21 Markets as Systems

In the first few chapters of this book we presented the fundamental question modern economics was created to answer: can the awesome power unleashed by the industrial revolution be left to the push and pull of the marketplace, or does it need to be steered by the conscious intervention of government or some other institution acting on behalf of society? Adam Smith thought that markets could do most of the job on their own, and subsequent generations of economists have struggled to identify the precise conditions under which Smith's Invisible Hand could be expected to function. Much of what they discovered has been summarized in the Market Welfare Model and the many adjustments and caveats we have examined in our tour of microeconomics.

Up to this point, however, we have looked at markets one at a time, largely in isolation from one another. We have looked at the market for coffee, for instance, as an example of supply and demand and also bargaining power, but separately from the other markets that interact with it, such as the labor markets for workers in the coffee sector and for consumers of coffee, or the markets for equipment used to process and ship each year's harvest. Much can be learned by putting individual markets under a microscope, but much is lost. In this final chapter we will step back to look at the entire system of markets and ask what the invisible hand hypothesis would mean at this level, and whether it would be justified. Our topic is the market economy as a whole.

Markets, we may recall, are simply the result of adding up individual two-sided exchanges. Two parties come to an agreement over a transaction—how much to buy or sell, at what price, and with what stipulations of quality, timing etc.—and their decisions are combined with thousands or millions of others to produce market prices and quantities. The question of whether the market system as a whole functions in the social interest largely boils down to whether desirable social decisions, such as those about what to produce and in what way or how the proceeds should be divided up, can be made by adding up very large numbers of one-on-one agreements.

Without going into the matter more technically, we can see the outlines of a likely debate. On the one side, leaving economic choices to a multitude of

P. Dorman, *Microeconomics*, Springer Texts in Business and Economics,
DOI 10.1007/978-3-642-37434-0_21,

individual agreements has the advantage of decentralization, spreading decision-making to a great many people. This in turn has the potential to draw on their diverse skills, attitudes and sources of information. It also is noncoercive in the sense of relying on negative freedom, as discussed in the appendix to Chap. 6. These are all attractive features, and for many they are appealing enough to make them devotees of "free market economics", no matter what the other drawbacks might be.

On the other side, society is more than just a sum of separate individuals; it has social and political structure, cultural values, shared physical spaces and the other qualities we find in the commons. What guarantee is there that, if people pursue only their personal self-interest, these social factors will be taken into account? Also, isn't it likely that, by not incorporating the indirect consequences of our choices, the sum of individual agreements will be very different from what society as a whole might prefer? We saw a stark example of this in the prisoner's dilemma, where the pursuit of self-interest can lead to worse outcomes for everyone compared to what they could get by making common decisions for the group. Moreover, individuals have highly unequal bargaining power, and the distribution of rewards generated by markets may be incompatible with commonly held standards of justice and fairness. Thus, one might begin this chapter with a deeply-held bias against markets.

In the pages to come we will see what contemporary economics has to add to this debate. It turns out that many of these issues are much better understood now than they were a few decades ago. The controversy over how much economic decision-making can safely be left to markets continues, but at a higher and more constructive level.

We will begin by setting forth a criterion for evaluating economic outcomes that has played a central role in modern social theory, **Pareto optimality**. Next we will turn to a stripped-down version of the economic analysis of markets as a system, which goes under the name **general equilibrium theory**. We will look at it from positive and normative angles, and then consider how it is put into practice in the form of **computable general equilibrium** (CGE) models. After this we will look at the qualifications that have been introduced by modern research: the **General Theory of the Second Best**, the problem of indeterminacy of equilibrium due to false trading, and the positive and normative issues raised by the existence of multiple general equilibria. Only then will we return to the core question of market analysis and the invisible hand for a final summing up.

21.1 Economic Efficiency and the Pareto Principle

In Chap. 6 we looked at the normative side of economics from the vantage point of a single market using the framework of net social benefits. This was the basis for the Market Welfare Model, which sets forth the conditions under which a particular market equilibrium can also be regarded as producing the best possible outcome for society. Unfortunately, this approach cannot be used to analyze an entire system of markets, so we will have to develop another one. Our immediate objective is to produce a formal definition of economic efficiency to serve as our criterion.

Let's begin with a relatively straightforward problem, determining whether a particular production system, such as a steel mill, is efficient. We could say that it meets this standard if it is not possible to produce any more output without using at least some additional inputs. This corresponds to the basis for the Production Possibility Curve presented in Chap. 3. By the same token, we could describe an allocation of goods between two people as efficient if it is not possible, through a reallocation, to improve the utility of one person without reducing the utility of the other. It is really the same idea, but using the production of utility rather than goods as the objective, achieved by rearranging outputs rather than inputs.

It is not difficult to translate this concept into familiar terms; for instance, a particular technology for producing steel may be regarded as efficient because there is no other way to produce as many ingots of the same quality without using more iron, or coal, or labor, or some other input. Similarly, the allocation of meals to two customers at a restaurant is efficient if it is not possible to switch plates so that both are happier with what they're eating. If each is eyeing the other's order and they don't make the switch, there is a loss of potential efficiency in terms of utility.

So the concept of efficiency is clear enough when we are examining small pieces of the economy, but how can it be applied to a world with millions of goods and billions of people; that is, how can we determine whether an entire economy is efficient? To be realistic, we would have to specify two different types of efficiency, static and dynamic, where **static efficiency** refers to how efficiently goods and resources are allocated in an economy at a particular moment in time, and **dynamic efficiency** refers to how well the economy is organized to grow over time. To keep matters as simple as possible in this chapter, we will focus only on the first of these. Our tool will be the concept of Pareto optimality, named for the early twentieth-century Italian sociologist Wilfredo Pareto. The main ideas are summed up in Fig. 21.1 on the following page.

Each axis measures the utility level of an individual in a two-person economy; the points in the quadrant represent economic outcomes: levels of production of goods and their distribution between the two potential consumers. The downward-sloping curve extending from one axis to the other is a **utility possibility frontier**; any point on or within it represents a feasible state of the economy. In particular, we are interested in points A, B, and C—how do they rank?

The principle of **Pareto optimality** states that one allocation is preferred to another if no individual is worse off and at least one individual is better off. Clearly that criterion can be used to show that A is preferable to B, since both individuals have higher utility at A, but what can we say about the ranking of A and C, or even B and C for that matter? In the first case both points are Pareto optimal, and that's the end of the story; neither can be said to be better or worse than the other by this criterion. As for the second, the inability to rank is paradoxical. After all, there is no point on or within the utility possibility frontier that would be Pareto preferred to C,

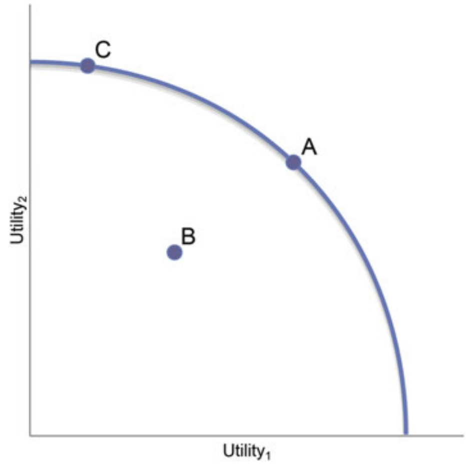

Fig. 21.1 A utility possibility frontier with Pareto rankings. Maximum potential combinations of utility for individuals 1 and 2 are shown by the curved line. Points *C* and *A* are Pareto optimal; neither can be made better off at such a point without reducing the utility of the other. *B* is worse than *A*, but the Pareto principle can't tell us how *C* and *A* or even *C* and *B* compare

since the only way one person could be made better off is by making the other worse off, whereas many points (such as A) would be Pareto preferred to B; nevertheless, C is not Pareto preferred to B, since individual 1 is better off at B than at C. This suggests the fundamental weakness of Pareto optimality as a criterion for normative efficiency: it doesn't supply a ranking over all, or even over a majority, of potential comparisons.

To remedy the situation, economists have advanced a more flexible version in the form of **potential Pareto optimality**. An allocation is potentially Pareto preferred to another if the individuals who gain from the first could fully compensate those who lose from it (so that after compensation they would be no worse off than under the second) and still have some of their gains left over. Suppose, for instance, that the utility depicted in Fig. 21.1 can be measured in dollars, and that each individual has the same dollar-to-utility exchange ratio. The transfer of equal amounts of utility/dollars from person 1 to person 2 or from 2 to 1 would take the form of movement along a line whose slope is -1 (U_1/U_2). Starting from C, for example, it is clear that potential compensation could result in point D, as shown in Fig. 21.2.

The loss of utility experienced by individual 2 due to paying compensating is exactly equal to the gain by individual 1 for receiving it, so both points could be regarded as equally ranked. But now we also have ways to rank A versus C and B versus C: A is potentially Pareto preferred to C since it is preferable to D, and by the same logic C is preferred to B. Allowing for the possibility of compensation has given us a complete ranking of all the points. For this reason, the principle of potential Pareto optimality is far more powerful than its more rigid, compensation-blind cousin. For purposes of practical implementation, economists are usually willing to assume that dollars are a close proxy for utility, so that maximizing the dollar value of an allocation is equivalent to maximizing its ranking according to the potential Pareto principle.

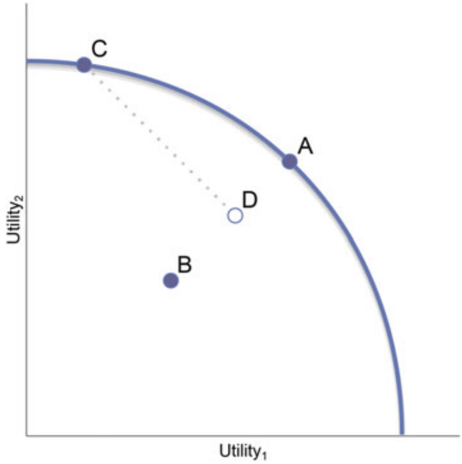

Fig. 21.2 A utility possibility frontier with potential Pareto rankings. The addition of the potential Pareto equality between *C* and *D* makes it possible to rank all four points: *A* best, *C* and *D* equal, *B* worst

It is important to emphasize that the compensation envisioned under the potential Pareto principle is strictly hypothetical. In general, this compensation does not actually get paid; in fact, if it did, potential Pareto optimality would be transformed into Pareto optimality pure and simple. This is because, having received their compensation, those who would otherwise lose from a redistribution would now be at least as well off as they were originally, while those who paid compensation would (if the redistribution were potentially Pareto preferred) be better off. In fact, this is what usually happens in a market transaction. Goods or services are redistributed from one owner to another; in return, compensation is paid. The ability of one side of the transaction to compensate the other suggests that all voluntary transactions are Pareto-improving and do not need the additional boost of the potential Pareto criterion.

To sum up so far: the Pareto principle covers a minority of potential alterations of the economy, but its standard is met in most or all market exchanges. For other possible economic interventions, such as those that might result from nonmarket methods like government regulation, it is necessary to have a comprehensive system of evaluation like the potential Pareto principle. In either case we still encounter all the difficulties with preference theory that we surveyed in Chap. 11, of course. If the general framework of utility analysis, with its assumptions about human behavior and well-being, are incorrect, as they appear to be, the particular techniques of Pareto and potential Pareto analysis have little to add. Nevertheless, since they play a central role in general equilibrium theory and economic approaches to policy analysis, we will keep them in view.

21.2 General Equilibrium

Consider two markets, one for labor and the other for rice. Let's imagine that labor is needed to harvest and process rice and that rice is important for the diet of workers. As the price of labor goes up, it becomes more expensive to produce rice. On the other hand, with higher wages workers will want to buy more rice. If we take just one market into consideration, say rice, it is not difficult to determine what we mean by market equilibrium: it is where the amount of rice people want to buy equals the amount producers want to sell at a common price. We could say something similar about the labor market, assuming that it works along the lines of supply and demand. (Chap. 16 gave us some reasons to doubt this.) But how do we know whether *both* markets could be in equilibrium simultaneously? After all, it may be that the wage that equalizes the supply of and demand for labor produces a cost and a demand for rice at which no equilibrium is possible. On what basis would we presume that the concept of equilibrium could apply to a system of interconnected markets and not just to one market in isolation from the rest?

This question haunted Leon Walras, a French economist of the second half of the nineteenth century. In 1871 he published a treatise setting out the results of his analysis. Walras approached the issue mathematically, expressing each individual market equilibrium as a single equation. In this equation, the quantities supplied and demanded of a good are determined by its price, along with the influences coming from all the other markets. Putting all these equations together, Walras created a system for simultaneously determining all the prices, much in the same way one would solve a system like this:

$$\begin{aligned} x + 2y &= 10 \\ 3x - y &= 9 \end{aligned} \qquad (21.1)$$

(Check your math: $x = 4$, $y = 3$.) This system can be solved because there are two equations and two unknowns. Walras' system had an indefinitely large number (n) of equations, but also same number of unknown prices. (Well, almost: the number of independent equations is $n-1$, so Walras can solve only for all the other prices relative to one given price.) Once the prices are calculated, the quantities of each good can be calculated too. The result would be a general equilibrium: the prices would satisfy the requirement that supply equals demand in each market, and they would do this simultaneously for all the markets in the economy.

Or so he thought, but the problem turned out to be more complicated than this. The little two-equation example above can be solved because it is linear; neither x nor y has an exponent, nor are they multiplied by each other. Try to solve, for instance, a different example:

$$\begin{aligned} x^2 + y &= 4 \\ x - y^4 &= 1 \end{aligned} \tag{21.2}$$

Simple methods don't work any more; in particular, one cannot say in advance that there is only one solution for such a system.

A second problem is that supply equals demand only for goods that are produced. Many goods are not produced, either because there is no demand at a price that suppliers would be willing to sell at or because there is no supply at a price that buyers would be willing to pay. In such cases what we have is not an equality (between supply and demand), but an inequality (one more than the other). Therefore, to represent a whole economy, Walras' system must include both equations and inequalities. The mathematics has become more difficult.

These two problems were solved only in the 1930s by mathematicians and economists employing tools that were not available to Walras half a century earlier. Finally, in the 1950s a complete specification, incorporating time, location and the full range of economic goods, produced and unproduced, was published by Kenneth Arrow and Gerard Debreu. It was seen at the time as a major intellectual triumph, one that provided a theoretical foundation for the entire field of economics.

What these theorists produced could be described in several ways. First, they had given a mathematical proof that an entire economy taking the form of a system of markets would have one and only one general equilibrium. This had a powerful implication: if you know the given aspects of an economy—the supply of resources available to it, the preferences of its members, and the initial distribution of resources among these members—you can determine the general equilibrium that should result. In this way the economy can be predicted and explained. To put it somewhat differently, markets, operating through the forces of supply and demand, are capable of fully regulating an economy, determining its precise outcomes over time. This is exactly what Walras had been searching for.

Second, they produced proofs of two normative propositions that have come to be known as the **Fundamental Theorems of Welfare Economics**. These are that, in the absence of market failures, (1) a general equilibrium is always Pareto optimal, and that (2) any possible Pareto optimal state of the economy can be arrived at through a general equilibrium provided that the initial distribution of resources is properly adjusted. It will take a little explanation to make these clear.

First, consider that, in a sense, the deck has already been stacked in favor of Theorem 1 because Pareto optimality rather than some other criterion, such as potential Pareto optimality, has been chosen as the basis for making evaluations. As discussed above, every market exchange is Pareto-improving: the final outcome must be Pareto preferred to the initial situation, otherwise the exchange wouldn't have taken place. Therefore the only additional burden this theorem takes on is demonstrating that *all* such Pareto improvements will occur—that in general equilibrium, and with no market failure, there are no such mutually advantageous exchanges that fail to take place. But, as we will see in slightly more detail in a moment, the model of general equilibrium economists work from assumes exactly what must be assumed in order for this theorem to hold, so it does. As a result, in the narrow but precise terms of the Pareto principle we can say that a system of markets in general equilibrium is economically efficient—again, if there is no market failure.

The second point is more subtle. Suppose in our labor and rice example we arrived at a general equilibrium in which supply equals demand in both markets, but wages are low and workers are hungry. We might prefer some alternative outcome where all workers are well-fed. This might in turn lead us to want to interfere with the market, to require that wages rise or that the price of rice be held down. Such interventions would, of course, mean abandoning the principle of market regulation for some other form of economic control.

The second theorem presents an alternative approach. It notes that there are many potential Pareto optimizing outcomes possible, such as points A and C in Fig. 21.1, depending on the initial distribution of resources in the society. If we can put our finger on one of the outcomes on the utility possibility frontier that we regard as sufficiently equitable, we can arrive at it without interfering with the way markets work, but only by changing the initial distribution. In other words, pick any efficient result and it is possible to use markets to get there, provided you do the right redistributions at the beginning. For many economists, this is an argument against price controls and other forms of public regulation and in favor of progressive taxes and income transfer programs. The advantage of the strategy of redistribution, they argue, is that by letting markets operate freely we can arrive not only at outcomes that have the balance we are looking for, but which are also efficient (as measured by the Pareto principle). In our labor and rice case, for example, this could mean redistributing some money from owners of rice-growing operations to their workers, so that demand for rice would increase, owners would invest in more harvesting equipment, rice workers would become more productive, and then their wage would rise. (This is a purely hypothetical account, and, without knowing the production and consumption responses on all sides, one couldn't say for sure what redistribution would lead to. The Second Theorem tells us there is *some* redistribution that can do the trick, but it doesn't tell us which one.)

But both the positive (single, predictable equilibrium) and normative (fundamental theorems of welfare economics) properties of general equilibrium depend on a large number of assumptions that must hold regarding how markets operate, and this is really the most important point: the general equilibrium theorists have shown us what these conditions are. The list is largely the same as the one that applies at the single market level in order to invoke the Market Welfare Model: no market failures like public goods or missing markets, no asymmetries or other distortions due to information, rational behavior (in the economic sense of Chap. 3) by every participant in the economy, and conditions in every market to achieve a single, market-clearing equilibrium. In other words, general equilibrium comes with no greater guarantee than its single-market relative, and in fact it comes with far less, since these conditions must hold throughout the economy. Since it is inevitable that they won't, practical guidance depends on how much deviation is thought to be "too much" for the positive and normative conclusions to be acceptable.

21.3 Applied General Equilibrium Models

General equilibrium theory is highly abstract, but it has given rise to a new approach to economic prediction in the form of computable general equilibrium modeling. The use of CGE has become widespread in recent years, so it is useful to know how it works.

The theories we have been describing try to represent every market operating in the economy simultaneously. Incorporating every worker and consumer, every owner of every resource, and every good or service produced, strictly applying such theories would be far beyond the calculating power of any computer. To transform this approach into workable applied models, tremendous simplification is required.

To begin with, the economy is reduced to a manageable number of markets. For instance, the most widely used model in international trade, developed by the Global Trade Analysis Project (GTAP) at Purdue University, provides 57 sectors in its most recent incarnation. This simplification is achieved by combining a large number of specific markets into one large aggregate. To take one example, depending on the desired level of detail, a CGE modeler might work with "agriculture" or perhaps a four-way division into "grains", "fruits and vegetables", "meat and dairy" and "fibers" instead of the thousands of agricultural markets that a real economy would have. Resources too must be grouped together: it is common, for instance, to use only four or five categories to incorporate natural resources, capital and different types of labor.

Next, instead of the millions of individuals and households whose decision-making drives real world markets, the model would include a single representative individual for each general role. For instance, in modeling savings decisions, a CGE model might assume that all households are the same, each with an average budget and average savings and consumption inclinations. At most, it would have one such type of household for each major sector or income class of the economy under study, like agriculture or "middle income".

If the model is concerned with international issues like trade and global environmental impacts, location has to be taken into account. This is done by specifying specific countries or groups of countries. The previously mentioned GTAP model, for example, uses 87 countries or other locational groupings.

Finally, the supply and demand factors operating within markets would be greatly simplified. Expectations of future prices, for example, might be held constant, or they might depend on recent and current prices according to some simple mathematical formula. Problems of quality differences among goods, strategic competition between firms and other complexities would mostly be suppressed. Most CGE models depend on a vision of market equilibration that looks little like the rough-and-tumble of normal life; at most they might highlight a market or two for more realistic treatment.

Although it gets ahead of our story somewhat, mention needs to be made of the macroeconomic dimension of CGE models. These include such results as the level of unemployment, the government's budget surplus or deficit and the trade balance of a national economy. Many techniques have been tried, but none are satisfactory at the present time. The most common approach among CGE modelers is to simply exclude macroeconomic effects, so that the outcomes listed above are assumed to remain constant. For instance, in models of international trade the usual approach is to assume that changes in trade policy have no effect on any country's trade balance (the difference between imports and exports), but only on the composition of trade (the particular goods imported or exported). In addition, it is assumed that changes in trade policy will have no effect on the number of workers employed in any country. While assumptions like these can be criticized, and frequently are, they result from the simple supply and demand framework (general equilibrium) the models are based on. The tension between supply and demand analysis and other approaches to the study of macroeconomics will be a thread running through the second volume of this text.

Once the model has been made simple enough, the next step is **calibration**. Each equation has parameters that govern the relationship between key variables, like prices, and outcomes, like quantities bought and sold. These relationships should produce results that look like those in the actual economy. For instance, one such relationship is the elasticity of demand, the percentage change in quantity demanded divided by the percentage change in price. If a particular equation represents the automobile sector, the elasticity of demand in it should correspond to the actual relationship between changes in auto prices and number of vehicles sold. Calibration is the process of setting all these parameters so that the model approximates the real world as much as possible.

The final step is to enter the data that tells the model what the initial state of the economy looks like: the incomes of the households, the production levels for each sector, initial prices and any other variables of interest like investment rates and international trade flows. This sets a starting point for running the model. Now the researcher is in a position to propose a change to one or more of the parameters. Suppose, for instance, that a law is passed that raises production costs in one sector; how will this ricochet through the economy to produce a new general equilibrium? Altering this one factor and recalculating the model will provide an answer.

CGE models have found many uses. They are popular in the analysis of trade policy, such as changes in tariffs or other trade barriers. They are used to estimate the ultimate effect of a tax change, particularly where it might effect one segment of the economy directly but the others indirectly. As you might expect, there have been many CGE analyses of policies to mitigate global climate change, since both the policies, like carbon taxes, and the effects of climate change themselves are likely to have large impacts throughout entire economies. Would increasing the cost of burning hydrocarbons like coal and oil raise or lower the incomes of farmers? How could you find out without doing some form of modeling that takes into account the main direct and indirect linkages between energy and other sectors?

Despite their increasing prominence, CGE models are not viewed favorably by all economists. The simplifications required to make these models tractable inevitably suppress much of the unpredictable dynamics of real events; the future is seldom as much like the past as the models predict. General equilibrium itself is a doubtful proposition given all the conditions that must hold if it is to arise as the models assume, and, as we will see, new questions regarding general equilibrium have emerged since the classic version of the 1950s. Above all, these models have a weak track record. There are few cases where CGE modelers have been able to successfully predict the results of economic policy changes or other shocks, and many of the well-known models have generated predictions that were nearly the opposite of what actually happened. In their defense, proponents of CGE have asked, if not this then what? Models of single markets or other small pieces of the economy make even greater assumptions, since they hold everything else constant, so what other direction is there to go in? Their hope is that, as computing power and modeling sophistication advance, their methods will become more effective. Critics feel that the shortcomings of general equilibrium as a basis for modeling will vitiate any improvement in technique.

21.4 The General Theory of the Second Best

The first significant qualification to general equilibrium theory, although it was not seen this way at the time, appeared in 1956 in the form of an article written by the Canadians R. G. Lipsey and Kelvin Lancaster. It called into question what mathematicians would call the asymptotic properties of general equilibrium—whether the properties of such an equilibrium, like the two Fundamental Theorems, apply more fully the closer one gets to it. This is an important question for any ideal arrangement in the social and natural sciences, because the real world seldom matches an ideal model perfectly, but instead moves closer to or further from it.

Suppose, for instance, a farmer in a dry region wants to know how much irrigation is required for growing a crop like sunflowers. It turns out that a particular amount of water, distributed in a particular way over the season, is optimal and yields the largest, healthiest crop. If water is expensive, however, the farmer might want to settle for a bit less than this. The crucial piece of information would be the relationship between water availability and crop growth as a little less than the optimal amount is used. Normally, we would expect that, if the water is nearly optimal so also will be the crop. If not—if a small decrease in water or some other input could cause a large decrease in the sunflower harvest—this is crucial information to have. It is this issue, transferred to the realm of economics, that is illuminated by the General Theory of the Second Best.

To return for a moment to the perspective of a single market, the Market Welfare Hypothesis presents a vision of the best possible state of affairs, where the marginal benefit to society of producing an extra unit of something is exactly equal to its marginal cost. Moreover, when all the assumptions of the model are adhered to, the equilibrium price, determined solely by the market, will convey both pieces of information, marginal benefit and marginal cost, perfectly. It is almost inconceivable, however, that this will be true for each and every market; surely there will be some which are out of adjustment for some reason, either because the assumptions do not hold (e.g. the presence of external costs or benefits) or the market is unable to reach equilibrium. The theory of market failure tells us that markets with such distortions may need surgery in the form of public intervention. But what about the otherwise undistorted markets connected to it via the linkages of a market system? Suppose the MWM conditions appear to hold for market A, but not for market B to which it is linked. Does this mean that only B needs attention, and that A can be safely ignored?

The general answer is no. It comes to us as The General Theory of the Second Best, which can be stated in this way: if the conditions for optimality are violated in one market, attainment of the second best generally requires that they be violated in at least one other market as well. By second best we mean the best that can be achieved given that the economy-wide conditions for the first best ($P = MB = MC$) are not attainable. The proof of this proposition requires more math than we can deploy here, but the intuition behind it is surprisingly simple. Consider the following:

Suppose you give me the best possible directions from the Eiffel Tower in Paris to the Coliseum in Rome, and that they take the form of "follow this road 2 miles, then turn right and continue for 8 miles, turn right again and continue for 300 miles" and so on. They are the best instructions in the sense that, if I follow them exactly, I will arrive at the right destination via the shortest possible route. I take these directions and almost follow them perfectly. Unfortunately, instead of turning right on the first turn, I turn left. If I continue to follow each succeeding instruction to the letter I might end up not in Rome but, say, Warsaw. There is a lesson here. Once I made my first mistake I was no longer on the first-best route. Worse, by continuing as if everything were just fine I failed to follow even a second-best (or tenth-best) route. On the contrary, once I had veered from the ideal path, my best course would have been to violate at least one further instruction, so that my itinerary, while not as good as the initial one you gave me, is the best possible given my earlier mistake. This additional adjustment could be as simple as turning around and returning to the original itinerary, or it could entail an entirely different itinerary, but the general point holds: one unprogrammed turn requires another.

So it is with economics. The Market Welfare Hypothesis can be understood as a set of instructions which tell us to arrange each market such that $P = MB = MC$. If the economy deviates from this prescription in one market, however, it will usually be necessary to make some offsetting change in another market to compensate. Here is a directly relevant example from current economic debates. Suppose that the price of oil is "too low". This might be because environmental externalities (which increase the true social cost of producing and distributing oil) are not taken into account, or because the governments which own most of the world's oil supplies give insufficient attention to potential future scarcities, or perhaps because our grandchildren, who will inherit a world largely depleted of this resource, are not represented in today's markets. If one or more of these factors are at work, the first-best option would be a substantial increase in the price of oil. But this may not be possible, either for political reasons, or because the macroeconomic effect, like recession, would be too severe. That puts us in the situation of searching for the second-best option. One measure that might be taken—in fact, it is a policy that the US government has followed for nearly four decades—would be to mandate, by law, that auto company product lines meet minimum standards for fuel efficiency. Such a policy is obviously an infringement on the freedom of buyers and sellers in the automobile market to pursue the dictates of what they see as their individual rationality.

If we look more closely at the fuel-efficiency policy it becomes apparent why a breakdown in the optimality properties of the oil market might demand an offsetting intervention in the car market. One of the assumptions for the Market Welfare Model to hold is that the demand curve reflect the true marginal social benefit from acquiring another unit of the good under consideration. In our example this means that the amount consumers are willing to pay for different types of cars truly reflect the benefits they provide. But in making their choices, consumers are incorporating in their calculations the price of gas. If gas were priced correctly, internalizing all externalities, this would not be a problem. We are assuming, however, that the price of gas does not reflect its true cost, and so consumers are not giving sufficient weight to fuel efficiency when they choose between cars. Thus they will tend to purchase larger gas guzzlers, imposing on society the costs of too much air pollution and a too-rapid depletion of oil reserves. To put it differently, the mispricing of oil has created a problem not only in the oil market, but in the car market as well. If we are unable to correct the price of oil, the second-best alternative is to compensate by intervening in the car market. (Whether this is best done by fuel efficiency standards, tax incentives on different types of cars, subsidy of public transit, or some other device is another matter entirely.)

Taken literally, the general theory of the second best applies everywhere. This is because there are many markets in which the assumptions of the Market Welfare Hypothesis do not hold, and because, directly or indirectly, all markets in the economy are interconnected; so these distortions are disseminated throughout the entire system. Consider one more example: Everyone knows that the labor market, which sets the price of labor for different types of work, is highly imperfect. We see discrimination by race and gender, barriers to the ability or willingness of workers to move between jobs, lack of compensating wage differentials, excess supply (unemployment) and many other features, discussed in Chap. 16, that suggest that wages do not represent the "true" costs to society (whatever that might mean!) of employing people. But all industries employ workers, and the prices they charge for their products depend in part on what they have to pay the people who produce them. If the wages are not right, neither are the prices of the products.

In practice, a large percentage of economic policy is concerned with second-best-type tinkering. Since so many prices in our economy are out of sync with true marginal benefits and costs, we must do a lot of ad hoc adjustment to minimize the resulting economic irrationality. On the other hand, it is certainly true that in many markets the distortions arising from linkages with the rest of the economy are minor, and the gains to be made by intervening are not worth the potential harm of overriding the market mechanism. Ultimately, where to draw the line—where to say, "This is an urgent problem requiring a policy to achieve the second best", and where to decide that the market works well enough—is a question of values. In our car example, for instance, everything depends on the initial assumption that oil is seriously mispriced. If oil is only slightly mispriced there may be no need to do anything about it. Reasonable people, of course, can disagree about whether a particular price distortion is serious or not.

If there is a single lesson to be learned from the theory of the second best, it is this: to evaluate the case for intervention versus laissez-faire in any market, it is necessary to look at the most closely-linked markets as well. The economy is an interconnected system, not a mere jumble of parts.

Recall also the point raised at the beginning of this section, about the farmer wondering whether or not small deviations from the optimal irrigation policy will lead to correspondingly small changes in the harvest. The General Theory of the Second Best can be seen as giving economic analysts just this sort of information. Any economy can move closer to being optimal in the sense that more markets adhere to pricing based on marginal cost and willingness to pay, but *further* in terms of the value of its outcomes to society, since second best may require more, not fewer, deviations from the marginal cost/benefit principle. Thus, even though a world in which all the assumptions of general equilibrium and perfect competition hold might be best of all, it would be a mistake to think that the more this world is approximated the better off we will be. This is a highly paradoxical result, which says that what works in its pure or perfect form may be a false objective in a world of compromises.

Second best analysis, however, should not be undertaken at such an abstract level, but in relation to the specific distortions that can't be removed in the economy. It is not enough to just invoke the principle every time you want to make an intervention; there has to be a clear connection between the intervention in one market and the distortion in another.

21.5 Out-of-equilibrium Trading and the Indeterminacy of Equilibrium

Real markets are seldom if ever in equilibrium; so the predictive force of theories based on equilibrium depends on the mechanisms that tend to move markets toward their equilibrium state. We have already examined the forces that do this at the level of a single market. If, for instance, the price is initially below its equilibrium level buyers will experience a shortage: the amount supplied will be less than the amount demanded. There will be a tendency for the price to be bid up, thereby stimulating additional production, and this process will continue until an equilibrium is reached. The story would not change a bit if the initial price were very much lower than equilibrium or only slightly lower; the same forces would be at work, but with greater or less intensity, and the same ultimate equilibrium would set the direction to which prices and quantities were heading. If we begin with a price that is too high, a similar mechanism works, except in the other direction: the bidding down of prices, the decrease in production. It does not matter what the initial

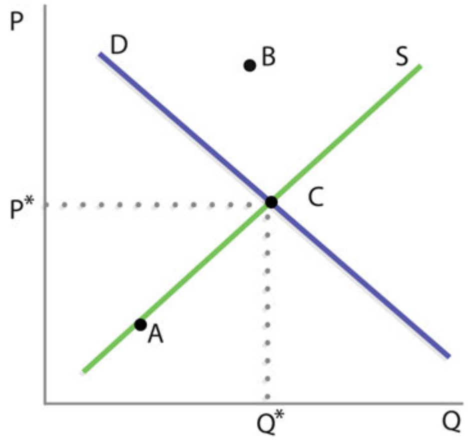

Fig. 21.3 In a single market, the equilibrium doesn't depend on the starting point. Point *C*, where the price is P* and the quantity is Q*, is the equilibrium and eventual destination for the market, whether it begins at point *A*, point *B* or somewhere else

situation is in this analysis of a single market: the equilibrium toward which the market tends is the same. Recall that this characteristic is enshrined in the third condition required for the Market Welfare Model to hold. It is portrayed visually in Fig. 21.3 above.

This property is clearly important for economics, since it establishes that, once we know the supply and demand curves, we know the one and only equilibrium. *This property does not hold in general equilibrium.* On the contrary, the initial out-of-equilibrium state of the economy plays a crucial role in determining which equilibrium the economy as a whole will end up at, as was shown in a series of articles published separately in the 1970s by Gerard Debreu, Rolf Mantel and Hugo Sonnenschein.. The mathematics behind this result are complex, but the idea is straightforward. In the single market case the only out-of-equilibrium possibilities concern prices and quantities that may be too high or too low. There is no reason why these starting points (A and B in Fig. 21.3) should affect the shape or locations of the demand and supply curves since they are not ceteris paribus conditions. The same cannot be said for out-of-equilibrium conditions in the economy as a whole

Suppose, for example, that the price of agricultural products is below equilibrium. This is reflected not only in the market for these goods, but all other markets as well. Farmers, after all, will have less income if their products are underpriced, and this means there will be less demand for the things farmers buy. Farm equipment manufacturers will face a shift in their demand curves, and there will be less demand for other farm-related goods. The ripple effects continue, as the consequences of these second-round effects cycle through the economy. Ultimately, if no other disturbances arise, a general equilibrium would be reached, but it would not be the same equilibrium as the one that would result from, say, an initial situation in which agricultural prices were too high. The curves themselves, and not just the momentary prices and quantities, have been altered.

Simply put, we can contrast single-market and general equilibrium as follows: In a single market, equilibrium is determined by supply and demand only. In general equilibrium, the outcome also depends on the initial and every subsequent out-of-equilibrium state of the economy: which prices are too high, which too low, whose income depends on which prices, and what the preferences are of these individuals. Clearly this presents a problem for both the positive and normative dimensions of economic analysis. On the positive side, the general equilibrium an economy tends toward cannot be predicted from supply and demand only, nor even from the initial state of the economy, since it also depends on continued out-of-equilibrium ("false") trading that may further shift the market conditions on which equilibrium depends. On the normative side, the equilibrium that exerts its gravitational force at any moment in time is only one of a vast number of possible equilibria, each the result of purely arbitrary factors, like the temporary mistakes traders may make. Even if the first two assumptions of the Market Welfare Model hold throughout the economy, so that the demand curves reflect marginal benefits and the supply curves marginal costs, there is no reason to suppose that society is better off under one of these equilibria than another.

How serious is this problem for the effectiveness of the market as an allocative device? It is difficult to say. If these potential equilibria are close together—if, to follow the previous example, the equilibrium resulting from initially high agricultural prices is broadly similar to that resulting from initially low prices—the difficulty is not too great. On the other hand, it is possible to imagine situations in which the effects of out-of-equilibrium adjustment are of great importance. Take the case of the great California Gold Rush of the late 1840s.

One could imagine that, just prior to the gold frenzy, there was an "equilibrium" distribution of Euro-Americans within the US. On the basis of their preferences for where to live, the technologies for transportation available to them and their wealth, a certain number would have chosen to move to California, with the rest moving to other regions or staying put. Then a false gold strike was announced. Based on misinformation, many thousands flocked to the west coast. They didn't find gold, but many found economic opportunities catering to the gold prospectors, as well as all the others catering to the caterers, etc. Once it became apparent that there were no fortunes to be made sifting sediments from California riverbeds, a few '49ers returned home, but most stayed. Some were unable to afford passage home, but most stayed because the new economy spurred by the Gold Rush changed their opportunities: they had now adjusted to a new "equilibrium" set of locations. The result was that the demand to live in California had shifted irreversibly; its effects continue down to the present. Was this ultimately for the better? Does the fact that the course of history was changed as a result of a mass delusion affect your answer?

21.6 Interaction Effects and Multiple Equilibria

The market is an institution that structures and processes human interaction. Workers compete for jobs, and the labor market determines who will be employed and at what wage. Firms compete for consumers, and markets determine which goods earn their producers a profit, and who will end up purchasing them. When we draw demand and supply curves, we are incorporating two types of social interconnections: the making an agreements between buyers and sellers, and competition between buyers or between sellers for the privilege of making such agreements. At the level of a single market, and assuming that these are the only relationships between people that should affect the final outcome, we can infer a single equilibrium.

The situation changes dramatically when individuals affect one another in a variety of ways, only some of which can be channeled through markets. First, as we have already seen, there arises the possibility of external benefits and costs, which directly violate the first two conditions of the Market Welfare Model. But that is not all. Nonmarket interaction, the connections between people or the goods they own that occur outside the boundaries of the marketplace, often give rise to multiple market equilibria.

Here is a simple example at the level of a single market. Suppose we consider the market for storefront space in a shopping mall. The mall owner wants to rent it for as much money as possible, given the demand from those who want to establish retail businesses. The retailers, on the other hand, want to get the best rents they can, also taking into consideration the advantages of location. At first, we might imagine that these interests give rise to the normal downward-sloping market demand curve and upward-sloping market supply curve.

But let us add one additional factor. The desirability of renting any particular property depends on the number of other stores in the mall, since that will affect the amount of foot traffic. No one, after all, wants to be the only store-owner in a "dead" mall. So, if we begin with an empty mall, only a very low rent would attract more than a few retailers to move in. As additional stores open up, however, the location becomes more attractive, and so the mall owner can now raise rents and still find willing buyers. Eventually the pool of retailers looking to relocate is exhausted, so additional inducements must be offered to rent still more space to new individuals and companies. At this level of demand, the demand curve is once again downward-sloping.

The demand relationship discussed in the previous paragraph, combined for simplicity with a perfectly elastic marginal cost curve, appears as in Fig. 21.4 on the next page.

To keep everything as simple as possible, let's assume that the mall owner is "nice" (maybe it is a nonprofit or a government agency) and wants to set the rental price at marginal cost as a perfect competitor would, and not according to the monopoly strategy outlined in Chap. 13. There would then be two possible equilibria. First, if we start at a low level of Q, to the left of the diagram, the equilibrium quantity is Q_1. But if the initial situation is one where many stores already occupy the mall and Q is on the right side of the diagram, Q_2 is the equilibrium. The best strategy for the owner, if the mall is new and initially unoccupied, is to set a rental price below P*, even though it entails a short-term loss. After the number of stores moving in exceeds Q_2, begin to raise the price until it reaches the equilibrium level MC.

Already in this extremely simple example we can see several points that characterize the general theory of interaction effects. First, there is more than one

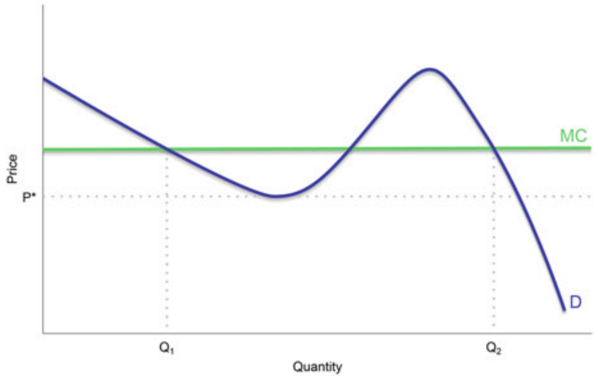

Fig. 21.4 Demand and marginal cost for rental space in a shopping mall. Demand for space at a shopping mall depends on the number of other stores already there. At the same rental price there could be a low demand if the quantity is already low, and a high demand at higher levels of total Q. To set MC = P at the higher level of demand, the mall owner should set a price below P* until Q exceeds Q_2, then raise the price back to MC

equilibrium, in the sense that there is more than one price/quantity combination from which no one would make small deviations. At any starting point in the vicinity Q_1, if buyers and sellers consider small changes in how much they are willing to buy or rent at, they will converge on this amount. The same applies to Q_2. Why do we specify only "small" changes? Because this is how most real-world markets work: each individual participant makes up only a small piece of the market and is able to change quantities demanded or supplied by a only tiny percentage.

Second, if there is more than one equilibrium we can no longer invoke the Market Welfare Model. A single equilibrium is one of the conditions, and the reason should be obvious: there is usually only one optimal allocation, but there are now two which qualify as equilibria. In fact, one of the equilibria in Fig. 21.4 (Q_1) looks pretty dreary, inasmuch as it would leave the mall mostly empty. A useful way to think about this point is that, if there is only one equilibrium in an economy, finding it is very important, but if there are many equilibria it is finding the *right* one that matters.

Third, in the absence of additional information we cannot predict which equilibrium will be chosen by the market. Suppose, for instance, that the mall owner is unable to do more than make very small adjustments in the rent. (Again, this corresponds to the situation in a large market with many participants, where none are able to have much effect.) This means that if the initial situation happens to be on the left side of the diagram, market forces will push it towards Q_1, or to Q_2 for the right side. In this way, the arbitrary force of history, which determines where we "begin" at any moment, makes itself felt. It may be, however, that the mall owner is not willing to leave matters to fate, but instead hires an economist to determine (after a long, expensive study) that, even if the initial situation happens to be the first equilibrium, the owner can do better by getting to the second through the strategy of a temporary price cut. What this means, however, that such a solution to the problem of blind historical inertia succeeds only by replacing the market mechanism with a planning process—replacing incremental adjustments of individual buyers and sellers with a blueprint for an overall, systemic shift.

Finally, note that the source of the problem is that retailers interact not only through the market, but also in the corridors of the mall, in the sense that each retailer's decision to open a store affects the decision of the others by attracting shoppers, and this effect is not directly incorporated in the market. In other words, there is a market in rentals but not in the effect that one rental has on another; this takes the form of a network externality. This is a general point about the relationship between nonmarket interaction and multiple equilibria.

From the standpoint of social theory, we have arrived at an interesting moment in the discussion—toward the end of the nineteenth century, to be exact. In the hundred years following Adam Smith, the notion that individual and collective rationality are reconciled via the invisible hand was the mainstay of educated public opinion in much of Europe and the United States. All social phenomena, it was believed, could be explained by reference to individuals acting in their own personal interest, and if their choices were uncoerced—if markets were free—the case for laissez-faire was self-evident. In the 1880's and '90's, however, a new generation of social scientists emerged to challenge this world view. A society, they said, is more than the sum of its member individuals; it also consists of the interrelationships between them. And not all of these connections occur through the market. On the contrary, Emile Durkheim, often regarded as the founder of modern sociology, felt that individuals are bound together by common intellectual and cultural forces, which they experience as coercive. What I feel, and therefore what I choose, depends on what you feel and, in fact, on what we as a society feel together. Max Weber, the German historian and sociologist from roughly the same period, identified other bonds between people, cemented in the social relations of family, work, government and other institutions. For him, the modern world was less and less an arena of free individual choice, and more a "iron cage" in which the comforts of life are purchased at the cost of massive social regimentation. In both cases, and for other pioneers of the social sciences, individuals were not seen as isolated atoms, bouncing off one another in the frictionless space of free markets; rather they were interconnected in many ways that the market could hardly recognize, much less organize.

The modern social sciences now investigate many systems of nonmarket interaction. There is the kinship system, incorporating family and other relationships of "private" life. There are political systems, in which individuals affect one another through their places in the hierarchies of law and governance. There are the human ecologies of urban life, where geographic proximity becomes the basis for neighborhoods and larger communities. There are cultural and discursive systems, such as the mass media, through which the values and perceptions of individuals are both controlled and given creative force. To this modern economics adds the strategic interaction modeled by game theory: the interplay between principles and agents, cooperators and defectors that people the institutions that make up our economy. And above all there is the natural environment itself, a maze of interconnected systems and cycles, binding together all living and nonliving things in a web of mutual causation, a complex of interaction we considered in the previous chapter. If these perspectives on the interconnectedness of people and things are valid, we would expect to see many instances in which the choices made in the marketplace also trigger significant nonmarket interactions, with the result that market equilibrium is no longer uniquely determined, and factors other than market forces must be taken into account to determine the actual course of events.

Here is a simplified example, which is not intended to represent the full complexity of the issue. Individuals might live predominantly in dense urban areas in which systems of mass transit, such as busses and trains, will best serve their needs. Or most might live in spread-out, more sparsely-populated suburbs for which the private automobile is the vehicle of choice. Each situation may have associated with it a unique equilibrium which reflects the rational decisions of all individuals concerned. If housing becomes a little too decentralized in the first scenario, some people living far from the transit lines may decide to move back into the city, restoring equilibrium. In the second scenario, a family in the suburbs may decide it needs a second car, so that more family members can have access to the places they need to get to. That too may represent movement toward equilibrium. In both cases a system of markets in transportation, housing, and other goods might serve as an efficient mechanism for achieving Pareto superior outcomes, i.e. the best allocation of consumer dollars to housing and transportation in either instance.

But what about the more fundamental question of which scenario we will inhabit? Since each scenario has its own market equilibrium, markets alone cannot do the job of choosing between them. Some other mechanism must be at work. Most likely, it is historical inertia: we will continue to live in the urban/mass transit world if this is what we have inherited from the past, or the suburban/automobile world if this is the initial reality. While real world differences between transportation and residential networks, for instance between parts of Europe and the United States, have more complex causes, the example can serve to illustrate the logic of multiple equilibria under which markets would fall short of both their positive and normative roles. That is, they would not offer a sufficient explanation for the what, how, and for whom of production, nor provide a mechanism for achieving socially rational outcomes. Note that this latter point holds irrespective of which scenario you think conforms to social rationality.

The culprit in this example, responsible for the multiplicity of equilibria, is nonmarket interaction. In fact, there are quite a few candidates for Most Important Nonmarket Interaction in the Field of Urban Development and Transportation. A short list would include the following: obligations (or the lack thereof) to remain close to family members, which can affect and be affected by where people choose to live; the role of neighborhood institutions (social capital), based on stable residential patterns, in making cities more liveable and desirable; and, of course, the effect that housing has on the demand for transit, along with the effect of transportation choices on the demand for housing (the possibility of complementarities). Each of these enter separately into individual market decisions about where to live and what to buy, but the effects they have on one another are not accounted for in markets. The result is a situation in which individuals can exercise individual rationality, but there can be no presumption that the combined outcome is collectively rational.

Economic geography, the field from which the previous example was drawn, has been revolutionized by the use of models incorporating interaction effects and generating multiple equilibria. Another realm in which this approach has become central is the study of financial markets. In an atmosphere of high stakes and great uncertainty, traders are profoundly affected by each others' choices. One very simple case might go like this: each trader might want to follow the lead of the majority in the market; if most of the other traders are bulls, then it is better to be a bull, and otherwise a bear. (Bulls are traders who think the market is headed up; bears think it is headed down.) Why? For one thing, if the trader is working with other people's money there is a great risk of taking a loss, but especially when others are not losing. Then the fund owners might ask, why are you losing when everyone seems to be coming out ahead? The extra cost of being the exceptional loser might outweigh the potential benefit of being a winner when others are losing. If so, it makes sense to go with the crowd. A similar result would occur if each trader is susceptible to being persuaded by the prevailing wisdom simply because it is prevailing. (There is a large literature in social psychology supporting the existence of this tendency.) Going with the crowd, if it becomes the most common strategy, readily creates multiple equilibria. If the crowd thinks the market is moving up, it is in the interest of individual traders to make bets based on moving up. That's one equilibrium. Or the crowd could turn bearish, and each individual trader chooses to go along with that too, creating a second equilibrium. There is quite a bit of evidence that actual financial markets oscillate between multiple equilibria in something like this fashion: that's one reason it's difficult to predict where the markets are headed in advance.

We are now in a position to return to the argument made in the previous chapter regarding complex ecological interaction. Suppose two factories, one making aluminum and the other beer, are located beside a small lake; let's call them A and B for short. Each of them uses the lake to dispose of chemical wastes. This is a problem for the fish who live in the lake, and also for a local sport-fishing resort: the fewer fish the lake can support, the less money will be made by the resort. To make matters very clear, let's suppose that the lake is actually owned by the resort, which has the right to prohibit any form of pollution if it chooses, and that the effects on the resort's profits are the only ones that need to be considered. It would seem we have a perfect setup along the lines of Coase to arrive at a market solution for what would otherwise be an externality. There are no missing markets, and the resort is in a position to compare the marginal cost of each form of pollution at each level to the willingness of each factory to pay for the right to pollute; these bargains should, as we read in Chap. 15, lead to a market equilibrium which is also an efficient allocation of resources—in this case, pollution in paper- and chip-making and water quality for the fish.

Now we will put some numbers to the problem. Each factory has its own effect on the fish based on the tons of waste emitted per week according to Table 21.1, where the effect can be thought of as a financial measure of potential damage to the fish stock as experienced by the resort, for instance in thousands of dollars.

Table 21.1 Separate effects of two factories on lake fish at different pollution levels, measured in tons of waste

Tons of waste	0	1	2	3	4
Factory A	0	1	3	6	10
Factory B	0	2	5	9	14

Table 21.2 Combined effects of pollution from two factories on lake fish (additive)

		Factory B				
		0	1	2	3	4
Factory A	0	0	2	5	9	14
	1	1	3	6	10	15
	2	3	5	8	12	17
	3	6	8	11	15	20
	4	10	12	15	19	24

The rows indicate the possible pollution levels, in tons of waste, from Factory A, the columns the possible pollution levels of Factory B, and the cells reflect the damage to fish resources.

Table 21.3 Marginal costs of pollution from two factories (additive effects)

	0→1	1→2	2→3	3→4
Factory A	1	2	3	4
Factory B	2	3	4	5

What is the combined effect of these two types of pollution? In the simplest case we might just add them together, so that the total cost to the resort would be given in Table 21.2.

Each cell is calculated by adding the sum of the two effects, A + B. For instance, if A emits 2 tons and B emits 3, the highlighted cell shows the sum (3 + 9), the two effects from Table 21.1. The next step is for the resort to calculate the marginal costs of each kind of pollution. This can be determined, as it turns out, directly from Table 21.1, and the calculations are given in Table 21.3.

The marginal cost of one type of pollution is unaffected by the amount of the other in the lake. For instance, if B emits 1 ton, and if A goes from 2 tons to 3, the effect rises from 5 to 8, for a marginal cost of 3. If B is emitting 2 tons and A goes from 2 to 3, the effect rises from 8 to 11—still 3. This means that the resort owner can deal with each source of pollution separately. So let's add the factories' side of the story; this

Table 21.4 Potential social surplus derived from Table 21.2

		Factory A				
		0	1	2	3	4
Factory B	0	−4	−2	−1	−1	−2
	1	−2	0	1	1	0
	2	−1	1	2	2	1
	3	−1	1	2	2	1
	4	−2	0	1	1	0

will provide the basis for bargaining. In the interest of continued simplicity, suppose that the marginal cost of reducing pollution is constant for both factories, 3 for A and 4 for B. That is, it costs A 3 units of money to cut its pollution by one ton per week, no matter how much it is currently emitting, and B's cost per ton is a steady 4.

This means we have all the information we need to describe what economically efficient bargains would look like. Begin with the negotiation between the resort owner and A. To allow the first ton of waste the resort would require a payment of at least 1. This would be in the interest of A, since it would cost 3 to stop it. The next ton would cost the resort an additional 2, so it would still be in the interest of A to offer more than that to continue dumping it. One more ton would just barely pass the test, since now the additional cost of pollution exactly equals the additional cost of abating it. The result is that 3 tons would be emitted, with A paying the resort something between 6 (the continuing damage to the fish) and 12 (the money saved by A for not having to eliminate these tons).

Next would come the negotiation between the resort and B. By the same logic they would arrive at 3 tons as well. The result would be that 15 units of damage to the fish would be allowed, while the factories would save ($3 \times 4 + 4 \times 4 = 28$) in pollution control costs. Is this socially efficient? We can answer that question by constructing a table whose cells reflect the social surplus from each amount of pollution control, defined as the benefit of reducing pollution to that level minus the cost. The benefit we can calculate as 24 (the cost of maximum pollution) minus the damage at each cell in Table 21.2; the cost is the combined cost of pollution control to A and B. This would give us Table 21.4.

It turns out that the pair of agreements does maximize the potential net social benefit from pollution control (although in this example the measurement of pollution is too lumpy to narrow down the range of four optimal possibilities). There is, as in Coase's original analysis, only one social optimum and bargaining between the affected parties arrives at it.

Now let's see what happens when pollution, operating through an ecosystem, is interactive. Again keeping to simplicity, let's change the combined effect of the two forms of pollution to $A \times B$, making it multiplicative rather than additive. This will mean that marginal effect of each type of waste will depend on the level of the other. Table 21.5 gives us the new joint impacts.

From this we can calculate two tables of marginal costs, for A and for B:

Suddenly things have become much more complicated. Each row in Table 21.6 is computed from the corresponding row or column from Table 21.5. For instance, if B is dumping 2 tons per week we are in the column headed "2" in Table 21.4.

Table 21.5 Combined effects of pollution from two factories on lake fish (multiplicative)

Factory A \ Factory B	0	1	2	3	4
0	0	0	0	0	0
1	0	2	5	9	14
2	0	6	15	27	42
3	0	12	30	54	84
4	0	20	50	90	140

Table 21.6 Marginal costs of pollution on lake fish (multiplicative)

Factory B \ Factory A	0→1	1→2	2→3	3→4
0	0	0	0	0
1	2	4	6	12
2	5	10	15	20
3	9	18	27	36
4	14	28	42	56

Factory A \ Factory B	0→1	1→2	2→3	3→4
0	0	0	0	0
1	2	3	4	5
2	6	9	12	15
3	12	18	24	30
4	20	30	40	50

Thus, as A goes from no tons to 1, damage goes from 0 to 5, then from 5 to 15 and so on. The row and column these numbers come from are highlighted.

The first impression we get is that the resort owner's task has become much more complicated. Since each marginal cost depends on the amount of the other pollution, each negotiation now depends on the other. (We can imagine the resort owner running back and forth between two offices where bargaining is taking place.) But the problem is actually more difficult still. Suppose, for instance, that the resort begins discussions with A. A is willing to pay as much as 3 per ton to continue polluting. So the resort owner might think, "Maybe B will be willing to purchase the right to dump 1 ton, so I should consider myself as being in the second row of A's marginal cost table. This means that, at a maximum price of 3 per ton, I should sell the right to dump one ton, since at 2 tons the marginal cost rises to 4, out of A's price range." So the deal is made, and then negotiations between the owner and B begin. If A is emitting just one ton, B will want to buy the right to dump 3 tons. That's a problem, since the agreement with A was based on the assumption that B would be emitting just one; now there is no room for an agreement with A. But if A is forced to end all pollution, the resort can now afford to sell the right to emit all four tons to B, since there is no marginal cost to B's pollution at all. The result is that the three parties would arrive at an equilibrium: no pollution from A and four tons from B. But the very opposite process could also occur. If a preliminary

Table 21.7 Potential social surplus derived from Table 21.5

		Factory B				
		0	1	2	3	4
Factory A	0	112	116	120	124	128
	1	115	117	118	118	117
	2	118	116	111	103	92
	3	121	113	99	79	53
	4	124	108	82	46	0

agreement with A leads to at least two tons of pollution from this source, no pollution by B can be allowed. But then it makes sense to sell all four tons to A and produce a *second* equilibrium: A pollutes four tons and B none.

In short, there are *two* Coase bargaining solutions to this problem, and both are equilibria. Is one equilibrium better than the other? Yes. Consider Table 21.7, which calculates net social benefit from pollution control in the same way Table 21.4 was calculated, except that now 140 rather than 24 is regarded as the maximum pollution cost. This makes the numbers much higher in Table 21.7, but it is the pattern that matters.

As we can see, both highlighted options maximize net social benefits *in the nearby area of the table*. Small movements away from either would reduce this benefit. Nevertheless, having A eliminate all pollution and allowing B to dump away is the better choice, since it is more costly for B to reduce its emissions. Since all of the social benefit is potentially available to the resort, either through a cleaner lake or more payments from the factories, if all of this information is publicly available the resort (after paying an economist a lot of money to sift through it) should choose the better option.

It is not likely, however, that the resort will have access to the factories' cost data. Instead, this will normally be revealed only during the process of negotiation, and, since negotiation must begin somewhere, the initial decision whether to limit primarily one sort of pollution or the other is likely to determine which option finally results. It will take a bit of luck for the resort to choose the right course of action.

Real world environmental issues are, of course, much more complicated than this deliberately simple example. There are often many sources of pollution, and their effects interact more complexly than a formula like A $\times$ B could possibly capture. There is seldom a single owner of a natural asset that needs to be preserved, and severe coordination and transaction cost problems are likely. The result is that potential markets, if we could set them up, face a multitude of equilibria with almost no possibility of selecting the best one. What is needed in such cases is an entity that can take in the entire scope of the problem, identify an overall solution and get all the parties involved to carry it out. Decentralized markets need to give way to centralized plans.

In the end, we should not be surprised by the pattern we've seen in these examples. Markets come to a collective decision by adding up individual choices. If there are important interactions between people or between the goods being exchanged, these will be left out of the process. Markets are thorough but myopic. They scour the economic landscape, looking for every small bit of improvement: possible exchanges that will put a resource in the hands of someone who wants it a little more. They proceed incrementally, however, one exchange at a time. If the situation calls for a coordinated approach, making many changes simultaneously because of the interconnections involved, markets stumble.

In very general terms we can see two types of coordinated action that economies need in addition to markets. One is public intervention—regulation, rule-setting, public enterprise—which was surveyed in Chap. 9. The other is the administrative coordination found in business organizations, and particularly entrepreneurship, the subject of Chap. 8. We don't often think of them as related, and they differ greatly according to which objectives they set for themselves and to whom they are accountable. Nevertheless, from the perspective of the theory of interaction effects, both embody a sort of planning that can span multiple equilibria and make choices among them. Fortunately, both forms are malleable: governments can become more entrepreneurial in style and corporations more responsive to democratic mandates. One of the main trends in modern politics is the search for hybrid economic and political forms that can answer society's need for innovation and coordinated action to respond to the challenge of interdependence in ways that are both democratic and economically efficient.

21.7 A Final Summing Up

We began with the Invisible Hand and we end with it. Adam Smith conjured up this image, but he left it for later generations of thinkers to determine whether it had any validity. Microeconomics as we understand it today is the result of painstaking efforts to identify the precise conditions under which markets could be expected to serve the larger social interest.

In very general terms, we have considered markets under three different sorts of lenses. We have looked at them as social institutions, as a multitude of processes that generate the prices and quantities we see for particular goods and resources, and now as a single interlocking system of allocation.

As social institutions, markets begin in metaphor, seeing all of economic life as an expression of two-sided exchange. An idealized realm of "the market" is culled from the diverse experiences societies have with markets that have evolved in different ways. This idealization sacrifices realism, but in return it offers powerful analytical tools, such as supply and demand analysis and the Market Welfare Model. Markets can also be scrutinized by methods based on other metaphors, like the prisoner's dilemma and bargaining power.

Much of this book is dedicated to a close examination of the way markets function in specific contexts: how shifts in supply or demand affect equilibrium for commodities like coffee, the role of market failure, the social dimension of work, inequality and poverty, and the interplay between economic, political and ecological factors. This has given us a more realistic sense of how markets actually function and what consequences they have. The Invisible Hand reappears as a benchmark in many of these cases, rather than as a force that can be relied on to operate on its own.

Now we have taken a very wide view of markets as they might constitute an entire system. The main message of this chapter is that a market system, unaided by other forms of organization, would be too limited. It would not take interaction effects into account, it would be too vulnerable to the random effects of error and out-of-equilibrium trading, and it would not necessarily function better in closer proximity to the stipulations of the Market Welfare Model than further from them. At this very general level the Invisible Hand simply cannot operate. But this should not surprise us, since no actual economy functions on the basis of markets alone; all depend on the multiple systems of allocation we first examined in Chap. 3. In particular, government and large-scale business are as fundamental to modern economies as markets. What general equilibrium theory offers us is not a challenge to the basic features of modern capitalism, but a basis for seeing more clearly what tasks have to be accomplished and how we should measure success.

The Main Points

1. An economic situation is Pareto optimal if it is not possible to improve one person's utility without reducing someone else's. If it is possible, the reallocation that achieves this is called a Pareto improvement. The problem with Pareto optimality is twofold. First, there is a vast number of allocations that are Pareto optimal (or efficient), and the principle gives no guidance regarding which should be chosen. Second, it is possible that a Pareto nonoptimal allocation could be regarded as preferable to a Pareto optimal one.
2. One remedy for these shortcomings is to adopt the criterion of potential Pareto efficiency. This considers whether, in switching from one allocation to another, it is possible for those who gain to fully compensate those who lose and still experience an improvement. If so, the second allocation is potential Pareto preferred to the first. An allocation is potential Pareto optimal if there is no other allocation potential Pareto preferred to it. One caveat should be borne in mind when thinking about the potential Pareto principle: the compensation it depends on is hypothetical and rarely offered in real life.
3. The various versions of the Pareto principle pertain to the static efficiency of an economy. Dynamic efficiency is about how successful an economy is at growing over time.
4. General equilibrium is a state in which all the markets that constitute an economy are in equilibrium simultaneously with respect to one another. They are linked because price and quantity outcomes in one market are normally factors that will affect outcomes in many other markets. In this way, an

economy can be considered a large, complex, highly interlinked system of individual markets. It is not obvious that such a system would in fact have an overall equilibrium.

5. The problem was framed in the late nineteenth century by Léon Walras, who represented the market system as a set of equations in which participants make offers to buy or sell. General equilibrium occurs when supply equals demand in the entire system of simultaneous equations. Technical difficulties in demonstrating this solution were overcome only in the twentieth century.
6. Given a number of supporting assumptions, theorists have established two fundamental theorems of welfare economics, that (1) a general equilibrium is always Pareto optimal, and that (2) any possible Pareto optimal state of the economy can be arrived at through a general equilibrium provided that the initial distribution of resources is properly adjusted.
7. In recent years there has been a large increase in the use of applied general equilibrium models. These reduce the number of markets and participants to mathematically convenient levels and, by basing equations on real-world data, calculate what general equilibrium ought to arise. Applied general equilibrium models are controversial, however: they require heroic assumptions to be mathematically tractable, and there is little evidence at this point that they improve forecasting.
8. The normative properties of general equilibrium are undermined to some extent by the General Theory of the Second Best. This holds that, if one component of an ideal allocation is violated, at least one other needs to be violated as well in order to arrive at a second-best solution. In practical terms, this implies that, even if a general equilibrium would be the best possible state for the economy, "closer" to this equilibrium is not necessarily better than "further" from it.
9. Modern research into general equilibrium theory has demonstrated that the process of arriving at an equilibrium can alter the equilibrium itself. This means that it is not possible to predict where the economy will end up without knowing its initial, out-of-equilibrium state and having detailed knowledge of the adjustments that will be made between the initial state and the ultimate equilibrium. In practice, this significantly reduces the predictive power of general equilibrium models, including the current generation of applied models.
10. The existence of a single general equilibrium for an economy depends on the assumption that its elements—individual participants and the goods and services they buy and sell—do not interact outside the bounds of the market-place. If there are such interaction effects, such as social, cultural, political or ecological interconnections, it is likely that there will be multiple potential equilibria. The implications are both positive and normative: it is more difficult to predict where the economy is headed, and additional information would be required to determine whether any particular equilibrium that eventuates is preferred to others that do not.
11. General equilibrium theory is the branch of economics that speaks most directly to the Invisible Hand hypothesis. On the basis of this theory we can

now identify with some precision the assumptions that must hold in order for this hypothesis to be vindicated. Realistically, it is extremely unlikely that all the needed assumptions will be in place for any actual economy, so the question is whether the deviations from Invisible Hand properties are significant enough to warrant corrective action.

▶ Terms to Define

Calibration
Computable general equilibrium
Dynamic efficiency
Fundamental Theorems of Welfare Economics
General equilibrium theory
General Theory of the Second Best
Multiple equilibria
Nonmarket interaction
Out-of-equilibrium trading
Pareto optimality
Potential Pareto optimality
Static efficiency
Utility possibility frontier

Questions to Consider

1. Can an economy with slavery be Pareto optimal? Can it be potentially Pareto optimal? In each case, to answer "yes", what exactly would have to be shown?
2. In this chapter it was asserted that general equilibrium theory presents the modern version of Adam Smith's Invisible Hand hypothesis. Is that entirely true? Does the claim that an ideal market economy has a single general equilibrium that adheres to the Two Fundamental Theorems of Welfare Economics encompass everything Smith meant by his invisible hand metaphor? Be specific, reviewing Smith's argument from earlier in this book.
3. Defenders of CGE modeling argue that the inability of such models to predict the actual effects of policy changes should not matter, since in real life many factors are always changing simultaneously, and not only particular economic policies model-builders are interested in. Do you agree? What should consumers of economic models—those who, like most of us, use them as potential sources of advice—expect them to offer?
4. In every developed country, and in many of the poorer countries as well, the government subsidizes agriculture by making payments to farmers. Economists sometimes criticize these subsidies on the grounds that the revenue farmers get should be determined by the willingness to pay of consumers for their products and nothing more—the logic of the Market Welfare Model. In their view, farmers are being encouraged to oversupply the market. Defenders of these subsidies might argue that they represent a second-best response to mispricing

in other aspects of the economy. Do you agree? If you do, what particular forms of mispricing are entailed, and how might they be offset by farm subsidies. How likely is it that the General Theory of the Second Best actually, and not only potentially, justifies existing agricultural policies?

5. As we saw in Chap. 5, one reason for the coffee crisis of the early 2000's was the rapid increase in supply during the 1990s. Since then, as the price of raw coffee beans plunged, many producers began ripping out their coffee trees and switching to other products; this should eventually cause prices to return to something like their previous level. Does this expansion and then decline of production reflect out-of-equilibrium production decisions by coffee producers? If so, could it have long-lasting effects on the equilibrium size and distribution of the global coffee sector?
6. Suppose workers, when deciding whether to take a job, are strongly influenced by how well that job pays in comparison to the other jobs they know about. Could this lead to more than one equilibrium wage in the labor market? Explain. What form of nonmarket interaction is involved?