Multiplicity and Uniqueness in Generalized Regime Change Games[1]

Mehdi Shadmehr[2]

[1]I am grateful for comments of Dan Bernhardt, Raphael Boleslavsky, and Stephen Morris.
[2]Institute for Advanced Study, Princeton. E-mail: mshadmeh@gmail.com

Abstract

Regime change global games have become popular models in comparative politics because they ensure a unique equilibrium when coordination concerns tend to generate multiple equilibria. This paper provides an intuition for this uniqueness in a generalized regime change game, and shows why public information (e.g., that the regime is relatively weak) can generate multiple equilibria.

Keywords: Global Games, Regime Change, Revolution, Protest, Common knowledge, Public Information.

1 Introduction

Regime change games with binary-action are a class of global games, where players must decide between a safe action and a risky action that only pays off if the fraction of players who take that action exceeds a threshold. Critically, this threshold is uncertain and players have noisy private signals about it. Regime change games have been used extensively in the study of protest, rebellion, or revolution in political science, because they yield a unique equilibrium in strategic settings where coordination concerns generate multiple equilibria. Examples of published papers in top political science journals include Boix and Svolik (2013), Casper and Tyson (2014), Little et al. (2015), Rundlett and Svolik (2016), Shadmehr and Bernhardt (2017), and Tyson and Smith (2018). There are many other working papers or economics publications that use regime change games to study different aspects of revolutions, including Chen et al. (2016), Edmond (2013), Loeper et al. (2014), and Shadmehr (2017).

Given that the main reason to use regime change global games is that they yield a unique equilibrium, it is essential to pin down a clear intuition for this uniqueness result, and to investigate whether uniqueness extends to other "revolution technologies," where the likelihood of regime change is not so stark as to be a step function, where the revolution succeeds when the fraction of players exceeds a threshold, and fails otherwise. Applying the results of Morris and Shin (2003), this paper extends the commonly used regime change game to settings where the likelihood of regime change is weakly increasing in the fraction of revolters, and weakly decreasing in an uncertain parameter that captures, e.g., the regime's strength; it also presents an intuition for the uniqueness of symmetric monotone equilibrium and for why public information that creates common knowledge can upset the uniqueness result.

Symmetric monotone equilibria take a simple form: players take the risky action (revolt) when their signals are below a threshold (so that they believe the regime is relatively weak), and take the safe action when their signals are higher. Let x^* be such a threshold, so that a player i with signal x_i revolts if and only if $x_i \leq x^*$. Assuming continuity, the marginal player whose private signal is exactly x^* must be indifferent between revolting and not revolting. The question of uniqueness and multiplicity then amounts to: How many different values of x^* make this marginal player indifferent given that other players take the same strategy?

Two conflicting forces are in play. First, when x^* increases, the marginal player believes that the regime is stronger, reducing his incentives to revolt. If this was the only force, then necessarily only one x^* could make the marginal player indifferent, and the equilibrium would necessarily be unique—under some auxiliary assumptions. However, there is also a second countervailing force. When x^* increases, it means that other players are more likely to revolt, which tends to raise the marginal player's incentive to revolt. But whether or not this force actually raises the marginal player's incentive to revolt hinges on his *belief* about the fraction of other players who revolt. The key relationship that one can establish is: absent public information about the regime's strength, the marginal player with signal x^* *always*, no matter what x^* is, believes that this fraction is equally likely to be any number between 0 and 1. That is, the marginal player's equilibrium belief about the fraction of players who revolt is always uniformly distributed between 0 and 1. In this sense, the marginal player's *strategic uncertainty* does not depend on his signal. Because this belief does not hinge on the value of x^*, the second force is moot, and the equilibrium is unique. Public information upsets this belief independence from x^*, opening the door for the possibility of multiple equilibria when the second, countervailing force is sufficiently strong.

A particular form of public information is when some public news reveals that the regime is not too strong, in the sense that the regime's strength is below some threshold. Then, similar to the previous case without any public information, the marginal player with signal x^* believes that the fraction of players who revolt is uniformly distributed. However, in contrast, now the lower bound of the support for this uniform distribution is increasing in x^*. This implies that higher x^* raises the marginal player's expectation of the fraction of revolters, thereby bringing back the second force (and with it the potential for multiple equilibria) that was absent with no public information.

2 Model

A continuum of citizens, indexed by $i \in [0, 1]$, must decide whether or not to revolt. Revolting has an opportunity cost of C. If the revolution succeeds, a participant obtains a benefit of $B > C$, but if it fails, he is punished, receiving a payoff of $S < C$. Let $Q \in [0, 1]$ be the likelihood that the revolution succeeds, so that a citizen's expected payoff from revolting versus not revolting is:

$$Q\ B + (1 - Q)\ S - C = Q(B - S) + (S - C) = (B - S)\ \left(Q - \frac{C - S}{B - S}\right).$$

Normalizing $B-S$ to 1 and defining $c \equiv C-S/B-S \in (0, 1)$, without loss of generality, we can assume that the payoff from revolting is 1 if the revolution succeeds, 0 if the revolution fails, and the payoff from not revolting is $c \in (0, 1)$. Let $a_i \in \{0, 1\}$ be citizen i's action with $a_i = 1$ corresponding to revolting and $a_i = 0$ corresponding to not revolting. Then, a citizen's problem is:

$$\max_{a_i \in \{0,1\}} \; a_i \; Q + (1 - a_i) \; c.$$

The probability of regime change Q is endogenous, and depends on the fraction of revolters and the regime's strength. In particular, let $n \equiv \int_{i=0}^{1} \; a_i \; di$ be the measure of revolters, and let $\theta \in \mathbb{R}$ capture the regime's strength, with higher θ corresponding to stronger regimes. We maintain mild assumptions on Q:

Assumption 1 *The probability of regime change $Q(n, \theta) : [0, 1] \times \mathbb{R} \to [0, 1]$ is a function of the fraction of revolters, n, and the regime's strength, θ, such that:*

1. *Action and State Monotonicity. $Q(n, \theta)$ is weakly increasing in n and weakly decreasing in θ.*
2. *Limit Dominance. There exists $\underline{\theta}, \bar{\theta} \in \mathbb{R}$, with $\underline{\theta} < \bar{\theta}$ such that:*

$$Q(1, \theta) < c \text{ if } \theta > \bar{\theta} \quad \text{and} \quad Q(0, \theta) > c \text{ if } \theta < \underline{\theta}.$$

3. *Continuity and Expectation Monotonicity: Let $F(\cdot)$ be an atomless cdf with full support on $\mathbb{R}$. $\int_{n=0}^{1} Q(n, \theta - F^{-1}(n)) \; dn$ is continuous and strictly decreasing in θ.*

If $Q(n, \theta)$ is continuous and differentiable with $\partial Q(n, \theta)/\partial \theta < 0$, then part 3 of Assumption 1 is automatically satisfied. However, the most commonly used class of $Q(n, \theta)$ shown in equation (1) below is discontinuous, with $\partial Q(n, \theta)/\partial \theta = 0$ almost everywhere. We call this the Bang-Bang revolution technology:

$$Q^b(n, \theta) = \begin{cases} 1 & ; n \geq \theta \\ 0 & ; n < \theta. \end{cases} \tag{1}$$

In contrast, $Q(n, \theta) = 1/(1 + e^{\theta - n})$ is continuous and differentiable with $\partial Q(n, \theta)/\partial \theta < 0$, and satisfies Assumption 1.

The limit dominance assumption means that when the regime's strength θ is sufficiently high, even if all citizens revolt ($n = 1$) so that security resources have to be divided among many tasks, still the likelihood of a failed revolution is high enough that it is not worth for a citizen to revolt. That is, no matter what others do, a citizen's optimal action is not to revolt. Conversely, when the regime's strength θ is sufficiently low, even if no one protests ($n = 0$) so that security resources can be concentrated, still the likelihood of regime change (for exogenous reasons) is high enough that it is optimal for a citizen to revolt. That is, a citizen has a *dominant strategy* to revolt. The absence of limit dominance opens the door for multiple equilibria. To see this suppose that $Q(0.5, \theta) > c$ for all θ. This implies that there is always an equilibrium in which all citizens revolt ($n = 1$ in equilibrium). Uncertainty about θ does not change this equilibrium: whatever anyone thinks that θ is, if all revolt, the revolution succeeds. Similarly, if $Q(0.1, \theta) < c$ for all θ, then there is an equilibrium in which no one revolts, and the revolution fails. Again, uncertainty about θ has no bearing on this equilibrium. Combining these observations, if $Q(0.1, \theta) < c < Q(0.5, \theta)$ for all θ, then we already know what, independent of the nature of uncertainty about θ, there will always be at least two equilibria: one in which all revolt and one in which no one does.[1]

Limit dominance by itself does not guarantee a unique equilibrium. The coordination nature of the game, that $Q(n, \theta)$ is weakly increasing in n can create multiplicity even when limit dominance is present. The classic Bang-Bang revolution technology is a well-known example:

Proposition 1 *Suppose $Q(n, \theta) = Q^b(n, \theta)$, and θ is known.*

- *If $\theta \leq 0$, then there is a unique equilibrium in which everyone revolt.*
- *If $\theta > 1$, then there is a unique equilibrium in which no one revolts.*
- *If $0 < \theta \leq 1$, then both above equilibria exist.*

3 Incomplete Information: Only Private Signals

Now, suppose θ is unknown and uncertain and citizens have private information about θ. In particular, θ is distributed with (improper) uniform prior over $\mathbb{R}$, and each citizen i receives a private signal $x_i = \theta + \epsilon_i$, with $\epsilon_i \sim iid\, F(\cdot)$, where $F(\cdot)$ is an atomless cdf with full support on $\mathbb{R}$ and corresponding pdf $f(\cdot)$. The strategy of a citizen i, $s_i : \mathbb{R} \rightarrow \{0, 1\}$ is a mapping from his signal into a decision of whether or not to revolt. We focus on symmetric monotone strategies, so that

$$s(x_i) = \begin{cases} 1 & ; x_i \leq x^* \\ 0 & ; x_i > x^*. \end{cases}$$

[1]See Bueno de Mesquita (2010, 2014) for discussions of limit dominance in revolution models.

It is easy to show that the best response to a monotone strategy is a monotone strategy. Thus, in a symmetric monotone equilibrium with equilibrium cutoff x^*, a player i who receives signal $x_i = x^*$ must be indifferent between revolting and not revolting. Recognizing that for a given θ, $n(\theta) = Pr(x_j \leq x^*|\theta) = F(x^* - \theta)$, we have:

$$\begin{aligned} c &= E[Q(n^*, \theta)|x_i = x^*] \\ &= \int_{\theta=-\infty}^{\infty} Q(F(x^* - \theta), \theta)\ f(x^* - \theta)\ d\theta \\ &= \int_{z=0}^{1} Q(z, x^* - F^{-1}(z))\ dz, \end{aligned} \tag{2}$$

where we have used the change of variables from θ to $z = F(x^* - \theta)$. Assumptions 1 on Q delivers a unique solution for x^*.

Proposition 2 *There is a unique equilibrium in which citizens use the monotone strategy with the cutoff point x^*, where x^* is the unique solution to equation (2).*

Corollary 1 *With the classic Bang-Bang revolution technology $Q(n, \theta) = Q^b(n, \theta)$, we have: $x^* = 1 - c + F^{-1}(1 - c)$. Further, $Q(n, \theta) = 1$ if and only if $\theta \leq \theta^* \equiv 1 - c$.*

4 Public News

Now, suppose citizens observe some public information, indicating that the state is relatively week. In particular, suppose it becomes common knowledge that $\theta \leq \hat{\theta}$. To maintain limit dominance, we assume that $\hat{\theta} > \bar{\theta}$, so that $Q(1, \hat{\theta}) < c$. That is, even though θ is below a threshold, this threshold is still large enough so that around it citizens have a dominant strategy not to revolt. If this was not true, then there was always an equilibrium in which everyone revolts and the revolution succeeds. To ensure that even in the presence of new public information that $\theta \leq \hat{\theta}$, higher private signals x_i indicate higher θ in expectation, we also impose an additional assumption that is satisfied by all commonly used distributions.

Assumption 2 *$F(\cdot)$ is log-concave.*

Given that others take a cutoff strategy with cutoff x^*, citizen i's net expected payoff from revolting is:

$$E[Q(n^*, \theta)|\theta \leq \hat{\theta}, x_i] - c = \int_{\theta=-\infty}^{\infty} Q(F(x^* - \theta), \theta)\ f(\theta|\theta \leq \hat{\theta}, x_i)\ d\theta - c,$$

where $f(\theta|\theta \leq \hat{\theta}, x_i)$ is the pdf of θ conditional on $\theta \leq \hat{\theta}$ and x_i. Because $f(\theta|\theta \leq \hat{\theta}, x_i)$ still satisfies MLRP, the expectation is still decreasing in x_i, with the limits being larger than c when x_i is very negative, and smaller than c when x_i is very large. Thus, the best response to a monotone strategy is still a monotone strategy. Thus, equation (2) becomes:

$$\begin{aligned}
c &= E[Q(n^*,\theta)|\theta \leq \hat{\theta}, x_i = x^*] \\
&= \int_{\theta=-\infty}^{\infty} Q(F(x^*-\theta),\theta)\ f(\theta|\theta \leq \hat{\theta}, x_i = x^*)\ d\theta \qquad (3)\\
&= \int_{\theta=-\infty}^{\hat{\theta}} Q(F(x^*-\theta),\theta)\ \frac{f(\theta|x_i = x^*)}{Pr(\theta \leq \hat{\theta}|x_i = x^*)}\ d\theta \\
&= \int_{\theta=-\infty}^{\hat{\theta}} Q(F(x^*-\theta),\theta)\ \frac{f(x^*-\theta)}{1-F(x^*-\hat{\theta})}\ d\theta \\
&= \int_{z=F(x^*-\hat{\theta})}^{1} \frac{Q(z, x^*-F^{-1}(z))}{1-F(x^*-\hat{\theta})}\ dz. \qquad (4)
\end{aligned}$$

Equilibrium Characterization. Let $I(x^*, \hat{\theta}, c) \equiv E[Q(n^*,\theta)|\theta \leq \hat{\theta}, x_i = x^*] - c$, and note that the indifference condition is: $I(x^*, \hat{\theta}, c) = 0$. Any solution x^* to this indifference equation constitutes an equilibrium. Obviously, when $\hat{\theta} \to \infty$, equation (4) simplifies to (2), and we have a unique equilibrium. Another sufficient condition for uniqueness is that when $I(x^*, \hat{\theta}, c) = 0$, the derivative (assuming it exists) of $I(x^*, \hat{\theta}, c)$ w.r.t. x^* is negative, so that $I(x^*, \hat{\theta}, c)$ can only cross zero from above, and hence only once. When Q is differentiable, differentiating the left hand side w.r.t. x^* at the equilibrium yields:

$$\begin{aligned}
\left.\frac{\partial E[Q(n^*,\theta)|\theta \leq \hat{\theta}, x_i = x^*]}{\partial x^*}\right|_{\text{at eq}} &= -\int_{z=F(x^*-\hat{\theta})}^{1} \frac{|Q_2(z, x^*-F^{-1}(z))|}{1-F(x^*-\hat{\theta})}\ dz \\
&+ \frac{f(x^*-\hat{\theta})}{1-F(x^*-\hat{\theta})}\ \{c - Q(F(x^*-\hat{\theta}),\hat{\theta})\},
\end{aligned}$$

where we have used equation (4) to replace for c. Observe that the first term is negative, and the second term is positive due to our limit dominance assumption. When $\hat{\theta}$ is sufficiently large, there is a unique equilibrium. The reason is that $\frac{f(x^*-\hat{\theta})}{1-F(x^*-\hat{\theta})}$ approaches 0, while $\{c - Q(F(x^* - \hat{\theta}),\hat{\theta})\}$ is bounded, and hence the second term shrinks to zero. In contrast, the first term remains negative and bounded away from 0. When $\hat{\theta}$ is smaller, multiple equilibria can arise. For example, for $Q(n,\theta) = Q^b(n,\theta)$ and $F(\cdot) = N(0,1)$, there are multiple equilibria in some ranges of c.

5 Multiplicity and Uniqueness

What underlies equilibrium multiplicity and uniqueness in regime change games? Why the fact that everyone knows $\theta \leq \hat{\theta}$ can overturn our simple uniqueness result when there was no common knowledge about θ? In equation (2), the right hand side was strictly decreasing in x^*, implying a unique symmetric monotone equilibrium. The presence of a threshold on θ significantly complicates the behavior of the expectation for the marginal player with signal $x_i = x^*$. We provide two forms of intuition. First, consider the mechanical effects. Consider equation (3). When x^* increases, (1) everyone else are more likely to act, raising a player's incentives ($F(x^* - \theta)$ increases), but (2) the marginal player's belief $f(\theta|\theta \leq \hat{\theta}, x_i = x^*)$ moves to the right, making him believe that θ is higher and reducing his incentives to act. When $\hat{\theta} = \infty$, the latter effect dominates. But when $\hat{\theta} < \infty$ restricts how pessimistic the marginal player can get (note that $f(\theta|\theta \leq \hat{\theta}, x_i = x^*)$ puts zero weight on $\theta > \hat{\theta}$), the second effect is weakened, opening the door for non-monotonicities in x^* and multiple equilibria.

Now, we provide a deeper intuition. Consider the fraction of players who act at a given θ: $n^*(\theta) = Pr(x_j \leq x^*|\theta) = F(x^* - \theta)$. Different players receive different signals, hence have different posteriors about θ, and hence different posteriors about $n^*(\theta)$. First, observe that when θ has an improper prior, there is no difference between x_i and θ, and hence one can think of θ as a signal of x_i. Then,

$$Pr(\theta \leq A|x_i) \quad = \quad Pr(x_i - \epsilon_i \leq A|x_i) = Pr(x_i - A \leq \epsilon_i|x_i) = 1 - F(x_i - A). \tag{5}$$

Let $H(\bar{n}|x_i) = Pr(n^*(\theta) \leq \bar{n}|x_i)$, and consider the marginal player's posterior belief about the fraction of players who act:

$$\begin{aligned} H(\bar{n}|x_i = x^*) \quad &= \quad Pr(n^*(\theta) \leq \bar{n}|x_i) = Pr(F(x^* - \theta) \leq \bar{n}|x_i) \\ &= \quad Pr(x^* - F^{-1}(\bar{n}) \leq \theta|x_i) = 1 - Pr(\theta < x^* - F^{-1}(\bar{n})|x_i) \\ &= \quad F(x_i - x^* + F^{-1}(\bar{n})) \quad \text{(from (5))} \\ &\overbrace{=}^{x_i = x^*} \quad F(F^{-1}(\bar{n})) = \bar{n}. \end{aligned} \tag{6}$$

That is, the marginal player with signal $x_i = x^*$ *always* believes, no matter what x^* is, that the fraction of players who act is uniformly distributed on $[0, 1]$:

$$n^*(\theta)|x_i = x^* \sim U[0, 1]. \tag{7}$$

Curiously, some of the algebra to show this result resembles the familiar processes of (i) generating a general random variable from the uniform distribution, or (ii) showing that p-values are uniform under the null. I do not know whether a deep connection exists, but a fundamental intuition can be derived from reinterpreting the algebra. Consider the following thought experiment. Suppose after the signals are realized player i with signal x_i is interested to know his percentile in the population—in the usual order. Because there is a continuum of citizens, when i is in the $\bar{n}$th percentile, it means that $\bar{n}$ percent of the population have signals below him. Thus, i's question of which percentile he is in is the same as the question of what percentage of the population have a signal lower than him. Thus, for a citizen with signal $x_i = x^*$, knowing $H(\cdot)$ above is the same as his belief about his percentile in the population.[2] Now, one can perform the algebra above, or just argue: because i has no information about θ, he does not have any information about his ranking in a realization of signals. Thus, his belief must be uniform.

Thus, higher x^* only makes the marginal player with signal x^* directly more pessimistic about θ without changing his belief about the fraction of players who revolt. Thus, raising x^* always reduces his incentives to revolt, yielding a unique equilibrium. The crux of this logic is equation (5) that hinges on the absence of any prior information about θ. Without it, the chain of argument in equation (6) breaks down (at the fifth equality), and the marginal player's posterior belief about the fraction of player who act is not *always* uniform anymore. Critically, now it varies with x^*, underlying the conflicting force that tends to raise his incentive to revolt. In particular, when $\theta \leq \hat{\theta}$, mirroring the above calculations, one can show that now the marginal player believes that:

$$n^*(\theta)|x_i = x^* \sim U[F(x^* - \hat{\theta}), 1]. \tag{8}$$

To see this, observe that because $\theta \leq \hat{\theta}$, we know that:

$$n^*(\theta) = Pr(x_j \leq x^*|\theta) = F(x^* - \theta) \in [F(x^* - \hat{\theta}), 1).$$

In particular,

$$F(x^* - \hat{\theta}) \leq n^*(\theta) \text{ for all } \theta \leq \hat{\theta} \quad \Rightarrow \quad Pr(n^*(\theta) \leq F(x^* - \hat{\theta})) = 0 \text{ for all } \theta \leq \hat{\theta}.$$

[2] What is i's belief that he is in the $\bar{n}$th percentile? If i knew θ, then he could calculate $F(x_i - \theta)$, which would be his percentile. Because he does not know θ, he believes that the probability that he is in the $\bar{n}$th percentile is $Pr(F(x_i - \theta) \leq \bar{n}|x_i)$. In particular, a citizen with signal x^* believes that he is in $\bar{n}$th percentile with probability $Pr(F(x^* - \theta) \leq \bar{n}|x_i = x^*) = Pr(n^*(\theta) \leq \bar{n}|x_i = x^*)$, which we calculated in (6).

But what about the distribution of $n^*(\theta)$ for higher values? Equation (8) asserts that it is uniform. To show this, we mirror our earlier calculations. Now, equation (5) becomes:

$$Pr(\theta \leq A|x_i, \theta \leq \hat{\theta}) = \begin{cases} 1 & ; A \geq \hat{\theta} \\ \frac{1-F(x_i-A)}{1-F(x_i-\hat{\theta})} & ; A \leq \hat{\theta}. \end{cases} \tag{9}$$

And, we show in the Online Appendix that the calculations that yielded (6) become:

$$H(\bar{n}|x_i = x^*, \theta \leq \hat{\theta}) = \begin{cases} 0 & ; \bar{n} \leq F(x^* - \hat{\theta}) \\ \frac{\bar{n}-F(x^*-\hat{\theta})}{1-F(x^*-\hat{\theta})} & ; \bar{n} \geq F(x^* - \hat{\theta}). \end{cases} \tag{10}$$

We end by observing that if we allow the noise in private signals to go to zero, players with such a precise private information effectively ignore their public information, and the problem represented in (4) and (8) is transformed back to (2) and (7). To see this, adjust the noise structure by introducing a scale parameter σ for noise, so that $x_i = \theta + \sigma\epsilon_i$. Then, $F(x^* - \hat{\theta})$ in (8) becomes $F(\frac{x^*-\hat{\theta}}{\sigma})$. From our limit dominance assumption, in any equilibrium where noise is very small, obviously, $x^* < \hat{\theta}$. Therefore, $\lim_{\sigma \to 0} F(\frac{x^*-\hat{\theta}}{\sigma}) = 0$. That is, (8) becomes (7), and the equilibrium is characterized by equation (2) again.

The intuition for this point is that when the noise is very precise, Bayesian players put negligible weight on their imprecise, public information, and hence act as if they have no public information. This observation can be made more generally and formally. When the noise in private information approaches zero, then (5) will still be true (as opposed to, e.g., (9)), and hence the marginal citizen's belief will remain uniform—because one can replicate the steps leading to equation (6) as we show in the Online Appendix.

These results beg the question of whether there is any difference between the case with no public information versus the case with public information, but vanishingly small noise. Consider the common Bang-Bang revolution technology. As Corollary 1 shows, with no public information, the regime collapses whenever it is sufficiently weak: $\theta < \theta^* = 1 - c$. Our above calculations imply the same result holds when there is public information, but the noise is vanishingly small. However, in this latter case, when $\theta < \theta^*$, so that the regime is going to collapse, then almost everyone will revolt. The reason is that given their very accurate signals, players almost perfectly estimate that the regime will collapse, and hence revolt. This result may seem odd. But one must bear in mind that the uniqueness is robust in the sense that even when the noise is not almost zero, as far as it is not too large, the equilibrium remains unique, in which the regime collapses whenever the regime's strength θ is lower than some threshold that is close to $1 - c$.

6 References

Boix, Carles, and Milan Svolik. 2013. "The Foundations of Limited Authoritarian Government: Institutions and Power-Sharing in Dictatorships." *Journal of Politics* 75: 300-16.

Bueno de Mesquita, Ethan. 2010. "Regime Change and Revolutionary Entrepreneurs." *American Political Science Review* 104: 446-66.

Bueno de Mesquita, Ethan. 2014. "Regime Change and Equilibrium Multiplicity." Mimeo.

Casper, Brett, and Scott Tyson. 2014. "Popular Protest and Elite Coordination in a Coup d'etat." *Journal of Politics* 76: 548-64.

Chen, Heng, Yang K. Lu, and Wing Suen. 2016. "The Power of Whispers: A Theory of Rumor, Communication and Revolution." *International Economic Review* 57: 89-116.

Edmond, Chris. 2013. "Information Manipulation, Coordination and Regime Change." *Review of Economic Studies* 80: 1422-58.

Little, Andrew, Joshua Tucker, and Tom LaGatta. 2015. "Elections, Protest, and Alternation of Power." *Journal of Politics* 77: 1142-56.

Loeper, Antoine, Jakub Steiner, and Colin Stewart. 2014. "Influential Opinion Leaders." *Economic Journal* 124: 1147-67.

Morris, Stephen, and Hyun Song Shin. 2003. "Global Games: Theory and Application." In *Advances in Economics and Econometrics, Eighth World Congress, Vol. I,* ed. by Dewatripont, Hansen, & Turnovsky. Cambridge University Press.

Rundlett, Ashlea, and Milan Svolik. 2016. "Deliver the Vote! Micromotives and Macrobehavior in Electoral Fraud." *American Political Science Review* 110: 180-97.

Shadmehr, Mehdi. 2017. "Capitalists in Revolution." Mimeo.

Shadmehr, Mehdi, and Dan Bernhardt. 2017. "When Can Citizen Communication Hinder Successful Revolution?" *Quarterly Journal of Political Science* 12: 301-23.

Tyson, Scott, and Alastair Smith. 2018. "Dual-Layered Coordination and Political Instability: Repression, Cooptation, and the Role of Information." *Journal of Politics* 18: 44-58.

Online Appendix

Proof of Corollary 1: We provide the proof for the case of $x_i = \theta + \sigma\ \epsilon_i$, where we have added the scale σ, which is set to 1 in the text. Observe that (2) becomes:

$$\begin{aligned} c &= E[Q(n^*,\theta)|x_i = x^*] \\ &= \int_{\theta=-\infty}^{\infty} Q(F(\frac{x^*-\theta}{\sigma}),\theta)\ \frac{1}{\sigma}\ f(\frac{x^*-\theta}{\sigma})\ d\theta \\ &= \int_{z=0}^{1} Q(z, x^* - \sigma\ F^{-1}(z))\ dz, \end{aligned} \tag{11}$$

Now, observe that,

$$Q^b(z, x^* - \sigma\ F^{-1}(z)) = \begin{cases} 1 & ; z \geq x^* - \sigma\ F^{-1}(z) \\ 0 & ; z < x^* - \sigma\ F^{-1}(z). \end{cases}$$

For a given x^*, let z^* be the unique solution to $z = x^* - \sigma\ F^{-1}(z)$. Then, $z \geq x^* - \sigma\ F^{-1}(z)$ if and only if $z \geq z^*$. Then, from equation (11),

$$c \;=\; \int_{z=0}^{1} Q(z, x^* - \sigma\ F^{-1}(z))\ dz = \int_{z=z^*}^{1} dz = 1 - z^* \quad \Rightarrow \quad z^* = 1 - c.$$

Thus, $x^* = z^* + \sigma\ F^{-1}(z^*) = 1 - c + \sigma\ F^{-1}(1-c)$. Further, from the derivation of equation (11), recall that $z = n^*(\theta)$ and $\theta = x^* - \sigma\ F^{-1}(z)$. Thus, $n^*(\theta) \geq \theta$ if and only if $z \geq z^*$. Because $\theta = x^* - \sigma\ F^{-1}(z)$, we have $z \geq z^*$ if and only if $\theta \leq \theta^* = x^* - \sigma\ F^{-1}(z^*) = 1 - c$. □

Proof of Equation 10:

$$\begin{aligned} H(\bar{n}|x_i = x^*, \theta \leq \hat{\theta}) &= Pr(n^*(\theta) \leq \bar{n}|x_i, \theta \leq \hat{\theta}) \\ &= Pr(F(x^* - \theta) \leq \bar{n}|x_i, \theta \leq \hat{\theta}) \\ &= Pr(x^* - F^{-1}(\bar{n}) \leq \theta|x_i, \theta \leq \hat{\theta}) \\ &= 1 - Pr(\theta < x^* - F^{-1}(\bar{n})|x_i, \theta \leq \hat{\theta}) \\ &= \begin{cases} 0 & ; \bar{n} \leq F(x^* - \hat{\theta}) \\ \frac{F(x_i - x^* + F^{-1}(\bar{n})) - F(x^* - \hat{\theta})}{1 - F(x^* - \hat{\theta})} & ; \bar{n} \geq F(x^* - \hat{\theta}) \quad \text{(from (9))} \end{cases} \\ &\overbrace{=}^{x_i = x^*} \begin{cases} 0 & ; \bar{n} \leq F(x^* - \hat{\theta}) \\ \frac{F(F^{-1}(\bar{n})) - F(x^* - \hat{\theta})}{1 - F(x^* - \hat{\theta})} & ; \bar{n} \geq F(x^* - \hat{\theta}) \end{cases} \\ &= \begin{cases} 0 & ; \bar{n} \leq F(x^* - \hat{\theta}) \\ \frac{\bar{n} - F(x^* - \hat{\theta})}{1 - F(x^* - \hat{\theta})} & ; \bar{n} \geq F(x^* - \hat{\theta}). \end{cases} \end{aligned}$$

□

Proof that equation (5) holds if the prior is not uniform, but noise is vanishingly small: Public news simply changes the prior from uniform to some other distribution. Let $g(\theta)$ be that distribution. Then, as shown in Morris and Shin (2003),

$$\begin{aligned}
Pr(\theta \leq A | x_i = x^*) &= \int_{\theta=-\infty}^{A} pdf(\theta | x^*) d\theta \\
&= \int_{\theta=-\infty}^{A} \frac{pdf(x^*|\theta)\ g(\theta)}{\int_{-\infty}^{\infty} pdf(x^*|\theta)\ g(\theta)\ d\theta}\ d\theta \\
&= \int_{\theta=-\infty}^{A} \frac{f(\frac{x^*-\theta}{\sigma})\ g(\theta)}{\int_{-\infty}^{\infty} f(\frac{x^*-\theta}{\sigma})\ g(\theta)\ d\theta}\ d\theta \\
&= \int_{z(A)}^{z=\infty} \frac{f(z)\ g(x^* - \sigma\ z)}{\int_{-\infty}^{\infty} f(z) g(x^* - \sigma\ z)\ dz}\ dz \\
&= 1 - F(z(A)) \quad \text{(in the limit when } \sigma \to 0) \\
&= 1 - F\left(\frac{x^* - A}{\sigma}\right) \\
&= 1 - Pr(x_i \leq x^* | \theta = A).
\end{aligned}$$

□