When Does Repression Work? Collective Action in Social Networks

David A. Siegel Florida State University

Empirical studies reach conflicting conclusions about the effect of repression on collective action. Extant theories cannot explain this variation in the efficacy of repression, in part because they do not account for the way in which social networks condition how individual behavior is aggregated into population levels of participation. Using a model in which the population is heterogeneous in interests and social influence, I demonstrate that the extent to which repression reduces participation, and the extent to which an angry backlash against repression increases participation, depends critically on the structure of the social network in place; this implies the need for greater empirical attention to network structure. To facilitate the model's empirical application, I focus on broad qualitative network types that require comparatively little data to identify and provide heuristics for how one might use qualitative network data to derive quantitative hypotheses on expected aggregate participation levels.

Repression is the process by which powerful actors attempt to deter a population from participating in a collective action that threatens them, such as protest, dissent, or rebellion. It is perhaps most commonly thought of in the context of state repression, taking the form of "disappearances" and imprisonment, poll taxes and water cannon. Yet substate actors can repress as well, as when insurgents target state collaborators, or terrorist groups target civilians. Further, repression need not be immoral: party whips, union bosses, and states facing terrorist threats all have occasion to function as repressive entities.

To limit repression when it is immoral, and to use it more effectively when it is not, one must understand when and how it works. However, even in the context in which repression has been most frequently explored—state repression—the literature provides no consensus on its functioning. Indeed, studies indicate that repression can increase, decrease, or have no effect on levels of dissent (e.g., Francisco 1996, 2004; Gupta, Singh, and Sprague 1993; Gurr and Moore 1997; Hibbs 1973; Lichbach and Gurr 1981; Muller 1985; Rasler 1996), depending on context (Davenport 2007). To address this indeterminacy, I present the first model incorporating social context that predicts how levels of participation in a collective action respond to repression.[1] Mobilization in the model occurs via interpersonal interactions mediated by social networks. Effective repression, ineffective repression, and backlash are all observed, and the model's explicit statements of the contextual and behavioral factors leading to these outcomes begin to answer the question of "When does repression work?"

Consideration of social networks makes the analysis possible. A robust literature supports the notion that individuals do not make complex and potentially dangerous decisions independently of their fellows; their decisions depend on considerations of safety, fairness, reputation, information, and influence that are fundamentally related to the actions of others (e.g., McAdam 1986; McAdam and Paulsen 1993; McClurg 2006; Ohlemacher 1996; Opp and Gern 1993; Snow, Zurcher, and Eklandolson 1980). Specifically, this literature indicates that people care most about the actions of those close to them, those within their social networks. This may be apparent in the domains of fairness, reputation, and influence, but even things like information or relative safety are affected by networks. For example, many people might be dissenting, but how many are planted by the state (Petersen 2001)? One can only rule out those one knows and trusts, and these are by definition the individuals in one's network.

[1]The appendix may be found online at both the author's website (http://myweb.fsu.edu/dsiegel/) and www.journals.cambridge.org/jop.

The Journal of Politics, Vol. 73, No. 4, October 2011, Pp. 993–1010 doi:10.1017/S0022381611000727

ISSN 0022-3816

Extant scholarship on repression has recognized that individuals respond to others' actions when considering whether to act themselves, but has generally considered these others in the aggregate only (e.g., Francisco 1995; Kuran 1991; Myers and Oliver 2005; Oliver and Myers 2002; Rasler 1996).[2] As such, it fails to account for theoretical evidence that the structure of social interactions—whose behavior one considers—will matter in determining the level of participation one might expect in a collective action, even absent repression (e.g., Centola and Macy 2007; Chwe 1999; Fowler 2005; Gould 1993; Kim and Bearman 1997; Marwell and Oliver 1993; Rolfe 2005; Siegel 2009). One explanation for this omission is the strong data requirements in much social network analysis (e.g., Fowler et al. 2007). Even careful social analysis might only yield group-level structure (e.g., Petersen 2001), particularly when repression is present. Looking at the behavior of movement leaders only (e.g., Lichbach 1987) or treating individuals' decisions to mobilize as independent and divorced from the likelihood of success of the cause (DeNardo 1985, 56–57) avoids this issue, allowing direct comparisons between different tactics of rational movement leaders. However, these approaches fail to tackle the collective action problem of why people follow their leadership and rise up, despite the clear risks and uncertain benefits.

The model I present in the second section considers both individual-level and society-wide factors in assessing the impact of repression on participation. Indeed, I find that the centrality of the interaction between social structure and individual motivations in predicting participation makes it *necessary* to include both factors if one desires to understand repression's effect. In addition, though some answers to the question of "When does repression work?" rely upon fine details of social networks, there are also generalities across broad network types; these types require less data to identify. Relevance in both data-rich and data-poor regimes greatly increases the scope of the model's substantive applicability.

[2]All four elements of Lichbach's ontology of solutions to the collective action problem of a rebel involve social conceptions of order and "relationships among rebels" (1995, 20), but social networks are not explicitly modeled. Oliver (2002) suggests the importance of networks in understanding repression, but explores only aggregate behavior in her models. DeNardo considers mobilization through ideological and organizational recruiting. The former has individuals acting independently, while the latter notes the importance of "incentives like bonds of friendship" (1985, 45) but offers no explicit model of these bonds.

To understand better its causal logic, I build the model in two stages in the second section and provide corresponding results in the subsections of the third section. The first stage explores behavior within social networks under repression. Individuals have heterogeneous motivations to participate in some collective action and respond to the behavior of others within their local social networks. As Kuran (1991) assumes in the aggregate, the more people participating in one's local network, the more likely one is to participate. I use a network typology drawn from Siegel (2009), consisting of four social networks defined by qualitative characteristics: Small World, Village/Clique, Opinion Leader, and Hierarchical networks. Each is meant to mirror a common large-scale social structure; together they form a typology that, while not exhaustive, covers a wide array of cases. Small World networks stand in for suburban/urban societies, in which information flows quickly and connections are made widely. Village/Clique networks mirror rural/less well-developed societies or tightly insulated cliques, in which most people respond to their immediate (largely nonoverlapping) social circles. Opinion Leader networks possess social leaders—those with great social influence—who drive behavior directly via their numerous connections to followers. Hierarchical networks transmit leaders' influence downward through deputies to the followers at the bottom.

As individuals respond to each others' actions they are also exposed to repression, which involves the removal of participants from the network. The model allows for variation in the strength of repression along a continuous scale. It also provides a dichotomous representation of repressive technology: targeted or random. Targeted repression falls first upon participants with greatest social influence, while random repression afflicts all with equal likelihood.

I find that both the strength and technology of repression alter expected levels of participation, but conditionally on the network in place. For example, the model makes the novel prediction that a society whose leaders are unified in their desire for mass participation will be highly resistant even to repression targeted directly at these leaders. In contrast, the same repressive technology can swiftly crush participation in Opinion Leader and Hierarchical networks that lack unified leaders. This and other predictions derived from characteristics of the model's broad network types are relevant even in situations where data are sparse and/or largely qualitative in nature. The model can thus be used in a predictive fashion, not only as a tool in understanding what has happened once more detailed data have been acquired.

In the second stage of the model, because the violent repression of a loved one is likely to elicit an additional emotional response, I introduce anger and fear to the model. These act as, respectively, incentives or disincentives to participate that arise in response to the removal of another within one's social network. They illustrate conditions under which backlash might be dependent on *who* was subjected to repression, as observed in Kaplan et al. (2005). Further, the way in which social networks can multiply individual anger, turning it into aggregate backlash, provides a strong practical reason to avoid violent repression.

In the fourth section I outline a procedure for applying the model to substantive cases, providing heuristics for the use of qualitative network data to derive quantitative hypotheses on expected aggregate participation levels. This section also serves as a brief review of major results. I then conclude. An online appendix provides further technical details, as well as a concrete illustration of the model's empirical application in the case of voting in the January 2005 Iraqi Legislative elections.

Model

The model is broken up into two parts exploring: (1) behavior in social networks under repression and (2) psychological responses to repression. In each case every individual in a population must make a series of decisions as to whether or not to participate in some collective action.[3] As the model is not analytically tractable, I utilize computer simulation to derive results. A third subsection briefly discusses the simulation and analysis methodology. A precise specification sufficient to replicate the model may be found in the online appendix.

For clarity of exposition, in what follows I focus largely on a single dependent variable: the maximal rate of participation achieved in a population. For behavior that happens a single time and then is done—voting, for instance—this captures the total number of people voting. For behavior that takes place over time, with individuals stopping and starting again—protest, for instance—this illustrates how "hot" the protest got. Due to the latter interpretation in particular, this dependent variable has strong substantive import.

Behavior in Social Networks under Repression

The model's actors comprise a finite population of N individuals.[4] Each person's motivations toward participation in the model are separated into two disjoint components. The first, termed *net internal motivation*, encompasses all factors relating to one's desire to participate in some collective action that do not depend on the participation of others. Examples of these factors include moral certainty in the cause, general disaffectedness with society, or the opportunity cost derived from missing work. I assume that there are many such factors that affect individuals' decision making within the population and that each of these factors is distributed across the population in some unknown fashion. Acting on the suggestion of a central limit theorem, I assume that the net of all these factors will be drawn from a normal distribution.[5] I call each person i's net internal motivation b_i, the mean of the normal distribution from which they are drawn b_{mean}, and its standard deviation b_{stdev}.

The second component of motivations is one's *net external motivation*, which I denote $c_{i,t}$ for each individual i, at each time t. It covers all factors relating to one's desire to participate that depend on the participation of others. I assume that $c_{i,t}$ is increasing in the observed participation of others within one's local social network; hence it captures network effects.[6] Note that, because internal and

[3]As specified here, the decision is a binary one: participate or "stay home." However, this particular dichotomy is not necessary. As long as there are two actions of which only one is subject to repression, the model applies. For example, individuals could be choosing between violent and non-violent participation (DeNardo 1985; Lichbach 1987) and only one tactic might be repressed because, for instance, there might be a norm against repressing non-violent protest, or the state might lack the capability to repress violent action such as terrorism effectively. In either scenario the model would address the relative spread of each tactic.

[4]The model in this subsection draws from Siegel (2009), particularly in network typology and choice of representative parameter values. Figure 1 uses the same parameters as does Figure 2 of that article.

[5]As discussed in the appendix, violations of this logic affect only the normality assumption, not the basic idea upon which this model rests—that each individual has some net internal motivation.

[6]Networks are defined as the set of people who influence one's decisions in this fashion. Though others may be visually observed, they could, for example, be government plants (Petersen 2001). The logic of the model holds as well when individuals respond only to others' influence, and do not directly observe participation, as long as those doing the influencing participate themselves when the time comes.

external motivations are assumed disjoint, *any* change arising from the actions of others in the network must *only* alter external motivations; only $c_{i,t}$ responds to the behavior of others in the network. However, as detailed below, internal motivations can change as well via psychological responses.

Putting these two motivations together yields a simple decision rule: An individual i participates at a given time t if and only if $b_i + c_{i,t} > 0$, i.e., if and only if her net motivation to participate is positive. Since the left-hand side of this inequality is increasing in others' participation, this rule implies that the more people who participate, the more one wants to do so as well. This may be due, for example, to increased safety-in-numbers (e.g., Kuran 1991), to increasing shame from violating norms of fairness by shirking (e.g., Gould 1993), or to the increasing possibility of ingroup punishment if one were to shirk.[7]

In each period of the model, all individuals in the population decide whether or not to participate based on the above decision rule. Between periods, they update their net external motivations according to information about the participation of others within their local networks during the preceding period. The model assumes that they utilize the linear updating rule $c_{i,t+1} = \lambda c_{i,t} - (1-\lambda)\left(1 - lpr_{i,t}\right)$, where $lpr_{i,t} \in [0,1]$ is the local participation rate for individual i at time t.[8] New external motivations are thus functions of both old external motivations and the present social context and are increasing in the proportion of participants in one's social network (i.e., $lpr_{i,t}$). The weight $\lambda \in [0,1]$ dictates the degree to which individuals use new participation information in their decisions, responding to their fellows' actions. Higher values indicate less responsiveness to local participation levels.

To avoid hard-wiring participation into the model, I initially set external motivations at their minimum, given the above rule: $c_{i,0} = -1$. This implies that $c_{i,t}$ increases from -1 to a maximum of 0 as $lpr_{i,t}$ increases from 0 to a maximum of 1. As the b_i are unbounded, there will be rabble-rousing types (Granovetter 1978) with $b_i > 1$ who always participate regardless of their fellows; all immediate participation arises from these individuals. There will also be "wet-blankets" with $b_i \leq 0$ who will never participate under any circumstances—for example, those who derive a great deal of power, prestige, or money from the status quo. Depending on the distribution of internal motivations and the structure of the network, sometimes the updating dynamic in this model leads to a cascade, where substantial aggregate participation levels are achieved, while other times there is no cascade and participation levels remain low.

In line with Siegel (2009), social networks are represented by a typology of qualitative network structures that mirror commonly observed empirical networks. Since little data are necessary to discern network type, scholars with limited network information can use hypotheses from the model that relate only to network type to draw specific conclusions about the impact of network structure on their empirical cases. Figure 1 provides a visualization of all four network types: the Small World, the Village (or Clique), the Opinion Leader, and the Hierarchical Network.[9] All ties are assumed symmetric: anyone you influence also influences you.[10] Network structure is constant apart from changes directly resulting from the removal of individuals. Unless noted otherwise below, I assume that one's net internal motivation is uncorrelated with one's location in the network.

The Small World network (Watts 1999) is used here to correspond to life in cities and suburbs. Individuals' local networks overlap significantly, so that people connected to each other are likely to share other connections as well. Yet each person also may have a number of "weak ties" (Granovetter 1973), defined here as connections that link socially distant people. This gives individuals a chance to influence people outside their own clusters, allowing for the swift spreading of information across the network. These networks can form when tight-knit groups of friends or family disperse upon, for example, transition to college or forced migration. One parameter dictates average connectivity—the number of people to whom one is connected—while a second determines the frequency of weak ties.

[7]This rule also implies that the model applies less well to scenarios such as anonymous public goods provision in which social pressure or influence is less likely to overcome the free rider problem. Formally, the core model does not allow for others' participation to reduce one's motivation to participate directly, though the fear response considered below does allow this to occur indirectly after repression.

[8]While only one possible rule of many, linear updating has the benefit of providing the necessary dependence on $lpr_{i,t}$ without requiring much cognitively of individuals.

[9]Though I focus on these networks' qualitative properties and suppress the one or two parameters that define each network, these may be estimated with more complete data. All parameters are described in the appendix.

[10]Note that symmetric ties do not imply symmetric influence within the network. An opinion leader has far more influence on her followers than any follower has on her.

FIGURE 1 Network Typology

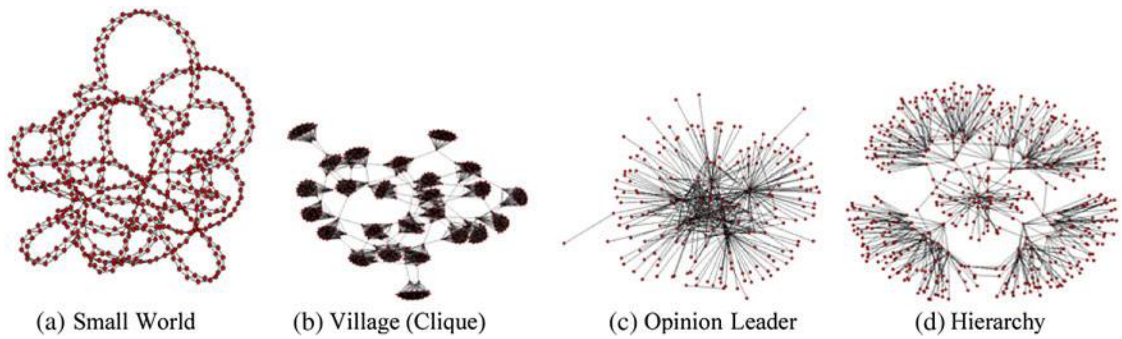

The Village network is meant to mimic villages, small towns, and cliques, in which everyone knows everyone else within the social unit. In addition, individuals may have weak ties to people in other villages; local elders who regularly meet in multivillage conclaves might possess these. I assume for these networks that such individuals are few, and so paths of influence are strongly clustered within social units. One parameter dictates the size of these social units, while a second determines the frequency of weak ties between units.

In these two networks individuals have roughly equal connectivity, and thus approximately equal levels of social influence. In contrast, in an Opinion Leader network most people have one or two connections, while a few—the opinion leaders—have many.[11] One parameter determines the level of leader influence, which depends on both the number of opinion leaders and the number of connections each has. Due to the importance of the opinion leaders, in addition to the case where internal motivations are uncorrelated with position in the network, I also consider cases where social leaders have uniformly high (positive correlation) or uniformly low (negative correlation) internal motivations. These characterize the extremes of leaders' interests and so demarcate the range of behavior to be expected.

The Hierarchy is the final network in the typology. While the power of leaders in the Opinion Leader network lies in their greater connectivity, in the Hierarchy it lies in their privileged placement at its top. The backbone of the Hierarchy is a series of levels expanding exponentially in size. Individuals are connected to one superior and a set number of subordinates; this number is one parameter of the network and determines the influence of the leaders. To a point, more leaders in the top few levels tend to make leaders overall more influential. As individuals at the same organizational level will often work closely with each other, a second parameter dictates the likelihood that one influences others within a given level. When this likelihood is low, followers at the bottom of the Hierarchy tend to have little influence relative to leaders at the top, while when it is high they can be very influential. As in the Opinion Leader network, the importance of leaders in the Hierarchy leads me to consider leaders with positively, negatively, and uncorrelated internal motivations.

If social networks are the pathways across which people are mobilized, repression acts to disrupt the mobilization dynamic by removing participants and cutting these pathways. It is exacted on the population by a unitary repressive entity that is external to the network. This may be, for example, a state, a militia, or a terrorist organization.[12] This entity has as its sole interest the minimization of societal participation in a collective action such as protest, rebellion, or voting. I assume a particular type of repression: the removal of participating individuals from the network via killing, imprisonment, exile, rendition, or the like. All removed individuals also have their network connections to others cut. As, by

[11]This type of network is commonly used to apply to the structure of the internet, and its behavior under repression, described below, exhibits some similarities to the internet under attack (Albert, Jeong, and Barabasi 2000). Simple versions of such networks have also been termed "star" or "wheel" networks (Gould 1993).

[12]As our interest is in the effect of repression on participation, rather than the behavior of the repressor, the assumption of a unitary actor is not restrictive. Multiple weaker repressors, assuming they each use the same technology, can be treated as a single stronger one in the model.

assumption, new network connections do not form, this implies that repression reduces overall network connectivity.[13] Repression thus eliminates both one's direct participation and one's indirect impact on others' participation via one's influence on others' external motivations.[14]

Repression in the model is limited in strength. Since repression equates to removal here, this limitation is operationalized as the rate at which individuals are removed.[15]

Networks such as the Opinion Leader network and the Hierarchy contain leaders with asymmetric influence. One might think—correctly, it turns out—that exactly *who* gets repressed is very important in such networks. Thus, the repressive entity has two different removal technologies within the model: random and targeted. The former randomly removes participants from the network, while the latter removes participants with the highest connectivity first. Thus, the model contains two dimensions of variation in repression: a continuous dimension corresponding to the strength of repression, and a dichotomous dimension corresponding to the technology of repression.

This completes the first model stage. Results derived from this stage in the third section focus on changes in the dependent variable of the maximal participation rate achieved as a function of the independent variables of the strength and the technology of repression, the type and the structure of the network, and the motivations of the population.

Two Psychological Responses to Repression

In the preceding analysis we assumed that violent repression altered individuals' decision making only via its effect on their external motivations. Yet, watching a friend or family member imprisoned or killed is traumatic and might alter individuals' internal motivations by instilling fear of or eliciting anger toward the repressor.[16] As a final piece of the model, I explore the changes in participation levels that arise when individuals grow either angry or fearful as those within their local networks are removed. I implement anger via a positive incentive to participate, added to one's net internal motivation, that is applied when someone to whom one is connected is removed from the network. Fear is modeled the same way, only using a negative incentive. The more individuals removed, the greater these incentives become.[17] Results derived from the psychological model focus on changes in aggregate participation as a function of the magnitude of the anger or fear response, the strength of repression, and the form of the network.

Producing Comparative Statics

Analysis of the model relies on simulation to overcome the problem of intractability. Each simulation run (or history) begins with the creation of a network and the distribution of individuals' internal motivations within the network. After initialization, in every period the following sequence of actions occurs: (1) individuals update their external motivations and decide whether or not to participate, (2) anger or fear effects may be applied, and (3) repression is applied. Each element in this sequence occurs simultaneously for all individuals in the population, continuing until no individual has changed her participation status for 50 consecutive periods. Maximal participation rates reported here are the average

[13]While this assumption is clearly more valid in some circumstances than others, preliminary work allowing an "underground" to form endogenously does not produce radically different results from those offered here.

[14]Milder forms of repression such as water cannon, or even "hearts and minds" disincentives, could be considered within the model by adjusting individuals' net internal motivations. Siegel (2011) shows that the functional response of maximal participation to this sort of mild repression is similar to that to seen in this article, justifying our focus on removal. This focus is beneficial for two reasons. First, it avoids the difficulties inherent in measuring within repressive regimes the impact of disincentives on individuals' motivations. Second, examining the impact of removal allows us to explore something we could not have under mild repression: changes in the structure of society induced by repression.

[15]A repressive entity that was itself a strategic actor might choose to vary this rate in response to circumstances, for example, in the presence of an angry response or if there were a cost to removal. This would of course alter the maximal level of participation. However, the model still predicts the level of participation to be expected under a given level of repression, even when the level of repression is not optimal. In that sense, it provides a necessary piece of the utility calculation that a strategic actor would use.

[16]The model of this subsection is in some sense a formalization of the idea of micromobilization in response to repression (McAdam 1988; Opp and Ruehl 1990). In this light, Figure 4 displays the differential response to strong and to weak repression observed empirically in Olivier (1991) and Khawaja (1993).

[17]I incorporate these emotion-induced incentives in order to explore the role of network structure on the aggregate effects of these incentives, not in order to achieve psychological veracity. The functional form for the psychological responses, described in the appendix, reflects this and is deliberately simplistic.

maximal rates over 1,000 simulation histories, each with an initial population of 1,000 individuals.[18]

I use two complementary methods for deriving interpretable results from the model despite the frequency of nonlinearity and nonmonotonicity in its outcomes. The first involves randomly sampling parameters from the relevant parameter space for each network type and then fitting average maximal participation rates to either a linear interactive or generalized additive regression model (GAM; Beck and Jackman 1998). Due to the many important interactions and nonlinearities present in model outcomes, however, the results from these analyses can be difficult to interpret; further, random sampling is unlikely to capture sufficient detail in regions of nonlinearity or nonmonotonicity. GAM regressions do provide a rough measure of the importance of input parameters, and I have included the most relevant in the appendix, referencing them in the text when appropriate.[19]

The second method involves theoretically directed, sequential parameter sweeping to derive comparative statics. Because this produces more intuitive and detailed results, the text focuses on this technique, which I describe more fully in the third section of the appendix. Sequential parameter sweeping is a method for better understanding a complex computational model with several parameters, and necessitates building the model in stages. At the first stage, only the most basic model is analyzed, ideally containing only one or two input parameters. These are slowly varied across their full ranges, and the model's outcome computed for each set of parameter values. This produces full comparative statics for the effect of these parameter values on the outcome measure in this simplified model. If no pattern can be deduced then the method ends here; however, frequently—with the aid of extant theory—one can identify regions of the parameter space in which the model's outcomes vary similarly in response to variation in the parameters. For example, increasing parameter A might always increase the outcome variable in one region, but always decrease it in another.

If one can identify such regions the model can be made more complex, adding one or two more parameters. One sweeps these parameters slowly across each of the identified regions; comparative statics on second-stage parameters hold for all first-stage parameters in a given region. This process continues until no regions can be identified at some stage of complexity. While not guaranteed to discover all possible interactions, this method, particularly when used in concert with sampling methods and regression, does produce substantial detail about the functioning of the model that would not be otherwise available.

In this model, a total of five stages were examined: (1) aggregate behavior absent networks or repression; (2) behavior in networks absent repression; (3) behavior under repression absent networks; (4) behavior in networks under repression; and (5) behavior in networks under repression with psychological responses. The next section details analysis of the fourth and fifth sections. The first two stages match those in Siegel (2009), and I rely on that article's analysis.[20] Two important facts from that work are relevant for the analysis of the last three stages. First, the space spanned by the trio of parameters $\{N, b_{mean}, b_{stdev}\}$ can be broken up into three regions, within each of which network structure acts similarly. These are denoted *motivation classes*, and called individually *weak*, *intermediate*, and *strong*. Because populations within the *weak* class participate rarely even absent repression, I consider only the *intermediate* and *strong* classes in this paper.

Second, networks without leaders (Small World and Village) may be described in terms of: (1) their levels of connectivity; and (2) whether their number of weak ties is either less-than-optimal, optimal, or greater-than-optimal in terms of how well they encourage participation. Optimal is defined as the parameterization that yields the highest level of participation in the intermediate class, all else equal, which may not occur at maximum connectivity. Networks with leaders may be described in terms of: (1) the level of influence of their leaders, (2) the correlation of motivations with network positions, and (3) the level of influence of their followers (only for the Hierarchy).

[18]Multiple instantiations of the model produce a bimodal distribution of outcomes, since, as noted earlier, the model's dynamics tend to produce either near-cascades or relatively little participation in any given run. As such, simple standard errors do not well capture the decrease in the variance of the average maximal rate as the number of runs increases. Consequently, plots display only the average maximal rates themselves, without error. However, there is little variation in average maximal rates upon increasing the number of runs past 200, implying plots display stable average values. The model's qualitative results are robust to varying the number of individuals; the third section of the appendix discusses the quantitative effects of varying N.

[19]Additional simulation data beyond that included in this article or its appendix were taken and helped inform the analysis in the next section. These are available from the author upon request. For further discussion of computational methodologies in general, see de Marchi (2005) and Miller and Page (2007).

[20]Results in Siegel (2009) continue to hold with this more general external motivation updating function.

As our focus is on network structure, details of the third stage are relegated to the fourth section of the appendix. The main lesson here is that the rate of updating, λ, matters only in comparison to the rate of removal; faster updating requires stronger repression to achieve similar outcomes. The reason for this is straightforward: nonrabble-rousing types need time to internalize the actions of other participants, increasing their external motivations, before participating themselves. Their participation draws more people in, and so on, leading to high levels of participation.[21] If repression occurs quickly relative to updating, then the repressive entity can pick out early participants before others have had time to internalize their actions. The movement never gets going, and the maximal participation rate is low.[22] If the updating rate is fast compared to the rate of removal, however, then repression will be insufficient to halt the resulting rapid increase in external motivations, leading to high participation levels (as well as a great deal of bloodshed if removal continues apace). Since a period in the model can cover any amount of real-world time,[23] only relative rates are substantively important within the model. Accordingly, no generality is lost in fixing λ.

When Does Repression Work?

As the model's description did, this section proceeds in stages. The first subsection discusses the interaction of network type, and network parameterizations within types, with repression and individual motivations. It opens with a discussion of model dynamics and then proceeds to cross-network comparisons of comparative statics. The second subsection analyzes the impact of individuals' growing either angry or fearful due to the removal of those important to them. The fourth section briefly summarizes major results.

[21]The model thus produces fundamentally path-dependent dynamics (cf. Page 2006), in that the early sequence of events can substantially change later outcomes.

[22]Populations of strongly individualistic people, influenced little by their connections, thus require significant time to build toward movements, and even weak repression is sufficient to minimize participation.

[23]E.g., a week from Monday to Monday, to mirror the mass demonstrations on 13 consecutive Mondays in Leipzig (Lohmann 1994).

The Impact of Network Structure: Dynamics

Given the path-dependent nature of behavior in the model, it helps to begin with the model's dynamics in order to understand how network structure alters the efficacy of repression. Consider a clique of seven friends attempting to mobilize despite being subjected to harsh repression by the state. Assume that one of these friends is a rabble-rouser, and that the removal rate is slow. At first only the rabble-rouser will participate, but her action will eventually draw in two others. The state may remove the rabble-rouser, but the remaining two participants will be sufficient not only to encourage each others' continued participation, but possibly also to draw in a third individual. Even if this is as far as it goes, the state must remove at least one more person to end mobilization and may have to remove two.

Now consider the exact same scenario, save instead of a clique, the network is arrayed as a star (a small Opinion Leader network), with the rabble-rouser at the center. Now the rabble-rouser has great influence, and full participation will be reached as long as no wet-blankets are present. How participation evolves from this point depends on the technology of repression. If repression is targeted the central figure is removed, collapsing the network and leading all present to cease participating. If repression is random, then a peripheral figure will likely be removed first. In this case no one else's participation status changes, as only the central rabble-rouser was connected to the removed person. In the end, several people might be removed before the central figure is caught and mobilization stops. Thus, *different removal technologies applied to different networks can lead to vastly different outcomes.*

Figure 2 is a large-scale dynamic representation of these two simple networks. The top row displays the change from period 1 (2a) to period 100 (2b) in an (uncorrelated) Opinion Leader network due to targeted repression. In such networks, if a sufficient number of opinion leaders decide to act, participation can spread very quickly. Yet, if removal is targeted at these same leaders, it can quite effectively squash participation. The period 100 image is the network after targeted removal has ended all participation, and its structural effects are clear visually: the central region, in which are located the leaders, is substantially diminished. Participating opinion leaders were the primary targets of removal, while the nonparticipating leaders still around actually *support* the goal of repression due to their depression of the

Figure 2 Change in Network Structure under Targeted Repression

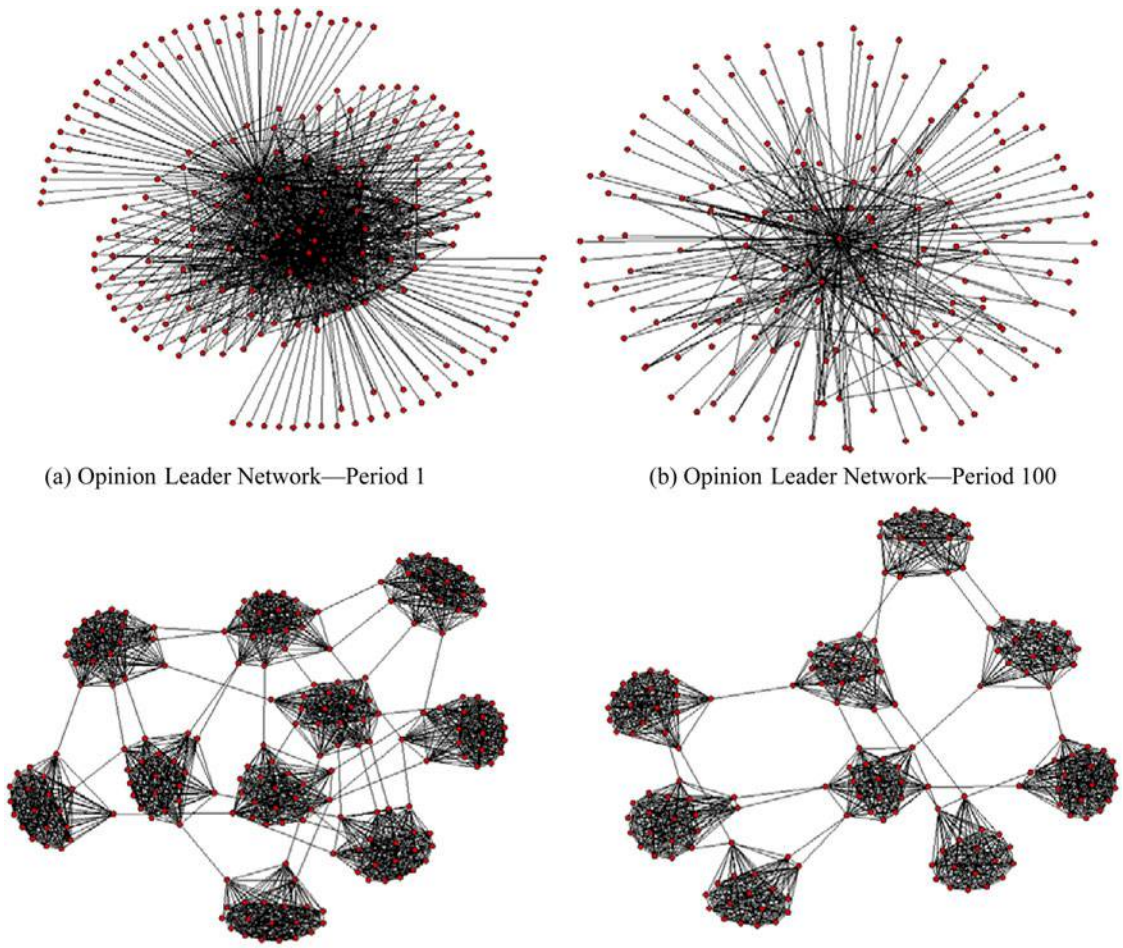

(a) Opinion Leader Network—Period 1

(b) Opinion Leader Network—Period 100

(c) Village Network—Period 1

(d) Village Network—Period 100

external motivations of their followers. All told, the maximal level of participation achieved here is under 24%, and a total of 36 people out of 200 were removed before participation was pushed down to zero.

The bottom row of Figure 2 displays the change in a Village network between the same two periods, also under targeted removal. Here the structural impact of repression is more localized. Where there were once 10 villages of 20 people each, now there are only nine. An entire village has simply been wiped off the map by the repressive entity. Importantly, this is *not* a consequence of the removal technology; targeted removal is less likely to remove an individual from a village the smaller is the village. Instead, the steady destruction of a single village—amounting to genocide if the village encompasses an ethnic group—is indicative of the model's dynamics. While Opinion Leader networks spur action widely via their leaders, participation within Village networks spreads first and easiest within villages. As such, over half of the 38 people removed from the network by period 100 *all came from the same village*, and destroying that village limited the maximal participation rate to 11%.

The contrast between the effect of repression on these two network types is striking. Consider the context of a repressive regime, determined to hold on to power regardless of the cost in human life. In an Opinion Leader network, a small number of leaders each hold significant sway over a large number of followers. Proregime leaders sit down when the call to mobilize comes, and in so doing encourage others to stay home as well. Antiregime leaders ignore the risks and suffer for it, losing their lives. Their brave actions do spur their followers on for a time, but without the leaders' continued support the followers gradually return to their homes. To the extent that leaders of both varieties are scattered throughout the nation, the struggle and loss that the antiregime leaders experience will play out on a national stage, and few will be unaware of it. History, being written by the victors, may demonize such rebellious leaders over time, but their story will be remembered and their elimination will change the fabric of society, as their former followers must now look elsewhere. This same tale played out in a Village network is far less dramatic for the great majority of citizens. Various rabble-rousers may disappear from the population, but the real action takes place in the single village that is brutally wiped off the map. There no one remains to tell the tale, and the regime's brutality is more easily covered up. In time others may forget about the violence entirely and write a history minimizing the entity's repressiveness.[24]

The Impact of Network Structure: Comparative Statics

Sample paths such as those above are illustrative of the way networks respond to repression, but they do not tell the whole story. Participation levels in any history depend not only on gross quantities of interest like network type and repression technology, but also on the specific location of individuals within the network and which particular people are removed. To surmount these specificities and derive general relationships across network types, we must average over many sample paths; these averages are our comparative statics.

The six plots of Figure 3 display the effect of harsh repression within networks. Each row illustrates an important commonality in the response of networked behavior to repression. The top row displays Small World and Village networks, which behave similarly under repression, conditional on their levels of connectivity and of optimality of weak ties. Taken together, they convey the importance of network structure (comparing lines in each plot) and motivation class (comparing plots) on the efficacy of repression. The middle row displays Opinion Leader networks with influential leaders and illustrates both the ability of targeted repression to disrupt networks without unified leaders (comparing lines in Figure 3c), and the power of unified leaders in resisting repression (comparing lines in Figure 3d to those in 3c). The bottom row makes the same point in a Hierarchy when the masses are not influential (Figure 3e), but shows that influential masses can introduce a vulnerability into the network (Figure 3f). In each plot, the vertical axis displays the average maximal participation rate, and the horizontal axis displays the removal rate; repression increases along the horizontal axis.

Comparison of Figures 3a and 3b immediately reveals one simple commonality: *as the rate of repression increases, the differences between network parameterizations decrease*. Increase the intensity of repression enough by ramping up the rate of removal and not only will participation be squashed, it won't matter what network is in place. This response to repression occurs in all network types—all lines eventually converge far enough to the right. Note that this is *not* due to complete breakdown of the network from removal. In fact, if removal is fast enough, then increasing the removal rate further *decreases* the total number of people removed. This is due to the path-dependent nature of the model; sufficiently fast removal cuts out early participants before they have a chance to be influential, diminishing the importance of the network connections over which they could have been influential.[25]

Not all movements are immediately crushed, of course, so it is likely that levels of repression less extreme than this are of substantive import. It is thus worth considering the primary difference between the plots. Consider the bottom two lines in Figure 3a, corresponding to an optimal number of weak ties in the strong class of the Small World network under random removal. As we can see, lines corresponding to both levels of connectivity indicate high participation levels absent repression. As the rate of repression increases, levels of participation decrease, again in a similar fashion in both lines. The same is true in Figure 3b, corresponding to the intermediate class of the Village network under targeted removal. But now compare the two plots. The lines in Figure 3b plunge abruptly, indicating that even low rates of removal are extremely effective at reducing participation, while those in Figure 3a decline much more gradually.

Why might this be? It is not due to the difference in removal technology. Lacking social leaders, both Small World and Village networks respond similarly to random and targeted removal. Nor can it be attributed to network-specific effects: the same difference can be observed when comparing the strong and

[24]Rosewood, Florida in 1923, and Tulsa, Oklahoma in 1921 illustrate that stories such as this are not limited to state repression within autocracies. See also accusations of ethnic cleansing in the American South (http://www.austinchronicle.com/gyrobase/Issue/story?oid=oid%3A456310. Last accessed 2/21/08.).

[25]The fifth section of the appendix contains additional support for this result.

FIGURE 3 Network Structure, Motivation Class, and Repression Technology

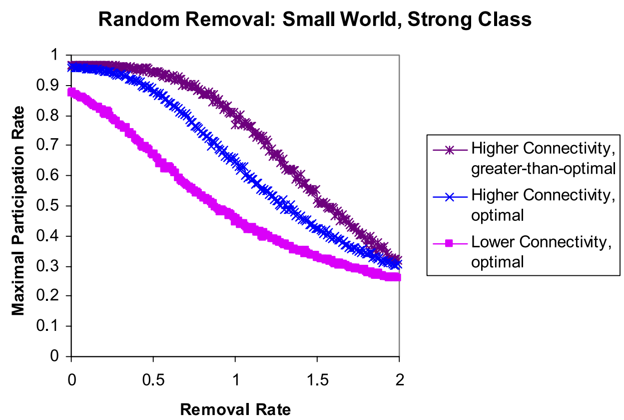

(a) Network Structure Alters Efficacy of Repression

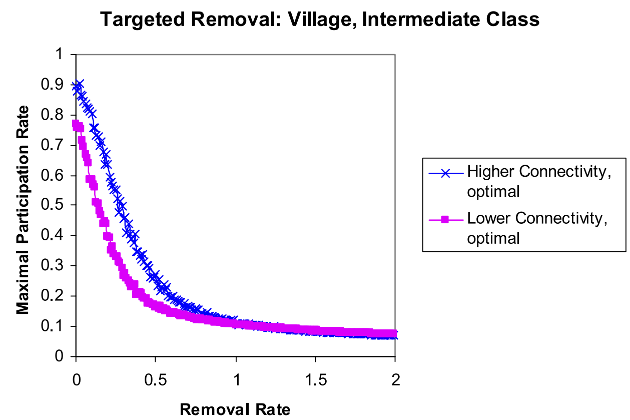

(b) Lower Motivation Class, More Effective Repression

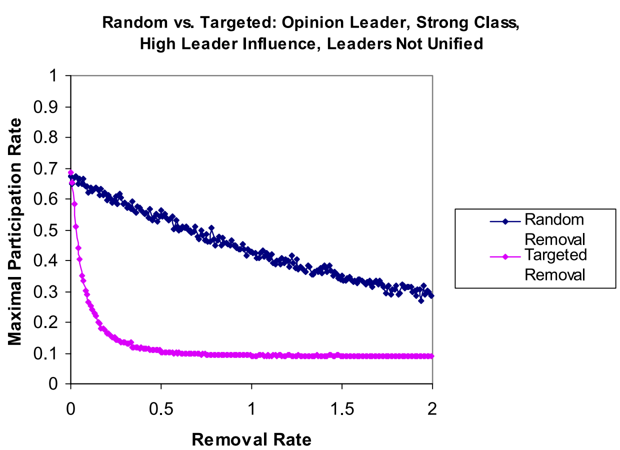

(c) Non-Unified Leaders Are Vulnerable to Targeting

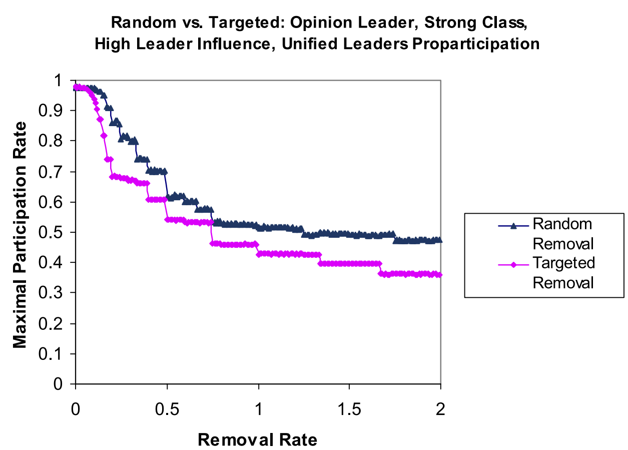

(d) Unified Leaders Are Resistant to both Repression Techs

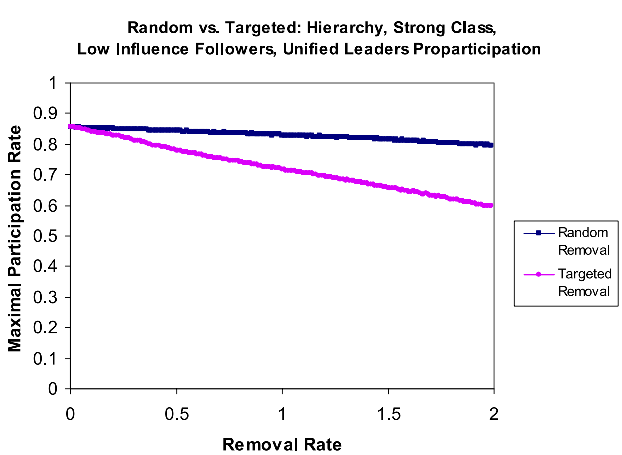

(e) Unified Leaders Are Resistant to both Repression Techs

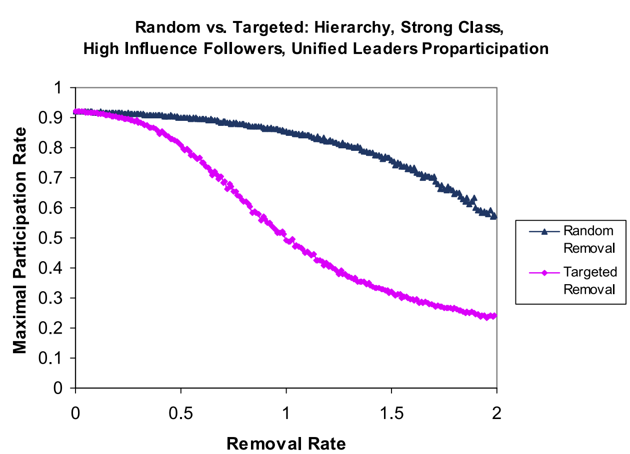

(f) Influential Followers Create Vulnerability to Targeting

the intermediate motivation classes in a Small World network, or the same two in a Village network. This leaves motivation class as the culprit. The strong class is far more robust to removal than the intermediate class, even when both produce high levels of participation absent repression. Participation in the intermediate class is fragile, in that it depends vitally on a comparatively small number of rabble-rousers and other high internal-motivation types in the network. Removing these individuals has a massive depressive effect on participation. In contrast, in the strong class there are many willing to take up the cause with little incentive from others' actions, providing a bulwark against repression.

This result is an example of a more general regularity: *networks that rely on very specific parameter configurations to achieve significant levels of participation are more vulnerable to repression*. Repression disrupts the delicate balance that allows such networks to spur participation efficiently and successfully. A second example of this regularity can be seen in the top line in Figure 3a, corresponding to a greater-than-optimal number of weak ties. This parameterization yields a level of participation absent repression identical to that observed for the optimal number of weak ties, but it proves significantly more robust to removal. Optimality of weak ties implies that the network has just the right number, with no redundancies, and removal has a significant effect on this fragile system. A greater-than-optimal number of weak ties may not be efficient—and actually lowers participation levels in the intermediate class—but it does provide a measure of protection against repression.

Though networks with optimal numbers of weak ties and networks in the intermediate class are fragile for similar reasons, they are not equally so; disrupting the configuration of weak ties has a lesser depressive effect on participation. Thus: *In data-poor regions corresponding to either a Small World or a Village network, obtaining general information about the level of dissatisfaction—one possible proxy for the distribution of internal motivations—is likely to be more important in predicting the response to repression than detailed information about network connections would be.*[26]

The next four graphs illustrate a closely related point: *networks that rely on a handful of individuals to achieve significant levels of participation are more vulnerable to targeted repression, but less vulnerable to random repression.*[27]Consider first Figures 3c and 3d, corresponding to Opinion Leader networks. In both plots, we consider only influential leaders, as less influential leaders produce far lower levels of participation absent repression, leading trivially to less of a role for repression. When leaders have many ties, they play a vital role in achieving high levels of participation. By the same logic as above, then, removing them should have a substantial depressive effect on participation even at low levels of repression. This is borne out in Figure 3c. If repression targets leaders, as in the lower line, the maximal participation rate achievable in the society plummets, even if the removal rate is slow. As most people are not leaders in this network, however, random removal usually only picks out uninfluential followers. As a consequence, networks in which participation is driven by a few individuals are generally robust to random removal, as indicated by the upper line. Channeling Machiavelli, we might instruct our princes that it is comparatively ineffective to repress an Opinion Leader network if a lack of intelligence or capability prevents the targeting of leaders. Conversely, individual followers subject to repression should not despair at overthrowing the oppressive regime, if they know their opinion leaders are safe from or resistant to repression.

The story gets better for the participating populace in Figure 3d. With leaders now possessing uniformly high motivations, participation levels are not only very high, but also extremely robust to repression, *even targeted repression*. The additional support from leaders that positive correlation entails provides this robustness, in a more potent way than greater-than-optimal weak ties did in the Small World network described earlier. With numerous leaders all spreading participation among their followers, removal of a few of them cannot stem the tide. Machiavelli's message is stronger here: a unified group of network leaders is difficult to deny, and only the most powerful (and brutal) of repressive regimes should attempt it. Conversely, a population possessed of such leaders can look forward to the success of their cause, though a determined (and most brutal) repressor will exact a steep toll before this success is achieved.

The Hierarchical network plots displayed in Figures 3e and 3f show a similar response to repression, with one twist: interconnections within levels of the hierarchy provide the opportunity for the masses to influence the level of participation in the network. In both plots, leaders are uniformly highly motivated to participate. In Figure 3e the masses are not influential, and the behavior of the Hierarchy under repression is much the same as it was in the Opinion Leader network: we observe high levels of participation that are robust to both technologies of repression. When, as in Figure 3f, followers do possess significant influence via intralevel network connections, they increase overall participation absent repression, helping the leaders' message to

[26]This is borne out in the GAM regressions displayed in Figures A4 and A5 of the appendix.

[27]This is borne out in the GAM regressions displayed in Figures A7 and A8 of the appendix.

disseminate. However, this introduces a vulnerability into the network as well. As repression increases, it inhibits or eliminates the conduits that pass down the leaders' messages to their subordinates, leaving followers on their own. When they are poorly interconnected, this does not greatly alter participation levels. When followers are influential, in contrast, removing ties to the leadership enhances the role of one's connections to other followers. Since followers are less motivated than their leaders here, the result is a steeper decline in participation than is the case with less influential followers. Targeted repression, which on average goes after conduits to the leadership first, is even more effective than random repression for this reason. The extra connections cost the network its robustness to targeted repression.

Anger and Fear

Thus far I have assumed that people respond only indirectly to the removal of their fellows. Yet their removal might occasion emotional responses as well. Figure 4 displays the outcome of adding anger or fear to the core behavioral model under random removal within the intermediate class,[28] across a sampling of network types and parameterizations. The horizontal axes here display the level of anger (moving right on the axis) or fear (moving left on the axis); the vertical axes display the maximal participation rate, averaged over 1,000 runs.

Figure 4a considers a simplified world in which all individuals are connected to all others (a "Fully Connected" Network) in order to develop a baseline measure of the effect of anger and fear. Each line corresponds to a different rate of removal, with the horizontal baseline indicating participation levels absent removal, and thus absent fear or anger effects as well. As might be expected, *anger increases participation and fear decreases it, with the effect mitigated at faster removal rates.* Less obvious is that, *if anger is strong enough, participation levels can be higher under repression than absent it. Individual anger at local repression endogenously enables aggregate backlash.* Further, for weak repression, comparatively little anger is needed to achieve backlash against the repressive entity.

Now turn to Figure 4b, which displays the effect of increasing anger or fear on two different parameterizations of a Village network. As seen on the y-axis, both parameterizations display little participation absent a psychological response. This symmetry is broken dramatically when emotions come into play, however. Though both parameterizations are exposed to the same rate of repression, anger only enables significantly increased participation when there is an optimal number of weak ties in the network (the solid line), not when the number is less than optimal (the dotted line). And significant is an understatement: a little bit of anger produces a threefold increase in participation.

To understand why this happens, it is necessary to return to the dynamics of participation spread within a Village network. As participation within such networks spreads first within a given village, and then more slowly to other villages, anger that arises due to removal will largely be localized in the village from which the individual was removed. This raises internal motivations in that village, which leads to more participation in that village, which increases external motivations, which leads to even more participation. This cascade effect occurs in both parameterizations, but without sufficient weak ties, the cycle ends within that village hotbed. Thus anger has little aggregate effect when network structure doesn't allow it to spread. However, once there is a sufficient number of weak ties, anger-driven participation can spread throughout the network rapidly enough to overwhelm repression and trigger a backlash.

This logic generalizes from the Village network; *anger and fear only generate aggregate backlash when the individuals most directly affected by them have sufficient ties to people further afield.* This is what drives the outcome illustrated in Figure 4c as well, which displays an (uncorrelated) Opinion Leader network. Again anger enables a substantial increase in participation, and here the effect is even more striking. When the removal rate is relatively slow, minimal anger can bring maximal participation up to near 100%, a three-fold increase; when it is fast the increase jumps to sixfold. Further, backlash is achieved for even low levels of anger, particularly when removal is slow. The leaders' high connectivity is the cause of this: whenever a leader is removed it makes a large number of people angry; conversely, a leader's anger over the removal of a follower spreads easily to the rest of the population. Consequently anger alters aggregate behavior swiftly and efficiently. So swiftly, in fact, that the resulting backlash makes nearly any form of repression unwise from purely practical considerations. As such,

[28]Using the strong class instead would increase the role of fear in an analogous set of plots to Figure 4.

Figure 4 Network Structure Multiplies Effect of Anger and Fear

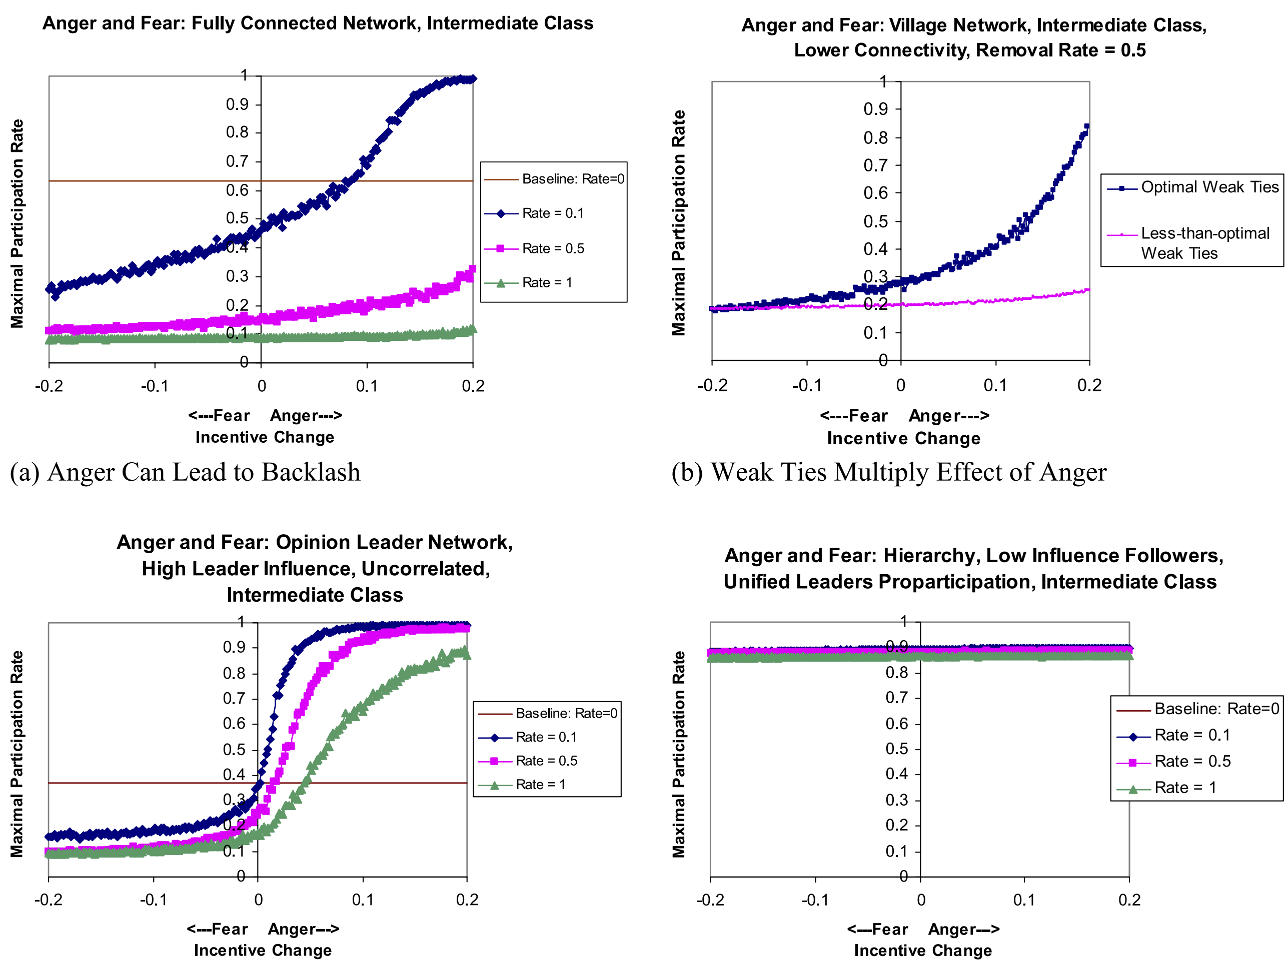

(a) Anger Can Lead to Backlash

(b) Weak Ties Multiply Effect of Anger

(c) Nonunified Leaders Greatly Multiply Effect of Anger

(d) Unified Leaders Overcome Fear and Anger

the message here is clear. *If a network configuration affords the swift spread of participation, the mere threat of an angry response to repression should be sufficient to rule it out entirely.* This theoretical result matches that shown empirically in Kaplan et al. (2005): Israel's killing of terror suspects (who share a network) leads to increased recruitment and subsequently to increased attacks (participation), while the killing of Palestinian civilians (more poorly connected to the terror network) does not have this effect.

There is one exception to this, however, and it arises among network configurations that are robust to repression. Such networks also prove robust to things like anger and fear, as seen in Figure 4d. For example, when the leaders are aligned against the repressive entity, as they are in this network, even added fear is insufficient to counter leader influence and quell mass participation.

Application of the Model and Summary of Results

In order to apply the model, it is first necessary to identify the type of network. The degree to which additional information is useful depends on the network in place. Table 1 provides, for each network type, a few heuristics for identifying the network and its properties, and an ordered list of the relative importance of additional information. In all cases additional information refines predictions as to the efficacy of repression and the level of expected participation, but in most cases one needs only network type and some idea as to the nature of repression (strength or technology) to produce a rough quantitative estimate of repression's effect. This is what drives the earlier claim that the model has comparatively mild data requirements: one need

Table 1 Heuristics for Model Application

Network Type	Identification Heuristics	Order of Information
Small World (SW)	Most Know Someone Far Away, Info Travels Quickly	Repressive Strength, Motivation Class, Average Connectivity, # Weak Ties
Village/Clique (V/C)	Social Groups Are Clumped, Info Travels Poorly btwn Cliques	Repressive Strength, Motivation Class, Average Connectivity, # Weak Ties
Opinion Leader (O-L)	Few Leaders Drive Opinion, Info from Common Sources, Skewed Dist of Connections	Leader Influence, Leader Motivations, Repressive Tech, Repressive Strength, Motivation Class
Hierarchy (H)	Rigid Chain of Influence, Few Superiors, Many Subordinates, Defined Organizational Structure	Leader Influence, Leader Motivations, Repressive Tech, Repressive Strength, Motivation Class, Follower Influence

not possess data on all connections in the network to understand how effective repression will be.

There are seven potential pieces of information listed in Table 1 that might be useful in predicting the effect of repression. Two of these, the strength and technology of repression, act directly to reduce participation levels. Table 2 lists a few heuristics for measuring each of these and a summary of when increasing each is most likely to be effective. The other five of these—motivation class, average connectivity, leader influence, leader unity, and follower interests—act both directly on participation, and conditionally on participation by altering the efficacy of repression. Table 3 provides a couple of heuristics for measuring each of these, and a brief statement as to whether the most common role of increasing the parameter is to increase or decrease the efficacy of repression, in the absence of anger or fear. The third section provides a much fuller picture of the interaction between network structure, motivations, and repression.

An eighth piece of information, the presence of anger or fear, is not listed because it modifies the effect of all parameters. The presence of anger (or fear) can greatly increase (or decrease) participation under repression, but only if a sufficient number of weak ties exists in a Small World or a Village network, or if an Opinion Leader or a Hierarchical network configuration is not robust to both technologies of removal.

To see how these heuristics for the model's application may be applied in practice, one can consult the sixth section of the appendix. There I provide a concrete example, utilizing qualitative data drawn significantly from Patel (2005), to "fit" the networks consisting of the followers of Grand Ayatollah Ali Sistani and Muqtada al-Sadr into the framework presented in this article. Doing so allows for a novel prediction of the differential levels of turnout under repression between the followers of each Shi'ite leader during the January 2005 Iraqi Legislative Elections, using only data available prior to the elections. I show in the appendix that each network may be modeled as a hierarchy, with the primary difference between the two the degree to which the upper echelons of the hierarchies are unified in their motivations. Sistani's network's leadership could be considered to be substantially more unified in its motivations than Sadr's, implying, as seen in Figure A9 of the appendix, that Sistani's network would be more robust to the random repression occurring in Iraq at the time, and produce

Table 2 Measurement and Direct Effect of Repression Parameters

Information	Measurement Heuristics	When Increasing Most Effective
Repressive Strength	State Capacity, Forced Migrants, State Killings, "Disappearances"	Non-unified O-L,H, Few Weak Ties in SW,V/C
Repressive Technology	Military/Police Capacity, State Domestic Intel, Civil Liberties	Non-unified O-L,H, Few Influential Leaders

Note: Increasing Repressive Strength equates to removing more people per period. Increasing Repressive Technology equates to using Targeted rather than Random Repression.

Table 3 **Measurement and Conditional Effect of Network and Motivation Parameters**

Information	Measurement Heuristics	Effect on Efficacy of Repression
Motivation Class	Public Opinion/Dissatisfaction	Decreases Efficacy
Av Connectivity	Survey Questions, Observation	Decreases Efficacy
Leader Influence	#, Stated Importance of Leaders	Decreases Efficacy
Follower Influence	Level Socializing of Subordinates	Increases Efficacy
Leader Unity	Public Statements, Elite Survey	Greatly Decreases Efficacy

Note: Increasing Motivation Class equates to moving toward the Strong Class, in which more individuals are predisposed to participate. Increasing Average Connectivity implies larger individual networks, on average. Increasing Leader Influence increases the number and average connectivity of leaders, who have many connections. Increasing Follower Influence increases the average intra-level connectivity in a Hierarchy. Increasing Leader Influence moves from non-unified leaders to those unified either pro- or anti-participation. (Both decrease the efficacy of repression, though the former also decreases participation on its own.)

higher levels of participation than Sadr's network, particularly in the presence of repression. The outcome of the elections matches this prediction, with Sistani's objectives—voting and voting for the UIA—receiving far more support than Sadr's—not voting or voting for the National Independent Cadres and Elites. Though sufficient data do not exist to distinguish between causal explanations, as we lack data on prior support for each of Sistani's and Sadr's objectives, variation in the violence directed against voters, and rates of voting of supporters of each leader, the outcome of the elections is certainly consistent with the theory presented here and illustrates the power of the model in drawing novel quantitative aggregate predictions based on qualitative data.

Conclusion

We have seen that answers to the question "When does repression work?" depend fundamentally on the structure of the social network that connects the population. Importantly, network structure interacts with the distribution of individuals' motivations within networks and the nature of repression in producing outcomes. Not only is understanding just one or two of these insufficient for predicting the efficacy of repression, one must also consider the conditioning effect of each on the others. While this leads to a more complex causal story, it is a necessary one for understanding when repression works.

As one example, the degree to which social leaders attempting to rouse their followers to action against the state find their effectiveness blunted by repression depends on their interests, their connections to their followers, and the repressive technology in place. Unified leaders are proof against all but the strongest repression, as long as they have sufficient connections to their followers. Promovement leaders facing opposition from other leaders, in contrast, are vulnerable when repression targets them specifically: the combination of their elimination and the antimovement leaders' continued influence makes even weak repression extremely effective. If followers grow angry at a leader's loss, however, leaders need not be unified; the same network structure that gives them power multiplies the effects of even mild anger, leading to a powerful backlash against the state.

One explanation for why networks have not been more common in analyses of behavior to this point, despite their importance, is that detailed individual-level data on social connections under repression rarely exist. The model presented in this article circumvents this problem through the use of a typology of qualitative networks and a methodology by which relevant network factors may be identified. By linking qualitative factors with formal analysis, this article illustrates that scholarship in both veins may be productively joined.

The model's broad specification and mild data requirements enable its utility across a wide range of substantive cases. Scholars seeking to employ the methodology of this paper and utilize its results need only follow the procedure laid out in the fourth section; the online appendix provides a concrete example of such an application: the case of the January 2005 Iraqi Legislative elections. Though this case was somewhat stylized in that years of Saddam Hussein's oppressive regime made the network analysis fairly clean, this sort of prospective analysis is portable to other cases and will be particularly easily achieved when the networks in play share some clearly identifiable difference. In the case of the Iraqi elections it was the degree of commonality in the interests of social leaders. Elsewhere it might be different types of networks or different levels of network connectivity. The important thing to note here is that detailed network characteristics, not often available, are not always necessary. Thus, in addition to its contribution to the literature on repression, this article contributes to the social network literature not only by addressing removal from networks, but also by illustrating how hypotheses can be drawn from sparse network data.

Though the article has focused on repression, the importance of networks applies more broadly, and the methodology introduced here is equally useful in other substantive areas. Multiple overlapping networks, competing repressors, multiple forms of participation, and other variants of the model can all be analyzed, though one must take care to maintain a strong sense of the model's causality despite the increase in complexity while doing so. In time this approach will allow inroads into such questions as state collapse, once the loop is closed and the participation of people is allowed to affect the viability of the repressor.

Acknowledgment

The author thanks Dave Baron, Jon Bendor, Özge Kemahlioglu, David Laitin, Charlotte Lee, Matt Golder, Sona Golder, Will Moore, David Patel, Chris Reenock, John Scholz, Jake Shapiro, several anonymous reviewers, the Political Institutions Working group at FSU, and the Political Economics group at the Stanford GSB for helpful feedback.

References

Albert, R., H. Jeong, and A. L. Barabasi. 2000. "Error and attack tolerance of complex networks." *Nature* 406 (6794): 378–82.

Beck, Nathaniel, and Simon Jackman. 1998. "Beyond Linearity by Default: Generalized Additive Models." *American Journal of Political Science* 42 (2): 596–627.

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

Chwe, Michael S. Y. 1999. "Structure and strategy in collective action. *American Journal of Sociology* 105 (1): 128–56.

Davenport, Christian. 2007. "State Repression and Political Order." *Annual Review of Political Science* 10: 1–23.

DeNardo, James. 1985. *Power in numbers: The political strategy of protest and rebellion*. Princeton, NJ: Princeton University Press.

Fowler, James H. 2005. "Turnout in a Small World" in *The Social Logic of Politics: Family, Friends, Neighbors, and Workmates as Contexts for Political Behavior*, ed. Alan Zuckerman. Temple University Press. 269–287.

Fowler, James H., Timothy R. Johnson, James F. Spriggs, II, Sangick Jeon, and Paul J. Wahlbeck. 2007. "Network Analysis and the Law: Measuring the Legal Importance of Supreme Court Precedents." *Political Analysis* 15 (3): 324–46.

Francisco, Ronald A. 1995. "The Relationship Between Coercion and Protest: An Empirical Evaluation in Three Coercive States." *Journal of Conflict Resolution* 39 (2): 263–82.

Francisco, Ronald A. 1996. "Coercion and Protest: An Empirical Test in Two Democratic States." *American Journal of Political Science* 40 (4): 1179–1204.

Francisco, Ronald A. 2004. "After the Massacre: Mobilization in the Wake of Harsh Repression." *Mobilization: An International Journal* 9 (2): 107–26.

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

Granovetter, M. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78: 1360–80.

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

Gupta, Dipak K., Harinder Singh, and Tom Sprague. 1993. "Government Coercion of Dissidents: Deterrence or Provocation?" *Journal of Conflict Resolution* 37: 301–39.

Gurr, T. R., and Moore, W. 1997. "Ethnopolitical rebellion: a cross-sectional analysis of the 1980s with risk assessments for the 1990s." *American Journal of Political Science* 41 (4): 1079–103.

Hibbs, D. A. 1973. *Mass Political Violence: A Cross-National Causal Analysis*. NewYork: Wiley.

Kaplan, Edward H., Alex Mintz, Shaul Mishal, and Clausio Samban. 2005. "What Happened to Suicide Bombings in Israel? Insights from a Terror Stock Model." *Studies in Conflict & Terrorism* 28: 225–35.

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

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

Khawaja, Marwan. 1993. "Repression and Popular Collective Action: Evidence from the West Bank." *Sociological Forum* 8 (1): 47–71.

Lichbach Mark, I. 1987. "Deterrence or Escalation? The Puzzle of Aggregate Studies of Repression and Dissent." *Journal of Conflict Resolution* 31: 266–97.

Lichbach Mark, I. and Gurr Ted, R. 1981. "The conflict process—a formal model." *Journal of Conflict Resolution* 25 (1): 3–29.

Lichbach Mark, I. 1995. *The Rebel's Dilemma*. Ann Arbor: University of Michigan Press.

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

de Marchi, Scott. 2005. *Computational and Mathematical Modeling in the Social Sciences*. Cambridge: Cambridge University Press.

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

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

McAdam, Doug. 1988. "Micromobilization Contexts and Recruitment to Activism." *International Social Movement Research* 1: 125–54.

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. 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.

Miller, John H., and Scott E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, NJ: Princeton University Press.

Muller, E. 1985. "Income Inequality, Regime Repressiveness, and Political Violence." *American Sociological Review* 50 (1): 47–61.

Myers, Daniel J., and Pamela E. Oliver. 2005. "The Opposing Forces Diffusion Model: The Initiation and Repression of Collective Violence." *University of Notre Dame Department of Sociology Working Paper and Technical Support Series* Number 1999–01.

Ohlemacher, T. 1996. "Bridging people and protest: Social relays of protest groups against low-flying military jets in West Germany. *Social Problems* 43 (2): 197–218.

Oliver, P. E. 2002. "How Does Repression Work? Working Paper. University of Wisconsin, Madison.

Oliver, Pamela E., and Myers, Daniel J. 2002. "The Coevolution of Social Movements." *Mobilization: An International Journal* 8 (1): 1–24.

Olivier, Johan L. 1991. "State Repression and Collective Action in South Africa, 1970-1984." *South African Journal of Sociology* 22: 109–17.

Opp, Karl-Dieter, and Wolfgang Roehl. 1990. "Repression, Micromobilization, and Political Protest." *Social Forces* 69 (2): 521–47.

Opp, K. D., and C. Gern. 1993. "Dissident Groups, Personal Networks, and Spontaneous Cooperation - the East-German Revolution of 1989. *American Sociological Review* 58 (5): 659–80.

Page, Scott E. 2006. "Path Dependence." *Quarterly Journal of Political Science.* 1 (1): 87–115.

Patel, David Siddhartha. 2007. "Islam, Information, and Social Order: The Strategic Role of Religion in Muslim Societies." Ph.D. diss. Stanford University.

Petersen, Roger Dale. 2001. *Resistance and Rebellion: Lessons from Eastern Europe, Studies in Rationality and Social Change*. Cambridge and New York: Cambridge University Press.

Rasler, K. 1996. "Concessions, Repression, and Political Protest in the Iranian Revolution." *American Sociological Review* 61 (1): 132–52.

Rolfe, Meredith. 2005. "Social Networks and Threshold Models of Collective Behavior." Working Paper. University of Chicago.

Siegel, David A. 2009. "Social Networks and Collective Action." *American Journal of Political Science* 53 (1): 122–38.

Siegel, David A. 2011. "Nondisruptive Tactics of Suppression are superior in countering Terrorism, Insurgency, and Financial Panics." PLoS ONE 6 (4): e18545. doi: 10.371/journal.pone.0018545.

Snow, D. A., L. A. Zurcher, and S. Eklandolson. 1980. "Social Networks and Social-Movements - a Microstructural Approach to Differential Recruitment." *American Sociological Review* 45 (5): 787–801.

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

David A. Siegel is Assistant Professor at Florida State University, Tallahassee, FL 32306-2230.