efficiency has very much to do with transaction costs—a central notion in law and economics that we develop in Chapter 4. (Transaction costs are, in brief, the costs of effectuating a bargain.) The connection between the different efficiency norms and transaction costs is this: when transaction costs are so high that they make consensual, mutually advantageous bargaining unlikely or impossible, there must be some criterion other than consent for deciding on the efficient allocation of resources. The transaction costs of compensation—identifying the gainers and the losers and transferring resources from the gainers to the losers—may be greater than the difference between the benefits and the costs. In those circumstances, a bargain cannot take place, even though (leaving transaction costs to one side) the total gains from a reallocation exceed the total costs. Kaldor-Hicks efficiency allows us to speak about those reallocations in the absence of unanimous consent.

■ Optimality

We stress the point—unremarkable for economists, but striking for non-economists—that the optimal amount of anything occurs when social marginal cost and social marginal benefit are equal. An important implication of this point is that the optimal amount of many bad things is not zero. That is because it costs something to get rid of bad things. Much as we might like to eradicate all air and water pollution, the costs of doing so eventually far outweigh the benefits.

An example of an important law that flies in the face of this point is the Delaney Clause of the Food and Drug Act. That clause instructs the Food and Drug Administration to prohibit all food additives that pose any risk of cancer. This is a good matter for class discussion. There are, of course, lots of other examples. We would be extremely grateful to those of you who send us your examples.

■ Opportunity Cost

This is a central notion in microeconomic theory, and we commend the boxed example that is included in Chapter 2. However, we also highly recommend that you look at the article by Paul J. Ferraro and Laura O. Taylor, "Do Economists Recognize an Opportunity Cost When They See One?: A Dismal Performance from the Dismal Science," 4 *Contributions to Econ. Analy. & Pol'y.* 1 (2005). That article's findings about the inabilities of professional economists to compute opportunity cost are contained on the *Microeconomic Theory II* set of slides.

■ Further Examples

A marvelous example of an external benefit is the automobile theft prevention system known as LoJack.[1] Here are the stylized facts about that system. The LoJack company will sell its system to automobile owners for approximately $500. If an owner pays that fee, representatives of the company insert a secret device in the owner's car. Not even the owner knows where or what the device is. If the car is stolen, the owner calls the LoJack company, and they then turn on their equipment, which picks up a signal transmitted from the device they put into the car, helping the police to locate the car quickly. Because of the large capital expenditure involved for the company, LoJack is currently available only in large cities in the United States.

See Ian Ayres and Steven D. Levitt, "Measuring the Impact of Unobservable Victim Precaution: An Empirical Analysis of LoJack," 113 *Q.J. Econ.* 43 (1998), and then the brief summary of that article on the *Microeconomic Theory II* slides.

An interesting sidelight of the LoJack is how thieves and police have responded to it. Thieves recognize now that any car they steal might be LoJack-equipped and that if it is, they will be caught quickly. So, bright auto thieves in major cities where LoJack is available drive a stolen car to a safe place, park the car, and watch it to see if anyone comes to get it. The police now recognize that this is what thieves do; so, when the LoJack company notifies them that one of their protected cars has been stolen and is parked at a particular place, the police then go to watch that car, too.

How does the LoJack create an external benefit? Contrast the LoJack with a car alarm. Typically there is an external indication that the car is equipped with an alarm. Therefore, a potential thief can look at a car equipped with an auto alarm, recognize that that car is not a viable acquisition, and turn his attention to less well-defended cars. As a result, a car equipped does not confer a benefit on other cars. Indeed, it may make their theft more likely.

But a LoJack does not have this problem. Because one cannot tell whether a car is equipped with a LoJack, thieves may be generally deterred from auto thievery if they suspect that cars are equipped with the system. And, indeed, that is what the early experience with LoJack indicates. In some major cities there has been a significant drop in auto theft after LoJack became available.

We are grateful to Ian Ayres and Steve Levitt for bringing this matter to our attention.

If there is time, you might also discuss, or at least recommend, Ronald A. Coase's "The Lighthouse in Economics," which originally appeared in the *Journal of Law and Economics* and is reprinted in Coase's *The Firm, the Market, and the Law*. The lighthouse has been cited by John Stuart Mill, Paul Samuelson, and generations of economics teachers as an example of a public good—something that because of free riding consumers, private enterprise cannot produce in a socially optimal amount. But Coase demonstrates that in the United Kingdom until late in the 19th century a private company successfully operated lighthouses for commercial shipping.

There are lots of other pithy examples that will illustrate for law students the intriguing view of the world that economics brings. We have found that chapters from Steven Landsburg's *The Armchair Economist* are wonderfully stimulating methods of getting the students to see the world in an economic fashion.

Please don't hesitate to send us your examples so that we might share them with other teachers who use this book.

[1] The company's name is a play on the word "hijack," which means to stop someone and take something from him or her forcefully. Some thieves used to hijack trucks containing valuable materials. Several years ago some criminals began to stop automobiles and take them, usually at gunpoint, from their rightful owners. That practice was called "carjacking."

The London Congestion Charge

Ulen used this example of the application of pricing a non-market good—traffic congestion—to illustrate the ability of economics to help address societal problems. See the *Microeconomic Theory II* slides, and Jonathan Leape, "The London Congestion Charge," 20 *J. Econ. Persp.* 157 (2006).

Ulen has also assigned John Tierney's article, "The Autonomist Manifesto," from the September 26, 2004, *New York Times Magazine* as a great discussion item.

Finally, Ulen has found that Harold Winter's *Trade Offs* is a wonderful short introduction to the use of economics in the discussion of public policy items.

The Arrow Impossibility Theorem

Among many if not most law students, the hardest pill to swallow about law and economics is the elevation of efficiency to the status of a serious legal norm.[2] Many law students have come to law school to further their ability to foster social justice. They will react very negatively to a contention that that goal is misguided or that it should give way to efficiency as a legal norm. We are not ones to pander to our student audiences. Nonetheless, we try to make clear to them that there are sound scholarly reasons that modern economics focuses on efficiency rather than equity. Two of the most important are the First and Second Fundamental Theorems of Welfare Economics, which argue, essentially, that efficiency and equity are separable and the Arrow Impossibility Theorem.

One of us uses a handout on the Arrow Theorem that says the following:

The Arrow Impossibility Theorem addresses the issue of how society aggregates individual preferences about social matters (*e.g.*, about the distribution of income and wealth) into societal preferences. Suppose that these aggregations are made by majority voting. We could imagine that elections are devices for converting individual preferences into societal preference orderings: candidates announce their social welfare functions and the associated distribution on the utility-possibility frontier that they intend to pursue and voters then choose among the candidates, with that policy, social welfare function, or candidate winning that commands the highest number of votes.

Make the following five assumptions about this means of aggregating individual preferences into social preferences:

1. There is no dictatorship—*i.e.*, no one person's preferences determine the group choice.

2. Each individual has ordered all the alternatives according to her preferences and votes for that policy, social welfare function, or candidate that ranks highest in her preference ordering.

3. If every individual unanimously agrees on an alternative, then that alternative is indicated as the society's preference.

[2] Some authors, such as Kaplow and Shavell in *Fairness Versus Welfare*, make a serious case that efficiency (in the sense of improvements in individual welfare) should be the *only* legal norm. We discuss their claims at Web Note 1.1.

4. Each individual's choices are complete, transitive, and reflexive.

5. The preferences between any two candidates or policies depend on how people rank those two alternatives, not on how they rank other alternatives. (This is known as the axiom of the independence of irrelevant alternatives.)

For the purpose of illustrating the Theorem, let us assume that there are only three individuals in society and three policies, candidates, or social welfare functions. Suppose that the individuals' preferences among the three policies—call them *x*, *y*, and *z*—are as follows (with P indicating the relationship "is preferred to"):

Individual 1	*x* P *y*	*y* P *z*	*x* P *z*
Individual 2	*y* P *x*	*y* P *z*	*z* P *x*
Individual 3	*y* P *x*	*z* P *y*	*z* P *x*

Each individual has complete, transitive, and reflexive preferences over the relevant social choices. For instance, for Individual 2, *y* is preferred to *z*, and *z* is preferred to *x*, and so by transitivity *y* should be preferred to *x*, and it is.

What happens if we try to aggregate these individual preferences into a societal preference ordering by means of majority voting? Suppose that we begin with a choice between *x* and *y*, with the winner advancing to a runoff against policy *z*. Thus, letting S stand for the relationship "is socially preferred to," we may write that *y* S *x* because both Individuals 2 and 3 prefer *y* to *x*, while only Individual 1 prefers *x* to *y*. What now happens in the runoff election between *y* and *z*? Individual 1 votes for *y*; Individual 2 votes for *y*; and Individual 3 votes for *z*. Thus, *y* wins so that *y* S *z*, and *y* is the socially preferred policy.

Just for the sake of completeness, what would have happened if we had begun with the pairing *x* and *z*? In that case, Individual 1 would have voted for *x*, but the other two individuals would have voted for *z*, making *z* the winner. Thus, *z* S *x*. If we then advanced the winning policy, *z*, to a runoff against policy *y*, we already know that *y* would have won.

This means that *y* is *the* socially preferred alternative, regardless of the order in which the alternatives are considered. Matters seem to be in good order. Majority voting has converted completed, transitive, and reflexive individual preferences into complete and transitive *social* preferences. (Can you show that the social preferences are, in fact, transitive?)

But Professor Kenneth J. Arrow, a Nobel laureate in economics, demonstrated in *Social Choice and Individual Values* (1952) that this result did not always hold. That is, he showed that complete, transitive, and reflexive individual preferences cannot necessarily be converted into complete, transitive, and reflexive social preferences by means of majority voting that obeys the five assumptions listed previously. To see why, suppose that individual preferences over the three policies of candidate alternatives were as follows:

Individual 1	*x* P *y*	*y* P *z*	*x* P *z*
Individual 2	*y* P *x*	*y* P *z*	*z* P *x*
Individual 3	*x* P *y*	*z* P *y*	*z* P *x*

At first glance there appears to be very little to distinguish this set of individual preferences from the first set. (The only difference has to do with how Individual 3 feels about *x* and *y*.) In both instances each individual has complete, transitive, and reflexive preferences. Let us conduct an election among these policies or candidates to get the social preferences. If we begin with an election between policies *x* and *y*, *x* wins 2-1, so that *x* S *y*. Now pit *x* against the remaining policy *z*; *z* wins 2-1, so that *z* S *x*. It appears to be the case that *z* is the socially preferred policy.

But suppose that the first pairing is not *x* and *y*, but *z* and *y*. If we held an election between alternatives *z* and *y*, *y* wins 2-1. And we know that if we were then to hold an election between *y* and *x*, *x* would be determined to be the socially preferred winner. Finally, if we were to start our election by pitting *x* against *z*, *z* would win. If we were then to pit *z* against *y*, *y* would be determined to be the socially preferred policy.

There's clearly a problem here. We get three different socially preferred policies depending on the order in which we pair then initially. (This possibility of circular group preferences in majority voting was first noted by Condorcet (1743–1794) and is sometimes called the "Condorcet paradox.")

The problem is that majority voting may not give rise to transitive social preferences. We know that if the group preferences were transitive, then, because *z* S *x* and *x* S *y*, it should be the case that *z* S *y*. But notice that *y* S *z* because two people prefer *y* to *z*.

The gist of the Arrow Impossibility Theorem is that even though individual preferences are complete, transitive, and reflexive, group preferences determined through majority voting may not be. There is apparently no way to distinguish between those sets of complete, transitive, and reflexive individual preferences that will give rise to transitive social preferences and those that will not. The only method of guaranteeing transitive social preferences through majority voting is to relax one of the five assumptions made at the beginning. But it is difficult to see which of those five ought to be relaxed.

As you probably know, one of the implications of the Arrow Impossibility Theorem that social choice theorists have teased out is that there are circumstances in which one can achieve one's objectives by controlling the agenda. For example, in the second example given above (the one on page 6) suppose that you are Individual 3, whose preferences indicate that you prefer *z* to either of the alternatives. Further suppose that you, Individual 3, are in control of the voting procedures for your group of three. If so, then by suggesting that the first pairwise voting to occur will be between *x* and *y*, you can guarantee that *z* will be the ultimate selection of your group.

There is a wonderful real-world example of how a clever former law professor used this knowledge of the value of agenda control to achieve his desired end. See "The Flying Club," Chapter 3 in William H. Riker, *The Art of Political Manipulation* (1986). For more on the Arrow Theorem and some possible solutions to its dire predictions, see Daniel A. Farber and Philip P. Frickey, *Law and Public Choice: An Introduction* 38–62 (1991).

Please feel free to use that material. It may not take long to present it in class, and it may serve to indicate why economists are somewhat skeptical about aggregations of individual choices into social choices.

Suggested Answers

In-Chapter and End-of-Chapter Questions
Cooter & Ulen, *Law and Economics* (6th ed.)

Chapter 2

In-Chapter Questions

2.1. (p. 17) Remember that the slope of the line is the coefficient of x. When that coefficient is positive, there is a direct relationship between x and y. The second equation says that there is an inverse relationship between x and y: When x increases in value, y decreases in value.

2.2. (p. 19) Think about transitivity over time. Is it necessarily true that a hot dog at lunch is the same good as a hot dog at dinner on the same day? Why would this not be the case? It surely is not the case that our tastes do not depend on what has happened in the recent past (even in the distant past) and the context in which we find ourselves.

2.3. (p. 20) As you reduce x from x_0 to x_1, notice that if the amount of y remains constant at y_0, then the consumer will be on a lower indifference curve, U_2. To return the consumer to the higher level of well-being represented by U_3 the amount of y in bundle must increase. How much? Draw a vertical line upward from the intersection of y_0 with U_2 until you reach indifference curve U_3. Then, draw a horizontal line from that intersection point over to the y-axis. That will give you the amount of additional y that must be given the consumer to make up for the loss of x.

2.4. (p. 21) Increases and decreases in income will shift the curve, while a change in the price will cause the line to rotate.

2.5 (pp. 22-23) Compare the marginal costs with the marginal benefits. If the initial level is greater than the optimum, then marginal costs are higher than marginal benefits. That is, the cost of the final unit exceeds the benefit received from it. The decision-maker will continue to decrease the amount being maximized until marginal benefit equals marginal cost.

2.6. (p. 24) Another way to ask this question is, "At what point does society stop receiving benefits from additional units of pollution reduction?"

2.7. (p. 24) Refer to your answer to question 2.6. In Figure 2.7, show the area that constitutes the net benefits. That's the triangular area above the MC curve and below the MB curve.

The technological change that lowers the MC of achieving any given level of pollution will lower the MC curve – that is, it will rotate downward. Therefore, its intersection with the MB curve will be further to the right. So, the socially optimal amount of pollution will be *less* than what it was before the technological innovation.

If it turns out that the health or environmental costs of pollution are *greater* than previously thought, then the marginal benefit (MB) of less pollution is *greater* at every level of pollution. As a result, the MB curve shifts upward. The intersection of the MB and MC curves will, therefore, be further to the right, suggesting that the new social optimum will have *less* pollution.

2.8. (p. 40) Given the technology used by this firm, external costs will always exist under any positive level of production. Think of an extreme case: This might be one in which the social cost of the activity was *always* above the social benefit. First-degree homicide might be an example. In those instances, the socially optimal amount of the activity is 0. But for non-extreme cases, where there is a range in which the social costs are below the social benefits, a nonzero amount of the activity is socially optimal, with some external costs being generated.

End-of-Chapter Questions

2.1. Economists generally assume that economic decisionmakers are attempting to maximize something subject to constraints. Thus, consumers are assumed to maximize utility subject to an income constraint, and firms are assumed to maximize profits subject to a production function constraint. Maximization helps to posit the goal or end that the economic decision-maker seeks. Equilibrium is a state of rest, the condition from which no further endogenous change will occur. The notion of equilibrium helps to specify the point (e.g., the quantity of output produced, the amount of bananas consumed) toward which the maximizing behavior of economic decision-makers tends. [p. 52]

2.2. Efficiency describes a point of equilibrium with particular characteristics. A productively efficient equilibrium describes the condition of a firm or firms in which it is impossible to produce a given level of output at lower cost or, alternatively, to use a given combination of inputs to produce a greater level of output. An allocatively efficient equilibrium describes an equilibrium distribution of goods and services among consumers. A particular distribution of goods among consumers is allocatively efficient if it is not possible to redistribute the goods so as to make at least one consumer better off without making another consumer worse off. (That, of course, is the Pareto criterion.) [p. 52]

2.3. In microeconomic theory consumers are assumed to maximize utility subject to an income constraint. The income constraint or budget line is described by the consumer's income and the relative prices of goods and services. The consumer's utility is maximized when she achieves the highest attainable indifference curve. This occurs at the point at which the highest attainable indifference curve is tangent to the budget line. At that point, the benefit (i.e., the utility) from spending an additional dollar on any given good is equal for all commodities. That is, at the consumer's constrained maximum the marginal benefit (in terms of the increase in utility) of an additional dollar spent on any good is exactly equal to the marginal cost (i.e., the dollar) of any good. [p. 52]

2.4. The partner who has the children is assumed to have a utility function of $u = cv$, where c equals the weekly child support payment from the other partner, and v equals the number of days per week that the children spend with this partner. Initially, $c = 100$ and $v = 4$ days, so that $u = 400$. If the partner paying child support payments wishes to reduce the weekly support payment to $80, then the number of days that the children spend with the other partner must increase to 5 in order for utility to remain unchanged at 400.

2.5. Price elasticity of demand measures the responsiveness of consumer's quantity-demanded to changes in relative price. Mathematically, price elasticity is defined as the percentage change in the quantity-demanded of a commodity divided by the percentage change in the commodity's price. Because quantity-demanded and price move in opposite directions (when price declines, quantity-demanded increases; when price increases, quantity-demanded decreases), the sign of price elasticity will be negative. To avoid having to remember whether 25 is greater or less than 24, economists drop the sign from price elasticity and talk about it as if it were simply a positive number. The ranges of value of price elasticity of demand are inelastic (price elasticity has a value less than 1), unitary elasticity (price elasticity equal to 1), and elastic (price elasticity has a value greater than 1). When price elasticity is greater than 1, the percentage change in quantity-demanded is greater than the percentage change in price—for example, a 10 percent drop in price gives rise to a 23 percent increase in quantity demanded. Can you compute the price elasticity of demand for this good?

There is an interesting relationship between a commodity's price elasticity of demand and total consumer expenditures (price times quantity-demanded) on that commodity. When price elasticity of demand is less than one, an increase (decrease) in price will lead to an increase (decrease) in the total amount that consumers spend on the commodity. When price elasticity of demand is unitary, consumers will spend the same total amount on the commodity regardless of the price. When price elasticity of demand is greater than one, an increase (decrease) in price will lead to a decrease (increase) in the total amount that consumers spend on the commodity.

We shall use this last relationship between price elasticity and consumer expenditures in Chapter 12 to suggest a troubling relationship between drug addiction and crime.

2.6. The expression "There is no such thing as a free lunch" is one of the most famous quotes in all of economics. Its origins are interesting. In order to encourage customers to drink, taverns used to post a sign saying "Free Lunch." The implication was that those who had purchased drink were entitled to eat from the bar's buffet at no extra charge. Clearly the bar expended real resources providing the luncheon buffet, and the explicit and opportunity costs of these resources figured in the bar's determination of the costs of doing business and, therefore, in its calculation of its profits. But what about the bar's customers? Was the lunch really free to them if they purchased a drink? No. The bar included a charge for the costs it incurred in putting on the buffet in the prices it charged its customers for drink. The fact that the charge for the buffet is hidden in the price of drink does not make the lunch free.

On many commercial airplane flights, passengers used to be offered meals, music, and movies at no apparent price. Were these entertainments really "free"? If they were free, why do you think that many of these enhancements have disappeared from airline travel?

2.7. Use your answer from 2.6 to think about this issue. Hint: There is an opportunity cost to a vacation. What could you be doing with your time instead of relaxing?

2.8. Firms are assumed to maximize profits, which are defined as the difference between total costs (including opportunity costs) and total revenues. A firm maximizes profits by choosing that output level for which the marginal cost (the addition to total cost of the last unit of output produced) equals the marginal revenue (the addition to total revenue of the last unit of output produced). (Note that if we call marginal revenue "marginal benefit," then the rule to maximize profits by choosing the output level for which marginal cost equals marginal revenue is exactly

equivalent to our general maximization rule of equating marginal cost and marginal benefit.) If the firm finds itself producing an output level for which it is the case that marginal revenue exceeds marginal cost, then by producing more output it can add more to total revenue than to total cost and thereby in-crease its profits. Alternatively, if marginal cost exceeds marginal revenue, the firm should cut back on production: the revenue lost from lower output will be less than the cost savings.

2.9. In a perfectly competitive industry there are a large number of buyers and sellers, so large that no single buyer or seller can influence the market price by his or her individual decisions. Entry and exit of resources into and out of the industry is free. These are the core characteristics of perfect competition about which there is general agreement among economists. There are several other conditions about which there is less widespread agreement—for example, that all buyers and sellers have perfect and complete information and that products are homogeneous. It is often said that the stock market and the market for agricultural commodities (such as wheat) are examples of perfectly competitive industries.

Monopoly occurs when there is only one seller. (For more on monopoly see the suggested answer to the following question.) By comparison to a perfectly competitive industry, a monopolist produces too little and charges too much.

Oligopoly holds where there are only a few sellers, so few that they recognize the interdependence of their decisions. That is, what is optimal for firm A to produce or to charge for its output depends not just on market price, marginal revenue, and the firm's own costs but on what firms B and C produce and charge. This interdependence gives rise to strategic considerations that are best analyzed through the use of game theory and such notions as a Nash equilibrium. The television programming market may be an ex-ample of an oligopolistic market.

An imperfectly competitive industry shares some of the aspects of perfectly competitive and monopolized industries. As in perfect competition, there are a large number of sellers, although the number is not as large as in perfect competition. Entry and exit are free. As in monopoly, each seller has some limited market control over the consumers of his output. Sellers distinguish their product by brand name, quality, and other characteristics. The market for breakfast cereals is probably imperfectly competitive. [pp. 52-53]

2.10. A monopoly occurs when there is only one seller. There are technological conditions, known as "natural monopoly," that can give rise to monopoly. There are also social welfare conditions that may make a case for the government's granting a monopoly—as with the patent, copyright, and trademark systems. (See the discussion of intellectual property in Chapter 5).

A monopoly can be sustained only if other resources are prevented from flowing into the industry to set up competitive firms. This might happen when a monopolist gains control of the only input that can be used to produce a particular output. Far more common are cases in which the government permits the monopoly to endure by forbidding lawful entry into the industry. Where entry is not so restricted, it will eventually occur, leading to the demise of the monopoly.

A special kind of monopoly is a cartel, which is collusion among otherwise competitive firms that seek to operate as a joint profit-maximizing monopoly. An example is the Organization of Petroleum Exporting Countries (OPEC). OPEC illustrates very well the difficulties that monopolies have in forestalling entry. The more successful the monopoly initially is in raising the price above the competitive level, the greater the incentive to others to enter the industry. After OPEC's success in raising the price of oil to nearly $33 per barrel (from $3 per barrel) in

the early 1970s, the entry into the petroleum-extracting industry by other countries was substantial, so much so that the OPEC countries now account for less than half of world petroleum output. (Later, the high price of oil induced a substantial amount of cheating by OPEC member countries, too.)

2.11. We discussed the adverse effects of a price ceiling with respect to rent control. The price ceiling is set below the equilibrium price so that there is an excess demand for the price-controlled product. Suppliers can make more money supplying unregulated goods and services and so transfer their efforts and resources elsewhere. Those suppliers that remain attempt to keep their profit rates at normal levels by letting the quality of their product (and therefore its costs of production) decline and by insisting on secret payments from consumers over and above the controlled price. Consumers of the controlled product may also alter their behavior in inefficient ways: they may, for example, inefficiently substitute other goods and services for those whose price is controlled.

2.12. Refer to your answer to 2.11. Instead of placing a ceiling above which prices can-not rise, the government sets a floor below which they cannot fall. The effect of a price floor is that too much is produced and not enough demanded. The simplified outcome in the labor market is that minimum wage constraints cause unemployment—that is, an excess labor supply.

2.13. The interaction between a minimum wage law and a law forbidding firing would play out at many subtle levels. Among the many issues to consider are: the decision to hire new workers, the use of overtime, and the decision to declare bankruptcy.

2.14. All other things equal, the implicit price of divorce would fall under a no-fault divorce law, thus leading to more divorce. (Exactly how much depends, of course, on the price elasticity of demand for divorce.) This fall in price might also influence the decision to marry. Perhaps there would be less hesitation to marry, knowing that the decision to marry was easily reversed.

2.15. The effort associated with acquiring information concerning interest rates effectively raises the price of obtaining a loan. With lower costs faced by consumers, all else equal, one would expect more lending activity. The profits to the lenders depend on the cost of supplying the information to consumers and the market power they exhibited before the Act was passed. The more competition in the lending industry, the closer the interest rate should reflect the true risks and costs to the lender.

2.16. General equilibrium is the condition in which all markets are in simultaneous equilibrium. The conditions under which this will happen are that all markets are perfectly competitive and that there are no sources of market failure. (See the answer to the next question.) Given the resources available to the economy, it is not possible for the economy to be productively or allocatively more efficient than it is in general equilibrium. (As we argue in the text, distribution is another matter.)

2.17. Market failure can arise from five sources. The first is non-competitive market structures, such as monopoly (the condition of one seller) and monopsony (the condition of one buyer). We have already examined the reasons why monopoly leads to sub-optimal results. In general,

policies of regulation (in the case of natural monopoly) and of antitrust enforcement (in the case of collusion and artificially sustained monopoly) will help to correct for the social costs of monopoly. The second is external effects, costs and benefits involuntarily imposed on others. In the case of external costs, such as water pollution that results from a firm's discharge of waste chemicals into a nearby stream, the externality-generator does not take into account the costs he has involuntarily imposed on others. As a result, his output costs him less to produce than it should in terms of the social resources his production really uses. So, someone who is generating external costs produces too much. To correct for these social costs, he must be induced to "internalize" these external costs, in which case he will reduce his output to the socially optimal level. Policies of taxation and subsidization, of exposure to legal liability, and of prior regulation can help to minimize the social costs of externalities. (In Chapters 4 and 5 we discuss how property law can help in this internalization, and in Chapters 8 and 9 we show how tort law helps to internalize risks of injury.)

The third source of market failure is the presence of public goods. These are goods for which the cost of excluding non-paying consumers are so high that no private profit-maximizing firm can earn enough to justify producing the socially-optimal amount of the public good. The market fails in respect to public goods in that it produces too little of them. A policy of subsidization of private producers or of public provision of the public good can correct for this failure.

The fourth source of market failure is extreme informational asymmetries. Because of the special characteristics of information, where access to information is highly skewed or where the ability to process information is highly unequal, questions arise about the optimality of otherwise voluntary transactions. (In Chapters 6 and 7 we discuss how contract law, and in Chapters 8 and 9 we discuss how tort law, can help to correct for problems arising because of informational asymmetries.)

The fifth source of market failure is collective and coordinated decisionmaking. There are circumstances—nicely illustrated by the Prisoner's Dilemma in the text—in which individuals are led to do things that may be in their own narrow interest but are not in the interest of the collectivity or group of which they are a member. The prisoners in the famous game could both do better if they could coordinate their behavior and not confess, but because it so costly to engage in believable coordination, the prisoners both confess.

2.18. a. A swimming pool large enough to accommodate hundreds of people does not seem to have either of the characteristics of a public good non-rivalrous consumption or costliness of excluding non-paying consumers. A fence and charging an entrance fee through a central entrance dispose of the costliness issue. Only when the pool is congested does the amount of the pool that one person can consume depend on the amount that others consume. The swimming pool does not seem to be a public good, and it there-fore follows that a private profit-maximizing firm could provide the optimal pool.

b. A fireworks display is clearly a public good. Consumption is non-rivalrous in the sense that one person's enjoyment does not diminish the amount available for another to enjoy. And it is costly to exclude non-paying consumers; they may, after all, simply stand on a hill away from the display and enjoy it without charge. For those reasons it is not surprising that it is municipalities, not private firms, that provide fireworks displays on the 4th of July and that they are generally financed from general tax revenues, not from a user fee.

c. A heart transplant is clearly a private good. Nonetheless, many will shy away from the conclusion that this particular private good should be allocated according to the usual rules prevailing in the market for private goods—presumably because transplants are extremely costly. For an interesting discussion of this issue and of other issues where high cost is an important component of choice, see GUIDO CALABRESI & PHILIP BOBBITT, TRAGIC CHOICES (1978). We also have a discussion of the situation for organ transplants in a box in Chapter 5 and in the Web Notes.

d. Vaccination against a highly contagious disease is a private good but one that has such strong external benefits that the market will provide sub-optimal amounts of it – that is, in deciding whether or not to receive inoculation, people are likely to take into account only their own well-being and that of their immediate friends and neighbors and not the public health benefits that their inoculation will confer on others. Thus, either public provision of vaccination or public subsidization of private provision is called for.

e. A wilderness area could conceivably be a private good. But today wilderness areas and national parks are federally owned and operated, presumably on the theory that these areas should be held in trust for future generations and that private owners would not do that. But there are a few examples of private ownership of natural resources that do work. A private organization called the Nature Conservancy attempts to preserve bird sanctuaries, wilderness areas, and other significant natural resources by using the contributions of its members to purchase and hold these areas.

f. (Also, g and h.) All three examples of education given here are largely private goods with trace elements of public goods in the sense that they contribute to societal well-being by giving those trained a sense of worth and a set of skills. There is general agreement that elementary education has such strong external benefits (in that all citizens are better off when their fellow citizens can read and write and reckon) that it should be subsidized or publicly provided.

2.19. Pareto efficiency (or Pareto optimality) describes a situation of equilibrium from which it is impossible, given the economy's resources, to produce more of one commodity without producing less of another or to make one person better off without making another worse off. Although Pareto efficiency is to be desired, we saw that it does not define a unique distribution of resources among the members of society. The set of Pareto efficient distributions that results from voluntary exchange depends crucially on the initial distribution of resources among the members of the society. Different initial distributions lead to different Pareto optimal outcomes.

2.20. Although votes cannot be bought and sold in the marketplace, a market for votes does exist in the form of campaign contributions and logrolling (legislators trading votes). What benefits would follow from allowing legal and open market for votes rather than the current "hidden" market? More troubling, how would allowing an open market change the balance of power between the haves (who can supply money and votes) and the have-not's (who can only supply votes)? Would allowing this market to exist fundamentally change our concept of government?

2.21. A Pareto improvement is a change in which at least one individual is made better off, while no one is made worse off. The Kaldor-Hicks requirement recognizes that some changes will be very beneficial to one subset of society, while harming another subset. As long as the benefits outweigh the costs, Kaldor-Hicks views the change as a superior out-come. A Pareto improvement requires those who benefit to compensate (or to purchase the consent) of those who lose. A Kaldor-Hicks improvement does not require explicit consent by the losers, only that someone objectively verify that the benefits of the change exceed the costs.

2.22. A dominant strategy is a strategy that is optimal regardless of what the other player does. If both players have a dominant strategy, then there is a Nash equilibrium. A Nash equilibrium is a strategy where an individual player cannot do better, so long as the other players do not change their strategies. Some games have multiple Nash equilibria, and many games (the most common of which are the prisoner's dilemma games) have a Nash equilibrium that is not the Pareto efficient.