CHAPTER 20

THRESHOLD MODELS OF SOCIAL INFLUENCE*

DUNCAN WATTS
PETER DODDS

Introduction

The study of social influence in decision-making has a long history in social science, dating back at least a century, to when economic philosophers like Rae (1905) and Veblen (1912) became concerned with what Veblen labeled 'conspicuous consumption'—a form of consumption whose primary purpose is to signal wealth to others. Although one could, in principal, signal wealth in very direct ways—say by burning piles of cash in public—Rae and Veblen noted that the wealthy typically prefer to purchase mansions and luxury items. To be acceptable, in other words, conspicuous consumption depends not only on the scarcity of the goods in question, but also on their social desirability (Robinson 1961)—an elusive and at times arbitrary-seeming quality that, like clothing fashions (Barber and Lobel 1952; Simmel 1957; Crane 1999), is driven by individual tendencies of both imitation

* The authors gratefully acknowledge the support of the National Science Foundation (SES-0094162 and SES-0339023) and the James S. McDonnell Foundation.

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 476 of 734 March 7, 2009 22:49

and also differentiation. Analogous connections between micro-level social influence and macro-level social change have been made subsequently to account for a wide range of phenomena, including scientific trends (Sperber 1990), business management (Blumer 1969), consumer and cultural fads (Johnstone and Katz 1957; Aguirre, Quarantelli, and Mendoza 1988); voting behavior (Lazarsfeld, Berelson, and Gaudet 1968), the diffusion of innovations (Katz and Lazarsfeld 1955; Coleman, Katz and Menzel 1957; Bass 1969; Valente 1995; Strang and Soule 1998; Young 2006) and word-of-mouth marketing (Gitlin 1978; Weimann 1994; Earls 2003).

A separate—and even earlier—root of social-influence studies appears to have been inspired by the observations of crowd behavior by early writers like Charles Mackay (1852), Gustave Le Bon (1895), and Gabriel Tarde (1903). At the risk of oversimplification, we suggest that this early work inspired at least two broad streams of productive research that have continued along different lines, and in different academic disciplines, to the present. The first such stream, carried out largely within experimental social psychology, deals with the micro-level psychological mechanisms leading to social conformity, as well as its converse, differentiation. The second stream, meanwhile, has focused on the more macroscopic question of collective action—for example, the dynamics of social movements or the provision of public goods—and has been largely the domain of sociologists, economists, and political scientists. We comment briefly on each of these fields in turn.

Since the 1930s, social psychologists (Sherif 1937; Asch 1953; Bond and Smith 1996; Cialdini and Goldstein 2004) have investigated the origins and effects of majority influence over individual members of the minority, as well as the converse problem (minority influence over the majority). The various manifestations of conformity and compliance, moreover, have been attributed to a multiplicity of psychological mechanisms (Deutsch and Gerard 1955; Cialdini and Goldstein 2004)—principally, desire for accuracy; desire to affiliate; and desire for positive self-image. In practice, it is usually unclear which of these mechanism is responsible for observed behavior, or even if the different mechanisms are analytically distinct (Cialdini and Goldstein 2004). Nevertheless, the tendency (whether conscious or unconscious) of individuals to conform to group norms and behavior has been invoked to account for a wide range of social phenomena, including local variability in crime rates (Glaeser, Sacerdote, and Scheinkman 1996; Kahan 1997), economic conventions (Young and Burke 2001), 'bystander inactivity' (Cialdini 2001), obedience to authority (Milgram 1969), residential segregation (Schelling 1971), and herd behavior in financial markets (Shiller 2000; Welch 2000).

The related literature on collective action is concerned less with the individual-level psychology of group participation, and more with the conditions under which groups of individuals can coordinate to achieve collective goals. As a result, collective-actions studies tend to be theoretical, rather than experimental, and are often framed in terms of the relative costs and benefits of coordinated action, rather than in terms of conformity per se. Nevertheless, the same psychological

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 477 of 734 March 7, 2009 22:49

mechanisms of social influence, such as desire to affiliate with a cause, and inferring its likely success from the participation of others, appear to be at least in part responsible for individual decisions to join in, or abstain from, collective action. Variations on these arguments, therefore, have been invoked to account for the success or failure of social movements (Kim and Bearman 1997), political uprisings (Kuran 1991; Lohmann 1994), contributions to public goods (Oliver and Marwell 1985; Ostrom et al. 1999), and other forms of collective action (Granovetter 1978; Chwe 1999; Macy and Flache 2002).

Finally, while much of the social-influence literature is concerned with non-market behavior, such as social movements, conformity to reference groups, and fashion, a related body of work that has attracted the attention of economists deals with a class of technology markets that exhibit what have been called 'network externalities' (Katz and Shapiro 1985), or somewhat more generally 'network effects' (Liebowitz and Margolis 1998). Both terms are meant to imply that the utility to an individual of a particular product (e.g. a fax machine) or skill (e.g. a language) is positively related to the number of other compatible products in use; that is, the size of the relevant 'network' associated with the product or skill. Katz and Shapiro (1985) further differentiated 'direct' from 'indirect' network effects, where the former refers to physical networks, such as the telephone network, which can presumably be owned, and the latter refers to virtual or metaphorical networks, like the population of users of some particular computer-operating system for whom availability of auxiliary products like compatible software, as well as support services, may influence purchasing decisions.

Social influence is thus not a singular phenomenon, or even (yet) a well-defined family of phenomena, but rather a blanket label for a loose congregation of social, psychological, and economic mechanisms, including: identifying with, or distancing oneself from, certain social groups; avoiding sanctions; obeying authority; reducing the complexity of the decision-making process; inferring otherwise inaccessible information about the world; gaining access to a particular network; or reaping the benefits of coordinated action. Precisely what these different mechanisms have in common, and to what extent their differences, when they exist, can be overlooked for the purpose of constructing models of individual choice, ought therefore to be a matter of considerable interest to 'analytical sociology'.

20.1. Social-influence and threshold models

Answering this question in general, however, is extremely difficult—the class of decisions under consideration is simply too broad to tackle all at once. In order to make some concrete progress, therefore, we restrict our analysis of social influence

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 478 of 734 March 7, 2009 22:49

to the class of 'binary choices', by which we mean choices between precisely two discrete alternatives. Although simple, binary-choice models can illuminate the dynamics of a surprising range of decisions, from the trivial (to dress in costume for the party or to dress normally; to cross the street or to wait for the walk sign) to the consequential (to join a strike or to keep working; to leave a neighborhood or to remain). Even decisions that involve choices between more than two alternatives—for example, which car to buy or which movie to see—may sometimes be represented as a sequence of binary choices (Borghesi and Bouchaud 2006)—say, to buy a car or not; to look for new or used cars; to buy through a dealer or privately; to prefer sedans to SUVs; and so on. Moreover, as Schelling (1978) argued, even quite complex decision-making processes, such as drafting a treaty on climate change, often culminate in binary choices—for example, to sign or not to sign. Binary choices have therefore received considerably more theoretical attention than other kinds of choices, and a wide range of models have been proposed by sociologists and economists; for example, diffusion models (Bass 1969); segregation models (Schelling 1971); coordination models (Schelling 1973; Morris 2000); social-learning models (Bikhchandani, Hirshleifer, and Welch 1992); threshold models (Granovetter 1978), and generalized contagion models (Dodds and Watts 2004, 2005).

Unfortunately, this proliferation of models has not yet produced an equivalently encompassing theoretical framework; thus, it is often unclear how one should relate the assumptions or findings of similar-sounding models to one another. For example, some social-influence models (e.g. Friedkin 1998) assume, in effect, that when A influences B the total quantity of opinion between A and B is conserved (in the manner of a gas diffusing between chambers); thus A's strength of opinion must diminish in order for B's to increase. Other models, meanwhile, are motivated instead by analogy with the spread of an infectious disease (e.g. Bass 1969); thus, A can 'infect' B with a new opinion, while A's strength of opinion remains undiminished. Either one of these assumptions may or may not be justified in any particular circumstance, but both models cannot be valid descriptions of the same process. Likewise, models of social learning that are typically studied in the 'information cascade' literature in economics (Bikhchandani, Hirshleifer, and Welch 1992) are formally quite different from the class of threshold models that have been studied by Granovetter (1978) and others—in a nutshell, the former assume that the particular order in which an individual observes the actions of others is important, whereas the latter assumes that it is not. Once again, in any given situation, it is either the case that the order of signals matters (in which case an information-cascade model is appropriate) or it does not (in which case an influence-response function is)—both models cannot be equally relevant to the same application. In practice, however, authors consistently invoke the same motivating examples—crowd behavior, group conformity, voting, diffusion of innovations, social movements, and herding in financial markets—regardless of what kind of model they then proceed to analyze.

Clarifying what assumptions are required in order for any given model to be a valid representation of the phenomenon in question, and how the various assumptions of different models can be related would therefore be extremely helpful steps towards a theoretically consistent and substantively interpretable theory of social-influence processes. Recently, Lopez-Pintado and Watts (2008*b*) have proposed that a number of mechanism-specific models are equivalent in the sense that they can all be described in terms of an 'influence-response function'—a one-dimensional function that maps the number of others choosing alternative A versus B into a probability that the focal actor i will choose A. By making explicit the assumption that signal order is unimportant, this framework excludes certain existing classes of models (e.g. social-learning models), and presumably also certain interesting phenomena (i.e. those in which signal order does matter). In restricting itself to certain phenomena and not others, however—by identifying clearly what it cannot explain—the explanatory power of the approach is arguably increased. For example, it can be shown that superficially quite different kinds of choices—inferring quality from observations of others versus deciding whether or not to contribute to a public good—can be shown to correspond to similar influence-response functions, and therefore may result in similar kinds of collective dynamics.

In other recent work Dodds and Watts (2004, 2005) have further partitioned the class of influence-response functions according to the importance of interactions between successive signals. In epidemiological models of contagion, for example—models that have been invoked by Bass (1969) and others as models of the diffusion of innovations—successive contacts between 'infected' and 'susceptible' individuals result in a new infection with constant probability; thus, infection 'events' are treated independently of one another. In threshold models, by contrast—models that are also used to model the diffusion of innovations—the probability that an additional positive signal will trigger adoption depends extremely sensitively on how many other signals have been observed: just below the threshold, a single observation can increase the adoption probability from near zero to near one, where otherwise it will have little effect. Epidemiological and threshold models of contagion are therefore quite distinct with respect to their (again, typically unstated) assumptions regarding the mechanism by which influence spreads from one person to another. What Dodds and Watts showed was that these differences can be captured in the shape of the influence-response function, and that, in fact, an entire family of contagion models can be specified in between the two cases. The shape of the influence-response function, moreover, can have important consequences for the conditions under which contagious entities, whether diseases, products, or ideas, can spread.

Whether or not a particular domain-specific example of social influence can be adequately described in terms of an influence-response function, say, and if so what the shape of the corresponding function should be, are ultimately empirical questions. Although empirical progress in this area is limited, some

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 480 of 734 March 7, 2009 22:49

recent progress is promising. By reanalyzing aggregate-diffusion curves (Griliches 1957), for example, Young has recently suggested that the diffusion of hybrid corn in the USA during the 1940s is better explained by a threshold model of adoption than by a Bass-style model of diffusion (Young 2006). Leskovec, Adamic, and Huberman (2007), moreover, have attempted to reconstruct the individual-level influence-response functions themselves, using online recommendations for books and movies. In general, such experiments are extremely difficult to perform; thus, it is still the case that formal models of social influence suffer from a dearth of realistic psychological assumptions. Nevertheless, a successful experimental program must be predicated on asking the right empirical questions, and in this respect a systematic formal-modeling approach of the kind we describe here is worth pursuing, if only as a means to focus empirical attention on the assumptions and parameters of greatest importance.

Bearing in mind this last objective, the subspace of social-influence models that can be represented as influence-response-functions influence models in turn exhibits a number of dimensions that ought to be of interest to the field of analytical sociology. Early work by Schelling (1978), and later Granovetter and Soong (1983), for example, indicated that binary decisions for which the corresponding decision 'externality' is negative—that is, when others' choice of A makes one less likely to choose A—generate qualitatively different dynamics than when the influence is positive—a result that has been studied in greater detail recently (Lopez-Pintado and Watts 2008*b*). Models of social influence, moreover, tend to assume (often implicitly) that all actors involved are of the same kind, whereas in reality individuals may be influenced by a variety of actors—for example, peers, role models, media organizations, and high-profile individuals like critics, celebrities, or increasingly 'bloggers'—each of which may exert a different kind of influence, and may in turn be influenced differently.[1] Bloggers, for example, exploit online media to share information and opinions with a potentially global audience, yet frequently engage in two-way, unfiltered conversations with readers; thus, they may be expected to exert influence that is neither like word of mouth, nor like mass media. Some recent modeling work (Lopez-Pintado and Watts 2008*a*) has begun to consider the effects of combined media and interpersonal influence, but this aspect of social-influence modeling is as yet poorly explored.

Another important yet understudied element of social influence—and the one that we focus on here—is that of influence networks. If it matters that people pay attention to one another, in other words, then surely it might also matter who pays attention to whom. Unfortunately, in spite of recent interest in the topic (Newman, Barabási, and Watts 2006), it is still the case that little is known about the structure of large-scale social networks, let alone how influence propagates on them. Kossinets and Watts (2006), for example, have studied the evolution of a relatively large network of email exchanges among the members of a university population. But this network is, at best, a representation of who *talks* to whom—it

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 481 of 734 March 7, 2009 22:49

tells us very little about who *pays attention* to whom, or with respect to what. Similar problems arise with other empirical examples of social networks, which typically utilize some proxy for social interaction, like comembership of corporate boards (Davis 1991) or coauthorship of a scientific paper (Newman 2004), rather than the social interactions themselves. It is plausible, in other words, that two people who sit on a board together might influence one another's opinions, but it is equally plausible that they do not, or that they do so with respect to only some issues, or that they are also influenced by numerous other unrecorded social relations. In the absence of relevant empirical data, therefore, we instead posit a series of theoretical models of networks, all of which are simple and unrealistic, but some of which are more realistic than others.

In order to make progress, we further narrow our focus to a special case of influence-response functions—namely, deterministic-threshold functions, according to which individuals adopt a new 'state' (e.g. wearing seat belts or joining a political uprising) based on the perceived fraction of others who have already adopted the same state. We choose to study threshold models for the practical reason that the collective dynamics of threshold models is already well understood in certain limiting cases—in particular, the 'all-to-all approximation' (Granovetter 1978), in which all individuals are influenced equally by the states of all others. Although the main purpose of this chapter is to consider the dynamics of social influence on networks, it is nevertheless helpful to anchor our results by reviewing the main features of Granovetter's model (Sect. 20.2). In the spirit of analytical sociology, we then proceed systematically up the chain of complexity, reviewing first the dynamics of 'cascades' of influence on random networks (Watts 2002) in which each individual i is exposed only to a fixed neighborhood of k others, drawn randomly from the population. We then introduce two models of networks that advance on the random-network model by including some simple notions of group structure (Sect. 20.3), and consider how these changes affect the likelihood of cascades for different seeding strategies. Although with each step up this chain the tractability of the corresponding models decreases, we are nevertheless able to make progress by leveraging our understanding of the simpler models that we have already considered.

20.2. Influence cascades on complete and random networks

Inspired by Schelling's seminal work on neighborhood segregation and coordination games (1969, 1973), Granovetter (1978) proposed a novel method for analyzing

the outcomes of collective action when individuals are faced with a choice to adopt some new ('active') state—a behavior, belief, or even an innovation—or else to remain in their existing, ('inactive') state. Granovetter illustrated the model with the example of a hypothetical crowd poised on the brink of a riot. Because all involved are uncertain about the costs and benefits associated with rioting, each member of the crowd is influenced by his peers, such that each of them can be characterized by some threshold rule: 'I will join a riot only when sufficiently many others do; otherwise I will refrain'. Granovetter did not specify an explicit theory of human decision-making from which the threshold model could be derived, and, as we have discussed, other kinds of rules are clearly possible (Dodds and Watts 2004; Lopez-Pintado and Watts 2008*b*). For the purpose of this analysis, however, we will accept Granovetter's informal reasoning that under some circumstances at least, a threshold rule is a plausible rule of thumb for an individual to follow, and instead focus on the consequences for collective dynamics of changing the influence network—that is, who pays attention to whom.

20.2.1. Granovetter's 'all-to-all' model

With respect to this last question, Granovetter made the simplest possible assumption—namely, that every individual in the population pays attention equally to all others—which in network terms corresponds to a 'complete' or an 'all-to-all' network. He then completed the model by allocating to each individual in the population a threshold ϕ_i, according to some probability distribution $f(\phi)$ (see Figs. 20.1*a* and 20.1*b* for an example), where the value of ϕ_i is assumed to capture all the relevant psychological attributes of individual i with respect to the particular decision at hand, and the distribution $f(\phi)$ represents both the average tendencies and also the heterogeneity present in the population. Lowering or raising the mean of $f(\phi)$, for example, would therefore correspond to lowering or raising the general susceptibility of the population, while increasing or decreasing the variance would correspond to an increase or decrease in variability in susceptibility across individuals.

Commencing in a population in which some fraction a_0 is assumed to have been activated by some exogenous process (and the remainder of the population $1 - a_0$ is therefore inactive), at each subsequent time step t, each individual i compares a_{t-1}, the active fraction of the population during time step $t - 1$, with their own threshold f_i, becoming (or remaining) active if $a_{t-1} \geq \phi_i$. The fraction of the population a_t that is active at any time t can then be described simply in terms of the active fraction at the previous time step, a_{t-1}, and the one-dimensional 'map' $a_t = F(a_{t-1})$, shown in Figure 20.1(c). The function F can be derived easily by observing that at any point in time t, a_t is just the fraction of the population whose thresholds fall below a_{t-1}; thus F is given by $F(a_{t-1}) = \int_{\phi=0}^{a_{t-1}} f(\phi)\mathrm{d}\phi$, which

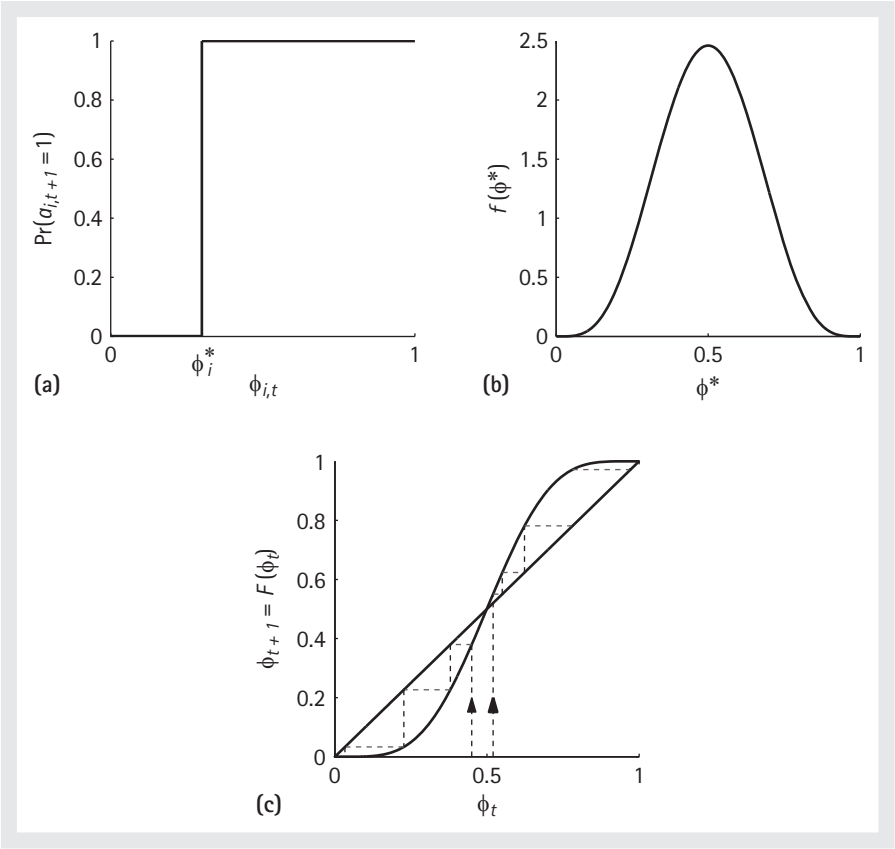

Fig. 20.1. Granovetter's threshold model of collective action

is simply the cumulative-distribution function of f (the threshold distribution) evaluated at a_{t-1}.

Once the system has been represented in this manner, its dynamics and equilibrium properties are surprisingly easy to compute using a simple graphical technique known as a 'cobweb' diagram—a common technique from the field of nonlinear dynamics (Strogatz 1994). Starting from any initial condition, a_0, the fraction of active individuals at the next time step, $t = 1$, is $a_1 = F(a_0)$ which can be found by drawing a vertical line on Figure 20.1(c) to intersect $F(\phi)$ at a_0. This fraction now becomes the input for the next 'iteration' of the map, which is achieved graphically by drawing a horizontal line from the map to the diagonal, also shown in Figure 20.1(c). The process now repeats, thus generating values for $a_2 = F(a_1)$, $a_3 = F(a_2), \ldots$, and so on until some equilibrium value a_∞ is reached, at which no further changes occur; that is, $a_\infty = F(a_\infty)$.

Equilibrium values of the map F are easy to identify graphically, as they are simply the points at which F intersects with the diagonal (i.e. where it maps

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 484 of 734 March 7, 2009 22:49

onto itself). Figure 20.1(c), for example, exhibits three equilibria—one at $F(0) = 0$, one at $F(1) = 1$, and one at an intermediate value $F(a_*) = a_*$. The last of these equilibria, however, is different from the first two: as can be verified using the cobweb technique, any initial condition will eventually converge on one of the first two equilibria, but will diverge from the intermediate value no matter how close to it one starts. The two extreme equilibria are therefore *stable*, whereas the intermediate value a_* is *unstable*. The equilibrium at a_* therefore acts as a kind of switch—small changes in the initial condition near a_* can result in polar-opposite outcomes, whereas away from a_* even large changes in the initial condition will converge on the same outcome.

Although extremely simple, this model already yields an important insight: that in the presence of social influence collective outcomes are not easily intuited from the individual attributes of the population itself.[2] In his hypothetical example of a crowd poised on the brink of a riot Granovetter observed that if the distribution of thresholds is precisely uniform—that is, one person will riot spontaneously, one person will join in when he observes one other rioter, another will join when he observes two others, and so on—the entire crowd will end up in the riot. This result, however, is exceedingly fragile with respect to perturbations in the distribution of thresholds (which in turn alters the number and nature of the corresponding equlibria). If, for example, no one has a threshold of three, and instead two individuals have a threshold of four, then the cascade will terminate after only three people have joined in. The two crowds would be indistinguishable in terms of their individual attributes, but their collective behavior could not be more different: rather than witnessing an all-out riot, an observer would see just three troublemakers jostling an otherwise orderly crowd. Consequently, small changes in individual preferences may lead to large and unpredictable system-level changes (a point that has also been made by Kuran 1991).

20.2.2. Extension to random networks

Recently, Watts (2002) has adapted Granovetter's threshold model to a network framework where, in contrast to the all-to-all assumption of Granovetter's model, individuals are assumed to be influenced directly only by a small subset of immediate 'neighbors'—a more realistic assumption for large populations. As with Granovetter's model, Watts (2002) considered a population of N individuals that is initially in a state of universal inactivity, and in which each individual i is randomly assigned a threshold ϕ_i^*, drawn from the distribution $f(\phi)$. Each individual is then assigned k_i neighbors, whom it both influences and is influenced by, where k_i is drawn at random from the distribution p_k. For the purpose of tractability, Watts considered networks in which neighbors were drawn at random—an unrealistic assumption in light of what is known about

real social networks, but one that provided a natural first step away from the all-to-all case. In the simplest case that Watts considered—that is, where all individuals have the same threshold—the population can therefore be characterized completely by just three parameters: N, ϕ^*, and k, where k is the mean degree (i.e. the average number of neighbors influenced directly by each individual). Even in such a simple case, however, the analysis is nontrivial and the results are counterintuitive.

At time $t = 0$ Watts assumed that the population would be 'shocked' by choosing some individual i_0 to be activated exogenously, meaning that their state is switched from inactive to active regardless of their threshold $\phi^*_{i_0}$ or their degree of influence k_{i_0}. If as a result of i_0's activation any of its neighbors' thresholds are now exceeded, those neighbors will also activate in the next time step, after which their neighbors may, in turn, activate, and so on, thus generating a 'cascade' of adoptions. At some time $T > 0$, where T is the duration of the cascade, some fraction of the population $S \leq 1$ will have been activated, and no more activations will be possible. (Once an individual is activated, we assume they remain so for the duration of the cascade.) Thus, the impact of every cascade can be quantified by its size S. Repeating this numerical experiment many times, it is possible to study the distribution of cascade sizes $g(S)$ for any particular population (as defined by the parameters N, ϕ^*, and k), and also to study the properties of $g(S)$ as a function of N, ϕ^*, and k.

Watts's (2002) main finding is illustrated in Figure 20.2: On random networks, 'global' cascades of influence can only take place within a certain region of the (ϕ^*, k) parameter space, called the 'cascade window', whereas outside this region cascades are typically small.[3] The extra condition of the cascade window is a major difference between Granovetter's all-to-all approximation and the more realistic case of sparse-influence networks, as now the success of a cascade depends not only on the individual attributes of the population (as captured by Granovetter in the threshold distribution), but also on the connectivity of the influence. Who pays attention to whom, in other words, is potentially every bit as important as how susceptible individuals are to a particular innovation or idea. The cascade window, moreover, has a particular shape, which Watts also explained in a manner that is illustrated in Figure 20.3. In brief, a cascade in its early stages can only spread via 'vulnerable' individuals who can be activated by only a single active neighbor. In order for a cascade to spread globally, therefore, the population must contain a connected network of vulnerable individuals that 'percolates', in the sense that it 'reaches' the entire population even though it may only be a small subset of the total (Stauffer and Aharony 1992). Global cascades can therefore occur if and only if the network contains what Watts called a 'percolating vulnerable cluster' (Watts 2002), but which might also be thought of as a 'critical mass,' meaning a relatively small population that, once activated, triggers a disproportionately large change in public opinion.[4]

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 486 of 734 March 7, 2009 22:49

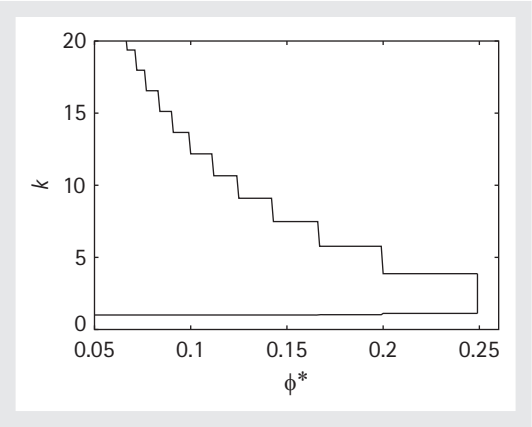

Fig. 20.2. The cascade window for random networks, where all members of the population have the same critical threshold ϕ^*, and k is the average degree of the influence network

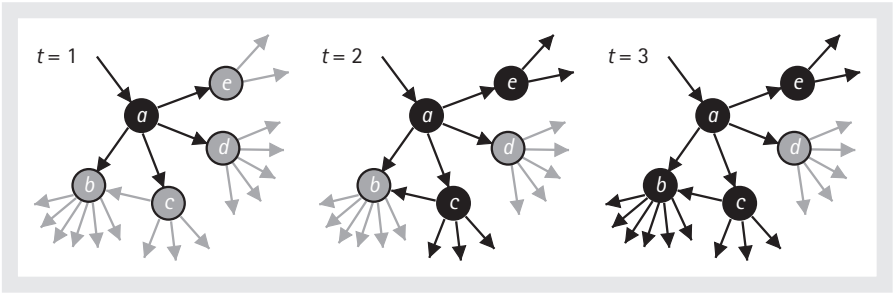

Fig. 20.3. Spread of influence throughout a network via vulnerable nodes. Black indicates a node or edge is active, and gray indicates inactive. All nodes in this example have a threshold $\phi^* = 0.18$, which means they are vulnerable if they have degree $k \leq 5$. At time $t = 1$ node *a* becomes active, as do its outgoing edges. At time $t = 2$ nodes *e* and *c* register that their thresholds have been exceeded and also become active. Node *b*, a nonvulnerable, switches on in time step $t = 3$, since now 2/8 = 25% of its neighbors are active

Since an individual's threshold is exceeded only when a specified *fraction* of its network neighbors are activated, the condition to be vulnerable is $\phi_i^* < 1/k_i$, where k_i is the degree of node i, Thus, influential (i.e. high-k_i) nodes are less likely than low-k_i nodes to be vulnerable. However, in order to *propagate* influence, an individual must be capable both of being influenced (which requires at least one neighbor) and also of influencing someone other than the source of the influence (requiring at least one additional neighbor); thus, only individuals

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 487 of 734 March 7, 2009 22:49

with $k_i \geq 2$ can contribute to the initial spread of influence. The upshot of these countervailing requirements is that a percolating vulnerable cluster—hence a critical mass—will only exist, and global cascades occur, when the average density k of the influence network is neither too low nor too high. For low k, although most nodes are vulnerable in the sense defined above, no large connected clusters exist, and cascades are confined to the small, isolated clusters in which they begin. On the other hand, when k is sufficiently large, the network will always exhibit a globally connected cluster (in graph-theoretic language, a giant component: Bollobas 2001), but too few of these nodes will be vulnerable. Lying in between these extremes is the cascade window, within which global cascades are possible.[5]

20.3. Cascades in networks with group structure

As interesting as they are from an analytical perspective, random networks are probably poor approximations of real social networks, for the simple reason that randomness overlooks the obvious importance of groups (Breiger 1974; Feld 1981; Blau and Schwartz 1984). People come together in well-defined, localized contexts—workplaces, schools, places of worship, clubs, and so on—that enhance the formation and maintenance of social connections. One might therefore expect not only that networks of influence relations will exhibit numerous characteristics of group structure, but also that these properties will have important consequences for the transmission of social influence across a network.[6] In this section we describe and analyze two models that emphasize, in different ways, the importance of social groups in the formation of influence networks.

20.3.1. Random-group networks

Consistent with our modeling strategy, our first class of networks with group structure constitutes only a modest departure from standard random networks (Newman, Strogatz, and Watts 2001), thus permitting us to test the effects of incorporating groups, while benefiting from our understanding of cascades on random networks. To build a random-group network we first create a standard random network with an average degree k_g (see Fig. 20.4). In two stages, we then replace each node in this network with a group of n_g nodes and then add edges between nodes: first, within each group, each pair of individuals is connected with

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 488 of 734 March 7, 2009 22:49

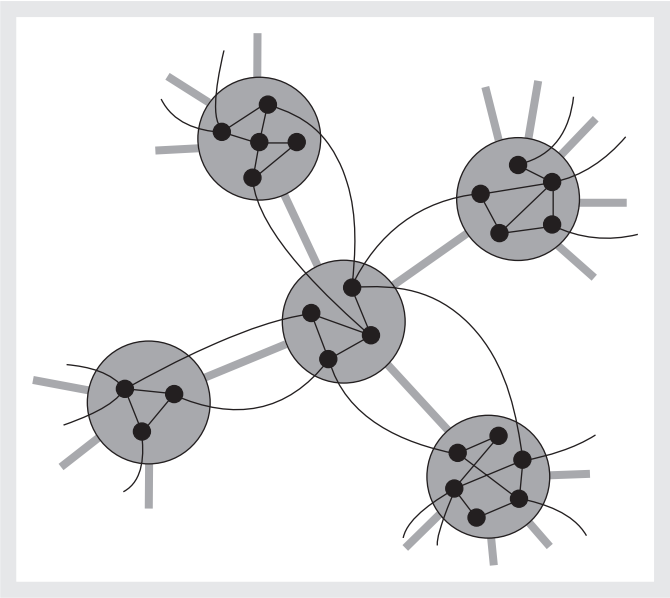

Fig. 20.4. Example of a random-group network. The gray disks and edges represent the underlying random network of groups

probability p, as per the typical construction of a standard random graph; and second, links between each pair of nodes that belong to adjacent groups on the underlying random network are created with probability q.

20.3.2. Generalized-affiliation networks

Our second class of networks with group structure is based on a model, first introduced by Watts, Dodds, and Newman (2002), that captures the effects of homophily (McPherson, Smith-Lovin, and Cook 2001) and group affiliation (Feld 1981) in determining social networks. In this model, each individual is allocated coordinates in each of H 'social dimensions' such as profession or geographic location, and 'social distance' d between two individuals is taken to be the minimum of all distances between their attributes. (For any given dimension, the distance between two attribute values is measured as the number of levels to the lowest common ancestor.) For example, in Figure 20.5 nodes i and j belong to the same group in the $h = 2$ attribute but are maximally far apart in the $h = 1$ attribute. The social distance d_{ij} is therefore 1 (we define two nodes sharing the same group to be a distance 1 apart) because i and j match in at least one attribute. Two individuals being socially 'close' is assumed to make a connection more likely; thus, for a

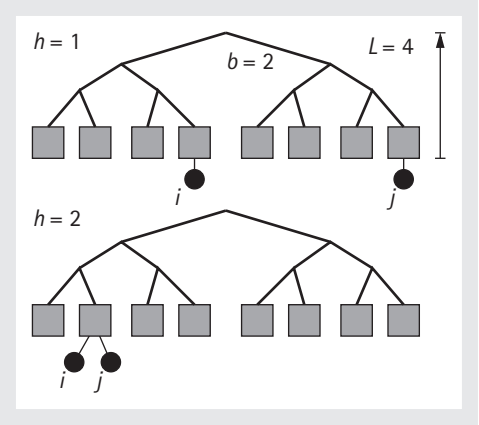

Fig. 20.5. Example of the generalized-affiliation model, where each node has attributes in two dimensions

population of individuals with identities assigned as above we realize a network by connecting each pair of individuals with probability $e^{-\alpha d}$ where d is social distance as defined above, and α is defined as the 'homophily parameter'—for high α, ties are almost always made between people who strongly match on at least one attribute, and for low α, ties may be made between more distant individuals. Finally, we depart from the original Watts, Dodds, and Newman (2002) formulation by allowing individuals who share a common acquaintance to connect with each other with probabilities β_1 and β_2, within and across groups respectively. The more acquaintances in common, the greater the chance of connection (Kossinets and Watts 2006); thus, we assume that any pair of unconnected individuals that are affiliated with the same group and have m mutual acquaintances within that group will have a connection probability of $1 - (1 - \beta_1)^m$.

20.3.3. Cascade-seeding strategies

In Figure 20.6 we show cascade windows for random-group networks, and generalized-affiliation-model networks, where, for purposes of comparison, we overlay the cascade windows of standard random networks (solid line) on those of the two networks with group structure. Our main interest here concerns seeding strategies that target multiple individuals simultaneously, in part because it is artificial to restrict seeding to a single individual, and in part because single-individual seeds do not exploit the presence of group structure. Nevertheless, we include the single-seed strategy in order to provide a direct comparison between

networks with group structure and the random networks studied previously. Figures 20.6*a* and 20.6*d* therefore show cascade windows for random-group networks and generalized-affiliation-model networks, for cascades that are triggered by a single-seed node, where the gray scale indicates the average size of cascades generated by a single, random node activated at time $t = 0$ (i.e. darker shade corresponds to larger cascades).

We find that both classes of networks involving group structure yield cascade windows different to that for standard random networks. For both classes of networks we see that the upper boundary of the cascade window exceeds that of standard random networks; that is, group-based networks can be vulnerable to activation cascades even when their average degree is significantly higher than that of the most vulnerable node. By enabling close-knit clusters of nodes to reinforce each others' adoptions, the introduction of groups therefore serves to push the upper bound of k well beyond its previous limit, where insufficiently many vulnerable nodes existed to form a critical mass. At the lower limit of the cascade window, however, we observe that the two classes of networks with group structure begin to differ: generalized-affiliation-model networks are similar to ordinary random networks; but, for reasons that we do not yet fully understand, the lower boundary for random-group networks is raised considerably.

The clear increase in the width of the cascade window in the presence of group structure also suggests a further question: Can seeds consisting of entire (albeit still small) groups trigger global cascades even under conditions where single-node seeds would fail? Figure 20.6 therefore also shows results for two other seeding strategies: (1) a random set of n_g nodes (Figs. 20.6*b* and 20.6*e*), and (2) a cohesive group of n_g nodes (Figs. 20.6*c* and 20.6*f*). A cohesive-group seed is a natural choice of seed in the current context as it uses the structure of the underlying network in an obvious manner—a marketer, for example, might provide a close-knit group of people with some free samples, or pay a team to use or endorse a particular product. It is also natural, however, to target an equal-size set of n_g individuals but scattered randomly throughout the network. Because it is not a priori obvious which of these two strategies will perform better, we compare them directly.[7] In both cases we also include (again for comparison purposes) the outline of the cascade window for the same seeding strategy employed on ordinary random networks (i.e. those studied previously).

On ordinary random networks the addition of larger seeds—created by first selecting a random node and then adding $n_g - 1$ of its nearest neighbors to the seed—can easily be shown to increase the frequency of global cascades when the network is inside the cascade window. It is also easy to show, however, that for standard random networks all three types of seed choices—cohesive-group, random-set, or a lone individual—lead to *exactly* the same cascade windows (Figs. 20.6(*a*)–(*f*)). The reason is that a random network either does or does not have a vulnerable cluster, and a finite seed in a random network with no vulnerable

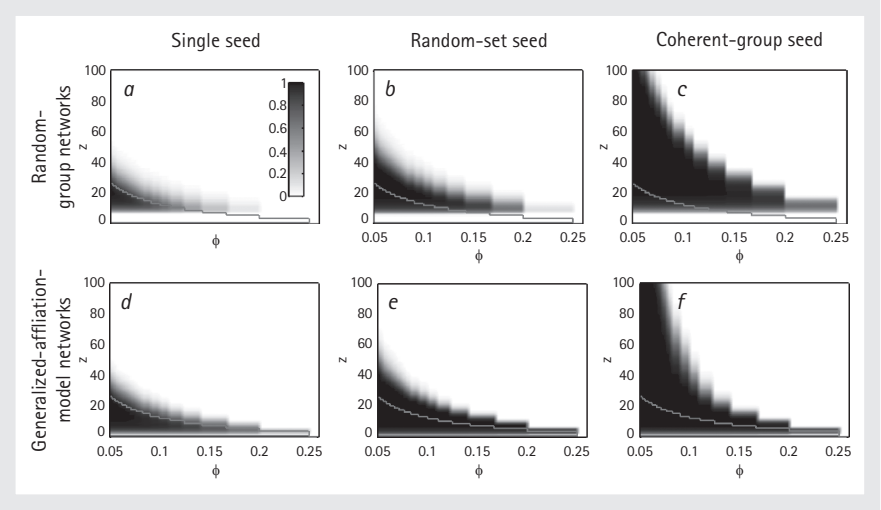

Fig. 20.6. Mean size of activation events for random-group networks and generalized-affiliation networks for three seeds: (*a*) single randomly chosen node; (*b*) a random set of $n_g = 10$ nodes; and (*c*) a cohesive group of $n_g = 10$ nodes. The solid line indicates the equivalent cascade window for standard random networks

cluster cannot generate a cascade. Within the boundaries of the cascade window, a larger seed, be it a cohesive group or a random set, has a better chance of hitting the percolating vulnerable cluster; thus, when global cascades are possible, the probability of generating a cascade is increased for larger seeds. At the same time, however, the nature of any finite seed does not alter the possibility of a cascade occurring in the first place; thus, the window remains invariant for random networks regardless of seeding strategy.

In networks with group structure the situation is different. First, we observe that for our two types of group-based networks activation cascades are generally more likely to occur over a broad range of k when the initial seed is a randomly chosen cohesive group of size n_g (Figs. 20.6*c* and 20.6*f*) as opposed to a randomly chosen set of n_g individuals (Figs. 20.6*b* and 20.6*e*). In particular, Figures 20.6*c* and 20.6*f* show that for both classes of networks, cohesive-group seeds can generate spreading when the average degree is an order of magnitude greater than that of the most connected vulnerable individual. Networks with group structure may therefore be highly vulnerable to influences initiated by cohesive groups, even when they are extremely resilient to social contagion, as we have modeled it here, when initiated by single seeds. Moreover, since we have assumed little knowledge of the networks in our simulations—only at the level of group membership—our results suggest easily implemented methods for increasing the spread of an influence in

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 492 of 734 March 7, 2009 22:49

real populations. Rather than targeting individuals who are thought to be influential (Katz and Lazarsfeld 1955; Weimann 1994), for example, a more successful strategy may be to target cohesive groups. Finally, these results also indicate that spreading is possible in group-based networks even when there are relatively few or even no vulnerable individuals (i.e. those individuals who are activated when only one neighbor is active), and certainly when there is no percolating vulnerable cluster.

Social influence in group-based networks therefore spreads in a way that is qualitatively distinct from spreading in all-to-all and random-network versions where, respectively, a nonzero vulnerable fraction and a percolating vulnerable cluster are needed for activation to spread. In the case of random networks, the largest vulnerable cluster can be interpreted as the *critical mass* of the system: when an individual in this cluster is activated, the rest of the cluster begins to follow in short order, whereupon nonvulnerables are also activated. Clearly, however, the vulnerable-cluster notion of critical mass on a network is insufficient to understand the dynamics of cascades in the presence of groups. Rather, it appears that when groups are the medium of transmission it is important to think of a critical mass in terms of the arrangement of vulnerable versus nonvulnerable *groups*, not individuals—a kind of 'renormalized' version of the previous conception. If an initial group is activated, the process within that group becomes self-reinforcing, since high levels of clustering within groups naturally maintain activation. Providing that neighboring groups have sufficient connections between them, activation will be able to spread, even when the individuals in question—when examined in isolation—would not appear vulnerable in the previous sense of being susceptible to activation by a single neighbor.

The vulnerability of group-based networks to social contagion is as yet a poorly understood phenomenon, but it clearly opens up new research questions as well as suggesting new possibilities for triggering, or preventing, cascades of social influence. As the results presented in this chapter make clear, changing the connectivity and topology of the influence network—even in possibly quite subtle ways—can have important implications both for the scale of cascades that may propagate throughout a population, and also the manner in which those cascades may be seeded. In addition to pursuing experimental studies of social influence at the level of individuals, therefore, we would argue that sociology in the analytical tradition espoused in this volume requires a more comprehensive theoretical understanding of the dynamics of social-influence networks—advanced, for example, through a systematic program of formal modeling—along with a tighter coupling between theory development and empirical testing. Aside from its interest to sociologists, moreover, a better understanding of the structure and dynamics of social-influence networks could be of value to marketers, public-health authorities, and indeed anyone concerned with affecting or understanding changes in public opinion, cultural beliefs, or social norms.

Notes

1. A number of diffusion models (e.g. Strang and Tuma 1993, Myers 2000) incorporate spatial and temporal heterogeneity; however, here we are making the somewhat different point that the actors involved in the diffusion process can themselves be heterogeneous.
2. At the broadest level, this insight is essentially the same as that derived from Schelling's much earlier work on residential segregation and coordination games (Schelling 1969, 1978). Nevertheless, Granovetter's model was considerably more transparent than Schelling's, allowing for individual heterogeneity and easy computation of equilibria; thus, it is Granovetter's model that we generalize here.
3. Strictly speaking, a 'global' cascade is one that occupies a finite fraction of the entire population in the theoretical limit of the population size $N \to \infty$, whereas nonglobal cascades are always finite in size. In practice, however, global cascades can be detected in finite populations simply by considering only cascades that exceed a prespecified cutoff size (where the precise choice of cutoff is unimportant).
4. Our use of the term is thus broadly consistent with Rogers (1995) definition as 'the point at which enough individuals have adopted an innovation so that the innovation's further rate of adoption becomes self-sustaining' (p. 313), but adds analytic power to the concept by specifying precise conditions under which a critical mass exists, regardless of whether any successful cascade is actually observed.
5. Recently, Whitney (2007) has shown that global cascades can also occur in a narrow region just above the upper limit of the cascade window, and that these cascades—which according to Watts's calculations should not take place—are driven by a slightly different mechanism than the one Watts proposed. Specifically, Whitney showed that cascades can occur outside of the theoretical cascade window as a consequence of triadic structures that he calls 'motifs', which can arise in sufficiently dense random networks. The local clustering implied by motifs can cause otherwise stable nodes to be vulnerable to the combined influence of two neighbors—a situation that does not arise in less dense random networks until the entire vulnerable cluster has been activated. Because motifs appear to matter only in a narrow region outside the cascade window, Whitney's findings are largely consistent with Watts's. Nevertheless, they suggest that even small variations away from pure randomness can lead to considerable additional complexity in the dynamics of social influence—a point we also make in Section 20.3.
6. Chwe (1999) has, in fact, made precisely this point in the context of somewhat smaller simulated networks than we consider here.
7. To ensure comparisons are fair, in all simulated networks the population size is $N = 10^4$, all groups are of size $n_g = 10$, and individuals have a uniform threshold of $f_i = 0.15$ (meaning nodes of degree six or less are vulnerable). The parameters for the generalized affiliation networks shown are $H = 2$, $b = 10$, $L = 4$, $f_i = 2$, $\mathrm{b}_1 = \mathrm{b}_2 = 0.5$.

References

Aguirre, B. E., Quarantelli, E. L., and Mendoza, J. L. (1988), 'The Collective Behavior of Fads: The Characteristics, Effects, and Career of Streaking', *American Sociological Review*, 53: 569–84.

Asch, S. E. (1953), 'Effects of Group Pressure upon the Modification and Distortion of Judgments', in D. Cartwright and A. Zander (eds.), *Group Dynamics: Research and Theory* (Evanston, III.: Row, Peterson), 151–62.

Barber, B., and Lobel, L. S. (1952), ' "Fashion" in Women's Clothes and the American Social System', *Social Forces*, 31: 124–31.

Bass, F. M. (1969), 'A New Product Growth for Model Consumer Durables', *Management Science*, 15: 215–27.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992), 'A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades', *Journal of Political Economy*, 100: 992–1026.

Blau, P. M., and Schwartz, J. E. (1984), *Crosscutting Social Circles* (Orlando, Fla.: Academic Press).

Blumer, H. (1969), 'Fashion: From Class Differentiation to Collective Selection', *Sociological Quarterly*, 10: 275–91.

Bollobas, B. (2001), *Random Graphs*, (New York: Academic).

Bond, R., and Smith, P. B. (1996), 'Culture and Conformity: A Meta-analysis of Studies Using Asch's (1952*b*, 1956) Line Judgment Task', *Psychological Bulletin*, 119: 111–37.

Borghesi, C., and Bouchaud, J.-P. (2006), 'Of Songs and Men: A Model for Multiple Choice with Herding', <http://www.lanl.gov/pdf/physics/0606224>, accessed 2008.

Breiger, R. (1974), 'The Duality of Persons and Groups', *Social Forces*, 53: 181–90.

Chwe, M. Suk-Young, (1999), 'Structure and Strategy in Collective Action', *American Journal of Sociology*, 105: 128–56.

Cialdini, R. B. (2001), *Influence: Science and Practice* (Needham Heights, Mass.: Allyn & Bacon).

——and Goldstein, N. J. (2004), 'Social Influence: Compliance and Conformity', *Annual Review of Psychology*, 55: 591–621.

Coleman, J. S., Katz, E., and Menzel, H. (1957), 'The Diffusion of an Innovation among Physicians', *Sociometry*, 20: 253–70.

Crane, D. (1999), 'Diffusion Models and Fashion: A Reassessment', *Annals of the American Academy of Political and Social Science*, 566: 13–24.

Davis, G. F. (1991), 'Agents without Principles? The Spread of the Poison Pill through the Intercorporate Network', *Administrative Science Quarterly*, 36: 583–613.

Deutsch, M., and Gerard, H. B. (1955), 'A Study of Normative and Informative Social Influences upon Individual Judgment', *Journal of Abnormal Social Psychology*, 51: 629–36.

Dodds, P. S., and Watts, D. J. (2004), 'Universal Behavior in a Generalized Model of Contagion', *Physical Review Letters*, 92: 218701.

—— (2005), 'A Generalized Model of Social and Biological Contagion', *Journal of Theoretical Biology*, 232: 587–604.

Earls, M. (2003), 'Advertising to the Herd: How Understanding Our True Nature Challenges the Ways We Think About Advertising and Market Research', *International Journal of Market Research*, 45: 311–36.

Feld, S. L. (1981), 'The Focused Organization of Social Ties', *American Journal of Sociology*, 86: 1015–35.

Friedkin, N. E. (1998), *A Structural Theory of Social Influence*, (Cambridge: Cambridge University Press).

Gitlin, T. (1978), 'Media Sociology: The Dominant Paradigm', *Theory and Society*, 6: 205–53.

Glaeser, E. L., Sacerdote, B., and Scheinkman, J. A. (1996), 'Crime and Social Interaction', *Quarterly Journal of Economics*, 111: 507–48.

Granovetter, M. S. (1978), 'Threshold Models of Collective Behavior', *American Journal of Sociology*, 83: 1420–43.

—— and Soong, R. (1983), 'Threshold Models of Diffusion and Collective Behavior', *Journal of Mathematical Sociology*, 9: 165–79.

Griliches, Z. (1957), 'Hybrid Corn: An Exploration of the Economics of Technological Change', *Econometrica*, 25: 501–22.

Johnstone, J., and Katz, E. (1957), 'Youth and Popular Music: A Study in the Sociology of Taste', *American Journal of Sociology*, 62: 563–8.

Kahan, D. M. (1997), 'Social Influence, Social Meaning, and Deterrance', *Virginia Law Review*, 83: 349–95.

Katz, E., and Lazarsfeld, P. F. (1955), *Personal Influence: The Part Played by People in the Flow of Mass Communications* (Glencoe, Ill.: Free Press).

Katz, M. L., and Shapiro, C. (1985), 'Network Externalities, Competition, and Compatibility', *American Economic Review*, 75: 424–40.

Kim, H., and Bearman, P. S. (1997), 'The Structure and Dynamics of Movement Participation', *American Sociological Review*, 62: 70–93.

Kossinets, G., and Watts, D. J. (2006), 'Empirical Analysis of an Evolving Social Network', *Science*, 311: 88–90.

Kuran, T. (1991), 'Now Out of Never: The Element of Surprise in the East European Revolution of 1989', *World Politics*, 44: 7–48.

Lazarsfeld, P. F., Berelson, B., and Gaudet, H. (1968), *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign* (New York: Columbia University Press).

Le Bon, G. (1895), *The Crowd* (New Brunswick, NJ: Transaction).

Leskovec, J., Adamic, L. A., and Huberman, B. A. (2007), 'The dynamics of viral Marketing', *ACM Transactions on the Web*, (1), at <http//:www.tweb.acm>, accessed 2008.

Liebowitz, S. J., and Margolis, S. E. (1998), 'Network Effects and Externalities', in P. Newman (ed.), *The New Palgrave Dictionary of Economics and the Law* (New York: Macmillan), 671–5.

Lohmann, S. (1994), 'The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989–91', *World Politics*, 47: 42–101.

Lopez-Pintado, D., and Watts, D. J. (2008*a*), 'Mass Media Versus Word of Mouth', Paper, Institute for Social and Economic Research and Policy, Columbia University.

—— (2008*b*), 'Social Influence, Binary Decisions, and Collective Dynamics', *Rationality and Society* Forthcoming.

Mackay, C. (1852), *Memoirs of Extraordinary Popular Delusions and the Madness of Crowds* (London: Office of National Illustrated Library).

McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001), 'Birds of a Feather: Homophily in Social Networks', *Annual Review of Sociology*, 27: 415–44.

Macy, M. W., and Flache, A. (2002), 'Learning Dynamics in Social Dilemmas', *Proceedings of the National Academy of Sciences of the USA*, 99: 7229–36.

Milgram, S. (1969), *Obedience to Authority* (New York: Harper & Row).

Morris, S. (2000), 'Contagion', *Review of Economic Studies*, 67: 57–78.

Myers, D. J. (2000), 'The Diffusion of Collective Violence: Infectiousness, Susceptibility, and Mass Media Networks', *American Journal of Sociology*, 106: 173–208.

NEWMAN, M. E. J. (2004), 'Coauthorship Networks and Patterns of Scientific Collaboration', *Proceedings of the National Academy of Sciences of the USA*, 101: 5200–05.
—— BARABÁSI, A.-L., and WATTS, D. J. (2006) (eds.), *The Structure and Dynamics of Networks* (Princeton, NJ: Princeton University Press).
—— STROGATZ, S. H. and WATTS, D. J. (2001), 'Random Graphs with Arbitrary Degree Distributions and their Applications', *Physical Review E* 6402: 026118.
OLIVER, P. E., and MARWELL, G. (1985), 'A Theory of the Critical Mass, 1. Interdependence, Group Heterogeneity, and the Production of Collective Action', *American Journal of Sociology*, 91: 522–56.
OSTROM, E. et al. (1999), 'Revisiting the Commons: Local Lessons, Global Challenges', *Science*, 284: 278–82.
RAE, J. (1905) *The Sociological Theory of Capital* (London: Macmillan).
ROBINSON, D. E. (1961), 'The Economics of Fashion Demand', *Quarterly Journal of Economics*, 75: 376–98.
SCHELLING, T. C. (1969), *Neighborhood Tipping* (Cambridge, Mass.: Harvard University Institute of Economic Research).
—— (1971), 'Dynamic Models of Segregation', *Journal of Mathematical Sociology*, 1: 143–86.
—— (1973), 'Hockey Helmets, Concealed Weapons, and Daylight Saving: A Study of Binary Choices with Externalities', *Journal of Conflict Resolution*, 17: 381–428.
—— (1978), *Micromotives and Macrobehavior* (New York: Norton).
SHERIF, M. (1937), 'An Experimental Approach to the Study of Attitudes', *Sociometry*, 1: 90–8.
SHILLER, R. J. (2000), *Irrational Exuberance* (Princeton, NJ: Princeton University Press).
SIMMEL, G. (1957), 'Fashion', *American Journal of Sociology*, 62: 541–58.
SPERBER, I. (1990), *Fashions in Science: Opinion Leaders and Collective Behavior in the Social Sciences* (Minneapolis, Minn.: University of Minnesota Press).
STAUFFER, D., and AHARONY, A. (1992), *Introduction to Percolation Theory* (London: Taylor & Francis).
STRANG, D., and SOULE, S. A. (1998), 'Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills', *Annual Review of Sociology*, 24: 265–90.
—— and TUMA, N. B. (1993), 'Spatial and Temporal Heterogeneity in Diffusion', *American Journal of Sociology*, 99: 614–39.
STROGATZ, S. H. (1994), *Nonlinear Dynamics and Chaos: With Applications in Physics, Biology, Chemistry, and Engineering* (Reading, Mass.: Addison-Wesley).
TARDE, G. (1903), *The Laws of Imitation* (New York: Holt).
VALENTE, T. W. (1995), *Network Models of the Diffusion of Innovations* (Cresskill, NJ: Hampton).
VEBLEN, T. (1912), *The Theory of the Leisure Class: An Economic Study of Institutions* (New York: Macmillan).
WATTS, D. J., DODDS, P. S., and NEWMAN, M. E. J. (2002), 'Identity and Search in Social Networks', *Science*, 296: 1302–05.
—— (2002) 'A Simple Model of Information Cascades on Random Networks', *Proceedings of the National Academy of Sciences of the USA*, 99: 5766–71.
WEIMANN, G. (1994), *The Influentials: People Who Influence People* (Albany, NY: State University of New York Press).
WELCH, IVO. (2000), 'Herding Among Security Analysts', *Journal of Financial Economics*, 58: 369–96.

978–0–19–921536–2 Hedstrom-Main-drv (Typeset by SPi, Chennai) 497 of 734 March 7, 2009 22:49

Whitney, D. E. (2007), 'Exploring Watts' Cascade Boundary', Working Paper, Massachusetts Institute of Technology.

Young, H. P. (2006), 'The Spread of Innovations by Social Learning', working paper, Johns Hopkins University.

—— and Burke, M. A. (2001), 'Competition and Custom in Economic Contracts: A Case Study of Illinois Agriculture', *American Economic Review*, 91: 559–73.