

Chapter 17

Network Effects

At the beginning of Chapter 16, we discussed two fundamentally different reasons why individuals might imitate the behavior of others. One reason was based on *informational effects*: since the behavior of other people conveys information about what they know, observing this behavior and copying it (even against the evidence of one's own private information) can sometimes be a rational decision. This was our focus in Chapter 16. The other reason was based on *direct-benefit effects*, also called *network effects*: for some kinds of decisions, you incur an explicit benefit when you align your behavior with the behavior of others. This is what we will consider in this chapter.

A natural setting where network effects arise is in the adoption of technologies for which interaction or compatibility with others is important. For example, when the fax machine was first introduced as a product, its value to a potential consumer depended on how many others were also using the same technology. The value of a social-networking or media-sharing site exhibits the same properties: it's valuable to the extent that other people are using it as well. Similarly, a computer operating system can be more useful if many other people are using it: even if the primary purpose of the operating system itself is not to interact with others, an operating system with more users will tend to have a larger amount of software written for it, and will use file formats (e.g. for documents, images, and movies) that more people can easily read.

Network Effects as Externalities. The effects we are describing here are called *positive externalities*. An *externality* is any situation in which the welfare of an individual is affected by the actions of other individuals, without a mutually agreed-upon compensation. For example, the benefit to you from a social networking site is directly related to the total number of people who use the site. When someone else joins the site, they have increased your

welfare even though no explicit compensation accounts for this. This is an externality, and it is *positive* in the sense that your welfare increases. In this chapter, we will be considering the consequences of positive externalities due to network effects. In the settings we analyze here, payoffs depend on the number of others who use a good and not on the details of how they are connected. In Chapter 19, we will look at the details of network connectivity and ask how they affect the positive externalities that result.

Notice that we have also seen examples of *negative externalities* earlier in the book — these are cases where an externality causes a decrease in welfare. Traffic congestion as discussed in Chapter 8 is an example in which your use of a (transportation or communication) network decreases the payoff to other users of the network, again despite the lack of compensation among the affected parties. In the final section of this chapter, we will look at a direct comparison of positive and negative externalities in more detail.

It's important, also, to note that not everything is an externality — the key part is that the effect has to be *uncompensated*. For example, if you drink a can of Diet Coke then there is one less can of Diet Coke for the rest of the world to consume, so you decrease the welfare of others by your action. But in this case, in order to drink the can of Diet Coke you have to pay for it, and if you pay what it costs to make another can of Diet Coke, then you have exactly compensated the rest of the world for your action. That is, there is no uncompensated effect, and hence no externality. We explore the interaction of externalities and compensation further when we discuss property rights in Chapter 24.

17.1 The Economy Without Network Effects

Our canonical setting in this chapter will be the market for a good: we will first consider how the market functions when there is no network effect — that is, when consumers do not care how many other users of the good there are — and then we will see how things change when a network effect is present.

We want to analyze markets with a huge number of potential purchasers, each of whom is small enough relative to the entire market that he or she can make individual decisions without affecting the aggregate behavior. For example, each individual considering the purchase of a loaf of bread does so without worrying about whether her individual decision — all else remaining the same — will affect the price of bread. (Note that this is different from worrying about whether decisions made by a large number of people will have an effect, which they certainly can.) Of course, in real markets the number of consumers is finite, and each individual decision does have a very, very small effect on the aggregate. But each purchaser's impact is so small relative to the market that we can model individuals as not taking this into account when they make a decision.

Formally, we model the lack of individual effects on the aggregate by representing the

consumers as the set of all real numbers in the interval strictly between 0 and 1. That is, each consumer is named by a different real number, and the total mass of consumers is 1. This naming of the consumers by real numbers will be notationally useful — for example, the set of consumers with names between 0 and $x < 1$ represents an x fraction of the population. A good way to think of this model of consumers is as a continuous approximation to a market with a very large, but finite, number of consumers; the continuous model will be useful in various places to avoid having to deal with the explicit effect of any one individual on the overall population.

Each consumer wants at most one unit of the good; each consumer has a personal intrinsic interest in obtaining the good that can vary from one consumer to another. When there are no network effects at work, we model a consumer's willingness to pay as being determined entirely by this intrinsic interest. When there are network effects, a consumer's willingness to pay is determined by two things:

- intrinsic interest; and
- the number of other people using the good — the larger the user population, the more she is willing to pay.

Our study of network effects here can be viewed as an analysis of how things change once this second factor comes into play.

To start understanding this issue, we first consider how a market works when there are no network effects.

Reservation Prices. With no network effects, each consumer's interest in the good is specified by a single *reservation price*: the maximum amount she is willing to pay for one unit of the good. We'll assume that the individuals are arranged in the interval between 0 and 1 in order of decreasing reservation price, so that if consumer x has a higher reservation price than consumer y , then $x < y$. Let $r(x)$ denote the reservation price of consumer x . For the analysis in this chapter, we will assume that this function $r(\cdot)$ is continuous, and that no two consumers have exactly the same reservation price — so the function $r(\cdot)$ is strictly decreasing as it ranges over the interval from 0 to 1.

Suppose that the *market price* for a unit of the good is p : everyone who wants to buy the good can buy it at price p , and no units are offered for sale at a price above or below p . At price p , everyone whose reservation price is at least p will actually buy the good, and everyone whose reservation price is below p will not buy it. Clearly at a price of $r(0)$ or more, no one will buy the good; and at a price of $r(1)$ or less, everyone will buy the good. So let's consider the interesting region for the price p , when it lies strictly between $r(1)$ and $r(0)$. In this region, there is some unique number x with the property that $r(x) = p$: as




Figure 17.1: When there are no network efforts, the demand for a product at a fixed market price p can be found by locating the point where the curve $y = r(x)$ intersects the horizontal line $y = p$.

Figure 17.1 illustrates, since $r(\cdot)$ is a continuous function that strictly decreases, it must cross the horizontal line $y = p$ somewhere.

This means that all consumers between 0 and x buy the product, and all consumers above x don't — so an x fraction of the population buys the product. We can do this for every price p : there is an x depending on p that specifies the fraction of the population that will purchase at price p . This way of reading the relation between price and quantity (for any price the quantity that will be demanded) is usually called the (market) *demand* for the good, and it is a very useful way to think of the relation between the price and the number of units purchased.¹

The Equilibrium Quantity of the Good. Let's suppose that this good can be produced at a constant cost of p^* per unit, and that, as is the case for consumers, there are many potential producers of the good so that none of them is large enough to be able to influence the market price of the good. Then, in aggregate, the producers will be willing to supply any amount of the good at a price of p^* per unit, and none of the good at any price below p^* . Moreover the assumption of a large number of potential producers who can create new copies

¹In the language of microeconomics, the function $r(\cdot)$ describes the *inverse demand function*. The inverse of $r(\cdot)$, giving x in terms of p , is the *demand function*.

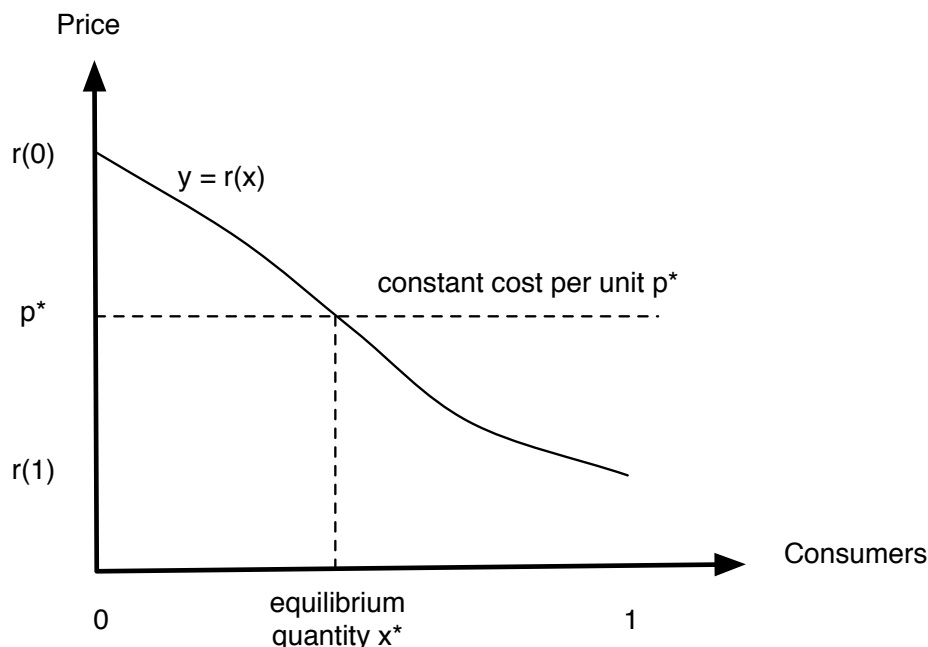


Figure 17.2: When copies of a good can be produced at a constant cost p^* per unit, the equilibrium quantity consumed will be the number x^* for which $r(x^*) = p^*$.

of the good at a constant cost of p^* implies that the price cannot remain above p^* , since any profit to a producer would be driven to zero by competition from other producers. Thus, we can assume a market price of p^* , regardless of the number of units of the good produced.² As above, cases in which p^* is above $r(0)$ or below $r(1)$ are not particularly interesting, since then either everyone or no one buys the good. Therefore, we assume that $r(0) > p^* > r(1)$.

To complete the picture of how the market operates without network effects, we now determine the *supply* of the good. Since p^* is between the highest and lowest reservation prices, we can find a unique x^* between 0 and 1 so that $r(x^*) = p^*$. We call x^* the *equilibrium quantity* of the good, given the reservation prices and the cost p^* . Figure 17.2 revisits Figure 17.1, including the cost p^* and the equilibrium quantity x^* .

Notice the sense in which x^* represents an equilibrium in the population's consumption of the good. If less than an x^* fraction of the population purchased the good, there would be consumers who have not purchased but who would have an incentive to do so, because of reservation prices above p^* . In other words, there would be “upward pressure” on the consumption of the product, since there is a portion of the population that would not have purchased but wished they had. On the other hand, if more than an x^* fraction of the

²Continuing with the microeconomic language, this is the long-run competitive supply for any good produced by a constant-cost industry.

population purchased the good, there would be consumers who had purchased the good but wished they had not, because of reservation prices below p^* . In this case, we'd have “downward pressure” on the consumption of the good.

One attractive feature of this equilibrium is that it is socially optimal (as defined in Chapter 6). To see why, let's consider the social welfare of the allocation, which we can think of as the difference between the total reservation prices of the consumers who receive a copy of the good and the total cost of producing the corresponding quantity of the good. Now, if society were going to produce enough of the good to give it to an x fraction of the population, then social welfare would be maximized by giving it to all consumers between 0 and x , since they correspond to the x fraction of the population that values the good the most. Which value of x would be the best choice? Since the contribution of a consumer x' to the social welfare is the difference $r(x') - p^*$, we can think of the social welfare, when consumers 0 through x get copies of the good, as the (signed) area between the curve $y = r(x)$ and the horizontal line $y = p^*$. It's signed in the sense that portions of the curve $y = r(x)$ that drop below $y = p^*$ contribute negatively to the area. Given this, we'd want to choose x so that we collect all the positive area between $y = r(x)$ and $y = p^*$, and none of the negative area. This is achieved by choosing x to be the equilibrium x^* . Hence the equilibrium quantity x^* is socially optimal.

We now introduce network effects; we'll see that this causes several important features of the market to change in fundamental ways.

17.2 The Economy with Network Effects

In this section, we discuss a model for network effects in the market for a good. We will follow a general approach suggested by Katz, Shapiro, and Varian [235, 368]; see also the writings of Brian Arthur [25, 27] for influential early discussions of these ideas.

With network effects, a potential purchaser takes into account both her own reservation price and the total number of users of the good. A simple way to model this is to say that there are now two functions at work: when a z fraction of the population is using the good, the reservation price of consumer x is equal to $r(x)f(z)$, where $r(x)$ as before is the intrinsic interest of consumer x in the good, and $f(z)$ measures the benefit to each consumer from having a z fraction of the population use the good. This new function $f(z)$ is increasing in z : it controls how much more valuable a product is when more people are using it. The multiplicative form $r(x)f(z)$ for reservation prices means that those who place a greater intrinsic value on the good benefit more from an increase in the fraction of the population using the good than do those who place a smaller intrinsic value on the good.

For now, in keeping with the motivation from communication technology and social media, we will assume that $f(0) = 0$: if no one has purchased the good no one is willing to

pay anything for the good. In Section 17.6 we will consider versions of the model where $f(0)$ is not 0. We will also assume that f is a continuous function. Finally, to make the discussion a bit simpler, we will assume that $r(1) = 0$. This means that as we consider consumers x tending to 1 (the part of the population least interested in purchasing), their willingness to pay is converging to 0.³

Since a consumer's willingness to pay depends on the fraction of the population using the good, each consumer needs to predict what this fraction will be in order to evaluate whether to purchase. Suppose that the price of the good is p^* , and that consumer x expects a z fraction of the population will use the good. Then x will want to purchase provided that $r(x)f(z) \geq p^*$.

We begin by considering what happens in the case when all consumers make perfect predictions about the number of users of the good; after this, we will then consider the population-level dynamics that are caused by imperfect predictions.

Equilibria with Network Effects. What do we have in mind, in the context of the current discussion, when we suppose that consumers' predictions are perfect? We mean that the consumers form a shared expectation that the fraction of the population using the product is z , and if each of them then makes a purchasing decision based on this expectation, then the fraction of people who actually purchase is in fact z . We call this a *self-fulfilling expectations equilibrium* for the quantity of purchasers z : if everyone expects that a z fraction of the population will purchase the product, then this expectation is in turn fulfilled by people's behavior.

Let's consider what such an equilibrium value of z looks like, in terms of the price $p^* > 0$. First of all, if everyone expects a $z = 0$ fraction of the population to purchase, then the reservation price of each consumer x is $r(x)f(0) = 0$, which is below p^* . Hence no one will want to purchase, and the shared expectation of $z = 0$ has been fulfilled.

Now let's consider a value of z strictly between 0 and 1. If exactly a z fraction of the population purchases the good, which set of individuals does this correspond to? Clearly if consumer x' purchases the good and $x < x'$, then consumer x will as well. Therefore, the set of purchasers will be precisely the set of consumers between 0 and z . What is the price p^* at which exactly these consumers want to purchase, and no one else? The lowest reservation price in this set will belong to consumer z , who — because of the shared expectation that a z fraction of the population will purchase — has a reservation price of $r(z)f(z)$. In order for exactly this set of consumers, and no one else, to purchase the good, we must have $p^* = r(z)f(z)$.

We can summarize this as follows:

³The assumption that $r(1) = 0$ isn't necessary for our qualitative results, but it avoids various additional steps later on.




Figure 17.3: Suppose there are network effects and $f(0) = 0$, so that the good has no value to people when no one is using it. In this case, there can be multiple self-fulfilling expectations equilibria: at $z = 0$, and also at the points where the curve $r(z)f(z)$ crosses the horizontal line at height p^* .

If the price $p^ > 0$ together with the quantity z (strictly between 0 and 1) form a self-fulfilling expectations equilibrium, then $p^* = r(z)f(z)$.*

This highlights a clear contrast with the model of the previous section, in which network effects were not present. There, we saw that in order to have more of the good sold, the price has to be lowered — or equivalently, at high prices the number of units of the good that can be sold is smaller. This follows directly from the fact that the equilibrium quantity x^* without network effects is governed by $p^* = r(x^*)$, and $r(x)$ is decreasing in x . The market for a good with network effects is more complicated, since the amount of the good demanded by consumers depends on how much they expect to be demanded — this leads to the more complex equation $p^* = r(z)f(z)$ for the equilibrium quantity z . Under our assumption that $f(0) = 0$, we've seen that one equilibrium with network effects occurs at price p^* and $z = 0$: Producers are willing to supply a zero quantity of the good, and since no one expects the good to be used, none of it is demanded either.

A Concrete Example. To find whether other equilibria exist, we need to know the form of the functions $r(\cdot)$ and $f(\cdot)$ in order to analyze the equation $p^* = r(z)f(z)$. To show how this works, let's consider a concrete example in which $r(x) = 1 - x$ and $f(z) = z$. In this case, $r(z)f(z) = z(1 - z)$, which has a parabolic shape as shown in Figure 17.3: it is 0 at $z = 0$ and $z = 1$, and it has a maximum at $z = \frac{1}{2}$, when it takes the value $\frac{1}{4}$. Of course, in

general the functions $r(\cdot)$ and $f(\cdot)$ need not look exactly like this example, but typically we expect to see something like the shape displayed in Figure 17.3.

Continuing with this concrete example, we can now work out the set of equilibria exactly. If $p^* > \frac{1}{4}$, then there is no solution to $p^* = r(z)f(z) = z(1-z)$ (since the right-hand side has a maximum value of $\frac{1}{4}$, at $z = \frac{1}{2}$), and so the only equilibrium is when $z = 0$. This corresponds to a good that is simply too expensive, and so the only equilibrium is when everyone expects it not to be used.

On the other hand, when p^* is between 0 and $\frac{1}{4}$, there are two solutions to $p^* = z(1-z)$: they are at points z' and z'' where the horizontal line $y = p^*$ slices through the parabola defined by $z(1-z)$, as shown in Figure 17.3. Thus there are three possible equilibria in this case: when z is equal to any of 0, z' , or z'' . For each of these three values of z , if people expect exactly a z fraction of the population to buy the good, then precisely the top z fraction of the population will do so.

There are two initial observations worth making from this example. First, the notion of a self-fulfilling expectations equilibrium corresponds, in a general sense, to the effects of aggregate “consumer confidence.” If the population has no confidence in the success of the good, then because of the network effects, no one will want it, and this lack of confidence will be borne out by the failure of people to purchase it. On the other hand — for the very same good, at the same price — if the population is confident of its success, then it is possible for a significant fraction of the population to decide to purchase it, thereby confirming its success. The possibility of multiple equilibria in this way is characteristic of markets in which network effects are at work.

A second observation concerns the nature of consumer demand in this case. Compared to the simple, decreasing curve in Figure 17.2, the curve in Figure 17.3 highlights the complicated relationship between the price and the equilibrium quantity. In particular, as the price p^* drops gradually below $\frac{1}{4}$, the “high” equilibrium z'' moves right (as in the simple model without network effects), but the “low” equilibrium z' moves left, toward smaller fractions of the population. To understand how these two equilibria relate to each other, we need to consider an important qualitative contrast between them, which we formulate in the next section.

17.3 Stability, Instability, and Tipping Points

Let’s continue with the example in Figure 17.3, and explore the properties of its equilibria. To begin with, it’s useful to work through the details of why values of z other than 0, z' , or z'' do not constitute equilibria. In particular, suppose that a z fraction of the population were to purchase the good, where z is not one of these three equilibrium quantities.

- If z is between 0 and z' , then there is “downward pressure” on the consumption of the

good: since $r(z)f(z) < p^*$, the purchaser named z (and other purchasers just below z) will value the good at less than p^* , and hence will wish they hadn't bought it. This would push demand downward.

- If z is between z' and z'' , then there is “upward pressure” on the consumption of the good: since $r(z)f(z) > p^*$, consumers with names slightly above z have not purchased the good but will wish they had. This would drive demand upward.
- Finally, if z is above z'' , then there is again downward pressure: since $r(z)f(z) < p^*$, purchaser z and others just below will wish they hadn't bought the good, pushing demand down.

This set of three different possibilities for the non-equilibrium values of z has interesting consequences for the equilibria z' and z'' . First, it shows that z'' has a strong *stability* property. If slightly more than a z'' fraction buys the good, then the demand gets pushed back toward z'' ; if slightly less than a z'' fraction buys the good, then the demand correspondingly gets pushed up toward z'' . So in the event of a “near miss” in the population's expectations around z'' , we would expect the outcome to settle down to z'' anyway.

The situation looks different — and highly unstable — in the vicinity of the equilibrium z' . If slightly more than a z' fraction buys the good, then upward pressure drives the demand away from z' toward the higher equilibrium at z'' . And if slightly less than a z' fraction buys the good, then downward pressure drives the demand away from z' in the other direction, down toward the equilibrium at 0. Thus, if *exactly* a z' fraction of the population purchases the good, then we are at equilibrium; but if the fraction is even slightly off from this, the system will tend to spiral up or spiral down to a significant extent.

Thus, z' is not just an unstable equilibrium; it is really a *critical point*, or a *tipping point*, in the success of the good. If the firm producing the good can get the population's expectations for the number of purchasers above z' , then they can use the upward pressure of demand to get their market share to the stable equilibrium at z'' . On the other hand, if the population's expectations are even slightly below z' , then the downward pressure will tend to drive the market share to 0. The value z' is the hump the firm must get over in order to succeed.

This view of the equilibria suggests a way of thinking about the price p^* . If the firm were to price the good more cheaply — in other words, to lower the price p^* — then this would have two beneficial effects. Since the parabola in Figure 17.3 would now be sliced by a *lower* horizontal line (reflecting the lower price), the low equilibrium z' would move left; this provides a critical point that is easier to get past. Moreover, the high equilibrium z'' would move right, so if the firm is able to get past the critical point, the eventual size of its user population z'' would be even larger. Of course, if p^* is set below the cost of production the firm loses money. But as part of a pricing strategy over time, in which early losses may




Figure 17.4: From a model with network effects, we can define a function $\hat{z} = g(z)$: if everyone expects a z fraction of the population to purchase the good, then in fact a $g(z)$ fraction will do so.

be offset by growth in the user population and later profits, this may be a viable strategy. Many firms do this by offering free trials for their products or by setting low introductory prices.

17.4 A Dynamic View of the Market

There is another way to view this critical point idea that is particularly illuminating. We have been focusing on an equilibrium in which consumers correctly predict the number of actual users of the good. Let's now ask what this would look like if consumers have common beliefs about how many users there will be, but we allow for the possibility that these beliefs are not correct.

This means that if everyone believes a z fraction of the population will use the product, then consumer x — based on this belief — will want to purchase if $r(x)f(z) \geq p^*$. Hence, if anyone at all wants to purchase, the set of people who will purchase will be between 0 and \hat{z} , where \hat{z} solves the equation $r(\hat{z})f(z) = p^*$. Equivalently,

$$r(\hat{z}) = \frac{p^*}{f(z)}, \quad (17.1)$$

or, taking the inverse of the function $r(\cdot)$,

$$\hat{z} = r^{-1} \left(\frac{p^*}{f(z)} \right). \quad (17.2)$$




Figure 17.5: When $r(x) = 1 - x$ and $f(z) = z$, we get the curve for $g(z)$ shown in the plot: $g(z) = 1 - p^*/z$ if $z \geq p^*$ and $g(z) = 0$ if $z < p^*$. Where the curve $\hat{z} = g(z)$ crosses the line $\hat{z} = z$, we have self-fulfilling expectations equilibria. When $\hat{z} = g(z)$ lies below the line $\hat{z} = z$, we have downward pressure on the consumption of the good (indicated by the downward arrows); when $\hat{z} = g(z)$ lies above the line $\hat{z} = z$, we have upward pressure on the consumption of the good (indicated by the upward arrows). This indicates visually why the equilibrium at z' is unstable while the equilibrium at z'' is stable.

This provides a way of computing the outcome \hat{z} from the shared expectation z , but we should keep in mind that we can only use this equation when there is in fact a value of \hat{z} that solves Equation (17.1). Otherwise, the outcome is simply that no one purchases.

Since $r(\cdot)$ is a continuous function that decreases from $r(0)$ down to $r(1) = 0$, such a solution will exist and be unique precisely when $\frac{p^*}{f(z)} \leq r(0)$. Therefore, in general, we can define a function $g(\cdot)$ that gives the outcome \hat{z} in terms of the shared expectation z as follows. When the shared expectation is $z \geq 0$, the outcome is $\hat{z} = g(z)$, where

- $g(z) = r^{-1}\left(\frac{p^*}{f(z)}\right)$ when the condition for a solution $\frac{p^*}{f(z)} \leq r(0)$ holds; and $g(z) = 0$ otherwise.

Let's try this on the example illustrated in Figure 17.3, where $r(x) = 1 - x$ and $f(z) = z$. In this case, $r^{-1}(x)$ turns out to be $1 - x$. Also, $z(0) = 1$, so the condition for a solution $\frac{p^*}{f(z)} \leq r(0)$ is just $z \geq p^*$. Therefore, in this example




Figure 17.6: The curve $g(z)$, and its relation to the line $\hat{z} = z$, illustrates a pattern that we expect to see in settings more general than just the example used for Figure 17.5.

$$g(z) = 1 - \frac{p^*}{z} \text{ when } z \geq p^*, \text{ and } g(z) = 0 \text{ otherwise.}$$

We can plot the function $\hat{z} = g(z)$ as shown in Figure 17.4. Beyond the simple shape of the curve, however, its relationship to the 45° line $\hat{z} = z$ provides a striking visual summary of the issues around equilibrium, stability, and instability that we've been discussing. Figure 17.5 illustrates this. To begin with, when the plots of the two functions $\hat{z} = g(z)$ and $\hat{z} = z$ cross, we have a self-fulfilling expectations equilibrium: here $g(z) = z$, and so if everyone expects a z fraction of the population to purchase, then in fact a z fraction will do so. When the curve $\hat{z} = g(z)$ lies below the line $\hat{z} = z$, we have downward pressure on the consumption of the good: if people expect a z fraction of the population to use the good, then the outcome will underperform these expectations, and we would expect a downward spiral in consumption. And correspondingly, when the curve $\hat{z} = g(z)$ lies above the line $\hat{z} = z$, we have upward pressure on the consumption of the good.

This gives a pictorial interpretation of the stability properties of the equilibria. Based on how the functions cross in the vicinity of the equilibrium z'' , we see that it is stable: there is upward pressure from below and downward pressure from above. On the other hand, where the curves cross in the vicinity of the equilibrium z' , there is instability — downward pressure from below and upward pressure from above, causing the equilibrium to quickly unravel if it is perturbed in either direction.

The particular shape of the curve in Figure 17.5 depends on the functions we chose in our




Figure 17.7: A “zoomed-in” region of a curve $\hat{z} = g(z)$ and its relation to the line $\hat{z} = z$.

example, but the intuition behind this picture is much more general than the example. With network effects in general, we would expect to see a relation between the expected number of users and the actual number of purchasers that looks qualitatively like this curve, or more generally like the smoother version in Figure 17.6. Where the curve $\hat{z} = g(z)$ crosses the line $\hat{z} = z$, we have equilibria that can be either stable or unstable depending on whether the curve crosses from above or below the line.

The Dynamic Behavior of the Population. In the 1970s, Mark Granovetter and Thomas Schelling used pictures like the ones in Figures 17.5 and 17.6 to model how a population might react dynamically to a network effect [192, 366]. Specifically, they were interested in how the number of people participating in a given activity with network effects would tend to grow or shrink over time.

To motivate the kind of question they formulated, let’s imagine that instead of evaluating the purchase of a discrete object like a fax machine, people in society are evaluating their




Figure 17.8: The audience size changes dynamically as people react to the current audience size. This effect can be tracked using the curve $\hat{z} = g(z)$ and the line $\hat{z} = z$.

participation in a large social media site — something where you chat with friends, share videos, or some similar activity. We are formulating the underlying story here in terms of participation rather than purchasing because the dynamics of participation are more fluid than the dynamics of purchasing: someone can change their mind about participation in a social media site from one day to the next, whereas purchasing a physical good is a step that isn't as naturally undone.

Despite the change in the motivating story, the model remains exactly the same. Each person x has an intrinsic interest in using the site, represented by a function $r(x)$, and the site is more attractive to people if it has more users, as governed by a function $f(z)$. Counterbalancing this, let's suppose that there is a fixed level of effort required to use the site, which serves the role of a “price” p^* (except that the price may consist of the expenditure of effort rather than money). Thus, if person x expects a z fraction of the population to want to participate, then x will participate if $r(x)f(z) \geq p^*$. This is just the same as the




Figure 17.9: Successive updates cause the audience size to converge to a stable equilibrium point (and to move away from the vicinities of unstable ones).

criterion we saw before.

Let's suppose that time proceeds in a fixed set of periods (e.g. days, weeks, or months) $t = 0, 1, 2, \dots$. At time $t = 0$, some initial fraction of the population z_0 is participating in the site — let's call this the initial *audience size*. Now, the audience size changes dynamically over time as follows. In each period t , people evaluate whether to participate based on a shared expectation that the audience size will be the same as what it was in the previous period. In terms of our function $g(\cdot)$, which maps shared expectations to outcomes, this means that $z_1 = g(z_0)$, since everyone acts in period $t = 1$ on the expectation that the audience size will be z_0 . After this, $z_2 = g(z_1)$, since in period $t = 2$ everyone will act based on the expectation that the audience size is now z_1 ; and more generally, we have $z_t = g(z_{t-1})$ for each t .

This is clearly a model in which the population is behaving in a myopic way — they evaluate the benefits of participation as though the future will be the same as the present.

However, it is an approximation that can be reasonable in settings where people have relatively limited information, and where they are behaving according to simple rules. Moreover, part of its value as an approximation in this case is that it produces dynamic behavior that closely corresponds to our notions of equilibrium: if the population follows this model, then it converges precisely to self-fulfilling expectations equilibria that are stable. We discuss the reasons for this next.

Analyzing the Dynamics. The dynamic behavior of the population can be analyzed in a way that is purely “pictorial” but nevertheless completely rigorous. Here is how this works, using a zoomed-in region of the curve $\hat{z} = g(z)$ in the vicinity of two equilibria as shown in Figure 17.7.

We have an initial audience size z_0 , and we want to understand how the sequence of audience sizes $z_1 = g(z_0)$, $z_2 = g(z_1)$, $z_3 = g(z_2)$, \dots behaves over time. We will do this by tracking the points (z_t, z_t) , as t ranges over $t = 0, 1, 2, \dots$; notice that all of these points lie on the diagonal line $\hat{z} = z$. The basic way in which we move from one of these points to the next one is shown in Figure 17.8. We start by locating the current audience size z_0 on the line $\hat{z} = z$. Now, to determine z_1 , we simply move vertically until we reach the curve $\hat{z} = g(z)$, since this gives us the value of $z_1 = g(z_0)$. Then we again locate the audience size z_1 on the line $\hat{z} = z$ — this involves moving horizontally from the point $(z_0, z_1 = g(z_0))$ until we reach the point (z_1, z_1) . We have therefore gone from (z_0, z_0) to (z_1, z_1) , following the evolution of the audience in the first time period.

This is the basic operation: for each time period t , we can determine the new audience size z_t from the current one z_{t-1} in the same way. We first move vertically from the point (z_{t-1}, z_{t-1}) to the point (z_{t-1}, z_t) (which lies on the curve $\hat{z} = g(z)$); we then move horizontally from the point (z_{t-1}, z_t) to the point (z_t, z_t) .

Figure 17.9 now shows what happens as we track this sequence of points, following how the audience size changes. When we’re following a part of the curve $\hat{z} = g(z)$ that lies above the diagonal line $\hat{z} = z$, the points move upward, converging to the nearest place where the two functions cross, which will be at a stable equilibrium point. On the left- and right-hand parts of the picture we show what happens to two other trajectories that start out from points where the curve $\hat{z} = g(z)$ lies *below* the diagonal line $\hat{z} = z$. Here, the sequence of points that track the audience size will move downward, again converging to the first crossing point it reaches; again, this will be a stable equilibrium. Notice that around the unstable equilibrium point in the figure, this means that the trajectories of points move *away* from it on either side — consistent with our view of unstable equilibria, there is no way to reach this point unless you start right at it.

Thus, this simple dynamics for updating the audience size — although it is based on myopic behavior by the population — illustrates how stable and unstable equilibria govern the

outcomes. Stable equilibria attract the population from both sides, while unstable equilibria act like “branch points,” with the audience size flowing away from them on either side.

17.5 Industries with Network Goods

The discussion and models thus far provide some useful intuitions about how an industry with network effects might be expected to evolve over time. Let’s discuss what we might learn at a qualitative level from these models. We’ll continue to use “audience size” for the fraction of the population that purchases a product.

Let’s start with Figure 17.3, and suppose that a new product is introduced with a high initial cost of production — in particular, suppose the horizontal line at height p^* is above the top of the parabola. In this case the only equilibrium is at audience size $z = 0$. If over time the cost of production falls then eventually a horizontal line drawn at p^* will intersect the parabola in two points, much like we see in Figure 17.3, and there will be three possible equilibria. But when p^* is large, near the top of the curve in Figure 17.3, it is likely that none of the good will be sold: to have any sales occur consumers would have to expect an audience size of at least z' , which will be large when p^* is large (near the fraction of the population where the parabola reaches its peak). Given that none were sold previously — when the cost was above the top of the curve — this seems an unlikely prediction. But as the cost of production continues to fall, the critical point decreases (as z' gets closer to 0), and an audience size of at least z' starts to seem more and more likely. Once consumers expect the good to be viable, with an audience size of at least z' , the stable equilibrium is in fact z'' . So as costs decline we would expect to initially see no sales, and then once purchases begin we would expect to see sales grow rapidly to the stable point.

Marketing a Product with Network Effects. How can a firm that wants to sell a product with a network effect use these insights to market its product? Suppose you are running a firm that is producing a new product subject to network effects; perhaps it’s a new piece of software, communication technology, or social media. The marketing of this product will not succeed unless you can get past the tipping point (at z'). Starting small and hoping to grow slowly is unlikely to succeed, since unless your product is widely used it has little value to any potential purchasers.

Thus, you somehow need to convince a large initial group to adopt your product before others will be willing to buy it. How would you do this? One possibility is to set an initial low, introductory price for the good, perhaps even offering it for free. This price below the cost of producing the good will result in early losses, but if the product catches on — if it gets over the tipping point — then your firm can raise the price and perhaps make enough profit to overcome the initial losses.

Another alternative is to attempt to identify fashion leaders, those whose purchase or use of the good will attract others to use it, and convince them to adopt the good. This strategy also involves network effects, but they are ones that cannot be studied at the population level. Instead we would need to identify a network of connections between potential purchasers and ask who influences whom in this network. We explore this idea in Chapter 19.

Social Optimality with Network Effects. We saw in Section 17.1 that for a market with no network effects, the equilibrium is socially optimal. That is, it maximizes the total difference between the reservation prices of the consumers who purchase the good and the total cost of producing the good, over all possible allocations to people.

For goods with network effects, however, the equilibria are typically not optimal. At a high level, the reason is that each consumer's choice affects each other consumer's payoff, and the consequences of this can be analyzed as follows. Suppose we are at an equilibrium in which the audience size is z^* . The consumer named z^* — the purchaser with the least interest in the product — has a reservation price of $r(z^*)f(z^*) = p^*$. Now, consider the set of consumers with names above z^* and below $z^* + c$ for some small constant $c > 0$. None of these consumers want to buy, since $r(z)f(z^*) < p$ for z in this range. But if they all did purchase the good, then all the current purchasers would benefit: the value of the product to each purchaser $x < z^*$ would increase from $r(x)f(z^*)$ to $r(x)f(z^* + c)$. The potential consumers between z^* and $z^* + c$ don't take this effect into account in their respective decisions about purchasing the good.

It is easy to set up situations where this overall benefit to existing purchasers outweighs the overall loss that consumers between z^* and $z^* + c$ would experience from buying the good. In such a case the equilibrium is not socially optimal, since society would be better off if these additional people bought the good. This example illustrates the more general principle that for goods with network effects, markets typically provide less of the good than is socially optimal.

Network Effects and Competition. Finally, let's ask what might happen if multiple firms develop competing new products, each of which has its own network effects. For example, we could consider two competing social-networking sites that offer similar services, or two technologies that do essentially the same thing, but where the value of each of these technologies depends on how many people use it. There are a number of classic examples of this from technology industries over the last several decades [27]. These include the rise of Microsoft to dominate the market for personal-computer operating systems, and the triumph of VHS over Betamax as the standard videotape format in the 1980s.

In such cases of product competition with network effects, it is likely that one product will dominate the market, as opposed to a scenario in which both products (or even more

than two) flourish. The product that first gets over its own tipping point attracts many consumers and this may make the competing product less attractive. Being the first to reach this tipping point is very important — more important than being “best” in an abstract sense. That is, suppose that if product A has audience size z , then consumer x values it at $r_A(x)f(z)$, while if product B has audience size z , then consumer x values it at a larger amount $r_B(x)f(z) > r_A(x)f(z)$. Let’s also suppose that each product can be produced at the same price. Then it seems reasonable to say that product B is the better product. But if product A is on the market first, and gets over its tipping point, then product B may not be able to survive.⁴

These considerations help provide some intuition for how markets with strong network effects tend to behave. Writing in the *Harvard Business Review* in 1996, Brian Arthur summarized the “hallmarks” of these markets in a way that reflects the discussion in the previous paragraph: “market instability (the market tilts to favor a product that gets ahead), multiple potential outcomes ([e.g.,] under different events in history different operating systems could have won), unpredictability, the ability to lock in a market, the possible predominance of an inferior product, and fat profits for the winner” [27]. It is not the case that a given market with network effects will necessarily display all these characteristics, but they are phenomena to watch for in this type of setting.

Of course, in our discussion of the dominance of product A over product B , we are assuming that nothing else changes to shift the balance after A achieves dominance. If the firm that makes product B improves its product sufficiently and markets it well, and if the firm that makes product A doesn’t respond effectively, then B may still overtake A and become the dominant product.

17.6 Mixing Individual Effects with Population-Level Effects

Thus far we have been focusing on models of network effects in which the product is useless to consumers when it has an audience size of 0; this is captured by our assumption that $f(0) = 0$. But of course one can also study more general kinds of network effects, in which a product has some value to a person even when he or she is the first purchaser, and its value then increases as more people buy it. We can think of this as a model that mixes individual effects (a person’s value for the product on its own) with population-level effects (the increased value a person derives when the product has a large audience size). In such a model, we would have $f(0) > 0$, and have $f(z)$ increasing in z .

We won’t attempt to cover all the ways of fleshing out such a model; instead we develop one general class of examples to illustrate how qualitatively new phenomena can arise when

⁴Exercises 3 and 4 at the end of this chapter offer simple models of this situation.




Figure 17.10: When $f(0) > 0$, so that people have value for the product even when they are the only user, the curve $\hat{z} = g(z)$ no longer passes through the point $(0, 0)$, and so an audience size of 0 is no longer an equilibrium.

we mix individual and population-level effects. In particular, we focus on a phenomenon identified in this type of model by Mark Granovetter [192], and which corresponds to an intuitively natural issue in the marketing of new products with network effects.

A Concrete Model. For our example, let's consider a function $f(\cdot)$ of the form $f(z) = 1 + az^2$ for a constant parameter a . We'll continue to use the simple example $r(x) = 1 - x$; so when the audience size is z , the value of the product to consumer x is

$$r(x)f(z) = (1 - x)(1 + az^2).$$

Now let's apply the analysis from Section 17.4 to this function to get the dynamic behavior of the market. We will assume that the price p^* is strictly between 0 and 1. When everyone expects an audience size of z , the fraction of people who actually use the product is $\hat{z} = g(z)$, where $g(\cdot)$ is defined as in Section 17.4:




Figure 17.11: The audience grows dynamically from an initial size of zero to a relatively small stable equilibrium size of z^* .

$$g(z) = r^{-1} \left(\frac{p^*}{f(z)} \right) \text{ when the condition for a solution } \frac{p^*}{f(z)} \leq r(0) \text{ holds; and } g(z) = 0 \text{ otherwise.}$$

As before, we have $r^{-1}(x) = 1 - x$. Since in our case $r(0) = 1$, $f(z) \geq 1$, and $p^* < 1$, this means that the condition for a solution $\frac{p^*}{f(z)} \leq r(0)$ will always hold. Plugging this into the formula for $g(z)$, we get

$$g(z) = 1 - \frac{p^*}{1 + az^2}.$$

When we plot this function $\hat{z} = g(z)$ together with the 45° line $\hat{z} = z$, we get something that looks like Figure 17.10.

Growing an Audience from Zero. In our earlier model with $f(0) = 0$, an audience size of zero was a stable equilibrium: if everyone expected that no one would use the product, then no one would. But when $f(0) > 0$, so that the product has value to people even when

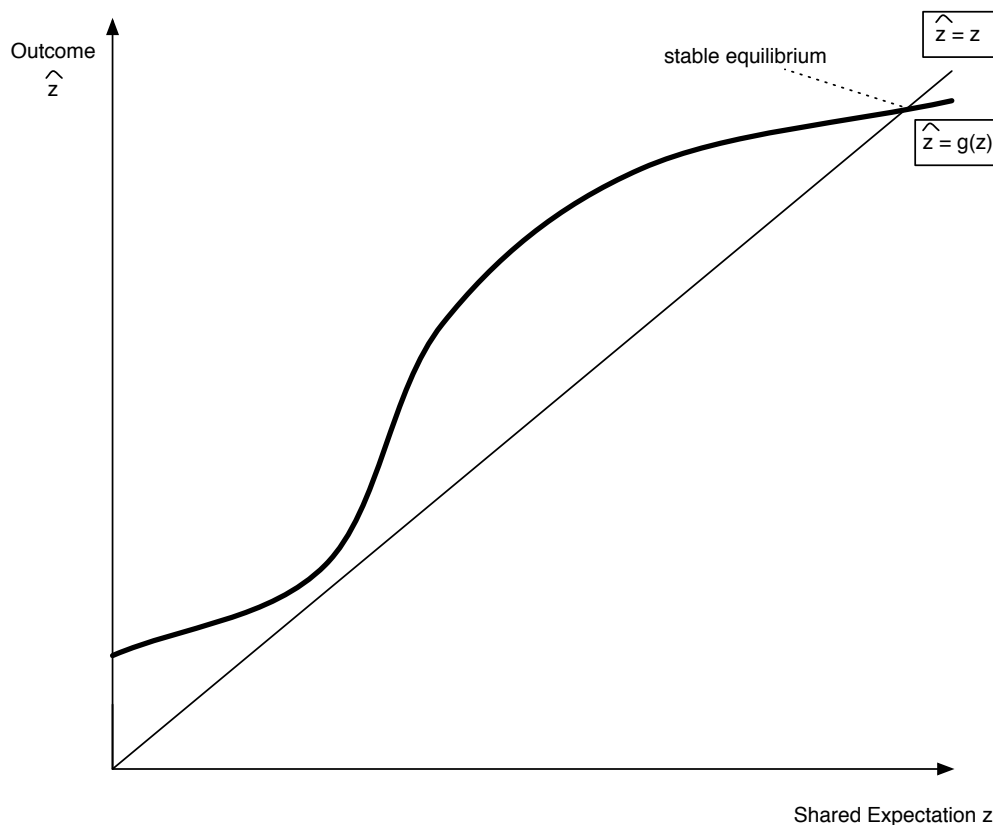


Figure 17.12: If the price is reduced slightly, the curve $\hat{z} = g(z)$ shifts upward so that it no longer crosses the line $\hat{z} = z$ in the vicinity of the point (z^*, z^*) .

they are the only user, an audience size of zero is no longer an equilibrium (when $p^* < 1$): even if everyone expects no one to use the product, some people will still purchase it.

As a result, it becomes natural to ask what happens when such a product starts at an audience size of zero, and we then follow the dynamics that were defined in Section 17.4. Figure 17.11 shows what happens when we do this: the sequence of audience sizes increases from $z_0 = 0$ up to the first point (z^*, z^*) at which the curve $\hat{z} = g(z)$ crosses the line $\hat{z} = z$. This is the stable equilibrium that is reached when we run the dynamics of the market starting from an audience size of 0.

Notice how the underlying story that we're modeling with this process has no direct analogue in the earlier model when $f(0) = 0$. There, because the product was useless if it had an audience size of zero, a firm marketing the product needed alternate ways to get over its tipping point at the low, unstable equilibrium in order to have any customers at all. But when $f(0) > 0$, the audience can grow from zero up to some larger stable equilibrium z^* through the simple dynamics in Figure 17.11. In other words, we're able to talk here about




Figure 17.13: The small reduction in price that shifted the curve $\hat{z} = g(z)$ has a huge effect on the equilibrium audience size that is reached starting from zero.

an audience that grows gradually and organically, starting from no users at all, rather than one that needs to be pushed by other means over an initial tipping point.

Bottlenecks and Large Changes. The firm marketing the product in our example, however, may well want more than what it gets in Figure 17.11. Although the audience grows to some size z^* on its own, there is a much higher stable equilibrium, shown in the figure at (z^{**}, z^{**}) , that would be much more desirable if only it could be reached. But starting from zero, the audience doesn't reach this high equilibrium z^{**} , because it is blocked by a “bottleneck” that stops it at z^* .

Here is where we get to the surprising phenomenon at the heart of this example: small changes in the properties of the market can cause enormous changes in the size of the equilibrium audience that is reached, starting from zero [192]. Suppose that the firm is able to lower the price p^* slightly, to some new value $q^* < p^*$. Then we get a new function $h(z)$

mapping shared expectations to outcomes,

$$h(z) = 1 - \frac{q^*}{1 + az^2},$$

which in turn defines a new dynamic process. As q^* is made smaller, the curve $\hat{z} = h(z)$ shifts upward until it no longer crosses the line $\hat{z} = z$ at all in the vicinity of the point (z^*, z^*) ; this is shown in Figure 17.12. However, $h(\cdot)$ still has a high stable equilibrium close to the high equilibrium (z^{**}, z^{**}) for the function $g(\cdot)$.

As soon as $h(\cdot)$ lifts enough that it no longer crosses $\hat{z} = z$ near (z^*, z^*) , the equilibrium audience size starting from zero changes abruptly and dramatically: it suddenly jumps from a value near z^* to a much higher value near z^{**} . There is a natural reason for this: as shown in Figure 17.13, the “bottleneck” at (z^*, z^*) has opened into a narrow passageway, and so now the dynamics starting from the point $(0, 0)$ can carry the audience all the way up to the stable equilibrium near (z^{**}, z^{**}) .⁵

This phenomenon shows how in models with network effects, small changes in market conditions can have strong, discontinuous effects on the outcome. The contrast between Figure 17.11 and Figure 17.13 relates to an important issue in the marketing of products with network effects. In Figure 17.11, the product reaches a small group of the most enthusiastic consumers (the ones with the highest values for the product), but it fails to make the leap from this group to the much broader set of people — mainstream, less enthusiastic consumers — who could collectively push the audience size up to the higher equilibrium at z^{**} . However, once the price is lowered very slightly, making the product slightly more attractive to everyone, a passage is opened that enables the success of the product with its most enthusiastic consumers to carry over to this larger mainstream set, driving the equilibrium audience size up to a much larger value.

17.7 Advanced Material: Negative Externalities and The El Farol Bar Problem

In different contexts, we have now analyzed situations with both negative externalities (traffic congestion and the Braess Paradox) and positive externalities (goods with network effects, in this chapter). The settings for these analyses each contained a number of details designed to capture their respective contexts: in our discussion of negative externalities we had the complexity of an underlying network through which the traffic flowed; with positive externalities, we had a heterogeneous population with diverse reservation prices, all reacting to a common market price.

⁵It is not hard to find specific numbers that cause this effect to happen; for example, you can try $f(z) = 1 + 4z^2$ and $p^* = 0.93$. In this case, the equilibrium audience size starting from zero is around 0.1. If we then lower the price slightly to $q^* = 0.92$, the equilibrium audience size starting from zero jumps to around 0.7.

But even after eliminating all these details, reducing the problems to simpler forms, the phenomena that surround negative and positive externalities remain quite different at a more fundamental level. In this section we consider some of these contrasts using stylized, simple examples that enable us to highlight the differences more clearly. In the process, we will also consider the question of how individuals coordinate on equilibrium behavior in each setting.

Simple Scenarios with Negative and Positive Externalities. As our simplified setting for negative externalities, we use the widely-studied *El Farol Bar problem* created by Brian Arthur [26]. The problem is named after a bar in Sante Fe that used to have live music every Thursday evening. In the formulation of the problem, the bar has seating for only 60 people, and so showing up for the music is enjoyable only when at most 60 people do so. With more than 60 people in attendance, it becomes unpleasantly crowded, such that it would be preferable to have stayed home. Now, unfortunately, there are 100 people each week who are interested in going to the bar, and they all share the view that going is only worthwhile when at most 60 people show up. How does each person reason in a given week about whether to go or to stay home — knowing that everyone else is reasoning about this decision as well?

The El Farol Bar problem describes a situation with a very simple negative externality: the payoffs to participating in the underlying activity (going to the bar) decrease as the number of participants increases. And despite the appealingly simple statement of the problem, it creates a situation in which the reasoning problem faced by the participants is very complex. To illustrate where some of these complexities come from, it's useful to compare it to a simple analogous situation that contains a positive externality.

In this parallel scenario, let's imagine a division of a large company consisting of 100 people, in which the management is encouraging the employees to use a particular corporate social-networking site as part of their workflow. The management would like each employee to create an account and maintain a presence on the site to facilitate increased interaction across the division. The employees each believe that this would be worthwhile provided that enough people in the division participate in the site; otherwise, the effort required would not be worth it. Thus, each employee wants to use the social-networking site if at least 60 other employees do so as well. (So counting the employee herself, this means the total number of employees using the site should be strictly greater than 60 in order for it to be worth the effort.) This is closely analogous to the scenarios we've considered in earlier sections of this chapter, concerning goods with network effects — although here, instead of each individual having a distinct reservation price, there is simply a common interest in participating provided the audience size is large enough. (Also, the size of the population is finite rather than infinite.) In a different way, the analogy to the El Farol Bar problem should also be clear: in the scenario of the social-networking site we have a positive externality in

17.7. ADVANCED MATERIAL: NEGATIVE EXTERNALITIES AND THE EL FAROL BAR PROBLEM

which a participation level above 60 is good, while in El Farol, we have a negative externality in which a participation level above 60 is bad.

These two examples have been designed in such a way that one exhibits only negative externalities, and the other exhibits only positive externalities. It's important to keep in mind, of course, that many real situations in fact display both kinds of externalities — some level of participation by others is good, but too much is bad. For example, the El Farol Bar might be most enjoyable if a reasonable crowd shows up, provided it does not exceed 60. Similarly, an on-line social media site with limited infrastructure might be most enjoyable if it has a reasonably large audience, but not so large that connecting to the Web site becomes very slow due to the congestion. To keep the discussion here as clean as possible, we are keeping the two kinds of externalities separate, but understanding how they work in combination is an important topic of ongoing research [229]. We also consider a simple way of combining the two effects in Exercise 2 at the end of this chapter.

Basic Comparisons between the Two Scenarios. The contrasts between the two scenarios — El Farol and corporate social-networking — translate into significant differences in how we should expect people to behave. Let's first think about these differences informally; later we'll carry out the analysis in more detail.

Reasoning about the social-networking scenario using what we've seen earlier in this chapter, we find that there are two very natural equilibria the group of 100 people could exhibit. If everyone participates, then everyone has an interest in participating; similarly, if no one participates, then no one has an interest in participating. (There are also other more complex equilibria, as we'll see later, but these two all-or-nothing equilibria are by far the two most natural.)

On the other hand, neither of these outcomes is an equilibrium in the El Farol Bar problem. If everyone were to attend, then everyone would have an incentive to stay home; and if no one were to attend, then everyone would have an incentive to attend. Instead, the equilibria have a more complex structure in which individuals need to break the underlying symmetry in such a way that some people attend and some stay home.

There is a second, essentially equivalent way to describe this contrast, using the idea of shared expectations from earlier in the chapter. In the social-networking scenario, if the individuals have a shared expectation that everyone will participate, then this expectation is self-fulfilling: everyone will in fact participate. On the other hand, if they have a shared expectation that no one will participate, this expectation too will be self-fulfilling. As we've seen, understanding such self-fulfilling expectations is a key part of reasoning about situations with positive externalities.

The negative externalities in the El Farol Bar problem, on the other hand, pose problems for shared expectations. In particular, individuals cannot have a fixed, shared expectation

of the audience size at El Farol that will be self-fulfilling. If everyone expects an audience size of at most 60, then everyone will show up, thereby negating this prediction. Similarly, if everyone expects an audience size above 60, then everyone will stay home, negating this prediction too.⁶

These are fundamental contrasts: with positive externalities, there exist self-fulfilling expectations and a natural set of outcomes to coordinate on; with negative externalities, any shared expectation of a fixed audience size will be self-negating, and the individuals must instead sort themselves out in much more complicated ways. Given this complexity, the El Farol Bar problem has become a testing ground for a variety of models of individual behavior. We now describe some of these models and styles of analysis in more detail.

Nash Equilibria in the El Farol Bar Problem. First, let's consider how to model the El Farol Bar problem as a game that will be played once by the set of 100 people. (We could imagine that instead of the bar having music every Thursday, it is simply hosting a single concert, and everyone needs to decide in advance whether to attend.) Each person has two possible strategies, *Go* (to the bar) or *Stay* (home), and his payoffs are as follows.

- If he chooses *Stay*, then he receives a payoff of 0 in all outcomes.
- If he chooses *Go*, then he receives a payoff of $x > 0$ when at most 60 people choose *Go*, and a payoff of $-y < 0$ when more than 60 people choose *Go*.

There are many different pure-strategy Nash equilibria for this game. Our discussion above makes clear that there is no equilibrium in which all players use the same pure strategy, but any outcome in which exactly 60 people choose *Go* and 40 people choose *Stay* is a pure-strategy Nash equilibrium. Of course, it is far from clear how the group would settle on this set of heterogeneous strategies, since at the outset of the game, they are all identical. We will return to the question of how heterogeneous strategies might arise later in this section.

There is, however, an equilibrium in which all players behave symmetrically, and this is through the use of mixed strategies, in which each player chooses *Go* with the same probability p . In this case too, there are some subtleties. It would be natural to guess that the shared probability p in this mixed-strategy equilibrium would be 0.6, but this is not necessarily the case. Instead, p depends on the payoffs x and $-y$: following the reasoning we saw in Chapter 6, we need to choose p so that each player is indifferent between choosing *Go* and choosing *Stay*. This will ensure that no one has an incentive to deviate from randomizing between the two alternatives.

⁶As Brian Arthur notes, this latter possibility is a reflection of the same phenomenon that the baseball player Yogi Berra invoked when he quipped about a popular restaurant, "Nobody goes there anymore; it's too crowded" [26, 105].

Since the payoff for *Stay* is always 0, this means that we need to choose p so that the expected payoff from *Go* is also 0. Therefore, we need to choose p so that the equation

$$x \cdot \Pr[\textit{at most 60 go}] - y \cdot \Pr[\textit{more than 60 go}] = 0 \quad (17.3)$$

holds. Using the fact that

$$\Pr[\textit{more than 60 go}] = 1 - \Pr[\textit{at most 60 go}],$$

we can rearrange Equation (17.3) to get

$$\Pr[\textit{at most 60 go}] = \frac{y}{x + y}. \quad (17.4)$$

So in order to have a mixed-strategy equilibrium, we must choose p so that Equation (17.4) holds. When $x = y$, choosing $p = 0.6$ will work [212]. But suppose that x and y are different; for example, perhaps the music at El Farol is pleasant, but the nights on which it is crowded are truly unbearable, so that y is significantly larger than x . In this case, p must be chosen so that the probability at most 60 people go is very high, and so p will be significantly less than 0.6. Since the expected number of people attending is $100p$, this means that in expectation, significantly fewer than 60 people will be showing up. So with $y > x$, the bar will be significantly underutilized in the mixed-strategy equilibrium, due to the shared fear of overcrowding.

The existence of this mixed-strategy equilibrium is a useful counterpoint to our earlier informal discussion about the difficulty of forming a shared expectation in the El Farol Bar problem. It's true that any shared expectation consisting of a fixed number representing the audience size — the kind of shared expectation we used earlier in this chapter — will be negated by what actually happens. But if we allow more complicated kinds of expectations, then in fact there is a shared expectation that will be self-fulfilling — this is the expectation that everyone plans to randomize their decision to attend the bar, choosing to go with the probability p that makes Equation (17.4) come true.

Analogies with Equilibria in Related Games. To get some intuition about the equilibria we've found, it's useful to compare them to the equilibria of some related games.

First, suppose we model the corporate social-networking scenario from earlier in this section as a similar one-shot game: the two possible strategies are *Join* or *Don't Join* (the site); the payoff to *Don't Join* is always 0, the payoff to *Join* is y when more than 60 people join, and the payoff to *Join* is $-x$ when at most 60 join. In this case, corresponding to what we saw in our informal discussion earlier, there are just two pure-strategy equilibria: one in which everyone chooses *Join*, and one in which everyone choose *Don't Join*. Interestingly, the same mixed-strategy equilibrium that applied to the El Farol Bar problem also holds here: if everyone chooses *Join* with a probability p for which

$$-x \cdot \Pr[\textit{at most 60 join}] + y \cdot \Pr[\textit{more than 60 join}] = 0 \quad (17.5)$$

then everyone is indifferent between joining and not joining, and so this is an equilibrium. Since Equations (17.3) and (17.5) are equivalent (they are simply negations of each other), we get the same value of p that we did in the El Farol Bar problem.

Since games with 100 players are inherently complex, it's also instructive to ask what the two-player versions of these games look like. Specifically, in the two-player version of the El Farol Bar problem, each player wants to attend the bar as long as the other player doesn't; in the two-player social-networking game, each player wants to use the site as long as the other one does. Scaled down to this size, each of these corresponds to one of the fundamental games introduced in Chapter 6: the two-player El Farol Bar problem is a Hawk-Dove game, in which the two players try to make their actions different, while the two-player social-networking scenario is a Coordination game, in which the two players try to make their actions the same.

Each of these games has pure-strategy equilibria, as well as a mixed-strategy equilibrium in which the players randomize over their two available strategies. For example, in the two-player version of the El Farol Bar problem, the payoff matrix is shown in Figure 17.14.

		Player 2	
		<i>Stay</i>	<i>Go</i>
Player 1	<i>Stay</i>	0, 0	0, x
	<i>Go</i>	x , 0	$-y$, $-y$

Figure 17.14: Two-Player El Farol Problem

The two pure-strategy equilibria consist of one player choosing *Go* while the other chooses *Stay*. For the mixed-strategy equilibrium, each player chooses *Go* with a probability p that causes the expected payoff from *Go* to be equal to 0:

$$x(1 - p) - yp = 0,$$

and hence $p = x/(x + y)$. By analogy with the multi-player version, p is not equal to $1/2$ unless $x = y$. Also, just as there will be random fluctuations in the actual attendance at the bar around the mean of $100p$ in the multi-player game, there is significant variation in how many people choose *Go* in the two-player version as well. Specifically, with probability p^2 both players choose *Go*, and with probability $(1 - p)^2$ both choose *Stay*.

Repeated El Farol Problems. While the existence of a mixed-strategy equilibrium in which all 100 people follow the same strategy is an important observation about the El Farol Bar problem, it can't be the whole story. It's not clear why, or whether, a group of people would actually arrive at this mixed-strategy equilibrium, or at any other particular equilibrium or pattern of behavior from among the many that are possible. Once the group is playing an equilibrium, no one has an incentive to deviate — that's what it means for the

17.7. ADVANCED MATERIAL: NEGATIVE EXTERNALITIES AND THE EL FAROL BAR PROBLEM

behaviors to be in equilibrium. But how do they coordinate on equilibrium behavior in the first place?

To formulate models that address these questions, it is useful to think about a setting in which the El Farol game is played repeatedly. That is, suppose that each Thursday night the same 100 people must each decide whether to go to the bar or stay home, and each receives a payoff of x (from going as part of a group of at most 60), $-y$ (from going as part of a group of more than 60), or 0 (from staying home). Each person also knows the history of what has happened on each prior Thursday, so they can use this information in making their decision for the current Thursday. By reasoning about decisions over time in this repeated El Farol game, we might hope to see how a pattern of behavior gradually emerges from rules that take past experience into account.

There are a number of different formalisms that can be used to study the repeated El Farol game. One approach is to view the full sequence of Thursdays as a dynamic game of the type studied in Section 6.10, in which players choose sequences of strategies — in this case, one for each Thursday — and correspondingly receive payoffs over time. We could consider Nash equilibria in this dynamic game and see whether the play of the game eventually settles down to repeated play of some equilibrium of the one-shot El Farol game. Essentially, we would be asking if equilibrium play by sophisticated players in the dynamic game converges to something that looks simple. Although learning can go on during the play of an equilibrium in a dynamic game [175], this approach can't fully answer our underlying question of how the individuals come to play a Nash equilibrium at all — since the learning here would be taking place within a Nash equilibrium of the larger dynamic game.

An alternate approach is to ask what might happen if the players are potentially more naive. A useful way to think about how players — sophisticated or naive — behave in a repeated game is to decompose their choice of strategy into a *forecasting rule* and a choice of action given the forecasting rule. A forecasting rule is any function that maps the past history of play to a prediction about the actions that all other players will take in the future. Forecasting rules can in principle be very complex. An individual could take all past behavior into account in generating a forecast and he may also forecast that how others will behave in the future depends on how he behaves now. Alternatively, a very naive player may forecast that each other player will simply use a fixed action forever. From an individual's forecasting rule, we can then make a prediction about his behavior: we assume that each individual behaves optimally given his forecasting rule. That is, he chooses an action that maximizes his expected payoff given whatever he forecasts about the behavior of others.

For the repeated El Farol game, most attention has focused on forecasting rules that work with audience sizes: a given forecasting rule is a function mapping the sequence of past audience sizes to a prediction about the number of other people who will go to the bar on the upcoming Thursday. (Thus, each forecasting rule produces a number between 0 and 99

when given a history of past audience sizes.) This is a bit less expressive than a forecasting rule that predicts, for each other person individually, whether they will go to the bar or stay home, but it captures the main quantity of interest, which is the total number of people who show up. For an individual using any such forecasting rule, his choice of action is easy to describe: he goes to the bar if his forecasting rule produces a number that is at most 59, and he stays home if it produces a number that is 60 or more.

In keeping with our informal discussion earlier in the section about self-fulfilling and self-negating expectations, we observe first of all that if everyone uses the same forecasting rule for the audience size, then everyone will make very bad predictions. In any given situation, either this common forecasting rule will predict 59 or fewer others in attendance, or at least 60 others in attendance. In the first case, everyone will show up, and in the second case, everyone will stay home; in both cases, the forecasting rule is wrong. So to make any progress, we will need for players to use a diversity of different forecasting rules.

A long line of research has considered how the group behaves when they use different classes of forecasting rules; the goal is to understand whether the system converges to a state in which, on any given Thursday, roughly 60% of the agents produce a forecast that causes them to go the bar, and roughly 40% produce a forecast that causes them to stay home (e.g. [26, 104, 167]). This investigation has been carried out both mathematically and by computer simulation, with some of the analysis becoming quite complex. In general, researchers have found that under a variety of conditions, the system converges to a state where the average attendance varies around 60 — in other words, providing near-optimal utilization of the bar over time.

While we won't go into the details of this analysis, it is not hard to get some intuition for why an average attendance of 60 arises very naturally when agents select from a diversity of forecasting rules. To do this, we can analyze perhaps the simplest model of individual forecasting, in which each person chooses a fixed prediction k for the number of others who will show up, and uses this prediction every week. That is, he will ignore the past history and always predict that k other people will be in attendance. Now, if each person picks their fixed value of k uniformly at random from the 100 natural numbers between 0 and 99, what is the expected audience size at the bar each week? The audience will consist of all people whose forecasting rule is based on a value of k between 0 and 59, and the expected number of such people is 60. Thus, with this very naive forecasting, we in fact get an expected attendance of 60 each Thursday, as desired.

Of course, this analysis is based on people who make extremely naive forecasts, but it shows how diversity in the set of forecasting rules can naturally lead to the right level of attendance. One can also ask what happens when individuals select random forecasts from a more complex space of possibilities, in which the prediction is based on the past several audience sizes. Under fairly general assumptions, an average attendance of 60 continues to

hold, although establishing this is significantly more complicated [104].

17.8 Exercises

1. Consider a product that has network effects in the sense of our model from Chapter 17. Consumers are named using real numbers between 0 and 1; the reservation price for consumer x when a z fraction of the population uses the product is given by the formula $r(x)f(z)$, where $r(x) = 1 - x$ and $f(z) = z$.
 - (a) Let's suppose that this good is sold at cost $1/4$ to any consumer who wants to buy a unit. What are the possible equilibrium number of purchasers of the good?
 - (b) Suppose that the cost falls to $2/9$ and that the good is sold at this cost to any consumer who wants to buy a unit. What are the possible equilibrium number of purchasers of the good?
 - (c) Briefly explain why the answers to parts (a) and (b) are qualitatively different.
 - (d) Which of the equilibria you found in parts (a) and (b) are stable? Explain your answer.

2. In Chapter 17, we focused on goods with positive network effects: ones for which additional users made the good more attractive for everyone. But we know from our earlier discussion of Braess's Paradox that network effects can sometimes be negative: more users can sometimes make an alternative less attractive, rather than more attractive. Some goods actually have both effects. That is, the good may become more attractive as more people use it as long there aren't too many users, and then once there are too many users it becomes less attractive as more people use it. Think of a club in which being a member is more desirable if there is a reasonable number of other members, but once the number of members gets too large the club begins to seem crowded and less attractive. Here we explore how our model of network effects can incorporate such a combination of effects.

In keeping with the notation in Chapter 17, let's assume that consumers are named using real numbers between 0 and 1. Individual x has the reservation price $r(x) = 1 - x$ before we consider the network effect. The network effect is given by $f(z) = z$ for $z \leq 1/4$ and by $f(z) = (1/2) - z$ for $z \geq 1/4$. So the network benefit to being a user is maximized when the fraction of the population using the product is $z = 1/4$, once the fraction is beyond $1/4$ the benefit declines, and it becomes negative if more than $1/2$ of the population is using it. Suppose that the price of this good is p where $0 < p < 1/16$.

- (a) How many equilibria are there? Why? [You do not need to solve for the number(s) of users; a graph and explanation is fine.]

(b) Which equilibria are stable? Why?

(c) Consider an equilibrium in which someone is using the good. Is social welfare maximized at this number of users, or would it go up if there were more users, or would it go up if there were fewer users? Explain. [Again no calculations are necessary; a careful explanation is sufficient.]

3. You have developed a new product which performs the same service as an established product and your product is much better than the established product. Specifically, if the number of users of the two products were the same, then each potential purchaser's reservation price for your product would be twice their reservation price for the existing product. The difficulty that you face is that these are products with network effects and no one wants to use more than one of the two products. Currently, every potential purchaser is using the established product. Your cost of production and your competitor's costs of production are exactly the same and let's suppose that they are equal to the price at which your competitor's product is sold.

If all of the potential purchasers switched to your product the maximum price that you could charge (and still have all of them buy your product) would be twice the current price. So clearly you could make a nice profit if you could attract these potential purchasers. How would you attempt to convince users to switch to your product? You do not need to construct a formal model of the situation described in this question. It is sufficient to describe the strategies that you would try.

4. In the model of network effects that we covered in Chapter 17 there was only one product. Now let's ask what might happen if there are two competing products which both have network effects. Assume that for each product:
- (a) If no one is expected to use the product, then no one places a positive value on the product.
 - (b) If one-half of the consumers are expected to use the product, then exactly one-half of the consumers would buy the product.
 - (c) If all of the consumers are expected to use the product, then all consumers would buy the product.

Using an analysis of network effects, describe the possible equilibrium configurations of numbers of consumers using each product and briefly discuss which of these equilibria you would expect to be stable and which you would expect to be unstable. You do not need to build a formal model to answer this question. Just describe in words what might happen in this market.