

Social Networks and the Diffusion of Economic Behavior

So you've decided you want to go to the movies with a friend. How will you go about deciding what to see? You'll probably start by asking a few friends what they have seen. Once you've gotten a few opinions, you'll have to consult the friend accompanying you, who also most likely spoke to several of her friends, to reach a final decision. If there is one particular movie that nearly everyone you know has already seen, you'll probably be tempted to see it, even if only to find out what all the buzz is about. In effect, the decisions of your friends will probably have a major and possibly even determining impact on your own choice.

The movie example is just one of many settings – ranging from decisions of taking up smoking to the chance of catching a computer virus – in which the choices of your social, professional, and familial network may have a powerful influence on your own behavior. In general, there are a plethora of causes for this phenomenon. Most obviously, acquaintances' decisions may directly affect the usefulness of certain choices. In the movie example, part of the enjoyment from a film is not just seeing it, but also exploiting it as a topic of conversation with friends. Or take learning a language. The benefits of this difficult undertaking depend on how many people one will end up interacting with who speak that language. Second, it may be that people learn useful information from their peers, which in turn affects what products or actions people choose. Initial word of mouth, for instance, can play a major part in determining whether a newly published novel will tread the path of commercial success or failure. Learning can also be very important in determining whether people are aware of various consumer opportunities, such as social welfare programs and student loans. Third, there can be direct contagion effects from certain behaviors. This is particularly relevant when thinking of the spread of computer viruses and human diseases. The more infected friends that a given individual has, the more likely that individual is to become infected. Deliberately-designed marketing plans also exploit similar paths, with such schemes as free calls to other cell phone users who have the same provider.

The unifying feature of all these examples is that individuals' local social networks can have a major impact on their lives. A person's behavior is most heavily influenced by the other people with whom they are in contact on a regular basis. Simple supply and demand style analyses are no longer appropriate for analyzing behavior in real-world settings. Studying such situations requires an understanding of the basic social network structure at hand and its interplay with economic decisions.

Social Networks and Tipping Points

To illustrate how this approach represents a potentially essential change in market modeling, consider the decision of an individual of whether or not to buy a new product (e.g., a software package), where the benefits to that individual from using that product increase with the number of friends that the individual has who also use the product. In such a situation, individuals can not be treated simply as clones. People will differ in a variety of ways that affect their decisions of whether or not to purchase the product. In a

practical sense, they might face different challenges when trying to acquire and use the new product, or they simply might have different propensities to like or dislike the product. From the social perspective, they might differ in terms of the number of friends with whom they interact, thus making the product more valuable to certain individuals.

The most readily apparent result of this added social factor is that, unlike in simple supply and demand models, it allows for multiple equilibria. This means that the final outcome of a process may end up with any of several different proportions of the population eventually adopting the product. For instance, it could be that nobody buys the product, and since nobody has any neighbors who have bought the product, the product fails. Alternatively, the product can end up with some isolated adoption, so that there are pockets of adopters, but the product only spreads sporadically. It could also be that it enjoys widespread adoption through a nontrivial percentage of the population. All of these three types of outcomes may be possible for a single given product, depending on the initial adoption rates; unless a certain initial threshold or "tipping point" is reached in terms of first adopters of the product, the product may not spread significantly. Moreover, these ideas, when applied to "products" like education, can help us to understand persistent differences of behavior across different social networks and provide a new angle on how poverty traps can emerge.

The eventual diffusion of adoption – the final percentage of the population that the product reaches – is dependent on several key parameters. Clearly, it is dependent upon the basic attractiveness and relative cost of the product. Lowering the cost and/or increasing the attractiveness of the product will, even according to traditional supply and demand models, help the product to diffuse in the initial stages. These changes, however, do not guarantee widespread adoption – the eventual outcome of the diffusion process depends on the social structure itself.

To understand how the social network structure matters, let us consider a couple of illustrative comparisons that we can make between differing social structures. First, we can think of looking at the average number of friends or interactions individuals have. This turns out to differ dramatically across populations, as it is contingent on gender, ethnicity, age, profession, and even on how one the network connections – through friends, co-workers, or acquaintances. For example, in a variety of studies looking at World Wide Web networks of sites (where two sites are connected if at least one of the two links to the other), the average number of neighbors is on the order of magnitude of 5. In contrast, data on romantic connections at representative high schools suggest that the average number of partners that students have per year is lower than 1.

The precise underlying social structure is important in terms of how the population reacts to different products, diseases, and innovations. Beyond how many connections individuals have on average, social networks also differ in terms of the variation across individuals. Is it that everyone has twenty friends, or do some have ten and others thirty? How different is the behavior likely to be of those who have ten friends, from those with twenty and those with thirty? As the settings change, different patterns emerge. For

example, there are higher variations in the number of connections that individuals maintain in some professional networks than in some purely social (friendship) networks.

When Popularity Breeds Consent

To see exactly how the variations in social structure impact the diffusion process, let us concentrate on situations where individuals' likelihood to adopt the behavior (e.g., buy the product) increases with the number of friends they have who have done so. Suppose that we randomly seed the population with some initial adopters of the product. People who have the most friends are the most likely to have friends that are initial adopters. Thus, all else held equal, people with more friends are more likely to then become adopters at a second stage. As we continue to follow the adoption process over time, we see a general propensity for people with greater numbers of friends to adopt earlier on and eventually at higher overall rates (presuming that people with more friends are similar to those with fewer friends along other dimensions). This prediction is consistent with many empirical studies of diffusion, such as the classic study of the adoption of a new drug by doctors in a famous early study by Coleman, Katz and Menzel (1966). 125 internists, pediatricians, and general practitioners in Illinois were followed over the span of 17 months in 1953-55. Their decisions of when to begin prescribing a new antibiotic, tetracycline, were recorded. After 8 months, those who were in touch with 3 or more doctors who had adopted the drug were more than twice as likely to adopt than those not in touch with any adopting doctor. At the end of the study, after 17 months, those who were in touch with 3 or more doctors who had adopted were approximately 20% more likely to adopt than those who were not in touch with an adopter.

The fact that, all else held equal, more highly connected individuals have higher propensities to adopt suggests a natural comparison between social structures. For instance, if we examine denser social networks that have increased connectedness of individuals, then we expect lower tipping points—how easy it is to get the process started—and higher eventual rates of adoption of the product or behavior. The diffusion has a sort of *social multiplier* associated with it. What this means is that once one makes it more likely to have some individuals adopt, it then becomes more likely that their neighbors will also adopt, and so on. In this situation, but slightly increasing the density of social networks, we see a dramatic change in the properties of the system as we leap from one setting where there is no diffusion to another where there is widespread diffusion. As such, it is possible that a relatively small change in the density of social interactions within a population can dramatically affect the end result.

Figure 1 illustrates these sorts of effects. It depicts the connection between the fraction of adopters at one period and the fraction of adopters in the period that follows. Intersections with the 45 degree line correspond to "equilibria" of the system, since they identify points at which the population gets stuck—in other words, when the number of adopters remains the same across periods. If the current fraction of adopters is at a point where the curve is above the line, then it leads to more adopters in the next period. If the

current fraction of adopters is at a point where the curve is below the line, then it leads to fewer adopters in the next period. Looking at the example depicted in Figure 1, we see that there are three intersections - one at 0 and two above 0. In particular, the population can remain with no adopters for an indefinite amount of time, but can also sustain a positive fraction of adopters at a steady state. The intermediate intersection point corresponds to a tipping point – if even a tiny fraction of adopters were added at this intermediate stage (say, through a promotional giveaway of the product), the curve would move above the 45 degree line, indicating that the number of adopters would continue to grow until the next intersection point was reached.

A denser social network changes the relationship between the fraction of adopters across time, and shifts the curve to the dotted line. Interestingly, the dotted line now has only one intersection point with the 45 degree line at full adoption (100 percent of the population adopts). This is the case when a sufficient number of links is added to make everyone adopt at the only possible steady state.

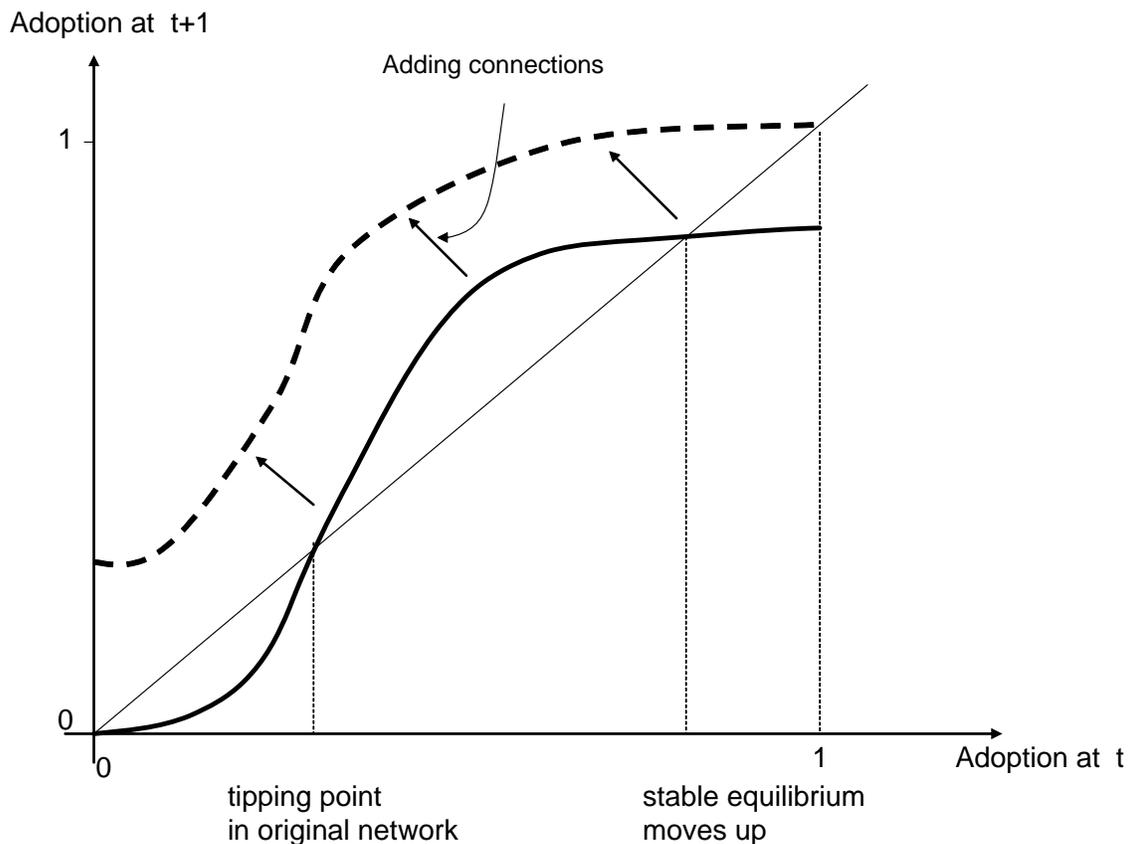


Figure 1: Dynamics of Adoption

The Popular Crowd and the Loners

As mentioned above, social networks not only differ in terms of their density (how many friends or acquaintances any given individual is likely to have), but also in terms of how much variation there is in connectedness across individuals. This variability can also have a profound impact on the diffusion of a behavior throughout a social network, although it is a bit more subtle than the direct effect of increasing the density of the social network.

To get some feeling for this, consider two different social networks, each having the same average number of connections, but one being a “mean-preserving-spread” of the other. That is, one has the same number of connections but its corresponding distribution of connections is more unequal among the population, so that there are relatively more people with very few connections and relatively more people with very many connections. Roughly, we have taken a network that was very regular in structure, and changed it to one that looks more like a “hub-and-spoke” network. This changes the diffusion properties as follows: the people with a large number of friends are more likely to be earlier adopters, regardless of adoption that we seed the network with; and people with very low numbers of friends are less likely to be adopting at any given adoption rate.

These two categories of people interact with each other and affect the overall adoption rate. For instance, if the highly connected individuals become adopters at the beginning, then they serve as conduits, and their high adoption rates mean that the less connected individuals become more likely to have friends who have adopted. The adoption then spreads throughout the network.

Whether the more variable network or the more regular network has lower tipping points and higher adoption rates (as appearing in the figure corresponding to the addition of connections) actually depends on readily identifiable curvatures (convexity measures) of the adoption process. In other words, the crucial determinant is how one’s likelihood of adoption depends on the variability of the connectedness of her neighbors (the number of friends her friends have). This dependence determines whether the increased rate of adoption among the highly connected is high enough to counteract the decrease in the initial adoption rate among the less connected.

The Hype Builds

The modeling that we have discussed of economic behavior within a social context not only makes predictions about things like tipping points and eventual adoption rates, but also allows one to predict diffusion over time. As already indicated, we see patterns indicating that those who are more connected adopt earlier and at higher rates overall. But there are also specific time dependent patterns of the diffusion of behavior. For example, one phenomenon that has been well-documented in diffusion studies is what is known as an “S-shaped” adoption curve: adoption starts out slowly, then gains

momentum and increases in speed, and then eventually slows down again. This has been observed in scenarios as diverse as the first purchases of electronic devices to the first uses of hybrid corn seeds among farmers.

Diffusion of a behavior exhibits such S-shaped patterns for intuitive reasons that we can understand from the social-network perspective. The initial adopters are those who choose the behavior for reasons largely independent of the social context. They might be recipients of a sample product, or those who simply want to try a product regardless of what their social neighbors do. This slowly gets the adoption process started. Consequently, social diffusion and the social multiplier kick in. The initial adoption leads those with many friends to start adopting. This increased adoption now means that more individuals have friends who have adopted which leads to more adoption, and so forth. The property that the more neighbors there are choosing an action, the more desirable that action becomes (formally termed as social complementarities) assures that there will consequently be an increase in the rate of spread. Eventually, the process begins to play itself out, until the remaining people who have not adopted the behavior are those with high costs or low benefits from the product or behavior or those who have a low number of friends. Thus, the process eventually hits diminishing returns and slows down.

While diffusion of products and innovations is a well-studied area, the mathematical modeling necessary to tie explicit social network structure back to the diffusion of economic behavior is just being developed. Results from the models that we have outlined here (and which are studied in detail in the references below) provide progress in that direction, helping to explain an array of empirically observed phenomenon ranging from patterns of disease contagion to hybrid corn adoption in rural areas of the U.S., to financial bank runs. The analysis also leaves some fascinating open questions for further exploration – how are networks formed in the first place? Empirically, do agents connect to friends similar to them in particular dimensions, and if so, why? What do agents know about the networks they participate in? These are but few questions this investigation suggests.

Matthew O. Jackson is a Professor of Economics at Stanford University and Leeat Yariv is an Associate Professor of Economics at the California Institute of Technology.

They thank co-authors Antoni Calvo-Armengol, Andrea Galeotti, Sanjeev Goyal, Brian Rogers, and Fernando Vega for collaboration on related projects, and Nilesh Vashee for helpful editorial comments.

Some of the models underlying the discussion above can be found in the following papers: Jackson, M.O. and L. Yariv (2006), ``Diffusion of Behavior and Equilibrium Properties in Network

Games'' Caltech HSS Working Paper # 1264, forthcoming in the American Economic Review (papers and proceedings). Jackson, M.O. and L. Yariv (2005), ``Diffusion on Social Networks,' ' Economie Publique, 16(1), 69-82. Jackson, M.O. and B. W. Rogers (2006,) ``Identifying the Impact of Interaction Structures on Diffusion Properties Through Stochastic Dominance Relationships," forthcoming in Advances in Theoretical Economics, BE Press. Galeotti, A., S. Goyal, M.O. Jackson, F. Vega-Redondo, and L. Yariv (2005), ``Network Games,' ' Caltech HSS Working Paper.