2013

An Agent-Based Model Of Centralized Institutions, Social Network Technology, and Revolution

Michael D. Makowsky
Johns Hopkins University

Jared Rubin
Chapman University, jrubin@chapman.edu

Recommended Citation

Makowsky, Michael D., and Jared Rubin. "An Agent-Based Model of Centralized Institutions, Social Network Technology, and Revolution." *PloS one* 8.11 (2013): e80380.
DOI:10.1371/journal.pone.0080380

An Agent-Based Model of Centralized Institutions, Social Network Technology, and Revolution

Michael D. Makowsky*
Department of Economics
Towson University

Jared Rubin
Department of Economics
Chapman University

October 2011

Abstract

Recent uprisings in the Arab world consist of individuals revealing vastly different preferences than were expressed prior to the uprisings. This paper sheds light on the general mechanisms underlying large-scale social and institutional change. We employ an agent-based model to test the impact of authority centralization and social network technology on preference revelation and falsification, social protest, and institutional change. We find that the amount of social and institutional change is decreasing with authority centralization in simulations with low network range but is increasing with authority centralization in simulations with greater network range. The relationship between institutional change and social shocks is not linear, but rather is characterized by sharp discontinuities. The threshold at which a shock can "tip" a system towards institutional change is decreasing with the geographic reach of citizen social networks. Farther reaching social networks reduce the robustness and resilience of central authorities to change. This helps explain why highly centralized regimes frequently attempt to restrict information flows via the media and Internet. More generally, our results highlight the role that information and communication technology can play in triggering cascades of preference revelation and revolutionary activity in varying institutional regimes.

JEL Codes: C63, Z13, D83, D85, D71, H11
Keywords: preference falsification, revolution, protest, network technology, agent-based model

* We'd like to thank Robert Axtell, David Levy, and participants in a workshop at Baylor University for helpful comments.

Introduction

Recent uprisings in Egypt, Libya, Tunisia, and other parts of the Arab world have come quite unexpectedly to most observers. Although the seeds of discontent had been sown for decades in these countries, public anti-government displays barely existed. Such rapid changes in publicly displayed preferences are not a new phenomenon; precedents include the fall of Communism in the Eastern bloc, the end of apartheid in South Africa, and the civil rights movement (Kuran 1989; Kuran 1991a; Kuran 1991b; Lohmann 1994; Kuran 1995a; Kuran 1995b; Wright 1999).[1]

We argue that economies containing two features – highly centralized power and widespread information and communication technology (ICT) – are conducive to massive and rapid preference revelation. We define power centralization as the ability of one actor to impose multiple sanctions on individuals. Examples include national and localized sanctions in autocracies, economic and religious sanctions in theocracies (such as Iran), or political and legal sanctions against dissidents (as in Burma).[2] The ability of central authorities to impose sanctions on individuals, coupled with heterogeneous citizens whose true preferences are hidden, can calcify a society – leaving it stuck at sub-optimal equilibria despite changes to individual preferences. These sub-optimal equilibria can be escaped in a cascade of sudden preference

[1] For more on the mechanisms underlying rapid changes in publicly displayed preferences, see (Granovetter 1978; Schelling 1978; Oliver, Marwell et al. 1985; Macy 1991; Banerjee 1992; Bikhchandani, Hirshleifer et al. 1992; Young 1993; Lohmann 1994; Centola, Willer et al. 2005; Bicchieri 2006; Siegel 2009; Willer, Kuwabara et al. 2009; Ellis and Fender 2011; Rubin 2011).

[2] Goldstone (2011), in an article on the Arab Spring, notes that "Sultanistic governments" are particularly susceptible to revolutions. His definition of Sultanistic is very similar to our definition of centralization (italics ours): "Such governments arise when a national leader expands his personal power at the expense of formal institutions. Sultanistic dictators appeal to no ideology and have no purpose other than maintaining their personal authority. They may preserve some of the formal aspects of democracy--elections, political parties, a national assembly, or a constitution--but they rule above them *by installing compliant supporters in key positions* … Behind the scenes, such dictators generally amass great wealth, which they use to *buy the loyalty of supporters and punish opponents*. … The new sultans control their countries' military elites by keeping them divided. Typically, the security forces are separated into several commands (army, air force, police, intelligence)--*each of which reports directly to the leader. The leader monopolizes contact between the commands*, between the military and civilians, and with foreign governments*, a practice that makes sultans essential for both coordinating the security forces and channeling foreign aid and investment*."

revelation initiated by a shock which encourages some individuals to publicly reveal their preferences, which in turn alters norms and triggers a cascade.

Centralization can encourage individuals to publicly lie about their privately-held preferences because those who transgress centralized authorities incur sanctions over numerous dimensions. For example, if one breaks religious dictates in Iran, they may suffer consequences in the afterlife as well as economic consequences in the present. Such societies are prone to cascades of preference revelation if preferences are inter-connected; that is, if individuals derive utility from conforming to the actions of others (Granovetter 1978; Kuran 1995). A cascade can occur after a shock encourages some to reveal their privately-held preferences, which encourages others to do so, and so on. We argue that ICT facilitates this process; in order for the cascade mechanism to occur, people have to know how others are acting. Indeed, we find that cascades of revolutionary activity are more likely to occur in centralized regimes, but *only* when networks are large. In small network societies, shocks are less likely to trigger a cascade when power is highly centralized. Further, the thresholds at which systems tip towards institutional change are reduced by increasing the range over which agents can add to their social networks. This helps explain why highly centralized regimes (e.g., Libya, China, Burma, North Korea) frequently attempt to restrict information flows; when their citizens are weakly connected, the probability of a revolutionary cascade arising decreases.

The popular notion that innovations in ICT are helping to facilitate the social and institutional changes we are witnessing in real time (Shirky 2011) is not without its detractors (Morozov 2011). Whether or not these technologies are party to the *initiation* of revolutions, what is less understood is the capacity for these innovations, such as Facebook and Twitter, to build upon social change - both helping to *propagate* a cascade of preference revelation and support the public protest that emerges from it and *facilitating its translation* into institutional change. We argue that widespread ICT supports efforts to challenge authority by encouraging the public revelation of preferences. We model such actions in an agent-based framework to provide a better understanding of the mechanisms connecting political institutions, ICT, and revolutions.

Revolution and the Agent-Based Framework

The basic model linking institutional centralization and rapid, revolutionary change was proposed in Rubin (2011). It consists of heterogeneous agents who face costs when their actions differ from i) their internal preferences (or "bliss points"), ii) an endogenous social norm, iii) a "central authority", and iv) a "non-central" authority. Both authorities face costs from diverging from the citizenry, and the central authority can impose a cost on the non-central authority. The degree of centralization is increasing in the latter cost. We model centralization in this manner to highlight the idea that centralized power works through institutional conduits. For example, the religious hierarchy in Iran has power to impose political sanctions because the leading political authorities face significant costs from disobeying their dictates. Likewise, most autocrats (such as in pre-revolution Egypt) impose multifarious sanctions through the military. In such a regime, the military is the "non-central" authority and the autocrat's degree of centralization hinges on how costly military authorities view choosing actions which defy the autocrat.

The model suggests that citizens falsify their preferences in favor of the central authority (i.e., make public expressions different from their internal preferences) when centralization is greater since they face multiple costs from transgression. Preference falsification can unravel, however, when a widespread shock alters the costs citizens face. If the shock is large enough, some citizens reveal their preferences, which alters the social norm, which itself encourages more citizens to reveal their preferences. A cascade can result, entailing a vastly different equilibrium of public expression. We refer to the difference between the actions of citizens and authorities as “protest” and the change in the pre-shock and post-shock actions of authorities as “institutional revolution”.

A cascade of preference revelation is dependent on the means of social transmission. Social norms change only as people are made aware that the modal behavior in their social network is changing. Network structure has been shown to be relevant to the way in which behavior spreads through populations in game theoretic proofs (Allen and Gale 2000; Lee and Valentinyi 2000; Morris 2000), network theory (Golub and Jackson 2011), computational simulations (Bonabeau 2002; Epstein 2002; Cowan and Jonard 2004; Delre, Jager et al. 2007), and social experiments (Centola 2010; Fowler and Christakis 2010).

We test the interactions between ICT, institutional centralization, and revolutionary activity with an agent-based model (ABM). Within our ABM we construct a population of autonomous, heterogeneous citizens whose rules-based decisions depend on, and in turn influence, the decisions of both their fellow citizens and the authorities that govern their artificial world. Agents occupy unique, randomly assigned spaces on a two-dimensional lattice, interacting with the members of their directed social network. We execute model simulations initializing the model and spinning it forward over discrete time steps. Macroscopic social patterns emerge from the interacting decisions made by agents over the course of a simulation (Epstein and Axtell 1996; Epstein 2006). We conduct experiments exploring the sources of these patterns, running the model tens of thousands of times over a variety of model parameterizations.

The Model

The model is a repeated game played over *T* discrete steps in which *M* citizens engage in a game with a central authority (C) and a non-central authority (N). The central authority moves first, then the non-central authority, and finally the citizens. Within the set of citizens, the order of activation is randomized at each time step. The two authorities choose an action that maximizes their utility function, based in part on the average citizen action from the previous time step.

The model is constructed on a 40 by 40 lattice with associated directed network graph.[3] Agent social networks are subsets of selected agents from within their social radius *r*, which is the "Moore Neighborhood" of radius *r*. The Moore neighborhood is the square of surround cells on the lattice. The surveying agent is not included in his own neighborhood. The lattice is torus shaped and wraps at the edges, preventing edge effects. See Figure 1 below.

Figure 1: Citizen social radius on the lattice and network graph.

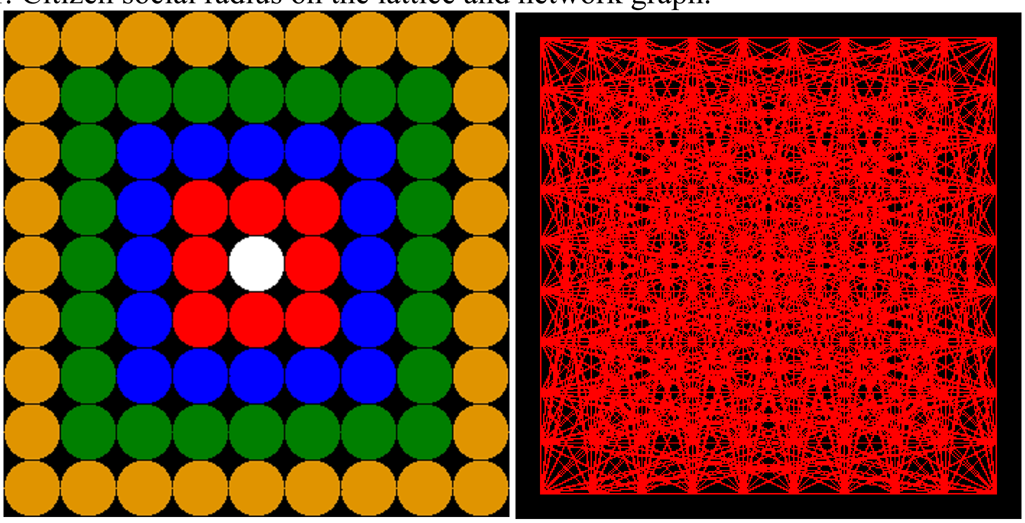

r = 1(red agents), r = 2 (red, blue), r = 3 (red, blue,green), r = 4 (red, blue, grean, orange)
The full lattice is 40 by 40 with 1600 agents.

Step (t = 0) Model Initialization. The model creates and places agents randomly, one per lattice coordinates (x,y). Agents are heterogeneous across bliss point and are given random values from a normal distribution. All agents are initialized with actions equal to their bliss points. Citizens

[3] The model program was written using the MASON simulation Java library (Luke, Cioffi-Revilla et al. 2005).

exist on a two-dimensional toroidal lattice. Their social network exists as a directed graph. Citizens actively form connections by choosing to connect to the *n* agents within radius *r* on the lattice whose actions a_{-i} are closest to their own intrinsic bliss point b_i.

Table 1 indicates the order of action. The model analyzes situations in which the preferences of some citizens differ exogenously from those of the authorities, so actions could represent varying levels of freedom of speech, press, or religion, publicly expressed dissatisfaction with the government or religious authorities, or public opinion on social issues.

Table 1: Order of action (within step)

1. Central Authority($\bar{a}_{citizens,t-1}$)
2. Non-central Authority($a_{central,t}$, $\bar{a}_{citizens,t-1}$)
3. Citizens($a_{central,t}$, $a_{noncentral,t}$, $\bar{a}^{\Omega}_{citizens,t}$)†

†$\bar{a}^{\Omega}_{citizens,t}$ is the mean of the most recent actions taken by agents in the acting agent's social network. It includes a mixture of agents whose most recent actions were taken in the current step and agents whose most recent action was taken in the previous time step.

Citizens face three costs. Two of these costs are a function of the distance between the citizen's action (a_j) and the actions of the two authorities, (a^N and a^C). These costs are increasing in the size of the violation and represent the costs (or punishments) associated with breaking a religious dictate, breaking a law, violating a political norm, and the like.

The third cost is a function of the distance between their action ($a_{j,t}$) and the average action of other citizens within their social network ($\bar{a}^{\Omega}_{citizens,t}$). This norm is a property of the system that emerges from the interacting decisions of all of the agents. Each citizen j maximizes the following utility function in each period:

(1) $$U_{j,t} = -w_1\left(a_{j,t} - b_j\right)^2 - w_2\left(a_{j,t} - \bar{a}^{\Omega}_{citizens,t}\right)^2 - w_3\left(a_{j,t} - a^N_t\right)^2 - w_4\left(a_{j,t} - a^C_t\right)^2,$$

where w_k is a weighting parameter for k={1,2,3,4}. The utility function is concave, containing only a global maximum. This allows us to employ the relatively simple golden mean search optimization algorithm (Press 2002) in all agent utility maximization.

We define centralization as the weight (γ) that the non-central authority places on conforming to the central authority's dictates. The greater this weight is, the more influence the central authority has over the citizenry since they face multiple costs from transgressing the central authority. The central and non-central authorities have bliss points b^C and b^N and maximize the following utility functions in each period:

(2) $$U_{j,t}^{C} = -w_1^C(a_t^C - b^C)^2 - w_2^C(a_t^C - \bar{a}_{t-1})^2$$

(3) $$U_{j,t}^{N} = -w_1^N(a_t^N - b^N)^2 - w_2^N(a_t^N - \bar{a}_{t-1})^2 - \gamma(a_t^N - a_t^C)^2$$

The fixed parameters employed in the ABM, which are constant across all model realizations, are reported in Table 2.

Table 2: Model parameters

Parameter	Context/Related Function	Value
M	# Citizens	1600
R	Social radius	{1,2,3,4}
γ	Centralization	{0, 0.25,…4}
S	Shock fraction	{10%, 20%…100}
N	Social network size	8
	All other utility function weights	0.5
t_{shock}	Time step for social shock	$t = 20$
Bliss points	Central, Noncentral, mean citizen	{0.0, 0.5, 1.0}

*The lattice is a 40 by 40 torus with 100% agent density and no overlapping agents. Each combination of run parameters was simulated 50 times (n = 42,000)

Citizens first choose their social network and then choose their action. The choice of social network involves a survey of all of the citizens, Θ_i, within their social radius, r, such that $\Theta_i \subset M$. Citizens fill their network using a homophilous selection mechanism (McPherson, Smith-Lovin et al. 2001; Macy and Willer 2002; Golub and Jackson 2011), ranking other citizens within their neighborhood and forming connections to the n other citizens, $\Omega_i \subset \Theta$, whose most recent actions are the closest to their personal bliss point.

$$
(4) \qquad
\begin{aligned}
&\Theta_j = \left\{ j_1 , j_2 \ldots j_n , \ldots j_K \right\} \\
&\Omega = \left\{ j_1 , j_2 \ldots j_n \right\} \\
&\overline{a}^{\Omega}_{citizens} = \frac{1}{n} \sum_{j=1}^{n} a_j
\end{aligned}
$$

Agent social networks are governed by two factors: the portion of the lattice over which an agent may search for agents to add to her social network and the total number of agents they will choose to include in their network. The portion of the lattice they will search is a function of their location on the lattice and the network radius parameter, *r*. In the two dimensional lattice of the model, the set Θ_i, will include $(2r+1)^2 - 1$ agents. The number of agents chosen in each social network, *n*, is a fixed exogenous parameter in the simulation experiments conducted.

Step t = 20

At t = 20, a shock hits a portion of the citizenry. This shock increases the weight that they place on their bliss point vis-à-vis the social and institutional costs. The shock represents a change in the relative weights associated with acting in favors of one's internal preferences; whether this happens through a rise in one's weight on their intrinsic preference or a fall in the cost of sanctions will yield the same results. In the case of Egypt in 2011, this shock can be interpreted as a portion of the population viewing the revolutionary events occurring in Tunisia and deciding that fighting the regime is more worthwhile relative to costs than it had been prior to the events in Tunisia.

Experiment Results

The model was simulated over a range of parameterizations varying the levels of centralization, agent network radius,[4] and the percentage of the population affected by shock. Each parameter combination was simulated 50 times, creating a total sample of 42,000 simulations.

Preference Falsification

Prior to the "Arab spring" of 2011, it was extremely costly for most citizens of Egypt, Libya, Tunisia, Syria, Bahrain, and a host of other nations to express anti-government preferences. The most prominent cost one could face was direct reprisal from the government. Detention without trial, partial jurisprudence, and beatings of dissenters were common enough to discourage such outward expression. Worse still, it was dangerous to express anti-government opinions even to seemingly close relations. In a state where most people publicly expressed favor for the government, it was difficult to discern who *actually* favored the government and who was *pretending* to favor the government. It is clear *ex post* that preference falsification – defined here as the difference between one's expressed and intrinsic beliefs (Kuran 1987; Kuran 1995) – was rampant in the Arab world prior to the 2011 revolts.

The primary path through which centralization of sanctioning ability affects revolution is through preference falsification. In highly centralized societies, citizens are more likely to falsify their preferences since the sanctions from expressing anti-government views are greater. This means, however, that the expressed preferences are more likely to unravel in a cascade after the shock, since the changing of the social norm encourages some to express actions closer to their internal preferences, which in turn encourages more to act close to their bliss point, and so on.

The role that social networking technology played in the Arab revolts has been widely reported. Some Western media outlets dubbed these movements the "Twitter revolutions" for the role that social networking played in mobilizing and coordinating the protests. Social radius is

[4] For an analysis of agent network radius using basic network statistics, see Appendix A.

also related to preference falsification. When individuals have fewer people with whom they can connect, they are less likely to run into people that have similar views, and thus the social norm that they follow is less likely to resemble their own preferences. Indeed, the media narrative surrounding the Arab revolts indicated that Twitter and Facebook were important in part for merely letting people know that others had similar views towards the government.

Our model confirms these relationships between preference falsification, social radius and centralization. Figure 2 shows the degree of preference falsification over varying degrees of centralization and social radius. It reveals that preference falsification is increasing in centralization and decreasing in social radius. Darker areas indicate greater preference falsification.

Figure 2: Preference falsification prior to the shock (t=20)

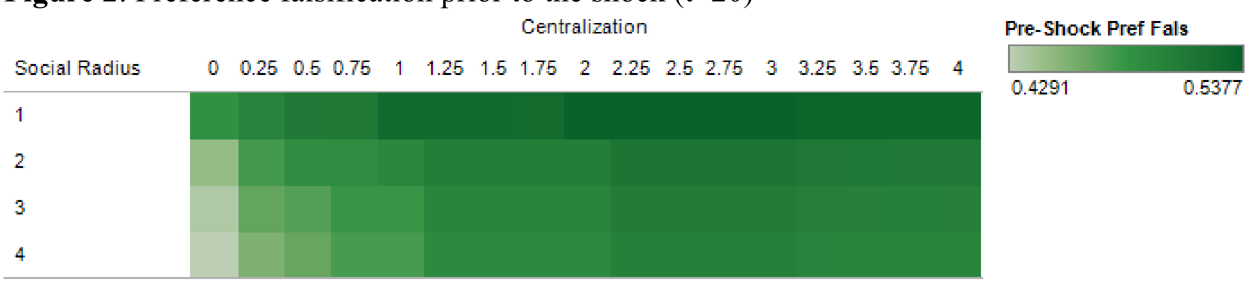

Consider next the relationship between social radius and preference revelation cascades. Preferences are falsified to a greater degree in low radius societies, suggesting a higher probability of a cascade. However, information travels faster in high-radius societies, which increases the probability of a cascade.

Figure 3 (upper) maps the difference in pre- and post-shock preference falsification over three dimensions: centralization, social radius, and the fraction of citizens affected by the shock. Blue areas indicate less difference in preference falsification (with the darkest blue being the least), red areas indicate greater difference in preference falsification (with the darkest red being the most) and larger squares within each panel indicate a greater social radius. The lower portion of Figure 3 charts the change in preference falsification over centralization, using only observations from middle range sized shocks (40% to 60% of the population).

Figure 3: (upper region) Average difference in preference falsification before (t=20) and after (t=40) the shock; (lower region) Average difference in preference falsification over centralization, shock fraction 0.4 – 0.6 only

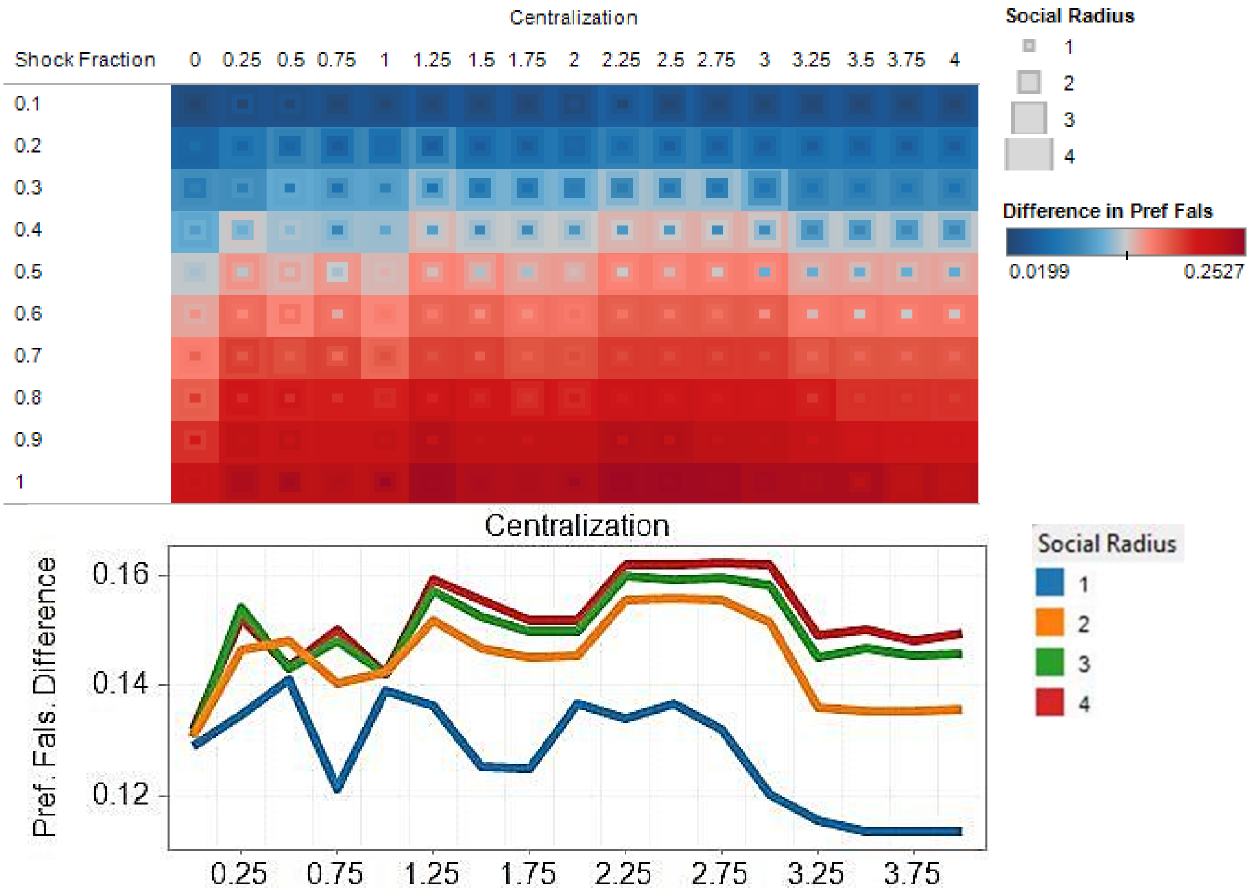

Several results are apparent in Figure 3. First, the impact of a shock on preference falsification is considerably larger when social radius is larger. All else being equal, the change in preference falsification is increasing with social radius. Second, the differences between the biggest and smallest social radii are largest in the "middling" shocks (40%-60%). When shocks are small (≤ 30%), not enough of the population is hit by the shock to start a cascade regardless of social radius. When shocks are systemic (≥ 70%), enough people are affected by the shock that the preference revelation mechanism which works through networks is not needed to transmit the shock and thus social radius matters less. In other words, social media are important devices for instigating preference revelation cascades, but *only* when the shock is not too large or too small. Figure 3 also reveals how the effect of centralization on preference falsification is affected by the social radius, especially in the "middling shock" ranges. Figure 3 (lower) suggests that at the smallest social radius, "middling" shocks do not reach enough of the population to affect actions through the changing of norms. Hence, those that do not have the shock transmitted to them will not change their actions much in highly centralized regimes, as the costs of transgressing the central authority are high. This result again highlights the importance of social media in the unraveling of preferences. Without the Internet, Facebook, Twitter, or other forms of mass communication outside the hands of the government, it takes a *really* large shock for a cascade to emerge in a centralized regime. Social media dramatically reduces the threshold level of the shock needed for a cascade.

Homophilous selection within populations with greater network range allows agents to create personal networks that are more conducive to behavior that deviates from the preferences of authorities. The limitations on homophilous selection in small-network societies encourage preference falsification, since there is limited spread of citizens' true preferences throughout the population. With larger social radii, however, there is an inverted-U shaped relationship between centralization and change in preference falsification. In highly decentralized economies, preference falsification unravels little because there was little of it in the first place. Meanwhile, preferences in highly centralized economies unravel little because the penalties for transgressing the central authority are severe enough that some do not join the cascade even after one has emerged. In the middle values of centralization, these two conflicting phenomena do not cancel each other out and larger cascades emerge. Preference falsification prior to the shock is great enough in the middle range where a cascade may occur, but centralization is not so great where the downside to joining a cascade discourages preference revelation.

Preference falsification can lead to a revolution only when a cascade of preference revelation occurs. This is more likely to happen when two conditions are met: i) preferences are highly falsified prior to the shock, and ii) network radius is large. Hence, cascades should occur more frequently in highly centralized societies, ceteris paribus, since preferences are falsified to a greater degree. It is possible, however, for an economy to be “too centralized” for a cascade to arise. If the sanctions from transgressing the central authority are too severe, too few agents will reveal their preferences following the shock and no cascade ensues.

Protest

In this section, we analyze the conditions which encourage protest, which we define as the difference between the average citizens’ action and the actions of the central authority. In other words, our protest measure indicates how far citizens are willing to openly transgress the central authority’s dictates.

Figure 4 shows the average level of protest prior to the shock. Not surprisingly, protest is decreasing in centralization and increasing in social radius. One would expect much more anti-government protest in an unpopular decentralized regime than in an unpopular centralized (autocratic) regime. Likewise, one would expect more public protest in high-radius societies than in low-radius societies.

Figure 4: Average protest prior to the shock (t=20)

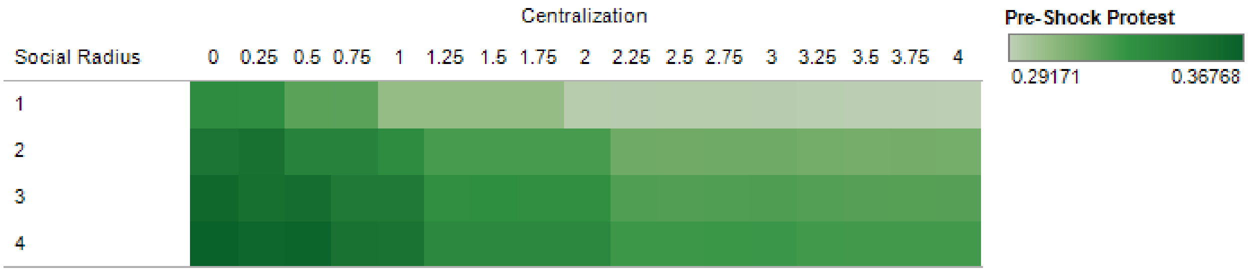

This figure provides a baseline for our analysis of protest. We are concerned with both the average amount of protest as well as change in average protest level after the shock. Both are relevant to the conditions that can precipitate rapid, revolutionary activity.

Figure 5 maps post-shock protest over centralization, shock size, and social radius. The comparative statics are straight-forward: protest is increasing in social radius and the portion of the population affected by the shock, and it is decreasing in centralization. This is what we expect given the intuition espoused above. Yet, the level of protest after the shock does not tell the entire story. If citizens were already in the street protesting against the central authority prior to the shock, their presence after the shock may not be nearly as revolutionary as the case where there was no protest prior to the shock but some protest after. For this reason, we are also concerned with the change in protest before and after the shock, which is mapped in Figure 6.[5]

[5] We ran similar tests for the top 5% of the bliss point distribution. These citizens are the ones who initiate the cascade – the "vocal minority" who are the first to change their actions after the shock. The results do not differ much from the ones reported above, except for the trivial result that it takes a smaller shock to encourage this group to change their level of protest. These results are available upon request.

Figure 5: (upper region) Average protest after the shock (t=40); (lower region) Average protest over centralization, shock fraction 0.4 – 0.6 only

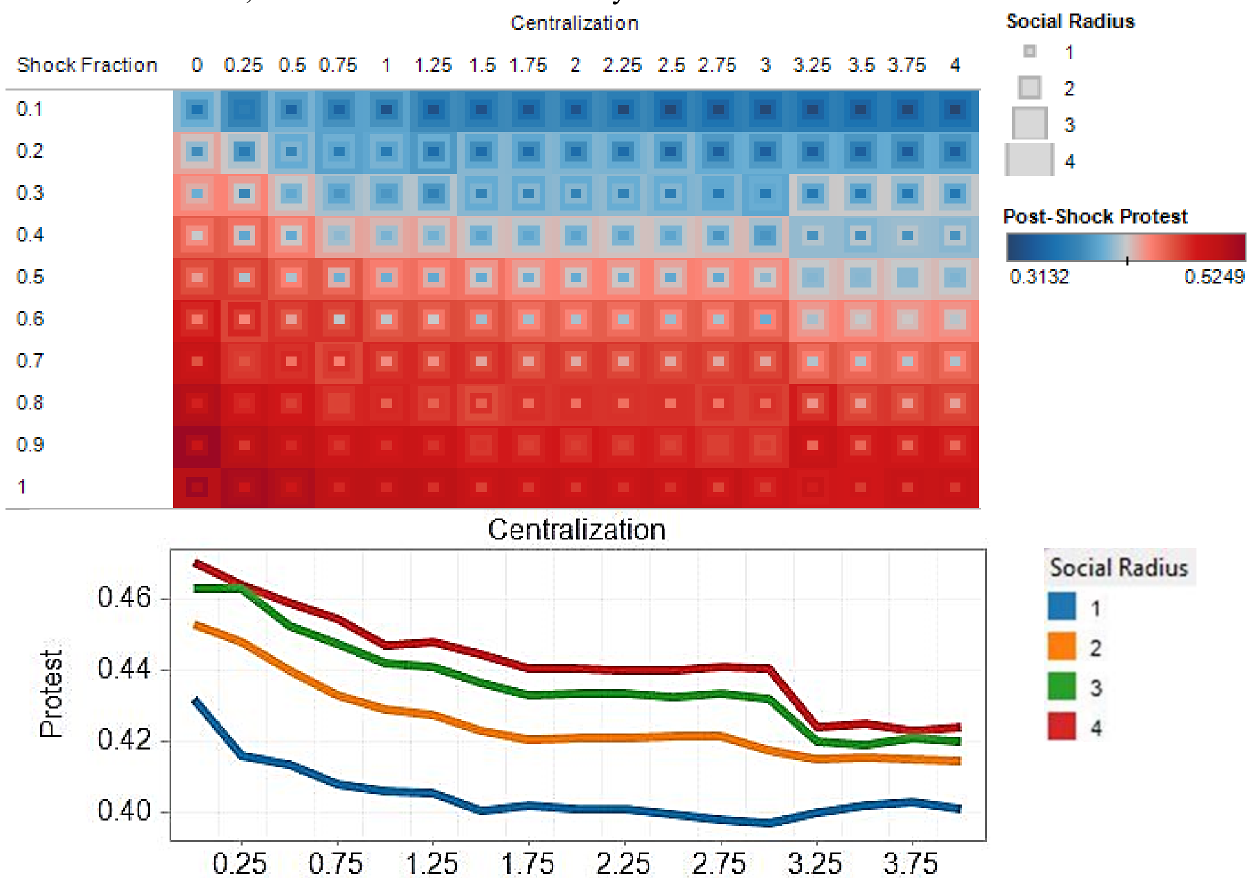

Two key results are apparent in Figure 6. First, the relationship between centralization and the change in protest takes on an inverted-U shape in large network-range economies (social radius = 3 or 4). In decentralized economies, sanctions from transgressing the central authority are low, and social pressures play a relatively muted role in encouraging or discouraging protest. As sanctions increase, a cascade is more likely to emerge, and citizens are more likely to express their displeasure with the central authority. In sufficiently centralized economies, the punishment from transgression is large enough that a cascade does not emerge, even as a minority of agents changes their behavior. This was seen in the previous section, which showed that the change in preference falsification was relatively small in highly centralized societies. On top of this, preferences are so falsified in highly centralized economies prior to the shock that the central authority's pre-shock action is not far from his bliss point. After the shock, there is more room for the central authority to cede some power to the citizenry, decreasing the overall level of protest. This is indicative of the actions taken by the central authorities in Syria Bahrain, Saudi Arabia, and Oman soon after the revolts spread from Tunisia to Egypt and Libya. Soon after protests began, Syrian President Assad accepted the resignation of the government and decreed nationality to thousands of Kurds, Bahrain's king gave out 1,000 dinars ($2,625) to each Bahraini family, Saudi Arabia's King Abdullah increased spending on a range of social programs, and the Sultan of Oman granted new lawmaking powers to councils and increased minimum wages by 43 percent.

Figure 6: (upper region) Average difference in protest before (t=20) and after (t=40) the shock; (lower region) Average difference protest over centralization, shock fraction 0.4 – 0.6 only

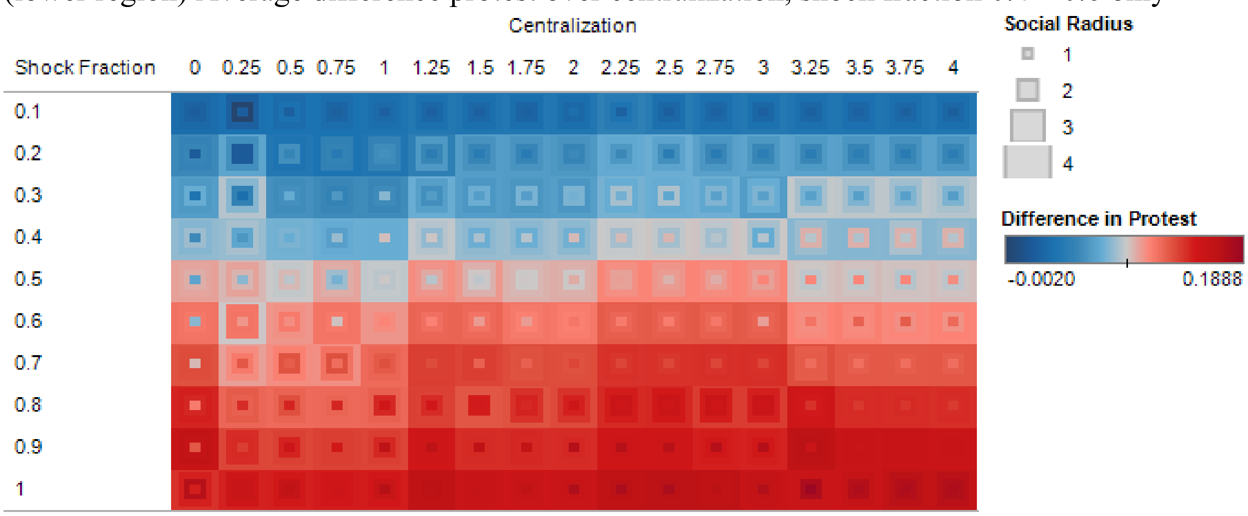

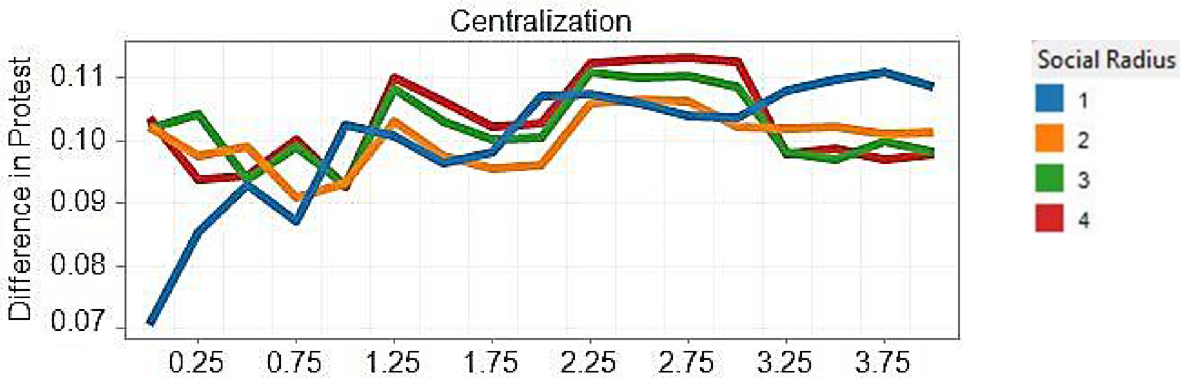

Secondly, in the middling shock range (40%-60%), the difference in protest is *decreasing* in social radius in highly centralized economies (centralization ≥ 3). Despite the fact that low network range societies have less absolute protest, the change after the shock is greater in these economies. There is very little incentive to protest prior to the shock when networks are weak and institutions are highly centralized, so the post-shock protests result in a very different world from the one that existed prior to the shock.

Yet, the larger change in protest does not necessarily lead to a revolution. As we show in the next sections, the change in protest is smaller in large-network range economies because *both* the population and the central authority change their actions significantly after the shock. This finding is relevant to the events which took place early in the Syria and Bahrain uprisings, where some protests occurred (and were eventually beaten back by force) but the governments attempted to get ahead of the protests by calling for freer elections and other reforms prior to protests escalating to the level of the Egyptian or Libyan revolts.

Social Revolution

We are also concerned with how the actions of the citizens change in response to the shock. We denote the degree of this change a "social revolution", since it indicates how publicly expressed preferences react to the shock. Figure 7 maps the degree of social revolution over the parameter space. There are numerous interesting results that emerge. First, social radius is an important predictor of social revolution in the model, especially when the shock is in the middle range (40%-60%). In this range, the difference in social revolution between low- and high-radius societies is considerably larger. As noted before, these "middling" shocks have a large effect on outcomes *only* when they are adequately transmitted through the population by the social mechanism. Since social transmission is much more likely to occur when social radius is large, the citizens are much more willing to revolt against the government in high-radius societies. It is *not* the case that preferences change, but instead the preferences of a large enough portion of the population are *revealed*, which encourages a cascade of further preference revelation.

Figure 7: (upper region) Average social revolution (mean citizen action$_{40}$ – mean citizen action$_{20}$); (lower region) Average social revolution over centralization, shock fraction 0.4–0.6 only

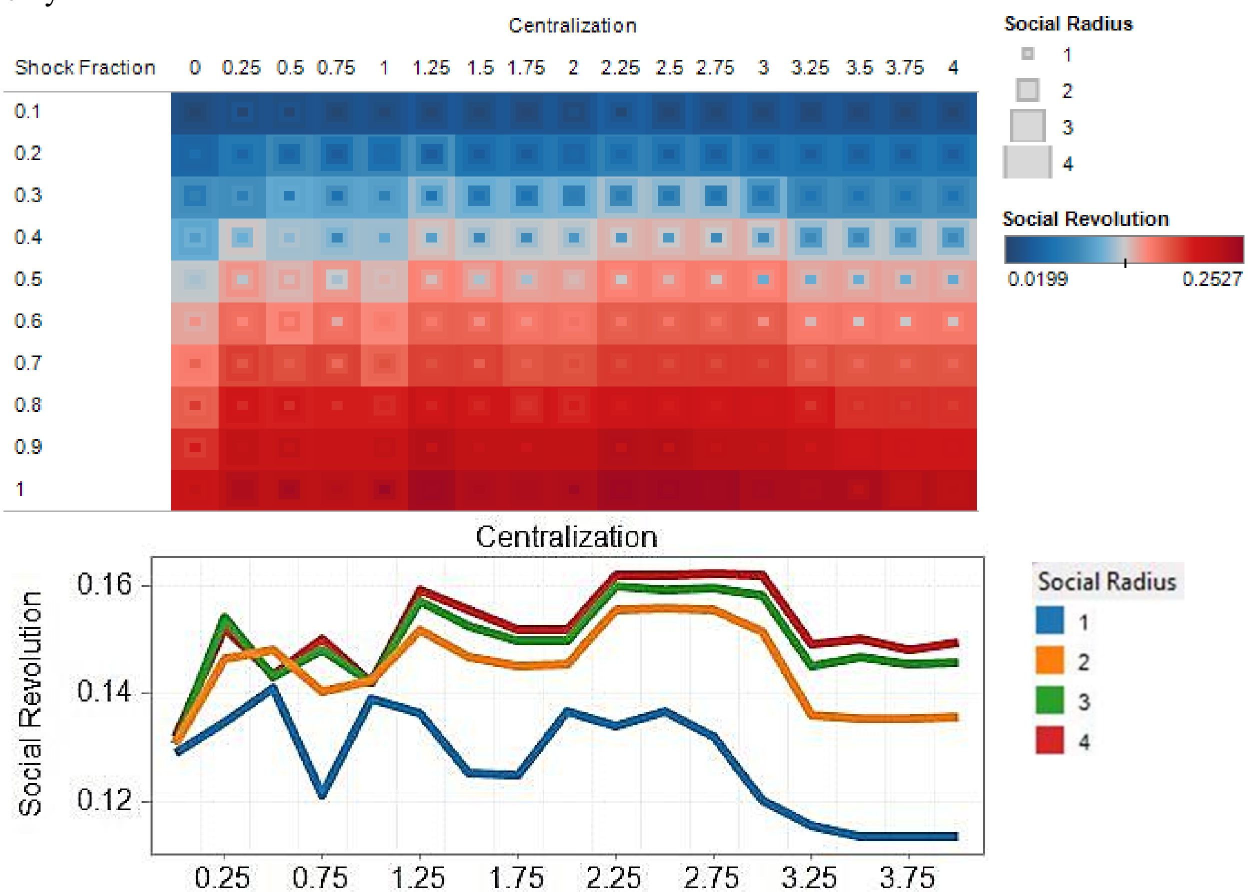

Focusing on the “middling shocks” in Figure 7, it is clear that the relationship between centralization and social revolution is not linear and is dependent on social radius. At the smallest social radius, social revolution is generally *decreasing* in centralization. This is not surprising; if the shock is not transmitted to the population *and* the cost of transgressing the authority is great, then citizens will change their actions less. This result helps explain why highly centralized regimes – not just in the Arab world but in China, Iran, North Korea, and Burma, to name a few – attempt to restrict to flow of information at almost any cost. By keeping their citizens weakly connected, they greatly reduce the possibility of a social revolution.[6]

Conversely, at larger social radii, social revolution has an inverted-U shape with respect to centralization. The intuition underlying this result is similar to that espoused in previous sections. At extremely high levels of centralization, many citizens choose not to change their behavior, even as they observe others in their network doing so. As such, the sanctioning power of a heavily centralized authority prevents the actions of an active minority of protestors from triggering a cascade and tipping the system. Meanwhile, a cascade may never emerge at lower levels of centralization because citizens were falsifying their preferences to a much lesser extent prior to the shock, and hence the social revelation mechanism is not necessary. This provides a rationale for why revolts in Syria and Bahrain did not succeed in a similar manner to those in Tunisia and Egypt. In Syria and Bahrain, the central authority maintained control of the military and was willing to use it to suppress revolt, thus increasing the cost of revolt. This was not the case in Tunisia or Egypt, where a fractured military did not support the leader in the face of

[6] In lower subfigures of Figure 6 through Figure 10, it is notable that the figures are rarely monotonic. This complexity in the model is symptomatic of sensitivity to initial conditions. The initial conditions in question are the differing arrangement of heterogeneous citizens across the lattice.

massive protests (Anderson 2011; Goldstone 2011).[7] Employing our definition of centralization, this simply means that the governments of Syria and Bahrain were more centralized than those in Tunisia and Egypt (where the governments lost control of the military), since the Syrian and Bahraini central authorities had control over more types of sanctions. This encouraged *more* social revolt in Syria and Bahrain than it would have if these countries were decentralized (since preferences were massively falsified), but *less* social revolt than in less centralized countries.

Since the social revolution metric considers actions across time periods, we also analyze the change in citizens' actions *relative* to their level of protest prior to the shock. These results are reported in Figure 8. The level of pre-shock protest is an appropriate weight because it measures how "quiet" the streets are prior to the shock. A change in the mean citizens' action is much more revolutionary if it occurs in a society where there is little protest prior to the shock relative to a society where citizens are protesting loudly prior to the shock.

[7] Goldstone (2011) notes that quick and massive changes may occur in centralized regimes when the military is fractured, because "the elite and military officers have every reason to hide their true feelings until a crucial moment arises, and it is impossible to know which provocation will lead to mass, rather than local, mobilization. The rapid unraveling of sultanistic regimes thus often comes as a shock."

Figure 8: (upper region) Average social revolution weighted by protest (mean citizen action_{40} – mean citizen action_{20}); (lower region) Average social revolution weighted by protest over centralization, shock fraction 0.4–0.6 only

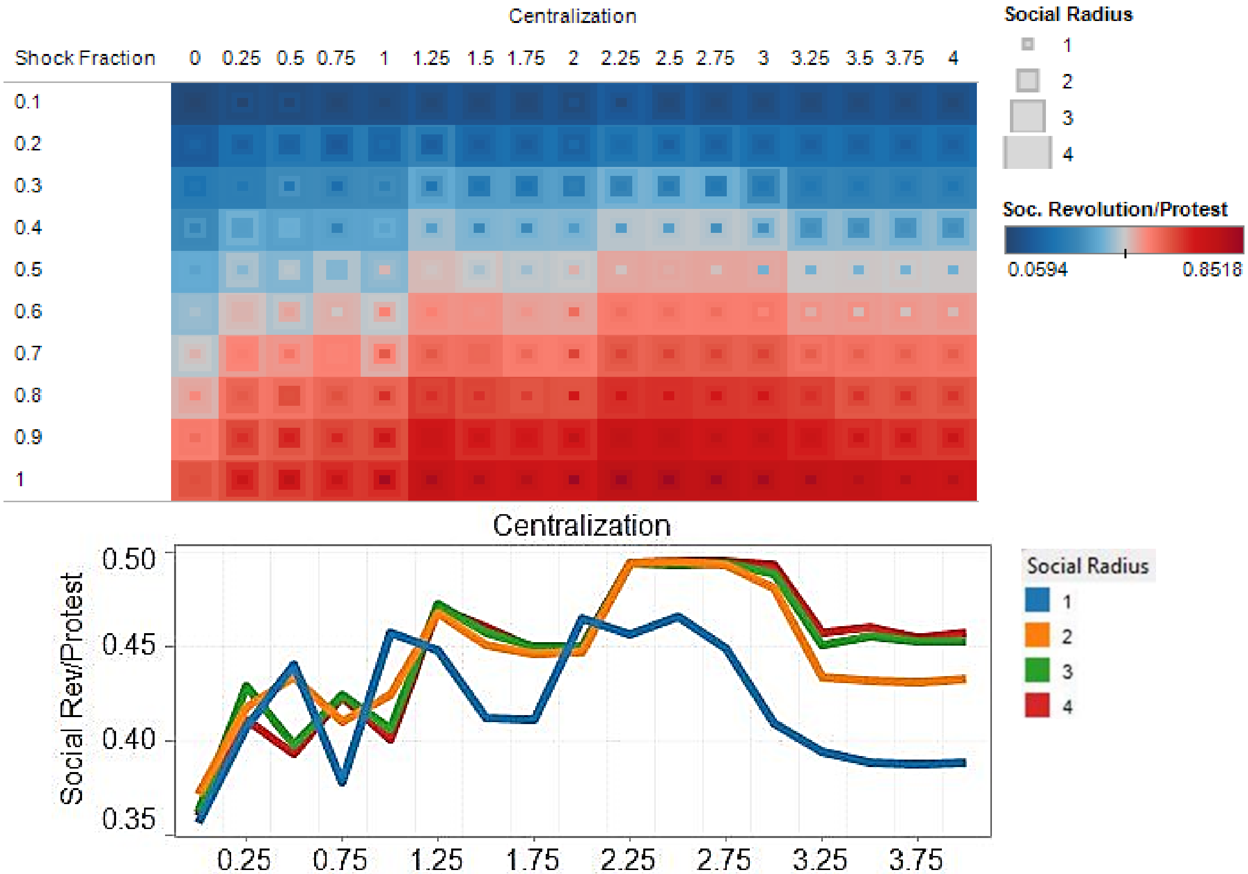

Many of the same results emerge when social revolution is weighted by protest as when it is not weighted by protest. However, one interesting, non-obvious result is that social radius leads to *less* social revolution in decentralized societies (centralization ≤ 1) but *greater* social revolution in centralized societies (centralization > 2). In decentralized, high-radius societies, citizens are likely to express the preferences prior to the shock. While the shock may encourage further expression of preferences, this is based in an society where such expression was the norm, so less change in publicly expressed preferences results. This is largely why modern, high connectivity[8] democracies see fewer massive, rapid uprisings than autocracies; in the former,

[8] See the Appendix for a discussion of connectivity statistics and how they relate to social radius in the model.

preferences are widely expressed, so while a shock may change actions, it rarely leads to a preference revelation cascade.

In centralized economies, however, high-radius is associated with greater social revolution. Even though a greater radius does not encourage preference revelation prior to the shock (since the associated sanctions are large), the shock facilitates an unraveling of actions to a much greater extent and hence social revolution is exacerbated.

Institutional Revolution

While the central authority influences outcomes through its ability to sanction deviant individual behavior, it is nonetheless beholden to the choices made by the citizen population. We denote the degree to which the central authority responds to the shock as "institutional revolution." Figure 9 and Figure 10 highlight the degree of institutional revolution over different parts of the parameter space, both unweighted (Figure 9Figure 9) and weighted by pre-shock protest (Figure 10).

Figure 9: (upper region) Average unweighted institutional revolution (Central action$_{40}$ – Central action$_{20}$) over centralization; (lower region) Average institutional revolution over centralization, shock fraction 0.4–0.6 only

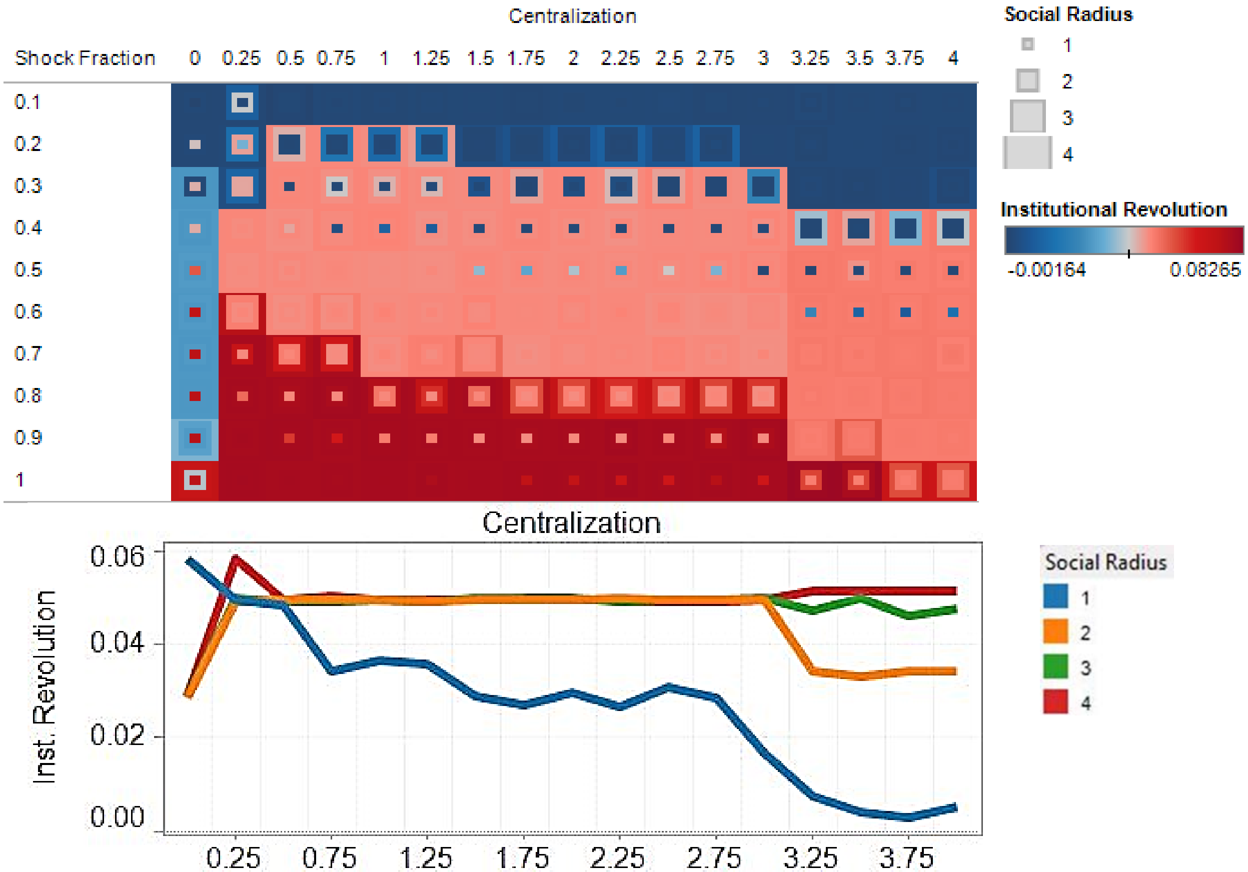

Figure 10: (upper region) Average institutional revolution weighted by protest in t=20 over centralization; (lower region) Average institutional revolution weighted by protest over centralization, shock fraction 0.4–0.6 only

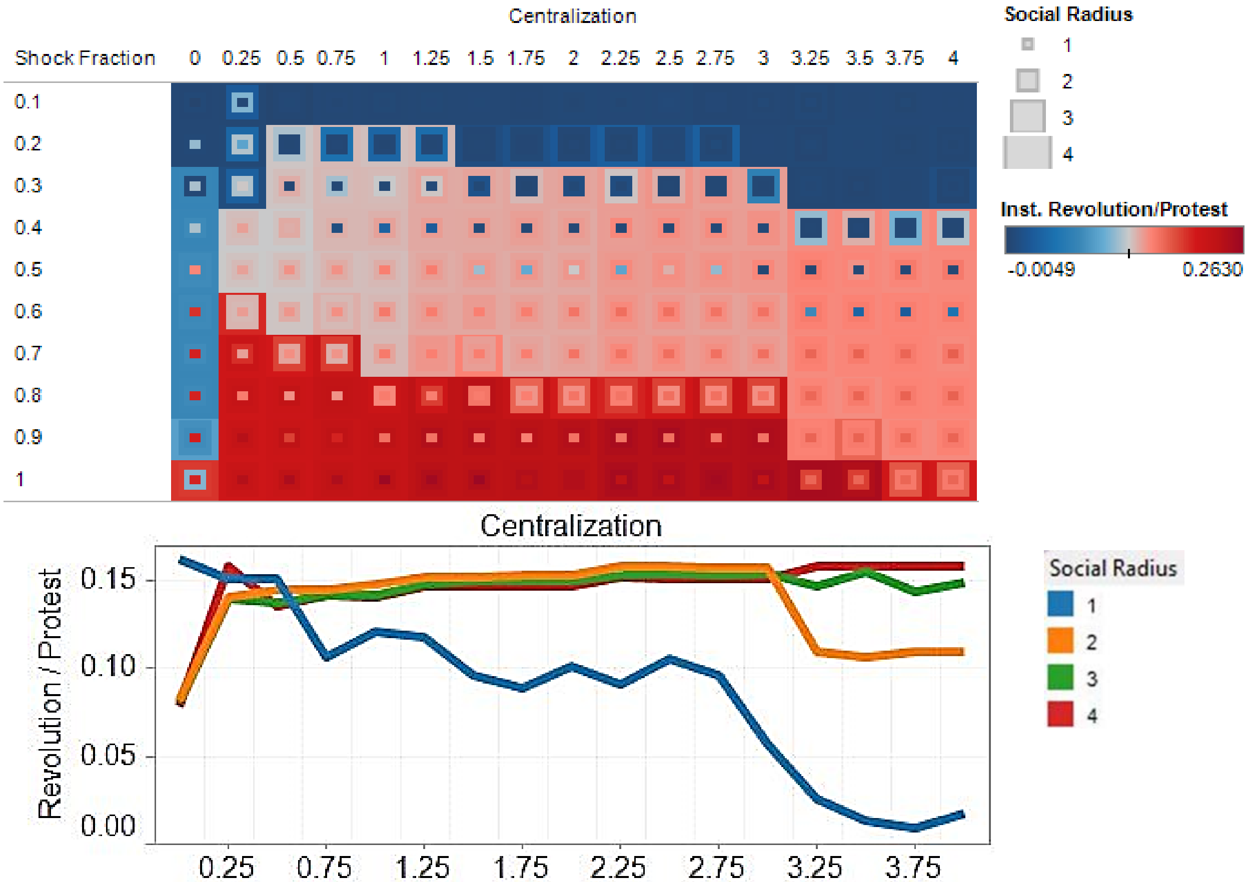

The cascade thresholds are visually striking in Figure 9 and Figure 10. We observe the parameter combinations that are sufficient to tip the system towards either a shallow or steep cascade that pushes the central authority towards significant change. In the lower subfigures of Figure 9 and Figure 10, we observe that within the middling shock ranges (40%-60%), the degree of institutional revolution is sharply *decreasing* with centralization if the social radius is small, while mildly increasing with centralization if the social radius is large.

Institutional revolution is dramatically curtailed in small radius societies (radius = 1 or 2) when centralization is sufficiently strong. Or, put another way, tipping the system towards institutional change in heavily centralized societies with low social radius requires a social shock that affects nearly the entire population. The lower region of Figure 10 suggests that in small-radius societies, highly centralized authorities (centralization ≥ 3) change their actions much less than decentralized authorities. A number of factors lead to this result. For one, the absolute levels of protest (Figure 5) and social revolution (Figure 7) are much lower in low radius societies. Figure 5 and Figure 7 also indicate that protest and social revolution decline much more steeply in low radius societies as centralization increases. Moreover, institutional revolution from middling shocks is slightly *increasing* in centralization in high radius societies (radius = 3 or 4). Unlike low-radius societies, which "tip" towards less institutional revolution at sufficiently high centralization levels, high-radius societies have slightly more institutional revolution at high centralization levels. At first blush, this seems inconsistent with previous results. If preference falsification and protest change less in highly centralized societies than in moderately centralized societies, (Figure 3 and Figure 6 show an inverted U-shape at large social radii), why is institutional revolution increasing in centralization in high-radius societies?

In such societies, the pre-shock actions of all players favor the central authority to such an extent that the authority was in a much better position (relative to a low-centralization authority) to cede ground to the citizenry following a cascade of preference revelation. Despite the fact that some citizens do not join the cascade in highly centralized economies for fear of retribution, the pre-shock action of the central authority is *so oppressive* vis-à-vis the citizens' bliss points that the central authority has more incentive to relax some of the restrictions following a shock that is transmitted to a broad swath of the population. The high radius permits the information transmission mechanism to spread the shock, and the central authority responds by changing its action to a less oppressive one. This provides an explanation for the actions of leaders in Libya, Syria, and Bahrain, where some protests occurred (and were eventually beaten back by force) but the governments attempted to get ahead of the protests by calling for freer elections and other reforms *prior* to protests escalating to the level of the Egyptian or Libyan revolts. While the success of the revolts in these countries is still in question, the leaders were much more willing to give ground than Tunisian and Egyptian leaders were prior to their ouster.

In both Figures 9 and 10 we can see two shock fraction thresholds beyond which point small and large institutional change can be observed on the part of the central authority. In Figure 9 and Figure 10 there are sharp demarcations between blue (zero change), white (mild change), and red (significant change) regions of the heat map. These demarcations reveal thresholds – the minimum shock size necessary to trigger a cascade causing institutional change. These minimum shock thresholds are decreasing with social radius. This effect can be observed more clearly in Figure 11. The amount of institutional revolution is monotonically increasing with shock size, but in a discontinuous manner. The threshold values which trigger large amounts of social change are both proceeded and followed by extended plateaus over which great shock size does not result any (significant) additional institutional change. The shock fraction thresholds at which these sharp increases in institutional change occur are decreasing as network radii increases.

Figure 11: Average unweighted institutional revolution (Central action$_{40}$ – Central action$_{20}$) over shock fraction (Centralization (γ) =3.0)

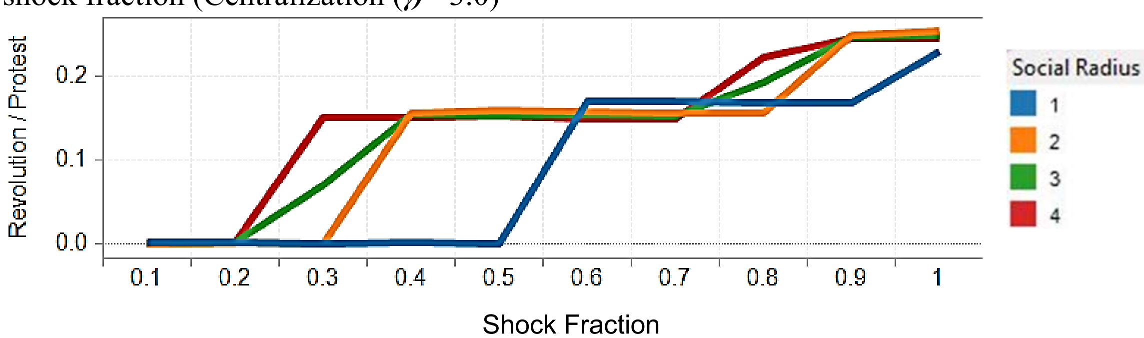

Figure 11 indicates that the stability of centralized institutions is decreasing with social radius. It also reveals the threshold below which social shocks can go completely unnoticed, having no visible effect on the behavior of the central authority. In Figure 11, when social networks have a radius of 1, the authority will not change at all until a social shock reaches 50% of the population. If the network radius increases to 2, however, much smaller shocks can tip the system towards significant institutional change. This result offers the theoretical possibility that recent revolutionary phenomena, such as the Arab Spring, are not the product of unusually pervasive social shocks, but rather authorities that are now vulnerable to much smaller shocks (shocks that would have previously gone unnoticed) because of the increased availability of modern social networking technology.

Conclusion

Our results suggest that highly centralized regimes may seem tranquil but are highly susceptible to revolution, especially in large network-range economies. This sheds light on the institutional, technological, and social mechanisms facilitating the recent spread of revolutionary activity in the Arab world, highlighting the reasons underlying differing institutional responses to public discontent in Tunisia, Egypt, Libya, Syria, and Bahrain. An agent-based model highlights the role that information and communication technology play in triggering cascades of preference revelation in centralized societies. We show that network range reduces the minimum shock that is sufficient to effect institutional change, though this result is contingent upon the degree to which institutions are centralized. We find that revolutions are more likely to occur as societies become more centralized when network range is large but less likely to occur in centralized societies when network range is small. While the citizen population reacts to exogenous social shocks in a steady, linear manner, central authorities exhibit a far more punctuated behavioral pattern, changing their behavior only when an exogenous shock is sufficiently large to the tip the system towards significant change. The thresholds at which central authorities change their behavior are significantly reduced by increased social network reach. These results point towards a world where heavily centralized authorities are more likely to move towards the preferences of the general population in societies with increased access to modern ICT. At the same time, these results also reveal the incentive for central authorities to limit citizen access to ICT, including the internet and social networking.

References

Allen, F. and D. Gale (2000). "Financial Contagion." Journal of Political Economy **108**(1): 1-33.

Anderson, L. (2011). "Demystifying the Arab spring: parsing the differences between Tunisia, Egypt, and Libya." Foreign Affairs **90**(3).

Banerjee, A. V. (1992). "A Simple Model of Herd Behavior." Quarterly Journal of Economics **107**(3): 797-817.

Bicchieri, C. (2006). The grammar of society : the nature and dynamics of social norms. New York, Cambridge University Press.

Bikhchandani, S., D. Hirshleifer, et al. (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades." Journal of Political Economy **100**(5): 992-1026.

Bonabeau, E. (2002). "Agent-based modeling: Methods and techniques for simulating human systems." Proceedings of the National Academy of Sciences of the United States of America **99**(Suppl 3): 7280-7287.

Centola, D. (2010). "The Spread of Behavior in an Online Social Network Experiment." Science **329**: 1194-1197.

Centola, D., R. Willer, et al. (2005). "The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms." American Journal of Sociology **110**(4): 1009-1040.

Cowan, R. and N. Jonard (2004). "Network structure and the diffusion of knowledge." Journal of Economic Dynamics and Control **28**(8): 1557-1575.

Delre, S. A., W. Jager, et al. (2007). "Targeting and timing promotional activities: An agent-based model for the takeoff of new products." Journal of Business Research **60**(8): 826-835.

Ellis, C. J. and J. Fender (2011). "Information Cascades and Revolutionary Regime Transitions." Economic Journal **121**(553): 763-792.

Epstein, J. M. (2002). "Modeling civil violence: An agent-based computational approach." Proceedings of the National Academy of Sciences of the United States of America **99**(Suppl 3): 7243-7250.
Epstein, J. M. (2006). Generative social science : studies in agent-based computational modeling. Princeton, Princeton University Press.
Epstein, J. M. and R. Axtell (1996). Growing artificial societies : social science from the bottom up. Washington, D.C., Brookings Institution Press.
Fowler, J. H. and N. A. Christakis (2010). "Cooperative behavior cascades in human social networks." Proceedings of the National Academy of Sciences **107**(12): 5334-5338.
Goldstone, J. A. (2011). "Understanding the revolutions of 2011: weakness and resilience in Middle Eastern autocracies." Foreign Affairs **90**(3).
Golub, B. and M. O. Jackson (2011). "How Homophily Affects the Speed of Contagion, Best Response and Learning Dynamics." Working Paper, http://www.stanford.edu/~bgolub/papers/homophily.pdf
Granovetter, M. (1978). "Threshold Models of Collective Behavior." American Journal of Sociology **83**(6): 1420-1443.
Kuran, T. (1987). "Preference Falsification, Policy Continuity and Collective Conservatism." Economic Journal **97**: 642-665.
Kuran, T. (1989). "Sparks and prairie fires: A theory of unanticipated political revolution." Public Choice **61**(1): 41-74.
Kuran, T. (1991). "The East European Revolution of 1989: Is it Surprising that We Were Surprised?" The American Economic Review **81**(2): 121-125.
Kuran, T. (1991). "Now Out of Never: The Element of Surprise in the East European Revolution of 1989." World Politics **44**(1): 7-48.
Kuran, T. (1995). "The Inevitability of Future Revolutionary Surprises." American Journal of Sociology **100**(6): 1528-1551.
Kuran, T. (1995). Private Truths, Public Lies: The Social Consequences of Preference Falsification. Cambridge, MA, Harvard University Press.
Lee, I. H. and Á. Valentinyi (2000). "Noisy Contagion without Mutation." The Review of Economic Studies **67**(1): 47-56.
Lohmann, S. (1994). "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91." World Politics **47**(1): 42-101.
Luke, S., C. Cioffi-Revilla, et al. (2005). "MASON: A Multiagent Simulation Environment." SIMULATION **81**(7): 517.
Macy, M. W. (1991). "Chains of Cooperation: Threshold Effects in Collective Action." American Sociological Review **56**(6): 730-747.
Macy, M. W. and R. Willer (2002). "From Factors to Actors: Computational Sociology and Agent-Based Modeling." Annual Review of Sociology **28**(ArticleType: research-article / Full publication date: 2002 / Copyright © 2002 Annual Reviews): 143-166.
McPherson, M., L. Smith-Lovin, et al. (2001). "Birds of a Feather: Homophily in Social Networks." Annual Review of Sociology **27**(ArticleType: research-article / Full publication date: 2001 / Copyright © 2001 Annual Reviews): 415-444.
Morozov, E. (2011). The net delusion : the dark side of internet freedom. New York, Public Affairs.
Morris, S. (2000). "Contagion." The Review of Economic Studies **67**(1): 57-78.

Oliver, P. E., G. Marwell, et al. (1985). "A Theory of Critical Mass. I. Interdependence, Heterogeneity and the Production of Collective Action." American Journal of Sociology **91**: 522-556.
Press, W. H. (2002). Numerical recipes in C++ : the art of scientific computing. New York, Cambridge University Press.
Rubin, J. (2011). "Centralized Institutions and Sudden Change." Working Paper.
Schelling, T. C. (1978). Micromotives and Macrobehavior. New York, W. W. Norton Company.
Shirky, C. (2011). "The Political Power of Social Media. Technology, the Public Sphere, and Political Change." Foreign Affairs **vol. 90**(no. 1): 28–41.
Siegel, D. A. (2009). "Social Networks and Collective Action." American Journal of Political Science **53**(1): 122-138.
Wasserman, S. and K. Faust (1994). Social network analysis : methods and applications. Cambridge ; New York, Cambridge University Press.
Willer, R., K. Kuwabara, et al. (2009). "The False Enforcement of Unpopular Norms." American Journal of Sociology **115**(2): 451-490.
Wright, G. (1999). ""The Civil Rights Revolution as Economic History"." Journal of Economic History **59**(2): 267-289.
Young, H. P. (1993). "The Evolution of Conventions." Econometrica **61**(1): 57-84.

Appendix

Distributions of Preferences and Actions

Figure 12 and Figure13 present the distributions of bliss points and actions at time steps 20 and 40 from single runs of the model under different parameterizations. Both sets of distributions are from runs of the model with shocks covering 20% and 80% of the population, with centralizations parameters (γ) of 1.0 and 3.0. Figure 12 includes runs with agent social radius of 1 andFigure 13 Figure 13 includes runs with a social radius of 4.

Figure 12: The distributions of agent bliss points (red), actions at t = 20 (blue) and t = 40 (yellow) at social radius = 1

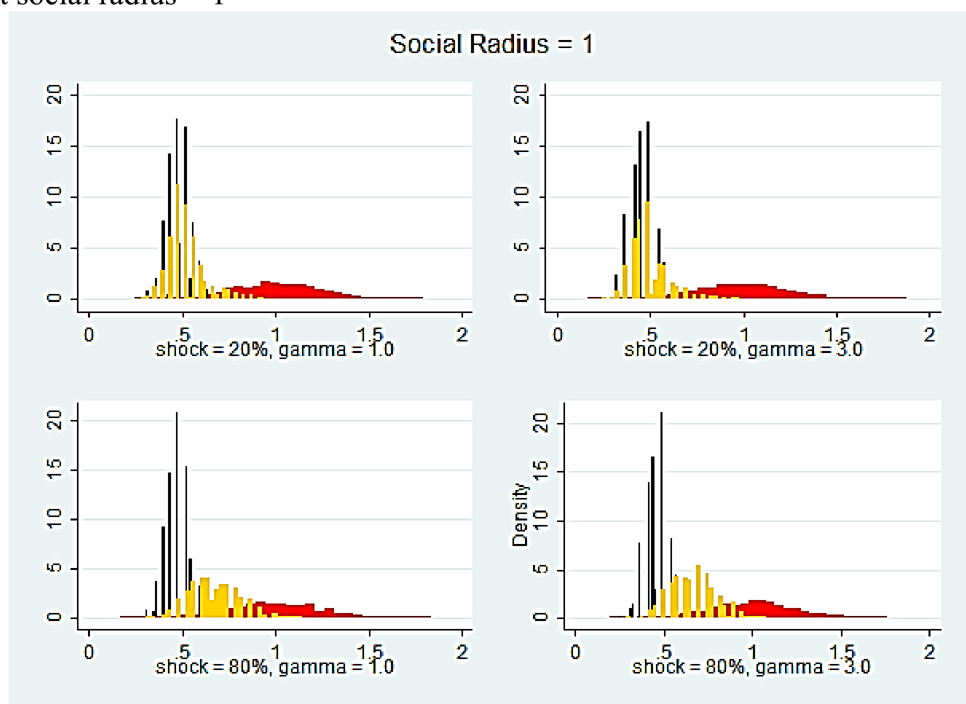

Figure 13: The distributions of agent bliss points (red), actions at t = 20 (blue) and t = 40 (yellow) at social radius = 4

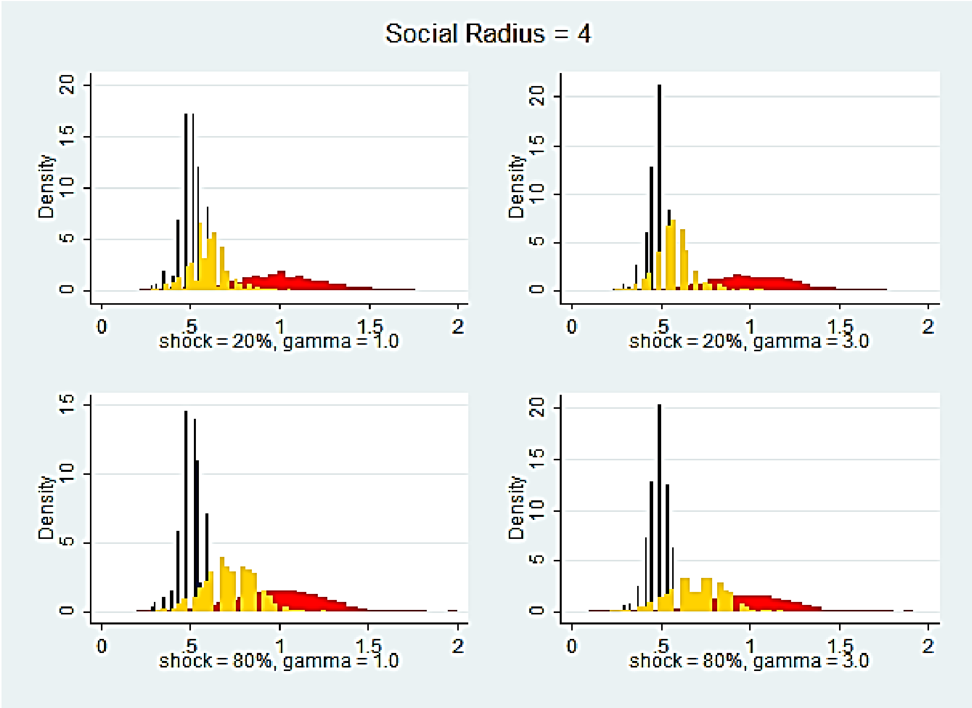

Alternative experimental construct

As a robustness check, we ran the model experiments using an alternative network exploratory scheme. Instead of increasing the radius of the of Moore neighborhood of agent networks across runs of the model, we created "mixed radius models" where a percentage of the population has "large radius" networks (r = 4) and the rest have "small radius" networks (r = 1). Varying the percentage of "large radius" across runs, we generated qualitatively identical results to models varying the network range of the entire population. Figure 14 presents the results for institutional revolution using the alternative network experiment structure.

Figure 14: Institutional revolution using alternative network

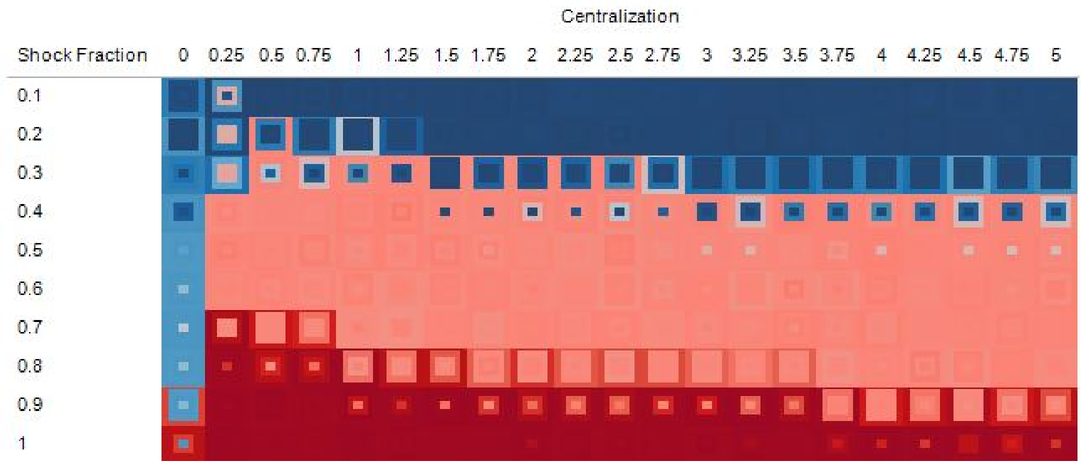

As can be seen in Figure 14, the qualitative structure of the results is nearly identical to our original experimental structure. For the nature of investigation, increasing the percentage of the population with long range social networks is largely equivalent to incrementally increasing the radius of the entire population.

Statistical Analysis of Network Radius

The changing ICT landscape is captured in the model with the social network radius parameter. The character of these changing social networks is further explored using rudimentary network statistical analysis. We make use of three measures: *node connectivity*, *mean geodesic*, and *local clustering coefficient.* The node connectivity of a graph is the minimum number of nodes that must be removed make it disconnected (i.e., leaving at least one agent stranded without a connection to any other agent). The mean geodesic is the average shortest path between any pair of nodes in the network. The local clustering coefficient measures the average ratio of connections to a node to number of possible connections. As agent neighborhoods move towards becoming complete (fully-connected) graphs, the local clustering coefficient increases. For further discussion of these network statistics, see Wasserman and Faust (1994).

In Figure 15, we chart the network statistics before and after the shock from an experiment over the different social network radii. All of the runs include a shock affecting 50% of the agent population and centralization (γ) = 2. The experiment was run 10 times with each network radius setting (n = 40). Connectivity (both before and after the shock) is increasing with network radius. At the same time, the mean geodesic is decreasing with network radius, as measured both before and after the shock. Clustering before the shock is steadily increasing with network radius. Clustering after the shock, however, has a less linear relationship with network radius, making a considerable jump from $R = 2$ to $R = 3$, but the two are nonetheless positively correlated. The connectedness of agents, both globally (in terms of both node connectivity and mean geodesic) and locally (clustering), is increasing with network radius, and in turn supports the cascades of social change observed in the model.

Figure 15: Network clustering coefficient, connectivity, and mean geodesic over network radius, t=20 and t=40