CHAPTER 5

Market Power: Imperfect Competition and Strategic Behavior

5.1 Monopoly and Monopsony: When One Seller or Buyer Sets Total Quantity and Price

5.1.1 Motivation and Guiding Questions

The market diagrams used so far in this book have many buyers and sellers, but what happens when a single enterprise controls the market? Our individual-choice diagrams show what drives the size and scale of each individual enterprise, suggesting the possibility that one might grow large enough to be the only seller or the only buyer at some place and time.

We use the term *market power* to mean the potential ability of just one seller or buyer to control the entire quantity sold in a particular market. Agriculture and food systems are vulnerable to market power because manufacturing and distribution enterprises have much greater economies of size and scale than family farms and individual households. In many places around the world, whole communities have just one buyer or seller for some important goods and services. Why does market power arise? What outcomes can we expect from this kind of imperfect competition, and how might the resulting market failure be addressed through policy interventions?

Our economic analyses refer to individual markets, each showing a specific community or population interacting at one place and time. Every analytical diagram is drawn based on prior knowledge of that situation, which then determines how supply, demand and trade opportunities are specified. The term *monopoly* refers to markets with just one seller, and *monopsony* refers to markets with just one buyer. The two are symmetrical: both types of market power rely on being just one enterprise buying or selling in a community. As

W. A. Masters and A. B. Finaret, *Food Economics*,
Palgrave Studies in Agricultural Economics and Food Policy,
https://doi.org/10.1007/978-3-031-53840-7_5

we will see, opportunities to trade with others and thereby increase quantities can eliminate market power. The ability of one seller or buyer to control quantity depends on their own scale relative to the market, so market power can arise with just one enterprise in a small town, a larger company in a region or country or a multinational entity serving the whole world.

By the end of this section, you will be able to:

1. Describe how scale effects and innovation create opportunities for market power;
2. Derive marginal revenue curves from demand curves faced by a monopoly seller, to show what quantities they would choose to gain the highest possible level of profit;
3. Derive marginal expenditure curves from supply curves faced by a monopsony buyer, to show what quantities they would choose to gain the highest possible level of profit; and
4. Use diagrams to show how differences in elasticities of supply and demand affect the markup and profits obtained when using market power to restrict quantity.

5.1.2 Analytical Tools

The underlying source of market power is increasing returns to size or scale of individual enterprises discussed in Chapter 2. Increasing returns often involve lumpy or indivisible inputs, such as one person or one machine, which fit together with other people and machines in ways that benefit from close coordination within an enterprise. The result is a high fixed cost of setting up the enterprise relative to its marginal cost of expanding, leading to differences between that marginal cost and the enterprise's total cost of operation, and hence its average cost per unit bought or sold. When a single enterprise serves the entire market at lower average cost than if there were multiple enterprises, it is called a *natural monopoly*.

Natural monopolies arise where and when it is more efficient to concentrate production in a single enterprise, using one set of fixed costs to reach many customers at low marginal costs. As we will see, natural monopolies are often regulated as public utilities or provided directly by government as public goods addressed in Chapter 6. In this chapter, we focus on private enterprises, using economic principles to see how their choices affect their own revenue, expenditures and profits.

The scale effects that create market power involve equipment and personnel working together in a single enterprise, often using some kind of specialized knowledge or trade secrets. Every enterprise involves learning from experience, building skills and information over time. The spread of that knowledge is among the most important externalities in agriculture and food systems. Knowledge spillovers help other people adopt valuable innovations, and government funding can help discover and share the most helpful kinds of knowledge, but some innovations arise only through learning by doing within an enterprise. The inventions and specialized knowledge of private enterprises have long been protected by governments using a rights-based approach, using legal restrictions on how ideas can be used. These instruments include privacy protections and labor laws that protect trade secrets, as well as patents and trademarks that confer specific *intellectual property rights*.

Intellectual property is the glue that holds together many enterprises, providing 'intangible' assets that complement their equipment and personnel. Some enterprises hold patents, through which they disclose a specific invention that they can then prevent others from using for a fixed period of time, typically 20 years. Many more enterprises keep trade secrets that may not ever be disclosed, and use trademarks to establish a brand identity that can last for centuries. All intellectual property is a kind of fixed cost, allowing large enterprises to grow and prevent the entry of competitors who might erode their market power.

Relative Scale of Enterprises in Agriculture and the Food Sector

Food purchase decisions are made by individual households, and farming is predominantly a family enterprise, but scale effects often lead to a few large enterprises around them. The resulting hourglass shape in the number of enterprises is illustrated in Fig. 5.1.

The hourglass in Fig. 5.1 illustrates how there are often just a few input suppliers selling to many farm households, and those farm households then sell their output to a few enterprises that trade, transform and distribute food to consumers. The names listed are modern examples with global operations, but the diagram could be used to help understand local agriculture and food systems at any place and time.

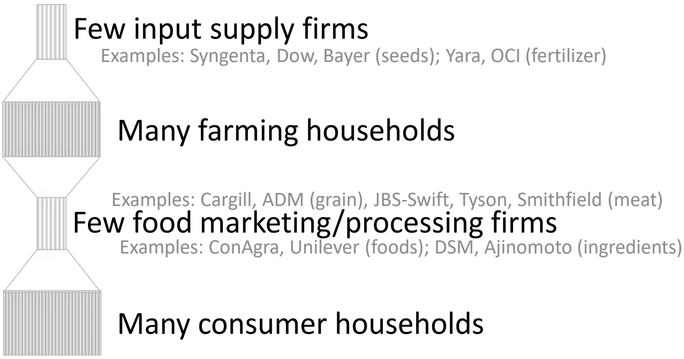

Fig. 5.1 Scale economies in agrifood systems create opportunities to exercise market power

Each enterprise that buys from or sells to farmers, which we can call an *agribusiness*, typically specializes in a specific kind of input in particular locations. The earliest agribusinesses in human history include grain mills, powered by water or wind and sometimes donkeys or horses walking in circles, grinding cereals into flour to serve dozens or hundreds of farmers in their vicinity. Other ancient kinds of agribusiness described in historical records include specialized makers and distributors of tools and equipment, and transport or storage providers in rural areas. Over time, enterprises grew to supply increasingly specialized seeds and other inputs. In each case, local farmers decide whether to do each thing within their own household or to buy that service from an agribusiness which might serve many farmers in their area.

Enterprises that serve consumers, which we might collectively call food businesses, have similar specializations. Food businesses operate at various scales. They often start small as family operations that grow and change as they discover sources of increasing returns and ways to expand. The names shown in Fig. 5.1 are food manufacturers like Unilever and ingredient makers like Ajinomoto, but retailers, restaurants and food service providers can also grow to enormous scale. Grocery chains and restaurants sometimes grow under a single brand name like Walmart or McDonald's, and sometimes grow as a conglomerate of multiple brands. Enterprises can grow through licensing as well as ownership, as for example Starbucks licenses its name and trade secrets to local operators and also directly manages some outlets for which it is both owner and operator.

The hourglass shape of Fig. 5.1, showing a small number of enterprises serving many farmers and many consumers, could be drawn at any geographic scale. Historically, small areas would be served by local enterprises, with agribusinesses serving a few dozen or hundreds of farmers, and food businesses serving hundreds of thousands of individual customers. Over time, increasing specialization and declining costs of transport has expanded the geographic scale of many enterprises. Whether their market is a small village or the entire world, one or more enterprises can potentially use their scale to exercise market power.

We use the term *monopoly* to describe a market with just one seller and the less common term *monopsony* when there is just one buyer. The two are symmetrical, so both kinds of market power are sometimes called monopoly power. But distinguishing between monopoly and monopsony is useful because food businesses can potentially exercise both at the same time. For example, a large dairy processor and distributor might become a monopsonist in buying raw milk from farmers and a monopolist in selling dairy products to consumers. Their potential market power is 'two-sided', similar to online platforms for food delivery that could potentially become the only intermediary between restaurants and customers. It is also possible for two large enterprises with market power to face each other, for example if an ingredient is made by just one seller and sold to just one food manufacturer, which would be a strategic interaction of the kind analyzed in the next section of this chapter. For now we turn to monopoly and then monopsony, showing how each can be understood using a similar kind of analytical diagram.

Monopoly Sellers, Marginal Revenue and Price Discrimination

To see how monopolies decide their quantities produced, we can go back to our toy model of the Alphabet Beach fish market. In this setting we know the names and details of each producer and consumer so can readily imagine what a monopolist would do using Fig. 5.2.

The stepwise supply and demand curves of Fig. 5.2 allow us to consider what would happen if Fio and Gio merged into a single enterprise. They might form a household that pools their resources, or be siblings in a family business, or just meet regularly to agree on what to do. Because this pooled Fio-Gio fishing enterprise controls set the entire quantity sold and earns all of the revenue from sales, their joint decisions differ from when Fio and Gio decided individually, when they did not take into account how their sales affected the other.

The earnings of the combined Fio-Gio enterprise from each unit sold are shown in the table on right of Fig. 5.2. In this initial scenario we consider the usual case in which the Fio-Gio enterprise cannot distinguish among buyers and prevent them from exchanging with each other. Each fish is identical so there is only one price, based on the community's marginal willingness to pay along the demand curve. For example if the monopolist sells just one fish, they can post a price of 9 and Ana will buy, but if they want to sell two fish they would have to reduce the price to 7 so that Bob will buy as well. The monopolists cannot prevent Ana from buying at the same price they offer to Bob, however, so the marginal revenue that a monopolist receives from additional sales is much less than the price received.

Monopolists like the Fio-Gio enterprise take account of the reduced price they get from a given customer like Ana when they decide to seek additional

In Alphabet Beach Village, if Fio and Gio merged into one fishing enterprise, what would they do?

Monopolists choose the total quantity produced, and typically sell at the same price to everyone.

A monopolist's total revenue is quantity times price. They maximize profits by expanding until *marginal revenue* from each additional unit sold falls below its marginal cost of production.
In this case, the unified Fio-Gio enterprise would stop at selling two fish and use only Fio's boat, since using Gio's boat to sell more fish yields marginal revenue (MR) below its marginal cost (MC).

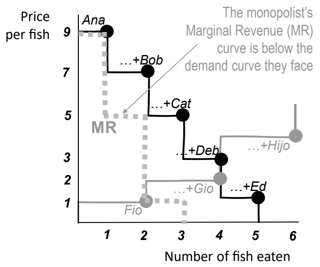

	Demand		A monopolist's perspective		Supply	
	Qty. sold	*WTP*	Total Revenue	Marginal rev. (MR)	*MC*	Producer
A	1	9	1x9 = 9	9-0 = 9	1	Fio
A+B	2	7	2x7 = 14	14-9 = 5	1	Fio
A+B+C	3	5	3x5 = 15	15-14 = 1	2	Fio+Gio
A+B+C+D	4	3	4x3 = 12	12-14 = -2	2	Fio+Gio
A+B+C+D+E	5	1	5x1 = 5	5-12 = -7	4	F+G+H

When Fio and Gio merge and control the whole market, adding the 3rd fish has a marginal cost of 2 but a marginal revenue of only 1, so they would prefer to sell only Fio's catch.

Fig. 5.2 Monopolists can earn excess profits by restricting production sales to customers like Bob. Selling one fish gave the Fio-Gio enterprise total revenue of 9, and selling two fish gave them total revenue of 14. The marginal revenue of the second fish was therefore 5. Raising quantity sold to three allows Cat to buy as well, but the price they can get falls to 5 and total revenue is 15, so the marginal revenue from their third fish is only 1. The marginal cost for the Fio-Gio enterprise to catch that third fish is 2. Monopolists who seek the highest level of total revenue minus total cost would produce only up to the quantity where marginal revenue is above marginal cost. If the Fio-Gio enterprise did catch a third fish, they would soon realize that was a mistake, and cut back to only two. They would use only Fio's fishing gear and share the resulting income.

The astonishing arithmetic of market power shows why a joint enterprise with both Fio and Gio would choose to produce less than if Fio and Gio worked independently. By merging with Fio, it is possible for Gio to make more by not fishing at all, as long as Fio shares the proceeds from the two fish they sell. The dynamics of their partnership is addressed in Section 5.2 where we introduce strategic interactions between two people. For now we focus on the unexpected logic of how and why monopolists sell less together than if they were separate enterprises along their supply curve.

To see market power graphically, we plot the incremental earnings from each fish on the seller's *marginal revenue* (MR) curve in Fig. 5.2. That curve is much steeper than the demand curve, and the monopolist's highest total income is where MR meets S. The marginal revenue curve is steep because each additional unit sold reduces the price received on all the items sold. Marginal revenue determines the income received by the monopolist but is not itself a demand curve. At the quantity selected by the monopolist where their marginal revenue meets or falls below their marginal cost, they can sell along the D curve at the consumer's willingness to pay.

The scenario shown in Fig. 5.2 is the baseline scenario for most monopolists, but only because they cannot distinguish well enough among buyers to charge each one a different price. Competitive sellers have no incentive or opportunity to differentiate among buyers, because they receive the entire price paid by the marginal buyer. Once an enterprise gains market power, however, they have a very strong incentive to find a way to sell at a higher price to buyers with a higher willingness to pay as shown in Fig. 5.3.

In the extreme benchmark case shown in Fig. 5.3, the Fio-Gio combined enterprise offers a differentiated fish to each buyer, and is somehow able to charge them the consumer's entire willingness to pay. One might imagine, for example, that Fio and Gio have prior knowledge that Ana is wealthy and would pay up to 9 for fish cut a certain way and delivered at a particular time, they might do that and sell one fish at 9. If they also knew that Bob would pay 7 for fish cut a different way, they might do that and thereby sell one at 9 to Ana and also one at 7 to Bob.

The result of charging each customer their entire willingness to pay is that marginal revenue equals demand (MR = D), and the monopolist can keep

Monopolists can sometimes charge higher prices to customers with greater demand

Successful price discrimination requires segmenting the market, charging each type of customer a different price based on their own willingness to pay.
Completely perfect price discrimination would allow monopolists to expand production to the perfectly competitive level, collecting all available consumer surplus as monopoly profits.

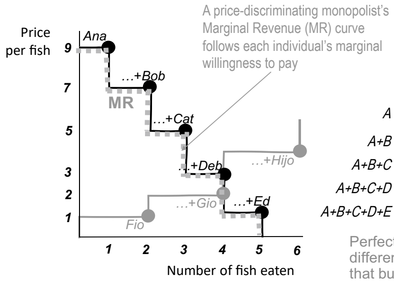

	Demand		A perfectly price-discriminating monopolist		Supply	
	Qty. sold	*WTP*	Total Revenue	Marginal rev. (MR)	*MC*	Producer
A	1	9	9 = 9	9-0 = 9	1	Fio
A+B	2	7	9+7 = 16	9-16 = 7	1	Fio
A+B+C	3	5	16+5 = 21	21-16 = 5	2	Fio+Gio
A+B+C+D	4	3	21+3 = 24	24-21 = 3	2	Fio+Gio
A+B+C+D+E	5	1	24+1 = 25	25-24 = 1	4	F+G+H

Perfect price discrimination requires selling the same fish for a different price to each buyer, *and* knowing exactly how much that buyer is willing and able to pay.

Fig. 5.3 Monopolists can earn even more excess profits through price discrimination

selling to each additional buyer until demand equals their marginal cost (MC) along their supply curve. On Alphabet Beach the Fio-Gio enterprise catches both of Gio's two fish, because the marginal cost of each is 2, and they can sell one to Cat for a price of 5 and another to Deb for a price of 3. This restores the perfectly competitive quantity of 4, but it is not worth expanding further. If Hijo were to join with Fio and Gio, their additional fifth fish would have a marginal cost of 4 but a maximum price of 1 from Ed.

The benchmark cases shown in Figs. 5.2 and 5.3 show the two mechanisms by which enterprises with market power can take advantage of becoming the only seller of something to a group of buyers. The first mechanism is quantity restriction, as they cut back on quantity sold to where MR meets S, and sell at the community's marginal WTP for that quantity along D. The second mechanism is price discrimination, as they try to sell each unit for that individual buyer's WTP, in which case they can sell a larger quantity out to where WTP meets S.

Total revenue for the Fio-Gio enterprise is shown in each table, and their total cost is readily seen by adding up marginal costs of each fish. With quantity restriction, the enterprise's total income is 12 (total revenue of 14 minus total cost of 2). With perfect price discrimination, by charging each buyer their entire willingness to pay, the enterprise's income is 19 (revenue of 25 minus total cost of 6). Both levels of total revenue for the Fio-Gio enterprise are far above their combined earnings prior to merging. When working independently, the competitive market led to a price between 2 and 3. Fio sold two fish and had producer surplus between 2 and 4 (total revenue of 4 to 6, minus total cost of 2), while Gio also sold two fish and had producer surplus between 0 and 2 (the same total revenue as Fio, but total cost of 4), so their combined revenue in the competitive market ranged from 2 to 6.

The results of Figs. 5.2 and 5.3 show clearly how every producer would like to be the only seller of their product for a particular market. In the Fio-Gio example, they go from combined earnings in the range of 2 to 6 when competing with each other, to joint earnings of 12 when they practice quantity restriction, and joint earnings up to a maximum of 19 when they achieve price discrimination. The field of *marketing* is devoted to understanding how companies can gain and exercise some degree of market power, which they call *pricing power,* and perhaps also achieve some degree of price discrimination. From an economics perspective, when companies become monopolies and restrict quantity, there is clear inefficiency because quantity is below the point where marginal costs just equal marginal benefits. If companies begin as a monopoly, their ability to price discriminate enables a larger quantity to be sold, although they also use that to take a larger share of the available consumer surplus.

Many businesses are able to achieve some degree of market power, for at least some of their products, in specific settings where they have few competitors. They would then have an opportunity to raise profits by restricting quantity, but an even stronger incentive to raise profits more through price discrimination. To see these decisions it was helpful to use our toy model of Alphabet Beach. For more general cases it is preferable to draw straight supply and demand curves in our stylized diagrams, which allow us to see the symmetrical case of monopsonies.

Monopsony Buyers and Marginal Expenditure

What if there is only one buyer, instead of only one seller? Markets with a single buyer are called a *monopsony*, and the buyer in a monopsony is called a *monopsonist*. As illustrated by the hourglass in Fig. 5.1, monopsony power can sometimes be exercised by agribusinesses that buy from farmers. This is especially common for products like raw milk that have significant scale economies in processing, and high transport costs for farmers to reach competing processors in other locations. Switching to stylized diagrams with straight lines for visual clarity, we can compare monopoly and monopsony in Fig. 5.4.

The left panel of Fig. 5.4 shows the same story as Fig. 5.2, but with linear MR and demand curves. The diagram shows how we can derive the exact MR curve from demand, with notation showing how one could use algebra and calculus to show that a linear demand curve leads to a linear MR curve whose slope is exactly twice that of the demand curve, as each additional unit sold reduces price received by the monopoly seller.

The right panel introduces the mirror image of MR, which is the marginal expenditure (ME) curve for price paid by the monopsony buyer. When the monopsonist buys each incremental unit along the sellers' supply curve, they raise the price they pay for the other units as well. In the case of a dairy monopsony, for example, they might be able to buy some raw milk from a few nearby farmers for a low price, but if they want to buy more they must offer a higher price to everyone.

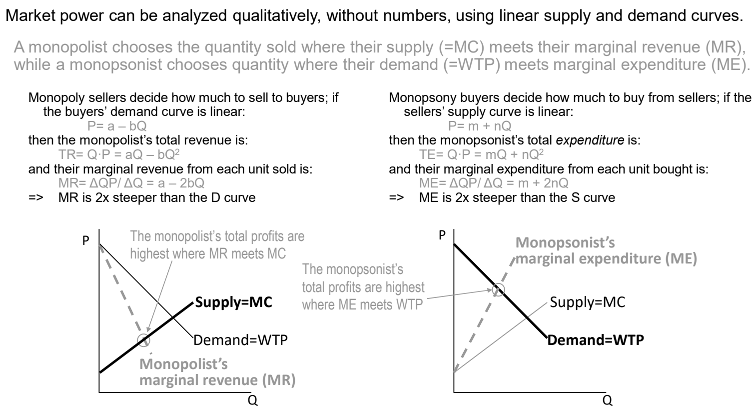

Fig. 5.4 Monopolies and monopsonies with simplified linear demand and supply curves

The circled points in Fig. 5.4 indicate where enterprises with market power stop adding additional units of quantity, because their S = MR for monopolists, and D = ME for monopsonists. The price at which they can sell that quantity, in the simple case without price discrimination, is shown in Fig. 5.5.

The symmetrical panels of Fig. 5.5 show how each kind of market power permits the charge or enterprise to earn higher profits than would be possible with a competitive market structure. Both show the simple case where only one price prevails, so the monopoly seller restricts quantity to Q_m so they can sell at P_m despite an additional unit costing only MC_m, and similarly the monopsony buyer restricts quantity to Q_m so they can buy at P_m despite having a willingness to pay for an additional unit of WTP_m. In both cases, if they were able to use price discrimination, they could increase quantity beyond Q_m and earn even more profits. Perfect price discrimination would potentially

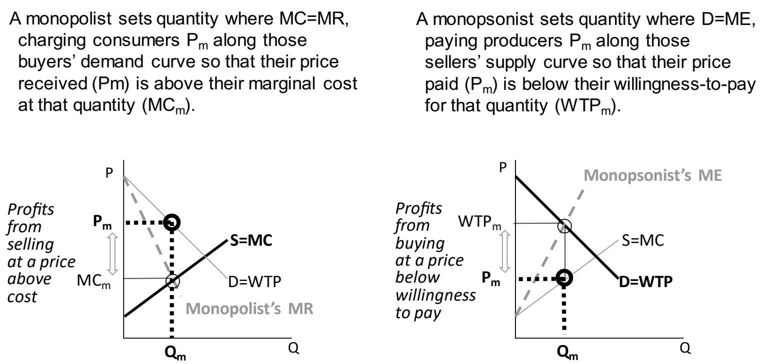

Fig. 5.5 Monopoly and monopsony both allow firms to raise profits by restricting quantity

allow them to sell all the way to where S meets D, capturing all of the profits shown for each unit.

Like all of our two-dimensional diagrams, this analysis of market power illustrates only general principles. In addition to the marginal costs shown here, a more complete analysis would take account of the fixed costs that create scale effects in the first place, and also take account of the complex detail around any particular case study. Before introducing two specific examples, it is helpful to add the areas of economic surplus gain or loss to the diagram.

Impacts of Market Power on Economic Surplus, Equity and Efficiency
Market power benefits enterprises that have it, at a cost to society. A monopoly seller's profits come at the expense of consumers along their demand curve, and a monopsony buyer's profits come at the expense of people who sell to it along their supply curve. We can see the relative magnitudes of these changes in Fig. 5.6.

The shaded areas and letters shown in Fig. 5.6 are gained and lost from market power relative to the perfectly competitive benchmark. In markets with scale economies, there is typically no actual policy instrument that could achieve perfect competition, but showing the effects of enterprises that restrict quantity in this way reveals what is at stake.

On each panel of Fig. 5.6, quantity restriction opens up area A that is gained by the enterprise at the expense of others, so that is purely an equity effect. In contrast, areas B and C measure the efficiency loss of producing less than this market's potential to generate economic surplus through additional units for which willingness to pay exceeds demand. The entire areas AB is lost by the population facing the monopolist or monopsonist. Area C is the additional loss of economic surplus when the enterprise cuts back on quantity. As in our previous analyses of market response, elasticities of S and D determine the

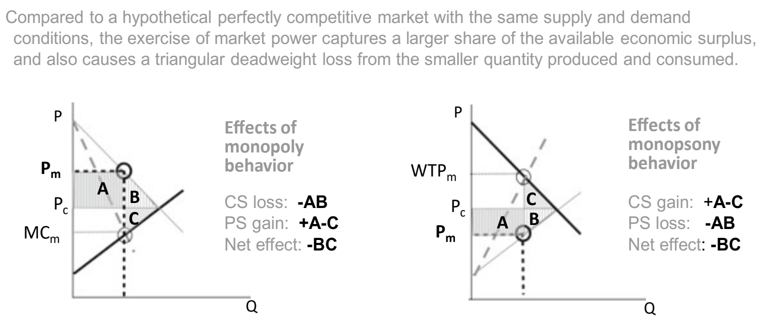

Fig. 5.6 Market power alters income distribution and also reduces total economic surplus

relative sizes of these areas, and especially the magnitude of deadweight loss BC relative to the equity effect A.

Market power can arise for a variety of reasons and might allow an enterprise to make high profits on some products in some locations, while other parts of the business are highly competitive. Each opportunity to be the only seller or buyer might be temporary, as others notice the high profits to be made and enter the market. Whether any particular enterprise actually has significant market power is difficult to determine, but it is useful to see two examples from recent U.S. history to illustrate specific aspects of how monopolies might arise and operate.

Market Power Can Be Obtained by Innovation: Walmart in the 1970s and 1980s

The first example is chosen to illustrate how an enterprise might gain economies of scale over time, using the example of Walmart as sketched in Fig. 5.7.

Walmart is a useful example because the roots of its initial success can be described in terms of a few familiar technologies that offered clear scale economies for retailing across the U.S. As shown on the left side of Fig. 5.7, Walmart was founded in 1962 grew into a chain at the start of the computer era, establishing one of the first interconnected systems of electronic inventory control in the 1970s. That network allowed inventories at all locations to be centrally monitored in real time, while competitors were still using much more expensive methods including periodic closure to physically count everything on the shelves. Walmart then became among the first users of several new

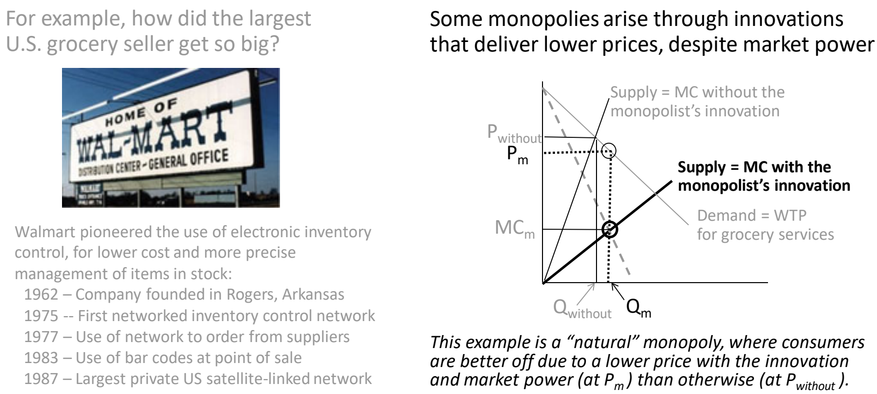

Fig. 5.7 Monopolies can arise from innovation, lowering costs through economies of scale *Source*: Timeline extracted from Jianfeng Wang [2006], "Economies of IT systems at Wal-Mart: an historical perspective." *Journal of Management Information and Decision Sciences,* 9[1]: 45–66

techniques for store management, each with high fixed cost but low marginal cost of expanding to new locations throughout the 1980s and 1990s.

The actual cost and pricing structure of any real business is enormously complex, but the basic principle of innovation and scale economies is drawn on the right of Fig. 5.7. One might imagine an initial competitive market of many small but expensively operated enterprises, operating at the high initial price and low total quantity. If an innovator successfully drops the marginal cost of supply low enough, it can be attractive to consumers even if it restricts quantity to Q_m and charges at P_m. In the case of Walmart, quantity restriction is seen in the way that its new stores were initially located relatively far apart across rural America. If Walmart were a public utility like the post office, they might have rolled out a larger number of stores closer to each other, as long as the marginal cost of each location was lower than willingness to pay and the enterprise could cover its fixed costs. Adding new locations would help customers reduce their travel time, but would have reduced the profitability of existing locations so Walmart had a smaller number of locations in the 1990s than its cost advantages might have allowed.

Market Power Can Be Obtained Legally, Including Through Protection from Trade

The second example is chosen to illustrate the potential role of government in allowing or preventing producers from joining together to operate as a monopoly. This example is particularly instructive because international trade is involved. Markets with trade usually cannot be monopolized, so creating market power in this case was possible only because the government could control trade in support of its efforts to help producers, as shown in Fig. 5.8.

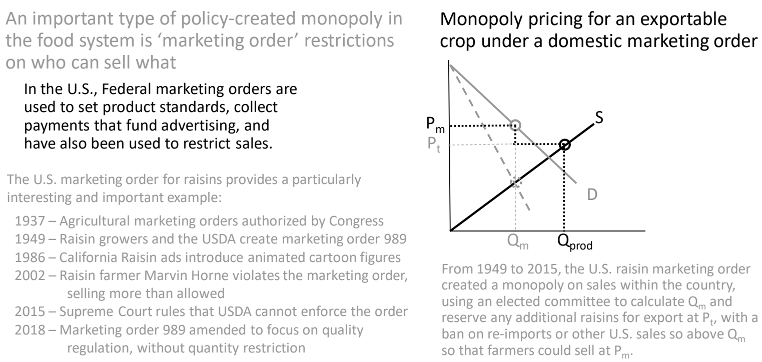

Fig. 5.8 Monopolies can arise from legal protections, as in a marketing board *Source*: Timeline adapted and extended from Dean L. Lueck [2016], "The curious case of Horne v. Department of Agriculture: good law, bad economics?" *NYU Journal of Law & Liberty*, 10: 608–625

The timeline sketched on the left side of Fig. 5.8 describes how the U.S. Department of Agriculture (USDA) is authorized by the law to help producers of a specific crop join together to regulate sales. These marketing orders allow registered growers to form an organization whose governing board is empowered to set standards and in some cases also limit quantities sold. In this case, from 1949 to 2015, marketing order number 989 allowed the raisin board to decide the total quantity allowed to be sold inside the U.S. each year. Each year the board would estimate demand, take account of production costs and attempt to find the quantity Q_m at which the price P_m would yield the highest total income for the farmers they represent, accounting for the fact that they could also export raisins at price P_t. The board also took into account fluctuations in supply and trade prices by managing storage, building up or drawing down their stockholding to provide additional control over Q_m to earn the highest possible farm income over time. Given the possibility of exports, farm income is shown in Fig. 5.8 as the entire producer surplus from the supply curve up to the U.S. price P_m for quantity Q_m, and then between the supply curve and the export price P_t for the quantity exported between Q_m and total production Q_{prod}.

For the USDA-supported raisin board to maintain higher prices inside the U.S. than elsewhere, for example across the border in Canada, they needed to restrict reimports of the quantity exported. That aspect of enforcement was administratively easy to accomplish, as import restrictions are a routine aspect of trade law. A more difficult challenge was to allocate shares of Q_m among farmers to sell at P_m, given that any additional quantities could be sold at the lower P_t. In practice, like many organizations in this situation, the raisin board allocated each farmer a share of Q_m based on their past production. If just one farmer were to sell more than their allotted share of Q_m at P_m, there might not be much decline in price along the demand curve. If multiple farmers did so, the price for all growers would eventually fall to P_t.

As shown in the timeline, raisin farmers generally obeyed the marketing order for many decades. Growers elected the board which decided Q_m and obtained P_m, using government regulation to prevent other farmers from entering which would have reduced the price. Over time, individual farmers might seek to increase their share of Q_m, and in 2002 one grower decided to do so on the grounds that a government-supported restriction on quantity sold was unacceptable to them. The case attracted the attention of people who wanted to limit government regulations in general, and they appealed the case all the way to the U.S. Supreme Court which ultimately ruled in favor of the farmer's right to sell as much as they wished. Thereafter the marketing board could no longer set quantities to raise prices, so its work is limited to quality standards and other functions.

The higher income earned by operating as a monopolist allows the group of farmers to behave as if it were a single enterprise, for example by advertising to promote the brand. The example of U.S. raisin farmers and their marketing board is famous in part because in the late 1980s, the board paid for an advertising campaign with cartoon figures known as California Raisins who formed a band playing popular songs. The fictional band's animated music videos were wildly successful, and although actual raisin sales did not rise enough to justify continuing the campaign in the 1990s, the idea of cartoon raisins remains vivid in American popular culture.

The economic aspects of the raisin board's story is worth telling in this book for many reasons. First, there is the human drama of organizing people for any collective purpose, because each individual then has an incentive to break away and take advantage of others having followed the rule. We will return to that in the next section of this chapter. Second, there is the way that our raisin example shows the role of trade restriction in making market power possible. Third, there are important aspects of the story involving human health and government decision-making, as policies adopted for one purpose can work against other interests, sometimes in ways that may remain unknown even to well-informed people but could be revealed by economic analysis.

The nutrition and health aspect of the marketing board story is important because raisins (among other fruits) are promoted in the USDA's own Dietary Guidelines for Americans. During much of the period shown in Fig. 5.8, the nutrition services of USDA were actively promoting fruit consumption for health reasons, even as the quantity sold was actively being restricted by the marketing arm of the USDA. Different political forces drive the two arms despite them being housed in the same agency. Even if higher-level decision-makers in government were aware that one arm of the USDA was restricting sales even as another arm sought to increase them, there might have been little they could do about that contradiction. The political balance of forces driving each policy was in a kind of equilibrium between the government's diverse constituencies, and there was little reason for anyone to devote the time and effort it might take to alter the outcome.

For economics generally, an important aspect of the marketing board story helps us understand that actions by individuals and groups have unintended consequences, and that economic analysis reveals those effects without needing to know anything about what people are thinking, or how they use what they earn in pursuit of their own objectives. In writing this section of the book we do not have or need any particular knowledge about the motivations of the raisin farmer and his supporters who financed the lawsuit that ended the board's quantity restrictions. They may have believed that government restrictions were harming them, or they may have been willing to sacrifice future earnings in pursuit of other goals. Economic analysis is useful only to show what decisions provide the largest total gains relative to costs in a particular setting, recognizing that each person can and will have multiple motivations for what to do with the income they might earn.

Finally, the raisin story brings us back to the analysis of market power, and whether the board's most important concerns actually involved quantity restriction at all. As we have seen, an even more valuable source of pricing power would have been price discrimination. Even before the marketing board was formed, raisin farmers had formed a cooperative called Sun-Maid to provide joint marketing services, one aspect of which was to differentiate branded raisins from the same food in generic packaging. Then in the late 1980s a major focus of the marketing board was to invest in advertising for all kinds of raisins. Building a generic California Raisins campaign might have aimed to shift the demand curve outward to raise Q_m, but it could also have aimed to make the demand curve steeper so as to reach a higher P_m even with no change in quantity.

Profits from Market Power Depend on Price Elasticities

The role of consumers' demand elasticities in allowing monopolists to charge high prices is illustrated in Fig. 5.9.

The two panels of Fig. 5.9 show two monopolists with identical supply curves and an identical quantity sold. The two monopolists differ only in the demand curves they face, and consumers' elasticities of response to their choice of quantity sold. Comparing the two figures shows how the steeper, more inelastic curve on the left offers greater potential pricing power. With linear curves the slope of each MR curve is exactly twice the slope the corresponding demand curve, so the difference between the two MR slopes is exactly twice as large as the difference between the two demand curves. When consumers have relatively inelastic demand on the left, the monopolist can charge P_m and earn a much larger markup over their marginal costs MC_m than the otherwise identical monopolist on the right who has the same monopoly position but faces more elastic demand, and hence lower profits based on the smaller gap between P_m' and MC_m'.

The comparison shown in Fig. 5.9 helps explain why enterprises invest heavily in trying to become monopolists for things whose demand is always price-inelastic, and also helps explain why enterprises with some market power

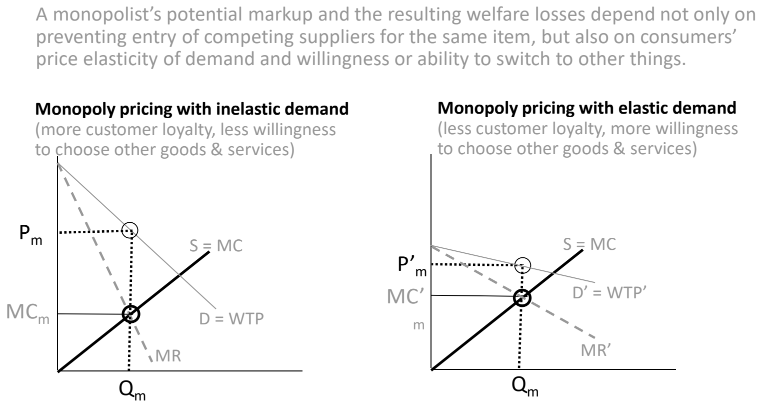

Fig. 5.9 Inelastic demand raises a monopolist's pricing power

often seek to make demand for their products as inelastic as possible by focusing their advertising and other business practices on brand loyalty and repeat purchases. In contrast, advertising that is targeted towards expanding quantities sold is more widely used in more competitive settings, and can make elastic demand curves even flatter by attracting purchasers who have more other options.

Figure 5.9 can also be used to see the demand curve along which a price-discriminating monopolist could charge, if they successfully differentiated their product to sell at high prices for consumers with high willingness to pay. More inelastic demand raises the potential profits from price discrimination, creating strong incentives for enterprises with some degree of market power to find ways of selling otherwise similar items to different people at different prices. For example, a food manufacturer could sell the same product under a premium brand with expensive packaging and advertising, while simultaneously selling it under a generic label at a lower price. Similarly, a grocery store could charge more to online shoppers who value convenience, and restaurants charge higher margins on alcohol and prestige items for diners who are willing to pay for that.

The economic mechanisms by which enterprises with market power can sometimes earn high profits also reveal how competition can work to bring those profits back down, as challengers see opportunities to enter and compete in newly profitable market segments. Product differentiation can attract other enterprises specializing in premium brands, leading to segmented markets for each kind of product. With market segmentation, enterprises aiming for the high value-added segment compete along their supply curve with high costs of marketing, packaging and other services, while enterprises aiming for the high-volume segment compete with low prices of that same product in generic form. Every market is defined spatially as well, as a monopolist's quantity restrictions in one location will attract a larger volume of sales from competitors elsewhere, creating geographic patterns of spatial competition. And competition also occurs over time, with the pricing power of monopolists is limited not only by existing competitors at each place and time, but also by the threat of future entry. Even longstanding monopolies might eventually be disrupted, and some monopolies are in contestable markets with few barriers to entry so they must behave competitively to deter competition that would displace them entirely.

Measuring Market Power

Each market diagram shows quantity and price for a specific product quality, at a particular place and time. When applying these models to any real-life situation, economists must specify the extent of the market being analyzed in terms of the product characteristics, time period and population whose supply, demand and trade opportunities are shown in the model.

In this textbook we show economic principles graphically in two dimensions, so our models in this section are limited to monopoly and monopsony.

By definition, real-life examples of just one buyer or seller arise only when the market is defined narrowly around one enterprise's specific product, place and time. For example, we might draw the market for groceries in a given neighborhood, helping to explain how a single big supermarket might behave differently from many small shops when serving the same population with a given demand curve. Studies can sometimes measure market power at that level of granularity, but the available data usually defines markets more broadly to include a whole sector or segment of the food system, for example as the number of different grocery chains that might potentially compete with each other over a given region.

The number of enterprises serving a market segment or sector is often reported directly, but enterprises differ greatly in size. For example, a given city might be served by one to three superstores or chains, and then a larger number of small and medium-sized enterprises. The likelihood that an enterprise can exercise market power in any part of the sector or segment that it serves could depend on its geographic scale or range of products offered. For example, if one large grocery chain serves a whole region, it might have monopoly power only in a few neighborhoods or product categories where smaller chains and independent shops cannot compete. Instead of counting enterprises, analysts typically use data on volume sold to compare market shares.

The market share of an enterprise is its fraction of sales. For sufficiently uniform products this can be defined in quantity terms, such as a dairy processor's share of all raw milk sold in a state each year. Every gallon of raw milk is similar enough in quality that volume could be measured in weight (pounds or kilograms) or volume (gallons or liters). In other markets, different enterprises sell a variety of differentiated products so their volume is measured by the value of sales in monetary terms. Having defined a market category, for example all dairy products, analysts compare the value of sales by each enterprise to the sum of sales by all enterprises in the market as a whole.

Market shares are often expressed as *concentration ratios*, focusing on the few largest enterprises that might be able to monopolize some part of the market. The largest market share is the C1 ratio, and sum of shares over the two or three largest enterprises would be the market's C2 or C3 ratio. A typical approach is to focus on C4. For example, in Britain during the 2010s the four largest supermarket chains by market share were Tesco (31%), Sainsbury's (17%), Asda (17%) and Morrisons (13%), for a combined C4 share around 78% in 2011. Then the expansion into Britain of two small-format chains from Germany, Aldi and Lidl, and also a new online-only retailer, Ocado, reduced the shares of all top-four retailers, bringing the C4 ratio to 66% in 2023. This kind of data typically comes from private firms that specialize in marketing strategy such as Kantar or Nielsen.

Market power can potentially be exercised within market segments, so economists are often interested in degrees of concentration across the entire distribution of enterprises. For example, two markets may have the same C4 ratio, but very different degrees of concentration among the top four, and potentially also different concentration among the other smaller enterprises. To capture that aspect of concentration, analysts can use the sum of squared market shares which is the same method used in environmental sciences to measure lack of biodiversity over a whole population of organisms. Among economists, the sum of squared market shares for each firm is known as the *Hirschmann-Herfindahl index,* while measuring biodiversity using the sum of squared shares of each species in a population is known as the *Simpson index*. With just a single monopolist this index has a value of 1, and increasing competition or diversity drives the index towards zero. When all have equal shares, then the sum-of-squares index simply returns that share. For example, with two equal shares, the squared value of each share is $0.25 = 0.5 \times 0.5$, and the sum-of-squares returns $0.50 = 0.25 + 0.25$. From that baseline, increasing concentration raises the index. For example, 80–20 shares have an index of 0.68, and 90–10 shares have an index of 0.82 which is beginning to approach the monopoly status of a single seller. The magnitude of a Hirschmann-Herfindahl index depends on whether shares are reported in decimal form or as percentage points (0.25 or 25) and also depends on the number of enterprises included in the index, so values may be rescaled for use in different contexts.

Measuring concentration is only a first step to inform policy response. Whether concentration actually leads to the exercise of market power and profits at the expense of consumers or other enterprises depends on all the factors shown in our models, such as elasticities of response. Our diagrams could be drawn in two dimensions, for example using the MR and ME curves to identify quantity sold, because they focus on a single monopolist or monopsonist facing a market of many others who adjust along their demand or supply curves. The next section of this chapter shows interactions between just two decision-makers, with the resulting outcomes shown in a table of payoffs from each choice. More advanced game theory considers an even wider range of possible interactions among two or more actors, with each kind of interaction corresponding to a different market structure in the field of economics known as industrial organization, with great relevance to agricultural input supply and food businesses.

Policies to Address Market Power

A first kind of policy concerns mergers or acquisitions, which is the initial example used in this section when Fio and Gio joined together to raise price. That scenario involved no innovations to reduce cost. The only source of market power was the agreement between Fio and Gio to merge operations and raise both of their incomes by either quantity restriction (ending Gio's catch and sharing the profits) or price discrimination (selling each unit at the buyer's willingness to pay).

In the U.S., rules against otherwise legal businesses gaining market power are known as *antitrust* law, because they were introduced in the late nineteenth century specifically against merged businesses that were then known as trusts. In other countries, similar legislation is known as competition policy. Many antitrust efforts aim to limit mergers or break up existing enterprises, through administrative review and legal proceedings to assess whether larger size operations would generate sufficient cost savings to offset the dangers of market power. The primary focus of legal cases is usually whether larger enterprises can manipulate their own prices and quantity. Antitrust policy can also be used to address whether large enterprises actively stop others from competing with them, for example by preventing workers from switching employers. Criminal law also plays a role in antitrust policy, through rules that prohibit enterprises from making agreements with each other to manipulate prices or quantities and limit competition. That kind of price-fixing through cartels is a criminal offense in many sectors, but antitrust regulations are commonly waived for organizations designed to help farmers such as cooperatives and marketing boards.

A second kind of policy concerns the flow of innovations that might affect market power and enterprise scale. Innovations often have high fixed costs for invention and adoption, but then low marginal cost to deploy over each unit of production, thereby introducing a new source of increasing returns that reaches lowest total cost at a larger scale of operations than earlier enterprises. Some innovations allow many small enterprises to be formed, such as online platforms or shared kitchens and co-packers that help individuals start new food businesses, but then the facilitating platform itself could begin to exercise market power against those businesses. Policies to encourage innovation also introduce some kinds of market power deliberately, using patents to give inventors a temporary monopoly over their invention in exchange for disclosing it, as well as other protections to encourage research and discovery within private enterprises. These factors make market power dynamic and temporary, where the best remedy against market power by one enterprise may be to encourage formation of other companies using different techniques at different scales over time.

A third category of actions to address market power involve the institutions, infrastructure and policies that influence whether enterprises can be insulated from competition. Market power comes from enterprise scale relative to the extent of each market. One aspect of market size is geographic area. The evolution of food systems often involves a transition from competition among enterprises for local market power (for example, a country might have a hundred dairy processors but only two or three in each place, seeking monopsony power when buying milk from local farmers and monopoly power when selling to local consumers), to competition among larger enterprises serving a greater geographic extent (for example, two or three enterprises competing with each other over the entire country). Enterprises with local market power are often challenged by entrants from neighboring places, in ways that can be helped or hindered by government action.

A fourth category of policies about market power concerns product differentiation and demand for higher-quality products. If the only way for an enterprise to signal quality is their own brand identity, then they will have to invest heavily in marketing, packaging, advertising and other ways to convince people that their product actually has the desired quality attributes. Price itself can be a signal of quality, if people expect that low prices imply low quality, and expect that high prices and high incomes provide the seller an incentive to maintain high quality over time. Both marketing costs and price as a signal of quality make high-quality products unnecessarily expensive, especially when there are scale effects and inelastic demand that give a monopolist some pricing power on top of all their actual high costs of maintaining product quality.

Product standards enforced by governments and private associations have been an important aspect of food systems since the earliest historical records, driven by the fundamental problem that people can actually observe only a few aspects of food quality such as color, odor and taste. Some of the first recorded food standards in European history focused on preventing use of nonfood ingredients in bread and beer that would increase their weight or volume, soon followed by rules to maintain food safety of products such as milk and meat. The minimum quality regulations for all foods sold were soon complemented by quality standards to differentiate higher-priced versions of similar products, such as the first pressing of olive oil at mills that also extract lower-quality oil. Labeling then allows consumers to see what they could not otherwise detect for themselves, making it possible for markets to sustain competition for high-quality products.

Introducing and enforcing quality standards can help new entrants compete with established enterprises, lowering the cost to consumers of items at or above each level of quality. Establishing new standards is politically challenging in part because established businesses that already meet the standard do so with brand identities and high prices signaling their own high quality, while other businesses might need to incur significant added costs to meet the standard. By definition the attributes that need to be signaled cannot be seen or experienced directly, allowing critics to sow doubt about the scientific basis for each standard.

Establishment of organic product standards is a particularly important example in the U.S. and many countries, aiming to create a larger and more competitive market for items that meet those requirements. Introducing a separate standard for organic products leaves open the question of quality standards for conventional products. Each market segment, the policy options for quality signaling range from disclosure requirements such as the back-of-package nutrition facts panel, to requiring more visible marks such as front-of-pack warning labels, and direct regulation of product contents and advertising such as bans on harmful ingredients or rules against deceptive marketing.

5.1.3 Conclusion

Market power can potentially be used in ways that raise an enterprise's own profits at the expense of the society as a whole, in ways that can be remedied by policies to facilitate entry of competitors meeting sufficiently high-quality standards. This section shows how barriers to entry allow existing enterprises to increase their own profits in non-competitive ways, either restricting quantity to earn higher margins on the uniform product at the same price or by product differentiation and price discrimination. Exercising market power can be done by enterprises that are the only seller (a monopoly) or the only buyer (a monopsony). In either case, exercising market power yields higher profits when quantity response is more inelastic. When monopolies or monopsonies exist, an effective way to limit the resulting harm is to make response more price-elastic by ensuring that people have other options.

The fundamental source of market power in the food system is scale economies in activities other than farming and household food consumption, such as farm input provision or food manufacturing and distribution. Agricultural production is done mostly by owner-operated family farms, and food choices are made by individuals and households, while agribusinesses and food enterprises can often gain market power through larger scale relative to the extent of the market they serve. Markets are defined in terms of a specific product, place and time period, and power over that market can be exercised only as long as it remains protected from competitors through high barriers to entry. New competitors, if they arise, could begin as startups within the monopolized market, but when scale economies make total cost lower at larger volumes, competitors can also enter and compete across a wider range of products and larger geographic areas. One of the most important sources of competition and elasticity of quantities to limit market power is trade with other places, limiting ability of existing enterprises in any one place to profit at the expense of other people in their own society.

Once an enterprise has market power, they can use it to earn additional profits through either quantity restriction or product differentiation and price discrimination. In the food system some products have uniform quality, often because a government agency or voluntary organization has set and enforced the specifications for that type of product. Crops such as wheat or rice are generic commodities bought and sold competitively, but each bag or truckload can be substituted for another only to the extent that they all meet a written standard published by the government or some other organization. In the absence of such standards, each enterprise must try to signal quality through expensive branding and high prices, often allowing price discrimination by large enterprises or the emergence of segmented markets whose costs of signaling quality could be reduced through uniform standards. Any food attribute that cannot be seen or experienced by end-users, including many aspects of nutrition and health, contributes to market failure that can potentially be remedied by quality standards and certification for more competitive markets.

5.2 Strategic Behavior: Game Theory for Two-Person Interactions

5.2.1 Motivation and Guiding Questions

The choices we have seen so far explain and predict outcomes by looking at each individual's options, identifying the actions that would best achieve their goals. Can we generalize these insights to interactions between two people, where each takes the other's decision-making into account? What determines whether two people will cooperate with each other towards a common goal? Can we predict what circumstances might lead them to stop cooperating, or even harm the other person?

The choices we address in this section are *strategic* behavior, used in the same sense as '*strategic*' moves in chess or other settings. The predicted equilibrium outcome of such interactions requires a kind of analysis known as *game theory*, used in economics to address industrial organization and other topics where just two or a few actors interact with each other. In this book we introduce strategic behavior in its simplest context, which is a single interaction between two individuals each choosing from the same two options. That simplicity allows us to draw a *payoff matrix* of four possible outcomes, and shows how modeled choices respond to incentives in ways that are surprisingly informative about real-life interactions.

Our analysis of game theory, like the rest of this book, focuses on general principles applied to the food system, using analyses that can be presented graphically with almost no formal mathematics. For strategic interactions, we focus on the simplest possible framework, reduced to just two decision-makers choosing between two options. Understanding how and why people in different circumstances would choose to cooperate in a joint activity, or go it alone against the other person's interests, is deeply revealing about human behavior in general and especially in the food system.

By the end of this section, you will be able to:

1. Describe how game theory and Nash equilibrium are used to explain and predict the outcome of strategic interactions;
2. Use a 2×2 payoff matrix to identify the predicted outcome of an interaction between two people choosing between two options;
3. Describe how altering the payoff matrix influences behavior, as in the prisoner's dilemma created by police to elicit confessions, or a positive-sum game that elicits cooperation; and

4. Describe how social norms and commitment mechanisms can alter payoffs and influence outcomes, with examples such as climate policy and natural resources in agriculture.

5.2.2 Analytical Tools

Previous sections have explained outcomes in terms of each person's choice, with everyone else's choices shown by their willingness to pay along the demand curve, marginal cost along the supply curve or opportunities for trade with others. In perfectly competitive markets, each individual is a *price taker,* adjusting their quantity whether or not that alters prices paid or received by others. In a monopoly or monopsony, one individual sets the entire quantity and can be seen as a *price maker,* choosing whatever prices will meet their own goals. In this section, we ask what if an individual faces just one other individual and can take into account that both of them are making choices?

As we will see, game theory models of strategic behavior have many concrete applications to the real world with clear relevance to food systems. Game theory yields useful insights at many different scales, from relationships within a household to bargaining between enterprises and country governments. To help us follow how human decision-making leads to the outcomes we see, it is helpful to start with our toy model of Alphabet Beach, and then introduce a variety of other scenarios.

When we first saw the perfectly competitive model of Alphabet Beach fish market, an additional seller such as Gio entered as long as the additional buyers such as Cat and Deb could cover their costs of production. Gio's entry directly benefited Cat and Deb who got to eat Gio's fish and also lowered costs for Ana and Bob. The price reduction caused by competitive entry came at the expense of the other seller, Fio, leading us to model a scenario where Fio persuaded Gio to join together in a merged Fio-Gio enterprise. The result was a monopoly with a clear interest in either ending Gio's catch entirely to sell at a single higher price to both Ana and Bob, or differentiating among buyers to sell at a higher price to Ana than to Bob and also to Cat and Deb.

In this section we can look more closely at the relationship between Fio and Gio. They have a strong interest in collaborating, but what circumstances make them more likely to agree on a shared strategy and act as one, and what would make them more likely to compete with each other? The actual example of Fio and Gio is not well-suited to introduce game theory, because they have an unequal starting place. Fio is the low-cost supplier who will choose to catch fish whether or not Gio goes fishing. In the toy model, that gives Gio a strategic advantage in bargaining with Fio: Gio can credibly threaten to catch and sell fish which would reduce Fio's price, whereas Fio has no similar way to threaten Gio's income. Also we know that Fio earns more from fishing than Gio, but that might be because Gio has better options for other employment, further strengthening Gio's strategic advantage in their negotiation over a potential partnership.

Game theory models with asymmetric payoffs could potentially be used to predict the outcome of bargaining between Fio and Gio. We might discover, for example, that Fio is likely to voluntarily give Gio a majority vote in decisions and larger share of profits from the Fio-Gio enterprise, because that is the only way for Fio to ensure that Gio remains in their joint enterprise. To see how bargaining unfolds, however, it is far easier to start with the case of symmetrical bargaining, in which the two actors face the same choice between two options. Modeling strategic interactions between two equal partners ensures that we are focusing on when and how cooperation emerges spontaneously, without one person being forced by circumstances to accept the other's conditions.

The Payoff Matrix for a Symmetric Two-Person Interaction

The payoff matrices we analyze in this textbook are the simplest case, with the same two choices being made by two people, for whom the potential outcomes can all be arrayed in a two-by-two matrix. To distinguish between rows and columns we will call the two people X and Y, and each chooses yes or no. The consequences of each choice for the X person are shown in rows, and consequences for the Y person are shown in columns, forming the payoff matrix in Table 5.1.

In each cell of the payoff matrix there will be two numbers, separated by a backward slash \, denoting that the first number is the payoff to person X from their choice in that row, and the second number is the payoff to person Y from their choice in that column. Using the backward slash can be a useful reminder of which payoff goes to which person, as person X's choice is labeled on the left side of each row, and person Y's choice is listed at the top of each column.

In Table 5.1, person X's choices are shown in two rows, and Y's choices are in two columns. In the first row if X says yes, Y might say yes or no, and in the second row if X says no, person Y might say yes or no. In this simplest version of their interaction, there is only one time period, the two people choose simultaneously, and each person knows that both of them face identical payoffs. This setup is valuable because it isolates the core question of how each person's choices are influenced by the payoff to each outcome.

Table 5.1 Example of variables in a payoff matrix

		Person Y	
		Says yes	*Says no*
Person X	Says yes	X_{yes} \ Y_{yes}	X_{yes} \ Y_{no}
	Says no	X_{no} \ Y_{yes}	X_{no} \ Y_{no}

The economic principles used to predict choices in this context are known as *Nash equilibrium*, named after the American mathematician John Nash who characterized the problem as part of his PhD dissertation in 1950. Nash's insight was that even if each person chooses simultaneously, we can imagine that they have learned from experience in other contexts and want to avoid choices they might regret. The resulting 'no regret' equilibrium techniques then have very wide applicability to many other problems, and allows us to see how the payoff matrix drives the outcomes we are likely to observe as each person decides on their best choice given what they know about the other person's options.

Payoffs and Predicted Outcomes of the Prisoner's Dilemma

The idea of Nash equilibrium in a two-by-two payoff matrix is often described as a *prisoner's dilemma*, because that is how the interaction was first explained by mathematicians and economists in the 1950s. They chose this example in part for its realism, because it helps explain how detectives learned to solve crimes and how those accused of crimes have learned to respond, in situations described by Table 5.2.

The prisoner's dilemma shows a situation where the police have set the penalties and rewards for each action in a way that helps them solve crimes quickly. The dilemma they create for each prisoner starts with arresting two suspects who the policy believe might have been involved in a crime, and placing the suspects in separate cells. Each is offered the same options: a favorable outcome if they confess and explain the crime, or a heavy penalty if they deny involvement and are convicted. The payoff matrix in Table 5.2 shows the two prisoners' options. If both deny involvement then the police have no evidence, and neither can be convicted so they both walk free. If both confess their penalty might be −2, but longstanding police practice is to make a favorable offer such as +1 for the first to confess, and a harsh penalty such as −3 to suspects who deny involvement and are later convicted.

By setting these penalties, the police have created a dilemma by which each prisoner knows they would both be better-off if neither confessed, but each prisoner has no way of ensuring that the other does not choose to confess. John Nash provided the algorithm to solve this and many other game theory problems by identifying the best option for X depending on what Y does in each column, and the best option for Y in each row, then ruling out options that would be regretted.

Table 5.2 The payoff matrix in a prisoner's dilemma is designed to elicit confessions

		Suspect Y	
		Confess	*Deny*
Suspect X	Confess	−2 \ −2	+1 \ −3
	Deny	−3 \ +1	0 \ 0

With the example payoffs in Table 5.2, person X knows that if Y confesses, their own choices in the first column are between −2 and −3, and if Y denies involvement their choices are between +1 and 0. The payoffs created by the police thereby ensure that X has a clear incentive to confess no matter what Y does. The situation is symmetrical, so suspect Y is choosing between −2 and −3 if suspect X confesses, or between +1 and 0 if suspect X denies involvement. Both suspects have been given an incentive to confess, no matter what the other does. Police set payoffs in this way to make it more likely that prisoners will confess, because it is the only Nash equilibrium outcome of the situation.

The payoffs in Table 5.2 are just the smallest whole numbers in the simplest scenario needed to illustrate the idea of Nash equilibrium in the prisoner's dilemma context. In reality, this kind of interaction gets repeated many times, and the penalties or rewards actually offered vary widely depending on the context. Potential prisoners who might be detained and rewarded for confessing will evolve a strong norm of silence in these situations, including severe retaliation against anyone who they think might have cooperated with the police. Countless stories could be told about how people try to alter the payoff matrix to elicit the behaviors they want, and some of the best examples come from real-life situations in the food system.

Price Fixing in the Global Lysine Market: Ajinomoto and ADM in the 1990s

A famous historical case for which our two-by-two symmetrical payoff matrix provides helpful insights occurred in the 1990s, when the leaders of Archer Daniels Midland (ADM) in the U.S. and Ajinomoto in Japan decided jointly to restrict quantity and raise prices for lysine, an amino acid that they manufactured in large volumes as an ingredient for animal feed around the world, and also citric acid, an important ingredient in soft drinks and other products. These are standardized commodities so no price discrimination was involved. ADM and Ajinomoto were the two leading global suppliers, and a few other companies also had significant market share but were not in a position to quickly increase output if prices rose.

In October of 1996, ADM was convicted of colluding with Ajinomoto to limit quantity sold in the U.S., leading to a large fine and ultimately also prison sentences for three officials of the company when they were found to have agreed on how much each would sell at what price in the U.S. and elsewhere. Price fixing was determined to have begun in June 1992 for lysine and January 1993 for citric acid, and ended when the scheme was exposed in June 1995. The details of the case are fascinating, in part because it was unusually well-documented by an informant who was himself also convicted of defrauding the company. The story clearly reveals how incentives to limit quantity can tempt company managers into illegal activity as illustrated in the payoff matrix of Table 5.3.

The payoffs in the matrix in Table 5.3 are very roughly scaled to potential revenue gains from restricting supply of lysine and citric acid to the

Table 5.3 The hypothetical payoff matrix for two participants in a price-fixing conspiracy

		Company Y	
		Compete	*Restrict supply*
Company X	Compete	50 \ 50	200 \ 0
	Restrict supply	0 \ 200	150 \ 150

U.S. market, in millions of dollars per year during the 1990s. In this case 'compete' means to produce additional lysine until price received just meets marginal cost, while 'restrict' means to hold back on sales to where their jointly estimated marginal revenue meets their marginal cost.

Solving for the Nash equilibrium is done in the same way for these payoffs as for the prisoner's dilemma. For company X, the first column shows their payoffs if company Y chooses to compete, so their options are 50 or 0 and it is better for them to compete. Likewise if company Y chooses to restrict, the options for company X are 200 or 150 and again it is better for them to compete. This example is symmetric so company Y faces the same choice. For both companies, it would be better to compete than to restrict production, unless it is possible for company leaders to agree that they will both restrict supply.

In the historical case of ADM and Ajinomoto, it was ADM leadership that first contacted Ajinomoto and persuaded them to cut back on sales. In this and similar conspiracies, there were large profits to be made from jointly restricting supply, but also temptation to violate the agreement and return to competition, taking advantage of the other having temporarily restricted supply. Because these conspiracies are illegal, the agreement to restrict supply can only be enforced privately. For example, X might persuade Y to restrict supply by threatening to sell at a loss until Y is forced into bankruptcy. The two companies might also credibly guarantee that they will both restrict supply by inviting observers from the other company into their factories.

In the case of ADM and Ajinomoto, the conspiracy was undone by an employee of ADM who rebelled against his own company. Governments provide large incentives to individual informants who are willing to come forward, because company leadership could potentially sustain their conspiracy for many years. Eventually the payoff matrix would change, for example due to the entry of other companies, making it less profitable to maintain the criminal conspiracy, or one company might violate the agreement for other reasons such as a change of personnel. Then the market might return to competition without the need for antitrust action, but in the meantime consumers would have suffered great losses. In the ADM-Ajinomoto case, the reduced quantity of lysine led to a loss of efficiency due to slower livestock growth around the world.

Influence of the Payoff Matrix on Cooperative Behavior

The classic examples of a prisoner's dilemma and a price-fixing conspiracy might give readers the impression that the payoff matrix to strategic interactions always or often leads to unfortunate outcomes. In fact those two examples are used precisely because of the drama involved. Everyday interactions often involve payoff matrices that reward positive or pro-social acts of collaboration and cooperation. These are sometimes called positive-sum games, because the sum of values from acting together exceeds the value of not doing so, as illustrated in Table 5.4.

The example of Table 5.4 shows how a parent with two children might offer incentives that encourage collaboration between siblings, for example by setting up games that reward nicer play. In this scenario, the payoffs to child X are higher if they play nicely no matter how child Y responds, and similarly for child Y, the payoffs are higher if they play nicely no matter how child X responds. In a payoff matrix like Table 5.4, collaboration can be sustained without the need for external enforcement, as long as each child understands the situation and realizes that playing nicely is their best choice from the available options.

Many examples of self-sustaining cooperation arise every day, in workplaces and public interactions where people help each other simply because pro-social behavior is their best option, whether or not other people return the favor. The concept of Nash equilibrium helps us understand how a payoff matrix like that of Table 5.4 leads each person to play nicely even if the other acts selfishly. In experiments where people are given incentives of this type people occasionally deviate from the predicted Nash equilibrium, because they misunderstood their options or just had a bad day, but most people return to cooperative pro-social behavior once they realize it is preferable for them to do so.

The influence of payoffs on expected outcomes can be seen by anticipating the consequences of a small reduction in the payoff to playing nicely when the other is selfish. That change can tip the equilibrium away from pro-social behavior, for example when payoffs are as shown in Table 5.5.

The only change from Tables 5.4 to 5.5 is the reduction from 6 to 4 in the payoff to playing nicely when the other child is selfish. That small difference creates a situation where if Y plays nicely, X would also want to play nicely, while if Y is selfish, X would also want to be selfish. In this situation there are two Nash equilibria, and each might be equally likely. An equilibrium where both play nicely, and both children experience their highest payoff,

Table 5.4 A payoff matrix for self-sustaining collaboration

		Child Y	
		Play nicely	*Be selfish*
Child X	Play nicely	8 \ 8	6 \ 7
	Be selfish	7 \ 6	5 \ 5

Table 5.5 A payoff matrix with two Nash equilibria

		Child Y	
		Play nicely	*Be selfish*
Child X	Play nicely	8 \ 8	4 \ 7
	Be selfish	7 \ 4	5 \ 5

would arise when each understands and expects the other to do the same. But if one expects the other to be selfish, then children might choose to protect themselves so both would act selfishly.

Parents and others who can influence the payoffs might be able to shift incentives in ways that reward cooperation, but changes in the payoff matrix can also tip the balance away from cooperation as in Table 5.6.

The change illustrated by Table 5.6 is a lower payoff in the top-left corner. When the rewards to both playing nicely declines from 8 to 6, the payoff matrix is such that when Y plays nicely there is an incentive for X to be selfish, and when Y is selfish there is also an incentive for X to be selfish. That makes it likely that each would choose to protect themselves and act selfishly. With a payoff matrix of this type, some pairs of children might play nicely and experience the payoff in the top-left corner, but the situation is such that each will be tempted to act selfishly no matter what the other does, leading them both to the less favorable anti-social outcome in the lower-right corner.

Applying economic principles to strategic behavior using a payoff matrix reveals surprising truths about human behavior. We are used to thinking about our own choices as having been the best we could do under the circumstances. Economic analysis allows us to think that way about other peoples' behavior, revealing how incentives might be altered to improve outcomes. Understanding behavior as a Nash equilibrium is especially helpful for agriculture and the food system, where real-life payoff matrices are heavily influenced by nature and technology, limiting our options but also providing new insights about the causes of each outcome we observe.

When people work together in farming, fishing, hunting, cooking or other food-related activities, they are working in nature, using tools that reward cooperation to differing degrees. For example plowing with oxen requires two people, irrigation from a river requires upstream and downstream farmers to agree on water use and kitchen operations may be suited to working together

Table 5.6 A payoff matrix that discourages cooperation

		Child Y	
		Play nicely	*Be selfish*
Child X	Play nicely	6 \ 6	4 \ 7
	Be selfish	7 \ 4	5 \ 5

or alone. Nature and technology shape behavior in ways that have been richly documented by anthropologists, economic historians and other observers, but the examples in this book also show how payoffs can potentially be modified by other people who influence the setup of each interaction, such as the police who set penalties to elicit confessions in the prisoner's dilemma, or parents who set up games that encourage their children to play nicely together.

The strategic behavior elicited by each kind of interaction often becomes a habit or a cultural norm, especially when the same type of payoff matrix appears repeatedly in a person's life. It is impossible and also unnecessary for people to use John Nash's algorithm for everyday decisions. In real life people just learn from experience, including both our own experiences and the experiences of other people as communicated to us in stories and advice. Using economic principles to explain these habits and cultural norms is useful to help us change, by understanding why each kind of behavior arises and how we might alter incentives to elicit different behaviors in the future.

Repeated Games, Commitment Mechanisms and Incomplete Contracts
Most interactions do not take place just once, but occur in the context of repeated opportunities to cooperate or compete. In the food system, producers and consumers are interacting with a small number of other people every year near the location where they live and work. Farmers and family members need to help each other to survive, and workers everywhere need to collaborate for their enterprise to thrive.

In settings where pairs or groups of people have repeated interactions, patterns emerge that differ slightly from the simple two-by-two example. One important finding is the emergence of intertemporal *commitment mechanisms*. It can be extremely valuable for people to make credible commitments that they will in fact do something in the future, whatever the circumstances when the time comes. Farmers in small, isolated communities often commit to helping each other in the event of hardship, and share many things as a way of demonstrating their commitment. Food consumers can also benefit from advance commitments, for example by subscribing to an entire season of regular food deliveries from a community-supported farm, instead of buying only the items they want each week, as a way of ensuring that the farmer can start the season in confidence.

A related aspect of repeated games is the prevalence of *incomplete contracts*. When we first learn how incentives affect behavior, it is tempting to think that offering additional incentives for each specific kind of effort is usually helpful, but in the food system and other sectors many important relationships are left vague. Important contracts such as land rental agreements for farmers or employment contracts for restaurant workers, are often little more than a handshake and a price. These contracts are 'incomplete' in the sense that they do not specify much about what will be done in exchange for payment. Economists have often run experiments to test whether additional incentives such as pay-for-performance contracts yield better results and sometimes they do, but actual businesses typically revert to incomplete contracts when the experiment ends. One reason is that spelling out everything needed for a successful outcome is very difficult, and over time people find it is preferable to use a strategy that relies on self-motivation. Farmland owners and tenant farmers choose contracts that reward mutual trust, as do restaurant managers and employees, and incomplete contracts are helpful for that purpose.

Finally, a common finding of research on repeated interactions is known as the *folk theorem*, so called because it emerged as a common understanding in the field before studies demonstrated its general validity. The folk theorem states that repeated interactions tend to elicit more pro-social cooperative behavior than interactions that are limited in time. Versions of the folk theorem have been shown using game theory, and similar results have been found in experiments and field studies. The basic mechanism behind the folk theorem is that repeated interactions can offer greater rewards to cooperation and stronger options for retaliation against anti-social behavior. The prospect of repeated interactions is not always sufficient to elicit cooperation, in part because real-life interactions are not actually infinite in duration, but this insight helps explain how and why the duration of relationships is related to their outcomes.

Multi-person Games and the Tragedy of the Commons

The symmetric two-by-two payoff matrices shown in this book are the simplest kind of game with which to model strategic interactions. Extensions to asymmetric bargaining and a very wide range of special cases have been explored and solved mathematically for their Nash equilibria, and used to test how different people respond to incentives in experimental settings or real-life observations. One of the most difficult and important kinds of extension is towards multi-person games, such as all people in a household, every worker in a restaurant or ultimately all people in society.

Early explorations of multi-person games used computer simulations, for example evolutionary models for a population of individuals each of whom is of a fixed type that always plays the same strategy. These simulations can be repeated to assess how different strategies perform under diverse conditions, simulating the process of natural selection among different types of individuals. Simulations of this type are also known as agent-based models, because each person is an 'agent' in the sense of playing out a fixed strategy. The development of these simulation models traces the history of computing, towards increasingly complex models and also convenient user interfaces. Insights from such models are sometimes used in economics, but they differ from standard economics in that each individual's strategy is predetermined. In other words, the agents in evolutionary models lack the 'agency' we associate with actual people who make their own choice among multiple options.

Solving for a Nash equilibrium in settings where multiple people take account of each person's responses to other people's choices can be mathematically impossible. To describe different kinds of situations, game theorists have developed a large toolkit of various specifications designed around specific forms of interaction. In each model, introducing just a few additional actors or different options can be sufficient to yield predictions that are similar to the outcomes of a competitive market with many participants. Our toy model of Alphabet Beach fish market had only eight people in it, but the interactions between them give insights into outcomes for a village of eight hundred, a country of eight million or a planet of eight billion people.

An especially famous and important kind of multi-person interaction for agriculture and the environment is known as the *Tragedy of the Commons*, after the title of a 1968 essay in *Science* magazine by the American biologist Garrett Hardin. In that essay, Hardin uses the example of herders whose animals graze on public pastures, using the term 'commons' to mean land open to all in a community. Hardin explained how each herder would gain the full benefit of putting one more animal on the commons, while experiencing only a fraction of the cost imposed when the animal eats plants that would otherwise keep growing and be available for others to graze. Mathematically, each herder can be seen as gaining +1 for the value of each animal they add, while experiencing costs of n/N where n is the number of animals they own, out of the community's total of N animals on the commons. Hardin's essay spells out the human tragedy as each herder chooses to add one more animal to their own n, despite imposing a cost of n/N on every other member of their community, relentlessly driving down the available commons until no grass is available for anyone.

The tragedy of the commons that prompted Hardin's essay was human population growth, which was the subject of widespread concern in the 1960s. Garrett's essay ended with a call to restrict people's 'freedom to breed'. Similar ideas contributed to a campaign of forced sterilization in India during 1976–77, and China's one-child policy introduced in 1979. Garrett's arguments about population control were popular at the time but have since been discredited. As we will see in Chapter 10, human demographics did not turn out to be a fixed path to tragedy or to require forced reductions in fertility, in part because incentives changed leading to smaller family size as women gained more schooling and employment opportunities outside the home as well as other changes associated with economic development.

The most important tragedy of the commons today is undoubtedly greenhouse gas emissions and climate change. Warnings were issued but not heeded for decades, in part because emissions cause a global externality where each individual, company or country bears only a small fraction of its cost to all of humanity in future years. Innovations that sharply reduce the cost of generating and storing renewable energy now offer a path to rapidly decreasing use of fossil fuels, but only if existing equipment is rapidly replaced and new installations adopt the lower-cost new technologies. Changing incentives plays a large role in that energy transition, along with changes in net emissions of carbon and methane from agriculture or other aspects of sustainability and health in the global food system.

5.2.3 Conclusion

This section of the book introduced game theory and its applications to the economics of agriculture, food and nutrition. The toolkit of game theory shows how strategic behavior emerges and is sustained in different settings, based on the payoffs to each action. Using this framework, observed outcomes can sometimes be predicted and explained as an economic equilibrium between people.

In a strategic interaction, equilibrium behavior is the set of actions by each person that would be the best of their options, no matter what the others decide. This is an equilibrium in the sense that each person would not regret their choice. Many empirical studies have shown that people do indeed often choose the predicted strategies in both experimental and real-life settings.

The central question addressed in this section is how pro-social, cooperative behaviors that lead to the most efficient use of available resources can be sustained voluntarily. In many interactions people might act selfishly, missing out on opportunities for more favorable outcomes available if people collaborate. Joint efforts are often needed in agriculture and the food system, for example among farmers who need to cooperate when using a shared irrigation system, or restaurant workers who need to cooperate in the kitchen.

The examples in this section focus on choices where two people select between two options. That is the simplest possible kind of strategic interaction, allowing us to see how the equilibrium depends on relative payoffs to each action. When people see that their interactions will have payoffs similar to those of the prisoner's dilemma or a tragedy of the commons, they can anticipate the outcome and alter the terms of interaction to promote collaboration. The main finding of this section is that changes to the payoff matrix can lead to different behaviors. In real life, actions become habits and norms that people experience as personality and culture. To the extent that those behaviors evolved in response to past incentives, a change in payoffs can lead people to learn from experience and eventually adopt different habits, adding up to different norms for their entire community.

The example of changing incentives used in this section is a parent who oversees different games played by their children. Some games have relative payoffs whose equilibrium is for each child to play nicely, and some have relative payoffs whose equilibrium is for each child to be selfish. Parents who want children to play nicely often provide games that encourage cooperative play, and avoid games that encourage selfishness. We also saw the example of a game in which both equilibria are equally possible, in which case parents can set norms that nudge children towards nicer play.

The use of two-by-two games in this section, like our toy model of Alphabet Beach and all of the analytical diagrams in other chapters, aims to provide a toolkit of stylized models in which economic principles play out differently in different contexts. Economists can then choose the most suitable model for each situation, based on prior knowledge or research about which model would be the best fit to explain, predict and guide choices.

In real-world applications to the food system, specifying each economic model relies on contextual knowledge of how nature, technology and society determine the available options for agriculture and health. The next chapter completes our introduction to the economics toolkit by developing the main models used to understand changes in government policy, then the second half of this book turns to empirical data about those facts.