H i C N Households in Conflict Network

A Model of Protests, Revolution, and Information[1]

Salvador Barbera[2] and Matthew O. Jackson[3]

HiCN Working Paper 243

February 2017

Abstract:

A revolt or protest succeeds only if sufficient people participate. We study how potential participants' ability to coordinate is affected by their information. We distinguish four phenomena that affect whether information either encourages or inhibits protests and revolutions: (i) Unraveling: When agents learn about each others' types, some are discouraged by meeting partisans of the status quo. This can unravel, as even confident agents realize that enough supporters will be discouraged to preclude a successful revolution. (ii) Homophily: Learning someone else's type under homophily is less informative since that individual is more likely to be similar to the learner. This can lead people to be less confident of a revolution, but can also stop potential unraveling. (iii) Extremism: Meeting other protestors, and seeing pilot demonstrations or outcomes in similar countries, reveal not only how much support for change exists, but also from which constituencies it emerges. This can undercut a revolution if factions differ sufficiently in their preferred changes. (iv) Counter Demonstrations: partisans for the status quo can hold counter-demonstrations to signal their strength. We also discuss why holding mass demonstrations before a revolution may provide better signals of peoples willingness to actively participate than other less costly forms of communication (e.g., via social media), and how governments use redistribution and propaganda to avoid a revolution.

Keywords: Revolution, demonstration, protests, strikes, Arab Spring

JEL Classification Codes: D74, D72, D71, D83, C72

---

1 Introduction

In human societies, when enough people agree on the direction of desirable change that is not being directly taken by a government or private organization, there may be room to force it, at a cost. Prior to acting for change, those agents who want it need to learn about the strength of their group and even to let others, who are also dissatisfied but maybe not for the same reasons, know that they have sufficient forces and a common ground upon which to act. Demonstrations, strikes, and other more spontaneous forms of mass protests, provide such information. A successful mass protest informs dissenters about the size and commitment of their group, and signals to potential allies the possibility to gather a larger segment of society into a common movement.

In this paper, we provide a simple model of collective action, focusing on the informational role of mass protests. We study how the incentives to participate in them, and their eventual success, depends on the information that is available to agents at the beginning of a process of revolt, through direct communication between agents and via series of protests. We also discuss how, in spite of the improvements in social media and communication, demonstrations and protests remain as a differentiated and particularly revealing method to learn about the intensity of preferences, and the conditions under which agents are willing to take risks in favor of change.

The final result of protests, if successful, will be called a revolution, but it must be understood that our model is meant to encompass different types of events under this generic name. In some extreme cases, success is the overthrow of a government; but in others, it may be a significant change in the political scenario, enough to produce a desired change in policy, or even just gathering media attention to change a company's policy. In some cases, the "revolution" will involve violence while in others it may remain peaceful. And the success of a revolution today, as we have seen in many cases, may turn into a failure tomorrow, especially if the protesters are themselves a heterogeneous group whose interests may conflict right after they achieve their initial common goal.

Although there may be instances where the strength of a single mass mobilization, coupled with the weakness of the status quo defenders, may bring about change instantly, a revolution is most often the result of a succession of events. Hence, we are interested in dynamics, and on the consequences of information gathering through collective action as part of multi-period processes. Protests can gain or loose momentum, they can arise in one country after observing the success of similar groups in other areas of the world, and they may also be deterred in one place when its success in others is a premonitory warning to the status-quo defenders.

These distinctions and others can be found in our paper, framed within a unifying model, whose simplicity and flexibility enables us to investigate many issues in a common context rather than restricting the model to analyze some particular phenomenon.

Even if we touch on additional issues, the main contribution of our paper is to distinguish four ways in which information either encourages or inhibits protests and revolutions:

Unraveling: When agents learn about each others' types, some of them will be discouraged as they will meet partisans of the status quo. This then unravels, as even though others will not be directly discouraged from what they learn, they realize that enough others will be discouraged to preclude a successful revolution. Whether learning about others enables or inhibits a revolution depends on whether numbers of supporters are sufficiently large so that the loss of some discouraged potential protestors does not outweigh the improved confidence of those who have met others with similar preferences for change.

Homophily: Homophily reduces the content of information. By having high homophily, learning about someone else's type is unlikely to add new knowledge since that individual is more likely to be similar to the learner. Thus, people learn little about the broader society's preferences as they know that their friends are not representative of the population's preferences. Thus, by reducing information content, homophily makes it harder to hold revolutions in cases in which learning was necessary to enable a revolution, but it makes it easier to hold revolutions in cases in which learning would unravel the revolution. Thus, homophily provides a new and important angle on when protests and revolutions may succeed.

Extremism: When potential protestors meet, or sequences of demonstrations are held, or a population learns about outcomes in similar countries, people may observe not only about how much support for change exists, but also from which constituencies that support emerges. This helps potential revolutionaries forecast what will happen after the revolution. If there are many extremists whose new agenda might not be preferable to the status quo, learning about their numbers can lead moderates to back away from a revolution that they might have otherwise supported.

Counter Protests and other Government Responses: Government, or their partisans, may wish to hold counter-protests to signal the strength of support for the status quo. In some settings, if an initial protest in favor of change leaves some doubts as to the size of the support, counter-protests can become important in fully revealing the preferences of the population and can inhibit an eventual revolution. Governments can also respond by trying to manipulate beliefs and sow doubt via propaganda, increase the costs of protests, or buy off some of the disenchanted.

We begin the formal part of the paper in Sections 2 and 3, by describing a one-shot model in which the members of a population must simultaneously decide whether to participate in a "revolution" (protest, strike, etc.) in ignorance of other agents types. The revolution is successful if sufficient numbers participate, but not otherwise. We present this model first, since it is a useful benchmark and building block for our main analysis. This sort of model is standard as it is a basic coordination game among a population of players, and has been extensively studied in the global games literature (e.g., see Angeletos, Hellwig, and Pavan (2007)). We provide some basic comparative statics for the model.

Then, in Sections 4 and 5, we move to our main analysis to analyze how information that potential revolutionaries get via meeting each other, and via the building up of demonstrations or protests, can prompt or deter "revolutions".

We first examine, in Section 4.1, how communication between potential protestors matters in either enabling or disabling a revolution. We consider situations in which people get to observe the type of another agent in society.[1] We show that this produces two countervailing effects. One, is that agents who meet with another person who prefers change are now more confident about the size of the potential revolution. The second is that agents who meet someone who prefers the status quo are now discouraged about the size of the potential revolution. The subtle aspect here is that even when there are many supporters of change, and so most of them are encouraged by what they see, they still know that some supporters will end up being discouraged which will reduce the participation. Thus, communication can help some become more confident, but it also can thin the numbers of those confident enough to show up for the revolution. So, we show that: (i) In some cases communication can enable a revolution that could not occur otherwise. This happens in cases in with poor prior information about the numbers of potential revolutionaries. (ii) In some cases communication can make a revolution impossible, even though it would have been possible without communication. This happens when there is stronger prior information about the numbers of potential revolutionaries, but not an overwhelming number of them relative to the threshold number that they need. We then go on to show that if people can communicate with sufficient numbers of agents, eventually reaching full knowledge of their numbers, revolutions occur if and only if they will be successful. Thus, we find an important, and new insight regarding how small amounts of information can actually make revolutions impossible - while no information or large amounts of information would lead to successful revolutions.[2]

In Section 4.2, we discuss how equilibria change when agents' sources of information are biased by homophily - so that agents who support change are most likely to be talking with others who feel the same, biasing the sample of communication that they receive about the society. When we know that most people with whom we interact are similar to us, then this makes it harder for us to judge how large the group who supports change might be. Our friends are not a representative sample of the population. This affects both of the forces just mentioned above above. Most people we meet will share our views, and so we learn little about the population from talking to them. On the other hand, that means that few of us might end up learning about how many people have other views and so we might not be

---

[1] Communicating types is incentive compatible in this setting.

[2] An interesting, but very different signalling phenomenon in a voting setting is studied by Lohmann (1993). She examines costly political action prior to voting, when voters are trying to estimate a state variable about which alternative is best to vote for, and her effects are based on the fact that only agents who are extreme enough take political action, which does not provide full information about the state and may confound it. Here, our effect ends up being a strategic one: agents know that some supporters will be discouraged and hence will not participate, and then through the strategic complementarity of the revolution discourages even those with strong information from participating.

discouraged from participating. Which of these effects dominates depends on our priors and the correlation of types with the state, and the degree of homophily.

Next, in Section 5, another of our main results concerns when it is that holding a demonstration before a revolution can be important in enabling the revolution. Without such a demonstration, moderate and cautious supporters may have insufficient information about the probability of success to participate. Demonstrations can help in settings in which moderates are sufficiently cautious, but stronger supporters are willing to demonstrate, even facing the costs of failure, to get the process rolling. Such a process can be efficiency enhancing, as then the full revolution only occurs in cases in which it is most likely to be successful.

Our model also shows why simple polls or cheap talk are not enough (Section 5.2). Social media have been very critical in helping coordinate demonstrations, but they cannot substitute for demonstrations, since they are 'cheap-talk' and do not involve the costly signaling that demonstrations provide.[3] A natural setting is one in which there are many people who would prefer change, but also in which many of them are not willing to pay the personal costs of being an active part of a revolution. They may communicate their support, but fail to turn out when action is needed. Holding a somewhat costly demonstration is a filtering device, which then signals whether there are sufficient numbers of people who are willing to act for change, not just cheer it on. Thus, holding a demonstration before a revolution can be a necessary step to enabling the revolution.

We also provide a picture of the successive steps through which mass demonstrations may build up and have different types join, including the case of contagion among countries with similar political structures (Section 5.3).

In Section 6, we discuss how what might happen after a revolution affects the possibility of having the revolution. Communication, demonstrations, and outcomes in similar countries, can not only inform people about overall number of dissenters, but also about the presence of different factions. This can inform those factions and what might happen if the revolution succeeds. This leads to trade offs, as added confidence in overall support for revolution can be undercut by worries about which faction might hold power after a revolution. Sufficient competition or distrust among factions can preclude the revolution.

In Section 7, we discuss actions of governments. In Section 7.1, we analyze how a government learns from protests and can also want to hold *counter*-demonstrations to learn about the support for and against a change in policy. We also briefly discuss what governments may do to prevent revolutions - for instance using propaganda, or redistributing income (Section 7.2).

There is a vast literature, both theoretical and empirical, on the subject of collective action and the building up of mass demonstrations. These phenomena have been analyzed from different angles, and in reference to different countries and circumstances. Indeed, there are many realities to take into account: the types of governments against which demonstra-

---

[3]See Little (2016) for a discussion of how improved technology changes both agents' knowledge of others' preferences, and also enables better coordination regarding where and when to hold protests.

tions are held, the means through which concerned agents may receive and send information, the varying objectives of agents that agree on the need to change the government, but may disagree on the alternative to set in place. The Arab Spring and the role of social networks ad cell phones have raised new important questions, activated the literature and provided opportunities for empirical tests. It is not possible to survey the literature on the subject here, but let us discuss some key references with more in the text as we proceed.

An early precursor on coordination games is Granovetter (1978). Other important studies of collective action and mobilization build upon the herding literature of Banerjee (1993) and Bikhchandani, Hirshleifer, and Welch (1993). Some of these papers examine sequential observations and how these affect voting, a politician's decision, or a collective action (e.g., see Chwe (1999), Lohmann (1993, 1994ab, 2000), Bueno de Mesquita (2010), Kricheli, Livne, and Magaloni (2011), Loeper, Steiner, and Stewart (2014), Little (2016), Battaglini (2016), Shadmehr and Bernhardt (2016ab)). The importance of information is central to all of these papers. At a high level, there is a common theme that there can be inefficiencies in outcomes due to imperfect information aggregation. The closest overlap is with Kricheli, Livne, and Magaloni (2011) who analyze a two period model in which the first period turnout informs second period activists about whether they should try a revolution.

Thanks to the simplicity of our model – which admits a rich study of a number of different issues regarding how and when information enables or inhibits a revolution, all within one model – we can study each of the four issues of information that we outlined above in a unified manner, which shows under what conditions these help or hinder ultimate revolutions. This adds to the understanding of the role of information in collective action. The unraveling argument that we present here is different from the herding and other arguments in the literature about information failures. Also, our analysis about the impact of homophily, of the size of the extremist population, and the uses of counter-demonstrations, are new to the literature.

2   A Static Model as a Building Block

We begin by describing a one-shot model in which a population must simultaneously decide whether to participate in a protest, strike, or revolution, etc., in ignorance of other agents' types. We present an analysis of this model of collective action first, since it is a useful benchmark and building block for our main results.

2.1   The Players

A continuum of citizens of mass 1 are indexed by $i \in [0, 1]$.
They have a choice to participate in a revolution or strike, etc.

In terms of the basic model, we will use the term 'revolt' but the model obviously has many applications.

The collective action is successful if at least a fraction $q \in (0, 1]$ of the population participates. If fewer than q participate, then the action fails.

2.2 Uncertainty

$\omega \in I\!R$ is the state of the world, which can encode information about the value of the revolution and what fraction of the population would gain from the revolution, and so forth.

There is a prior distribution over ω, denoted G - and agents do not directly observe ω.

$\theta_i \in I\!R$ is the type of agent i, which is the private information of that agent.[4]

The distribution over types depends on the state of the world and is denoted $F(\theta_i | \omega)$. We treat these as if they are independent across agents conditional upon the state, which is technically convenient but has some measurability issues that are easily handled as the limit of a finite model.[5]

We assume the standard ordering property on information:[6] conditional upon θ_i, the distribution on ω and others' types are both increasing in θ_i in the sense of strict first order stochastic dominance. Thus, higher types of an agent lead that agent to expect higher types of other agents.

2.3 Payoffs

An agent gets a value from the the revolution as a function of whether it is successful or not and whether the agent participates or not. All of these payoffs can be type and state

---

[4]We could allow the states and types to be multidimensional and more complicated. The advantage of one dimension is that what we ultimately care about is whether an agent is sufficiently unhappy with the government would revolt. More dimensions would involve partial orders, but the story would basically be the same - some people are unhappy enough to revolt and others are not, and the agents are trying to learn about the relative fractions and potential for success.

[5]For a discussion of the issues of a continuum of agents having independent observations see Feldman and Gilles (1985) and Judd (1985). In our model, the independence is not really needed, and so a very easy way of formalizing the signals for our purposes is as follows. Uniformly at random, draw i_0 from $[0, 1]$ - this will be the agent who gets the lowest signal in society. Then let $\theta_i = F^{-1}(i - i_0 | \omega)$, where $F^{-1}(\cdot | \omega)$ is the inverse of $F(\theta_i | \omega)$, and we take $i - i_0$ modulo 1, so that if $i < i_0$, then we set $i - i_0 \equiv i + 1 - i_0$. So, we randomly pick an agent to have the lowest signal, and then just distribute the signals then in a nondecreasing way for the rest of the agents with higher labels, and then wrap around beginning again at 0. This results in the right distribution of types without any measurability issues and the independence of types is not needed for our results, as agents only care about the population behavior rather than any particular agent's behavior.

[6]See Milgrom (1981).

6

dependent, and are given by the following table.

$$
\begin{array}{ccc}
 & Success & Failure \\
Participate & a(\theta_i, \omega) + V_i(\theta_i, \omega) & b(\theta_i, \omega) - C_i(\theta_i, \omega) \\
NotParticipate & a(\theta_i, \omega) & b(\theta_i, \omega)
\end{array}
$$

Here, $a(\theta_i, \omega)$ is the value that an agent gets if the revolution is successful, regardless of whether the agent participates or not, and this can depend on the agent's type and the state. Similarly, $b(\theta_i, \omega)$ is the value that an agent gets if the revolution fails, regardless of whether the agent participates or not, and this can depend on the agent's type and the state. The values, $V_i(\theta_i, \omega)$ and $C_i(\theta_i, \omega)$ then are the additional value and cost that an agent gets from participating in the revolution as a function of whether it is successful or fails. Generally, C_i will be negative, but V_i could be positive or negative.

Note that this is strategically equivalent to the following payoff matrix:

$$
\begin{array}{ccc}
 & Success & Failure \\
Participate & V_i(\theta_i, \omega) & -C_i(\theta_i, \omega) \\
NotParticipate & 0 & 0
\end{array}
$$

The strategic equivalence is due to the fact that the only thing that motivates an agent to participate is the difference that they experience from participating or not, as a function of whether the revolution is successful or not.

Since V_i can already encode relevant heterogeneity in the population via θ_i, from a strategic perspective only V_i/C_i matters and so it is without loss of generality for the strategic analysis to normalize the model so that $C_i = C > 0$ for all i. We still keep C as a variable, as we wish to consider cases in which a government adjusts the penalties for participating in a failed protest/revolution.

We presume that V_i is symmetric across agents - depending on their identity only via their type and thus drop the subscript i. We take V be nondecreasing in θ_i, ω, and increasing in at least one of the two arguments.

Thus, we consider games of the form:

$$
\begin{array}{ccc}
 & Success & Failure \\
Participate & V(\theta_i, \omega) & -C \\
NotParticipate & 0 & 0
\end{array}
$$

Let us mention two canonical cases:

2.3.1 (Correlated) Private Values

One case of interest is that of "private-values" so that $V(\theta_i, \omega)$ depends only on θ_i. In this case it is without loss of generality (adjusting distributions) to set $V(\theta_i, \omega) = \theta_i$, and so

payoffs are

$$
\begin{array}{ccc}
 & Success & Failure \\
Participate & \theta_i & -C \\
NotParticipate & 0 & 0
\end{array}
$$

An interpretation of this case is that each citizen knows how unhappy he or she is with the government - which is the θ_i. Here, the state of the world ω captures how unhappy the overall population is via the distribution of θ_i's. Agents, via Bayes' rule, can infer how unhappy the rest of the world is by inference given that higher states, ω's, lead to a higher distribution over θ_i's. So, if an agent is very unhappy, then she infers that it is likely that ω is high and so it is then likely that other agents are unhappy too.

2.3.2 Common Values

Another case of interest is where $V(\theta_i, \omega)$ depends only on ω. In this case, if preferences are symmetric, then it is without loss of generality (adjusting distributions) to set $V(\theta_i, \omega) = \omega$, and so payoffs are

$$
\begin{array}{ccc}
 & Success & Failure \\
Participate & \omega & -C \\
NotParticipate & 0 & 0
\end{array}
$$

This case is one in which agents do not really know whether they would like to have a successful revolution – that is governed by a state ω. For instance, agents might not know how competent or corrupt the government really is, or what might replace it. Each agent has a signal θ_i which is some noisy information about the state, and so they must infer ω via Bayes' rule from their own types.

For our purposes, it is not really important which formulation we use as they all have similar effects: agents with higher θ_i's are more optimistic that there is a high payoff from participation and that other agents feel the same. So, they all have the same basic structure of equilibria: agents with types or signals (θ_is) above some threshold participate, and others do not. Thus, we first state that general result, and then we specialize to the model with private values, for a clean and intuitive analysis.

2.4 Strategies and Best Responses

A strategy for player i is a function $\sigma_i : \mathbb{R} \to \Delta(\{0, 1\})$, which specifies a probability of participating, $\sigma_i(\theta_i) \in [0, 1]$, as a (Lebesgue measurable) function of an agent's type. Let σ denote the profile of strategies.[7]

---

[7]We work with strategies that are also Lebesgue measurable as a function of the agents' labels. Generally, the equilibria will naturally depend only on agents' types and not their labels, and so this is not really a restriction.

Let $p_\sigma(\theta_i)$ denote i's beliefs that at least a fraction q of the other agents will participate, conditional on other players playing according to σ and the agent seeing θ_i.

Given the continuum, an agent is never pivotal in determining whether there is a fraction of at least q of the population who participate, and so this is a straightforward calculation.

Then expected payoff to participation is then

$$p_\sigma(\theta_i)E[V(\theta_i,\omega)|\theta_i] - (1 - p_\sigma(\theta_i))C,$$

and the payoff from non-participation is 0, and so it is a best response to participate if and only if

$$\frac{E[V(\theta_i,\omega)|\theta_i]}{C} \geq \frac{1 - p_\sigma(\theta_i)}{p_\sigma(\theta_i)} \quad \text{equivalently} \quad p_\sigma(\theta_i)E[V(\theta_i,\omega)|\theta_i] \geq (1 - p_\sigma(\theta_i))C \quad (1)$$

Note that, given the ordering of types and preferences, $\frac{E[V(\theta_i,\omega)|\theta_i]}{C}$ is strictly increasing in θ_i.

2.5 Equilibria

We examine Bayesian equilibria of the game. Later in the paper, when we consider dynamic versions of model, we examine weak perfect Bayesian equilibria, which reduce to Bayesian equilibria in a one-shot game. So, whenever we say 'equilibria' we are referring to weak perfect Bayesian equilibria.

2.6 Existence

As this is a coordination game, there often exist multiple equilibria. For instance, nobody participating is always a strict equilibrium: if none of the other agents participate then the revolution will surely fail and so it is a best response not to participate. However, in many cases there also exist participatory equilibria.

In particular, we focus on the class of equilibria in which agents play monotone strategies: their probability of participating is non-decreasing in θ_i. Given the increasing preferences and ordering on information, such equilibria always exist. Nonetheless, there do exist other equilibria, although for generic distributions these will be the only equilibria.[8]

---

[8]For an example of a non-monotone equilibrium consider a common values setting with $\omega = 2, 3$ with equal probability and $C = 1$; and such that $\theta_i = \omega$ so that all agents know the state. In this case, regardless of q, there is always a 'best' equilibrium in which all agents participate, and there is a worst equilibrium in which no agents participate, in either state. However, there is also a non-monotone equilibrium in which all agents participate if $\theta_i = \omega = 2$ and none participate if $\theta_i = \omega = 3$.

PROPOSITION **1** *Symmetric and monotone equilibria exist. Each monotone equilibria can be described by a single threshold t (the same for all agents), such that an agent participates if $\theta_i > t$ and not if $\theta_i < t$. Monotone equilibria are all symmetric up to the possible mixing that occurs at t. Monotone equilibria can be ordered by their thresholds, with ∞ always being an equilibrium threshold.*

This follows from an application of Tarski's fixed point theorem, which establishes that equilibria form a complete lattice, which here is just ordered in terms of the thresholds. Given that the proof is standard, we omit it. The symmetry follows from the continuum of agents who have the same priors, and the fact that payoffs are monotone in types and states, so that higher types lead have higher expected payoffs from participation conditional on success.

So, we can represent monotone equilibria by thresholds t, such that an agent participates if $\theta_i > t$ and not if $\theta_i < t$. In cases with atoms in the distribution it is possible to have mixing at t.

In what follows, when we say 'equilibrium', we refer to a symmetric monotone equilibrium.

In the discrete private values setting, for the results and examples that we analyze below, there are generally at most two pure strategy equilibria: one in which nobody participates and another with maximum participation. In the continuum setting, there can exist a mixed strategy equilibrium if one allows for specific tie breaking rules when the fraction of demonstrators is exactly q; however, such an equilibrium is unstable - as any slight perturbation leads to the non-participation or the most participatory equilibrium. Since the comparative statics on the non-participation equilibrium are trivial, in what follows we focus our attention on a (most) participatory equilibrium exists and its properties.

3 A Discrete Private Values World

As we stated above, the main insights and the workings of equilibria are broadly similar for the private and common values worlds. The private values model focuses attention on the strategic aspects of the uncertainty - people are sure about whether they themselves are miserable, but unsure about how many others are willing to act, and so this makes for the clearest and most natural interpretations of the results.

3.1 The Simplified Setting

Consider a simplified version of the model in which types are either $\theta_H > 0$, called H types, or $\theta_L < 0$ called L types.

In particular, suppose that either $z > q \geq 1/2$ of the population are H types which happens with probability π, which we call the "High" state; or $1 - z < q$ of the population

are H types, which happens with probability $1 - \pi$, and we call the "Low" state.[9] This is pictured in Figure 1
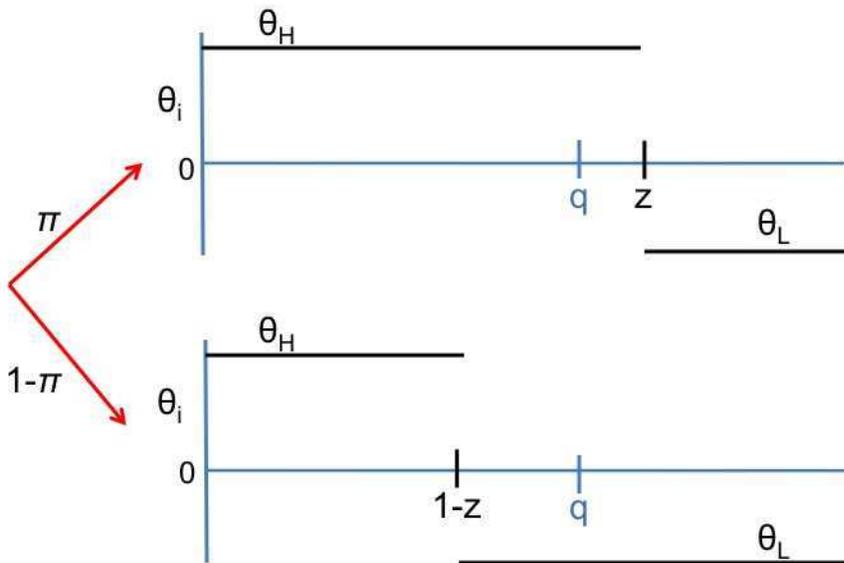

Figure 1: Two states, with the High state having probability π. The High state has more of the θ_H types (a fraction $z > q$) and the Low state has more of the θ_L types.

If a player is an H type, by Bayes' Rule her conditional probability on the "High" state is

$$\frac{\pi z}{\pi z + (1 - \pi)(1 - z)}$$

3.2 The Multiplicity of Equilibria

Again, as this is a coordination game, there can be a multiplicity of equilibria. That multiplicity has been extensively studied in the global games literature (e.g., see Angeletos, Hellwig and Pavan (2007)) and in the protest literature (Bueno de Mesquita (2010)).

In the private values settings that we study here the multiplicity is very simple. There is always an equilibrium in which nobody participates. This is straightforward and so there is no point in saying more about it. In some situations there are also participatory equilibria. There is one in which all of the H types participate. This is also a strict equilibrium when it exists, and is the one that is of interest to us.

There can also exist variations on mixed strategy equilibria in which H types are exactly indifferent between participating and not.[10] In the world of a continuum, in order to get these

---

[9]It is direct to see that the symmetry between the fraction being z and $1 - z$ simplifies calculations, but does not alter the intuition behind our results.

[10]Either agents would have to mix, or it would have to be asymmetric.

equilibria to exist, one also has to put in place a rule that indicates the exact probability of the revolution being successful if exactly q agents show up (see Jackson, Simon, Swinkels, and Zame (2002)). Those equilibria are quite unstable, as slight perturbations of the actions lead best replies to converge either to the all participate or no participate equilibria.

Thus, in the game that we analyze here, the equilibrium structure is quite simple: the only stable and strict equilibria are the ones in which nobody participates or else in which all the H's participate. Therefore, our focus is on the one in which all H's participate as the other one is trivial.

We could follow the global-games literature to refine down to one equilibrium, but which of these two equilibria we would end up with would depend on how the uncertainty was specified (see Weinstein and Yildiz (2007)). So, there is little point in our setting in putting in such a refinement. Thus, we analyze the existence and structure of the participatory equilibrium, and note that the nobody-participates-equilibrium is trivial, and the mixed strategy equilibria are unstable.

Given that the pure strategy participatory equilibria that we analyze are strict, robustness to a variety of perturbations of the game follows directly.

This also keeps our model easy to understand and allows us to analyze many variations on it, which would become cluttered with no added insights if we added continuous signal spaces.

3.3   Equilibria with a Revolution

Let us examine when there exists a (pure) equilibrium with a revolution.

By (1) there exists an equilibrium where the H types all show up if and only if:

$$\theta_H/C \geq \frac{(1-\pi)(1-z)}{\pi z} \quad \text{equivalently} \quad \pi z \theta_H \geq (1-\pi)(1-z)C. \tag{2}$$

The existence of an equilibrium in which the H types have a high enough belief that they expect a positive payoff from showing up is pictured in Figure 2, as a function of π and z.

We emphasize that there are two requirements for the existence of an equilibrium in which H types all participate:

- it must be that $z \geq q$, as otherwise even in the High state there would not be enough H types to be successful, and

- it must be that beliefs of the H types put a large enough weight on the chance of success so that they are willing to participate, which is true if and only if $\pi z \theta_H \geq (1-\pi)(1-z)C$.
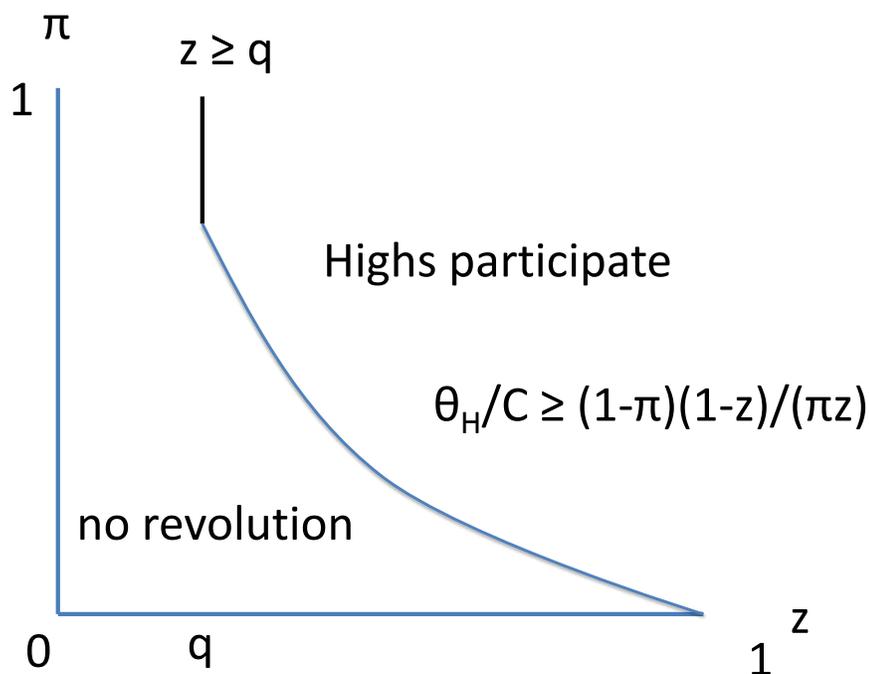
Figure 2: There is an equilibrium in which the H types participate if and only if the prior π and the correlation z are high enough.

The first constraint is that z lies to the right of the vertical segment at $z = q$ and the second constraint is that π and z are above the level curve at which $\theta_H/C = \frac{(1-\pi)(1-z)}{\pi z}$. If and only if both of these are satisfied does there exists an equilibrium in which H types participate. There always exists an equilibrium in which nobody participates.

The model produces some intuitive comparative statics, that follow directly from equation 2 and are pictured in Figure 3. We see that the range of values of π and z for which there is a revolutionary equilibrium shrinks as we decrease θ_H and/or increase C.[11]

There are H types in either state, and they act based on their beliefs conditional on the fact that they are a H type. So, they know that they still face a chance of failure as it is possible that it will be the Low state and there are just not enough H types to succeed.[12] So, H types participate but the revolution still fails whenever it happens to be the Low state; and thus the likelihood of success increases as the likelihood of the High state, π, increases.

---

[11]See Kricheli, Livne, and Magaloni (2013) for evidence that increased costs lead to fewer protests, but ones that are more likely to be successful when they occur.

[12]This is provided $z < 1$, as otherwise (if $z = 1$) types are fully correlated with the state and fully revealing and the analysis becomes trivial.
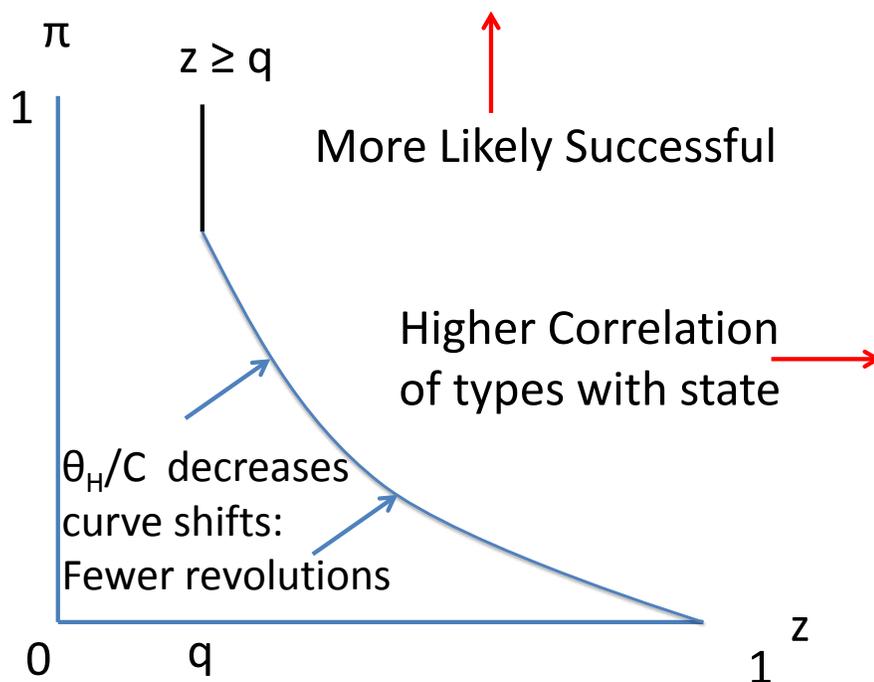
Figure 3: The range of values of the prior belief on the High state, π, and the correlation between types and the state, z, shrinks as the cost of the revolution increases or the value to H types from participating decreases. Also, as π increases, the likelihood of success increases, and as z increases there is a better match of the H types with the state.

Also, as z increases there is a higher correlation of the H types with the state: there are more H types who show up in the High state when the revolution is successful, and fewer who show up in the Low state when the revolution fails.

It is important to note that the change from no revolution to a revolution is discontinuous: as π, z, and θ_H/C pass a threshold we can go to a regime that experiences no revolutions to one that can have (large) ones.

4   Communication Prior to a Revolution

With the basic model in hand, we now expand to analyze how information affects the possibility of having revolutions.

We begin with the question of what happens when people get to see some information about how others feel about the regime.

4.1 Each Agent Sees One Other Agent's Type

We first consider what happens if each agent get to see one other agent's type, where that agent is chosen uniformly at random. So, each agent gets to talk to one other agent in the society and learn that agent's type.[13] This provides additional information to the agents, since now they have two signals about the state rather than just one.

First note, again, that since $\theta_L < 0$, we only have to analyze the H type's incentives in order to characterize equilibria, since L types never participate.

If a H type sees another agent of a H type, then by Bayes' Rule, the agent's belief that the state is 'High' is

$$\frac{\pi z^2}{\pi z^2 + (1-\pi)(1-z)^2}.$$

If a H type sees that the other agent is a L type, then by Bayes' Rule, the agent's belief that the state is 'High' is

$$\frac{\pi z(1-z)}{\pi z(1-z) + (1-\pi)z(1-z)} = \pi.$$

First, let us consider the case in which the prior belief is so high that even if an H type meets an L type, the H type is still convinced enough of the High state that the agent is willing to go to the revolution. By (1) there exists an equilibrium where the H types show up *regardless of signals* if and only if:

$$\pi\theta_H \geq (1-\pi)C. \tag{3}$$

The above condition is more demanding than our previous equilibrium in the absence of any signals, since it is asking whether H types go even when getting another signal which is a low one. Another possibility is that only some of the H types are now willing to show up - those who get to see another H type.

For there to exist an equilibrium in which the H types only show up when they see another H type, two things are necessary: one is that they are sufficiently convinced of the High state that they are willing to show up: by (1) this requires that

$$\frac{\theta_H}{C} \geq \frac{(1-\pi)(1-z)^2}{\pi z^2} \quad \text{equivalently} \quad \pi z^2 \theta_H \geq (1-\pi)(1-z)^2 C. \tag{4}$$

The second requirement is that there have to be enough of these H types who also see other H types (in the High state) for the revolution to be successful. This requires that

$$z^2 \geq q,$$

---

[13]Note that it is incentive compatible for agents to tell each other their types, and so it is without loss of generality to simply assume that types are observed when two agents meet. This does depend on the payoff normalization in our model. If we allowed the θ_L types to still prefer the protest to be successful, but not want to participate, then that would induce them to lie. Alternatively, as long as people can observe each others' types as depending on some demographic variables (e.g., employment, income, etc.) then the types would be at least partly observable.

since z^2 of H types will see another H type in the High state.

These two different sorts of equilibria are pictured in Figure 4.
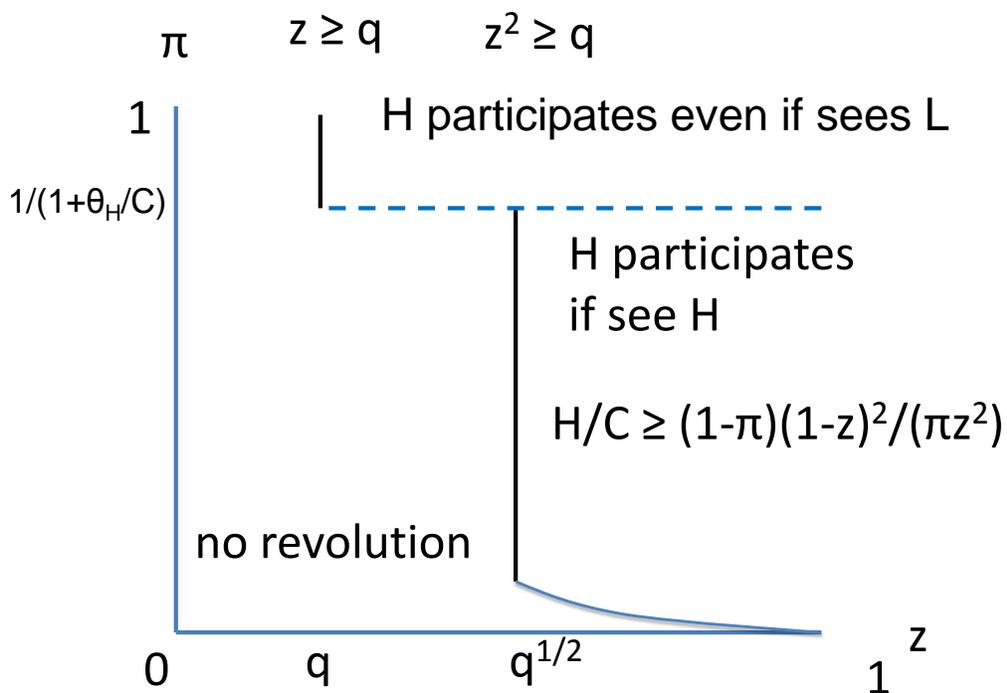

Figure 4: There are two regions of equilibria.

Comparing this to the no information case, Figure 5 shows the equilibrium structures for the two settings:

In Figure 6 we see that information helps the revolution when π (the prior prob of the High state) is low and when types are sufficiently correlated with the state and so seeing another H type is very informative. In contrast it hurts the revolution when the correlation between types and the state is lower and so many people can see others that have low signals and become discouraged and so that even types who see others who are high know that too few people will show up for it to be successful.
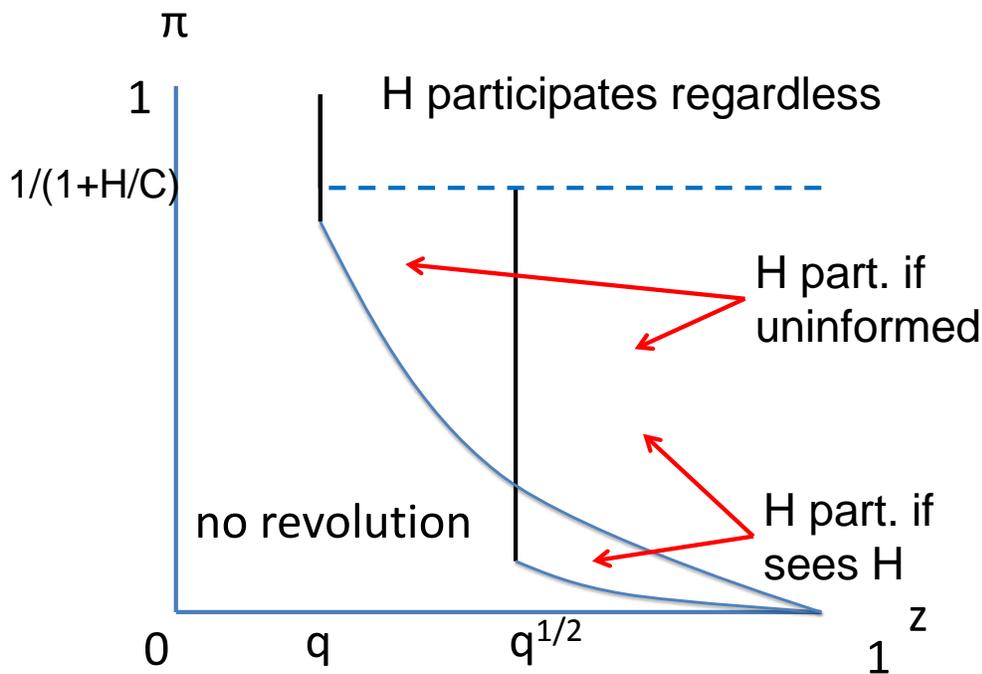
Figure 5: Five different regions: no revolution, always a revolution regardless of what info is, only a revolution if don't see signals, only a revolution if see signals and both are high, revolution if see signals or not - but only H types that see another H type show up.

So, to summarize, we see four different possibilities:

- With a high enough prior on the High state, there exists an equilibrium in which the H types to show up regardless of what they observe from the other type, in which case it would have been an equilibrium for them to show up without seeing another agent's type. Here the equilibrium is the same as not observing anything, as the prior is strong enough so that information does not influence the agents' decisions. This happens if (3) holds (and $z \geq q$). In this case, it also would have been an equilibrium for all H types to show up without any information, and so there is no change in equilibrium structure in this parameter region.

- Next, there is a region in which there was an equilibrium for H types to show up without any information, but with information it is no longer an equilibrium for the H types to show up *even if they see another H type*. Here the equilibrium fails not because those who see two H types are not convinced enough about the High state, but instead because they know that they are too small a fraction of the society to be successful. Here information is damaging for the H types as it would have been an equilibrium for them to show up if they did not see another agent! In this region the H
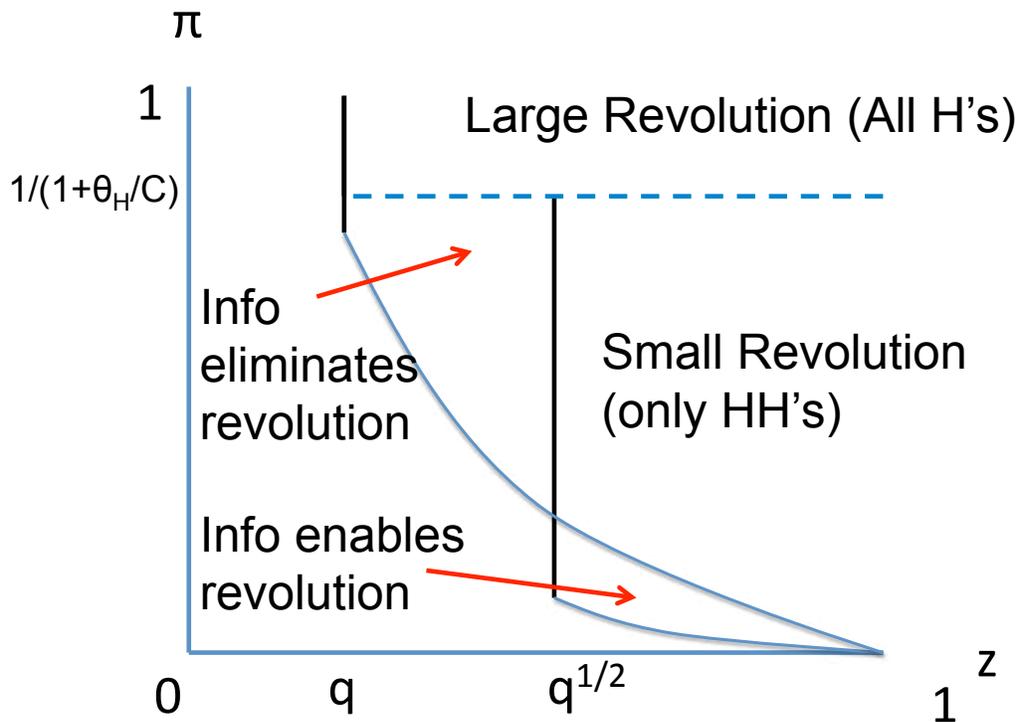
Figure 6: Sometimes information aids the revolution and other times it blocks it

types are ex ante worse off and the L types are better off. This happens if $z^2 < q < z$ while (2) holds.

- Next, there is a region in which there is an equilibrium in which the H types show up if and only if they see that the other agent is an H type. This breaks into two pieces.

  – One part of this region is where it would also have been an equilibrium for them to show up without seeing anything. Here the equilibrium is now changed, as fewer H types show up in the High state and also in the Low state, but the revolution is still successful in the High state and not the low. The H types are better off ex ante, and the L types are indifferent. Ex post, some H types are better off and others worse off in this setting than in the no information case, and overall they are better off ex ante. This happens if $z^2 > q$ and (4) holds, as does (2), while (3) do not.

  – The other part of this region in where it is an equilibrium for the H types to show up if and only if they see that the other agent is of the H type, but it would not have been an equilibrium for them to show up without seeing anything. Here the equilibrium is now changed, as seeing the other type enables H types to show up as they are now surer of the state, while without the information they would not have been able to have a revolution. Again, the H types are better off ex ante,

18

and the L types are worse off. This happens if $z^2 > q$ and (4) holds, while (2) does not.

We should point out that the basic intuition that having some information can disrupt a revolution, as it will inevitably discourage some higher types, extends to more general payoffs. For instance, the same result holds in a common values version of the model, as well as hybrids. Basically, seeing a low type lowers the high type's beliefs about the state regardless of the specifics of private versus common values, and so makes her more pessimistic. Knowing that some high types will be discouraged then means that even the more optimistic agents now know that their numbers are reduced.

4.2   Homophily and Networks

The previous analysis considers a case in which an individual meets another person chosen uniformly at random from the population. However, as we know, in many contexts people that we talk with are those around us in our networks and local communities. People are substantially more likely to interact with others who are similar to each other, not only in some base characteristic, but also in preferences and political views.[14]

To capture this, let us consider a variation on the above setting in which we allow for homophily. A very easy way to adjust the model to include a bias in meetings, is to allow that a fraction $h \in [0, 1]$ of matches that would be have been between highs and lows under uniform random are instead with highs matched to highs and lows to lows.

If $h = 0$ then there is no homophily and matching is uniformly random, while if $h = 1$ then highs always meet highs and lows always meet lows.

In terms of information, when $h = 1$, there is no information in a partner's type as it is then the same as the agent's own type regardless of the agent's type. The informativeness of the signal is highest when $h = 0$. However, given the non-monotonicities in equilibrium, the effect of homophily on equilibrium can be ambiguous, as we now show.

In particular, the probability of an H type seeing another H type with homophily $h \in [0, 1]$ is $z^2 + z(1 - z)h$ in the High state and $(1 - z)^2 + z(1 - z)h$ in the Low state.

This leads to a new constraint for the equilibrium in which a H type is willing to participate if and only if seeing another H type. These are straightforward variations on the previous analysis, just using Bayes' rule. The necessary conditions for an equilibrium (presuming that a θ_H meeting a θ_L type would not participate) are then:

$$\frac{\theta_H}{C} \geq \frac{(1 - \pi)[(1 - z)^2 + z(1 - z)h]}{\pi[z^2 + z(1 - z)h]},$$

and

$$z^2 + z(1 - z)h \geq q.$$

---

[14]For background on this empirical observation, termed "homophily", see McPherson, Smith-Lovin and Cook (2001) and Jackson (2008).

Note that the second inequality gets easier to satisfy as h increases, while the first one gets harder to satisfy as h increases: this is the tradeoff as homophily is increased. Homophily decreases information, making the individual incentive to participate harder to satisfy, but also leads to fewer agents who are discouraged by meeting L types. Which effect dominates depends, again, on the relative prior and correlation of types with the sate.

This leads to the adjustment in the equilibrium structure as pictured in Figure 7:
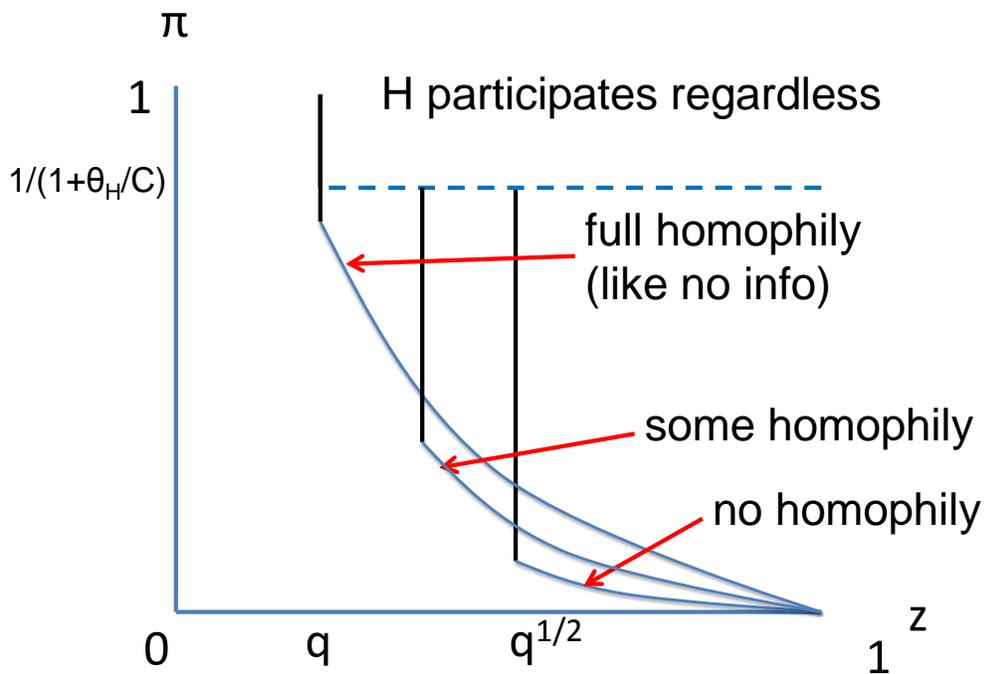
Figure 7: Homophily (assortativity in meetings) changes the equilibrium structure.

So, we see that higher homophily increases the region of having a revolution if the prior is high enough, since more highs will see high signals and be willing to join, but higher homophily reduces the region for low priors and high z since it decreases the information contained in a meeting.

4.3   Seeing Many Other Agents' Types

Next, we consider what happens in the same setting when agents get to see many other randomly chosen agents' types.

PROPOSITION 2 *For any $z \geq q$, π, and θ_H/C, there exists a number of signals above which there is an equilibrium which involves protests conditional upon sufficient fraction of H types being observed. Moreover, as the number of others observed increases, the fraction of H types participating in the 'High state goes to 1 and the fraction of H types participating in the 'Low state goes to 0: the protest is perfectly effective in the limit.*

In the proposition, by the law of large numbers agents will eventually be sure of the state, and so there exists an equilibrium in which agents who are H types show up whenever their posterior is above a threshold, and in the limit they are almost always successful.

The interesting aspect, putting this result together with our analysis of just seeing one other agent's type, is that information can be non-monotonic: small amounts of information can be disruptive, while large enough amounts of information are always enhancing.

As such, we might expect that technological advances that allow agents to learn about the opinions of greater numbers of others to eventually lead to more accurate demonstrations. As people learn about greater number of others the correlation of the size of the demonstration with the state will increase. This is consistent with empirical background on this sort of effect, as in Breuer, Landman and Farquhar (2012) and Farrell (2012), as well as Manacorda and Tesei (2016), Pierskalla and Hollenback (2013), and Steinert-Threlkeld, Mocanu, Vespignani and Fowler (2015).

Demonstrations of nontrivial size may become more or less frequent depending on the parameter region, but then much more likely to be successful when of large size. We can solve for some aspects of the equilibrium in more detail.

An individual now gets to see m random other individuals' types. We now can see how many signals they must see before they are willing to participate.

Let t be the threshold so that if at least t other H types observed, then they are willing to participate.

There are two constraints that need to be satisfied in order to have an equilibrium where some people participate. One is that some agents end up with high enough beliefs that it is the High state. The other is that enough of them participate to be successful. So, collective action is feasible only if these threshold intervals overlap.

Let us examine first the lower bound $\underline{t}(m)$ on the threshold that can convince an agent to participate.

If a player is of type θ_H and sees t out of m other H types, then the conditional probability on the state that z of the population are of the H type is

$$\frac{\pi b(t+1, m+1, z)}{\pi b(t+1, m+1, z) + (1-\pi)b(t+1, m+1, 1-z)}$$

where $b(t, m, z)$ is the binomial probability of seeing t positives out of m trials that are positive with probability z. So to get an agent to act (presuming the agent expects success conditional upon the High state) requires:

$$\theta_H/C \geq \frac{1-p_i}{p_i} = \frac{(1-\pi)b(t+1, m+1, 1-z)}{\pi b(t+1, m+1, z)} = \frac{(1-\pi)(1-z)^{2t+1-m}}{\pi z^{2t+1-m}}.$$

21

Solving this with equality allows us to deduce $\underline{t}(m)$

$$\underline{t}(m) = \frac{m-1}{2} + \frac{\log\left(\frac{\theta_H \pi}{1-\pi}\right)}{2C \log\left(\frac{1-z}{z}\right)}.$$

Next, we solve for $\bar{t}(m)$. In order to have the fraction of agents who show up be at least q conditional upon the High state it must be that

$$(1 - B(t-1, m, z))z \geq q,$$

where $B(t-1, m, z)$ is the c.d.f. of the binomial distribution (so the probability that there are $t-1$ or fewer other H types out of the m observed when drawn with probability z). Thus,[15]

$$\bar{t}(m) = B_{m,z}^{-1}\left(1 - \frac{q}{z}\right) + 1.$$

Note that in the limit, $\bar{t}(m) \to zm$, while $\underline{t}(m) \to m/2$, and so eventually $\bar{t}(m) > \underline{t}(m)$.

As a numerical example, let $z = 2/3$ and $q = 1/2$ $\theta_H \pi/[C(1-\pi)] = .8$. Here the non-montonicity of information is clear: we have an equilibrium with H types participating if $m = 0$ or if $m = 2$, but *not* if $m = 1$.

Note that if we add homophily to this setting, then this slows the rate of informativeness of our signals. If I get to meet a hundred people, but more than ninety percent of them are my friends and so very similar to me in terms of political views, then that is almost like meeting just ten people. Also, to the extent to which people do not fully understand the homophily around them, they will believe that the world is more like themselves and so not properly update. Thus, the interaction rates before people really learn about the world might need to be very high, and to have low homophily in order for people to become well-informed.

5 Dynamics: A Two-Period Version of the Model

So far we have examined information revelation as agents meet some other people from the population. Another important informational channel is having mass demonstrations in which agents protest. These can be important precursors to a strike or revolution as they signal information in a much broader and more revealing way than people just seeing the preferences of a few friends.

---

[15]Here, the inverse of B is rounded downwards, so it is the largest value of t for which $B(t-1, m, z) < 1 - \frac{q}{z}$, which then assures that t is the smallest values for which the chance that at least t H types are observed is at least q/z.

5.1 The Informational Role of Demonstrations before a Revolt

To study the role of demonstrations, we enrich the model so that there are two periods and three types. The types are $\theta_L, \theta_M, \theta_H$. Now the values are purely private, and the highest value types simply are more disadvantaged by the current government, and the moderates types would also prefer to overthrow the government if it is possible, but are harder to convince to join the revolution since they are not as dissatisfied as the higher types.

There are two states. In the 'High' state $1 - z$ of the population are θ_L and $z/2$ are θ_M and $z/2$ are θ_H, while in the 'Low' state z of the population are θ_L and $(1 - z)/2$ are θ_M and $(1 - z)/2$ are θ_H.

So, this is exactly the same as our first model, except that we have split the H types equally into moderates and highs. This allows us to see the value of having protests before the revolution. This is pictured in Figure 8
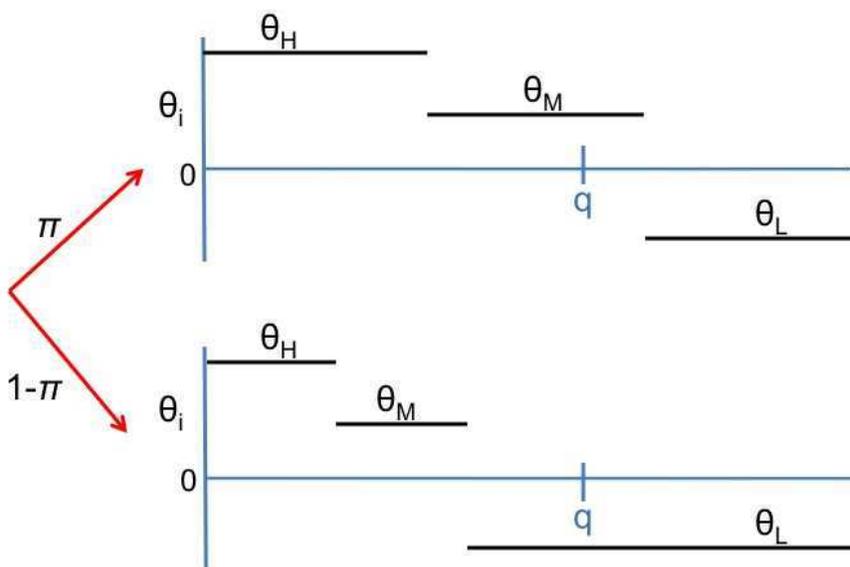
Figure 8: Two possible states. By seeing how many θ_H types turn out at a protest, the state is revealed, which can enable a revolution.

So, there is a first period in which the population can hold a demonstration, and the a second period in which they can hold the revolution. They can skip the first period if they wish, but it signals information about the state.

Let the cost of having participated in a protest or revolution if the revolution is not ultimately successful depend on the period, and for the first period be c and the second period be C.

Let us consider a case in which

$$\frac{\theta_M}{C} < \frac{(1 - \pi)(1 - z)}{\pi z},$$

23

but $z \geq q$.

So, without any additional information, the moderates are too frightened/pessimistic to participate in the revolution.

However, note that if

$$\frac{\theta_H}{c} \geq \frac{(1-\pi)(1-z)}{\pi z},$$

then it is possible to have the revolution.

The highs are willing to demonstrate in the first period. If $z/2$ of them show up, then the moderates learn that it is the High state and the revolution takes place in the second period. If only $(1-z)/2$ of them show up in the first period, then the demonstration is a failure and there is no revolution in the second period.

This illustrates the possibility of having successive demonstrations, where people learn about how many people are dissatisfied by observing the size of the turnout, and more extreme individuals protest earlier, enabling more moderate types to assess the state and join later if things look strong enough.

It should be clear that with richer heterogeneity one could build richer versions in which protests gradually escalate over time, and which several successive protests are needed, over time and/or geography, before sufficient certainty is reached to hold a successful revolution.

5.2 The Difference Between Polls and Demonstrations

One question that we have not yet addressed, but is important, is why one needs demonstrations at all in a world where people can hear about how others feel via polls and/or social media. In the above example, why do they still need to turn out at a demonstration in order to convince the population to revolt rather than just expressing their preferences in a poll or on some social platform?

The answer is that demonstrations involve costs - and so agents must be sufficiently willing to participate to overcome those costs. Having many agents willing to pay those costs can signal to others that there is enough of the population willing to take costly action, that the revolution has a chance of succeeding. In contrast, polls and social media may involve much lower costs, and so agents simply saying that they support change does not indicate that they would be willing to act if needed. This is illustrated in the following example.

Suppose that payoffs are of the following form:

	Success	Failure
Participate	θ_i	$-C$
NotParticipate	a_i	0

Here, agents who have $a_i > 0$ and $\theta_i > -C$ would like to see the revolution succeed. However, those who have $a_i > \theta_i$ have a dominant strategy not to participate. These are non-activist people who prefer to have others participate, but still would like to see change.

In this sort of setting, if one holds a poll to see who favors change, the agents who have $a_i > \theta_i > -C$ and $a_i > 0$ will say that they favor change. However these people cannot be counted upon to show up for the revolution when it is needed. Thus, the poll does not differentiate between people who favor change and those who support it enough to do something about it. In contrast, a demonstration can be costly to show up for, and so can screen out the non-activists and give a more accurate assessment of agents who are willing to act for change.

Thus, demonstrations can be essential for successful further action and change in ways that polls and other sorts of media posting and cheap-talk might not.[16]

5.3   The Arab Spring

Another variation on the above example is one in which there are not two periods, but instead two correlated countries. If one country has a large enough turnout in its revolution, then other country's population may learn about their own state and revolt as well.

Let us consider our original setting, but the only difference is that there are now two countries. The have the same probability of a High state, designated by π, but differ in the value and costs to H types, and the correlation of types with the state. We use the obvious notation: $(\theta_{H1}, C_1, z_1, q_1), (\theta_{H2}, C_2, z_2, q_2)$

The states of the two countries are correlated, with the correlation in High states being $\rho \geq 0$. In particular, the probability of the high or Low states for the respective countries are given by:

$$
\begin{array}{ccc}
 & High_2 & Low_2 \\
High_1 & \pi^2 + \rho\pi(1-\pi) & \pi(1-\pi)(1-\rho) \\
Low_1 & \pi(1-\pi)(1-\rho) & (1-\pi)^2 + \rho\pi(1-\pi)
\end{array}
$$

Let us suppose also that $z_1 \geq q_1$ and $z_2 \geq q_2$, so that both countries can have successful revolutions in their respective High states.

Suppose that

$$\frac{\theta_{H1}}{C_1} \geq \frac{(1-\pi)(1-z_1)}{\pi z_1}$$

but

$$\frac{\theta_{H2}}{C_2} < \frac{(1-\pi)(1-z_2)}{\pi z_2}.$$

This is a world in which country 1 is sufficiently unhappy, or convinced of the High state, that a revolution is possible for that country on its own, while country 2 fails to satisfy that constraint, and so would only be willing to revolt if they are sufficiently convinced. In fact, the data on the Arab Spring collected by Brummitt, Barnett, and D'Souza (2014), who find

---

[16]Note that in a very repressive regime - that penalizes people who even say they support change - then it would be possible for that to provide a costly signal. However, that would only work if sufficiently many people are able to express their opinions, and such very repressive regimes may also censor information about any opposition.

a significant correlation between the unemployment rate in countries and the date of first protest (e.g., Tunisia had higher unemployment than Egypt than Syria, and the first date of protests occurred in that order - and they analyze fifteen countries in total).

In this case, if country 1 holds its demonstration/revolution, then country 2 can learn about the state, provided there is sufficient correlation.

In particular, some direct calculations of the posterior conditional on success in country 1 (together with the appropriate variation of (2)) show that if

$$\rho \geq \frac{\frac{C_2(1-z_2)}{z_2 \theta_{H2}} - \frac{\pi}{1-\pi}}{\frac{C_2(1-z_2)}{z_2 \theta_{H2}} + 1},$$

then there is an equilibrium with contagion.

6 Extremists and Forecasting the Post-Revolution World

The analysis so far has been on situations in which the forecast of what might happen after the revolution does not depend on the state. The state determines whether the revolution succeeds or not, but if it is successful, then the forecast of what will happen was not state-dependent.

In many situations, however, participation may depend on what people expect to happen after the revolution – which involves their expectations of what a new government will be like.[17]

Such an enrichment of the model can add to the analysis of all of the situations we have discussed so far: meeting and learning types of other agents, observing demonstrations, and observing the outcomes from other countries. Any of these information revelations can include not only information about number of dissenters and likelihood of success, but also about the size of relative factions of potential revolutionaries and who might emerge in power after the revolution.

To explore how potential conflict after a revolution can affect the revolution, let us consider the following variation on our basic model.

Suppose that there are now three types: θ_L support the government and never want to participate, θ_M are moderates who will support a revolution, but only if they are the majority of the revolutionaries and get to impose a moderate government after a successful revolution; and extreme types θ_E, who want a revolution whenever it would be successful regardless of the next government.[18] In particular, participating moderate agents get θ_M if the revolution is successful and there are more moderates than extremists, and get $-C$

---

[17] See Shadmehr (2015) for an analysis of an endogenous agenda as part of a revolution. Our example here presumes that there is no ability to commit to what will happen after the revolution.

[18] See Acemoglu, Hassan, and Tahoun (2015) for a description of conflicts between different revolutionary groups during the Arab Spring.

otherwise. Extremists get θ_E if the revolution is successful and they outnumber moderates, $\alpha\theta_E$ if the revolution is successful and moderates outnumber extremists, and $-C$ if it fails.

In particular, moderate types prefer to participate in the revolution only if the fraction of moderate and extreme types exceeds q, but *also* only if the fraction of moderates exceeds the fraction of extreme types.

The state ω is a list, $\omega(\theta_L), \omega(\theta_M), \omega(\theta_E)$, of the fractions of the population that are of the corresponding types.

There are three states $\omega \in \{\omega^L, \omega^M, \omega^E\}$:

- Low state: $\omega^L(\theta_M) + \omega^L(\theta_E) < q$, so the revolution will fail even if moderates and extremists participate.

- Moderate state: $\omega^M(\theta_M) + \omega^M(\theta_E)$, but $\omega^M(\theta_M) < q$ and $\omega^M(\theta_E) < q$ (so the revolution will succeed if and only if both moderates and extremists participate), and moderates outnumber extremists, $\omega^M(\theta_M) > \omega^M(\theta_E)$.

- Extreme state: $\omega^E(\theta_M) + \omega^E(\theta_E) \geq q$, but $\omega^E(\theta_M) < q$ and $\omega^E(\theta_E) < q$ (so the revolution will succeed if and only if both moderates and extremists participate), and extremists outnumber moderates $\omega^E(\theta_M) < \omega^E(\theta_E)$.

There are different equilibrium possibilities depending on the prior probabilities of the states, π^L, π^M, π^E. Here we focus on the case without communication, although the extension to communication is straightforward and parallels that above.

In order for a revolution to be possible, it must be that the moderates place a high enough probability on the moderate state (conditional on being a moderate), while the extremists place a high enough probability on both the moderate and the extreme state. In particular, it is straightforward to check that the necessary conditions for having a revolution are that

$$\theta_M/C \geq \frac{\pi^L\omega^L(\theta_M) + \pi^E\omega^E(\theta_M)}{\pi^M\omega^M(\theta_M)}.$$

and

$$\theta_E/C \geq \frac{\pi^L\omega^L(\theta_E)}{\alpha\pi^M\omega^M(\theta_E) + \pi^E\omega^E(\theta_E)}.$$

This can allow a revolution to take place in both the moderate and extreme case, provided the prior on the moderate state is high enough relative to the extreme state for the moderates.

It is easy to see how this then enhances the analysis of meeting others, seeing demonstrations before a revolution, and seeing the outcome in other countries. If any of those processes reveal sufficient likelihood that it is the extreme state (or that the revolution would fail), then the moderates would no longer participate. Thus, the conditions for the revolution to succeed require sufficiently high prior information, or revelation of a high likelihood, of it being the moderate state. Again, information could be either encouraging or disruptive to the revolution, depending on the state and prior probabilities.

For example, extending Section 5, in which two stages of demonstrations can enable a revolution, we could also view that example's High types as the extremists. The composition of extremists versus moderates in the High state then matters. We could split that state into two sub-states: one in which the extremist Highs are in the majority of those who favor change, and the other in which the moderates are in the majority of those who favor change. This makes for interesting dynamics, as if the first period demonstration shows that there are too many extremist High types, then the revolution would fail, as the moderates would prefer to avoid an extremist state.[19] The equilibrium thus then only successful in the second period if enough people - but not too many extremists - show up in the first period demonstration. Similarly, if a revolution in a correlated country turns too extreme, it may discourage a nearby population from revolting.

7 Counter Demonstrations and other Government Responses

Let us discuss how governments might react to demonstrations, either after the fact or by trying to deter them.

7.1 Counter Demonstrations

Demonstrations can be useful in signalling to the government the level of support for a policy change, and counter-demonstrations can be useful in signalling the level of support for keeping the current policy.

We make this point in the context of a setting where the population consists of three groups that can differ in their preferred policies, since three groups is just enough to allow for variation in which is the most preferred policy and also to allow different sizes of demonstrations and counter-demonstrations to non-trivially signal the state.

There are three equal-sized groups in the population. The groups can either support change or no-change, which we denote by $\mathcal{C}$ and $\mathcal{N}$. The groups can also either be strong supporters or weak supporters - in terms of how much they prefer their choice to the opposite choice. We denote these by $\mathcal{S}$ and $\mathcal{W}$.

In terms of preference parameters:

- $\theta_{\mathcal{CS}} > \theta_{\mathcal{CW}} > 0$

- $\theta_{\mathcal{NS}} = -\theta_{\mathcal{WS}}$, and

- $\theta_{\mathcal{NW}} = -\theta_{\mathcal{CW}}$.

---

[19]For an illuminating but different discussion of how information revelation by a government about potential counter-policies can affect revolutions, see Shadmehr and Bernhardt (2016c).

So, a group's preference is one of four types: $\mathcal{CS}, \mathcal{CW}, \mathcal{NW}, \mathcal{NS}$. A state is listed as a triple of each group's preference type. So, for instance, $(\mathcal{CS}, \mathcal{CS}, \mathcal{NS})$ indicates that the first two groups both strongly prefer change while the third group strongly prefers no change.

With three groups and four types for each group, this leads to 64 possible states. To simplify the exposition, we focus on just four states - which capture the main ideas. Obviously, the analysis extends to including all 64 states depending on the prior probability on the various states, provided there is can be some uncertainty after a first demonstration, and sufficient likelihood that a counter-demonstration will resolve that uncertainty when it arises. The main point that counter-demonstrations can be useful for learning holds in the more complicated setting, but then specifying all of the possible priors for which this holds becomes intractable, so we just illustrate the point for one possible prior that has weight on four possibilities.

In particular, we presume that one group prefers change, one group prefers no-change, and the remaining group is the only one with that could be on either side - so there are two "partisan" groups whose direction of preference is known, just not their intensity, and one "pivotal" group which could have any preference and whose preference always determines the direction of the majority preference. We focus on 4 key states.

State 1 $(\mathcal{CS}, \mathcal{CS}, \mathcal{NS})$

State 2 $(\mathcal{CS}, \mathcal{CW}, \mathcal{NW})$

State 3 $(\mathcal{CW}, \mathcal{NW}, \mathcal{NS})$

State 4 $(\mathcal{CS}, \mathcal{NS}, \mathcal{NS})$

This is pictured in Figure 9

Let these four states be equally likely.

The optimal policy (in terms of a utilitarian goal of maximizing total welfare) is change in states 1 and 2, and no change in states 3 and 4. In states 1 and 4, all have strong preferences but either have 2/3 of the population in favor of change or in favor of no change. In states 2 and 3, there is a mixture of weak and strong preferences, but the strong preferences are always on the side of a majority, and so the preference on the majority side is stronger than the minority.

PROPOSITION **3 (Counter-Demonstrations)** *Suppose that $\theta_{\mathcal{CS}} > C/2$. Then there exists an equilibrium in which:*

- *a demonstration is held by all $\mathcal{CS}$ types.*

- *a counter-demonstration is held by $\mathcal{NS}$ types if there is a demonstration in which only 1/3 of the population shows up.*
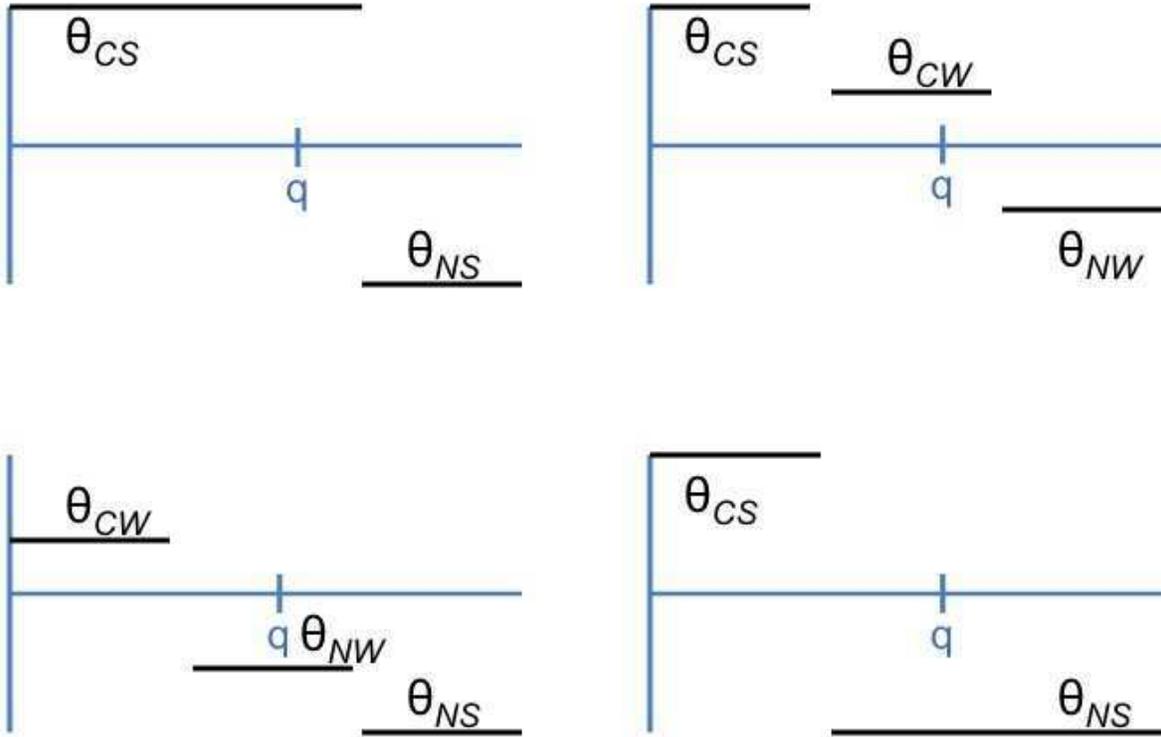
Figure 9: The four states that we focus on. In the two on the top, change is overall utility maximizing, while in the two on the bottom no-change is overall utility maximizing.

- *weak types never show up to a demonstration or counter-demonstration.*

- *there is a successful revolution (or the government voluntarily enacts change) if either 2/3 of the population shows up at the original demonstration, or if there is a counter-demonstration and nobody shows up to that. Otherwise, they do not make any change.*

The four resulting cases are pictured in Figure 10.

The reasoning behind the proposition is straightforward and so we simply explain it here. The possible outcomes under the prescribed strategies are:

- If 2/3 show up, then it must be state 1 and change is enacted. There is no use for a counter-demonstration.

- If 1/3 show up, it could be either state 2 or 4. After the counter-demonstration:

  - If 2/3 show up to counter-demonstration then it must be state 4, and there is no change.

  - If 0 show up to counter-demonstration then it must be state 2, and change will be enacted either via a revolution or via the government.
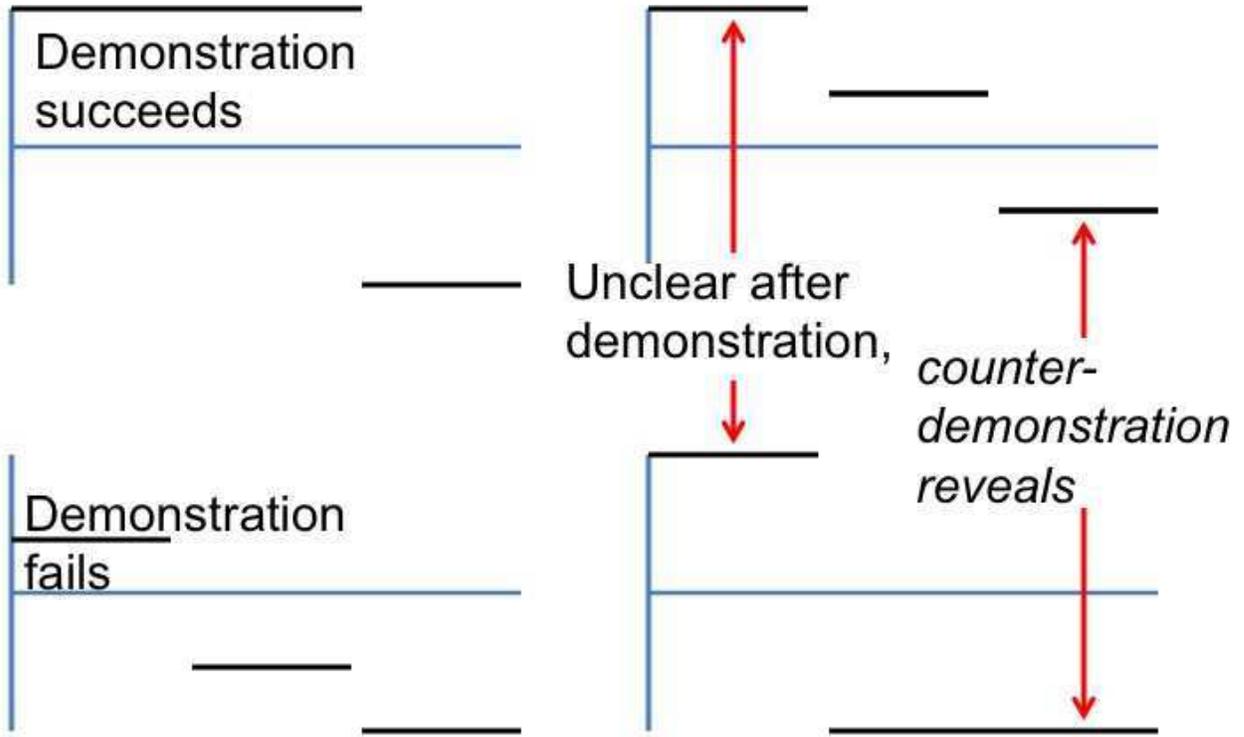
Figure 10: The counter-demonstrations are needed in the two states on the right in which there is a middle level of turnout at the original demonstration. That makes it clear that it is one of the two states on the right, but does not distinguish the state. The counter-demonstration then reveals whether there is a large support for no change, and so distinguishes the two states. The combination of the demonstration and counter-demonstration fully distinguishes among the four states.

- if 0 show up, then it must be state 3 and change is enacted. There is no use for a counter-demonstration.

The incentives for the groups to demonstrate or counter-demonstrate are clear:

The first group, whenever it has strong preferences would like to demonstrate since it has a two-thirds chance of eventual success. The necessary and sufficient condition for it to want to demonstrate in equilibrium is that $\theta_{\mathcal{CS}} > C/2$.

The third group clearly wants to counter-demonstrate they are $\mathcal{NS}$ types, since they know it is state 4 and they will be successful. They do not want to counter-demonstrate when they are $\mathcal{NW}$ types since then they know it is state 2 and they will fail.

The second group always gets its most preferred outcome by showing up to a demonstration or counter-demonstration when they are strong but not weak, and so they have no reason to change their strategy.

This example shows how counter-demonstrations can reveal a state and be useful in learning the state.

Note also that the example is fully symmetric - which group holds the first demonstration and which counter-demonstrates could also be reversed. In this case, since a natural status-quo is no change, it seems more natural to have the group supporting change be the one to hold the first demonstration and to bear some risk in doing so. But the example works in either way.

7.2 Other Actions by Governments

A government can change the world from being one in which there is an equilibrium with a revolution to one in which there is not, by affecting the various parameters.[20] This presumes that the government would like to avoid a revolution and keep the status quo.

Let us examine some of those behaviors.

7.2.1 Costs

Most directly, by increasing the cost to failed revolutionaries (increasing C), the government can make the conditions for a revolution harder to satisfy. For instance, in the base model, it is sufficient to raise C to a point at which

$$\theta_H/C < \frac{(1-\pi)(1-z)}{\pi z}$$

to avoid the revolution. Correspondingly, there are values of C that prevent revolution for different levels of information.[21]

7.2.2 Information Control and Homophily

The government can also suppress and censor information. As we saw, having only a few meetings with others, or if those meetings are mostly with own type then this can lessen the chance that people have to learn about the number of others who support change. By limiting information flows, especially across groups or geography, so that most interactions are limited and local, one could shift an equilibrium to preclude a revolution. As we have seen however, it could also work the other way in cases in which the prior beliefs are strong enough – by encouraging information exchange one could end up undercutting the support for a revolution and preclude it. Which policy a government would want to undertake would depend on the information structure.[22]

---

[20]For important analyses of governments and propaganda as well as censoring and other informational distortions in models that are very different from ours, see Edmond (2013) as well as Egorov, Guriev, and Sonin (2009), Little (2012), and King, Pan and Roberts (2013).

[21]For a model of repression, see Shadmehr and Boleslavsky (2016).

[22]See Luo and Rozenas (2016) for more discussion of informational control by a government.

7.2.3 Propaganda

The government could also bias information via propaganda.[23] Propaganda is interesting in that it does not have to convince all of the potential revolutionaries that revolution is a bad idea or that the state is Low, but instead it just needs to convince enough of them so that the remaining types know that they will no longer have sufficient numbers to be successful. For instance, if more than $z - q$ of the potential revolutionaries are convinced by the propaganda, then the revolution cannot succeed, regardless of whether the remaining H types are convinced or not.

Thus, propaganda can be disruptive even if it only convinces a small subset of the population that they should not take part in a revolt. This could happen by convincing people that they stand no chance of success, for instance, by inflating the estimates of how many θ_L types there are in the population; or by convincing people that they are better off than they are, or better off than what would happen after a revolultion, etc.

7.2.4 Redistribution

Finally, the government could also redistribute resources. Again, the government does not have to redistribute resources to all of the potential revolutionaries, they simply need to buy enough of them off to discourage the rest - so they just need to please $z - q$ of the H types. They can produce some very unhappy parts of the population, provided that they make the middle range sufficiently happy that they will no longer revolt.

Specifically, suppose that redistribution by the government is observable and that the government knows the state (so it knows the condition of the whole population). Thus, whenever the government does redistribute income, then the population knows it is the High state. So, it is clear that in that case they must pay at least θ_H to a fraction $z - q$ to avoid the revolution. The equilibrium must be one in mixed strategies. To see this note that if it were a pure strategy equilibrium, then it would be one in which the government only redistributed in the High state. But then when seeing no redistribution, agents would infer it is the Low state and not revolt. In that case, the government would not need to redistribute in order to avoid the revolution. Thus, the redistribution must be in mixed strategies. In order for this to make sense with a continuum of agents, we then allow agents to correlate their strategies, so that H types revolt with some probability p when not seeing redistribution. The probability of redistribution is then just enough to make agents indifferent conditional on seeing no redistribution, and the probability of revolt is just enough to keep the government indifferent between being overthrown and paying the redistribution.

---

[23]For a different views of information manipulation in the face of social coordination, see Edmond (2016) and Little (2016b).

8 Concluding remarks

We have provided a model that serves as a basis for the investigation of how information and learning affect the possibility of having successful revolutions.

We have shown four ways in which information can be either enabling or disruptive: (i) by encouraging some but discouraging others from participating, (ii) in settings with homophily, by weakening the content of information, (iii) by gaining information about the number of extremists in a society who might replace the status quo with an undesirable policy, and (iv) by triggering counter-protests that reveal support for the status quo.

We have shown that there are non-monotonicities so that small amounts of information can actually discourage enough of the population to make success impossible. We have also shown how demonstrations can provide important information, both within and across countries, that can help make revolutions possible, and increase the likelihood of their success.

Our model is deliberately simple, which makes many intuitions very clear and allows us to analyse a number of questions within one model - providing a more holistic view of what is needed for collective action to succeed, and should provide a basis for further studies of collective action.

We have focused on the coordination issues and the role of information. There can also be public-good aspects and free-riding behavior in protests and revolutions that we have not modeled here and could be interesting to combine with the coordination issue.[24]

It can also be that the intensity of a demonstration or revolution matters more continuously rather than just passing a threshold.

Finally, the feedback between politics and demonstrations is something that is deserving of much more study. This can fit into a more general study of the endogeneity of governments.[25]

References

Acemoglu, Daron; Georgy Egorov, and Konstantin Sonin. (2012). "Dynamics and Stability of Constitutions, Coalitions, and Clubs," *The American Economic Review*, 102(4), 1446-1476.

Acemoglu, Daron; Tarek A. Hassan and Ahmed Tahoun (2014). "The Power of the Street: Evidence from Egypt's Arab Spring". Mimeo.

Aghion, Philippe; Alberto Alesina and Francesco Trebbi (2004). "Endogenous Political Institutions," *The Quarterly Journal of Economics*, 119(2), 565-611.

---

[24]For an interesting paper on free-riding in protests, see Cantoni, Yang, Yuchtman, and Zhang (2017).

[25]E.g., see Aghion, Alesina, Trebbi (2004), Barbera and Jackson (2004), Acemoglu, Egorov, Sonin (2012).

Angeletos, George-Marios; Christian Hellwig and Alessandro Pavan (2007). "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks". *Econometrica*, 75(3), 711-756.

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

Barbera, Salvador, and Matthew O. Jackson. (2004). "Choosing How to Choose: Self-Stable Majority Rules," *The Quarterly Journal of Economics*, 119(3), 1011-1048.

Battaglini, Marco (2016). "Public Protests and Policy Making". forthcoming in the *Quarterly Journal of Economics*.

Bikhchandani, Sushil; David Hirshleifer and Ivo Welch (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades". *The Journal of Political Economy*, 100(5), 992-1026.

Bikhchandani, Sushil; David Hirshleifer and Ivo Welch (1998). "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades". *The Journal of Economic Perspectives*, 12(3), 151-170.

Breuer, Anita; Todd Landman and Dorothea Farquhar (2015). "Social Media and protest Mobilization: Evidence from the Tunisian Revolution". *Democratization*, 22(4), 764-792.

Brummitt, Charles D.; George Barnett and Raissa M. D'Souza (2015). "Coupled Catastrophes: Sudden Shifts Cascade and Hop among Interdependent Systems". *Journal of the Royal Society Interface*, 12(112), 1-12.

Bueno de Mesquita, Ethan (2010). "Regime Change and Revolutionary Entrepreneurs". *American Political Science Review*, 104(3), 446-466.

Cantoni, Davide; David Y. Yang, Noam Yuchtman, Y. Jane Zhang, (2017). "Are Protests Games of Strategic Complements or Substitutes? Experimental Evidence from Hong Kongs Democracy Movement" mimeo.

Chen, Ying and Hulya Eraslan (2014). "Rhetoric in legislative bargaining with asymmetric information". *Theoretical Economics*, 9, 483-513.

Chwe, Michael Suk-Young (2000). "Communication and Coordination in Social Networks". *Review of Economic Studies*, 67, 1-16.

Edmond, Chris (2013). "Information Manipulation, Coordination, and Regime Change". *Review of Economic Studies*, 80 (4): 1422-1458

Edmond, Chris (2016). "Information Manipulation and Social Coordination", mimeo.

Egorov, Georgy; Sergei Guriev and Konstantin Sonin (2009). "Why Resource-poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data". *American Political Science Review*, 103(4), 645-668.

Farrell, Henry (2012). "The Consequences of the Internet for Politics". *Annual Review of Political Science*, 15, 35-52.

Feldman, Mark and Christian Gilles (1985). "An Expository Note on Individual Risk without Aggregate Uncertainty". *Journal of Economic Theory*, 35, 26-32.

Granovetter, Mark (1978). "Threshold Models of Collective Behavior". *The American Journal of Sociology*, 83(6), 1420-1443.

Jackson, Matthew O. (2008). *Social and Economic Networks*, Princeton, NJ: Princeton University Press.

Jackson, Matthew O., Leo K. Simon, Jeroen M. Swinkels, and William R. Zame. (2002). "Communication and equilibrium in discontinuous games of incomplete information." *Econometrica*, 70(5): 1711-1740.

Judd, Kenneth L. (1985). "The Law of Large Numbers with a Continuum of IID Random Variables". *Journal of Economic Theory*, 35, 19-25.

King, Gary; Jennifer Pan and Margaret E. Roberts (2013). "How Censorship in China Allows Government Criticism but Silences Collective Expression". *American Political Science Review*, 107(2), 326-343.

Kricheli, Ruth; Yair Livne and Beatriz Magaloni (2011). "Taking to the Streets. Theory and Evidence on Protests under Authoritarianism". Mimeo, Stanford University.

Little, Andrew T. (2012). "Elections, Fraud, and Election Monitoring in the Shadow of Revolution". *Quarterly Journal of Political Science* 7, 249-283.

Little, Andrew T. (2016). "Communication Technology and Protest," *Journal of Politics*, 78(1), 152-166.

Little, Andrew T. (2016b). "Propaganda and Credulity," mimeo.

Loeper, Antoine; Jakub Steiner, Colin Stewart (2014) "Influential Opinion Leaders," *The Economic Journal*, 124, (581), 1147-1167.

Lohmann, Susanne (1993). "A Signaling Model of Informative and Manipulative Political Action". *The American Political Science Review* 87(2), 319-333.

Lohmann, Susanne (1994a). "Information Aggregation Through Costly Political Action". *The American Economic Review*, 84(3), 518-530.

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

Lohmann, Susanne (2000). "Collective Action Cascades: An Informational Rationale for the Power in Numbers". *Journal of Economic Surveys*, 14(5), 655-684 .

Luo, Zhaotian and Arturas Rozenas (2016). "Ruling the Ruling Coalition: Information Control and Autocratic Governance," mimeo, NYU.

Manacorda, Marco and Andrea Tesei (2016). "Liberation Technology: Mobile Phones and Political Mobilization in Africa". Mimeo.

McPherson, Miller; Smith-Lovin, Lynn and James M. Cook (2001). "Birds of a feather: Homophily in social networks," *Annual Review of Sociology*, 27, 415-444.

Milgrom, Paul (1981). "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380 - 391.

Myatt, David P. (2016). "A theory of protest voting." *The Economic Journal*, 1-41.

Pierskalla, Jan H. and Florian M. Hollenbach (2013). "Technology and Collective Action: The Effect of Cell Phone Coverage on Political Violence in Africa". *American Political Science Review*, 107(2), 207-224.

Shadmehr, Mehdi (2015). "Extremism in Revolutionary Movements". *Games and Economic Behavior*, 94, 97-121.

Shadmehr, Mehdi and Bernhardt, Dan (2016a). "Information Aggregation and Coordination". mimeo.

Shadmehr, Mehdi and Bernhardt, Dan (2016b). "Vanguards and Revolution". mimeo.

Shadmehr, Mehdi and Bernhardt, Dan (2016c). "When Can Citizen Communication Hinder Successful Revolution?" mimeo.

Shadmehr, Mehdi and Boleslavsky, Raphael (2016). "Institutions, Repression and the Spread of Protest". mimeo.

Steinert-Threlkeld, Zachary C.; Delia Mocanu; Alessandro Vespignani and James Fowler (2015). "Online Social Networks and Offline Protest". *EPJ Data Science*, 4(1), 1-9.

Weinstein, Jonathan, and Muhamet Yildiz (2007). "A structure theorem for rationalizability with application to robust predictions of refinements." *Econometrica* 75(2): 365-400.