CHAPTER

8

Viscosity Models and the Diffusion of Controversial Innovation

David Krackhardt

As March has indicated in his foreword to this book, the promise of computational models in organizational theory has not been fulfilled. One reason is that too much was promised, such that computational models could not possibly deliver. I have argued elsewhere that computational models are particularly good at developing theory—that is, suggesting the logical consequences of a set of assumptions (Krackhardt, 2000). Most human-generated theories are limited to hopelessly simplistic linear additive assumptions about social phenomena, phenomena we all accept to be complex, dynamic, with feedback loops that make systemic behaviors very difficult to predict from a linear model. But, computational models do not prove these theories they help develop; they are not empirical by that standard. They merely help the researcher to answer logical "what if" questions. Expectations that computational models can *demonstrate* or *prove* anything beyond theory building is asking too much of them and will lead to disappointment.

The other downfall of computational models has been that they have often concluded the obvious, conclusions that could have been derived easily by a human's limited linear thinking. However, there have been historical examples to counter this problem, such as March and Cohen's Garbage Can model. It is in this spirit, of developing a theory of diffusion of innovation that incorporates nonlinear dynamics and of developing a theory that has nonintuitive implications for the diffusion process, that I write this chapter.

*Krackhardt, David
2001 "Viscosity Models and the Diffusion of Controversial Innovations" in Dyanmics of Organizations: Computational Modeling and Organizational Theory, Alessandro Lomi and Erik R. Larsen (eds.), 2001, MIT Press, pp. 243-268.

It has long been acknowledged that the diffusion of innovations is a social process. That is, new ideas, new technologies, new management practices are diffused through a set of relationships that exist or emerge among actors within an organization or social system (Allen 1977, Price 1965). As Rogers (1982) noted in his comprehensive review, thousands of studies have been conducted on innovation diffusion, but there have been relatively few that have shown *how* the social structure affects the diffusion process.

There have been notable exceptions to this main thrust, however (Coleman, Katz, and Menzel 1966; Becker 1970; Burt 1973; Burt 1980; Carley 1990).These seminal pieces have emphasized how diffusion can be mapped out through structural features of the social system. For example, Coleman, et al., found that diffusion of a prescription drug was heavily influenced by direct social ties among physicians. Burt (1987), on the other hand, reanalyzed the same data and concluded that the diffusion mechanism was better explained by physicians mimicking each other when they were structurally equivalent (connected to the same set of third parties) rather than connected to each other directly.

A different stream of research has explored a "threshold" model of diffusion (Granovetter 1978). In these models, it is assumed that individuals influence each other toward adopting but that they have different thresholds—that is, some individuals will adopt after only a small proportion of their alters has adopted, while others will not adopt until a large proportion of their alters has adopted. With such a set of assumptions, one can model the diffusion process and recreate the standard "S" curve that is commonly associated with the diffusion process (Granovetter 1983, Granovetter 1978, Granovetter and Soog 1988). T. W. Valente's (1996) recent work in this area is an excellent adaptation of this approach. His approach focused on the ego network for predicting adoption rather than on the density of adoptions throughout the overall network.

Valente empirically showed that early adopters were much more likely to adopt in the face of few neighbors adopting, and conversely "laggards" adopted only after a relatively high proportion of neighbors had adopted. Again, diffusion is mapped out as a function of the structure of relationships among the adopters of various types.

A third stream of work that has hallmarked the diffusion literature has employed computer simulation as a modeling technique. The advantage of the computer is that one can model very complex systems that analytic or simple theoretical models cannot handle (Krackhardt, 1999). Thus, one can conclude that, if a set of assumptions holds (and is modeled appropriately), the diffusion rate should take on a particular shape. This work is epitomized by Carley's structuralism models of group formation and diffusion (Carley 1991, Kaufer and Carley 1990, Carley and Wendt 1991). These models assume that knowledge and beliefs diffuse as a function of what the actors have learned from each other in the course of structured interactions. Carley has found that interesting and complex dynamics can result by restricting access between two groups of individuals (Carley 1991). Another stream of work that explores structural features of diffusion success is that of Abrahamson and Rosenkopf (1993, 1997). They have used computer simulations to explore how structural conditions affect the adoption rates of innovations that have negative consequences for the organization.

A common theme among all of this diffusion research has been to assume that innovations will eventually diffuse throughout the population. Indeed, it is common to explicitly restrict the realm of interest in diffusion studies to cases in which the innovation successfully dominated the organization:

> The diffusion of innovations is the process by which a few members of a social system initially adopt an innovation, then over time more individuals adopt until all (or most) members adopt the new idea.... [Valente, 1996:70]

Innovations do not always succeed in diffusing, however. One of the most notorious examples of a "good" idea that refused to diffuse was the original PC, first proposed by a group of R&D engineers at IBM long before any commercial versions were available. The PC promoters were housed in the guts of IBM's research center in Tennessee. Twice they tried to promote the idea within IBM's structure, and twice they were defeated, allowing Apple to gain a substantial advantage in the market. It was not until years later when the PC developers were transplanted to a separate location in Florida that the PC flourished as an IBM product. Following on the work of Abrhamson and Rosenkopf mentioned earlier, the question I would like to propose is, why do some innovations succeed when others fail to diffuse in a social system? More specifically, under what conditions will such innovations diffuse and under what conditions will such innovations be stopped?

The Nature of Innovation Diffusions

Before specifying what conditions might be influential in this process, it is useful to differentiate types of innovation diffusion processes based on the ease with which they are accepted and adopted. First of all, consistent with Carley's work, I consider innovations to be inherently ideational. That is, actors adopt innovations because they come to believe that it is beneficial to do so. Therefore, the process of innovation diffusion is one of converting people into the belief that the innovation is in fact a good idea.

Given this ideational premise, I suggest that there are two forces that can lead to such a conversion. First, there is the exogenous inherent quality of the innovation itself. This force stems from the intrinsic strength (or failing) of the innovation. It is exogenous in that the evaluation that leads to possible conversion is not the result of political, influential, or other endogenous social forces; rather, that the decision to adopt or not is based solely on the intrinsic merits of the innovation itself. That is, conversion happens because each person independently and objectively recognizes a better mousetrap when one sees it. I will refer to this evaluation/conversion process as "rational" to emphasize its objective and nonsocial process of evaluation (I do not mean to imply anything about the decision maker's utility function here). This "rational" evaluation results in one of two extreme behaviors: either those who are made aware of the innovation immediately see the superior value in the innovation and adopt it; or, those who are made aware of the innovation immediately see its inferior value and reject it in favor of the status quo or some other alternative.

The other force is a social one, wherein the innovation's value is not so clearly determined by external or objective measures. This force suggests either that the quality of the innovation is ambiguous or that even if it is objectively demonstrable people can be swayed through a social or political process to a counter position. Thus, in such cases, innovations are valued through a dynamic social process, wherein people influence each other as to their evaluation of the "true" value and as to whether the innovation should be adopted. Innovations that are subject to this process of evaluation and conversion may be termed *controversial*. The conversion force is quite different from the "rational" ones because potential adopters' minds can be changed back and forth as they are exposed to different social forces from supporters on one side to the detractors on the other.

While it is an empirical question beyond the scope of this chapter to verify, I will go boldly out on a limb to suggest that most innovations fall into this latter, controversial category. Most innovations have identifiable advantages, which supporters of the innovation can point to, addressing some need or shortfall in the status quo. But the flip side is that these innovations threaten the status quo. Consequently, there will be detractors. Few innovations are so clearly and markedly inferior or superior to the status quo as to overcome this inherent conflict. Thus, most innovations are controversial and subject to a social process of convincing others of their superiority over the status quo.

Early information processing models (e.g., March and Simon, 1958) have made similar claims. Organizations accrue competencies by executing stable routines to perform frequently-encountered tasks. People learn by doing; they get stuck in this rut. Innovation based on a new technology they are unfamiliar with becomes a threat to their knowledge base. Thus, even in cases where technological superiority can be easily recognized by an outsider, the innovation may be resisted by the insider whose job must change because of it. For example, France has proven to be one of the slowest among developed countries to adopt the internet. The reason for this is probably that the country has had Minitel since the early 1980s, which is in many ways a lesser technology. However, people have experience using Minitel and trust it to perform many of the functions that the web offers, and so adoption of the web has been slow.

Others have made similar claims. For example, Cohen and Levinthal's notion of absorptive capacity suggests that innovation occurs not simply because the technology is superior but because the firm has invested in R&D's ability to understand and absorb (or perhaps accept) this new technology. Another example is the work of Brian Arthur and his colleagues whose collective work has underscored the critical role of social structure and context in the diffusion process even though the innovation may have clear technological superiority. Thus, a strong argument can be made that most innovations are in fact controversial, despite their apparent objective appeal.

Having said that, it may still be that some innovations are more likely to be seen as controversial than others. Indeed, innovations on organizational procedures and routines, such as re-engineering and TQM, are almost universally regarded as controversial. Other more technologically based innovations, such as the Xerox™ photocopying process, may appear as less controversial. While these distinctions are not hard and fast, I will restrict myself in the remainder of this chapter to controversial innovations as the main focus of interest.

We can plot the progression of all three types of innovation diffusion patterns resulting from both rational and social forces in terms of the rapidity with which they diffuse or retreat. To do this, we note that for any given time period, *t,* there will exist a proportion of the population that will "agree" with the innovation—that is, they hold the belief that the innovation is a good idea. Since we are concerned here with ideational aspects of innovation, I will call such people adopters and not differentiate them from people who agree with the innovation but have not yet behaviorally acted on it. Further, I will assume for the purposes of simplicity that people either agree or disagree with the innovation; that is, I assume everyone can be classified as either an adopter or a nonadopter.

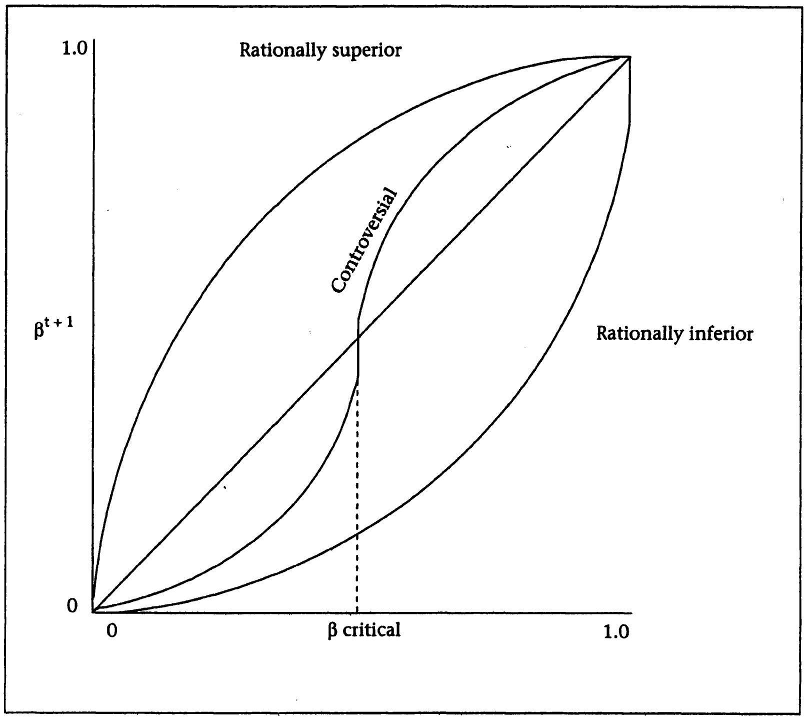

Figure 1. A plot of the progression of all three types of innovation diffusion patterns.

If we designate the proportion at time *t* who are adopters as β^t then we can plot the dynamics of this process for all three innovation types as in figure 1. The proportion of adopters at time *t* (β^t) is represented on the horizontal axis, while the proportion of adopters at time $t + 1$ (β^{t+1}) is represented on the vertical axis. Any point along the 45° line would represent a steady state solution, where no changes occur in the proportion of adopters in the population from one time period to the next. Points above the line would indicate that the proportion of adopters was increasing over time; points below the line would indicate that the proportion of adopters was decreasing over time.

In this figure, we see the fundamental difference between rational and controversial innovation diffusion processes. In the top curve, the "rationally superior" innovation is plotted. In this case, no matter what proportion of believers you start with (as long as there is at least one adopter in the population), each subsequent time period will have a higher proportion of converts to the innovation, until everyone is an adopter. The bottom curve represents the "rationally inferior" innovation. In this case, no matter what proportion of exogenously-determined adopters you start with, there will be fewer during the next time period, until all the adopters have converted to nonadopters.

The controversial innovation, however, has no automatic outcome. In this case, there are three possible outcomes, as represented in the S curve in the middle of figure 1. There is a saddle point on the curve, marked β^t_{crit}, located on the 45° line, indicating that each subsequent time period will have the same proportion of adopters. In this case, people may convert from adopters to nonadopters, but to the extent that they do they will be exactly balanced by the conversion rate of nonadopters to adopters, so that no net gain or loss in adoption proportions is observed. The intuition here is that, since adoption is a function of the social forces around you (rather than personal and objective experience with the innovation itself), the number of proponents and the number of opponents of the innovation are precisely balanced at β^t_{crit} so that neither side has the ability to convert more than the other side. Thus, an unstable equilibrium (standoff) is reached wherein both sides retain the same strength of support.

But, as we move off of β^t_{crit} we see that there is a pronounced effect on the outcome. Any starting point below β^t_{crit} results in successively fewer adopters in each time period, until extinction is reached. Any starting point above β^t_{crit} results in successively more adopters in each time period, until saturation is reached. Thus, the success of the innovation depends not on the quality of the innovation itself but rather the ability of the innovators to establish a critical mass of support for the innovation, β^t_{crit} after which they are virtually certain that the innovation will eventually dominate the social system.

One important observation should be made here: Once a critical threshold of density of adopters is reached, one could argue that the remaining people are simply acting "rationally" by adopting, seeing that the result is inevitable and not wanting to be left out of the wave of the future. While this is indubitably possible, such a process is outside the scope of the model I am proposing in this chapter. For simplicity's sake, this model incorporates only a local conversion process, as a function of who interacts with whom.

In terms of dynamics and equilibria, the two rational innovations are trivial in their solution. The controversial innovation, on the other hand, is both interesting and complex. But the question left unanswered is, how is it possible for the adopters to attain β^t_{crit} short of a massive infusion of exogenously determined support for the innovation? That is, is it possible for a smaller group of adopters to reach β^t_{crit} and "take over" the organization? The rest of this chapter is devoted to the answer to this question. Without relying on any assumptions of political motives or moves on the part of any of the actors in the system, I demonstrate that the. structure of interactions among the participants can itself lead to surprisingly stable and counter-intuitive results To demonstrate this, I draw heavily from the work of Boorman and Levitt (1980}, who asked a similar question about how is it possible for altruistic genes to propagate in a population of animals when the gene itself puts the host at greater risk for survival. While it is not my intent to explicate their model in this chapter, the curious reader will find their work to have a similar mathematical basis and very similar results to those I uncover for this problem.[1]

The Model

I start with a set of simple axioms about the diffusion process at the micro level in our social system. I will use the organization as a metaphor for any social system to which this diffusion may apply. The organization is populated by an arbitrarily large number of two kinds of persons, adopters (those who currently believe in the innovation) and nonadopters (those who currently do not believe in the innovation). Within any given time period, each person actively seeks out a set of others within the local part of the organization in which they currently find themselves and confers with those others on their beliefs about the innovation. If among those others they find agreement with them on the value of the innovation, they will retain their own belief about it. If, on the other hand, they find themselves surrounded by those who disagree with them, then they will tend to convert to the other belief (change from being an adopter to being a nonadopter, or vice versa).

To formalize this model, let us make the following assumptions:

Assumption 1: Each individual adopter (a person who believes in the value of the innovation) searches randomly through L_a others to find another like-minded individual. Each individual nonadopter (a person who believes that the innovation does not have value) searches randomly through L_n others to find another like-minded individual. I assume that the innovators are more likely to proselytize the status-quo oriented nonadopters than the converse; therefore, $L_a > L_n$.[3]

Assumption 2: (The Asch assumption): If in the process of this search, an individual finds at least one other individual who agrees with them (i.e., an adopter interacts with one other adopter, or a nonadopter interacts with one other nonadopter), then the individual will retain their current belief. This assumption acknowledges the work of Asch

{1951}, who found that it only required one person to agree with the subjects of his experiments to allow them to retain their beliefs, no matter how many confederates disagreed with the subjects.

Assumption 3: If an adopter fails to find at least one other adopter in the course of their search, then the adopter will convert to being a nonadopter with probability σ. This is the probability of conversion from adopter to nonadopter for those adopters who find themselves isolated. Again, Asch's work supports this assumption.

Assumption 4: If a nonadopter fails to find at least one other nonadopter in the course of their search, then the nonadopter will convert to being an adopter with probability τ. This is the probability of conversion from nonadopter to adopter for those nonadopters who find themselves isolated.

These four assumptions are all that are necessary to drive the shape of the S curve (the "controversial innovation" curve) in figure 1. In particular, these values will determine how large the threshold β^t_{crit} is. The larger the threshold to be surmounted, the more difficult it will be to successfully diffuse the innovation throughout the entire organization. The larger L_a is relative to L_n, the lower will be β^t_{crit} the larger τ is, the lower will be β^t_{crit} the larger σ is, the higher will be β^t_{crit}. But, no reasonable values of these parameters will permit an arbitrarily small minority of innovators to convert, through this minimalist process, the entire organization.

Organizational Viscosity

I now introduce the idea of structure into our organization. Embedded in the prior set of assumptions is the notion that each individual performs a random search through the entire organization. In fact, people are usually confined in their interactions to those more locally accessible (Simon 1962). It has been argued elsewhere that structural differentiation can easily affect the diffusion process (Hagerstrand 1967), although our argument will rely on social structural forces rather than spatial ones such as Hagerstrand used. For simplicity, I will assume the organization can be partitioned into subsets, which I will call groups, and that people only search within their own group. Again, for simplicity, each group will be assumed to contain the same number of people. Diffusion within each group, then, simply becomes a smaller problem of the one tackled earlier. Within any group, if the proportion of adopters exceeds β^t_{crit} then the adopters will win over that group; otherwise, they will become extinct within that group.

Carley (1991), in her study of signal analysis, showed that the presence or absence of migratory links between groups, and the length of the chain connecting two groups, determined the rate and pattern of adoption. Likewise, if we assume some small degree of mobility among groups, we offer the opportunity of the problem becoming much more complex and interesting than one encounters in the case where everyone is free to interact with everyone else in the organization.

The restrictions on mobility will be monitored through two mechanisms. First, only certain pairs of groups will be characterized as exchanging individuals at all. The pattern of these exchange possibilities will constitute the structure of the permissible flow of people in the organization. Second, the rate at which individuals are allowed to migrate along these exchange paths from group to group will be controlled by a parameter (v), a rate which reveals the extent of organizational *viscosity*.[4] For purposes of this chapter, we will restrict ourselves to some simplifying assumptions:

Assumption 5: For each period t, prior to individuals searching for like-minded others, a certain fraction of the group's inhabitants will migrate to another group. The migration rate from group i to group j will be given as $M_{ij} = vS_{ij}$, where S_{ij} is a symmetric adjacency matrix designating the structure of possible exchange relations between all group pairs, and v is the fraction of the group that migrates to the adjacent group each period.

It is important to point out that people are not being modeled individually here. Rather, the group is the unit of analysis, and the measure of interest is the proportion of adopters within the group. Moreover, all groups are assumed to be the same size, and this size does not change over time since movement between groups occurs equally in both directions.

Dynamic Details

There are two steps in the model. First, a certain proportion of individuals migrate to their new groups. Second, after a migration has taken place, a certain proportion of individuals may convert to being either adopters or nonadopters.

Step 1: Migration

We now have enough information to determine the dynamics of this diffusion process. Let β_i^t be the proportion of adopters in group i. Let M_{ij} be the proportion of group i that migrates to group j in one period (and vice versa). I will refer to all groups that directly exchange people with group i as groups *adjacent* to i (i.e., all groups j such that $M_{ij} \neq 0$).

The fraction of adopters who left due to migration is the sum of the product of migration rate for each adjacent group and the proportion