Social Networks and Collective Action

David A. Siegel Florida State University

Despite growing attention to the role of social context in determining political participation, the effect of the structure of social networks remains little examined. This article introduces a model of interdependent decision making within social networks, in which individuals have heterogeneous motivations to participate, and networks are defined via a qualitative typology mirroring common empirical contexts. The analysis finds that some metrics for networks' influence—size, the prevalence of weak ties, the presence of elites—have a more complex interaction with network structure and individual motivations than has been previously acknowledged. For example, in some contexts additional network ties decrease participation. This presents the potential for selection bias in empirical studies. The model offers a fuller characterization of the role of network structure and predicts expected levels of participation across network types and distributions of motivations as a function of network size, weak and strong ties, and elite influence.

Individuals do not make political decisions in a vacuum. Across social science, a wealth of empirical evidence illustrates the ways in which social interactions can alter choice. One can see this most clearly in the political participation (e.g., Huckfeldt and Sprague 1995; Kenny 1992; Leighley 1990; McClurg 2003) and social movement (e.g., Chong 1991; Kuran 1991; McAdam 1986; Petersen 2001) literatures, but the importance of others' actions in one's decisions is also evident in substantive areas as diverse as regulatory enforcement (e.g., Scholz and Wang 2006), the diffusion of democracy (e.g., Gleditsch 2002), and the adoption of new organizational, policy, or technological innovations (e.g., Berry and Baybeck 2005). Yet, despite the commonality of social, political, and economic networks' empirical importance, we still know little about how the structure of these networks affects aggregate political outcomes, and less about how network structure interacts with individual motivations. The theory I develop here elucidates the causal role of social structure in shaping group outcomes, allowing the prediction of aggregate participation levels from qualitative network factors.

A deeper understanding of how social structure translates into aggregate outcomes is important for several reasons. For one, extant research on political participation indicates that one's network size affects the likelihood that one will engage in political activity (e.g., Lake and Huckfeldt 1998; Leighley 1990; McClurg 2003). Yet this is given as an average effect only; without a completely detailed network, we do not know if a group of people whose large network size is due to membership in a completely closed clique will participate more than a group of people each with smaller network size, but more disparate social connections. This has consequences beyond political participation as well. Individuals in the Senate are more highly interconnected than those in the House (Fowler 2006); should we expect a priori the former to display more homogeneous voting behavior than the latter?

We also know that the structure of networks helps determine both the availability of expertise and the levels of conflict and political sophistication to which one is exposed, all of which have an impact on the willingness to participate (e.g., Huckfeldt 2001; McClurg 2004, 2006; Mutz 2002). Clever experiments have verified that these results are robust to, among other things, the endogeneity of network formation (Klofstad 2007; Nickerson 2008), but we still do not know whether the large-scale structure of the network alters these results at all. To get at this question we must be able to ask the important counterfactual: how would the political outcome have been different had the network(s) been different?[1] Again, the

David A. Siegel is assistant professor of political science, Florida State University, 541 Bellamy, Tallahassee, FL 32306-2230 (dsiegel@fsu.edu).

The author thanks Dave Baron, Jon Bendor, Bill Berry, Ray Block, Matt Golder, Sona Golder, Jenn Jerit, Özge Kemahlioglu, Charlotte Lee, Scott McClurg, Will Moore, Chris Reenock, John Scholz, Jake Shapiro, Jeff Staton, the Political Institutions Working Group at FSU, and several anonymous reviewers.

[1]To be fair, this would be difficult even in the most careful empirical work, due to an absence of observable counterfactuals. For an interesting experimental take on this problem, see Kearns, Suri, and Montfort (2006).

American Journal of Political Science, Vol. 53, No. 1, January 2009, Pp. 122–138

 ISSN 0092-5853

consequences go beyond participation. For instance, how much weaker would the incentive to engage in violence have been for members of the global Salafi Jihad had existing cliques been weaker (Sageman 2004, 152–55)?

Better understanding the role of networks can also aid scholars studying collective action in determining how important network structure is to participation in their particular case. Knowing which aspects of network structure are likely to be most relevant can help to focus data collection, conserving limited resources. This also opens up new opportunities for substantive theoretical arguments and their empirical testing that have for too long been underexplored. For example, how did the structure of Iraqi networks affect voting in the January 2005 Iraqi legislative elections (Siegel 2008), or how did media coverage of George Allen's snafus interact with network effects to alter voter turnout in the 2006 Virginia Senate race?

To clarify the role of network structure and so allow questions like these to be addressed, I propose and analyze a model of aggregate behavior that includes the explicit pathways across which individuals influence each other. Local social networks delimit this influence: another's behavior influences you if and only if that person is within your network. I focus on a particular form of influence, in which individuals become more likely to participate the more others do as well, though the underlying cause of this influence (e.g., information, social pressure, sanctioning, or safety in numbers) can vary by substantive area.[2]

The model's networks address the frequent empirical condition that the rich data necessary to reconstruct an entire social network are often unavailable. A typology of networks, shown graphically in Figure 2 in the fourth section, is designed to represent commonly observed social structures: modern cities, villages or cliques, opinion leader networks, and hierarchies. Many results of the article are conditional only on network type, and these types are differentiated via rough, qualitative factors. For example, rigid bureaucratic hierarchies tend to admit lower levels of participation than the more fluid connections within modern cities, all else equal. However, each network type also is defined by one or two empirically estimable parameters; the model thus yields hypotheses testable with more complete, though still rough, data. For example, the ability of negatively predisposed social elites to suppress participation decreases as their followers become more interconnected, but this power of the "proletariat" is conditional on the structure of the hierarchy.

In addition to network-specific results, the model also offers more general insights. One, the relationship between network size and aggregate participation is conditional on the distribution of individual motivations in the population, as well as on network structure. People who have intrinsically low motivations—and thus need more urging to participate—can be discouraged by an excess of network connections, particularly if these ties are "weak" in the sense of Granovetter (1973). Besides illustrating the conditional nature of weak ties,[3] this suggests that increasing network size will have different effects in different varieties of participation. When participation requires little urging, size should be positively correlated to aggregate participation regardless of the structure of the network. When it is very costly, however, size is likely to have more mixed, network-dependent effects. This raises the issue of selection bias; if one only observes successful social movements or only considers certain forms of political participation, one risks overgeneralizing about the role of network size.

Two, like weak ties, the power of social elites depends strongly on the structure of the network in which they reside. The fact that elites are in a privileged position in the network does not imply they necessarily have a strong influence on aggregate levels of participation, as the hierarchy example above indicates. If elites cannot control their networks, therefore, their power to effect change may be greatly lessened.

As the model introduced here draws strongly on a tradition of mathematical sociology that has remained largely distinct from any literature in political science, the next section details the model's pedigree, with ties to the political participation literature as appropriate. The following two sections analyze the model in stages: first using three simple networks, then the four more complex networks in the typology. The fifth section offers a brief summary of major hypotheses drawn from the model in qualitative terms; a more technical description of these and the model itself can be found in an online appendix.

Interdependency and Social Structure

While literature in political science has only recently begun to consider in depth the role of social structure in understanding interdependent behavior, there is a broad and growing interdisciplinary literature on the topic. I

[2]With minor changes the model extends to issues of public goods provision in which people start becoming less likely to participate after a certain point.

[3]Centola and Macy (2007) make this point in the context of a Small-World network. My analysis indicates the phenomenon is general to different network types.

focus on that most directly related to the model's development and its primary application to participation in collective actions.[4] Even so, the volume of the literature implies I can only sample here; for more, see excellent reviews in Centola and Macy (2007), Oliver (1993), and Strang and Soule (1998).

In order to explore interdependency in networks one must first specify the nature of the interdependency. I follow that proposed in such seminal works as Granovetter (1978), Schelling (1978), and Marwell and Oliver (1993):[5] the more people who participate, the more likely it is that one will decide that it is in one's own best interest to participate as well. In Heckathorn's typology of collective actions, this model is thus closest to the assurance game (1996, 259).

The behavioral referents for this type of interdependency are varied. In the political participation literature, information transfer plays a major role in the relevance of networks (e.g., Huckfeldt 2001; McClurg 2006); this is a significant focus in formal studies of mass movements, protest, and rebellion as well (e.g., Lohmann 1994; Neill 2005). In general, information exchange allows people to update their beliefs about the costs and the benefits inherent in participation, and so change their decisions. Networks can also coordinate and transfer resources, which have an independent effect on one's willingness to participate (Verba, Schlozman, and Henry 1995).

Less common in political science, but more in sociology, is the idea that networks transmit direct influence, changing one's interests in and inherent motivations toward participation (e.g., Friedkin and Johnsen 1999; Gould 1993; Klofstad 2007; Oliver and Myers 2003). Closely related to direct influence are notions of reputation (e.g., Gerber, Green, and Larimer 2008; Klofstad 2007; Kuran 1995; Mutz 2002) and fairness (cf. Gould 1993), in which either negative (you worry about being punished) or positive (you worry about acting unfairly to others) social pressures encourage you to act or not to act. Finally, typically confined to the movement, protest, and rebellion literatures is the safety in numbers argument; you are safer the more others join your actions (e.g., Kuran 1991).[6]

These all act in the same way, in that others' participation increases one's own desire to participate, and so may be incorporated in the same model. However, a model covering multiple types of behavior makes it vital to include heterogeneous motivations. After all, it is difficult enough to assume that everyone responds the same way to information, let alone information *and* social pressure. However, as Granovetter and Soong (1983), Kim and Bearman (1997), and Yin (1998) show, heterogeneity in interests can by itself have a substantial effect on the expected level of participation in a collective action. For example, using a threshold model in which people participate only once a certain percentage of others are already taking part in the action, Granovetter (1978) shows that small differences in the distribution of individual thresholds can completely alter the outcome. To see how this works, note that if there are 10 people with thresholds arrayed from 0 to 90% in steps of 10, full participation results. Each new participant "tips" the next into joining the movement. If we change the person with a 10% threshold to a 20% one, however, the potential cascade stalls after the movement has only a single member.[7]

This brings us, finally, to network structure itself. As discussed in the previous section, individuals' decisions depend strongly on their social contexts; this has been noted particularly often in the sociological participation literature (e.g., Gould 1991; Hedström 1994; McAdam and Paulsen 1993; Ohlemacher 1996; Opp and Gern 1993; Snow, Zurcher, and Ekland-Olson 1980). Mississippi's Freedom Summer is an example of this; as McAdam writes: "Participants were much more likely than withdrawals to have had ties—especially strong ties—to other volunteers" (1986, 80). We might therefore expect the structure of people's connections to alter outcomes. Indeed, this is the case; both the pattern of network connections (e.g., Centola and Macy 2007; Chwe 1999; Gould 1993) and the position of individuals within networks (e.g., Borgatti and Everett 1992; Gould 1993; Kim and Bearman 1997) play a role in the decision of whether or not to participate.

In sum, then, there is extensive evidence that the structure of networks, in terms both of the pattern of connections and of the way in which individuals are distributed across them, alters aggregate outcomes. Further, there is a tradition of exploring this relationship theoretically within sociology that has rarely been addressed within political science despite its relevance. Rolfe (2005) and Fowler (2005) are two exceptions. Each addresses the spread of behavior within a single network type, varying its parameterization in order to make inferences about

[4]Cooperation (e.g., Ohtsuki et al. 2006) is a separate topic, though many results should generalize.

[5]I am considering only the "accelerating" regions of Marwell and Oliver's production functions here, as this article is more concerned with the problem of getting a participatory bandwagon going than with the inevitable "deceleration" when free-rider problems begin to take hold.

[6]This list is not exhaustive; see Centola and Macy (2007, 707–8) for a nearly nonoverlapping set.

[7]Centola and Macy (2007, 724) find that heterogeneity itself does not change some results on network structure; however, they appear to vary only the mean of the distribution of interests. Results below echo this finding, but illustrate the large effect caused by varying the standard deviation of interests in the population.

how large-scale network statistics such as average network size, path length, and clustering/density alter behavioral spread across all networks. In contrast, this article's contribution is to provide a general description of how network structure and individual motivations interact to determine participation in collective behaviors, in a manner not reliant on detailed knowledge of said network structure.

Basic Network Dynamics

Model

The core behavioral model this article employs rests upon two assumptions. One, individuals have varied motivations to participate. Two, individuals adjust their desires to participate over time, in response to the behavior of those to whom they are connected via local networks.

To keep the first assumption simple and widely relevant, I separate each individual's motivations to participate at a given time into two components. The first component, which I call *net internal motivation* and label b_i for each individual i, covers all potential motivations both for and against participation that are *independent* of the participation of others, from a driving need to effect social change to the opportunity cost of missing work while doing so. The second component, which I call *net external motivation* and label $c_{i,t}$ for each individual i at each time t, covers all potential motivations both for and against participation that are *dependent* on the participation of others. Note that these two components are mutually exclusive. Any motivation whose change over some time frame is a function of others' participation is by definition an *external* motivation; all others are necessarily *internal* motivations. Which motivations are in each component can vary by action. For example, a network might discuss the opportunity cost of turnout but not of protest; thus this motivation would be classified as external under turnout, but internal under protest. As long as the classification of motivations into nonempty components can be considered fixed over some time frame, the separation of motivations in this way is not a strong assumption.

Given these two components, the decision rule is simple: an individual i participates at time t if and only if $b_i + c_{i,t} > 0$, i.e., only if her net motivation to participate is positive.[8]

As motivations within the population are heterogeneous by assumption, they must be distributed initially in some fashion. Consider net internal motivation first. I assume these follow a normal distribution, with parameters b_{mean} and b_{stdev}. If one's net internal motivation—one's b_i—is the sum of elements distributed across the population independently of all other such elements, weighted in a particular way, then a central limit theorem implies that b_i is distributed normally, with some mean and standard deviation.[9] The two parameters of this distribution dictate the mean net internal motivation and the dispersion of the net internal motivation within the population. Higher mean values imply collective actions in which people are, all else equal, more likely to participate. Thus the model accounts for the possibility that different forms of participation—voting versus contacting officials, for instance—might behave differently (Leighley 1990; or, in the language of Centola and Macy 2007, that some contagions are simple, others complex).

To keep the focus on the impact of networks, I assume that nothing exogenous to the effects of networks occurs during the time frame over which people alter their motivations. This implies that all b_i remain constant, and only the $c_{i,t}$, which by definition cover all network-related factors, change.[10]

Now consider net external motivation. As stated earlier, the model assumes the more people within one's local social network who participate, the more that one desires to participate oneself. Call the local participation rate $lpr_{i,t}$ for individual i at time t. Then a parsimonious definition of net external motivation consistent with the behavioral referents discussed in the previous section lets $c_{i,t} = -(1 - lpr_{i,t})$. As the local participation rate for an individual increases from zero to its maximum of one, external motivations increase, making that individual more likely to participate. Under this definition net external motivations range from -1 to 0 for all time, and their initial values are $c_{i,0} = -(1 - lpr_{i,0})$, where the $lpr_{i,0}$ are the initial local participation rates. I follow the literature stemming from Granovetter (1978) and consider only the case where there is no participation at first, so that $lpr_{i,0} = 0$. If we define "rabble-rousers" as those people with b_i

[8]As written this rule is deterministic, but little changes if we let behavior be probabilistic. See the online appendix for details (http://myweb.fsu.edu/dsiegel/Research/Siegel_network_model_AJPS_appendix_final.pdf).

[9]Even if these assumptions do not all hold, this only implies that a normal distribution might not be appropriate or that it might not be common to all members of the population. The basic idea of each individual's having some net internal motivation—upon which this model rests—still holds.

[10]This has the added benefit of not washing out heterogeneity over time, so that individuals remain different in motivations even if all converge in action; compare to Gould (1993). It also allows for cascades which are not all-or-nothing affairs; compare to Neill (2005). It is, however, relaxed in Siegel (2008).

> 1 then, since $c_{i,0} = -1$, all initial participation is due to these people. Note that this definition is independent of network position; rabble-rousers may be network elites, or uninfluential members of a network.

The definition $c_{i,t} = -(1 - lpr_{i,t})$ dictates precisely the level of external motivations at every time given the local participation rate at that time.[11] One might argue that this is *too* parsimonious, lacking in the kind of behavioral generality claimed above. To illustrate that this is not the case, consider the following far more general rule for adjusting external motivations: one increases one's propensity to participate the more people in one's local network participate as well. Further, let us normalize the maximal value of $c_{i,t}$ to $-(1 - lpr_{i,t})$.[12] Then an array of specific behavioral rules including Bayesian (e.g., Lohmann 1994), "naïve" Bayesian (Neill 2005), or any one of a number of adaptive methods (e.g., Macy 1991a) converge over time to the simple definition for $c_{i,t}$ given here. Thus, as long as we consider only steady-state properties, we are well justified in opting for parsimony. Further, this simple rule does have some theoretical heft to it—it incorporates both positive influence (when many neighbors participate and $c_{i,t}$ is large), and negative influence (when few neighbors participate and $c_{i,t}$ is small; Mutz 2002).

The model is dynamic, and each realization begins with the assignment of internal and external motivations to individuals, and their placement within the appropriate network. In an online appendix I detail the method of network creation and network parameterization; here the focus is on the qualitative properties of the networks most relevant to the discussion of the results, and most useful for empirical work. In this section consideration is limited to three simple networks chosen for pedagogical reasons rather than realism; empirically relevant networks are discussed in the following section. The first is a Fully Connected network, in which everyone is tied to everyone else; this is the world of Granovetter (1978). Only the number of people, N, and the internal motivation distribution parameters, (b_{mean}, b_{stdev}), are variables here. The second is a Random network, visualized in the upper-left quadrant of Figure 1 below, in which individuals are connected to each other by chance. This type of network admits one additional parameter: the likelihood that any two individuals are connected. The third is a Ring network, visualized in the upper-right quadrant of Figure 1 (located later in the section), in which individuals are arrayed in a great circle and connected only to their neighbors. For Ring networks, the additional parameter is the distance along the ring to either side of an individual that makes up that person's local social network.

In every period after the network is created, individuals gauge participation within their local networks and use this local participation rate to adjust their net external motivation, as described above. This occurs simultaneously for all individuals. It is important to keep in mind that, while one might visually observe the participation of individuals outside one's local network, this does not affect $c_{i,t}$ because one is not tied to these outsiders, and so by definition receives no influence, credible information, or eventual punishment from them. Due to a lack of exogenous factors, no one has an incentive to stop participating once one starts. Thus external motivations either increase or stay the same in each period. After a number of periods specific to the particular realization of the network and the distribution of b_i, an equilibrium or steady-state is reached. At this point no more individuals are spurred to participate. In a Fully Connected network this model matches Granovetter's threshold model discussed in the previous section. Here $1 - b_i$ is i's threshold, and i participates if this threshold is less than the rate of participation in i's local network.[13]

This equilibrium value of participation is the dependent variable under consideration in this article. Given its importance, it helps to see how an equilibrium is derived in a simple, concrete setting. Let us begin by assuming that there are six people arranged in what is typically called a "star" formation, with one in the center and the other five tied only to her. Further, let internal motivations be such that the central person's b_i exceeds one, while everyone else's are between zero and one, exclusive. In the first period the points of the star will not participate, since the central individual is not yet participating, and so $b_i - 1 < 0$ for them at this point. The central person will, however, begin participating immediately; she is a rabble-rouser. Once this has occurred the local participation rates for all other individuals become one, since all are connected only to the center by assumption. This drives their net external motivation to the maximum of zero, causing them to participate as well. All individuals thus participate within two periods.

Now assume the same setting, save that the central, influential person is a less-motivated type, rather than a rabble-rouser. Further, let the sole rabble-rouser in the

[11] Note that the net external motivation is a positive linear function of the local participation rate; there are no diminishing marginal returns to, or net negative incentives arising from, others' participation.

[12] This amounts to setting a scale for $c_{i,t}$ and b_i, which we are free to do since the decision rule compares only b_i and $c_{i,t}$. The lack of bounds on the distribution of b_i (as compared to c_i) allows for rabble-rousers.

[13] This equivalence to a threshold model vanishes if we allow exogenous events like repression.

FIGURE 1 Random and Ring Networks (Top Row: Visualizations of the Networks) (Bottom Row: Comparison of Mean Participation Rate vs. Average Network Connectivity in Weak and Strong Motivation Classes)

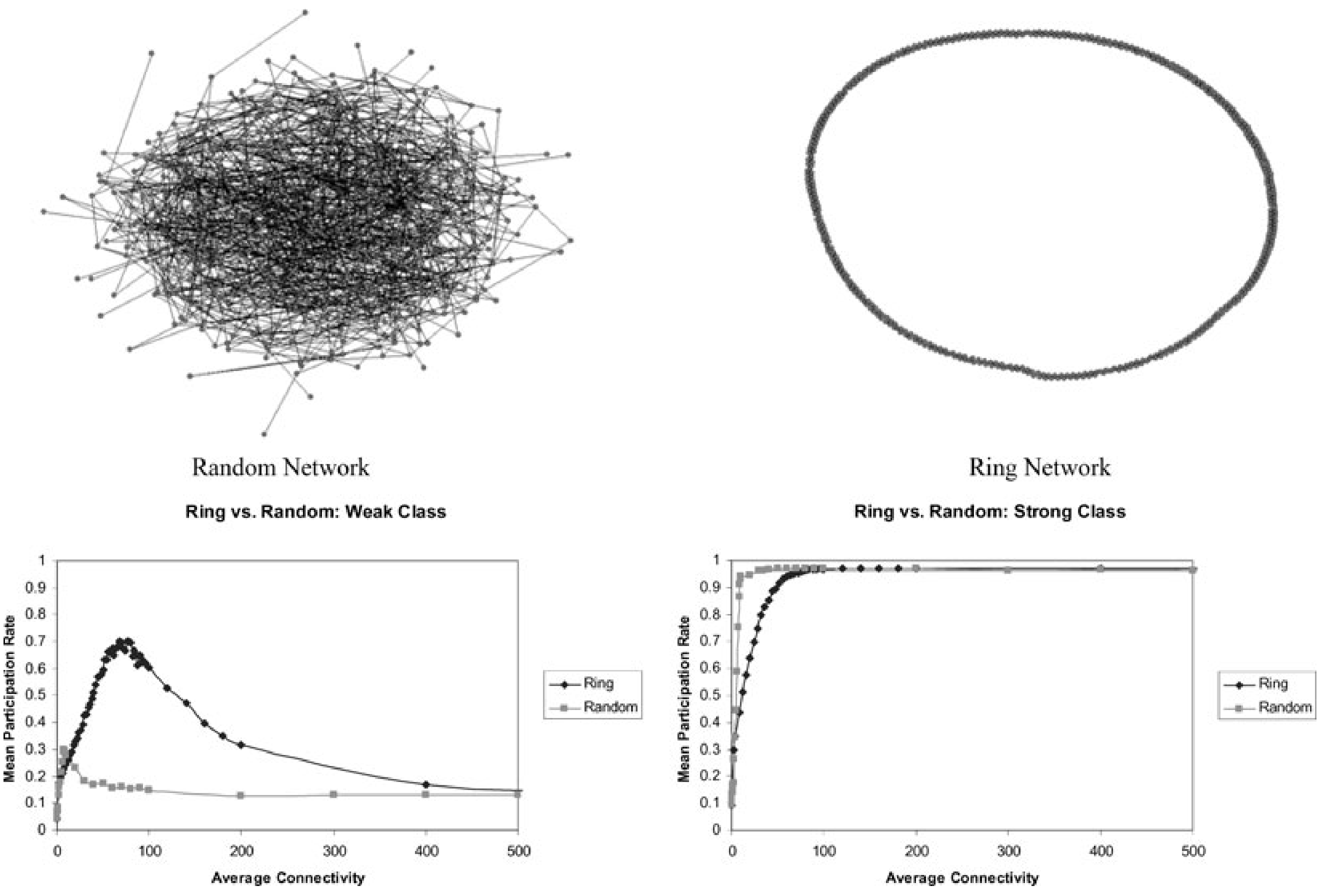

network be located at one of the points of the star. As before, the rabble-rouser participates in the first period. However, in this network configuration only one person is affected by this decision—the central person. One-fifth of his local network is participating, driving his external motivation to negative four-fifths.[14] If his internal motivation is greater than this he participates, triggering a cascade as in the earlier case. If it is not, however, he does not participate, and the movement stalls at one person. In short, initial distributions of motivations that produce identical levels of participation in the first period can lead to very different subsequent behavioral dynamics, due to the way in which the shape of the network structures the way individuals perceive others' participation. Though simple, this example is illustrative of the two paths the evolution of this model generally takes. In the first, participation spreads sufficiently widely so as to cause a cascade, leading to very high levels of participation in equilibrium. In the second, participation spread stalls at some point, and only lower levels are achieved. Which path is taken depends on both the distribution of motivations within the population and on the structure of network ties.

This complex relationship between equilibrium participation and both network structure and the distribution of internal motivations necessitates careful analysis. Accordingly, using a method I describe and justify at length in the online technical appendix, I build and analyze the model in discrete stages that permit direct computation of comparative statics. This technique reduces the number of parameters under consideration at any one time, but requires prior theoretical justification. As detailed in the online appendix, Granovetter and Soong (1986; see also Yin 1998 and Rolfe 2005) provide this for the Fully Connected network, showing that we can expect up to three equilibria in an infinite population. These correspond to regions of low, medium, and high participation. I use this theoretical finding to break up the space spanned by the parameters b_{mean}, b_{stdev}, and N into three

[14] $c_{i,1} = -(1 - 0.2) = -0.8$ in this case; see page 125. b_i must exceed the absolute value of this.

regions, which I denote "motivation classes." Each class, entitled "weak," "intermediate," and "strong," produces on average (in a Fully Connected network) low, medium, and high degrees of participation, respectively. These are best thought of as the rates at which cascades of participation are achieved. Because behavior under networks is similar within each motivation class, one can explore network parameterizations without simultaneously varying all parameters in a five-dimensional parameter space. The results below are a fraction of those taken; further ones can be obtained from the author upon request.

Results

Before speaking in terms of motivation classes, however, it is useful to discuss the independent effects of the parameters (b_{mean}, b_{stdev}) and N. The action of the first two can most easily be examined by assuming an infinite population and making a simple comparison between what happens in the absence of a network, and what happens when a Fully Connected network is present. In the former case, people only participate if $b_i > 1$, since external motivations do not adjust. The participation rate is thus $\bar{p} = 1 - \Phi(\frac{1-b_{mean}}{b_{stdev}})$, where Φ represents the standard normal CDF and $\bar{p}$ is the equilibrium global participation rate. In the latter case, people can learn or become influenced to participate due to the actions of others. The steady-state participation rate in this case is the point at which extant participation spurs no change in people's choices, or $\bar{p} = 1 - \Phi(\frac{(1-\bar{p})-b_{mean}}{b_{stdev}})$.

There are several things of note here. First, in both cases increasing b_{mean} increases participation levels. This is unsurprising—when people are more likely to want to participate, perhaps because the action is not very costly, they end up doing so more often. Second, the impact of the dispersion is nonmonotonic. Increasing b_{stdev} increases participation when the mean is low, but decreases participation when the mean is very high. This is largely due to its effects on the numbers of "rabble-rousers" and "wet blankets," those who will always and who will never participate, respectively.

Third, there is nearly always more participation when some form of a network is present than when none is, and never less. Individual action is sufficient to produce significant levels of participation only when the population is extraordinarily self-motivated.

Fourth, the *increase* in participation levels when a network is added is strongly nonlinear, and conditional on the motivation class. In one set of parameters, (0.6,0.23), connecting everyone produces an increase of only 5% in an infinite population, from 4% to 9%, while in another, (0.6,0.3), doing so increases participation by 88%, from 9% to 97%. Further, even when network connectivity is effective in spurring participation, some classes might benefit more than others. For example, in a third set of parameters, (0.6,0.25), connecting everyone increases participation by 94%, from 5% to 99%.

The lesson of this quick analysis will arise repeatedly in more complex settings. Increasing network size and altering network structure may have significant results in some settings but not in others; where any particular case falls depends on the distribution of internal motivations in the population. Interdependency thus enters in a fundamentally interactive way, and empirical models seeking to understand its impact must take this into account.

These conclusions are true not only for infinite populations, but also for finite ones. The difference is that smaller populations increase the model's sensitivity to heterogeneity. This has two effects. One, it means that we must discuss mean effects over multiple realizations of the model, rather than deterministic relationships. Two, increased randomness changes the likelihood that a cascade is achieved, thereby altering the average level of participation reached in the population. When a cascade is unlikely, shrinking the population makes "lucky" placements of a few rabble-rousers more common, increasing mean participation. When a cascade is likely, in contrast, shrinking the population makes "unlucky" draws, in which there are too few rabble-rousers or in which they are distributed poorly, more common, decreasing participation.[15] In terms of our motivation classes, then, increasing N increases participation in the strong class and decreases it in the weak class. Changing the population thus has a nonmonotonic impact on the level of participation achieved in equilibrium. This matches the results of Rolfe (2005), who finds the same relation under a different behavioral model.[16]

Having explored the impact of varying each of b_{mean}, b_{stdev}, and N independently and in concert, the next step is to understand the role of network structure. As stated earlier, the impact of structure on participation varies strongly by motivation class, but is similar within each class. Therefore, in the figures below I choose a single common set of three parameters to use as exemplars of

[15] These effects would diminish were there diminishing marginal returns to others' participation. I thank an anonymous reviewer for pointing this out.

[16] Given Olson's (1965) focus on group size, and empirical literature which suggests it matters to participation rates (e.g., Sandell and Stern 1998), population size as an independent variable has been surprisingly understudied in the literature on participation (cf. Oliver and Marwell 1988; Rolfe 2005).

each class.[17] Each member of the trio is located within a different motivation class, and so this design ensures that differences due to network structure, conditional on class, are readily apparent. However, one should keep in mind that the results displayed in the figures and discussed in the text hold far more generally than for just these three sets of parameters.[18] For a Fully Connected network, these parameters yield mean participation levels of 13%, 64%, and 97% on average, implying that they well represent the classes for which they are exemplars.

We turn now to two simple networks: the Random and the Ring. Below the visualizations of the Random and the Ring networks in Figure 1 are two plots summarizing the results of 1,000 realizations of the model for each set of parameter values. The graph on the left displays the weak motivation class for both the Random and the Ring networks, the one on the right the strong class. The intermediate class is similar in nature to the weak class for these networks and so is not shown. In these and in all such plots in this article, the y-axis corresponds to the dependent variable of interest, the equilibrium rate of participation, averaged over 1,000 realizations of the model. One realization of the model encompasses everything from the initial distribution of the population to the final recording of the equilibrium participation rate. The x-axes here display the average network size (or connectivity) for each individual, i.e., the average number of people to whom each individual is connected.

Together, these two plots illustrate the impact of increasing connectivity on the equilibrium rate of participation. Consider first the plot on the right that displays the strong motivation class. Increasing average connectivity in either network increases participation until it matches that obtained in a Fully Connected network. This is as one would expect—if one increases the average connectivity enough, eventually everyone is connected. What is less predictable a priori is the rate at which participation increases in connectivity. In the Random network this happens quickly; each additional connection greatly increases equilibrium participation, and an average connectivity of 40 almost always yields a cascade. In the Ring network this happens more slowly; eventually nearly everyone participates almost all the time, as in a Random network, but it requires an average connectivity of about 70 connections out of a population of 1,000. Thus, even holding motivations constant, network size is insufficient to predict participation; one must also know the manner in which the connections are arranged, the structure of the network.

Now consider the plot on the left, which displays results for the weak class. Though the plot displays the same two networks as in the strong class, the outcomes are very different in the weak class. The most important thing to note is that, while a Fully Connected network yields only 13% participation in equilibrium, both networks are able to achieve *higher* levels of participation with *fewer* average connections per person. These are the mixed effects described in the introduction: increasing average connectivity increases participation up to a point, but any further increases begin to decrease participation rates in equilibrium.

While formal models allow for increases, decreases, or constancy in participation as a function of connectivity, depending on the exact circumstances in play (e.g., Fowler 2005; Gould 1993; Kim and Bearman 1997; Macy 1991b; Marwell, Oliver, and Prahl 1988), empirical work generally supports the view that increased connectivity increases participation (e.g., Gould 1991; Putnam 2000; Snow, Zurcher, and Ekland-Olson 1980; Tilly 1978). Figure 1 offers an explanation: it depends on the internal motivations in the population. When rabble-rousing types abound, due either to a high mean level or a high dispersion of internal motivations, additional connections make it more likely that individuals are tied to them, increasing the likelihood that ties affect participation positively. At the extreme, in the strong motivation class, participation will monotonically increase in network size. In this case, adding social capital (in the form of network ties) will increase levels of participation, supporting arguments like Putnam's (2000).

In contrast, when rabble-rousing types, who participate regardless of the behavior of others in their local networks, are less common, increasing size will have a more mixed effect (Centola and Macy 2007).[19] This can lead to issues of selection bias in empirical work. Focusing only on successful actions or on a subset of political activities may yield a set of cases in which motivations are particularly high. This is of most concern in the social movement literature, in which data and interest are more likely to be substantial for the most successful movements; quickly abandoned movements rarely make headlines, after all.

[17]The trio of values is (0.6, 0.23,1000), (0.6, 0.25,1000), and (0.6, 0.3,1000). These and all other parameter settings necessary to replicate the figures shown here may be found in the online appendix.

[18]Note, however, that not all classes will be present for all parameter values. For example, when N is very large, few values of (b_{mean}, b_{stdev}) fall into the intermediate class, while when N is very small, most values do. This particular trio of values was chosen in part because all three classes are represented, making illustration easier.

[19]Note that this mixed effect is not simply a consequence of a population that cares less overall about the cause—the two plots in Figure 1 use *identical* mean motivation levels. Only the dispersions differ, implying that it is the interaction between structure and the location of individuals that matters.

However, the political participation literature is not immune; aggregating types of participation might produce false generalities about the impact of network size, since the distribution of motivations to vote might be very different than those to volunteer for a statewide campaign.

Motivations and network structure are thus intricately linked. To understand better how, consider again Figure 1. In the strong class, both networks achieve equally high levels of participation, but the Random one does so more efficiently, with fewer links. In the weak class, this is no longer the case. While the Random network again reaches its peak faster than the Ring, that peak is less than half as high as the Ring's. Now there is a true trade-off: efficiency but little participation, versus relative inefficiency with considerable participation.[20]

The difference between networks is due to the interplay of two fundamental properties. The first, associated with what is called the clustering coefficient, is the degree to which individuals' local networks overlap. In a friendship network, for example, this would mean that one's friends are also friends with each other. These clusters may be thought of as enclaves of participation, since shared experiences encourage a similarity of behavior within them. If one's connections are insufficient to spur one to participate, after all, they are less likely to spur another with very similar connections to participate. Small enclaves are necessary for the initial spread of participatory behavior, as they allow each individual rabble-rouser, by virtue of the way the net external motivation depends on the local participation rate, to have a substantial effect on the actions of those to whom she is tied.[21] Too big an enclave, however, can dilute the rabble-rouser's impact, particularly in the weak and intermediate classes, in which less of the population on average shares the rabble-rouser's motivations. This dilution accounts for the dip in participation seen in the weak class as average connectivity increases. Maximal participation occurs when the chance an added individual will be sufficiently motivated so as to spur participation exactly balances the watering-down effect this extra person's presence has upon one's perception of the local participation rate.

The second property, associated with what is called path length, is the ability of a behavior to spread across a network. Specifically, it is the ability of a behavior to spread out of the enclaves of participation to the wider network. Again there is the trade-off: improving the ability of a network to spread behavior also decreases the effectiveness of enclaves of participation. Spreading a behavior requires ties to individuals outside the enclave, but these dilute the power of the enclave, exposing its members to outside individuals who might not share their motivations.

To see how this trade-off accounts for the differences observed both within and between the plots in Figure 1, note that the Ring has a significant degree of overlap in local networks, and so encourages enclaves of participation. Any behavioral spread, however, must follow its rigid structure; thus it does not support quick spreading of behavior. In contrast, the Random network is in many ways the Ring's opposite. Connections can be made to anyone, and so behavior can spread easily across the Random network. This randomness also means that there is little overlap between local networks, though, and thus it does not encourage small enclaves well. In the weak motivation class encouragement of enclaves is more important than behavioral spread, as it is likely that few of those to whom one is connected will be participating early on, and so it is particularly important to maximize each rabble-rouser's impact. In the strong class, in contrast, early participation is more common, and so small enclaves are less needed to achieve significant participation levels in equilibrium. Accordingly, the Random network's greater ability to spread behavior allows it to do so more efficiently than the Ring, without much worry about diluting enclaves.

Complex Networks and Network Elites

Model

The simple networks of the previous section were helpful in understanding how network structure interacts with the distribution of motivations in determining levels of participation, but have less utility in making empirical predictions as to how participation varies by network type. After all, while one could get at the distribution of motivations with a survey of a sample of the population, one is rarely presented with a perfectly Random network, and perfect Rings are perhaps even more uncommon. One possibility would be to measure the network characteristics that most closely mirror the trade-offs identified in the previous section—clustering coefficients and path lengths—but these require detailed network data that are difficult to acquire, particularly in situations that are inherently risky such as mass protests.[22] Any theory that

[20] This would be of particular importance if making ties were costly.

[21] This is comparable to critical mass in Marwell and Oliver (1993), "bonding" social capital in Putnam (2000), and incubation in Rolfe (2005).

[22] Using a different behavioral model and a Small World network, Fowler identifies the nonlinear impact of average path length on

requires the construction of the entire social network will therefore be of limited use to scholars who only have qualitative network data. This last point is of particular importance, as often even the most careful social analysis of a particular locale can only discern group-level structure (e.g., Petersen 2001).

Thus, in this section I present a typology of networks, chosen not to be an exhaustive characterization of all network configurations, but rather a listing of commonly observed social structures that may be distinguished on qualitative grounds. I keep their description qualitative as well, in order to focus on the differences in outcomes that vary only by network type. This allows scholars with limited network data to draw concrete conclusions about the impact of network structure on their empirical cases. This also provides for general discussions about the power of elites, which are present in some of the networks I examine. For convenience, I summarize major results in the next section. The technical description in the online appendix may prove useful to scholars studying environments in which more complete data exist, as each network is defined by only one or two parameters estimable with incomplete information. Further information about some of the network types used here can be found in existing network surveys, such as Strogatz (2001).

Figure 2 provides a visualization of all four network types: the Small World, the Village (or Clique), the Opinion Leader, and the Hierarchical Network. For simplicity I assume that all ties between individuals in these networks are symmetric—anyone you influence also influences you—and constant throughout each realization of the model. The latter is valid as long as the pace of network formation is slow compared to the rate of behavioral spread. The former is valid for forms of influence that involve reciprocity. This is a reasonable assumption, as many cases of costly action are facilitated by the encouragement of mutual friendship or familial connections (e.g., McAdam 1986; McAdam and Paulsen 1993); one simply has a better idea of the motivations underlying the decisions of those with whom one has been in close contact. Note that symmetric ties do not imply symmetric influence within the network. It is far easier for an opinion leader to affect the behavior of her many followers than for any one follower to influence the leader's behavior.

The Small World network (Watts 1999) is used here to correspond to modern, reasonably dense cities and suburbs, in which there are no exceptional citizens who hold an inordinate amount of sway over their peers. Individuals have substantially overlapping networks, but each also has some chance to influence individuals outside these clusters. These networks form when a tightly regimented series of connections, such as childhood friendships, is perturbed, as when individuals move away and join new groups of friends. One parameter dictates average connectivity, as for a Ring, while a second determines the ease with which influence can spread across the network.

The Village network is similar, but more tightly clustered. It is meant to mimic small towns, villages, and cliques, in which everyone knows everyone else within the social unit, and all exert equal influence on each other. Only the rare person who spans multiple cliques, acting as a "social relay" (Ohlemacher 1996) who possesses "bridging," rather than only "bonding" social capital (Putnam 2000), is able to exert influence outside the unit. One parameter dictates the size of these social units, while a second determines the likelihood of influence outside of them.

In these two networks individuals generally have equal numbers of connections to others. The next two networks model situations in which social elites are present. The first elite network is the Opinion Leader network. Here most people have few connections, while a few—the opinion leaders—have many. A single parameter determines both the number of opinion leaders and the number of connections each has. Simple versions of such networks have also been termed "star" or "wheel" networks (e.g., Gould 1993).

The final network considered is the Hierarchy. While the power of elites in the Opinion Leader network lies in their greater number of network ties, the power of elites within the Hierarchy lies in their privileged placement at its top. Like the one described in Morris (2000), the backbone of the Hierarchy is a series of levels expanding exponentially in width. Individuals are connected to one person above them, and a number of people one level below them equal to the rate of expansion of the hierarchy. For example, if the expansion rate is 3, there is one person at the top, three people on the second level, nine on the third, 27 on the fourth, and so on. Those in the second level are all connected to the person at the top, and each also is connected to three people on the third level. Figure 2 displays a top-down view of the Hierarchy. Because individuals within a given level will often work in a spatially localized area, which is correlated with social closeness (Sandell and Stern 1998), I also add a second parameter, which dictates the likelihood that two people influence one another within a given level.

As shown in the example of the star network in the previous section, it is not just the shape of the network that

voting turnout, but even in this comparatively peaceful setting must concede that "no one knows the true average path length for a typical political discussion network" (2005, 15). This article offers a way other than detailed measurement to get at the relative impacts of real networks.

FIGURE 2 Network Typology (Network Visualizations)

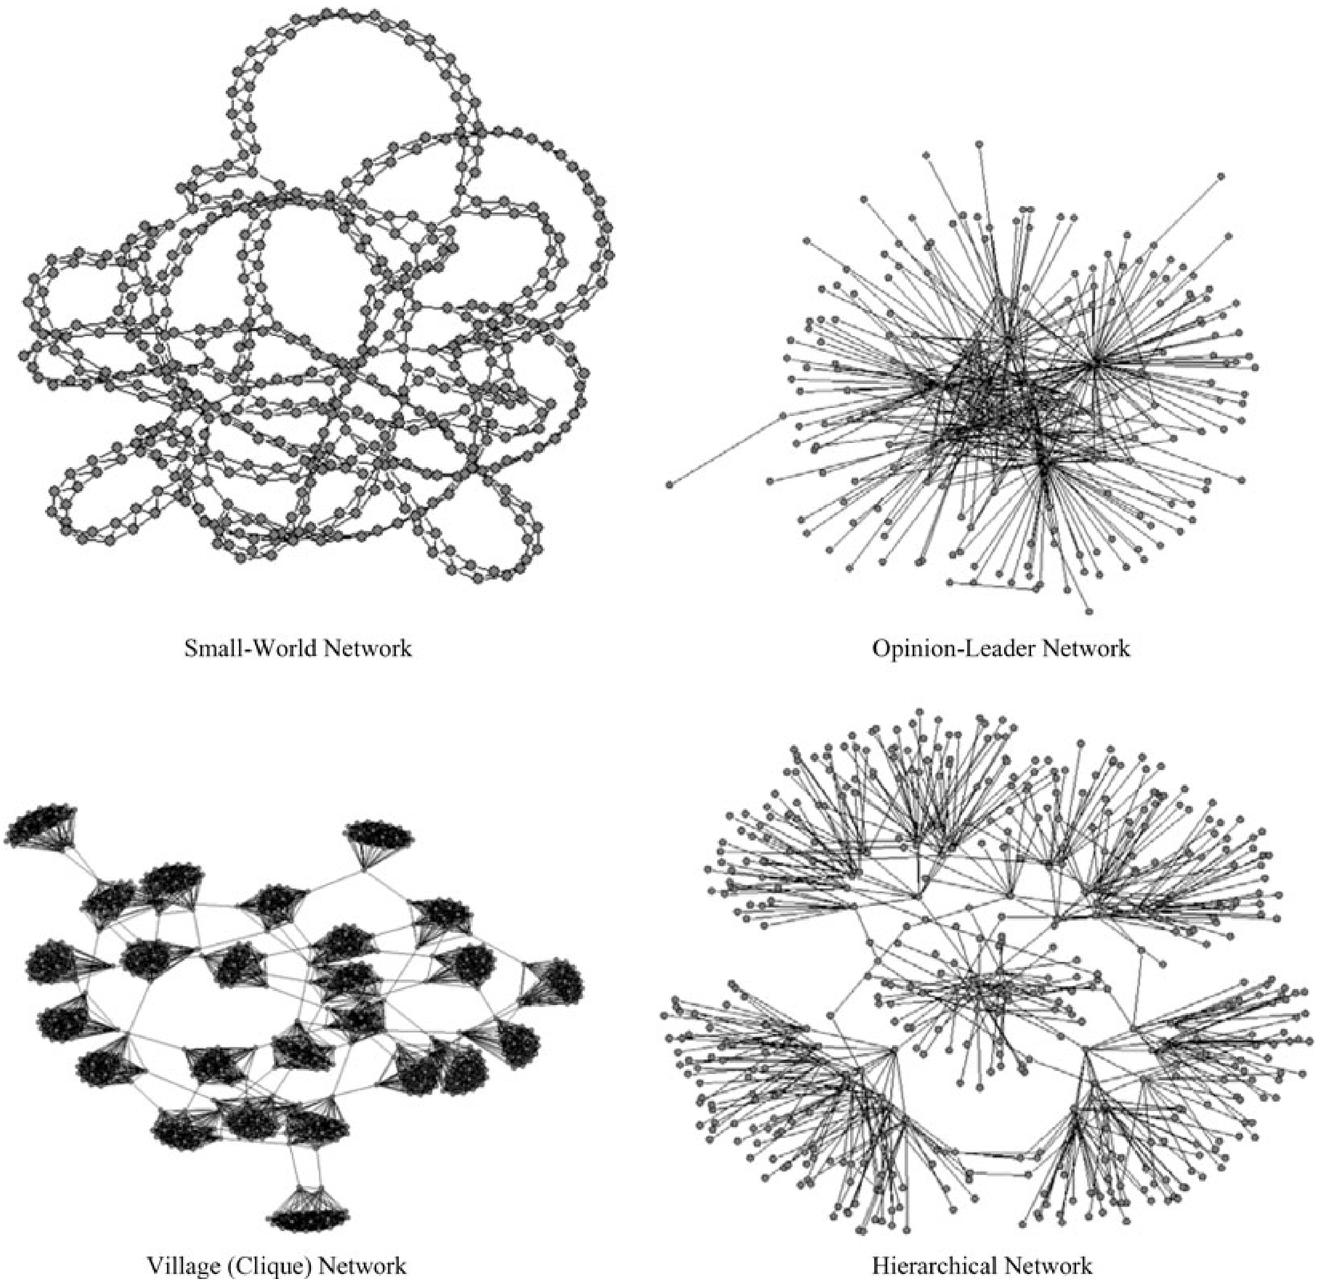

matters: different realizations of the distribution of internal motivations within the network can lead to different outcomes. Thus I also include a parameter that dictates the level of correlation between motivations and position within the network. When individuals have roughly equal influence—as in the Small World and Village networks—I assume that motivations are uncorrelated within position and so distribute individuals' b_i randomly across the network. When individuals have asymmetric influence—as in the Opinion Leader and Hierarchical networks—I also examine correlated motivations, in which elites have either uniformly high or uniformly low internal motivations. In these cases the distribution of motivations is as much a part of network structure as is the arrangement of connections.

Results

To facilitate comparison between the four networks, I will focus on the intermediate motivation class. In general, the

FIGURE 3 **Small World and Village Networks (Top Row: Mean Participation Rate vs. Ease of Behavioral Spread in Small-World Network in Weak and Intermediate Motivation Classes) (Bottom Left: Mean Participation Rate vs. Size of Villages in Hierarchy in Intermediate Motivation Class) (Bottom Right: Mean Participation Rate vs. Intralevel Connectivity in Hierarchy in Intermediate Motivation Class)**

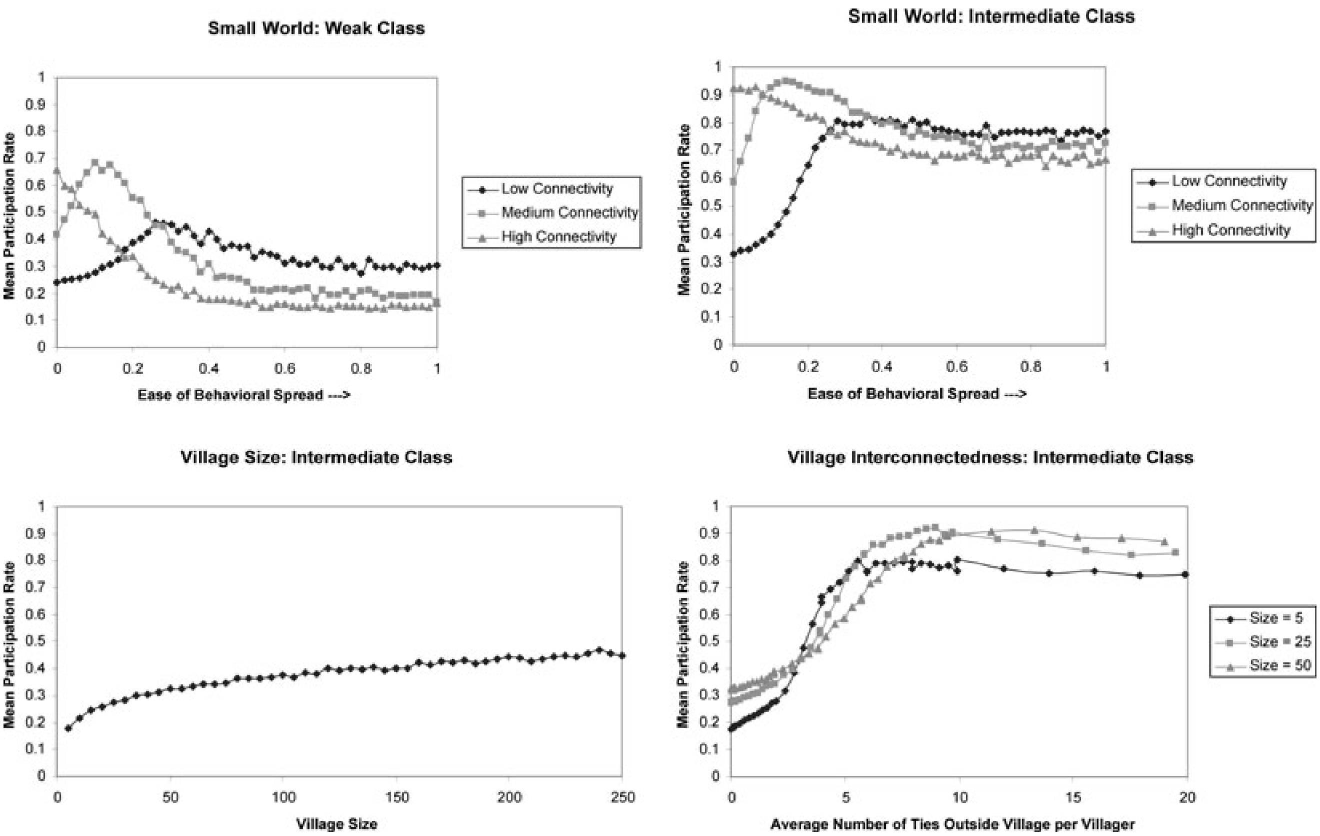

strong class displays increasing participation with increasing average connectivity in all networks, while in the weak class a narrower range of parameter values yields high levels of participation than in the intermediate class. As before the y-axes correspond to the average equilibrium rate of participation. Figure 3 presents results for both the Small World and the Village networks. This figure displays clearly the trade-off between the ease of spreading participation and the nurturing of enclaves of participation discussed in the previous section and additionally allows us to understand the strength of weak ties in this setting.

Consider first the top two plots, which display results for the Small World network. Different lines correspond to different levels of connectivity, while the x-axis is increasing in the ease of behavioral spread throughout the network. For every level of connectivity, there is an optimal level of the ease of behavioral spread that produces maximal participation in equilibrium. When connectivity is not high, this optimal level is greater than zero (which equates to the ease of spread experienced within a Ring), but also less than one (which equates to the ease of spread experienced within a Random network). This is just the trade-off discussed earlier in action: increasing the ease with which a network spreads participation also dilutes its ability to encourage small enclaves. At lower connectivities a Small World network benefits from increased ease in spreading behavior. At higher connectivities, however, the Small World network obtains less benefit from the faster spreading of participation, and the trade-off is no longer beneficial.

Since "weak" ties are the network structure that encourages quick behavioral spread in a Small World network, this trade-off translates directly into a statement about the conditional impact of weak ties. Adding weak ties is likely to have the greatest effect in two contexts: (1) when connecting the population in *any* way leads to more participation, as in the strong motivation class; and (2) when existing network ties are insufficient to spread participation. The first could hold because people are predisposed to participate; in the language of Centola and

Macy (2007), the behavior is a simple contagion. It could also hold because motivations are more widely dispersed in the population. The second could hold because average connectivity is too low, or, more obviously, when there are an insufficient number of weak ties. When either of these conditions obtains, weak ties can be very effective, in some cases more than doubling the rate of participation, but when neither do, adding weak ties can be detrimental, leading to substantial decreases in participation.

The second context can also obtain when network ties are too tightly clustered, so that behavioral spread between clusters is difficult. This is often the case in the Village network, as we can see in the bottom two plots of Figure 3. The leftmost plot displays participation levels at each village size when there are no connections between villages. This arrangement provides for excellent support of enclaves—each village is a perfect enclave—but has no allowance for behavioral spread, leading to low participation levels. In fact, the only reason that participation is increasing in village size is that increasing group size in the intermediate class increases participation here. This same plot in the weak motivation class (not shown) displays a nonmonotonicity in village size.

The plot on the right varies the average number of ties each villager has to individuals outside his village; the more such ties, the easier participation spreads between villages. Since the separated villages on their own are poor at spreading participation, these weak ties between villages are extremely effective in spurring participation, in some cases tripling the level observed. Further, the fundamental trade-off is almost completely resolved in favor of behavioral spread; many weak ties are needed before adding them begins to decrease participation, and this decrease is comparatively small.

Predicting which type of society yields greater participation, rural or urban, thus falls to the number of extra-village connections present. If there are few, urban societies will experience much higher participation levels; if there are many, rural societies can nearly match them. Note, however, that the patterns of participation within each society will be different. In particular, participation within a Village network will be more spatially varied: in some villages there will be significant participation, while in other villages little to none.[23]

Now turn to networks with elites, the focus of Figure 4. There is considerably less overlap of individuals' local networks in both Opinion Leader and Hierarchical networks, and so we would expect weak encouragement of enclaves and a resulting lesser degree of participation in equilibrium, all else equal.[24] This is indeed what we see when motivations are uncorrelated with network position, as in the upper-right plot and the central line of the upper-left plot. Consider first the latter, in which the x-axis is increasing in both the number and the average connectivity of elites. In comparison to the two networks without elites, an Opinion Leader network in which the elites do not have common motivations produces less participation on average. Indeed, such a network is often worse at producing participation than completely separated cliques. The plot displaying results for an uncorrelated Hierarchy tells a similar story; here each line is a different rate of expansion for the width of the hierarchy's levels, and the x-axis is increasing in the likelihood of connections within levels. Without intralevel connectivity the hierarchy does a poor job of spreading participation, again worse than separated cliques, as clustering for the formation of enclaves is minimal and pathways that could lead to behavioral spread are tightly constrained. Increasing connectivity within levels increases participation rates in much the same way as does increasing connections between villages, but to a lesser extent. The message here is straightforward: elites have power within the network by virtue of substantially dictating the rate of behavioral spread, but if they do not share common motivations this power is wasted, as the influence of more-motivated elites is mitigated by that of less-motivated elites.

What if they do have common motivations? The bottom-left plot and the bottom line in the top-left plot display equilibrium participation levels when motivations are negatively correlated with network position. Elites within these networks have little independent reasons to participate and substantially control the spread of behavior by virtue of their network positions. Consequently they severely dampen participation. This effect is particularly strong in the Opinion Leader network, where each elite can more directly influence many individuals. Yet the power of elites is not absolute—it can be mitigated almost entirely by the structure of the network. If the rate of expansion of the hierarchy is neither too fast nor too slow, the people at the bottom of the network—the proletariat, if you will—can counteract this elite influence if they have enough connections among themselves. The key here is to obtain a sufficiently large and

[23]It is also important to note that this is a ceteris paribus condition. A Village network with high average connectivity or containing individuals in the strong motivation class may very well yield considerably higher participation levels than a Small World network that differs on either dimension.

[24]Despite significant differences in the underlying models, results for the Opinion Leader network share many features with those found under Gould's (1993) star-shaped network. Such similarities suggest strongly that the results given in this article are not mere artifacts of the underlying behavioral model, but instead are fundamental properties of behavior within networks in general.

FIGURE 4 Opinion Leader and Hierarchical Networks (Top Left: Mean Participation Rate vs. Frequency of Elites in Opinion Leader Network in Intermediate Motivation Class) (Top Right and Bottom Row: Mean Participation Rate vs. Intralevel Connectivity in Hierarchy in Intermediate Motivation Class When Motivations Are Uncorrelated, Negatively Correlated, and Positively Correlated with Network Position)

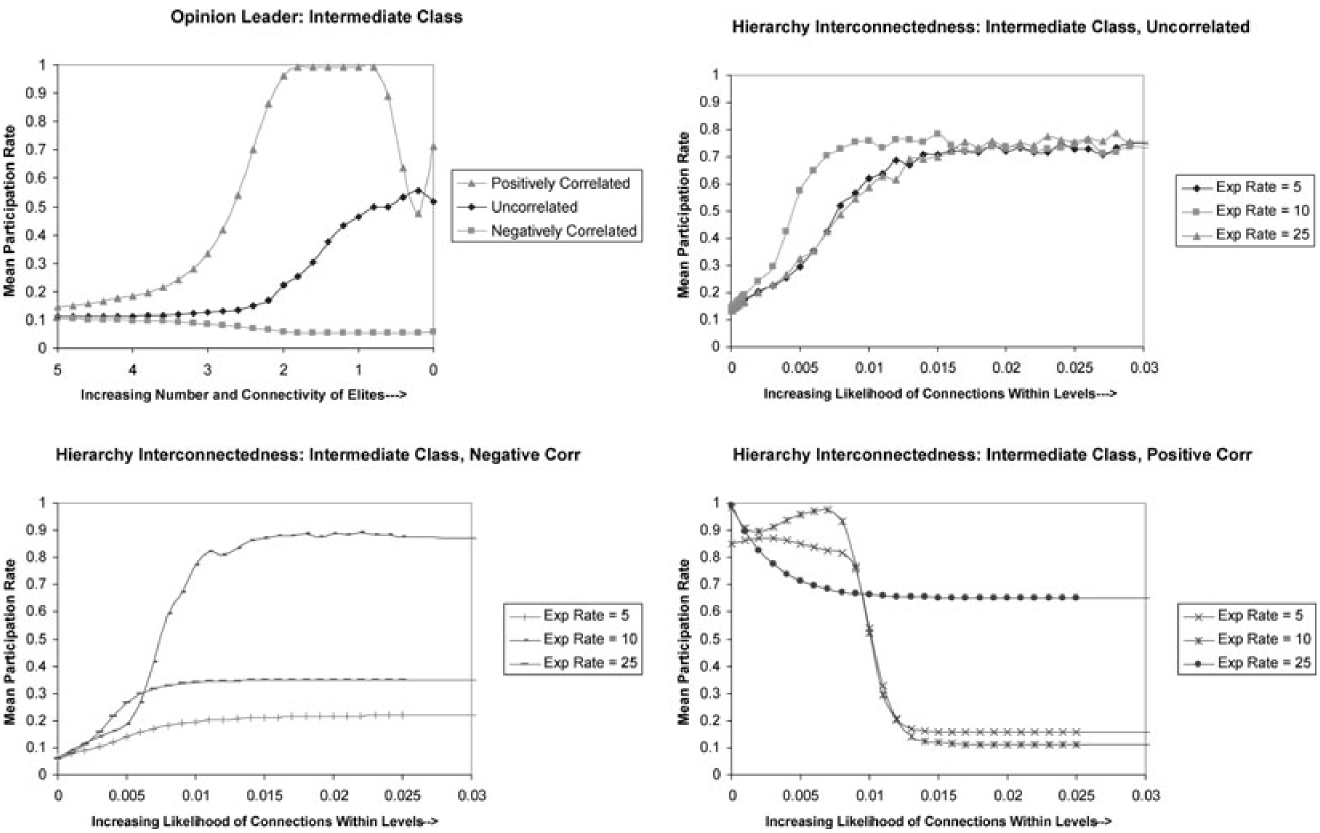

well-connected group of people at the bottom of the hierarchy who, due to the negative correlation of motivations with network position, are highly internally motivated to participate. If these requirements are achieved, the bottom of the hierarchy can spur the network on to very high levels of participation.

This same effect can be seen in reverse in the positively correlated case, displayed in the bottom-right plot and in the top line in the top-left plot. As long as elites have uniformly high internal motivations and a unique position within the network, their presence encourages near-total participation across a substantial range of network parameters. Diminish their power by making their position less unique, however, and their impact on participation falls rapidly. In the Opinion Leader network this occurs at the extremes, when there are either too few elites with too few connections, or too many elites, some of whom now are not quite so motivated. In the Hierarchy, this occurs when the proletariat gains too many interconnections and effectively forms its own power base. When this occurs participation levels fall for all but the widest hierarchies, which diminish this effect by having more elites in direct contact with those at the bottom.

Thus, the power of elites is conditional on both the structure of the network and the distribution of motivations within it. If we were to assume that highly motivated elites *want* others to participate as well, perhaps because they receive the lion's share of the benefits, then there is a moral here for elites. Direct control over the flow of influence, as in an Opinion Leader network, is optimal, as long as you can be sure that elites have uniform preferences. If such control is not available then a Hierarchy is often effective, but with a caveat. If we assume that connections within levels are not determined by elite choice, then great care must be taken in setting up the skeleton of the Hierarchy (as determined by its expansion rate), lest rabble-rousers spur participation you do not want, or wet blankets smother participation that you do want. This moral is of particular substantive relevance for entrenched government bureaucracies seeking to change norms, as the appointed upper echelon is often at odds with the career bureaucrats, mimicking the case of correlated benefits seen here. This also helps explain why some organizations (e.g., the army) adapt quickly to some imposed changes, while the exact same organizations adapt slowly to others. The form of the Hierarchy dictates whether the preferences of the elites matter most for behavioral spread, or if those of lower-ranked individuals hold sway.

Summary of Major Hypotheses

We have seen that some commonly used metrics for the importance of networks—size, the prevalence of weak ties, the presence of elites—have a more complex interaction with network structure and with the distribution of motivations in the population than has been previously acknowledged, leading in some cases to the potential for selection bias. However, by focusing on large-scale factors like network type and motivation class, it has been possible nevertheless to make predictions about expected participation levels that take these complexities into account. The list below summarizes the most prominent of these for the four empirically relevant networks of the typology, with pointers to the page(s) in which the justification for the prediction can be found. (For definitions of motivation classes and network types, see pages 127–28 and 131–32, respectively.)

- **Small World**: This network efficiently induces high levels of participation, which spreads quickly via a combination of strong and weak ties. In the strong motivation class, increasing ties of any form increases participation. In the weak and intermediate classes, increasing weak ties increases participation only when strong ties are not prevalent, and only to a point. The more strong ties the network has, the more adding weak ties decreases participation (pp. 133–34).
- **Village**: Similar to the Small World network, except as noted here. Behavior spreads first within and then between villages, leading to less efficiency and slightly less participation. Aggregate participation is dependent on the weak ties between villages, which are more important in prediction than is the size (number of strong ties) of each village. Weak ties more often encourage participation than in the Small World (p. 134).
- **Opinion Leader**: The number of elites (who have many connections, by definition), and the degree of elite conformity in motivations are more relevant in predicting participation than weak ties, network size, and even motivation class here. Up to a point, increasing the number of elites tends to increase participation. Behavior spreads outward from motivated elites to followers. When elites have uniformly low motivations, there is little participation; when their motivations are uniformly high, participation is near total. Between these extremes Opinion Leader networks generally permit lower levels of participation than Small World networks (pp. 134–35).
- **Hierarchy**: The relative importance of factors here mirrors that in the Opinion Leader network, with one exception: ties between people in the same level (which are generally "weak" here) can alter outcomes when elite motivations are uniform, for some widths of the hierarchy. When elites have uniformly low motivations, highly interconnected followers can produce in some cases significant levels of participation anyway (the "proletariat" revolt). When elites have uniformly high motivations, highly interconnected followers can in some cases reduce the level of expected participation to very low values (pp. 134–36).

Discussion

Though recent research in political science has begun to take greater account of the importance of social networks in determining participation (e.g., Huckfeldt 2001; Leighley 1990; McClurg 2006), the role of network structure has thus far been comparatively little examined (Fowler 2005; Rolfe 2005). Via a typology of network structures, defined by qualitative traits and requiring little data to discern empirically, the model analyzed in this article addresses this gap in the literature.

In particular, the model illustrates the conditional nature of network effects. While empirically important factors such as network size may strongly encourage participation under some distributions of individual motivations, under others it might have a neutral or even discouraging effect. Further, even when increasing average network size might initially produce additional participation, adding too many ties—particularly weak ties (Centola and Macy 2007)—may begin to discourage participation. This is due to a fundamental trade-off between encouraging participation via tightly connected enclaves, which possess "bonding" social capital, and spreading participation to other individuals via weak ties, which possess "bridging" social capital (Putnam 2000). Both are necessary to achieve significant levels of participation, but which has the more positive effect depends on the type of network in place, and the distribution of motivations in the population. This conditionality calls into question policies that focus on building social capital irrespective of conditions, such as USAID's Iraq Community Action Program.[25] Adding network ties when motivations are low, or when the true network type is a hierarchy with a motivated leadership, can lead to unwanted decreases in political participation.

A strength of the model is its ability to predict participation levels, which can help guide social-capital-based policies. These predictions are summarized in the previous section. Further, the model's analysis raises several more general empirical issues. The conditional nature of network structure implies the risk of selection bias. If one's empirical analysis is limited to a subset of (successful) collective actions or methods of political participation, one cannot be sure that one's results are due to network effects, or to less costly actions or more motivated populations. Analyses of social influence have a similar problem—the impact of highly connected individuals must necessarily be viewed in light of other elites' motivations and the larger structure of the network, and cannot be assumed simply from one's number of connections. Even highly influential elites can find their effectiveness plummet when the structure of the network and the interests of their followers collude to create a contrary power base.

References

Berry, William D., and Brady Baybeck. 2005. "Using Geographic Information Systems to Study Interstate Competition." *American Political Science Review* 99: 505–19.

Borgatti, Stephen P., and Martin G. Everett. 1992. "Notions of Position in Social Network Analysis." *Sociological Methodology* 22: 1–35.

Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113(3): 702–34.

Chong, Dennis. 1991. *Collective Action and the Civil Rights Movement.* Chicago: University of Chicago Press.

Chwe, Michael S. Y. 1999. "Social Mechanisms: An Analytical Approach to Social Theory." *Journal of Economic Literature* 37(2): 668–69.

Fowler, James H. 2005. "Turnout in a Small World." In *Social Logic of Politics*, ed. Alan Zuckerman. Philadelphia: Temple University Press, 269–87.

Fowler, James H. 2006. "Connecting the Congress: A Study of Cosponsorship Networks." *Political Analysis* 14(4): 456–87.

[25]URL: http://www.usaid.gov/iraq/accomplishments/cap.html; last accessed May 22, 2007.

Friedkin, Noah E., and Eugene C. Johnsen. 1999. "Social Influence Networks and Opinion Change." *Advances in Group Processes* 16: 1–29.

Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large-scale Field Experiment." *American Political Science Review* 102(1): 33–48.

Gleditsch, Kristian S. 2002. *All International Politics Is Local: The Diffusion of Conflict, Integration, and Democratization.* Ann Arbor: University of Michigan Press.

Gould, Roger V. 1991. "Multiple Networks and Mobilization in the Paris Commune, 1871." *American Sociological Review* 56(6): 716–29.

Gould, Roger V. 1993. "Collective Action and Network Structure." *American Sociological Review* 58(2): 182–96.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78(6): 1360–80.

Granovetter, Mark S. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83(6): 1420–43.

Granovetter, Mark S., and Roland Soong. 1983. "Threshold Models of Diffusion and Collective Behavior." *Journal of Mathematical Sociology* 9(3): 165–79.

Granovetter, Mark S., and Roland Soong. 1986. "Threshold Models of Interpersonal Effects in Consumer Demand." *Journal of Economic Behavior Organization* 7(1): 83–99.

Heckathorn, Douglas D. 1996. "The Dynamics and Dilemmas of Collective Action." *American Sociological Review* 61(2): 250–77.

Hedström, Peter. 1994. "Contagious Collectivities—On the Spatial Diffusion of Swedish Trade-Unions, 1890–1940." *American Journal of Sociology* 99(5): 1157–79.

Huckfeldt, Robert. 2001. "The Social Communication of Political Expertise." *American Journal of Political Science* 45(2): 425–39.

Huckfeldt, Robert R., and John Sprague. 1995. *Citizens, Politics and Social Communication: Information and Influence in an Election Campaign.* Cambridge: Cambridge University Press.

Kearns, Michael, Siddharth Suri, and Nick Montfort. 2006. "An Experimental Study of the Coloring Problem on Human Subject Networks." *Science* 313(5788): 824–27.

Kenny, Christopher B. 1992. "Political Participation and Effects from the Social Environment." *American Journal of Political Science* 36(1): 259–67.

Kim, Hyojoung J., and Peter S. Bearman. 1997. "The Structure and Dynamics of Movement Participation." *American Sociological Review* 62(1): 70–93.

Klofstad, Casey A. 2007. "Talk Leads to Recruitment: How Discussions about Politics and Current Events Increase Civic Participation." *Political Research Quarterly* 60(2): 180–91.

Kuran, Timur. 1991. "Now Out of Never—The Element of Surprise in the East European Revolution of 1989." *World Politics* 44(1): 7–48.

Kuran, Timur. 1995. *Private Truths, Public Lies : The Social Consequences of Preference Falsification.* Cambridge, MA: Harvard University Press.

Lake, Ronald La Due, and Robert Huckfeldt. 1998. "Social Capital, Social Networks, and Political Participation." *Political Psychology* 19(3): 567–84.

Leighley, Jan E. 1990. "Social Interaction and Contextual Influences on Political Participation." *American Politics Research* 18(4): 459–75.

Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades—The Monday Demonstrations in Leipzig, East Germany, 1989–91." *World Politics* 47(1): 42–101.

Macy, Michael W. 1991a. "Chains of Cooperation—Threshold Effects in Collective Action." *American Sociological Review* 56(6): 730–47.

Macy, Michael W. 1991b. "Learning to Cooperate—Stochastic and Tacit Collusion in Social-Exchange." *American Journal of Sociology* 97(3): 808–43.

Marwell, Gerald, and Pamela E. Oliver. 1993. *The Critical Mass in Collective Action: A Micro-Social Theory*. Studies in Rationality and Social Change. Cambridge: Cambridge University Press.

Marwell, Gerald, Pamela E. Oliver, and Ralph Prahl. 1988. "Social Networks and Collective Action—A Theory of the Critical Mass." *American Journal of Sociology* 94(3): 502–34.

McAdam, Doug. 1986. "Recruitment to High-Risk Activism—The Case of Freedom Summer." *American Journal of Sociology* 92(1): 64–90.

McAdam, Doug, and Ronnelle Paulsen. 1993. "Specifying the Relationship between Social Ties and Activism." *American Journal of Sociology* 99(3): 640–67.

McClurg, Scott D. 2003. "Social Networks and Political Participation: The Role of Social Interaction in Explaining Political Participation." *Political Research Quarterly* 56(4): 449–64.

McClurg, Scott D. 2004. "Indirect Mobilization: The Social Consequences of Party Contacts in an Election Campaign." *American Politics Research* 32(4): 406–43.

McClurg, Scott D. 2006. "The Electoral Relevance of Political Talk: Examining Disagreement and Expertise Effects in Social Networks on Political Participation." *American Journal of Political Science* 50(3): 737–54.

Morris, Stephen. 2000. "Contagion." *Review of Economic Studies* 67(1): 57–78.

Mutz, Diana C. 2002. "The Consequences of Cross-Cutting Networks for Political Participation." *American Journal of Political Science* 46(4): 838–55.

Neill, Daniel B. 2005. "Cascade Effects in Heterogeneous Populations." *Rationality and Society* 17(2): 191–241.

Nickerson, David W. 2008. "Is Voting Contagious? Evidence from Two Field Experiments." *American Political Science Review* 102(1): 49–57.

Ohlemacher, Thomas. 1996. "Bridging People and Protest: Social Relays of Protest Groups against Low-Flying Military Jets in West Germany." *Social Problems* 43(2): 197–218.

Ohtsuki, Hisashi, Christoph Hauert, Erez Lieberman, and Martin A. Nowak. 2006. "A Simple Rule for the Evolution of Cooperation on Graphs and Social Networks." *Nature* 441(7092): 502–5.

Oliver, Pamela E. 1993. "Formal Models of Collective Action." *Annual Review of Sociology* 19: 271–300.

Oliver, Pamela E., and Gerald Marwell. 1988. "The Paradox of Group-Size in Collective Action—A Theory of the Critical Mass 2." *American Sociological Review* 53(1): 1–8.

Oliver, Pamela E., and Daniel J. Myers. 2003. "Social Movements and Networks: Relational Approaches to Collective Action." In *Networks, Diffusion, and Cycles of Collective Action*, ed. Mario Diani and Doug McAdam. New York: Oxford University Press, 173–203.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups.* Cambridge, MA: Harvard University Press.

Opp, Karl-Dieter, and Christiane Gern. 1993. "Dissident Groups, Personal Networks, and Spontaneous Cooperation—The East-German Revolution of 1989." *American Sociological Review* 58(5): 659–80.

Petersen, Roger Dale. 2001. *Resistance and Rebellion: Lessons from Eastern Europe.* Cambridge, MA: Cambridge University Press.

Putnam, Robert D. 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon Schuster.

Rolfe, Meredith. 2005. "Conditional Choice: Social Networks and Decision-Making." Presented at the 2005 annual meeting of the Midwest Political Science Association.

Sageman, Mark. 2004. *Understanding Terror Networks.* Philadelphia: University of Pennsylvania Press.

Sandell, Rickard, and Charlotta Stern. 1998. "Group Size and the Logic of Collective Action: A Network Analysis of a Swedish Temperance Movement 1896–1937." *Rationality and Society* 10(3): 327–45.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior.* New York: Norton.

Scholz, John T., and Cheng-Lung Wang. 2006. "Cooptation or Transformation? Local Policy Networks and Federal Regulatory Enforcement." *American Journal of Political Science* 50(1): 81–97.

Siegel, David A. 2008. "When Does Repression Work? Collective Action in Social Networks." Working Paper. Florida State University.

Snow, David A., Louis A. Zurcher, and Sheldon Ekland-Olson. 1980. "Social Networks and Social-Movements—A Microstructural Approach to Differential Recruitment." *American Sociological Review* 45(5): 787–801.

Strang, David, and Sarah A. Soule. 1998. "Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills." *Annual Review of Sociology* 24: 265–90.

Strogatz, Steven H. 2001. "Exploring Complex Networks." *Nature* 410(6825): 268–76.

Tilly, Charles. 1978. *From Mobilization to Revolution.* Reading, MA: Addison-Wesley.

Verba, Sidney, Kay L. Schlozman, and Henry E. Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics.* Cambridge, MA: Harvard University Press.

Watts, Duncan J. 1999. *Small Worlds: The Dynamics of Networks between Order and Randomness.* Princeton, NJ: Princeton University Press.

Yin, Chien-Chung. 1998. "Equilibria of Collective Action in Different Distributions of Protest Thresholds." *Public Choice* 97(4): 535–67.