

A Model of Protests, Revolution, and Information

Salvador Barbera and Matthew O. Jackson *

Draft: March 2018

Abstract

A revolt or protest succeeds only if sufficient people participate. We study how potential revolutionaries' ability to coordinate is affected by what they learn from those around themselves. We begin by exploring the tradeoffs from such learning: some supporters of the revolution will meet other supporters and become more convinced that there is strong support for the revolution, while others will meet partisans of the status quo and become discouraged. We show that which effect wins is determined by the prior beliefs of the agents and the strength of the correlation of preferences across the population. With this understanding in hand, we then explore how homophily (situations in which people are more likely to meet others who have similar preferences) affects the possibility of a revolution. Homophily undercuts learning (at a quantifiable rate), but also has some other interesting effects. Under homophily, potential revolutionaries learn less about the support of the revolution from meeting another supporter of the revolution given that meetings are biased; but still learn nontrivially when they do meet a supporter of the status quo, and so learning becomes more asymmetric. We examine these effects on revolutions in detail. Beyond this, we also use our simple model to study the role of protests and counter-protests before a revolution, not only in revealing support for and against a revolution, but also in signaling relative sizes of extremist factions. 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). We close with some discussion of government actions, including propoganda, redistribution, and violent repression.

*Barbera is at MOVE, Universitat Autònoma de Barcelona and Barcelona GSE. He acknowledges financial support under grants ECO2014-53051-P, 2014SGR-515 and the Severo Ochoa Programme. Jackson is at the Department of Economics, Stanford University, Stanford, California 94305-6072 USA, and is an external faculty member at the Santa Fe Institute and a member of CIFAR. Emails: and jacksonm@stanford.edu. He gratefully acknowledges financial support under ARO MURI Award No. W911NF-12-1-0509. We thank Ernesto Dal Bo and Santiago Oliveros for very helpful discussions and suggestions. We also thank Daron Acemoglu, Avi Acharya, Dan Bernhardt, Ethan Bueno de Mesquita, Raphael Boleslavsky, Steve Callendar, Ying Chen, Georgy Egerov, Jon Eguia, Hulya Eraslan, Navid Hassanpour, Navin Kartik, Dániel Pass, Arturas Rozenas, Mehdi Shadmehr, Scott Tyson, and participants of various seminars and conferences for helpful comments.

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 other organization, there may be room to force it. However, revolutions and demonstrations can be quite risky, with substantial costs to participants if their actions fail - not just in terms of lost time and effort, but in some cases involving imprisonment, exile, or even death. Yet, there are also potential personal advantages, both moral and material, from having participated in successful revolutions. Hence, whether one unfolds hinges on the confidence that supporters of change have about the number of others are also willing to act to force change. As such, it can be very difficult to collectively act, even when such action would be successful.

In this paper, we provide a simple model of collective action, focusing on how people learn from each other about the level of support for a collective action. We study how people learn from interacting with others around them, as well as from things like a buildup of protests prior to a revolution. We also discuss how, in spite of the improvements in social media and communication, demonstrations and protests remain differentiated methods of signalling the intensity of preferences, and the conditions under which agents are willing to take risks in favor of change.

In what follows, we refer to the collective action as a “revolution”, even though it should be understood that our model encompasses different types of actions 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 a political or business 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.

A main contribution of our paper is to analyze how people learn from meeting other people. This has countervailing effects: some supporters of the revolution will be discouraged as they will meet partisans of the status quo. Others become encouraged as they meet other supporters and become more confident that there are enough supporters to have a successful revolution. If too many become discouraged, then this can disable a revolution; while in contrast, increased confidence from meeting other supporters can make it possible to have one when otherwise supporters would have been too uncertain to act. Which effect dominates depends in nicely quantifiable ways on: prior beliefs, the relative benefit of success compared to the cost of failure, and how correlated preferences are across the population. For example, with low levels of correlation of preferences the discouragement effect dominates, while with higher levels of correlation the encouragement effect dominates.

With this basic understanding in hand, we turn to studying how homophily impacts the potential for a revolution. Homophily refers to people being biased to meeting other people

with the same preferences. For instance, in a society stratified by social and economic classes, people may mostly only talk to others who have very similar backgrounds and circumstances, and thus similar preferences. This was an effect that is often cited as to why some people were surprised at the outcome of the Brexit vote.¹ In the extreme, if people only interact with others who are very similar to themselves, they learn nothing about the overall prevalence of different preferences in the society. 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 otherwise unravel them. Thus, homophily provides a new and important angle on when collective action may succeed.

These results are all framed within a very simple model. The flexibility of the model then enables us to investigate additional issues in a common context. For example, we also examine how extremism can undermine a revolution. 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. We also examine the role of repeated demonstrations, counter demonstrations and other government actions before a revolution. For example, a government, or their partisans, may wish to hold counter-demonstrations to signal the strength of support for the status quo. In some settings, if an initial demonstration in favor of change leaves some doubts as to the size of the support, counter-demonstrations 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 demonstrations, or buy off some of the disenchanted.

We begin the formal part of the paper in Sections 2 and 3, by describing the underlying one-shot model in which the members of a population must simultaneously decide whether to participate in a revolution (or 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 the 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 Section 4, we move to our main analysis to analyze how information that potential revolutionaries get via meeting each other.

We first examine, in Section 4.1, when it is that meetings between potential participants

¹For example, see “How to check if you’re in a news echo chamber and what to do about it” by Tom Stafford, December 12, 2016, *The Conversation*.

either enable or disable a revolution. We consider situations in which people get to observe the type of another agent in society.² 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) Communication can enable a revolution that could not occur otherwise when there is poor prior information about the numbers of potential revolutionaries. (ii) Communication can make a revolution impossible, even though it would have been possible without communication when there is stronger prior information about the numbers of potential revolutionaries, but too few of them to overcome the loss in their numbers due to the discouragement of some of them meeting partisans of the status quo. 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; and we show how this depends on the prior beliefs and the correlation of preferences across members of the society.³

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 estimate the size of the group who supports change. Our friends are not a representative sample of the population. This affects both of the forces just mentioned 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 fewer people can be discouraged from participating. Which of these effects dominates depends on our priors, the correlation of types with the state, and the degree of homophily, in quantifiable and intuitive ways.

We then explore – in Sections 5, 6, and 7 – a number of important issues that have

²Communicating types is incentive compatible in this setting.

³An 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.

been already discussed in the literature and show how our model allows for fresh, simple, and often complementary, explanations of such phenomena. In Section 5, we examine 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 5.3). 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.⁴ 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 5.4).

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).

1.1 Relation to the Literature

⁴See Little (2016) as well as Kiss, Rodriguez-Lara and Rosa-Garcia (2017) for discussions of how improved technology has changed agents’ knowledge of others’ preferences, and also enables better coordination regarding where and when to hold protests.

Collective action has been analyzed from various angles, both theoretically and empirically, and in reference to different countries and circumstances. Indeed, there are many realities to take into account: the types of governments against which demonstrations 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 and 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 and thresholds for action is Granovetter (1978) - a theme explored in detail by Hassanpour (2017) who examines how network position and connectivity can affect leadership in collective action. 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.

From this expansive literature two things are clear: crowds can get things wrong, and learning from others is important and can change the possibility of collective action for good or bad.

The main contribution of our paper comes from the simplicity of our model that allows us to get a picture of learning from others that goes beyond previous analyses in several regards.

First, our formulation allows us to see what happens when individuals learn from their neighbors. This plays out very differently from the cascades approach in which each person observes the previous ones and the usual herds or cascades can form, or the global games approach in which agents make deductions about the potential behaviors of others. Instead, in our model people are *simultaneously* learning from each other, but each with a *different* sample of the population. Such a model has not been analyzed before, and matches how many people learn. We show that this leads to countervailing effects that we are able to relate to prior beliefs and the correlation in preferences across the population.

Second, the simple formulation allows us to introduce homophily, an important phenomenon that has never been studied in the context of collective action. We show how homophily impedes learning and how its effects depend upon the prior beliefs and the correlation in preferences across the population

Third this is done within a unified framework that has many advantages, including the

possibility of providing a unified view of a large variety of other topics that have been analyzed piecemeal in the literature and in more complicated models, such as the role of demonstrations, counter-demonstrations. We also show how these interplay with the size of an extremist population, and government actions.

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 revolution (or protest, strike, 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.

We start with a description of a general form of payoffs as a base for future study, but then specialize the model to a simple case for much of our main analysis.

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.

Again, we use the term 'revolt' but the model obviously has many applications.

The revolution is successful if at least a fraction $q \in (0, 1]$ of the population participates. If fewer than q participate, then the revolution fails.⁵

2.2 Uncertainty

$\omega \in \mathbb{R}$ 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 .

$\theta_i \in \mathbb{R}$ is the type of agent i , which is the private information of that agent.⁶

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

⁵This is a sharp model of collective action with a discontinuous change above or below a threshold. In reality there may be grey areas close to the threshold, but it does not seem that complicating the model in that way would add much insight. Uncertainty on the part of agents can already be used to smooth equilibrium behavior, and this version of the model is very tractable.

⁶We 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.

technically convenient but has some measurability issues that are easily handled as the limit of a finite model.⁷

We assume the standard ordering property on information:⁸ 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 dependent, and are given by the following table.

	<i>Success</i>	<i>Failure</i>
<i>Participate</i>	$a(\theta_i, \omega) + V_i(\theta_i, \omega)$	$b(\theta_i, \omega) - C_i(\theta_i, \omega)$
<i>NotParticipate</i>	$a(\theta_i, \omega)$	$b(\theta_i, \omega)$

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 positive ($-C_i$ is negative), which represents the personal cost to a person of being caught in an unsuccessful revolution - for instance, being jailed, fined, executed, etc. On the other side, V_i captures the personal pleasure or pain that a person would feel from participating in a successful revolution. Clearly, someone who favors change will generally have a positive V_i - which is the warm glow from having been there and supported the successful cause when it was needed. This is the positive feeling that someone who marched alongside Ghandi in the Salt March, or alongside Martin Luther King during the civil rights movement, feels for having been part of a history-changing demonstration. Generally, someone against change would have a negative V_i as they would feel guilt or shame

⁷For 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.

⁸See Milgrom (1981).

for being part of the revolution that overthrew a regime that they supported - for instance, a racist might have gotten a negative payoff from participating in the March on Washington for Jobs and Freedom in 1963.⁹

The structure of this game is similar to that in the expressive voting literature (e.g., see Feddersen (2004) for a review). Here, the chance that any single agent will be pivotal in terms of making the revolution succeed rather than fail, in a large population and a nontrivial threshold, is negligible. In the continuum model it is in fact 0. Thus, we follow the usual approach from the voting literature and have $V_i(\theta_i, \omega)$ be the marginal utility for having participated, conditional upon success. *So this is not the payoff that an agent gets in terms of expecting to be pivotal, but instead the utility they get from knowing or being able to say that they participated in a revolution that was successful:* from having been one of those who stormed the Bastille, or protests in Tunisia, etc. It will almost never be the case that one more or one fewer revolutionary would make the difference. People react to some more basic utility from the action itself, which can be motivated in many different ways, just as discussed in the voting context (again, see Feddersen (2004) for discussion and references).

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

	<i>Success</i>	<i>Failure</i>
<i>Participate</i>	$V_i(\theta_i, \omega)$	$-C_i(\theta_i, \omega)$
<i>NotParticipate</i>	0	0

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 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:

	<i>Success</i>	<i>Failure</i>
<i>Participate</i>	$V(\theta_i, \omega)$	$-C$
<i>NotParticipate</i>	0	0

Let us mention two canonical cases:

⁹In some cases, people who favor the status quo might want to be seen as supporters of change if the status quo is overturned, so there are situations in which even some supporters of the status quo end up having a positive V_i . That is fine for the model, as we are agnostic on what drives V_i and just analyze its consequences.

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

	<i>Success</i>	<i>Failure</i>
<i>Participate</i>	θ_i	$-C$
<i>NotParticipate</i>	0	0

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

	<i>Success</i>	<i>Failure</i>
<i>Participate</i>	ω	$-C$
<i>NotParticipate</i>	0	0

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_i s) 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 $\sigma_i : \mathbb{R} \rightarrow \Delta(\{0, 1\})$, 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.¹⁰

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

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

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.

2.6 Existence

As this is a coordination game, equilibria exist and in fact there are generally 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 this section we provide an existence theorem for the general model, and in fact show that nicely behaved equilibria exist. We defer our discussion of the multiplicity and selection of equilibria to the version of the model that we analyze in detail (see Section 3.2).

These games have equilibria in which agents play monotone strategies: their probability of participating is non-decreasing in θ_i , and for the cases that we examine those will be the only equilibria. Generally, given the increasing preferences and ordering on information, such

¹⁰We 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.

equilibria always exist. Nonetheless, for some special cases there do exist other equilibria, although for generic distributions these will be the only equilibria.¹¹

PROPOSITION 1 *Equilibria exist, and in fact, symmetric and monotone equilibria exist. Each monotone equilibrium 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 is implied by 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 Bayesian equilibrium, or weak perfect Bayesian equilibrium when we go beyond the simultaneous move setting.

3 A Discrete Private Values World: Background Analysis

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.

¹¹For 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 $\theta_i = \omega = 3$.

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.¹² 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

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

The parameter z is an important one that captures how correlated people’s preferences are. A high level of z gives someone a high confidence that if they are type H then most other people are as well; while if z is closer to $1/2$, then knowing one’s own type does not tell one much about other people’s types.

3.2 The Multiplicity of Equilibria and Global Games

Again, as this is a coordination game, there can be a multiplicity of equilibria.

In the private values settings that we study here the multiplicity is very simple, and so our analysis provides a full characterization of the set of strict equilibria.

In our game, equilibria come in three flavors.

First, there is always an equilibrium in which nobody participates. This is straightforward, and it is clear that this is a strict equilibrium. If nobody else participates, then an agent’s payoff is strictly negative from participating as they pay the cost for certain.

Second, there for some parameter values there also exist participatory equilibria. There is one in which all of the H types participate. This is also a strict equilibrium when it exists.

¹²It 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.

Third, when there exists a participatory equilibrium, there also exist mixed strategy equilibria in which H types are exactly indifferent between participating or not.¹³ In the world of a continuum, in order to get those sorts of 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)).

The mixed strategy equilibria are unstable, as slight perturbations of the actions lead best replies to converge either to the all participate or no participate equilibria. Thus, we follow the literature and focus our attention on the pure strategy equilibria which are strict and stable (in the sense that best replies converge back to the equilibrium under a perturbation of strategies).

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 and the ones in which all the H 's participate. Therefore, to fully characterize these equilibria, we need only characterize when there exists an equilibrium in which H 's participate as we know the other is always present.

We do not take a stand on which of these two equilibria is more natural. In cases in which the unique equilibrium is non-participation, then the outcome is clear. In cases in which there exists a participatory equilibrium, one can find some conditions under which it is picked out via global games arguments. Since those arguments appear elsewhere,¹⁴ we do not repeat it here.

Thus, in what follows we analyze when it is that there exists a participatory equilibrium in addition to the non-participatory one.

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:

$$\theta_H/C \geq \frac{(1-\pi)(1-z)}{\pi z} \quad \text{equivalently} \quad \pi z \theta_H - (1-\pi)(1-z)C \geq 0. \quad (2)$$

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:

¹³With a continuum of agents, there exist a variety of such equilibria in which different agents use different mixtures - but the aggregate leads to exactly q agents showing up. Only some of those equilibria involve measurable strategies.

¹⁴That 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)). See also Weinstein and Yildiz (2007) for an argument that one could select either of the strict equilibria in our setting could be made the unique equilibrium with different introductions of slight noise into the model, depending on how that noise is modeled.




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.

- 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 $\pi z \theta_H \geq (1 - \pi)(1 - z)C$.

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 exist 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 .¹⁵

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.¹⁶ 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. 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.

¹⁵See 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.

¹⁶This 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.

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 Homophily and Communication Prior to a Revolution

With the basic model in hand, we can now turn to our main results and 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, and how that depends on how homophilistic the society is.

4.1 Each Agent Sees One Other Agent's Type

We first consider what happens if each agent gets to meet another agent and to see that other agent's type, where that agent is chosen uniformly at random – so that there is no homophily in the meeting process.

Each agent gets to talk to one other agent in the society and learn that agent's type.¹⁷

¹⁷Note 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 revolution to be successful, but not

This provides additional information to the agents, since now they have two signals about the state rather than just one.

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

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

The evolution of beliefs is pictured in Figure 4.¹⁸

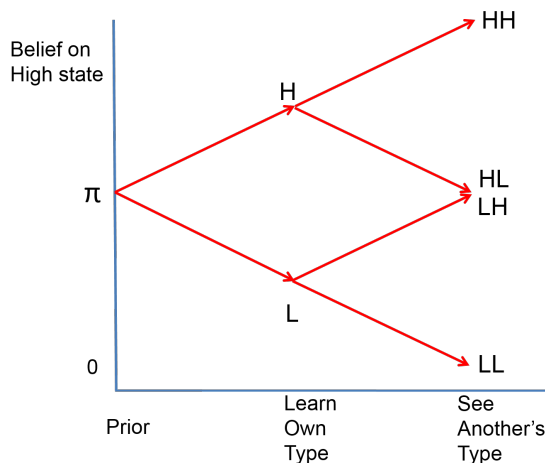


Figure 4: How an agent's beliefs evolve from the prior, to learning his or her own type, to meeting another agent and learning their type. For instance, HL indicates that an H type agent meets an L type agent.

Let's examine how this impacts an equilibrium. Since $\theta_L < 0$, we only have to analyze the H type's incentives in order to characterize equilibria, since L types never participate regardless of their beliefs.


As we see from Figure 4, some of the H types meet other H 's and become more sure that it is the high state, while others meet L 's and lose confidence.

Which effect dominates depends on how high the prior is and what the relative payoff from success is compared to the cost of failure.

There are several possibilities, pictured in Figure 5.

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.

¹⁸We thank Santiago Oliveros and Ernesto Dal Bo for suggesting Figures 4 and 5.



(a) Meeting another makes no difference: Beliefs are high enough regardless of whom an H meets .

(b) HL s are discouraged. Without meeting another they would have showed up, but now will not. This disables a revolution if $z^2 < q$.

(c) HH s are encouraged. Without meetings they would not have showed up, but now will. This enables a revolution if $z^2 > q$.

Figure 5: **The Effects of Meeting Another Agent:** Depending on the prior and the relative payoffs from the revolution, meeting another agent may either enable or disable a revolution.

First, as in panel (a) of Figure 5, the prior belief could be 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 whom they meet* if and only if:

$$\pi\theta_H \geq (1 - \pi)C. \quad (3)$$

This condition is more demanding than our previous equilibrium in the absence of meeting another agent, since it is asking whether H types go even when meeting an L and getting discouraged.

Second, as in panel (b) of Figure 5, it could be that there would have been a revolution if agents did not meet one another, but now after meeting others the H types who meet L 's are discouraged and no longer have enough confidence in the state to participate. The H types who meet H 's are encouraged, but if they are rational, then they know that some H 's will be meeting L 's and will no longer show up. If $z^2 < q$, then the HH types will not be numerous enough to have a successful revolution on their own, and so knowing this, even though they are confident in the state they are not willing to show up and so the revolution will no longer take place.

Third, as in panel (c) of Figure 5, it could be that there would *not* have been a revolution if agents did not meet one another, but now the H types who meet another H type are encouraged enough and numerous enough, to hold a revolution. 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

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

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

$$z^2 \geq q,$$

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

If we view how this works in terms of the space of parameters π and z , Figure 6 shows that there are three different regions.

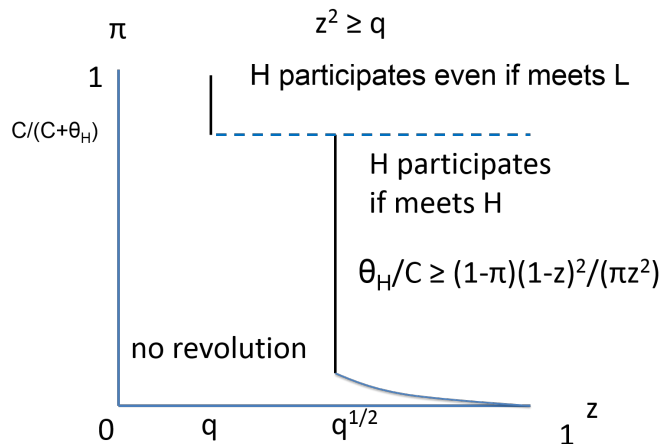


Figure 6: **Three Regions of Equilibria:** priors are so high that all H types show up regardless of whom they meet; or priors are in a region such that H types only show up if they meet an H type, and there are enough HH types to be successful ($z > \sqrt{q}$); and otherwise there is too little confidence or too few confident HH types to hold a revolution.

Comparing this to the no information case, Figure 7 shows the difference in equilibrium structures for the two settings:

In Figure 8 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 7: 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.




Figure 8: Sometimes information aids the revolution and other times it blocks it

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, 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 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, 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.¹⁹

To capture this, let us consider a variation on the above setting in which we incorporate homophily. Some of the meetings are biased towards own type.

A direct way to model this (effectively, without loss of generality) is to have a fraction $h \in [0, 1]$ of matches that would 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, and in between h tracks the relative bias towards meeting own types.

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.

Homophily makes people more likely to meet their own type. Thus, by meeting their own type they update their beliefs less as homophily increases. Meeting other types decreases under homophily, but the rate at which it decreases is the same across states and so the updating conditional on meeting an opposite type is unaffected. This is as in Figure ???. Nonetheless, fewer such meetings occur and so homophily's impact is to lower updating conditional on meeting same type and to lower the chance that other types are met, and so overall it slows learning in both ways.

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. This is a variation on the previous analysis, using Bayes' rule. It is summarized in the following proposition.

PROPOSITION 2 *If $\pi > \frac{C}{C+\theta_H}$, there is no effect of homophily on the existence of an equilibrium in which there is a revolution: H types participate regardless of whom they meet.²⁰ In contrast, if $\pi < \frac{C}{C+\theta_H}$, then there exists an equilibrium in which a revolution occurs if and only if:*

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

and

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

and the participants in the revolution are the H types who meet other H types.

¹⁹For background on this empirical observation, termed “homophily”, see McPherson, Smith-Lovin and Cook (2001) and Jackson (2008).

²⁰At the exact equality there exist a continuum of equilibria in which the H types that meet L types mix on participating as they are exactly indifferent.




Figure 9: Homophily lowers the updating of beliefs from an H meeting an H , but does not affect the updating of an H meeting an L (the relative odds of that happening across states are not affected by homophily). However it does lower the frequency with which H 's meet L 's regardless of the state

We have already provided the proof in the derivation from the previous sections and above.

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 state.

This leads to the adjustment in the equilibrium structure as pictured in Figure 10.




Figure 10: Homophily (assortativity in meetings) changes the equilibrium structure.

We see that higher homophily increases the region of having a revolution if the prior is

in a middle range, 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, which otherwise could have given more confidence to H types who meet other H types.

4.3 Under-Estimating Homophily

The above discussion presumes that agents are aware of how much homophily there is in their society. In many settings people may understand that they are more likely to meet others who are similar to themselves, but have no idea of how extreme homophily tends to be.²¹

Under-estimating homophily can inflate the confidence that agents derive from meeting others. With high levels of homophily, agents are very likely to meet others who have similar types. For instance, without knowing the extent of homophily most H types would meet other H types (even in the *Low* state) and could become over-confident that it is the *High* state. For example, in situations in which a revolution would not take place without meeting anyone, as in panel (c) of Figure 5, H types would become more confident and be willing to take part (provided $z^2 > q$), if they meet another H type without correcting for potential extreme homophily and so end up over-participating in the *Low* state.

4.4 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, presuming there is some chance of meeting the other types of agents so that homophily is not complete. For instance, for each meeting with probability h , H types are matched with other H 's and L 's with L 's, and then with probability $1 - h$ agents are matched at random so that they expect to meet other types in proportion to their weight in the population. This can be done independently over matches, so that if a person meets a sequence of people then that sequence will be biased towards own type (at a rate h), but each meeting is independent of the other meetings.

PROPOSITION 3 *For any $z > q$, π , and θ_H/C , there exists a number of signals above which there is an equilibrium which involves an H type participating in a revolution 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 revolution is perfectly effective in the limit.*

²¹For some discussion of the extreme levels of homophily in many settings, see McPherson, Smith-Lovin, and Cook (2001) and Jackson (2014, 2017).

We can also see directly how the speed of convergence of beliefs is affected by the presence of homophily. The usual rate of convergence (expected error of the posterior) would be proportional to $1/\sqrt{n}$, where n is the number of observations of draws observed from the distribution. That would be the rate with no homophily, by the Central Limit Theorem. Under homophily, however, the effective number of real observations will be $1+(1-h)(n-1)$, or roughly $(1-h)n$ (on average). Thus, the rate of convergence will be proportional to $1/\sqrt{(1-h)n}$ - which gets slower as h increases. So, homophily slows the rate of learning at a rate $1/\sqrt{1-h}$, which becomes infinite as homophily tends to 1.

We sketch the key idea behind the proof of the proposition.

As long as homophily is not complete ($h < 1$), the posterior belief on the *High* state will converge to 0 or 1 and be correct almost surely. This can be derived from Levy's Zero-One Law: The fraction of other agents that an H type will meet in the *High* state will differ from that in the *Low* state. The H type would meet a fraction p_H of H types in the *High* state in expectation, and some p_L in the *Low* state. As long as homophily is not complete ($h < 1$), $p_L < p_H$. An event that the H type meets a fraction of at least $(p_L + p_H)/2$ over an infinite sequence of meetings is a tail event, and its probability is then 0 or 1 (depending on the state). Thus, as an agent sees a sequence of meetings, the posterior that they have will converge to either 0 or 1 in the limit, almost surely (by Doob's Martingale Convergence Theorem). Thus, as the number of people agents meet becomes large, then the fraction of H types in the population whose posterior has high certainty on the correct state tends to 1. With enough observations, H agents can be sure that there will be a fraction above q who have sufficient belief on the *High* state to make the revolution successful with high probability in the *High* state.

Thus, even with homophily the expected long stream of people that an agent will meet will differ by the state, and so 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 high enough threshold, and in the limit they are almost always successful. For any finite number some agents will make mistakes, but that proportion vanishes in the limit.

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, including social media and cell phones, 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 Kiss, Rodriguez-Lara, and Rosa-Garcia (2017), 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:

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

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,²²

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

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

²²Here, 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 .

As a numerical example, let $z = 2/3$ and $q = 1/2 \theta_H \pi / [C(1 - \pi)] = .8$. Here the non-monotonicity 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 before a revolution. 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.

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 11.

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 .




Figure 11: Two possible states. By seeing how many θ_H types turn out at a protest, the state is revealed, which can enable a revolution.

Let us consider a case in which

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

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 Homophily and Demonstrations

Our discussion about the role of demonstrations focused on situations in which the prior beliefs of the moderates were insufficient to revolt without the information gained from a

demonstration. Other situations in which a demonstration helps are ones in which moderates meet others but are not willing to revolt. For instance, one case is such that there is sufficient homophily such that moderate beliefs remain too low to prompt willingness to participate even if they meet other supporters. Thus, demonstrations can help overcome homophily since people see the actions of many people, and they are sure of what others have seen.

5.3 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:

	<i>Success</i>	<i>Failure</i>
<i>Participate</i>	θ_i	$-C$
<i>NotParticipate</i>	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.²³

²³Note 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.

5.4 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. They 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:

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

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.

This is consistent with the data on the Arab Spring collected by Brummitt, Barnett, and D'Souza (2014), who find 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.

Beyond the Arab-Spring example, a careful analysis of the Swing Riots of the 1830s by Aidt, Leon, and Satchell (2017) provides convincing evidence of a contagion of riots; again much along the lines predicted here.

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.²⁴

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.²⁵ In particular, participating moderate agents get θ_M if the revolution is successful and there are more moderates than extremists, and get $-C$ 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

²⁴See 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.

²⁵See Acemoglu, Hassan, and Tahoun (2015) for a description of conflicts between different revolutionary groups during the Arab Spring.

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.²⁶ 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.

²⁶For an illuminating but different discussion of how information revelation by a government about potential counter-policies can affect revolutions, see Shadmehr and Bernhardt (2016c).

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

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 12




Figure 12: 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.

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 4 (Counter-Demonstrations) *Suppose that $\theta_{CS} > C/2$. 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.
- 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 13.




Figure 13: 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.

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.
- 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 $\theta_{CS} > C/2$.

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.^{27,28} 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,

²⁷For 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).

²⁸Events beyond a government's control that uncover its weaknesses can also change conditions, enhancing the possibility of a revolution - for instance, see González-Torres, Ada and Elena Espisito's (2017) discussion of Ebola and social unrest.

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.²⁹

7.2.2 Information Control, Seclusion, 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.³⁰

Our results on homophily also suggests that a revolutionary group might want to seclude its members. By allowing its members to possibly meet others who do not support the revolution, the group risks having its members doubt the possibility of success which could disrupt the revolution.³¹

7.2.3 Propaganda and Fake News

The government could also bias information via propaganda.³² 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 revolution, etc.

²⁹For a model of repression, see Shadmehr and Boleslavsky (2016).

³⁰See Luo and Rozenas (2016) for more discussion of informational control by a government.

³¹ This applies quite generally, and military forces and paramilitary groups are at times discouraged from interacting with populations that might raise doubts about their mission or the support for it.

³²For a different views of information manipulation in the face of social coordination, see Edmond (2016) and Little (2016b).

Noisy news sources can have the same effect as propaganda, effectively lowering the confidence that individuals have in the information that they receive as well as increasing the chance that some others may be discouraged. It can work in the same way as overt propaganda in that if people are concerned that a small subset of supporters of change may be discouraged, then that can unravel the revolution. Of course, it could work in reverse, convincing a small group that might not otherwise participate which then enables a revolution.

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.

8 Concluding remarks

8.1 Summary

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

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.

8.2 Testable Implications

Although our results provide for different possibilities – sometimes revolutions are precluded and other times they are not – those results are tied to specific parameters. Thus, the results are far from “anything can happen”. For example, it is for medium levels of z that meeting others can disable/unravel a revolution. The variable z is easily interpreted: How homogeneous is the society in question? In a country in which most agents are likely to feel the same impact of the government (e.g., most are in similar economic conditions, such as in Tunisia prior to its revolution, or France prior to its, etc.), then it is much easier to sustain a revolution even in the face of small numbers of meetings of others than in a country where there many different social classes who might have diverging preferences.

The comparative statics in terms of homophily are also testable. Levels of homophily, both in networks of interactions and in which media people pay attention to, are measurable and differ across societies. For instance, the caste structure in India leads to very high levels of homophily compared to other countries, in terms of interactions, political affiliations, and the news sources that people access.

Our results generally suggest that higher levels of homogeneity and of homophily, given bad enough conditions (so high enough payoffs) will lead to revolutions whereas lower levels of homogeneity or homophily could preclude a revolution. Our results also suggest that high homogeneity, bad conditions, and very high levels of communication, would make revolutions easier to have, all else held equal.

Our results also suggest that revolutions or successful collective actions are more likely to happen after relatively larger demonstrations, and less likely after small first demonstrations or large counter-demonstrations; again all else held equal.

As finding a data set that would allow one to test such hypotheses, especially causally, may be difficult, another option would be to first test many of our model’s predictions in a laboratory setting in which agents’ information, interactions, and payoffs can be observed and controlled.

8.3 Extensions

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.³³

We have provided our analysis in the context of correlated private values, but a similar analysis applies to more general affiliated and common value settings. The analog of homophily is still that people with similar types are likely to meet each other – so that people who are in close contact are likely to be getting similar information, and thus do not learn as much from meeting each other as meeting a uniformly random draw from the population.

Our focus in terms of learning has been on the population learning how many others are willing to support the revolution. In some instances, it might be that the key learning that goes on is whether sufficient numbers of the military and/or police would also support a revolution, and whether they would fire upon the population. This could easily be added to the model by having different roles in the population, and could explain why some governments try to isolate some of their key military from the general population.

In our analysis we have taken the meeting and homophily structures as exogenous. Control of social media by a government, as well as rules that limit internal movement, could be used to control the interaction structure within a society, and could be interesting to explore as another extension of the model.

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.³⁴

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.

Aidt, Toke, Gabriel Leon, and Max Satchell (2017) "The social dynamics of riots: Evidence from the Captain Swing Riots, 1830-31," mimeo: University of Cambridge.

³³For an interesting paper on free-riding in protests, see Cantoni, Yang, Yuchtman, and Zhang (2017).

³⁴E.g., see Aghion, Alesina, Trebbi (2004), Barbera and Jackson (2004), Acemoglu, Egorov, Sonin (2012), and Egorov and Sonin (2017).

- 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 and Konstantin Sonin (2017). "Elections in Non-Democracies". mimeo Northwestern University and University of Chicago.
- 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.
- Feddersen, Timothy J. (2004) "Rational choice theory and the paradox of not voting." *The Journal of Economic Perspectives* 18.1, 99-112.
- Feldman, Mark and Christian Gilles (1985). "An Expository Note on Individual Risk without Aggregate Uncertainty". *Journal of Economic Theory*, 35, 26-32.
- González-Torres, Ada and Elena Espisito (2017) "Epidemics and Conflict: Evidence from the Ebola outbreak in Western Africa," mimeo: U.C. Berkeley.
- Granovetter, Mark (1978). "Threshold Models of Collective Behavior". *The American Journal of Sociology*, 83(6), 1420-1443.
- Hassanpour, Navid (2017) *Leading from the periphery and network collective action*, Cambridge University Press.
- Jackson, Matthew O. (2008). *Social and Economic Networks*, Princeton, NJ: Princeton University Press.
- Jackson, Matthew O. (2014) "Networks in the Understanding of Economic Behaviors," *Journal of Economic Perspectives*, 28(4), 322.
- Jackson, Matthew O. (2017) *The Human Network*, Pantheon Books - Knopf Doubleday, forthcoming.
- 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.

- Kiss, Hubert Janos, Ismael Rodriguez-Lara, and Alfonso Rosa-Garcia (2017) "Overthrowing the dictator: a game-theoretic approach to revolutions and media," *Social Choice and Welfare*, 49: 329, 329-355.
- 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.