Web Appendix for When Does Repression Work?: Collective Action in Social Networks

David A. Siegel

The first section of this technical appendix details the order of operations within the model, the second the method of network construction and parameterization, and the third the methodology by which results given in the paper were derived. The fourth section details additional analysis of the role of the updating rate, λ, absent a complex network. The fifth section does the same for core behavior in networks under repression.

1 Order of Operations

Each model history (run) begins with: 1) the creation of a network containing N individuals; and 2) the assignment of net internal motivations b_i to each person, drawn from a normal distribution with mean b_{mean} and standard deviation b_{stdev}. The net internal motivation will most likely be the sum of random variables drawn from many different distributions, each corresponding to a different aspect of one's motivation. Assuming these draws are independent, and that their third central moments satisfy the Lyapunov condition, the sum of these random variables will be distributed normally, by Lyapunov's central limit theorem. Note that, even if one argues that there are few relevant motivations or that draws are not independent, this only affects the specification of the distribution from which the b_i are drawn.[1] The basic idea of each individual's having some net internal motivation—upon which this model rests—still holds. The b_i may be uncorrelated with network position, positively correlated (large b_i get assigned to the most well-connected people in the Opinion-Leader network or to those at the top of the Hierarchy), or negatively correlated (vice-versa). See Siegel (2009) for additional details and justification of these two steps.

In each period after this, the following sequence repeats: 1) Individuals update external motivations according to the equation $c_{i,t+1} = \lambda c_{i,t} - (1 - \lambda)(1 - lpr_{i,t})$, and participate in a given period if $b_i + c_{i,t} > 0$. All individuals update and decide simultaneously, and the values of the local participation rate ($lpr_{i,t} \in [0, 1]$) they use to do this are calculated between periods in step four. This particular weighting rule has few cognitive requirements, and its use of the parameter $\lambda \in [0, 1]$ allows direct variation in the speed with which individuals respond to others' actions. 2) For anger (fear), each person increases (decreases) her b_i by the parameter *Bchange* multiplied by a fraction equal to the number of social connections one has lost through removal divided by the total number of these losses across the population. All graphs displaying anger or fear effects have the value of *Bchange* on the x-axis.[2] 3) Participating individuals are removed under repression at a frequency equal to the *removal rate*. Fractional values of the *removal rate* correspond to removal slower than a person a period; e.g., a rate of 0.2 indicates one person removed every five periods. Under random removal, participants to be removed are chosen at random; under targeted removal, the most well-connected participants are removed first. 4) Participation rates are (re)calculated, both population-wide and locally for each individual. The latter calculation provides ($lpr_{i,t}$).

---

[1] Different distributions do lead to different aggregate levels of support (see, e.g., Granovetter and Soong 1986; Yin 1998) in the absence of repression; an investigation of the same under repression could be a productive avenue of further research.

[2] Note that this rule for anger (fear) has each person compare losses in her network to the total in the population, rather than to the total number of connections one ever had. While lacking somewhat in psychological realism (see footnote 21 in the text), this specification has the advantage of keeping constant the total level of motivation change in the population, which allows for easier understanding of casual pathways. All qualitative results discussed in the paper, however, are robust to an alternate specification in which people compare their losses to their potential losses.

Finally, if no individual has changed her participation status for fifty consecutive periods, the run ends and final data for that run are recorded. After 1000 (for comparative statics) or 200 (for GAMs) independent runs for a given model parameterization are completed, means and standard deviations of maximum participation levels over these runs are calculated. (Results do not differ between 200 and 1000 realizations; see the text for justification of the absence of error terms in the plots.) All simulation data were obtained via a JAVA program coded by the author.

2 Network Construction and Parameterization

All network ties are symmetric 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. The following is a brief description of the networks in the typology and their associated parameters. See Siegel (2009) for detail sufficient for replication.

Fully-Connected: Individuals respond to aggregate participation levels. **Small World**: The base is a ring in which people are connected to *Connection Radius* other individuals to both sides of them. This sets the average connectivity. A *Connection Radius* of 5 thus indicates a connectivity of 10. Each connection has a chance equal to the parameter *Rewire Probability* of being severed and reconnected randomly. This sets the frequency of weak ties. **Village**: The base is an array of equally-sized groups called villages or cliques that are of size *Village Size*; every possible connection within these is made. If the population does not divide evenly in this way, then any left over individuals are placed into a last, smaller village. Each individual also has some probability, called *Far Probability*, of being connected to another individual outside of her village. These probabilities, which determine the number of weak ties, are checked twice, so the true probability of any individual's being connected to a particular person outside her own village is equal to twice *Far Probability*. **Opinion-Leader**: Each individual is assigned a number of ties, k, according to the distribution $p(k) \propto k^{-\gamma}$, and connected randomly to this number of people. The parameter γ thus determines the characteristics of the network, with smaller values corresponding to greater leader influence, as there are more leaders with greater individual connectivity. These networks are also known as Scale-Free networks. **Hierarchical**: The base is a hierarchy determined by parameter *Expansion Rate* in which one individual is placed at the top, and each individual in the network is connected to a number of individuals below him equal to *Expansion Rate*, continuing until no more individuals are left in the population. Thus, while each level of the hierarchy before the last one contains a number of individuals equal to a power of *Expansion Rate*, the last level may have fewer than this if the total population does not divide appropriately. Higher values of this parameter tend to lead to more influential elites, at least to the point where the top does not become too watered down. Each potential tie between individuals within the same level also has a probability equal to *Level Connection* of being made. The higher this probability, the greater the influence of the masses.

3 Sequential Parameter Sweeping

As discussed in the text, deriving rigorous results from a computational model with more than a couple of parameters requires care. The method of parameter sweeping I use largely matches that described and justified in Siegel (2009). It differs only in the following respects. First, the number of parameters are greater in this model, requiring additional parameter regions to be identified. These

are described below. Second, I employ a supplementary generalized additive model (GAM) analysis (Beck and Jackman 1998). This entails sampling uniformly over the entire parameter space, and then fitting a GAM with these simulation data via a spline (or other) scatterplot smoother. This approach reduces problems of non-linearity and non-monotonicity in linear regression approaches to analysis, in that a GAM returns fitted values of each (flexible) function of the independent variables. These fitted functions are what are plotted in the GAM figures below. However, this approach does not avoid the problem of sampling appropriately in regions of potential non-linearity or non-monotonicity. Consequently, I display only results arising from parameter sweeping in the text.

The text identifies five stages used in sequential parameter sweeping. The first is the parameter space spanned by $\{N, b_{mean}, b_{stdev}\}$. Siegel (2009) describes the analysis and theoretical support behind splitting this into three regions. Higher values of b_{mean} increase participation levels in all regions. In line with limit theorems, increasing N reduces randomness in aggregate behavior, decreasing participation when it is unlikely, and increasing it when it is likely. Increasing b_{stdev} increases participation as long as b_{mean} is not too high. The representative parameter triples used for the intermediate and strong motivation classes are respectively $\{1000, .6, .25\}$ and $\{1000, .6, .3\}$. Note that these particular numbers are used for purposes of visual representation only; results discussed in the text as applying to motivation classes apply to all sets of parameters that fall into a given class.

The second stage is networks. Relying on Siegel (2009), I characterize each of the four network types according to the regions of the parameter space over which the model behaves similarly under repression. For purposes of visualization I choose representative parameter values for the figures. "Higher connectivity" lines correspond to a *Connection Radius* of 15 for a Small-World network, or a *Village Size* of 25 for a Village/Clique network; "Lower connectivity" lines correspond to values of 5 and 5, respectively. The number of weak ties that is optimal depends on the connectivity parameter, so I list the values used in pairs. For "optimal" Small-World networks, with the *Rewire Probability* second: $(15, .14), (5, .3)$. For the "greater-than-optimal" Small-World network: $(15, .7)$. For "optimal" Village networks, with the *Far Probability* second: $(25, .004), (5, .003)$. For the "less-than-optimal" Village network: $(5, .001)$. "High Leader Influence" lines correspond to a γ of 1.4 in an Opinion-Leader network. Finally, lines in Figures 3 and 4 for the Hierarchy all have positive motivation correlation; "High Influence Followers" lines correspond to the parameter pair $(10, .007)$ for (*Expansion Rate, Level Connection*), while "Low Influence Followers" lines use the pair $(10, .002)$. The single uncorrelated Hierarchy displayed is used in Figure 5 for the Iraq example; here the parameter pair for both lines is $(10, .007)$, so as to ensure comparability.

The third stage is repression absent networks; the next section details the analysis that allows us to fix the updating rate, λ, at 0.8. The fourth stage is repression within networks. I maintain variation across two dimensions of repression throughout this analysis: the continuous strength of repression parameter, called the *removal rate* and the dichotomous parameter corresponding to repressive technology (random or targeted). As there are now several parameter regions under consideration, I supplement parameter sweeping with a GAM analysis in the fifth section of this appendix. The fifth stage is anger or fear. This produces another dimension to explore. I focus on looking at the effect of psychological responses over levels of repression that are more and less effective. This analysis is the most difficult from the perspective of the method; however, the outcomes from adding the psychological effects are uniform and clear.

3

4   Rates of Learning and of Repression

In this section I provide further support for the claim made on page 16 of the text, that it is safe not to vary the updating rate λ in subsequent analyses. The two graphs in Figure A1 illustrate the connection between the rate of learning, measured by λ, and the strength of repression. The strength of repression is represented by the rate at which individuals are physically removed from the network. The x-axes of each plot display the strength of repression; faster removal rates indicate stronger repression. The y-axes display the maximal participation rate, averaged over 1000 simulation runs. The plot on the left displays the intermediate motivation class, the plot on the right the strong one. The random repressive technology and Fully-Connected networks are used in both plots, as the logic described below is not specific to complex networks.

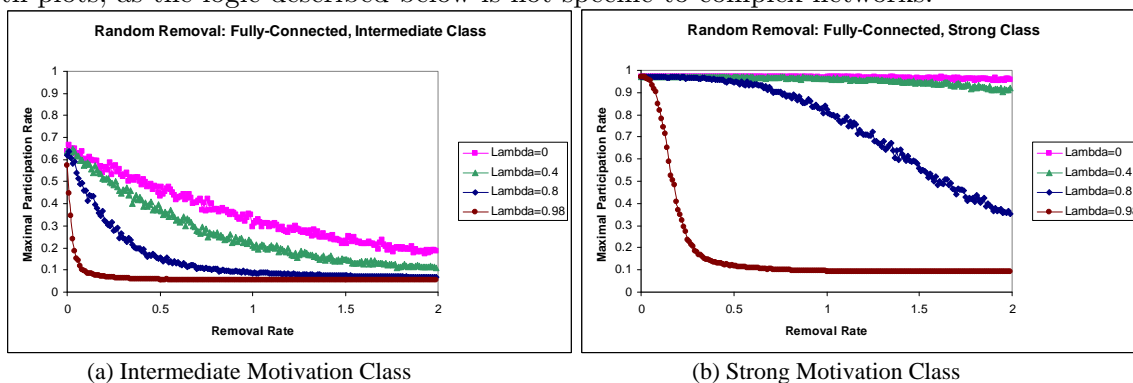(a) Intermediate Motivation Class   (b) Strong Motivation Class

Figure A1: Learning and Repression in the Fully-Connected Network

Figure A1 illustrates that when repression is very weak (near zero) or very strong (near two in the intermediate class; off the scale in the strong class) the learning rate matters relatively little. Absent repression λ has no effect on the maximal rate of participation, while sufficiently strong repression can quash any movement. In between, the more immediate updating is, the more robust the population is to repression. Slow updating enables repression because no one has time to internalize the actions of rabble-rousers. Fast updating limits repression as the base of observable people already participating rapidly becomes larger in each period. Updating and removal rates are thus important in relative terms; however, their importance in absolute terms is limited. We fix $\lambda = 0.8$, which allows exploration of a range of removal rates without removing too much of the population. Results hold for faster updating as well, with appropriate increases in removal rate.
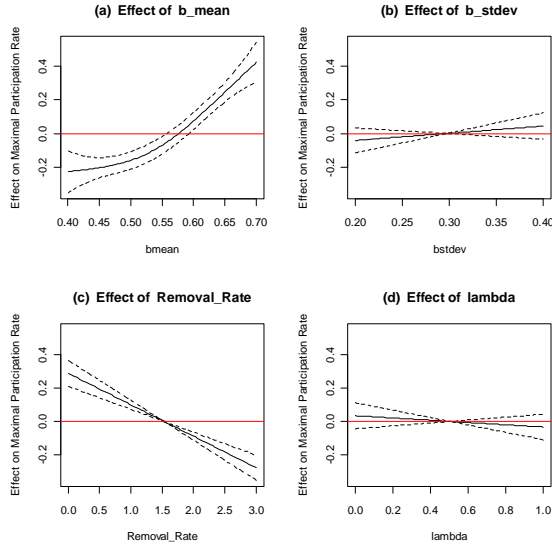
4

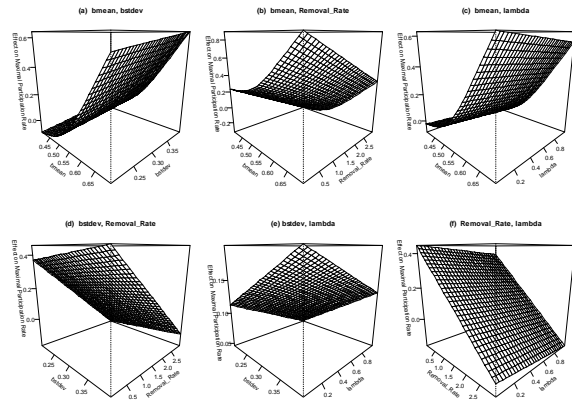Figure A2: Smooth Components of Fitted GAM (Fully-Connected)

Figure A3: 3D GAM Visualization (Fully-Connected)

Figures A2 and A3 provide additional support. They display graphically the outcome of a GAM regression of the following form: $maximum\,participation\,rate = \alpha + s_1(b_{mean}) + s_2(b_{stdev}) + s_3(removal\,rate) + s_4(\lambda)$. The GAM analysis utilizes spline smoothing with scale equal to 0.5, implemented via the `gam` command in R, part of the `mgcv` package. Figure A2 displays the smooth components of the fitted GAM. That is, each solid line plots the fitted functions s_i for each of the four independent variables, illustrating the effect of that parameter on the (average) maximal participation rate achieved. The dashed lines display 95% confidence intervals. We see first that all directional effects are as expected. Holding all other variables at their medians, increasing b_{mean} and b_{stdev} increase participation, though the latter effect was not significant at the relative low median for b_{mean} and high median for $removal\,rate$. Further, increasing $removal\,rate$ strongly decreases participation. λ seems to have little to no effect here, for the same reasons as b_{stdev}.[3]

Figure A3 shows the effect of pairs of parameters on the maximal participation rate. The parameter λ appears in (c), (e), and (f) in Figure A3. In Figure A3(f), for each level of $removal\,rate$, decreasing λ (making updating faster) increases participation, as found above. In (c) and (e), except for low levels of b_{mean} in which participation is virtually nonexistent anyway, decreasing λ shifts participation roughly linearly upward. Accordingly, its role is well understood, further justifying our no longer varying it. Note also that, in (a), increasing b_{stdev} increases participation for the levels of b_{mean} considered here.

5 Repression in Networks

The Effect of Strong Repression

---

[3]A less-smooth computer-chosen scale of .05 results in significance for all parameters. Specifically, high values of λ and low values of b_{stdev} significantly decrease participation. I display the smoother GAM solely because Figure A3 is easier to understand in this case.
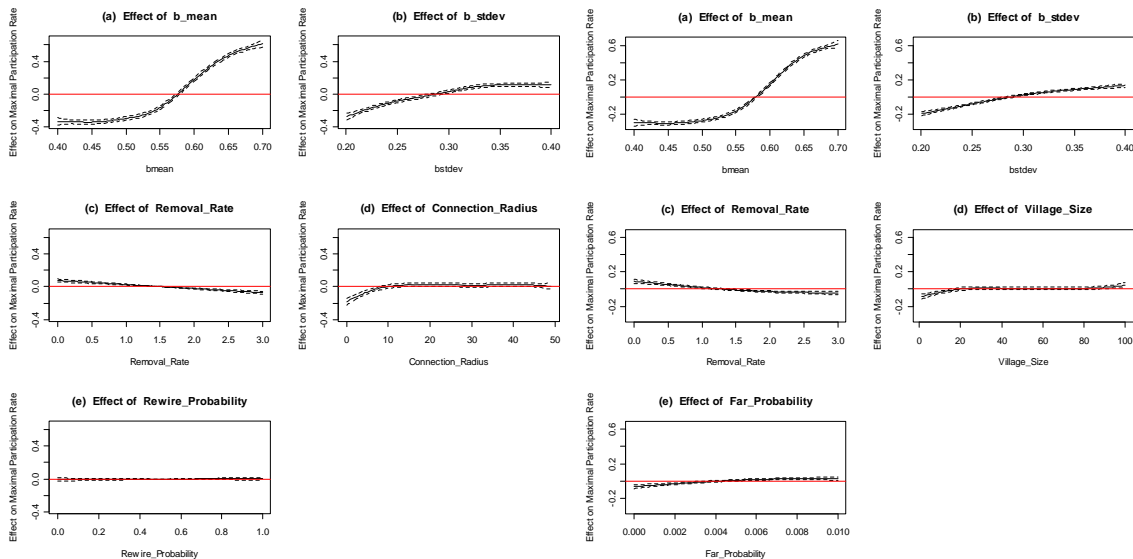
Figure A4: Smooth Components of Fitted GAM (Small-World)   Figure A5: Smooth Components of Fitted GAM (Village)

This subsection provides additional support for the claim (made on page 20 of the text) that sufficiently strong repression washes out differences between networks. GAM analyses indicate that across the range of the parameters sampled, all network parameters relevant to the Small World network, save *Rewire Probability*, and all network parameters for the Village network are significant predictors of the maximal participation rate, holding all other parameters at their medians. Figures A4 and A5 display the smooth components of the fitted GAM as in Figure A2.

Figure A4 covers Small World networks and illustrates that b_{mean} (a) and b_{stdev} (b) have positive significant effects, while *removal rate* (c) has a negative significant effect. Network parameters *Connection Radius* (d) and *Rewire Probability* (e) are less significant, owing to the other variables' being held at their medians. *Rewire Probability* in particular affects participation levels only over a subspace of the full parameter space, a fact captured much more easily in the results arising from parameter sweeping, and seen in Figure 3 in the text.

Parameters behave largely similarly in Figure A5, covering Village networks. The network parameter *Far Probability* has a more significant effect than the network parameter *Rewire Probability* of a Small World network due to the slightly greater role of weak ties in the Village network. Specifically, when weak ties between villages are very few, participation can't spread, so increasing *Far Probability* significantly raises participation at low levels of the parameter.

A GAM analysis (not shown) on a subset of the sampled simulation data using only points with strong repression (a *removal rate* between 2.75 and 3) has both motivational parameters remaining strongly significant, but only *Connection Radius* remaining significant out of the four network parameters, and even then it is significant only for low values. This is suggestive of the decreased role network parameters play as repression gets stronger, but it is an imperfect test. Better support can be found in Figure A6, which shows the mean number of individuals removed from the start to the end of participation (i.e., the total number removed until zero participation is again reached) for the parameter choices displayed in Figures 3a and 3b in the text. As is clear, while many people are removed over the full course of the movement, particularly in the strong motivation class, the number is decreasing for sufficiently quick removal.
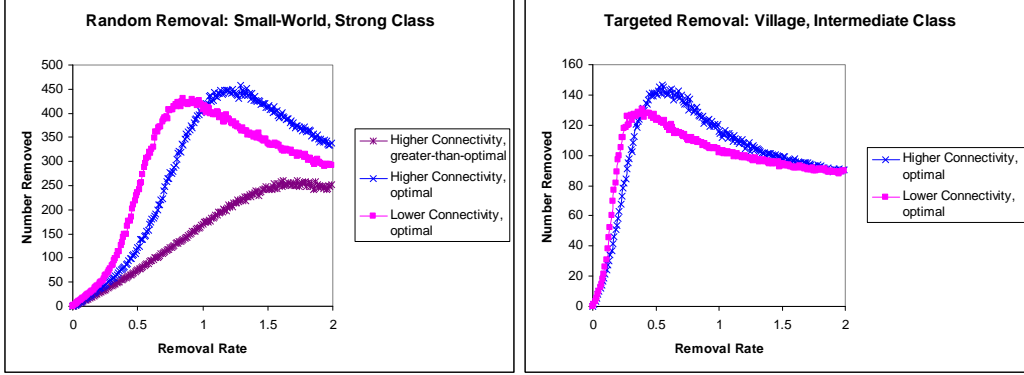
Figure A6: Total Number Removed in Figures 3a (left) and 3b (right)

Relative Importance of Network Parameters in Non-Elite Networks

Because the y-axes of the plots display partial effects, Figures A4 and A5 can also be used to display roughly the relative importance of the different parameters. We see that the absolute effects of varying the motivation parameters (in (a) and (b) of each figure) are greater than varying the network parameters (in (d) and (e) of each figure) under repression. This is in line with the claim made on page 23 of the text for Small World and Village networks, though of course the GAM results are only suggestive, since they hold only for other variables held at their medians.

Vulnerability of Elite Networks to Targeted Repression

Figures A7 and A8 display partial results of GAM analyses for the Opinion-Leader and Hierarchical Networks, respectively. All parameters proved strongly significant predictors of the level of maximal participation. As we see, targeted repression is more effective in each of these two network types, as seen in the negative and significant slopes of the lines in A7(g) and A8(g). Further, positive correlation of motivations leads to more participation (as seen in A7(e) and A8(e)), while negative correlation, on average, leads to less participation (as seen in A7(f) and A8(f)).
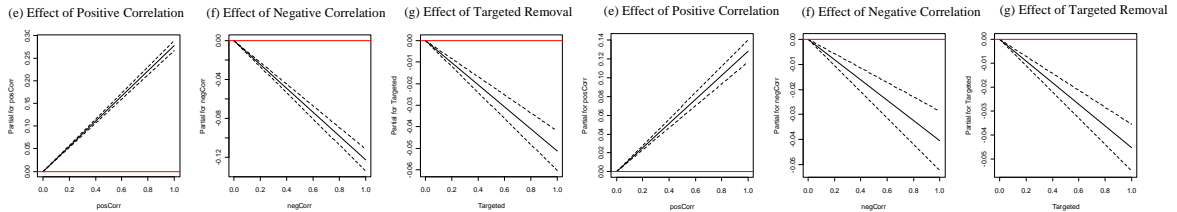

Figure A7: Linear Components of Fitted GAM (Opinion-Leader)   Figure A8: Linear Components of Fitted GAM (Hierarchy)

6   The January 2005 Iraqi Legislative Elections

To set up this section, consider the following quote from an article by Iraqi sociologist Faleh A Jabar, written shortly after the January 2005 Iraqi Legislative Elections:[4]

> Voters had expected worse. Early birds showed up before 9:00am to avoid attacks. The more audacious ventured out at what many suspected would be the most hazardous time. The bulk waited. Then, by midday, voters rightly guessed Salafi attackers had deployed all they had. The masses poured into voting stations, amazing themselves and the world.

---

[4]http://www.iraqitradeunions.org/wordpress/?p=192. Last accessed 12/16/10.

The logic behind this quote matches precisely that which motivated the model. A combination of a local insurgency and members of Al-Qaida in Iraq (jointly a repressive entity) attempted via largely violent means to deter voting in state elections (a collective action). Between the deadly violence,[5] and the risk of reprisals both immediately and in the future,[6] fear was prevalent. Rabble-rousers went out early to vote, but most people waited to see if they would be safe first, which required them to observe others they trusted returning from voting unharmed and unthreatened.

Though detailed networks of who trusted whom in Iraq during the time of the elections would be difficult to come by—and near-impossible before the elections themselves, when predictions of turnout would have been most useful—the model does not have such severe data requirements. The primary piece of information we need, as discussed in the previous section, is the type of network in place. For our purposes, understanding the sectarian nature of the salient cleavages during the January 2005 Iraqi elections will prove sufficient to make a novel prediction about differential levels of turnout between the followers of Grand Ayatollah Ali Sistani and Muqtada al-Sadr.

We begin with Iraq under Saddam Hussein. During this time the regime increased the level of danger inherent in widespread social ties through the use of a system of internal security agencies designed to stamp out any opposition to the state. As David Patel writes (2005; p. 3), "Every Iraqi today has stories about children denouncing parents, [and] neighbors settling local disputes through corrupt security officials." As a result, "the average Iraqi was intimately connected with and confided in only a few other Iraqis" and "[f]riendship networks, therefore, overlapped considerably, often consisting of a couple of brothers and cousins. These clusters of strong friendship networks were often unconnected to one another."

Thus, at the time of Saddam's fall Iraqi society was best represented by a Village network structure with small, sparsely interconnected groups. Even if individuals could observe the behavior of others in close spatial proximity, there were few trust bonds among this larger set, and so little reason to alter behavior based on the actions of anyone outside one's own small group. Participation is minimal in such cases, even absent repression, and so one should have expected little turnout by Sunnis, with or without a boycott.[7]

This accounts for Sunni networks, but ignores the Shi'ite religious hierarchies that had been in place during the Baathist era. Shi'ism within Iraq requires the emulation of figures of authority (*marja-e taqlid*), who are so designated due to their religious learning and perceived superior personal qualities. Individuals may freely choose which *marja' al-taqlid* (effectively Ayatollah) to follow, but unqualified people must subordinate their opinions to that of the *marja'*. Thus, with respect to each *marja'*, the influence networks of Shi'ites were much like star (small Opinion-Leader) networks, with the *marja'* at the center and the mass of self-selected followers at the points.

This accounts for some of the Shi'ite network structure, but not all; as Meir Litvak writes (p. 28):

> patronage networks linking the teacher and his former disciples who resided as 'ulama' in various Shi'i localities came close to the ideal type of radially connected network. In this model each member is directly linked to the central figure, i.e. the teacher, and members communicate with one another only through him.

---

[5] There were, for example, eight suicide attacks in Baghdad alone election day, killing 11 and leaving 47 wounded. Further, electricity and polling stations were being targeted in the run-up to the election, and insurgents in some areas promised to do the same to people who voted.

[6] Purple fingers, while an excellent symbol, were not likely to be helpful in avoiding recognition.

[7] From the turnout data provided below, one can estimate non-Kurdish Sunni turnout at around 15%.

Thus the upper echelon of the network is also a star network. An Ayatollah makes his desires known to his disciples and deputies, who pass these desires down to the great mass of society. Put together, we get the conclusion that "Each Ayatollah traditionally sits atop a pyramid-shaped patronage network and competes with others for students and followers." (Patel 2005; p. 5). This implies that each Ayatollah occupies the top spot in a Hierarchical network, with profoundly uninfluential followers due to the Baathist-era effect on civil society.

This covers the network type and level of influence of followers, but not the level of influence of the leaders, nor the correlation of leader motivations within Sistani's and Sadr's networks.[8] Sistani's network was well-established, being partially inherited from his prominent mentor Ayatollah al-Khoei. He controlled the upward mobility of his deputies, and thus the membership at each level of the hierarchy. Accordingly, it is safe to assume that leaders were influential in his network, and further that motivations within the network were positively correlated with position for any collective behavior he supported, since he could ensure that those at the top of the hierarchy at least outwardly expressed his favored behavior.

In contrast, Sadr's position was far more tenuous, deriving largely from previous emulation of his respected late father, Grand Ayatollah Muhammad Sadiq al-Sadr. His clerical rank was low, and it was unlikely that he had completed the study necessary to be an object of emulation. While he had been able to replace a few heads of local offices, "tremendous variation [was] observed in both rhetoric and action across Sadr offices." (Patel 2005; p. 11). Though the network structure he inherited was likely strong, giving leaders substantial influence, one could not consider leaders within the network as having positively correlated motivations over all collective actions. In cases where there was already wide agreement across Sadr's entire network, such as ending foreign occupation of Iraq,[9] this assumption might have been accurate, and indeed in these instances Sadr drew large numbers to protest. However, in the context of the elections Sadr's message was ambivalent and his followers divergent in intent,[10] and so we must treat leader motivations as being uncorrelated with position.

Next we consider repression strength and repression technology. Repression on election day was effectively random in nature, particularly with respect to Shi'ites. Importantly, no credible threat emerged to Sistani's life should people have turned out to vote. Beyond being not-insubstantial, the strength of repression is unclear; thus we vary it explicitly.

Last, there is motivation class, and the potential for anger or fear. To keep things simple, I assume no emotional response, given no data to support it. As for motivation class, while it is difficult to separate the distributions of internal motivations from actions taken under the influence of leaders, there seems little reason to assume that any particular group of Shi'a was inherently more motivated to participate than any other. Further, all had much to gain from representation.

---

[8]The argument for the relative strengths of Sistani's and Sadr's networks is taken from Patel (2005). Note that both leaders' networks—Sadr's in particular—have changed significantly since January 2005; this analysis is only intended to be valid around the time of the election, which is the focus of this section. That Sadr has been able to mobilize increasing support as he has consolidated control over his network, however, is unsurprising given the analysis presented here.

[9]A Zogby poll prior to the election (`http://zogby.com/news/ReadNews.cfm?ID=957`. Last accessed 12/16/10) found 69% of all Shi'ites favored "U.S. forces withdrawing either immediately or after an elected government is in place." The percentage among Sadr's following, largely marginalized within the transition government, was likely higher.

[10]From Juan Cole (`http://www.juancole.com/2005/01/allawi-pockets-will-not-be-able-to.html`. Last accessed 12/16/10): "Sadr himself has given such mixed signals that it would be hard for [his followers in Sadr City] to follow him if they wanted to. First he said he was neutral about the elections, then more recently that he opposes them. Some of his chief lieutenants have called for a boycott (Shaikh Bahadili in Basra), while others are actually standing for election [in the National Independent Cadres and Elites party]."

Accordingly, an assumption that both groups of Shi'a were either in the intermediate or in the strong motivation class seems reasonable.[11] Both assumptions yield similar results, though I display only the intermediate case here.

This completes the application of the model; note that all information used was available prior to the election. Figure A9 displays expectations derived from the model. As we can see, the model predicts that mobilization in favor of *any* of the goals of Sadr or his followers—voting for the National Independent Cadres and Elites party or not voting at all—would be less than that for Sistani's objectives—voting (which he made a religious obligation) and voting for the UIA—at any level of repression. Further, and more importantly given the measurable repression of the day, Sistani's support is predicted to be far more robust to repression than Sadr's. Thus a level of repression that barely affected Sistani's ability to mobilize his followers would completely crush Sadr's ability to do the same.
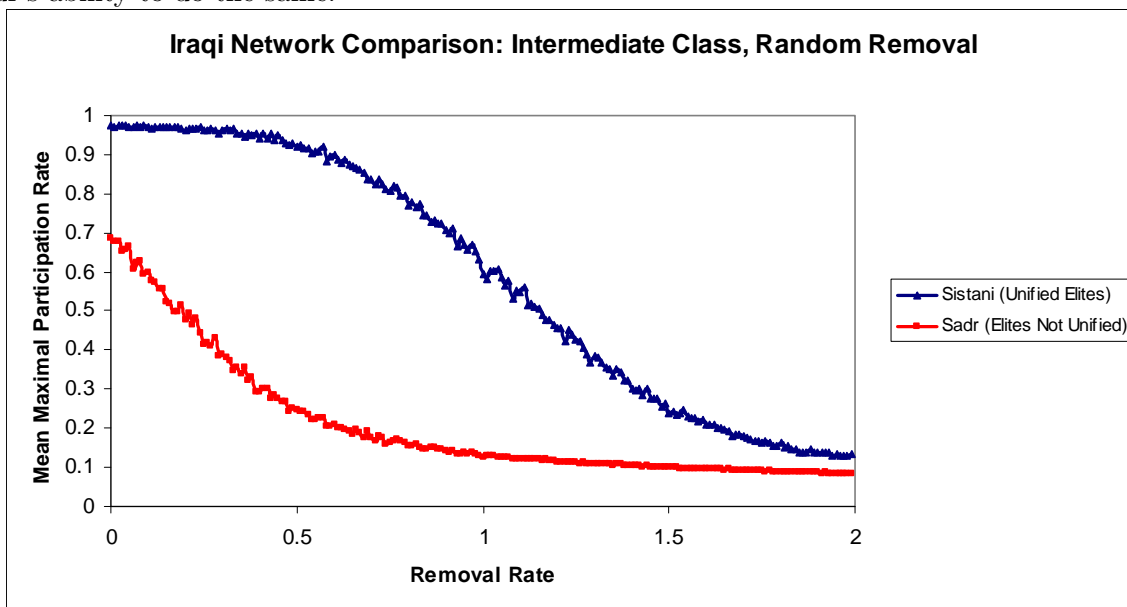

Iraqi Network Comparison: Intermediate Class, Random Removal

Legend:
- Sistani (Unified Elites)
- Sadr (Elites Not Unified)

(y-axis: Mean Maximal Participation Rate; x-axis: Removal Rate)

Figure A9: Turnout within Iraqi Social Subgroups

These predictions are in line with what occurred election day.[12] Turnout in seven predominantly Shi'ite southern provinces was approximately 70%, less than the 80% a prior Zogby poll had predicted,[13] but still quite high. Sistani's UIA was the clear beneficiary of this. In Baghdad, of which a third of the population resides in Sadr's locus of power, Sadr City, National Independent Cadres and Elites barely outdrew the Sunni party, The Iraqis, while the UIA did quite well.[14] Thus we see that even rough approximations of the local networks in play, along with an equally rough characterization of the nature of the likely repression, was sufficient in this case to predict the relative influences Sistani and Sadr would have on electoral outcomes.

---

[11] This is further supported by accounts noting that the residents of Sadr City, in which Sadr's base of support lay, were enthusiastic to vote. (http://www.abc.net.au/pm/content/2005/s1290550.htm. Last accessed 12/16/10.)

[12] http://aceproject.org/ero-en/topics/electoral-participation/turnout/updatedelectionresults.pdf/view. Last accessed 12/16/10.

[13] http://zogby.com/news/ReadNews.cfm?ID=957. Last accessed 12/16/10.

[14] While it is somewhat difficult to find a clear breakdown on turnout below the province level, given the overall low turnout of Sunnis, the Shi'ite makeup of Sadr City, the strong showing for the UIA in Baghdad, and the fact that Sadr City makes up about a third of Baghdad's population, it is safe to say that Sadr City's population provided substantial electoral support to the UIA, and so did not substantially boycott.

References

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

Litvak, Meir. 1998. *Shi'i Scholars of Nineteenth-Century Iraq.* New York: Cambridge University Press.

Patel, David Siddhartha. 2005. "Ayatollahs on the Pareto frontier: The Institutional Bases of Shiite Religious Authority in Iraq. *Working Paper.*

Yin, Chien-Chung. 1998. "Equilibria of Collective Action in Different Distributions of Protest Thresholds." *Public Choice* 97(4): 535-567.