Situating Social Influence Processes: Dynamic, Multidirectional Flows of Influence Within Social Networks

Winter A. Mason, Frederica R. Conrey and Eliot R. Smith

Pers Soc Psychol Rev 2007; 11; 279
DOI: 10.1177/1088868307301032

Situating Social Influence Processes: Dynamic, Multidirectional Flows of Influence Within Social Networks

Winter A. Mason
Frederica R. Conrey
Eliot R. Smith
Indiana University–Bloomington

Social psychologists have studied the psychological processes involved in persuasion, conformity, and other forms of social influence, but they have rarely modeled the ways influence processes play out when multiple sources and multiple targets of influence interact over time. However, workers in other fields from sociology and economics to cognitive science and physics have recognized the importance of social influence and have developed models of influence flow in populations and groups—generally without relying on detailed social psychological findings. This article reviews models of social influence from a number of fields, categorizing them using four conceptual dimensions to delineate the universe of possible models. The goal is to encourage interdisciplinary collaborations to build models that incorporate the detailed, microlevel understanding of influence processes derived from focused laboratory studies but contextualized in ways that recognize how multidirectional, dynamic influences are situated in people's social networks and relationships.

Keywords: *attitudes; social influence; social power; metatheory; group processes*

The study of social influence—the ways other people affect one's beliefs, feelings, and behavior—in large measure defines social psychology. Topics such as persuasion, conformity, and obedience deal obviously and directly with how other people influence one's thoughts and actions. More broadly, areas of social psychology such as social learning, relationship formation and maintenance, attitude and stereotype formation, group decision making, power, intergroup relations, and bargaining and negotiation also critically involve social influence. This topic truly lies at the heart of this scientific field.

Recognizing the importance of social influence, much social psychology research in the past 40 years, especially in the areas of attitudes and social cognition, has focused on processes of social influence, implementing tightly controlled laboratory studies to clearly define the impacts of manipulations on participants (Smith & Semin, 2004). The bulk of persuasion research, for instance, exposes participants to a persuasive message and then measures attitude or attitude change as well as potential mediators of the process. Similarly, research on conformity examines effects on the participants' own beliefs or attitudes, usually by exposing participants to experimenter-constructed information about what other people believe or desire.

Using well-controlled laboratory studies to isolate, manipulate, measure, and analyze social influence at the level of individual cognitive processes has many advantages, such as allowing for strong causal inferences. However, the approach is less informative about how social influence plays out in larger scale social contexts over time. Theorists in the area of complexity theory or

Authors' Note: Preparation of this article was supported by National Science Foundation Grants BCS-0091807 and BCS-0527249. Thanks to Elise Hall, Charlie Seger, Zak Tormala, and Stanley Wasserman for comments and advice on earlier versions of the article. Correspondence may be addressed to Eliot R. Smith, Department of Psychology and Brain Sciences, 1101 E. Tenth St., Indiana University, Bloomington, IN 47405; e-mail: esmith4@indiana.edu.

PSPR, Vol. 11 No. 3, August 2007 279-300
DOI: 10.1177/1088868307301032

"complex adaptive systems" (e.g., Resnick, 1994) have pointed out that simple, individual processes often combine to create complex systems with nonintuitive emergent properties. The laboratory-based, focused empirical approach is aimed at identifying the psychological processes that underlie people's judgment and behavior. However, no amount of intuitive consideration of those processes permits a thorough understanding of the complex, dynamic phenomena that can be observed when the processes are iterated across time and space (Wolfram, 2002).

Whereas social psychologists have typically taken a focused, process-oriented approach, researchers in other disciplines have taken a complementary approach, generally examining group- or population-level outcomes. For instance, economic, sociological, and political-science research on the diffusion of innovation often seeks to explain the proportion of a population that adopts a particular innovation (Valente, 1995) as opposed to the processes leading to change in any individual. Extensive research in sociology and physics has focused on social influence networks and the dynamics of opinion flow, measuring the speed of opinion convergence or the number of subgroups that emerge over time from a starting distribution of positions. In other words, this work examines important aggregate-level variables (such as the proportion of the population in a particular state) but bypasses the level of psychological process and ignores individual opinion trajectories.

We believe that an overall understanding of the nature of social influence will be attained only by integrating these two general approaches. Theorists in disciplines outside social psychology, recognizing that social influence is key to their concerns, are building models of influence largely without recognition of the extensive conceptual and empirical work done by social psychologists. Likewise, social psychologists' typical focus on microlevel processes will benefit from a broader consideration of the social contexts and networks in which influence occurs as well as the dimension of time. Because social influence is a central and even defining concern of social psychology as a scientific field, social psychologists should need little encouragement to play a central role in the emerging interdisciplinary effort to build an understanding of social influence in its broadest context. The major goal of this article is to outline a framework for building theories and models of social influence by conceptually integrating the extensive literature from social psychology, with its emphasis on understanding specific processes of influence, and the literature from other fields, with its focus on how processes interact over time to produce important aggregate outcomes.

In this article, we first discuss the two most fundamental features of the context of social influence: (a) multiple sources and targets of influence and (b) influence extending over time. Next we examine one potential consequence of extending social influence processes across individuals and time: potential collapse to a unanimous viewpoint. Many social influence models make it difficult to maintain realistic diversity of opinions as individuals interact and influence each other over time (Abelson, 1964). In the subsequent sections, we discuss four dimensions on which models of social influence differ, and in each case we devote special attention to how the dimension may help maintain the diversity of opinions. Finally, we end the article with our recommendations for specific types of models that we view as most promising and most worth further theoretical exploration, a discussion of methodological and conceptual tools that can enable such exploration, and guidelines for researchers to use those tools.

CONTEXTUALIZING SOCIAL INFLUENCE

Two fundamental components of social influence are sources and targets of influence and time. Typical models in social psychology consider only one source and one target at only one point in time. However, in important real situations, the spread of influence occurs through populations over a span of time with each individual serving as both a source and a target.

Multiple Sources and Targets

Models of social influence differ with respect to the types of influence they consider: in-person persuasion attempts, television advertisements, conformity to social norms in small groups, and so on. However, no model of social influence is complete without some explicit treatment of the structure and direction of links between sources and targets of influence. Many models in social psychology assume unidirectional influence links. A model of persuasion, for example, might suggest that Amy delivers a message to Brian, and Brian alters his opinion to be more similar to Amy's. When the assumption of a single source and single target is relaxed, however, it is easy to see that as the conversation continues Brian's opinion might influence Amy as well. In fact, their interchange might influence Amy and Brian's other friends, and so on. Social influence involving many sources and many targets occurs every time people converse with a group of friends, drink Brand X soda in public, or wear an "abolish capital punishment" t-shirt. Unlike many laboratory experiments, social influence in natural settings is inherently both multidirectional (involving multiple sources and multiple targets of influence) and dynamic (occurring over time).

Models that seek to adequately contextualize social influence processes must incorporate reciprocal and multidirectional influence pathways instead of tacitly assuming that one source influences one target only.

Some important cases of influence, such as influence by a televised advertisement on consumers, at first glance appear to be unidirectional. But even the apparent inertness of a televised message or a billboard is illusory. Research has established that much media influence is actually mediated by other people. That is, Amy sees an ad on TV and discusses it the next day with Brian and Candace. They may not have seen the ad themselves, or even if they had their reactions to it may be shaped by Amy's opinion. The effects of the ad on them (or on others with whom they converse in turn) is therefore indirect, influenced by the pre-existing opinions and other characteristics of individuals in the population—who therefore act as sources and channels as well as targets of influence, even in the case of media ads. The overall process has been termed the "two-step flow of communication" (Katz & Lazarsfeld, 1955; Lazarsfeld, Berelson, & Gaudet, 1944). Furthermore, although a specific media message itself appears to be static and uninfluenced by the recipients, in fact the content of media ads is heavily influenced by characteristics of the targeted population, which is taken into account in the process of message design and production. For example, ad attributes such as the use of nudity and even the nature of products that are advertised reflect existing attitudes of the population (e.g., nudity is much more common in ads in Europe than in the United States). Again, when the content of ads is shaped by the characteristics or attitudes of the target population, that population is a source of influence as well as a target.

Like an advertisement, a speaker at the podium might appear to have a unidirectional effect on his audience's beliefs and attitudes, but again the apparent independence of the influencer is misleading. During the question-and-answer period, audience views that are expressed explicitly or implicitly may influence the speaker's own thinking. More subtly, social psychological research has established that speakers regularly "tune" their communications in line with the known or expected attitudes of audiences. This tuning affects not only the messages they communicate but also the speakers' own later attitudes on the issue (McCann, Higgins & Fondacaro, 1991).

In short, unidirectional social influence may exist in special situations, but in most real-life circumstances, listeners and audiences affect speakers and advertising messages as well as the reverse. Models of influence in group decision making as well as models from outside social psychology typically assume that influence is reciprocal and multidirectional, a realistic assumption that allows researchers to investigate outcomes at multiple levels of analysis. Many influence models that do consider multiple actors, however, focus on aggregated outcomes such as the rate of spread of a new opinion through a population rather than on the psychological processes that underlie influence. Social psychology, with its expertise in such areas as biases, levels of processing, and social relationships, has much to bring to the study of influence over multiple actors.

Social Influence Over Time

The second fundamental component of social influence is time. Many social psychological models consider only a single time point: Amy urges Brian to drink Brand X cola, Brian acquiesces or refuses, and the story ends there. Like the focus of some models on a single transmitter and a single receiver of influence, the focus on a single moment of time is congenial to process-oriented laboratory studies. However, influence models that address only a single time point are limited because influence is an inherently dynamic process—particularly when one takes account of the reciprocal and multidirectional influences that can occur among multiple sources and targets. Thus, models that take a broader view and consider influence processes that dynamically unfold in time are likely to yield important insights.

Examples of influence processes that unfold over time include sleeper effects (Kumkale & Albarracín, 2004; Weiss, 1953), where attitude change fails to appear initially after an influence attempt but becomes evident at a later time as a discounting cue that originally limited persuasion is forgotten. Influence can evolve over time for another reason as well: because changes in one cognitive element (belief or attitude) may lead to changes in others. If social influence leads the target to change a given attitude (e.g., about the relative merits of Israeli versus Palestinian positions in their ongoing conflict), over time other related attitudes (e.g., views of the trustworthiness of media sources that take different perspectives on the conflict) may shift accordingly. In other words, because people's cognitions are linked and interdependent, often a change in one will lead to corresponding adjustments in others as described, for example, by cognitive dissonance theory (Festinger, 1957) or parallel constraint satisfaction models (Kunda & Thagard, 1996; Smith, 1996). In general, current views of cognitive processing (e.g., dual-process models; Smith & DeCoster, 2000) reject the picture of an attitude as a static, isolated representation that can be changed by a specific manipulation such as a persuasive message but otherwise sits inert. Instead, attitudes as mental representations actively participate in and bias ongoing processing, with potentially complex effects over time. For all these reasons, researchers

should be thinking of influence not as a phenomenon that occurs at a single point in time but as a process that has an extended time course.

Feedback and Other Dynamic, Nonlinear Processes

When both time and multiple sources and targets are included in an influence model, many possibilities arise for interesting dynamic behavior. An important example is feedback loops. If Amy influences Brian's attitude, then Brian may influence Amy in the same direction at the next moment, and so on. Such recurring, reciprocal influences can result in a positive-feedback spiral that will continue until some other process arrests it. This type of process is found in marital interactions within distressed couples (Gottman, 1994). When one member of the couple makes a negative comment, the other is likely to respond with further negativity, which is reciprocated by the first, and so on. The resulting feedback loop of negative affect is one characteristic behavioral pattern of such couples. Models assuming that influence in both directions combines only in a simple, additive fashion would be unable to capture this dynamic process.

A recent study illustrates important real-world implications of feedback because of multidirectional social influence over time. Cultural marketplaces such as movies, books, and popular music are characterized by high variance (a few huge successes, many flops) and unpredictability (the specific items that will be successful cannot be easily predicted). Salganik, Dodds, and Watts (2006) argued that this pattern is a result of positive feedback from social influence: As individuals influence each other, success and failure are both amplified, producing the observed pattern of unpredictable but high variation. To test their hypothesis, they set up an experimental marketplace, making songs available for downloading from a Web site. Some participants (visitors to the Web site) were in an independent, no-influence condition in which they could listen to songs and (if they chose to) download them—but without any indication of what others had done. Other participants were in a social influence condition in which they could also see how many previous visitors had downloaded each song, which might affect the participant's own decision about downloading. Results showed much greater variance in the social influence condition, with popular songs being more popular and unpopular ones less popular. In addition, success was less predictable in the social influence condition. This study illustrates the kinds of real-world effects of nonlinear feedback that can result from social influence over time.

One might assume that findings about influence derived from single-source, single-target, one-shot models and paradigms can be combined in some simple way to fully account for the complexities of dynamic social influence. However, there is no reason to believe that the effects of social influence repeated over time are necessarily linear and additive—that each time Amy wears Banana Republic clothing Brian becomes incrementally more positive about that brand. Brian's attitude may become exponentially more positive over time, or it may not change at all until the amount of exposure exceeds a threshold and then change dramatically, or exposure could have diminishing returns as he becomes inured to the positive impact of the excellent cut and high-quality fabrics. Models that generalize the results of one-shot laboratory studies in the most straightforward and direct way by assuming that each individual event has equal weight are unable to capture fascinating nonlinear social influence phenomena including radical conversion events and religious schisms.

There is some social psychological work that considers the effect of changes over time on influence. For instance, Prislin, Limbert, and Bauer (2000) had confederates change their position on an issue over the course of the study to make the participants' attitude (which stayed constant) become the majority position after initially being the minority position, or vice versa. They found an asymmetrical effect such that the decreased evaluation of the group when the participant's position changed from majority to minority was greater than the increase when the participant's position became the majority opinion. Gordijn, De Vries, and De Dreu (2002) used a more traditional laboratory paradigm to consider the effect of changing minority size and found that minorities that were increasing in size had a greater influence on targets' attitudes. These interesting and applicable results would be unknown without considering changes over time.

Considering influence among many sources and targets over time gives rise to fascinating questions about emergent effects, many of which can be studied using the well-controlled laboratory paradigms with which social psychologists are familiar. Does Amy's ability to persuade Brian depend on Amy's simultaneous responsiveness to Brian's opinion on another issue in the course of the same conversation or on Brian's successful or unsuccessful persuasion attempt aimed at Amy at a previous time? The norm of reciprocity might suggest an affirmative answer to such questions. How does a message from Amy that affects Brian's attitude end up affecting Candace when Brian later talks to Candace—could it have an equal or even greater impact on her than it did on Brian? When Amy as an influence source addresses Brian, Candace, and David as a group of targets, does she tune her influence to the most or least receptive target or to the average? Are people ever

made more confident of an opinion in an illusory way when they hear their own opinion coming back to them? Amy convinces Brian; Brian convinces Candace; Candace now agrees with Amy, reinforcing Amy's original view but not on the basis of the independent concurrence of multiple opinions. Questions such as these can only be framed and answered in the context of models that allow for the existence of multiple sources and targets that influence each other over time.

FOUR DIMENSIONS OF SOCIAL INFLUENCE MODELS

Potential Collapse of Opinions to Uniformity

Including multiple people and time in models of social influence is a crucial step in contextualizing social influence and identifying its potentially counterintuitive, emergent consequences for larger groups and populations. However, including multiple people and time in models of social influence leads to a problem. Abelson (1964) analyzed the class of models that assume that influence has a linear effect, causing a change in the target's attitude position toward the source by some constant fraction of the distance between them. That is, if Brian talks to Amy, her attitude will shift by some percentage in the direction of Brian's attitude. Although social psychologists are aware that influence does not inevitably result in assimilation or movement toward the advocated position (e.g., Brehm, 1966), as we will discuss later, most theories suppose that social influence typically brings the target's thoughts, feelings, or behavior more in line with those of the source.

The problem with this seemingly reasonable assumption, identified by Abelson (1964), emerges when the postulated influence processes are extended across multiple people and over time. Using linear differential equations, Abelson showed that if the network of influence is compact, meaning that there is at least one person with a direct or indirect path of influence to all other individuals, the only possible outcome is a collapse to complete uniformity of opinions. In other words, if social influence between two people makes the two become more similar to each other by a constant (even if small) amount, the process of many people influencing each other many times must result in everyone sharing the exact same psychological state. Abelson further demonstrated that in any compact model of assimilative social influence, the group must converge on an attitude within the range of the participants' original attitudes.

Both of these predicted aggregate-level outcomes are inconsistent with what is known about attitudes and social influence in groups. Although convergence to a unanimous opinion does occasionally occur, it is certainly not inevitable. Minority opinions often endure over time (see Nowak, Szamrej, & Latané, 1990), so there is meaningful variance in our individual psychological states that translates into meaningful variance in our attitudes and behaviors. Furthermore, the well-documented group polarization phenomenon demonstrates that when convergence does occur the group can converge on a position that is more extreme than any of the participants' original positions, and it often converges on a position more extreme than the initial group mean (Moscovici & Zavalloni, 1969).

Abelson (1964) realized that despite his clear theoretical prediction there are ways that diversity of opinion can be maintained. Armed with a dynamic perspective, we can begin to shed light on some of the processes that can contribute to the maintenance of diversity. The next four sections describe fundamental dimensions on which social influence models differ:

- the pattern of connectivity and influence assumed among individuals,
- the treatment of attitudes or behavioral responses as continuous versus discrete,
- the presence or absence of individual differences in private information, and
- the assumptions made about the assimilative versus contrastive directional effects of social influence.

For each dimension, we describe example models from within and outside social psychology that illustrate the various possibilities. We also discuss how each dimension may contribute to solving Abelson's problem, offering a mechanism for maintaining diversity of internal states in the population. Beyond their role in structuring the presentation of our review, the larger significance of these dimensions is that they describe the overall "space" of possible models of social influence. Each existing or potential model has a value on each dimension that locates the model at a particular position in the abstract space, thereby illustrating structural relationships of similarity and contrast among all such models.

Patterns of Connectivity

Social psychologists have typically left the understanding of how an attitude or behavior spreads in a population to sociologists or other researchers interested in higher, more aggregated levels of analysis, but the patterned flow of information in groups and populations can profoundly influence individual-level cognition. If Amy refuses to listen to Candace, single-source, single-target models of persuasion would predict that Amy's opinion would not be affected by Candace's. However, if Amy is best friends with Brian who is influenced by

Candace, Candace's opinion is quite likely to influence Amy's through their mutual friend, Brian. Similar dynamics apply when sexually transmitted diseases spread over multiple indirect links through social networks (Klovdahl, 1985; Klovdahl et al., 1994). The overall patterns of connection in a group of individuals can substantially affect the final distribution of attitudes or disease states or the outcome of a group's collective response, as shown by several studies (e.g., Gould, 1993; Macy & Skvoretz, 1998; Rojas & Howe, 2006; Rolfe, 2005). Thus, the large-scale implications of a particular set of assumptions about microlevel social influence processes depend crucially on the social structural pattern in which the processes are embedded.

Patterns of influence have been conceptualized in a variety of ways but perhaps the most useful conceptualization is the social networks framework (Wasserman & Faust, 1994). In a social network graph, individuals are seen as nodes and the channels of influence between people are seen as links or edges. This formalization allows researchers to use tools from graph theory to analyze patterns of human interconnection and to make quantitative comparisons between networks, to study properties of networks and of information passing through them and to make predictions about different kinds of networks.

From the social network perspective, models of social influence make assumptions that fall into four categories: all-connect networks, in which all individuals can directly influence each other; grid or lattice networks, in which each individual is influenced by a few immediate neighbors; heterogeneous networks, such as small-world networks; and dynamic networks, in which links between individuals can change. We discuss each in turn.

All-Connect Networks

Models with all-connect networks either explicitly assume that all individuals communicate with and influence each other or ignore the social network structure, tacitly assuming that all possible connections exist.

Social influence in small groups. Within social psychology, the study of group process and group decision making is concerned with how face-to-face groups of people communicate interactively and influence each other. Models of group processes, including Stasser's (1988) formal models, typically assume that social influence networks are all-connect, as does less formal work on group problem solving, opinion polarization, and jury decision making (Levine & Moreland, 1998). For example, Stasser (1988; Stasser & Titus, 1985) has empirically studied and modeled how individuals in groups share judgment-relevant information and reach group decisions. His model assumes that all individuals in a small group have the ability to communicate with each other and that each individual begins with both private information and information that is shared by others in the group. Much of this research focuses on determining the conditions under which individuals do and do not make their private information public and on the impact of those informational items on individual opinions and on the ultimate group decision. Stasser's formal and informal models take into account both the fact that every group member can be a source as well as a target of influence and that influence occurs over time throughout the group discussion.

Grid Networks

Although the assumption that all individuals can communicate with each other is useful for the study of specific topics such as jury decisions, all-connect networks are most appropriate for modeling small groups of individuals completing tasks in limited time frames. To study broader phenomena such as voting behavior, attitude formation, or conformity to social norms, it is often useful to think of individuals as being connected to only a limited number of others within a larger population. One such network is termed a *grid* or *lattice* network, in which each individual is influenced only by the states of a small number of neighbors.

Dynamic social impact model. Within social psychology, Nowak, Szamrej, and Latané's (1990) dynamic social impact model is one of the few formal treatments of social influence that occurs in populations over time. These researchers formalized Social Impact Theory (Latané, 1981), which presumes that the amount of influence depends on the distance, number, and strength (i.e., persuasiveness) of influence sources. Nowak et al. built on a class of models called *cellular automata* (Wolfram, 2002). In a cellular automaton, individuals are embedded in a grid with a fixed number of "neighbors" (typically eight). Each individual has a specific state, which in the context of opinion dynamics is taken to represent either the pro or the anti position. At each time tick, every individual chooses to stay the same or change its state (e.g., switch from pro to anti) by observing the states of its neighbors and applying a decision rule (e.g., take the state of the majority of neighbors). Nowak et al.'s model departs slightly from the typical cellular automata assumptions by postulating that each individual is influenced by more than just the eight immediate neighbors in the grid but that influence drops off as the square of the distance, so each individual is more influenced by close neighbors than by individuals further away in the grid.[1] With their assumptions, the researchers used simulations to demonstrate that individuals' opinions beginning in a random distribution

tend to form clusters over time, with pockets of minority opinions persisting despite the greater numbers of the majority. This model is an excellent demonstration of the value of contextualizing social influence: The initial pattern of opinions in an entire population can profoundly shape the degree to which each individual is influenced as well as the distribution of final opinions of the individuals.

Local patterns of connection as found in Nowak et al.'s (1990) social influence model and the similar Sznajd models (Stauffer, 2001) have repeatedly been observed to lead to clustering of opinions, an outcome that is less likely to emerge in all-connect networks. In an all-connect network where everyone is influenced by everyone else, everyone is subject to the same influence, pushing the group toward uniformity. For example, once a majority opinion emerges, every group member is subject to influence from that majority. In contrast, in a grid network each individual receives influence from a different subset of individuals (i.e., his or her particular neighbors). Thus, most individuals (all but those on a narrow dividing line separating regions of different opinions) may be surrounded totally or mostly by neighbors who hold the same opinion as themselves, leading to stably persisting, clustered patterns of opinion.

Heterogeneous Social Networks

All-connect and grid networks assume that everyone has the same pattern of connections as do everyone else (e.g., connections to everyone or to exactly eight neighbors). In contrast, heterogeneous networks do away with this assumption so that one individual may have many links to others while another only has one link or none at all. One individual may have all his or her links to other members of a small clique while another serves as a bridge between two distinct social groups. These features can make social network models more realistic and affect the way influence and information can flow within the group.

For instance, Bavelas (1950) and Leavitt (1951) found that the structure of communication networks (e.g., star, ring, all-connect) affected the performance of groups in problem-solving situations. Experiments in problem solving demonstrate that groups with limited contacts are better than more highly connected groups are at solving complex problems because they tend not to converge too quickly on the wrong answer (cf. Hutchins, 1991). However, on simple problems the quick flow of information between individuals leads all-connect networks to perform better (Mason, Jones, & Goldstone, 2006). Recent research in mathematical sociology and information science has demonstrated reliable differences in information transmission between different types of networks (e.g., Kleinberg, 2000). Research (Barabási & Albert, 1999; Watts & Strogatz, 1998) has also revealed that communication patterns between individuals form distinct and identifiable kinds of networks that have different implications for information transmission. For example, small-world networks allow the quick propagation of information across an entire connected network, whereas grid or other localized connection patterns make long-distance information transfer much slower.

Many concepts from the graph theory used to conceptualize social networks directly relate to how information or influence flows through social groups. The *diameter* of a graph is the longest distance between any two people, where distance is measured as the smallest number of links on the path connecting two people. The larger the diameter of a social network, the longer it will take for influence to reach everyone. In an all-connect network, the diameter is exactly one by definition, so information is disseminated very quickly. However, in many social networks such as office hierarchies, individuals are more sparsely connected. If the diameter is high, news about a fire in the mailroom or fake record keeping in accounting can take quite a while to reach the chief executive officer. The average distance between all pairs of nodes also relates to the spread of influence because the shorter the average distance, the faster information can travel (on average) from one person to the next. In real social networks, the popularized notion of "six degrees of separation" (Milgram, 1967) refers to the small distance between any two people in the world, more formally referred to as the *small-world* property.

Some networks, such as the links between Web pages, have a *scale-free* property (Barabási & Albert, 1999). This refers to the fact that there are many Web pages with just a few links to them but only a few Web pages are targeted by many links. Those few have the potential to be disproportionately influential over the entire network. To some extent, social ties follow this pattern as well. For instance, collaboration networks (defined by coauthorship) and citation networks also tend to have a scale-free distribution (Newman, 2001; Redner, 1998). Very few researchers publish with many people and are widely cited but many researchers are relatively obscure. Another feature of collaboration networks and other kinds of social networks is a high degree of *clustering*—meaning that two people who both have a link with a third are likely to be linked, themselves. Higher clustering in a network indicates high "cliqueishness," so attitudes also tend to cluster in these networks (cf. Latané & Bourgeois, 1996).

Because most naturally occurring social networks are heterogeneous, the social network framework and the tools for conceptualizing and analyzing such networks can make potentially useful contributions in the study of social influence. In particular, a rich vein of exciting research questions involve the relations of network

properties (such as clustering, small-world properties, and average distance) to dynamic aspects of the flow of social influence through the network. Despite the importance of such questions, few influence models actually assume that individuals are connected in heterogeneous networks.

Opinion dynamics models. Both physicists and sociologists have contributed to a class of social influence models termed *opinion dynamics models*, the most influential of which is Friedkin's (1998) structural theory of social influence. Slightly simplified, this formal model represents the opinions of a group of N people at a particular point in time as an $N \times 1$ vector and the strength of each individual's influence on each other individual as the entries in an $N \times N$ matrix. The entry in the ith row and jth column represents the strength of influence person i has on person j, which is assumed to be constant over time. This representation is quite flexible and allows for a variety of patterns of network patterns. Some individuals can be totally isolated (when all their influence parameters are zero) or the group can be a completely connected network in which each individual can directly influence everyone else. By iteratively multiplying the initial opinion vector by the influence matrix, it is possible to predict the trajectory of opinions as well as the final opinion for every member of the population. The focus of studies using the structural theory of influence is on the impact of the pattern of connections between individuals on the way opinions change and ultimately reach equilibrium in a group.

Dynamic Social Networks

The thoughtful reader at this point may have realized that actual social networks are not static but change over time. There are surprisingly few models that attempt to incorporate dynamic networks. To do so, a model must assume that individuals can strengthen or weaken (or create or discard) network links to other individuals, depending on certain criteria such as the extent of agreement or disagreement with those others. Thus, the modeled individuals will be able to shape their local social networks (the set of others they are connected to) as well as to change their opinion in response to influence from others who are connected to them.

Besides the obvious fact that patterns of connectivity in real-life social networks do change over time, there is another important reason for modeling dynamic rather than static networks. Start with the observation that "homophily" is pervasive in the social world: People tend to be similar to their friends or, more formally, behaviors and attitudes are generally correlated between individuals connected by network links (McPherson, Smith-Lovin, & Cook, 2001). What cognitive and social processes create homophily? It could be the result of influence flowing over pre-existing links. (Brian influences the political attitude of his friend Amy.) But it could also occur through the preferential linkage of people whose pre-existing attitudes agree. (Brian strikes up a friendship with Amy because they meet at a Democratic Club event and learn that they agree politically.) Finally, homophily could come about because people who are interlinked share similar structural positions in the overall network (e.g., they are all students or all working people), and the structural positions in turn give them certain interests and perspectives that lead to similar attitudes. This idea has been advocated by Burt (1978), although the effect of structural equivalence on homophily has also been questioned (Friedkin, 1994). To disambiguate causality in this situation, a model must take all these potential effects into account rather than assuming that similarity reflects solely social influence processes (Doreian, 2001). In fact, research (McPherson et al., 2001) suggests that social selection processes (the creation of network links in response to pre-existing similarity) are more important than social influence processes are (the flow of influence across pre-existing links). In such instances, a model that failed to take dynamically changing network links into account would overestimate the effects of social influence.

Axelrod's adaptive culture model. Axelrod's (1997) adaptive culture model is one important dynamic network model of social influence. In this model, each individual is represented by a fixed-length string of discrete symbols (e.g., 8 letters) that represent cultural traits (think of specific group memberships, attitudes, etc.). Individuals sit at fixed locations in a grid and are potentially able to interact with their neighbors, but the probability of interaction depends on the similarity of the two individuals. For example, if two neighbors have strings that match in 2 of the 8 positions, they would interact with probability 2/8. Interaction creates more similarity by replacing a randomly chosen element in one interactant's string with the corresponding element from the neighbor. Thus, similarity increases the likelihood of interaction and interaction itself creates more similarity. Iterating the model over time tends to result in the evolution of several distinct "cultures," regions of individuals sharing identical strings. The cultures do not affect each other because at their boundaries neighbors have completely nonoverlapping strings and therefore never interact.

Hegselmann and Krause's dynamic network model. Hegselmann and Krause (2002) presented the only other influence model involving a dynamic network of which we are aware. Their model is a generalization of

TABLE 1: Examples of Models That Assume Different Patterns of Network Connectivity

	Social Psychology	*Other Fields*
All-connect networks	Group decision models (Stasser, 1988)	Innovation diffusion (Granovetter, 1978)
Local grid networks	Dynamic social impact (Nowak, Szamrej, & Latané, 1990)	Information cascade (Bikhchandani, Hirshleifer, & Welch, 1992) Cellular automata Swarm intelligence (Kennedy & Eberhart, 2001; Moldovan & Goldenberg, 2004)
Heterogeneous networks		Structural theory of social influence (Friedkin, 1998) Innovation diffusion (Valente, 1995)
Dynamic networks		Culture model (Axelrod, 1997; Hegselmann & Krause, 2002)

Friedkin's (1998) model, relaxing Friedkin's assumption that the matrix of weights representing the influence of each individual on each other is constant across time. Instead, Hegselmann and Krause assumed that individuals accept influence only from others whose current opinions are within a certain threshold distance from their own current positions. For example, if Amy's position is 3 on a 1-to-9 attitude scale and her threshold for allowing influence from others is 2, she will be influenced by Brian if his position is between 1 and 5 on the scale (i.e., 3 ± 2) but not if his position is more extreme. Because each individual's opinion may change over time, so will the set of other individuals who can exercise influence, making this an instance of a dynamic network model. The model predicts that groups can end up converging to a common opinion or splitting into several subgroups holding differing opinions, depending on the initial attitude distribution and the threshold for influence from others.

Table 1 presents representative models from both social psychology and other fields, illustrating each type of assumed influence network.

Beyond a Specific Social Network

An individual does not have a unique social network. The structure as well as the size of a network differs dramatically depending on whether one is considering a network of hundreds of acquaintances, dozens of daily contacts, or a few truly close friends and kin. A model of social influence must consider the network that is relevant for the specific attitude or behavior of interest. For example, the set of others who might influence a person's opinion about a work problem will generally differ from the set of those who might influence decisions about how to invest for retirement, what candidate to vote for, or what treatment to seek for a medical condition. Investigation of the different networks relevant for different types of issues is in its infancy; Krackhardt and Porter's (1985) is the only study of which we are aware.

A related issue is that people may be able to obtain some information from others beyond their specifically connected group of friends or acquaintances—even aggregated information from an entire population. Amy knows the political attitudes of her friends but news reports may indicate the proportion of the vote from the entire city, state, or nation that went to her candidate. Brian knows that he and his friends smoke, but he may also realize that smoking is generally unpopular (through observations of others' behavior in crowds or reading about the enactment of antismoking ordinances). Candace may realize that her friends all drive gas-guzzling SUVs, but she can also observe that they are selling only slowly and with huge rebates, suggesting that they are becoming less popular in the overall auto marketplace. One interesting theory that seeks to incorporate social influence from this aggregate level as well as from a network of acquaintances is that of Blanton and Christie (2003). Their model predicts, for example, that people like Brian the smoker who agree with a majority of their social network but disagree with the majority in the larger population may seek a positively valued image as deviants, nonconformers, or free thinkers. Clearly, much remains to be learned about the ways influence from a local network and from larger populations combine to affect people's attitudes and behaviors, and this represents an important direction for future research.

The Role of Network Structure in Preserving Diversity of Opinion

The assumed social network structure, as one dimension on which models of social influence can vary, is a key determinant of a model's predictions regarding attitudinal diversity. The most straightforward means of maintaining diversity of opinion is to have some individuals in the network completely isolated from influence by others. Abelson (1964) pointed out that the inevitability of unanimity was only true in compact networks—if some part of the network is totally unconnected to

another part of the network, some diversity of opinion can be maintained. For example, in Friedkin's (1998) theory it is possible to isolate individuals or subgroups by setting the parameters of their potential influencers to zero. An isolated individual is simply not influenced by any of the other individuals in the model, and an isolated group is a group of individuals who are influenced by each other but not by any others outside the subgroup. This pattern may often be approximated in reality. For example, social identity theory suggests that this process often preserves differences of attitudes between groups (Tajfel, 1978). People tend to be open to influence by members of their ingroups but are little influenced by outgroup members (Mackie, 1986). Similar group-based patterns of influence arise in Axelrod's (1997) culture model in which individuals preferentially interact with likeminded others, actively preserving diversity. In the real world, there is evidence that charismatic leaders employ isolation from outsiders to preserve members' loyalty to cults (Lifton, 1991), and similar tactics are attributed to abusers in domestic situations (Wolfson, 2003).

Leaving aside the possibility of totally isolated individuals or subgroups, less extreme patterns of network connectivity can also assist the preservation of opinion diversity over time. A network can have what is termed *community structure* (Jin, Girvan, & Newman, 2001), in which distinct cliques of individuals are highly interconnected but have only a few connections to other cliques. Each individual in such a network can be in a subgroup where the great majority of his neighbors (members of the same clique) agree with him, providing opinion support and maintaining diversity even though the various cliques do have a few links between them. In this type of model, linear influence rules will still lead to opinion uniformity. But with a nonlinear influence rule (for example, a rule that individuals do not change their attitudes if a majority of their neighbors agree with them), the sparse connections between cliques can maintain attitude diversity because each individual has plentiful local support from his or her own clique.

Modeling Continuous Versus Discrete Responses

Turning from the network structure, the second dimension on which social influence models vary is their focus on continuous variables (often attitudes) or discrete choices (often behaviors) as the source and target of influence.

Continuous Attitudes

In part because goals of influence models in different fields differ, the models differ widely in their assumptions about the underlying nature of the entity that is subject to influence. Some models assume that social influence affects attitudes, which can vary along a continuous scale incorporating intensity as well as direction (pro or anti). An individual might be assumed to maintain a mental representation of an attitude on a 1-to-9 scale, for instance, where 5 is a neutral midpoint, 1 means strongly against, 4 means weakly against, and so forth. Most models of social influence in social psychology are of this type, assuming a continuous representation of an attitude. Representing attitudes as continuous has important implications for a model. It allows researchers to detect incremental changes in attitude. For example, if Amy is weakly anti–death penalty and Brian is strongly pro, a communication from Amy to Brian may shift Brian's attitude from 9 to 8 on a 9-point scale, with potentially meaningful consequences down the line even though Brian has not converted to Amy's side of the continuum. Considering an attitude as a continuous dimension also allows researchers to measure the discrepancy between an influencer and the target of influence, which has been shown to affect the degree of influence (Hegselmann & Krause, 2002; Sherif & Hovland, 1965).

Representing attitudes as continuous also makes theoretical sense. If everyone stands either at a definite pro or anti position on any given attitude, one would expect to see highly polarized distributions of opinion on important topics in the population, yet there is little or no evidence that population attitude distributions are nonnormal (Dimaggio, Evans, & Bryson, 1996). For all these reasons, social psychological models have generally assumed that attitudes can be represented as a continuous dimension. Just a few examples include the Theory of Reasoned Action (Fishbein & Ajzen, 1975), Fazio's (1986) model of an attitude as a continuously variable strength of association between an object and an evaluation, and Friedkin's (1998) sociological opinion dynamics model.

Discrete Responses

Other models, especially those developed outside psychology, depart from the assumption of an underlying continuous dimension and concern themselves only with predicting binary choices: yes or no, pro or anti, riot or don't riot, Pepsi or Coke, Republican or Democrat (see Table 2). These models are obviously applicable to situations involving naturally discrete choices such as voting (for or against a candidate) or consumer choice (purchasing or refusing to purchase a product). These models also capture the important idea that in many cases the only way that each individual's attitude is visible to (and able to influence) other individuals is through the individual's discrete behavioral choices. Amy can observe Brian buying a Mac rather than a Windows computer, for instance, but she probably cannot tell whether Brian's attitude toward a Mac is a 4 while his attitude toward Windows is a 3. Of course,

TABLE 2: Examples of Models That Assume Continuous Underlying Attitudes Versus Those That Model Discrete Behavioral Choices

	Social Psychology	*Other Fields*
Continuous attitude	Theory of Reasoned Action	Structural theory of social influence (Friedkin, 1998) Swarm intelligence (Kennedy & Eberhart, 1995)
Discrete choice	Dynamic social impact (Nowak, Szamrej, & Latané, 1990) Social decision scheme (Stasser, 1988)	Cellular automata Innovation diffusion (Granovetter, 1978; Valente, 1995; Walker & Wooldridge, 1995)

some behaviors naturally fall along a continuum and can be observed by others in quantitative detail and not just as a discrete choice. Brian may find out not just that Amy supports Habitat for Humanity but how many hours per month Amy spends building houses. Our point is not that all behaviors are discrete (in contrast to continuous attitudes) but simply that many forms of behavior that are importantly relevant to social influence are inherently discrete: buying a product, marrying a partner, attending a party, joining a church, voting for a candidate, and so forth.

Innovation diffusion models. An important class of models of discrete choices from economics and sociology addresses influence processes that underlie the spread of ideas or behavioral innovations in populations (Valente, 1995). Granovetter's (1978) model was an early formal treatment of innovation diffusion. In his model, individuals have varying thresholds for adopting a particular behavior such as joining a riot or purchasing a new consumer product. Each individual's threshold is based on a personal calculation of costs and benefits of performing the behavior, which will depend on how many others are doing the same thing. For example, different individuals may have different degrees of ideological commitment (which predispose them to riot) and different degrees of caution and risk aversion (which discourage riot participation). Seeing many others riot may both increase the perceived benefits and decrease the perceived risks of joining the riot. For this reason each individual can be described as having a threshold: The minimum number of other people performing the behavior that make this individual feel that the behavior is more rewarding than costly. An instigator is one whose threshold is 0%: He or she will riot even if nobody else does so. A second person with a low (1%) threshold may then join the first. Once a few people begin rioting, others with relatively low thresholds may join in, and finally even "respectable" citizens whose radicalism is low and cautiousness is high may join in if their thresholds of 80 or 90% are exceeded.

The structure of this model emphasizes the importance of the distribution of individual states or differences in the population. If all of the individuals have a threshold of 50%, no one will riot; whereas if a population of 100 individuals includes one individual with each possible threshold from 0% through 99%, then everyone will ultimately riot. Although the mean threshold in these two cases is exactly the same, the behaviors of the populations are very different. In models of this type, social influence (how many others are performing the behavior) is obviously the key determinant of any individual's choice but there is also a focus on individual differences: the assumption that individuals have distinct predispositions summed up in their participation thresholds. This model implicitly assumes that everyone can see everyone else's rioting behavior, making it an all-connect network. Valente (1995) extended the model to heterogeneous networks and found that an individual's network position can be as important as his or her threshold in determining whether he or she will adopt the innovation.

Besides innovation diffusion models, other models assuming discrete states include cellular automata (Wolfram, 2002) and the similar social-psychological model of dynamic social impact (Nowak et al., 1990). Within social psychology, models of social influence assuming discrete positions include social decision schemes (Stasser, 1999). These are mappings of initial distributions of discrete preferences in a group to the group's final decision. For example, a jury starts out with 8 favoring conviction and 4 acquittal, or a hiring committee has 5 favoring candidate X and 2 favoring candidate Y. In each case the social decision schemes model will permit estimation of the probability of each possible final group decision.

Multiple Representations

An obvious possibility that is not currently well represented in models of social influence is to model both a continuous underlying attitude dimension and a discrete behavioral choice. For instance, Urbig (2003) modeled the

dynamics of persuasion when an underlying continuous attitude can only be communicated as a discrete point on a scale. In the model, limiting individuals' abilities to express their continuous attitudes means that more minority opinions can persist in a group. The argument for considering attitude and behavior separately is straightforward. The attitude functions as the individual's internal, continually updated summary of the relevant inputs encountered over time (e.g., positive and negative items of information about a particular consumer product, weighted by their importance). The attitude is probably not directly observable by other individuals and hence cannot itself be a source of social influence. The individual's discrete behavioral choice can be modeled to depend on the attitude in specified ways, for example by adopting the assumptions of the Theory of Reasoned Action or other prominent social-psychological models of attitude-behavior relations. The overt behavior can be assumed to be observable to other individuals and therefore potentially to influence them. Models that distinguish between and model both attitude and behavior have the advantage of being able to simultaneously represent public behavioral conformity and private attitudinal diversity. Models lacking both of these representations cannot account for pluralistic ignorance (Miller & McFarland, 1987; Monin & Norton, 2003). This is a state in which an entire population behaves inconsistently with an attitude that they all share, for example by refraining from card playing although they privately approve of it, because they assume that everyone else disapproves.

Implications for Maintaining Variability

When combined with the assumption that social influence is linear and assimilative, continuous representations of attitudes lend themselves to the complete collapse discussed by Abelson (1964). Over time, all members of a mutually influencing population will inevitably end up with exactly the same attitude. Models assuming discrete states, on the other hand, will not share this fate because they cannot use linear influence rules. For example, in a cellular automaton model, each individual determines its discrete attitude by observing the behavior of its neighbors and applying a simple decision rule. An individual might use the rule that it should change its position to that of the majority of its eight neighbors. Because of the rule's nonlinear (threshold) nature, it is possible to maintain diversity even with purely assimilative rules such as this one. For instance, imagine a grid of individuals like a checkerboard where pro and anti individuals alternate. Each individual in this grid has four neighbors who are pro and four who are anti; he will never change his opinion because a majority of his neighbors never forms. While the average attitude of each individual's neighbors (and of the whole grid) is equal to 0.5, attitudinal diversity is maintained because the potential attitude states are limited to 0 and 1 and a local majority is required as the threshold for change. Thus, models that represent discrete behavioral choices (e.g., Nowak et al., 1990) may support attitudinal diversity (avoiding collapse to unanimity) simply as a function of the underlying nature of the representation. Whether the actual underlying mental representations are discrete or continuous is a question that social cognition approaches (which have long been concerned with the nature of attitude representation) seem well equipped to answer.

Variable Environmental Influences

Some social influence models assume that opinions depend on nothing but social influence. Amy may not care about a particular issue, so her only concern may be to make her opinion line up with those of Brian, Candace, and David. For example, Nowak et al. (1990) and other cellular automata models generally assume that each individual maintains its previous state (opinion position) or changes state solely as a function of the states of its neighbors.

Other models, in contrast, assume that social influence is just one among the forces that affect each individual's attitudes or decisions (see Table 3). Amy may want to take account of the views of her friends on tax increases but also to keep her opinion on the issue in line with her underlying political ideology, religious worldview, and financial self-interest.

Group problem solving models. Models that are essentially about social influence often go by some other name, such as *group problem solving* or even *swarm intelligence*. Group problem solving involves social influence in that information about one person's attitude or behavior has an impact on another person's attitude or behavior. This can occur through the simple provision of information about the merits of different potential solutions rather than argumentation or social pressure. Kennedy and Eberhart (2001) developed the swarm intelligence model, in which multiple individuals search for good solutions to a problem (e.g., a good recipe or combination of input materials and processing conditions to manufacture a product). Each possible solution is conceptualized as a location in a multidimensional space. Each individual tries various solutions, learning how good a result is obtained at each location and keeping track of the location where it has found its best solution to date. Each individual also gets

TABLE 3: Examples of Models That Assume Social Influence Is the Only Effect on Individual Attitudes Versus Those That Assume Personal Information or Attitudes Are Also Important (as Well as Social Influence)

	Social Psychology	*Other Fields*
Social influence only	Dynamic social impact (Nowak, Szamrej, & Latané, 1990)	Cellular automata
		Innovation diffusion (Granovetter, 1978)
		Structural theory of social influence (Friedkin, 1998)
		Culture model (Axelrod, 1997; Hegselmann & Krause, 2002)
Other sources of input	Theory of Reasoned Action (Fishbein & Ajzen, 1975)	Swarm intelligence (Kennedy & Eberhart, 2001)
	Group decision models (Stasser et al., 1989)	Information cascade models (Bikhchandani, Hirshleifer, & Welch, 1992)
		Sznajd models (Stauffer, 2001)

information from a small number of linked neighbor individuals about the locations they explore and the results they obtain. The individual tracks the best location sampled by these neighbors and on each move chooses for further exploration locations near its personal best-ever spot and its neighbors' best-ever spots. Thus, each individual obtains information from its physical environment by trying solutions and from its social environment via its neighbors' explorations, and both types of information affect the individual's choices of locations to explore. If one individual finds a particularly good solution, that information will flow to others in the population through the neighbor links and they will all converge on that good solution.

Grid models with external influence. The Sznajd models studied by Stauffer (2001) resemble the Nowak et al. (1990) model in assuming dichotomous attitudes held by agents who sit in a rectangular grid and are influenced by their neighbors. Unlike Nowak et al.'s, these models often incorporate an effect in addition to social influence, a single parameter that biases all agents equally in the direction of one alternative or the other (think of it as an ad campaign for one product). This model of external influence is highly simplified, though, in the assumption of an equal effect on all agents.

Information cascade model. Like the swarm intelligence model, the information cascade model of Bikhchandani, Hirshleifer, and Welch (1992) explicitly assumes that each individual obtains his or her own private information from the environment. In this model, each individual in turn obtains an item of private information relevant to a decision and then publicly announces his or her decision. Subsequent individuals are affected by the information conveyed by the previously announced decisions as well as their private information. For concreteness, suppose Amy has a biased coin that falls one way two thirds of the time and the other way only one third of the time. Amy challenges Brian, Candace, and David to determine whether the bias is toward heads or tails. Brian privately flips the coin, obtaining heads. Based on that evidence, the probability that the bias is heads is .67, so Brian announces he thinks it's heads. Next Candace privately flips the coin; say she gets tails. Now Candace can infer that one heads and one tails have been thrown. The odds that the bias is heads are now .50. Candace makes her guess and announces it before Amy moves to David, the next person in line. David will have Brian's guess, Candace's guess, and his own private coin flip on which to base his judgment—and so on.

An interesting aspect of the information cascade model is that if the first two flips agree (say the first two players both announce heads as their guesses), then David and all subsequent players should never rationally go against them, even if they obtain contrary private evidence (by flipping tails). The first two flips should outweigh the third person's private flip, and so the third person should guess the same as the first two, creating an even stronger situation for the fourth person and subsequent people to do the same. The outcome of such a cascade is that person after person ends up—perfectly rationally—agreeing with an emerging consensus even when his or her private opinion, gleaned from interaction with the environment, might indicate a different decision. (This situation has conceptual parallels with pluralistic ignorance, discussed earlier.) This model emphasizes the distinction between private information, which is gathered from the environment or from preexisting knowledge, and public information. It is easy to imagine that studying the ways these types of information are gathered and the ways they influence attitudes and behavior could tell social psychologists a great deal about how and when people are influenced.

Private information in group discussions. Within social psychology, Stasser and his colleagues (Stasser,

1988; Stasser & Titus, 1985) have investigated and modeled situations in which members of a discussion group come together bringing different sets of information (e.g., positive and negative attributes of political candidates or policy choices). The information can be arranged so that some individuals start with private information that leads them to have favorable attitudes while others start with information dictating unfavorable attitudes toward the object. These differences may be resolved as the group members discuss and share their information. Again, each group member's private information as well as social influence from others can potentially affect his or her decisions.

Differential Weights on Social and Personal Information

If individuals are assumed to have their own private information about or preferences for different attitudinal positions or behaviors, the possibility emerges that individuals will put different weights on social influence and private information. Models of innovation diffusion (Granovetter, 1978; Valente, 1995), described earlier, are of this sort. The relative responsiveness of each individual to private versus social influences is summed up in the participation threshold.

In social psychological research, several important variables have been found to influence the relative weights people give to their personal attitudes versus perceived social norms in making behavioral decisions. For example, behaviors performed in private are relatively more influenced by attitudes than by norms. Individuals who are high self-monitors (Snyder, 1974) pay more attention to social norms, low self-monitors to their attitudes. Formal models of social influence have rarely incorporated such moderator variables.

Implications for Maintaining Diversity

In models in which individuals obtain private information as well as information about others' choices, individuals' behaviors need not be uniform for several reasons. First and most obviously, the information that individuals sample from their environments is likely to differ from one person to another (as in the information cascade model). Second, the evaluation of different attitudes or behaviors may differ across individuals. For example, someone who is genetically a supertaster will generally dislike broccoli because its flavor is experienced as extremely bitter; such an individual is unlikely to be responsive to social influence from others who advocate eating lots of broccoli. Numerous theorists in social psychology have assumed that individuals have their own private attitudes because of differences in learning history, tastes and preferences, innate differences, and so forth even if they share the same decision-relevant information. Prominent examples are the Theory of Reasoned Action (Fishbein & Ajzen, 1975) and the Theory of Planned Behavior (Ajzen, 1991). Both of these models assume that people make behavioral decisions on the basis of their attitudes (personal likes or dislikes for an object) as well as on the basis of their subjective norms or perceptions of what important other people would like them to do—in other words, social influence. Models that maintain diversity of opinion by maintaining diversity of individual experiences command substantial empirical support; we know that different people have very different experiences of their worlds, shaped by their pre-existing beliefs and attitudes (Griffin & Ross, 1991). Models in which individual opinions are determined solely by social influence generally have a more difficult time maintaining diversity (Abelson, 1964), although the other factors we discuss here may make that possible.

Assimilative Versus Contrastive Social Influence

Most models of social influence both inside and outside social psychology assume purely assimilative influence. However, evidence shows that influence is not always assimilative: Under certain conditions people contrast against or move away from a position expressed or advocated by others. Amy's expression of her political opinion might be persuasive to Brian, or he might instead see her as parroting the uninformed and bigoted views of a talk-radio personality and move his own position even further away from hers.

To understand how social influence can lead to either assimilation or contrast, we must briefly review the reasons people respond to influence in the first place. Deutsch and Gerard (1955) initially drew a distinction between informational and normative social influence. Informational influence involves learning what other people think and treating it as valid information relevant to finding the correct answer to an issue. For example, if Amy observes that a majority of other consumers buy a particular product or endorse a particular candidate, that may be taken as indicating that the product or candidate is a good one, so Amy should prefer it as well. Normative influence, on the other hand, is based on others' power to socially reward agreement or punish deviance. For example, if most people in Amy's social circle endorse a particular candidate, she can probably expect social approval if she agrees or hostility or ridicule if she disagrees. Of course, several important variables such as the public versus private nature of one's opinion expressions moderate the importance of these two types of influence.

The informational–normative distinction is now recognized as oversimplified to some extent (see Hogg &

Turner, 1987; Turner et al., 1989). More recently, researchers (e.g., Cialdini & Trost, 1998) have insightfully reformulated the informational–normative distinction into a typology of the distinct goals that people might accomplish by conforming to others' opinions. One potential goal is to understand the world accurately (corresponding to informational influence); another is to maintain close and satisfying social relationships (corresponding to normative influence). These are not the only possible goals; a third is to maintain a positive self-image for oneself and for others (e.g., an image as a cooperative team player or as an independent-minded innovator).

Models Assuming Only Assimilative Influence

This conceptual analysis suggests that social influence in a context where accuracy is the chief goal (informational influence) will ordinarily be purely assimilative (e.g., Janssen & Jager, 2001). If other people tell you about a great new product or give you their honest opinions on a difficult problem you are trying to solve, there is little reason to disregard them and much reason to give those opinions some weight in your own decision. After all, the convergence of many people on a single opinion or judgment does frequently indicate its accuracy and adaptiveness.[2] Most theories of social influence outside of social psychology (e.g., information cascade, innovation diffusion, Friedkin's structural theory, and swarm intelligence models) have implicitly or explicitly assumed that influence is informational in nature, and so they have not even considered the possibility of contrast.

In such models, the flow of influence can equally well be conceptualized as the flow of information between people, with that information (e.g., about a tremendous new product) changing the behavior of each individual in the same direction. In fact, although we do not pursue the possibility here, models of the transmission of disease can be formally identical to information transmission models (Klovdahl, 1985). In either case, it is assumed that people are situated in social networks, and some are initial carriers of the infection or the information (rumor, etc.). Contact between an individual carrier and another person is assumed to lead, with some probability, to the second person becoming infected or knowledgeable and therefore becoming a carrier him or herself, able to further the spread through the entire linked network. The influence is by definition assimilative: Contact with a carrier makes the other person become similar to the carrier (i.e., become a carrier himself or herself).

Norm convergence model. Friedkin's (1998) model, Hegselmann and Krause's (2002) bounded confidence model, and Granovetter's (1978) innovation diffusion model, to name just a few, all assume that when influence occurs it is assimilative. Yet another example, not previously discussed, is Walker and Wooldridge's (1995) model of convergence on a social convention. The model assumes a situation in which several possible conventions (simple rules for behavior) could be applied in a population, all of which are equally effective so long as everyone is using the same one. An example would be the convention of driving on the right or left side of the road. In this model, each individual starts out with a random guess among the possible conventions, and then the individuals begin to interact. When two individuals interact, each learns what the other's current convention is, and each can switch to the other's convention depending on various possible decision rules (e.g., switch if the majority of other individuals encountered so far are using a particular convention). The foci of this model are on whether a single convention ultimately spreads through the entire population and on how long it takes to do so. Obviously this is a model of purely informational influence; individuals adopt the convention being used by another individual because to do so is adaptive (it is likely to put them in the emerging majority) rather than because other individuals might socially reject them if they do not.

Processes Creating Contrastive Influence

Social psychological theory and research suggest that when social rewards or positive self-identities are salient goals influence can be assimilative. The standard assumption regarding normative social influence is that groups will give social rewards to those who conform and perhaps punishments to those who deviate, and people can certainly be motivated to conform by the positive cooperative and team-spirited identities provided by conformity.

However, at times people seek to satisfy social reward and self-esteem goals by contrasting away from the influencer's opinion or behavior. Social rewards for nonconformity are unlikely to come from the immediate group itself; face-to-face groups generally do not treat deviants well (Cialdini & Trost, 1998). However, rewards may come from another social group that regards an individual's nonconformity as a praiseworthy form of standing up for principle in the face of social pressure to do otherwise. Nations honor their prisoners of war who hold out under enemy mistreatment, and religious movements remember martyrs who died for their faiths. People can also be motivated toward nonconformity if they value identities as independent-minded innovators rather than identities as conservative, solid team players. Social identity theory (Tajfel, 1978), self-categorization theory (Turner et al., 1987), and optimal distinctiveness theory (ODT)

TABLE 4: Examples of Models That Assume Social Influence Is Always Assimilative Versus Those That Also Allow Contrastive Influence

	Social Psychology	*Other Fields*
Assimilative influence only	Dynamic social impact (Nowak, Szamrej, & Latané, 1990) Group decision models (Stasser, 1988) Theory of Reasoned Action (Fishbein & Ajzen, 1975)	Structural theory of social influence (Friedkin, 1998) Swarm intelligence (Kennedy & Eberhart, 2001) Information cascade models (Bikhchandani, Hirshleifer, & Welch, 1992) Innovation diffusion (Granovetter, 1978)
Contrastive and assimilative influence	Social identity theory Optimal distinctiveness theory Psychological reactance (Blanton & Christie, 2003)	Seceder models (Dittrich, Liljeros, Soulier, & Banzhaf, 2000)

(Brewer, 1991) all offer motivational accounts of how people seek to maintain identities for their ingroups (and hence themselves) as distinctive from salient outgroups or others. Such motives can obviously lead to attitude contrast under specified conditions. ODT, for instance, suggests that if one perceives that his or her ingroup is too similar to others, that person may adjust his or her attitudes or behaviors away from those others to maintain the ingroup's distinctiveness.

A social psychological theory that includes contrastive influence of a slightly different sort is Brehm's (1966) theory of psychological reactance.[3] This theory suggests that when an action is forbidden people may be motivated to carry out the action anyway. So when a movie is "banned in Boston," people may try extra hard to see it even if they would not otherwise care about it. This process is interpreted as a reaction to threats to behavioral freedom, and it might not be too much of a stretch to view reactance as a way of reinforcing a positive self-identity as being in control of one's own actions. Similarly, researchers have observed a boomerang phenomenon in attitude change, noting that attitudes sometimes move away from persuasive messages that induce fear (Rogers & Prentice-Dunn, 1997), that originate in a minority (Wood, Ouellette, Busceme, & Blackstone, 1994), or that are seen as extremely discrepant from one's own position (Sherif & Hovland, 1965). Within our motivational framework, these effects might also be considered as examples of people's desires to avoid a negative identity (as being fearful or as being an adherent of a devalued minority position).

Models of Mixed Assimilation and Contrast

Finally, some models postulate interesting mixes of assimilative and contrastive social influence. ODT (Brewer, 1991) and self-categorization theory (Turner et al., 1987) are two examples, assuming that people are motivated to assimilate to their relatively small ingroup but also to maintain differentiation between that ingroup and salient outgroups. In self-categorization theory, this is termed the *metacontrast principle*, and it operates as group members construct a perceived prototype of their group. The prototype is not simply an average of group members' characteristics but is biased away from the attributes of relevant (often competing) outgroups. For example, if the outgroup is seen as having generally conservative attitudes, the ingroup prototype will be biased toward the liberal end of the spectrum rather than being a simple average of group members' attitudes. The prototype in turn takes on motivational force as group members seek to approach and conform to it—the result being that they tend to move away from the perceived position of the outgroup as well as toward the position of fellow ingroup members.

In somewhat similar fashion, Blanton and Christie (2003) differentiated between a local peer group or social network, the nexus of face-to-face interactions, and the broader culture or society. Social influence by the peer group is virtually always assimilative (as discussed earlier, research generally finds that friends tend to be similar; McPherson et al., 2001). But the self and peer group may be able to find shared positive identities either by assimilating to and upholding the values of the broader society or by contrasting from them and taking on a positively valued image as deviants and nonconformists. Teen subcultures such as Goths are particularly evocative examples of the latter.

Outside of social psychology, an intriguing class of "seceder" models also assume that people are motivated to seek distinctiveness—but not alone, only in company with others. In one such model (Dittrich, Liljeros, Soulier, & Banzhaf, 2000), every individual has a particular attitude position on a continuous scale. Each individual in turn looks at three randomly chosen others and selects the one whose attitude is the most different from the mean attitude of the three. This is a contrastive process, moving toward the extreme and away from the mean. The individual then moves himself near the position of that deviant individual. This is an

TABLE 5: Placement of Representative Models in the Four-Dimensional Conceptual Space

		Social Influence Only		*Other Sources of Input*	
		Assimilation Only	*Contrast and Assimilation*	*Assimilation Only*	*Contrast and Assimilation*
All-connect networks	Discrete	Innovation diffusion (Granovetter, 1978; Walker & Wooldridge, 1995)		Social decision schemes (Stasser, 1988)	
	Continuous		Seceder models (Dittrich, Liljeros, Soulier, & Banshaf, 2000; Jager & Amblard, 2005)		
Grid networks	Discrete	Dynamic social impact (Nowak, Szamrej, & Latané, 1990) Cellular automata		Information cascade (Bikhchandani, 1992) Swarm intelligence (Kennedy & Eberhart, 2001) Sznajd models (Stauffer, 2001)	
	Continuous				
Heterogeneous networks	Discrete	Innovation diffusion (Valente, 1995)		(Rojas & Howe, 2005)	
	Continuous	Structural theory of social influence (Friedkin, 1998)		(Janssen & Jager, 2001)	
Dynamic networks	Discrete Continuous	Culture model (Axelrod, 1997[a]; Hegselmann & Krause, 2002)			

[a]Agents sit in a fixed grid but interaction probabilities between neighbors change over time.

assimilative effect, moving toward the chosen individual. The individual behavior rule is essentially to become similar to someone who is different. In this regard all of these models, from self-categorization theory to ODT to Blanton and Christie's (2003) model to the seceder model, make broadly similar assumptions. Apparently people can find positive value in a deviant, nonconforming position but only in company with at least a few similar others. These and other models are organized in Table 4.

Implications for Maintaining Diversity

Adding the possibility of contrast or seeking for distinctiveness is one important way to maintain diversity of opinion in a model of social influence. Existing theories make very different predictions about when and how contrast will occur. These different models may all result in opinion diversity but predict quite different distributions of opinion. Researchers can now recommence the type of conceptual analysis Abelson (1964) performed on the subclass of purely assimilative models: using logical analyses, thought experiments, or (most likely) dynamic, multiagent computer simulation approaches to determine the consequences of different assumptions about the nature of social influence over time—results that can be compared to observed patterns of opinions in real social groups.

CONCLUSIONS

One general observation from our review is simply the large number of models of social influence that

exist—not only in social psychology but spread widely across other disciplines such as sociology, economics, political science, computer science, cognitive science, and even physics. We have described enough specific models to exemplify the major dimensions on which social influence models can vary. Many additional models, especially those that represent minor variations on other models, have not been covered in our review although they too fall within the overall structure depicted in Table 5. The large number of social influence models and especially the diversity of disciplines in which they have been developed indicate that social influence is a major concern across all the social sciences and beyond (especially if the structurally similar models of the spread of infection are also considered). What the majority of these models lack, however, is a deep understanding of individual-level psychological processes. Many, therefore, make empirically questionable simplifying assumptions (e.g., that people are rational and accuracy-seeking or that influence is always assimilative).

For social psychologists, one clear implication follows our conceptual review. If we believe that our field owns the topic of social influence, even that our field is defined by that topic, we should seek to be central participants in the development and testing of social influence models. To some extent, we have been: Our theories and research have yielded core insights into the processes of influence, the nature of attitudes versus discrete behavioral choices and their relationships, the factors that lead people to rely more on their private attitudes versus social influence, and the possibility of contrastive as well as assimilative influence. What our field has often failed to do is to contextualize our robust microlevel understanding of social influence processes by explicitly situating those processes in a social situation involving multiple individuals, interacting over time, linked in social networks of friendship and influence (Smith & Semin, 2004). Theorists from other fields have taken on the task of embedding social influence in such broader contexts, and they have often done so without social psychologists' participation. As a result, they have often employed simplistic, empirically inaccurate models of influence processes. One of our goals in this article is to call for cross-disciplinary conceptual integration to remedy this situation. If social influence is the name of our game, we should adapt and strengthen our methods and our models to meet those of theorists and researchers from across a wide swath of disciplines who are focusing on social influence. Together, we can work to understand how microprocesses of influence aggregate to produce broad social patterns of opinion distribution, rumor spread, majority decisions in groups, and even disease epidemics.

Although our review of specific models is not intended to be comprehensive, we have described and given examples of four fundamental dimensions on which social influence models vary. These dimensions constitute a conceptual framework for the entire class of possible social influence models, as shown in Table 5.

For each of the four dimensions, we identify a particular subset of models that we view as particularly promising or important for empirical or conceptual reasons.

1. Beyond simply including multiple influencers and targets over time, social influence models should make explicit the patterns of links between individuals. Specifically, we suggest that, because real-world connections (acquaintanceship, friendship, communication links, etc.) among individuals are most adequately characterized as networks (Wasserman & Faust, 1994), the most appropriate models for social influence will typically be heterogeneous or dynamic networks. This is in contrast to the more common assumptions that either (a) all individuals can equally influence all others (i.e., all-connect models) or (b) individuals sit in fixed locations and can influence and be influenced by only their close neighbors (i.e., grid models). Dynamic networks are especially ripe for exploration (although their study poses special empirical difficulties). Influence flows over network links to make linked individuals more similar; but, equally, people change their network links to make similar individuals more linked (Doreian, 2001).
2. We advocate modeling both attitudes (as continuous scales) and the resulting discrete behavioral choices. This approach has several advantages: The attitude serves to summarize the individual's inputs (socially provided or personally obtained information), whereas the discrete choice is the person's behavioral output that is visible to (and able to influence) others. Modeling only discrete behavioral choice while ignoring underlying attitudes may lead to artifactual preservation of diversity in the population. Continuous attitudes are more realistic models of mental representations but discrete behaviors often serve as the visible cues of internal states. Social psychological theories (the Theory of Reasoned Action and others) give us considerable leverage for modeling the intraindividual relationship between attitudes and behaviors. This approach also allows the modeling of conceptually important phenomena such as pluralistic ignorance.
3. Models should allow for individually obtained (private) information as well as individual differences in model parameters, such as the relative weight given to private tastes versus social influence. In other words, models should attempt to integrate social influence with other effects on individual decisions rather than to be models solely of social influence that assume people have no other nonsocial reasons to hold one opinion or another.
4. We believe that the most conceptually promising and empirically realistic models are those that allow for the possibility of contrast as well as assimilation resulting from social influence. As discussed earlier, various motives including reactance or the desire for distinctiveness can

drive people to move against a position taken by others rather than always to conform to others.

It is interesting to note that none of the models we have described have all four of the features we believe should be included. In fact, none of the models include both assimilation and contrast and sources of input other than social influence (empty last column of the table). Presumably this is because these are the most complex categories of models. And still other sources of complexity remain to be brought into the picture. For example, Jager and Amblard (2005) have begun to model the situation in which individuals have two attitudes (not just one), say an attitude toward a policy and an attitude toward a political leader. They are assumed to change the policy attitude either by thoughtful consideration of arguments, or simply by following the leader's guidance (if they approve of the leader), or by contrasting away from his recommendations (if they disapprove of him). Obviously, many other issues related to the interdependence of different attitudes, beliefs, and behaviors held by the same individual could be introduced into social influence models.

How can models incorporate all these various dimensions of complexity? We believe, with workers in several other disciplines, that agent-based modeling (ABM) techniques are the most suitable and appropriate for investigating dynamic models that include social networks (Macy & Willer, 2002; see Hastie & Stasser, 2000). Smith and Conrey (2007) recently presented an overview and introduction to ABM aimed specifically at social psychologists. The core idea behind this approach is that group-level outcomes of theoretical assumptions about intraindividual and interindividual processes are rarely obvious. When many individuals interact over time, their behaviors are interdependent, creating a complex, dynamic system that may have unpredictable (unexpected, emergent) outcomes. One way to elucidate the predictions of group-level outcomes from individual-level processes is through simulation. In ABM, many simulated individual agents act according to theoretically postulated behavioral rules and the interaction of the agents in the virtual environment generate group-level outcomes. For instance, Schelling (1969), seeking to understand the dynamics underlying residential segregation, simulated agents of two types living in a neighborhood, initially randomly intermixed. The behavioral rules were very simple: If an agent's neighbors consisted of more than a threshold percentage of agents of the other type, for example 50%, the agent moved to a new location. The counterintuitive result was that even when agents had a very high threshold or tolerance for outgroup neighbors, the final outcome was invariably almost completely segregated neighborhoods.

ABM is the tool of choice for social influence modelers outside of social psychology (e.g., Axelrod, 1997; Bikhchandani et al., 1992; Kennedy & Eberhart, 2001) and has been applied within our field as well, for example by Nowak et al. (1990). One important reason is that ABM is more flexible and general than are mathematically formulated models such as Friedkin's (1998) influence model (Smith & Conrey, 2007). For instance, many mathematically based models require simplifying assumptions about causation, such as "X causes Y but Y does not cause X," but causation in ABM is inherently multidirectional. Additionally, nonlinear effects (such as threshold dependent behavior) that in many cases make mathematical models intractable are easily incorporated into ABM.

Another advantage of the ABM approach is the ease with which one can test hypotheses. If one suspects that it is the range of initial attitudes rather than the strength of group norms that leads to group polarization, for example, one can simply vary the range of agents' attitudes while keeping the group norms constant, or keep the range constant while varying the strength of group norms. The results of the simulation under these different conditions can then be compared to experimental results. In fact, validation of ABM can be done at both the micro and macro levels (Moss & Edmonds, 2005), strengthening the model's falsifiability. One can ask

1. Does the model's assumption about individual agent behaviors match what is known about human social behavior? For example, how does knowing that a friend has bought a new type of consumer product generally influence people's intentions to purchase a comparable product themselves (Granovetter, 1978)?
2. Does the model's simulated outcome match what is observed in real human groups or populations? Do the model-generated results match aggregate data representing cumulative sales of new products over time?

Virtually all of the social-influence models described in this article can be validated or compared to data at both of these levels, individual and aggregate. Of course, a match at both levels increases confidence in the validity of the model. Validation of the assumptions regarding individual agent behavior is a task that is especially well suited for social psychology's most familiar and powerful research technique, lab-based experimental studies.

As an illustrative example, the ABM approach could assist in disentangling the sources of homophily within groups, an issue described earlier in this review. Homophily could in principle result from any of three processes: social influence that makes the attitudes of linked individuals become more similar, social selection processes by which people form links preferentially to

others with similar attitudes, and structural similarity in which linked individuals tend to have similar attitudes because they tend to occupy a shared position in the larger social structure (Burt, 1978). An agent-based model could be formulated, and multiple runs could be made with parameters varied to allow the operation of each of these processes alone or in combinations. The resulting distributions of attitudes could be compared to the real distributions of attitudes within groups. The agent-based model could be used not only to predict naturally occurring outcomes (e.g., attitude distributions, extent of observed homophily) but also responses to experimental manipulations (e.g., the introduction of a new group member with an extreme attitude). Again, such predictions can be compared to empirical data to validate the model. Validation of both individual-level processes (typically in laboratory experiments) and aggregate outcomes (in data collected from real groups) would provide strong support for the theoretical model represented by the agent-based model.

Ultimately, we hope that further work, with ABM techniques or others, can take social psychologists some distance further down the road toward fully adequate models of social influence in dynamic, multidirectional interactions while fully leveraging social psychology's traditional concern with empirically grounded descriptions of individual cognitive and communicative processes.

NOTES

1. The actual implementation of Nowak, Szamrej, and Latané's (1990) model differed from their statement of their theory. The theory involved an inverse-square law, with the amount of influence falling away rapidly with increasing distance between two individuals. However, for performance reasons in their program, they did not allow any influence from individuals beyond a certain threshold distance. Either of these specifications is essentially similar to a grid network pattern because each individual is influenced only by a set of the nearest agents (whether with a fixed cutoff over some distance or with a gradual falloff that approaches zero as distance increases).

2. An exception might be if a particular influence source was more often wrong than correct—either because the source is hopelessly misinformed or because it is hostile and trying to mislead. In such a case, contrasting against that opinion might be rational for one whose goal is to hold an accurate opinion.

3. Of course, there are also a number of social psychological theories related to contrast versus assimilation in judgment (e.g., on the effect of a salient reference point on the judgment of an object). These theories are not reviewed here because, unlike Brehm's (1966) theory of reactance, they are not theories focused on a perceiver's behavioral assimilation or contrast in relation to another person's attitude or behavior. It is true, of course, that in some cases judgmental assimilation or contrast might be part of an overall process that produces behavior.

REFERENCES

Abelson, R. P. (1964). Mathematical models of the distribution of attitudes under controversy. In N. Frederiksen & H. Gulliken (Eds.), *Contributions to mathematical psychology*. New York: Holt, Reinhart, & Winston.

Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, *50*, 179-211.

Axelrod, R. (1997). The dissemination of culture: A model of local convergence and global polarization. *Journal of Conflict Resolution*, *41*, 203-226.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*, 509-512.

Bavelas, A. (1950). Communication patterns in task-oriented groups. *Journal of the Acoustical Society of America*, *22*, 725-730.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*, 992-1026.

Blanton, H., & Christie, C. (2003). Deviance regulation: A theory of identity and action. *Review of General Psychology*, *7*(2), 115-149.

Brehm, J. W. (1966). *A theory of psychological reactance*. Oxford, UK: Academic Press.

Brewer, M. B. (1991). The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin*, *17*, 475-482.

Burt, R. S. (1978). Cohesion versus structural equivalence as a basis for network subgroups. *Sociological Methods & Research*, *7*, 189-212.

Cialdini, R. B., & Trost, M. R. (1998). *Social influence: Social norms, conformity and compliance*. New York: McGraw-Hill.

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *Journal of Personality and Social Psychology*, *51*, 629-636.

DiMaggio, P., Evans, J., & Bryson, B. (1996). Have Americans' social attitudes become more polarized? *American Journal of Sociology*, *102*, 690-755.

Dittrich, P., Liljeros, F., Soulier, A., & Banzhaf, W. (2000). Spontaneous group formation in the seceder model. *Physics Review Letters*, *84*, 3205-3208.

Doreian, P. (2001). Causality in social network analysis. *Sociological Methods & Research*, *30*, 81-114.

Fazio, R. H. (1986). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *The handbook of motivation and cognition: Foundations of social behavior* (pp. 204-243). New York: Guilford.

Festinger, L. (1957). *A theory of cognitive dissonance*. Oxford, UK: Row & Peterson.

Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention, and behavior: An introduction to theory and research*. Reading, MA: Addison-Wesley.

Friedkin, N. E. (1994). Structural cohesion and equivalence explanations of social homogeneity. *Sociological Methods & Research*, *12*, 235-261.

Friedkin, N. E. (1998). *A structural theory of social influence*. Cambridge, MA: Cambridge University Press.

Gordijn, E. H., De Vries, N. K., & De Dreu, C. K. W. (2002). Minority influence on focal and related attitudes: Change in size, attributions, and information processing. *Personality and Social Psychology Bulletin*, *28*, 1315-1326.

Gottman, J. M. (1994). *What predicts divorce*? Hillsdale, NJ: Lawrence Erlbaum.

Gould, R. V. (1993). Collective action and network structure. *American Sociological Review*, *58*, 182-196.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, *83*, 1420-1443.

Griffin, D. W., & Ross, L. (1991). Subjective construal, social inference, and human misunderstanding. *Advances in Experimental Social Psychology*, *24*, 319-359.

Hastie, R., & Stasser, G. (2000). Computer simulation methods for social psychology. In H. T. Reis & C. M. Judd (Eds.), *Handbook of research methods in social psychology* (pp. 85-116). New York: Cambridge University Press.

Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, *5*(3). Available at http://jasss.soc.surrey.ac.uk/JASSS/5/3/2.html

Hogg, M. A., & Turner, J. C. (1987). Social identity and conformity: A theory of referent information influence. In W. Doise & S. Moscovici (Eds.), *Current issues in European social psychology* (Vol. 2, pp. 139-182). New York: Cambridge University Press.

Hutchins, E. (1991). The social organization of distributed cognition. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 283-307). Washington, DC: American Psychological Association.

Jager, W., & Amblard, F. (2005). Uniformity, bipolarization and pluriformity captured as generic stylized behavior with and agent-based simulation model of attitude change. *Computational and Mathematical Organization Theory*, *10*, 295-303.

Janssen, M. A., & Jager, W. (2001). Fashions, habits and changing preferences: Simulation of psychological factors affecting market dynamics. *Journal of Economic Psychology*, *22*, 745-772.

Jin, M., Girvan, M., & Newman, M. E. J. (2001). The structure of growing social networks. *Physics Review E*, *64*, 381-399.

Katz, E., & Lazarsfeld, P. (1955). *Personal influence*. New York: Free Press.

Kennedy, J., & Eberhart, R. C. (2001). *Swarm intelligence*. San Francisco: Morgan Kauffmann and Academic Press.

Kleinberg, J. M. (2000). Navigation in a small world. *Nature*, *406*, 24.

Klovdahl, A. S. (1985). Social networks and the spread of infectious diseases: The AIDS example. *Social Science & Medicine*, *21*, 1203-1216.

Klovdahl, A. S., Potterat, J. J., Woodhouse, D. E., Muth, J. B., Muth, S. Q., & Darrow, W. W. (1994). Social networks and infectious disease: The Colorado Springs study. *Social Science and Medicine*, *38*, 79-89.

Krackhardt, D., & Porter, L. W. (1985). When friends leave: A structural analysis of the relationship between turnover and stayer's attitudes. *Administrative Science Quarterly*, *30*, 242-261.

Kumkale, G. T., & Albarracín, D. (2004). The sleeper effect in persuasion: A meta-analytic review. *Psychological Bulletin*, *130*, 143-172.

Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, *103*, 284-308

Latané, B. (1981). The psychology of social impact. *American Psychologist*, *36*, 343-356.

Latané, B., & Bourgeois, M. J. (1996). Experimental evidence for dynamic social impact: The emergence of subcultures in electronic groups. *Journal of Communication*, *46*(4), 35-47.

Lazarsfeld, P. F., Berelson, B., & Gaudet, H. (1944). *The people's choice: How the voter makes up his mind in a presidential campaign*. New York: Columbia University Press.

Leavitt, H. J. (1951). Some effects of certain communication patterns on group performance. *Journal of Abnormal Psychology*, *46*, 38-50.

Levine, J. M., & Moreland, R. L. (1998). Small groups. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., Vol. 2, pp. 415-469). Boston: McGraw-Hill.

Lifton, R. J. (1991). Cult formation. *Cultic Studies Journal*, *8*(1), 1-6.

Mackie, D. M. (1986). Social identification effects in group polarization. *Journal of Personality and Social Psychology*, *50*, 720-728.

Macy, M. W., & Skvoretz, J. (1998). The evolution of trust and cooperation between managers. *American Sociological Review*, *63*, 638-660.

Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, *28*, 143-166.

Mason, W. A., Jones, A., & Goldstone, R. L. (2006). *Propagation of innovations in networked groups*. Unpublished manuscript, Indiana University–Bloomington.

McCann, C. D., Higgins, E. T., & Fondacaro, R. A. (1991). Primacy and recency in communication and self-persuasion: How successive audiences and multiple encodings influence subsequent evaluative judgments. *Social Cognition*, *9*, 47-66.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*, 415-444.

Milgram, S. (1967). The small world problem. *Psychology Today*, *1*, 60-67.

Miller, D. T., & McFarland, C. (1987). Pluralistic ignorance: When similarity is interpreted as dissimilarity. *Journal of Personality and Social Psychology*, *53*, 298-305.

Moldovan, S., & Goldenberg, J. (2004). Cellular automata modeling of resistance to innovations: Effects and solutions. *Technological Forecasting and Social Change*, *71*, 425-442.

Monin, B., & Norton, M. I. (2003). Perceptions of a fluid consensus: Uniqueness bias, false consensus, false polarization, and pluralistic ignorance in a water conservation crisis. *Personality and Social Psychology Bulletin*, *29*, 559-567.

Moscovici, S., & Zavalloni, M. (1969). The group as a polarizer of attitudes. *Journal of Personality and Social Psychology*, *12*, 125-135.

Moss, S., & Edmonds, B. (2005). Towards good social science. *Journal of Artificial Societies and Social Simulation*, *8*(4). Available at http://jasss.soc.surrey.ac.uk/8/4/13.html

Newman, M. E. J. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences of the United States of America*, *98*, 404-409.

Nowak, A., Szamrej, J., & Latané, B. (1990). From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review*, *97*, 362-376.

Prislin, R., Limbert, W. M., & Bauer, E. (2000). From majority to minority and vice versa: The asymmetrical effects of losing and gaining majority position within a group. *Journal of Personality and Social Psychology*, *79*, 385-397.

Redner, S. (1998). How popular is your paper? An empirical study of the citation distribution. *European Physical Journal B–Condensed Matter*, *4*, 131-134.

Resnick, M. (1994). *Turtles, termites, and traffic jams: Explorations in massively parallel microworlds*. Cambridge, MA: MIT Press.

Rogers, R. W., & Prentice-Dunn, S (1997). Protection motivation theory. In D. S. Gochman (Ed.), *Handbook of health behavior research I: Personal and social determinants* (pp. 113-132). New York: Plenum.

Rojas, F., & Howe, T. (2006). *Contact patterns and aggregate opinion levels—Results from a simulation study*. Unpublished manuscript.

Rolfe, M. (2005). Social networks and simulations. In C. M. Macal, D. Sallach, & M. J. North (Eds.), *Proceedings of the agent 2004 conference on Social Dynamics: Interaction, Reflexivity, and Emergence* (pp. 483-499). Chicago, IL: University of Chicago Press.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*, 854-856.

Schelling, T. C. (1969). Models of segregation. *American Economic Review*, *59*, 488-493.

Sherif, M., & Hovland, C. I. (1965). Social judgment: Assimilation and contrast effects in communication and attitude change. Oxford, UK: Yale University Press.

Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology*, *70*, 893-912.

Smith, E. R., & Conrey, F. R. (2007). Agent-based modeling: A new approach for theory-building in social psychology. *Personality and Social Psychology Review*, *11*, 87-104.

Smith, E. R., & DeCoster, J. (2000). Dual process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, *4*, 108-131.

Smith, E. R., & Semin, G. R. (2004). Socially situated cognition: Cognition in its social context. *Advances in Experimental Social Psychology*, *36*, 53-117.

Snyder, M. (1974). Self-monitoring and expressive behavior. *Journal of Personality and Social Psychology*, *30*, 526-537

Stasser, G. (1988). Computer simulation as a research tool: The DISCUSS model of group decision making. *Journal of Experimental Social Psychology*, *24*, 393-422.

Stasser, G. (1999). A primer of social decision scheme theory: Models of group influence, competitive model testing and prospective modeling. *Organization Behavior and Human Decision Processes*, *80*, 3-20.

Stasser, G., Kerr, N. L., & Davis, J. H. (1989). Influence processes and consensus models in decision-making groups. In P. B. Paulus (Ed.), *Psychology of group influence* (2nd ed., pp. 279-326). Hillsdale, NJ: Lawrence Erlbaum.

Stasser, G., & Titus, W. (1985). Pooling unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, *48*, 1467-1478.

Stauffer, D. (2001). Monte Carlo simulations of Sznajd models. *Journal of Artificial and Simulated Societies*, *5*(1). Available at http://jasss.soc.surrey.ac.uk/5/1/4.html

Tajfel, H. (1978). *Differentiation between social groups: Studies in the social psychology of intergroup relations*. London: Academic Press.

Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D., & Wetherell, M. S. (1987). *Rediscovering the social group: A self-categorization theory*. Cambridge, MA: Basil Blackwell.

Turner, J. C., Wetherell, M. S., & Hogg, M. A. (1989). Referent informational influence and group polarization. *British Journal of Social Psychology*, *28*, 135-147.

Urbig, D. (2003). Attitude dynamics with limited verbalisation capabilities. *Journal of Artificial Societies and Social Simulation*, *6*(1). Available at http://ideas.repec.org/a/jas/jasssj/2002-38-2.html

Valente, T. W. (1995). *Network models of the diffusion of innovations*. Cresskill, NJ: Hampton Press.

Walker, A., & Wooldridge, M. (1995). Understanding the emergence of conventions in multi-agent systems. In *Proceedings of the First International Conference on Multi-Agent Systems* (pp. 384-389). Menlo Park, CA: Association for the Advancement of Artificial Intelligence.

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications. Structural analysis in the social sciences*. New York: Cambridge University Press.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of "small-world" networks. *Nature*, *393*, 440-442.

Weiss, W. (1953). A "sleeper" effect in opinion change. *Journal of Abnormal & Social Psychology*, *48*, 173-180.

Wolfram, S. (2002). *A new kind of science*. Champaign, IL: Wolfram Media.

Wolfson, L. B. (2003). A study of the factors of psychological abuse and control in two relationships: Domestic violence and cultic systems. *Dissertation Abstracts International*, *63*, 2794.

Wood, W. L. S., Ouellette, J. A., Busceme, S., & Blackstone, T. (1994). Minority influence: A meta-analytic review of social influence processes. *Psychological Bulletin*, *115*, 323-345.