# Network Effects

Carl Shapiro
Hal R. Varian*
University of California, Berkeley

September 1998 (revised: September 19, 1998)

Let us try to model network externalities using a simple demand and supply model. Suppose that there are 1000 people in a market for some good and we index the people by $v = 1, \ldots, 1000$. Think of $v$ as measuring the reservation price for the good by person $v$. Then if the price of the good is $p$, the number of people who think that the good is worth at least $p$ is $1000 - p$. For example, if the price of the good is \$200, then there are 800 people who are willing to pay at least \$200 for the good, so the total number of units sold would be 800. This structure generates a standard, downward-sloping demand curve.

But now let's add a twist to the model. Suppose that the good we are examining exhibits network externalities, like a fax machine or a telephone. For simplicity, let us suppose that the value of the good to person $v$ is $vn$, where $n$ is the number of people who consume the good—the number of people who are connected to the network. The more people there are who consume the good, the more *each* person is willing to pay to acquire it.[1] What does the demand function look like for this model?

If the price is $p$, there is someone who is just indifferent between buying the good and not buying it. Let $\hat{v}$ denote the index of this marginal

---

*Notes to accompany *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, 1998. Adopted from Hal R. Varian, *Intermediate Microeconomics*, 5th edition, W. W. Norton & Co., 1999. 

[1] We should really interpret $n$ as the number of people who are *expected* to consume the good, but this distinction won't be very important for what follows.

individual. By definition, he is just indifferent to purchasing the good, so his willingness to pay for the good equals its price:

$$p = \hat{v}n.$$

Since this "marginal person" is indifferent, everyone with a *higher* value of $v$ than $\hat{v}$ must definitely want to buy. This means that the number of people who want to buy the good is

$$n = 1000 - \hat{v}.$$

Putting equations and together, we have a condition that characterizes equilibrium in this market:

$$p = n(1000 - n).$$

This equation gives us a relationship between the price of the good and the number of users. In this sense, it is a kind of demand curve; if there are $n$ people who purchase the good, then the willingness to pay of the marginal individual is given by the height of the curve.

However, if we look at the plot of this curve in Figure 1, we see that it has quite a different shape than a standard demand curve! If the number of people who connect is low, then the willingness to pay of the marginal individual is low, because there aren't many other people out there that he can communicate with. If there are a large number of people connected, then the willingness to pay of the marginal individual is low, because everyone else who valued it more highly has already connected. These two forces lead to the humped shape depicted in Figure 1.

Now that we understand the demand side of the market, let's look at the supply side. To keep things simple, let us suppose that the good can be provided by a constant returns to scale technology. As we've seen, this means that the supply curve is a flat line at price equals average cost.

Note that there are three possible intersections of the demand and supply curve. There is a low-level equilibrium where $n^* = 0$. This is where no one consumes the good (connects to the network), so no one is willing to pay anything to consume the good. This might be referred to as a "pessimistic expectations" equilibrium.

The middle equilibrium with a positive but small number of consumers is one where people don't think the network will be very big, so they aren't

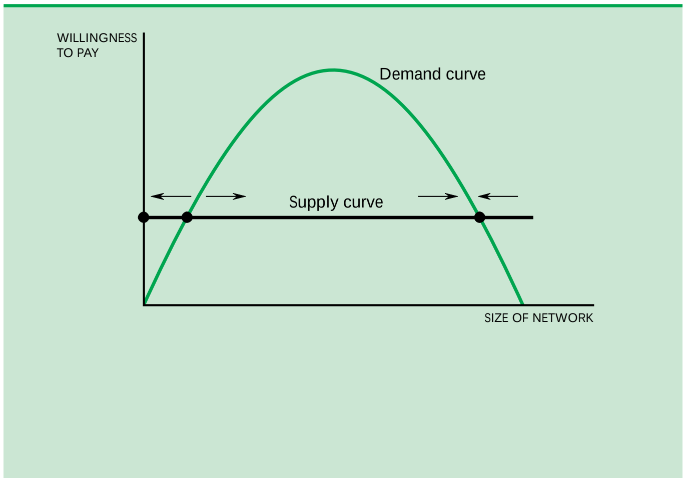


Figure 1: Network externalities. The demand is given by the curved hump, the supply by the horizontal line. Note that there are three intersections where demand equals supply.

willing to pay that much to connect to it—and therefore the network isn't very big.

Finally the last equilibrium has a large number of people, $n_H$. Here the price is small because the marginal person who purchases the good doesn't value it very highly, even though the market is very large.

# 1 Market Dynamics

Which of the three equilibria will we see occur? So far the model gives us no reason to choose among them. At each of these equilibria, demand equals supply. However, we can add a dynamic adjustment process to help us decide which equilibrium is more likely to occur.

It is plausible to assume that when people are willing to pay more than the cost of the good, the size of the market expands and, when they are willing to pay less, the market contracts. Geometrically this is saying that when the demand curve is above the supply curve, the quantity goes up and, when it is beneath the supply curve, the quantity goes down. The arrows in Figure 1 illustrate this adjustment process.

These dynamics give us a little more information. It is now evident

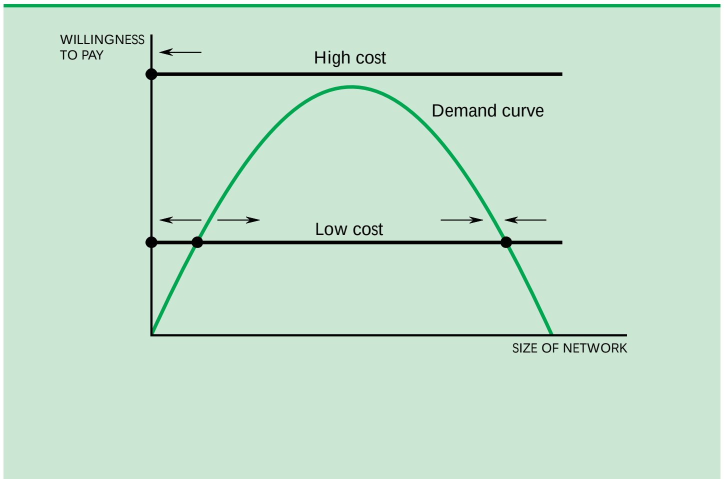


Figure 2: Cost adjustment and network externalities. When the cost is high, the only equilibrium implies a market of size zero. As the cost goes down, other equilibria become possible.

that the low-level equilibrium, where no one connects, and the high-level equilibrium, where many people connect, are stable whereas the middle equilibrium is unstable. Hence it is unlikely that the final resting point of the system will be the middle equilibrium.

We are now left with two possible stable equilibria; how can we tell which is likely to occur? One idea is to think about how costs might change over time. For the kinds of examples we have discussed—faxes, VCRs, computer networks, and so on—it is natural to suppose that the cost of the good starts out high and then decreases over time due to technological progress. This process is illustrated in Figure 2. At a high unit cost there is only one stable equilibrium—where demand equals zero. When the cost decreases sufficiently, there are two stable equilibria.

Now add some noise to the system. Think of perturbing the number of people connected to the network around the equilibrium point of $n^* = 0$. As the cost gets smaller and smaller, it becomes increasingly likely that one of these perturbations will kick the system up *past* the unstable equilibrium. When this happens, the dynamic adjustment will push the system up to the high-level equilibrium.

A possible path for the number of consumers of the good is depicted in Figure 3.

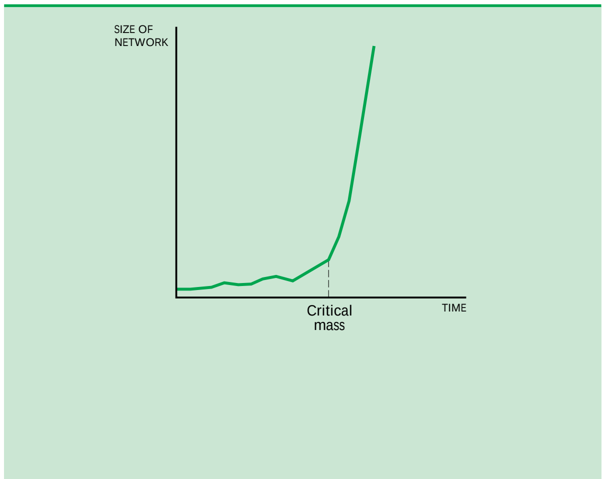


Figure 3: Possible adjustment to equilibrium. The number of users connected to the network is initially small, and increases only gradually as costs fall. When a critical mass is reached, the network growth takes off dramatically.

It starts out at essentially zero, with a few small perturbations over time. The cost decreases, and at some point we reach a critical mass that kicks us up past the low-level equilibrium and the system then zooms up to the high-level equilibrium.

A real-life example of this kind of adjustment is the market for fax machines. Fig 4 illustrates the price and number of fax machines shipped over a period of 12 years.[2]

[2]This diagram is taken from "Critical Mass and Network Size with Applications to the US Fax Market," by Nicholas Economides and Charles Himmelberg (Discussion Paper no. EC-95-11, Stern School of Business, N.Y.U., 1995). See also Michael L. Katz and Carl Shapiro, "Systems Competition and Network Effects," *Journal of Economic Perspectives*, 8 (1994), 93–116, for a nice overview of network externalities and their implications.

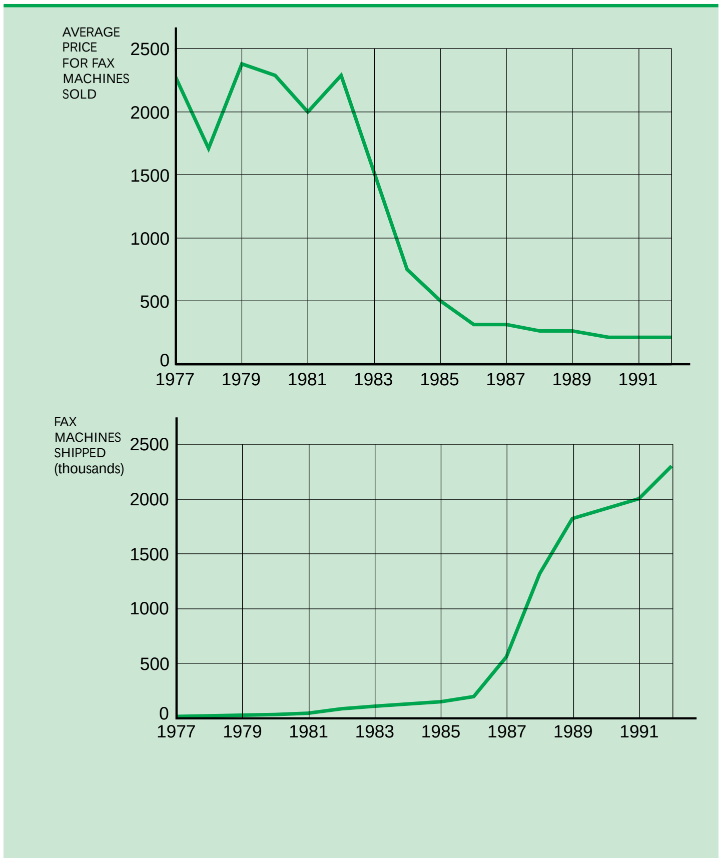


Figure 4: Fax market. The demand for fax machines was small for a long time since so few people used them. During the mid-eighties the price fell significantly and the demand suddenly exploded.