Regime Change and Equilibrium Multiplicity*

Ethan Bueno de Mesquita†

April 17, 2014

Abstract

At least since Schelling (1960), theorists have argued that mass uprisings are a coordination problem. As such, it is striking that much recent game theoretic work on mass uprisings focuses on models known as "global games of regime change", which typically yield a unique equilibrium. This raises questions about the validity of important arguments—such as Schelling's (1960) analysis of spontaneous revolutions or Weingast's (1997) analysis of the foundations of accountability—which rely on equilibrium multiplicity. I argue that the assumptions driving equilibrium uniqueness in global games are not compelling for models of mass uprisings. I also show that it is possible to model mass uprisings in a way that retains the attractive features of global games while avoiding such assumptions. The analysis highlights the importance of introducing uncertainty into models of mass uprisings in substantively motivated ways that do not, inadvertently, rule out deep arguments about the strategic logic of such phenomena.

*I have received helpful input from Scott Ashworth, Sandeep Baliga, Amanda Friedenberg, Cathy Hafer, B. Pablo Montagnes, Stephen Morris, David Myatt, Gerard Padro i Miguel, and Mehdi Shadmehr.

†Associate Professor, Harris School, University of Chicago, e-mail: bdm@uchicago.edu

1 Introduction

At least since Schelling (1960), theorists have argued that credible revolutionary threats are as much a problem of coordination as of collective action. The basic argument is this. Consider a complete information setting in which a regime falls if and only if enough people mobilize and in which people want to mobilize if and only if the regime will in fact fall. In such an environment, there are two pure strategy equilibria: one in which no citizens mobilize and one in which the full citizenry mobilizes. On this sort of analysis, mass uprisings occur only when citizens' expectations of one another are properly coordinated.

This logic is important for at least two reasons. First, it offers a theoretical framework within which we can make sense of the often seemingly spontaneous nature of mass uprisings (Schelling, 1960; Kuran, 1989; Hardin, 1996). Second, to the extent that credible revolutionary threats lie at the heart of self-enforcing political accountability, it provides a way to understand variation in governance outcomes, absent institutional variation (Weingast, 1997; Fearon, 2011).

In light of these traditions, it is striking that much recent game theoretic work on mass uprisings focuses on a class of models known as "global games of regime change" (Edmond, 2013; Egorov, Guriev and Sonin, 2009; Persson and Tabellini, 2009; Little, 2012; Boix and Svolik, 2013).[1] While these models incorporate much of the structure of coordination games, they typically yield a unique equilibrium. As such, this technology seems to rule out the sort of coordination arguments that have dominated much of our thinking about mass uprisings for half a century.

The purpose of this paper is to interrogate whether the assumptions that underly equilibrium uniqueness in global games of regime change are in fact well suited to the study of mass uprisings. A lot is at stake here. In particular, if we were to conclude that equilibrium uniqueness is in fact a natural feature of models of regime change, this would suggest that analyses in the tradition of Schelling (1960) or Weingast (1997) are misguided. However, I argue that in fact the assumptions that drive equilibrium uniqueness are not for models of mass uprisings. Further, I show that it is possible to model regime change in a way that retains the most attractive features of global games of regime change, while avoiding such assumptions.[2]

[1]Such games have been used to model many phenomena of economic interest, such as currency attacks, bank runs, and debt crises (see, for examples, Morris and Shin (1998, 2000, 2004); Rochet and Vives (2004); Goldstein and Pauzner (2005); Corsetti, Guimaraes and Roubini (2006); Angeletos, Hellwig and Pavan (2006, 2007); Guimaraes and Morris (2007)).

[2]My model is not unique with respect to this latter point. See, for example, Baliga and Sjöström (2004); Chassang and Padro i Miguel (2010); Bueno de Mesquita (2010); Shadmehr and Bernhardt (2011*a*).

1.1 The Regime Change Approach: Advantages and Disadvantages

I use the term *regime change game* to refer to an incomplete information coordination game in which each player chooses whether to attack a regime and the regime falls if and only if enough players attack. Players may be uncertain about how strong the regime is, the preferences of other players, or a variety of other factors.

As applied models of mass uprisings, regime change games offer an important advantage over complete information coordination games. In particular, the presence of uncertainty "smooths" players' best response correspondences in a way that facilitates the study of a variety of important substantive phenomena.

In a pure strategy equilibrium of a standard complete information coordination game, typically either all players participate or all players refrain from participating. In such a setting, equilibrium behavior is not responsive to small changes in the structural environment (e.g., opportunity costs, regime strength, geography). Further, strategic actors—such as the government, the opposition, a revolutionary vanguard, or the media—can only change outcomes insofar as they can shift society from one focal equilibrium to another. Hence, in such models, changes are discontinuous—a structural or strategic intervention either doesn't matter at all, or it radically shifts the course of events. This fact limits the scope of phenomena that can potentially be explained with complete information coordination models.

By contrast, in a regime change game, the presence of uncertainty typically implies that there are pure strategy equilibria in which players use cutpoint strategies—a player participates if her private information is favorable enough and does not participate otherwise. When players use cutpoint strategies, behavior can change continuously in response to local changes in the structural environment or the strategic behavior of other actors. In particular, small changes in players' beliefs or payoffs may cause small changes in the cutpoint thereby inducing small changes in the level of mobilization or the probability of the regime falling. As a consequence, scholars have been able to use regime change models to study how the risk of mass uprisings relates to a host of important strategic and structural factors including the media and censorship (Egorov, Guriev and Sonin, 2009; Edmond, 2013); revolutionary vanguards and provocateurs (Bueno de Mesquita, 2010; Baliga and Sjöström, 2012); repression and counterrevolutionary mobilization (Shadmehr and Bernhardt, 2011*a*; Smith and Tyson, 2014); elections and electoral fraud (Little, 2012); power sharing institutions within autocracies (Boix and Svolik, 2013); the presence of weapons, peace-keepers, or military asymmetries (Chassang and Padro i Miguel, 2010); and economic development

(Persson and Tabellini, 2009).[3]

Given the importance of the analyses it facilitates, the smoothness of equilibrium behavior in regime change games is clearly a desirable feature of the modeling technology. But the standard regime change model applied to mass uprisings—a global game of regime change—has a second important feature: equilibrium uniqueness.

Whether equilibrium uniqueness is a desirable feature of a game used to study mass political uprisings is debatable and worth investigating. Equilibrium uniqueness rules out arguments in the traditions of Schelling (1960) and Weingast (1997), which rely on multiplicity. If uniqueness is in fact a robust feature of natural models of mass uprisings, perhaps these analyses are wrong and we have learned something important. But if, instead, uniqueness is not a robust and natural feature of regime change models, then by adopting a model with a unique equilibrium as canonical, we run the risk of ignoring an important and deep part of the strategic logic of mass protests.

Thus, the regime change approach raises two questions. First, is equilibrium uniqueness in fact a natural and robust feature of incomplete information models of mass uprisings? To answer this question we must examine whether or not the assumptions that drive equilibrium uniqueness are realistic or natural. Second, if not, is it possible to retain the desirable smoothness properties associated with global games of regime change while jettisoning unrealistic assumptions that drive equilibrium uniqueness?

In what follows I argue that, indeed, the assumptions underlying equilibrium uniqueness are not natural for models of mass uprisings. Hence, arguments along the lines suggested by Schelling (1960) and Weingast (1997) must remain an important part of our understanding of the strategic logic of such phenomena. Moreover, as I show, there is no tension between the benefits of smoothness that derive from the regime change approach and equilibrium multiplicity.

2 The Basic Argument

A regime change game becomes a global game, and thus has a unique equilibrium, if it satisfies technical assumptions on the informational environment. Two closely related conditions are critical: two-sided limit dominance and thick tails (Carlsson and van Damme, 1993; Chan and Chiu, 2002; Morris and Shin, 2003; Frankel, Morris and Pauzner, 2003). Two-sided limited dominance means that players' beliefs assign positive probability to a state of the world in which it is a dominant strategy to attack the regime and assign posi-

[3] Dewan and Myatt (2007, 2008) apply a closely related technology to study leadership.

		$N < T$	$N \geq T$
Player i	$a_i = 0$	0	0
	$a_i = 1$	$-k$	$\theta - k$

Figure 1: Payoffs for a Representative Player in the Regime Change Game

tive probability to a state of the world in which it is a dominant strategy not to attack the regime. Thick tails means the probability assigned to those states by the common prior is sufficiently large.

To examine the two questions discussed above, I study a canonical regime change game under two different assumptions on the informational environment. One of these assumptions——that players are uncertain of the payoffs from overturning the regime—induces one-sided limit dominance and, as I show, generically yields multiple equilibria that differ in terms of the their level of mobilization and the probability of regime change (all else equal). The other assumption—that players are uncertain of the regime's strength—induces two-sided limit dominance. If the uncertainty is large enough (thick tails), the model has a unique equilibrium.

Both types of uncertainty are substantively plausible. And both give rise to locally smooth equilibria. However, I argue that the conditions that give rise to uniqueness under the second informational environment are in fact unrealistic for models of mass uprisings. Hence, the natural model has the desirable smoothness properties, but is also consistent with the literatures that follow the arguments of Schelling (1960) and Weingast (1997), which rely on multiplicity of equilibria.

I build on the canonical regime change game developed by Morris and Shin (2004) and Angeletos, Hellwig and Pavan (2006, 2007). There is a continuum of individuals (of mass 1), each of whom makes a binary choice, $a_i \in \{0, 1\}$. There is regime change if the measure of players choosing $a_i = 1$, labeled N, is greater than a threshold T. Choosing to participate imposes cost k on the participant. Regime change yields a payoff of θ to the participants and zero to non-participants. Payoffs for a representative player are given in Figure 1.

Within this model, we can think of T as a measure of regime strength and θ as a measure of anti-government sentiment or the quality of potential replacement regimes. (See Meirowitz and Tucker (2013) for a model where citizens are uncertain about the quality of replacement regimes.) It is important to note that θ is a benefit of regime change that is enjoyed only by people who participate in the revolution.[4] One might, then, think of θ as being related to the chance that participants in the revolution will be specially privileged by a future regime or as involving some sort of "warm glow" benefit derived from having actively participated in a successful revolution.

These payoffs induce a complementarity between players that would give rise to multiplicity under complete information In particular, let p be the probability player i assigns to at least T other players participating. Player i wants to participate if and only if:

$$p\theta \geq k.$$

Hence, the more likely player i believes it is that others will participate, the more willing she is to participate.

Now, consider the two informational environments in turn. First, suppose there is uncertainty over the threshold T. As Morris and Shin (2004) and Angeletos, Hellwig and Pavan (2006, 2007) show, if players' beliefs assign positive probability to $T < 0$ and assign positive probability to $T > 1$, then the game has two-sided limit dominance. In particular, if $\theta > k$. then if $T < 0$, it is a dominant strategy to participate and if $T > 1$ it is a dominant strategy not to participate.

Next suppose there is uncertainty over the payoff from regime change, θ. This informational environment induces only one-sided limit dominance. If $\theta < k$, then it is a dominant strategy not to attack the regime. However, no matter how large θ is, it is never a dominant strategy to attack to regime.[5]

Both forms of uncertainty are substantively plausible. However, there is an important sense in which the overall model with uncertainty over T is less plausible. To see this, consider the assumption that citizens assign positive probability to the case of $T < 0$. This says that citizens believe it is possible that the regime will fall even if no one participates in a mass uprising. On its own, this is not an unreasonable assumption—regimes may fall for all manner of reasons without a rebellion. But the assumptions of the model go one step further. They say that, in the event that the regime is so weak that it will fall regardless of the presence of a rebellion, if a single measure-zero person turns out to protest, she

[4]All the results are robust to allowing some portion of θ to be enjoyed by everyone when the regime changes. What is important is that there be some share of θ that only goes to the participants.

[5]Note, payoff uncertainty need not lead to one-sided limit dominance in all regime change games. Nor need threshold uncertainty always lead to two-sided limit dominance in all regime change games. They simply happen to in the canonical regime change game form studied here. See, for example, Carlsson and van Damme (1993) and Morris and Shin (1998) for games with payoff uncertainty and two-sided limit dominance.

derives the extra benefit (θ) from having participated in the rebellion that "led" to the regime falling, even though she was in fact irrelevant to the outcome. This combination of two-sided limit dominance and the payoff structure of the regime change game seems hard to motivate. But, as we will see, it is critical for uniqueness.

In what follows, I start by characterizing the equilibrium correspondence for the regime change game under each form of uncertainty.

I show that the game with uncertainty over the payoff of regime change always has an equilibrium in which no one participates. I then show that the game may also have equilibria with positive participation. These equilibria are locally smooth in the sense described above—players use cutpoint strategies and the cutpoints respond continuously to small changes in the structural environment. Moreover, these positive participation equilibria are generically non-unique. In particular, generically, the game either exhibits zero or two equilibria with positive participation (although only one of these positive participation equilibria is stable in a natural sense of the term).

I then consider the model with uncertainty over T. I report the canonical results for restrictions on parameter values that yield equilibrium uniqueness and also describe the equilibrium correspondence when those conditions do not hold.

Next I provide some intuition for why the two models differ in terms of equilibrium uniqueness. I show that the key functions characterizing equilibrium in these games can be decomposed into three substantive effects. Understanding how these effects interact differently in the two games is key for developing the intuitions. In particular, studying these functions confirms that it is indeed the unrealistic assumption on payoffs under the condition that the regime would have fallen even absent any participation in mass uprisings that drives uniqueness in the game with uncertainty over T.

Taken together, these results suggest that equilibrium uniqueness is not a natural feature of models of mass uprisings. Rather, natural models have multiple equilibria and are thus compatible with the literature deriving from the arguments of Schelling (1960) and Weingast (1997). Moreover, such models have the key advantage associated with global games of regime change—cutpoint equilibria that are locally smooth and are, thus, amenable to use in substantive models.

3 Payoff Uncertainty

Consider a game with the players, strategies, and payoffs described in Section 2 and Figure 1. Let θ be the realization of a normally distributed random variable with mean m and variance σ_θ^2. Each player receives a signal $s_i = \theta + \epsilon_i$, where each ϵ_i is the realization of a normally distributed random variable with mean zero and variance σ_ϵ^2. The random variables are independent.[6]

Label as Γ^θ the game in which players receive the signals just described and face the payoff matrix in Figure 1. Define the set $\mathcal{R} = \{(m, \sigma_\epsilon, \sigma_\theta, T, k) \in \mathbb{R} \times \mathbb{R}_+^4\}$. A particular instance of this game, with parameter values $r \in \mathcal{R}$, is $\Gamma^\theta(r)$. The value of r is common knowledge.

In a game $\Gamma^\theta(r)$, following a signal s_i, a player has posterior beliefs about θ that are normally distributed with mean $\overline{m}_i = \lambda s_i + (1-\lambda)m$ and variance $\sigma_\lambda^2 = \lambda \sigma_\epsilon^2$, with $\lambda = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\epsilon^2}$. Let Φ be the cumulative distribution function of the standard normal distribution, with associated probability density function ϕ.

I study symmetric, Bayesian equilibria in cutoff strategies. That is, profiles in which all players adopt the same strategy and that strategy takes the form, "choose $a_i = 1$ if and only if my signal crosses some cutoff." I refer to such an equilibrium as a *cutoff equilibrium*. If the cutoff rule is finite (i.e., the player participates if the signal crosses some finite cutoff), I refer to the equilibrium as a *finite cutoff equilibrium* (and similarly for an *infinite cutoff equilibrium*).

Suppose player i believes all players j participate if and only if $s_j \geq \hat{s}$. Then, for a given θ, player i anticipates total participation $1 - \Phi\left(\frac{\hat{s}-\theta}{\sigma_\epsilon}\right)$. Hence, player i believes regime change will be achieved if and only if $\theta \geq \theta^*(\hat{s}; r)$ with

$$1 - \Phi\left(\frac{\hat{s} - \theta^*(\hat{s}; r)}{\sigma_\epsilon}\right) = T \Rightarrow \theta^*(\hat{s}; r) = \hat{s} - \Phi^{-1}(1-T)\sigma_\epsilon. \tag{1}$$

From the perspective of a player who receives the signal s_i, and believes all other players use the cutoff rule $\hat{s}$, the probability of successful regime change is $1 - \Phi\left(\frac{\theta^*(\hat{s};r) - \overline{m}_i}{\sigma_\lambda}\right)$. Such a player will participate if

$$\left(1 - \Phi\left(\frac{\theta^*(\hat{s}; r) - \overline{m}_i}{\sigma_\lambda}\right)\right) \mathbb{E}[\theta | \theta \geq \theta^*(\hat{s}; r), s_i] - k \geq 0.$$

From standard facts about the expectation of a truncated normal random variable (see

[6]Shadmehr and Bernhardt (2011*b*) study a related, two player, regime change games with payoff uncertainty. Their focus is on conditions under which actions are strategic substitutes and under which actions are strategic complements.

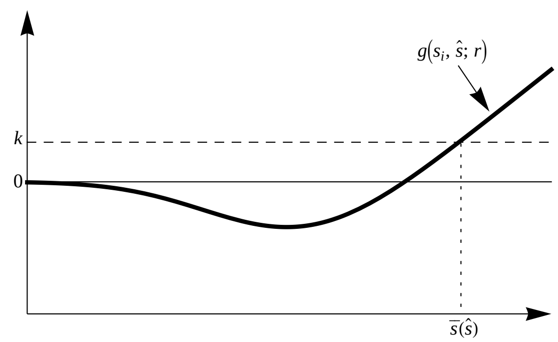

Figure 2: The function $g(\cdot, \hat{s}; r)$ crosses k exactly once.

Greene (2003), Theorem 22.2), a player who receives the signal s_i believes

$$\mathbb{E}[\theta|\theta \geq \theta^*(\hat{s};r), s_i] = \overline{m}_i + \sigma_\lambda \frac{\phi\left(\frac{\theta^*(\hat{s};r)-\overline{m}_i}{\sigma_\lambda}\right)}{1-\Phi\left(\frac{\theta^*(\hat{s};r)-\overline{m}_i}{\sigma_\lambda}\right)}.$$

So the conditions under which player i will participate, given that she believes all other players are using the cutoff rule $\hat{s}$, can be rewritten:

$$g(s_i, \hat{s}; r) \equiv \left(1 - \Phi\left(\frac{\theta^*(\hat{s};r) - \overline{m}_i}{\sigma_\lambda}\right)\right)\left(\overline{m}_i + \sigma_\lambda \frac{\phi\left(\frac{\theta^*(\hat{s};r)-\overline{m}_i}{\sigma_\lambda}\right)}{1-\Phi\left(\frac{\theta^*(\hat{s};r)-\overline{m}_i}{\sigma_\lambda}\right)}\right) \geq k. \qquad (2)$$

For a cutoff equilibrium to exist, player i must want to use a cutoff rule, given that she believes all others do so. Establishing this fact is subtle because her expected incremental benefit, $g(s_i, \hat{s}; r)$, need not be monotone increasing in s_i. (Because the normal distribution has the monotone hazard rate property, $\frac{\phi\left(\frac{\theta^*(\hat{s};r)-\overline{m}_i}{\sigma_\lambda}\right)}{1-\Phi\left(\frac{\theta^*(\hat{s};r)-\overline{m}_i}{\sigma_\lambda}\right)}$ is decreasing in $\overline{m}_i$.) Nonetheless, the following result establishes sufficient conditions for player i using a cutoff rule given that all others do. All proofs are in the appendix.

Lemma 3.1 *(i)* $\lim_{s_i \to -\infty} g(s_i, \hat{s}; r) = 0$

(ii) $\lim_{s_i \to \infty} g(s_i, \hat{s}; r) = \infty$

(iii) There is exactly one $\overline{s}(\hat{s})$ *satisfying* $g(\overline{s}(\hat{s}), \hat{s}; r) = k$. *For all* $s_i < \overline{s}(\hat{s})$, $g(s_i, \hat{s}; r) < k$ *and for all* $s_i > \overline{s}$, $g(s_i, \hat{s}; r) > k$.

Figure 2 illustrates the points made in Lemma 3.1. In particular, $g(\cdot, \hat{s}; r)$ crosses zero and becomes monotone increasing before it does so. Hence, given that a player i believes all other players use the cutoff rule $\hat{s}$, she too wants to use a cutoff rule. In equilibrium, player i must want to use the cutoff rule $\hat{s}$. So an equilibrium cutoff rule must satisfy

$$g(\hat{s}, \hat{s}; r) = \left(1 - \Phi\left(\frac{\theta^*(\hat{s}; r) - \hat{m}}{\sigma_\lambda}\right)\right)\left(\hat{m} + \sigma_\lambda \frac{\phi\left(\frac{\theta^*(\hat{s};r)-\hat{m}}{\sigma_\lambda}\right)}{1 - \Phi\left(\frac{\theta^*(\hat{s};r)-\hat{m}}{\sigma_\lambda}\right)}\right) = k, \tag{3}$$

where $\hat{m} = \lambda\hat{s} + (1-\lambda)m$ is the mean of the posterior distribution of a player whose signal was $\hat{s}$. That is, a player whose signal is right at the cutoff must be indifferent between participating and not (i.e., her incremental benefit must equal her incremental cost). It will be useful to have notation for the incremental expected benefit from participating to a player of type $\hat{s}$ given that she believes all others use the cutoff rule $\hat{s}$.

$$G^\theta(\hat{s}; r) \equiv g(\hat{s}, \hat{s}; r).$$

I now provide several results that will help characterize the number of finite cutoff equilibria.

Lemma 3.2 *A finite cutoff rule, $\hat{s}$, is a finite cutoff equilibrium of $\Gamma^\theta(r)$ if and only if it satisfies*

$$G^\theta(\hat{s}; r) = k. \tag{4}$$

Lemma 3.3 *For all $r \in \mathcal{R}$, $G^\theta(\cdot; r)$ has the following properties:*

(i) $\lim_{\hat{s}\to\infty} G^\theta(\hat{s}; r) = 0$.

(ii) $\lim_{\hat{s}\to-\infty} G^\theta(\hat{s}; r) = -\infty$.

(iii) $G^\theta(\hat{s}; r)$ *has a single peak.*

Now label as r_{-k} a collection of parameters $r \in \mathcal{R}$ with the fifth component (i.e., k) removed. Notice that k has no effect on the value of $G(\hat{s}; r)$. Let R^*_{-k} be the set of parameter values satisfying the following: For any (r_{-k}, k) with $r_{-k} \in R^*_{-k}$, $\arg\max_{\hat{s}} G(\hat{s}; r) \geq 0$.

Now we have the following result.

Theorem 3.1 *(i) For any $r \in \mathcal{R}$, the game $\Gamma^\theta(r)$ has an infinite cutoff equilibrium.*

(ii) For any r_{-k}, there is an open set $O_{-}(r_{-k})$ such that, for $k \in O_{-}(r_{-k})$, the game $\Gamma^{\theta}(r_{-k}, k)$ has no finite cutoff equilibria.

(iii) For any $r_{-k} \in R^{*}_{-k}$

(a) There is an open set $O_{+}(r_{-k})$ such that, for $k \in O_{+}(r_{-k})$, the game $\Gamma^{\theta}(r_{-k}, k)$ has two finite cutoff equilibria;

(b) There is exactly one k such that the game $\Gamma^{\theta}(r_{-k}, k)$ has exactly one finite cutoff equilibrium.

An implication of this result is that, for the following topological definition of genericity, the game Γ^{θ} generically has either zero or two equilibria in finite cutoff strategies.[7]

Definition 3.1 *Let X be a topological space. Then a set E that is a subset of X is* non-generic on X *if it is meagre on X—i.e., if it is the union of countably many nowhere dense subsets of X.*

Theorem 3.2 *Endow $\mathcal{R}$ with the relative topology induced by considering $\mathcal{R}$ as a subspace of $\mathbb{R}^5$. Then the game Γ^{θ} has either zero or two finite cutoff equilibria, except on a set of parameter values that is non-generic on $\mathcal{R}$.*

The logic of Theorems 3.1 and 3.2 is as follows. The function $G^{\theta}(\cdot; r)$ is single peaked and goes to $-\infty$ as $\hat{s} \to -\infty$ and to 0 as $\hat{s} \to \infty$. Thus, except in the non-generic case where $\max_{\hat{s} \in \mathbb{R}} G^{\theta}(\hat{s}; r) = 0$, if G^{θ} crosses zero once, it does so twice. Put differently, for any r_{-k}, there is one and only one k such that the game $\Gamma^{\theta}(r_{-k}, k)$ has only one finite cutoff equilibrium. Thus, generically, the game Γ^{θ} either has no finite cutoff equilibria (if k is high enough relative to the rest of r_{-k} so that $\max_{\hat{s} \in \mathbb{R}} G^{\theta}(\hat{s}; r) < k$) or has two finite cutoff equilibria (if k is low enough relative to the rest of r_{-k} so that $\max_{\hat{s} \in \mathbb{R}} G^{\theta}(\hat{s}; r) > k$).

These facts are illustrated in Figure 3. Notice that here, since a player participates if her signal is *above* the cutoff rule, the "stringency" of the cutoff rule is increasing to the right on the x-axis of Figure 3.

Theorem 3.1 also points out that, for all parameter values, the game Γ^{θ} has an infinite cutoff equilibrium. In this equilibrium, no player participates. As we will see, such an equilibrium does not exist for any parameter values of Γ^{T}.

[7] Genericity is not actually important here. This is just a convenient way of bracketing knife-edge cases.

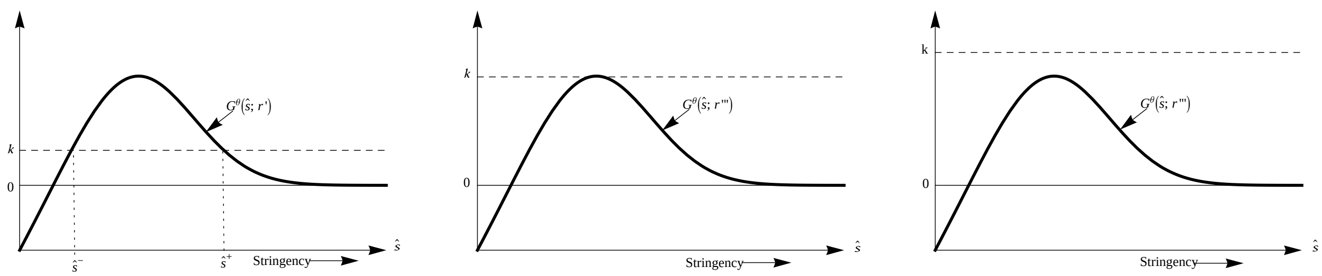

Figure 3: $G^{\theta}(\cdot; r)$ is single peaked in its first argument for all $r \in \mathcal{R}$. Thus, except on a non-generic set of parameter values, there are either two finite cutoff equilibria (as in the first cell) or no finite cutoff equilibria (as in the third cell). The non-generic case of a unique cutoff equilibrium is represented in the second cell. Moreover, for any collection of parameter values, there is an equilibrium with an infinitely stringent cutoff rule (i.e., no one participates).

3.1 Stability

The analysis above shows that, generically, the game Γ^{θ} has either one or three cutoff equilibria. For values of r where there is one cutoff equilibrium, the cutoff rule is infinite. For values of r where there are three cutoff equilibria, two involve finite cutoff rules and one involves an infinite cutoff rule. Thus, if there are any finite cutoff equilibria, there are generically two.

It is worth noting that the finite cutoff equilibrium with the larger (i.e., more stringent) cutoff rule can be thought of as unstable in a way that is analogous to the instability of the middle equilibrium of complete information games with strategic complements (Echenique and Edlin, 2004). To see this, imagine a simple learning dynamic, such as players playing best responses to the distribution of play in a pervious round. Suppose $\Gamma^{\theta}(r)$ has two finite cutoff equilibria. Label the higher (i.e., more stringent) equilibrium cutoff rule $\hat{s}^+$ and the lower (i.e., less stringent) equilibrium cutoff rule $\hat{s}^-$. Consider the equilibrium where all players use the cutoff rule $\hat{s}^+$. If play is slightly perturbed such that a few too many players participate, then players with types slightly lower than $\hat{s}^+$ want to participate, making more players want to participate, until everyone with a type greater than $\hat{s}^-$ is participating. Similarly, if a few too few players participate, then players with types slightly higher than $\hat{s}^+$ do not want to participate, making more players not want to participate, until no one is participating.

Given this instability, one might worry that the multiplicity described in Theorems 3.1

and 3.2, which will differentiate Γ^{θ} from Γ^{T}, is in some sense fragile.

However, this is not the case. As will be clear in the next section, the game Γ^{T} has open sets of parameter values for which there is a unique finite cutoff equilibrium, whereas Theorem 3.1 shows that a unique finite cutoff equilibrium is non-generic in Γ^{θ}. And Γ^{T} never has an infinite cutoff equilibrium, whereas Theorem 3.1 shows that Γ^{θ} always has an infinite cutoff equilibrium. Thus, even if one rules out the unstable finite cutoff equilibrium in Γ^{θ}, the equilibrium correspondences of the two games remain qualitatively different.

3.2 Local Smoothness and Comparative Statics

In the Introduction I argued that one advantage of the regime change game approach is smoothness of the equilibrium correspondence, which facilitates applications. It is clear that Γ^{θ} has this sort of smoothness.

To see this, fix parameter values such that a positive participation equilibrium exists. The cutpoint, $\hat{s}$ in such an equilibrium is characterized by:

$$\begin{aligned} G(\hat{s};r) \equiv & \left(1-\Phi\left(\frac{(1-\lambda)}{\sigma_\lambda}\hat{s}-\frac{(1-\lambda)m+\sigma_\epsilon\Phi^{-1}(1-T)}{\sigma_\lambda}\right)\right) \\ & \times\left(\lambda\hat{s}+(1-\lambda)m+\sigma_\lambda\frac{\phi}{1-\Phi}\left(\frac{(1-\lambda)}{\sigma_\lambda}\hat{s}-\frac{(1-\lambda)m+\sigma_\epsilon\Phi^{-1}(1-T)}{\sigma_\lambda}\right)\right)=k. \end{aligned}$$

Now it is straightforward from the implicit function theorem that $\hat{s}$ is locally differentiable in each of the parameters of the model. Hence, locally, behavior changes continuously with changes in regime strength (T), the costs of participation (k), the quality of private (σ_ϵ^2) and public (σ_θ^2) information, and prior beliefs about the payoffs from regime change (m), and so on.

3.3 Spontaneous Revolution

Despite the fact that the game Γ^{θ} is locally smooth, it provides two ways to understand spontaneous revolutions.

First, as established in Theorem 3.1, for a significant set of parameter values the game has multiple equilibria, including an equilibrium with zero participation and an equilibrium with positive participation. Hence, Schelling's (1960) argument about the possibility of spontaneous revolution due to a shift in focal equilibria holds in this model.

Second, as can be seen in Figure 3, while the equilibrium correspondence is almost-everywhere smooth, there is one discontinuity where the game goes from having only a zero-participation equilibrium to having both a zero-participation equilibrium and a positive participation equilibrium. At this point, a small change in parameter values could lead to a discontinuous change in mobilization, even without shifting the focal equilibrium.

To see this, fix parameter values other than k. Let's assume that players always select the equilibrium with the highest level of mobilization, given parameter values. Hence, there is no possibility for focal points here, since equilibrium selection is pinned down by the selection criterion. Nonetheless, it is possible to get spontaneous revolutions as a result of small changes in a parameter. Define

$$s^*(r_{-k}) = \arg\max_s G(s; (r_{-k}, k)).$$

Since G does not depend on k, s^* does not depend on k.

As is clear from Figure 3, if $k > G(s^*(r_{-k}); r_{-k}, k)$, then the game has only one equilibrium—zero participation. However, for any $k \leq G(s^*(r_{-k}); r_{-k}, k)$, the game has multiple equilibria. And, given the selection criterion, as k passes through $G(s^*(r_{-k}); r_{-k}, k)$ from above, the players will discontinuously shift from not participating at all, to positive participation and positive probability of regime change. This fact is illustrated in Figure 4. Hence, in addition to being consistent with Schelling's (1960) argument about spontaneous revolution due to shifts in focal points, this model also offers an account of spontaneous revolutions as a result of small changes to the structural environment, even without a shift in focal equilibria.

4 Threshold Uncertainty

Again consider a game with the players, strategies, and payoffs described in Section 2 and Figure 1. But now suppose θ is commonly known, but the threshold T is the realization of a random variable that is normally distributed with mean m and variance σ_T^2. Players receive signals $t_i = T + \xi_i$, where $\xi \sim \mathcal{N}(0, \sigma_\xi^2)$. Again, all random variables are independent.

Define the set $\mathcal{P} = \{(m, \sigma_\xi, \sigma_T, \theta, k) \in \mathbb{R} \times \mathbb{R}_+^4 | \theta > k\}$. An element of $\mathcal{P}$ is a vector of parameter values in which the cost of participation is lower than the payoff from certain success. I restrict attention to the set of parameter values $\mathcal{P}$ because if θ is (strictly) less than k, then it is a (strictly) dominant strategy not to participate for all players regardless of signal, an uninteresting case. (This assumption is standard.)

Label as Γ^T the game in which players receive the signals described above and face the

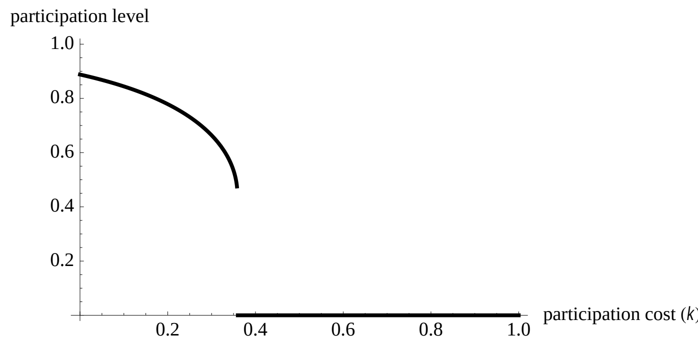

Figure 4: Equilibrium participation under the highest participation equilibrium as a function of k, for the case $\sigma_\theta^2 = 1$, $\sigma_\epsilon^2 = 1$, $T = 0.5$, $m = 0$, and $\theta = 0.5$. The figure shows that there is a critical threshold in k after which there is only a zero participation equilibrium. Thus, around that threshold, a small change in costs of participation leads to a discontinuous change in participation.

payoff matrix in Figure 1. A particular instance of this game, with parameters $p \in \mathcal{P}$, is $\Gamma^T(p)$. The value of p is common knowledge.

Consider a game $\Gamma^T(p)$ and posit a cutoff rule $\hat{t}$. If players use this cutoff rule, an individual will participate if $\xi_i \leq \hat{t} - T$. Given a state T, participation is $\Phi\left(\frac{\hat{t}-T}{\sigma_\epsilon}\right)$, so there is regime change if T is less than $T^*(\hat{t};p)$ given by:

$$\Phi\left(\frac{\hat{t} - T^*}{\sigma_\xi}\right) = T^*, \tag{5}$$

which can be rewritten

$$\hat{t} = \Phi^{-1}(T^*)\sigma_\xi + T^*. \tag{6}$$

A player who observes signal t_i has posteriors over T that are normally distributed with mean $m_i = \gamma t_i + (1-\gamma)m$ and variance $\sigma_\gamma^2 = \gamma\sigma_\xi^2$, with $\gamma = \frac{\sigma_t^2}{\sigma_t^2+\sigma_\xi^2}$. Thus, a player with signal t_i believes that the probability of victory, given all others use the cutoff rule $\hat{t}$, is

$$\Phi\left(\frac{T^*(\hat{t};p) - \overline{m}_i}{\sigma_\gamma}\right).$$

She will participate if

$$\Phi\left(\frac{T^*(\hat{t};p) - \gamma t_i - (1-\gamma)m}{\sigma_\gamma}\right)\theta \geq k.$$

The left-hand side is decreasing in t_i, so an equilibrium cutoff rule must satisfy

$$\Phi\left(\frac{T^*(\hat{t};p) - \gamma\hat{t} - (1-\gamma)m}{\sigma_\gamma}\right)\theta = k.$$

Substituting from Equation 6, this can be rewritten

$$G^T(\hat{t};p) \equiv \Phi\left(\frac{(1-\gamma)(T^*(\hat{t};p) - m) - \gamma\Phi^{-1}(T^*(\hat{t};p))\sigma_\xi}{\sigma_\gamma}\right)\theta = k. \tag{7}$$

Lemma 4.1 *A finite cutoff rule $\hat{t}$ is a finite cutoff equilibrium of the game $\Gamma^T(p)$ if and only if it satisfies $G^T(\hat{t};p) = k$.*

The above lemma gives necessary and sufficient conditions for the existence of a finite cutoff equilibrium. The next result establishes that these conditions can be met.

Lemma 4.2 *The following facts hold for all $p \in \mathcal{P}$:*

(i) $\lim_{\hat{t}\to\infty} G^T(\hat{t};p) = 0$

(ii) $\lim_{\hat{t}\to-\infty} G^T(\hat{t};p) = \theta$

(iii) For any finite $\hat{t}$, $G^T(\hat{t};p) \in (0,\theta)$

Taken together, these two Lemmata allow me to state the standard equilibrium uniqueness result for global games of regime change (Morris and Shin, 2004; Angeletos, Hellwig and Pavan, 2007). Let $\mathcal{P}^*$ be defined as follows:

$$\mathcal{P}^* = \left\{(m, \sigma_\xi, \sigma_T, \theta, k) \in \mathcal{P} | \sigma_\xi < \sigma_T^2\sqrt{2\pi}\right\}.$$

The restriction that σ_T^2 be sufficiently large is the "thick tails" assumption referred to earlier.

Proposition 4.1 *(Angeletos, Hellwig and Pavan, 2007) If $p \in \mathcal{P}^*$, then the game $\Gamma^T(p)$ has a unique finite cutoff equilibrium.*

Proposition 4.1 is illustrated in Figure 5. Since players participate if and only if their signal is less than the cutoff rule, moving to the right on the x-axis constitutes a *decrease* in the stringency of the cutoff rule. (This is the opposite of Figure 3.) When $\sigma_\xi < \sigma_T^2\sqrt{2\pi}$, the function G^T is everywhere decreasing in its first argument (i.e., increasing in stringency), so equilibrium uniqueness is guaranteed.

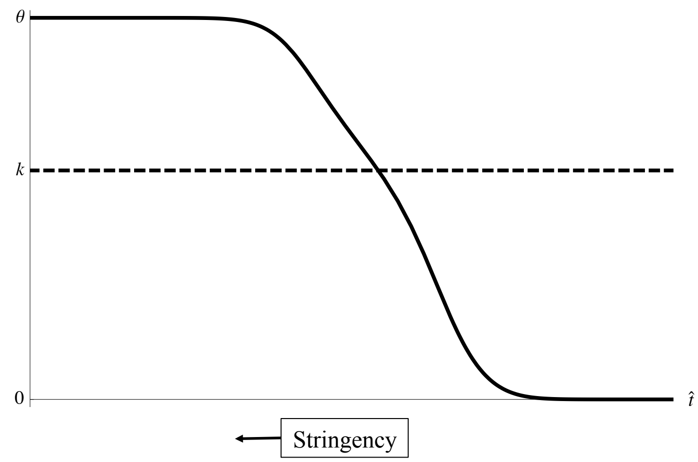

Figure 5: $G^T(\cdot;p)$ is decreasing (i.e., increasing in stringency) when $\sigma_\xi < \sigma_T^2\sqrt{2\pi}$. Thus, for For all $p \in \mathcal{P}^*$, the game $\Gamma^T(p)$ has a unique equilibrium.

What about the case where the standard assumptions that insure monotonicity of G^T do not hold? Even in that case, there is an open set of parameter values such that there is a unique equilibrium in cutoff strategies. The following lemma is the key step in the proof.

Lemma 4.3 *The function $G^T(\cdot;p)$ has either zero or exactly two critical points.*

To complete the characterization of the number of equilibria when uniqueness is not guaranteed, it will be useful to have a little more notation. Define the set $\mathcal{P}_{-k} = \mathbb{R} \times \mathbb{R}^3_+$. Now, for any $p = (m, \sigma_\eta, \sigma_T, \theta, k) \in \mathcal{P}$, let $p_{-k} \in \mathcal{P}_{-k}$ be p without its fifth component—i.e., be the quadruple $(m, \sigma_\eta, \sigma_T, \theta)$.

Theorem 4.1 *For any $p_{-k} \in \mathcal{P}_{-k}$, there is an open set $O \subset \mathbb{R}_+$ such that the game $\Gamma^T(p_{-k}, k)$ has a unique cutoff equilibrium for any $k \in O$.*

This result implies that $p \in \mathcal{P}^*$ is not necessary for uniqueness. For any collection of parameter values $p_{-k} \in \mathcal{P}_{-k}$ (i.e., even those not satisfying $\sigma_\xi < \sigma_T^2\sqrt{2\pi}$), there is an open set of costs of participation that imply a unique equilibrium.[8]

This situation, where there is equilibrium uniqueness for some, but not all, values of k, is illustrated in Figure 6. Here, the function G^T is not monotone. However, as shown

[8]It is worth clarifying that Theorem 4.1 is consistent with the claims in Morris and Shin (2004) and Angeletos, Hellwig and Pavan (2007) that, in my notation, $p \in \mathcal{P}^*$ is necessary and sufficient for uniqueness. They refer to uniqueness for all values of other parameters (e.g., for all m).

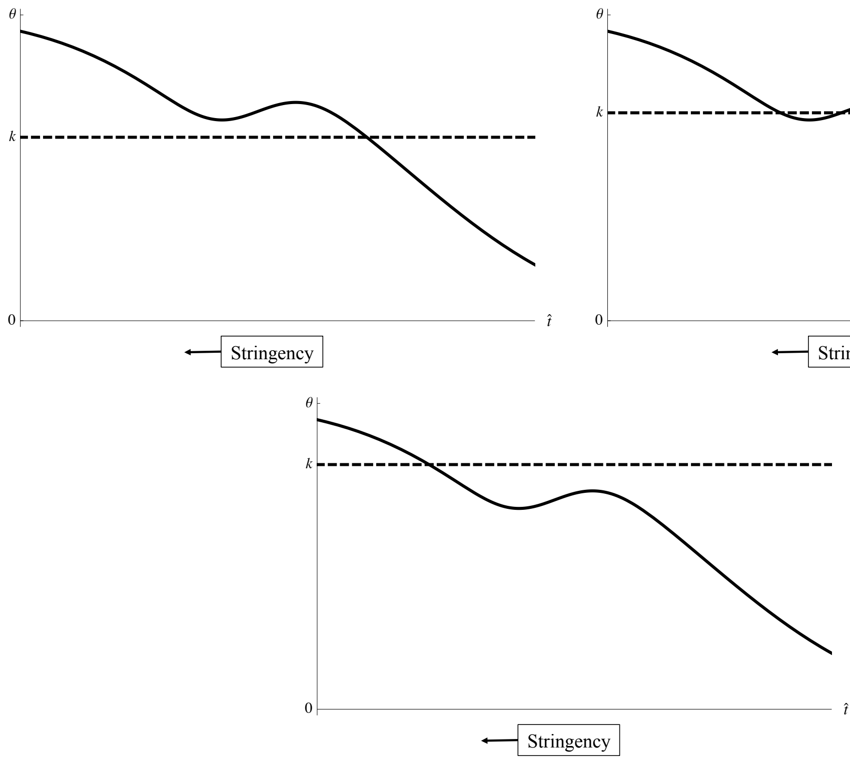

Figure 6: $G^T(\cdot; p)$ is non-monotone when $\sigma_\xi < \sigma_T^2\sqrt{2\pi}$. Except in knife-edge cases, $\Gamma^T(p)$ has one (panels 1 and 3) or three (panel 2) equilibria.

in Lemma 4.3, it has exactly two critical points. Hence, there is still an open set of costs such that there is a unique equilibrium. In the first panel, the costs are low enough that there is a unique equilibrium. In the second panel the costs are moderate, so there are three equilibria. In the third panel, the costs are high enough that there is again a unique equilibrium. (It is straightforward to show that a generic property of this game is that it has one or three finite cutoff equilibria.)

Finally, the fact that, in both limits, G^T is increasing in stringency, rules out the possibility of an infinite cutoff equilibrium in Γ^T.

Theorem 4.2 *The game $\Gamma^T(p)$ does not have an infinite cutoff equilibrium for any $p \in \mathcal{P}$.*

5 Discussion

Both Γ^T and Γ^θ are examples of regime change games whose equilibrium correspondences are locally smooth in a way that facilitates use in applications. Moreover, the analyses of these two games suggests two ways in which such models can be consistent with arguments, like Schelling's (1960) and Weingast's (1997), that rely on equilibrium multiplicity. First, if uncertainty is about θ, then the game (Γ^θ) always has an equilibrium with zero participation and, as long as participation costs are sufficiently low, also has a stable equilibrium with positive participation. Second, if uncertainty is about T, then the game (Γ^T) can still have multiple equilibria, as long as two conditions hold: (i) the prior distribution doesn't put too much weight on extreme realizations of government strength (thin tails, which is to say σ_T^2 not too big) and (ii) the participation costs are neither too high nor too low. (See Figure 6.)

The analyses provided, then, already show that there is no inherent tension between locally smooth equilibria and arguments that rely on multiplicity. The appearance of a tension simply comes from the fact that standard applications of regime change models to mass protest use a model like Γ^T and, further, make a thick tails assumption that generates equilibrium uniqueness. But uniqueness is by no means inherent to such models—multiplicity is restored by changing the locus of uncertainty or by relaxing the thick tails assumption.

Before turning to the issue of the locus of uncertainty, it is worth pausing for a moment to reflect on this point about thick tails. In many settings, one can think of various actors or institutions actually creating public signals about regime strength. Such information might come from opposition parties, the media, electoral outcomes, violent provocateurs, the international community, and so on. The analysis above highlights the fact that, if such sources of endogenous information are sufficiently informative, then multiple equilibria may exist even when the game satisfies two-sided limit dominance. (See Hellwig (2002), Morris and Shin (2003), and Angeletos, Hellwig and Pavan (2006) for related discussions.) Hence, many natural extensions of a regime change model to mass uprisings might endogenously generate equilibrium multiplicity even if the model primitives satisfy the thick tails assumption.

However, as shown in Section 3, even without such extensions, if the locus of uncertainty is about a parameter like θ that does not induce two-sided limit dominance, then the game naturally has multiple equilibria because it always has an equilibrium with no participation, in addition to admitting the possibility of a stable positive participation equilibrium. This raises one remaining question: is a zero-participation equilibrium in fact a natural feature of a regime change model or not? In the remainder of this section, I attempt to address this question by providing some formal intuition for why the game Γ^T never has a zero participation while the game Γ^θ always has one.

Consider the functions G^T and G^θ. For very low levels of stringency of the cutoff rule (i.e., low $\hat{s}$ or high $\hat{t}$) both functions go to values below k. However, for very high levels of stringency of the cutoff rule, G^T goes to $\theta > k$, while G^θ goes to $0 < k$.

This fact drives the difference between these games with respect to the existence of a zero-participation equilibrium. Because G^θ starts negative and ends going to zero, if it crosses $k > 0$ once it will cross it twice (except in the knife-edge case). Thus, multiplicity of finite cutoff equilibria is generic and there is always an infinite cutoff equilbrium with no participation. Because G^T is increasing in stringency everywhere except (at most) on a closed set (see the discussion surrounding Figure 6), it either crosses k once or three times (again, except in knife-edge cases). Thus, there is an open set of parameter values with uniqueness and there is never an infinite cutoff equilibrium.

So the key to seeing why a zero-participation equilibrium always exist in Γ^θ and never exists in Γ^T is understanding why G^θ and G^T behave so differently as the cutoff rule becomes very stringent. To do so, let's compare the effect of increased stringency of the cutoff rule on each of these two functions.

Recall these functions represent the expected incremental benefit to a player whose signal was right at the cutoff rule that all players are using. That is, they represent the expected incremental benefit to the marginal participant. The functions can be rewritten as follows:

$$G^T(\hat{t};p) = \Phi\left(\frac{T^*(\hat{t};p) - \gamma\hat{t} - (1-\gamma)m}{\sigma_\gamma}\right)\theta$$

and

$$G^\theta(\hat{s};r) = \left(1 - \Phi\left(\frac{\theta^*(\hat{s};r) - \lambda\hat{s} - (1-\lambda)m}{\sigma_\lambda}\right)\right)\left(\lambda\hat{s} + (1-\lambda)m + \sigma_\lambda \frac{\phi\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)}{1-\Phi\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)}\right).$$

In order to compare the responses of G^T and G^θ to increases in stringency, I compare $-\frac{dG^T(\hat{t};p)}{d\hat{t}}$ to $\frac{dG^\theta(\hat{s};r)}{d\hat{s}}$, since increased stringency involves decreasing the cutoff rule in Γ^T but increasing the cutoff rule in Γ^θ.

$$-\frac{dG^T(\hat{t};p)}{d\hat{t}} = \phi\left(\frac{T^*(\hat{t};p) - \gamma\hat{t} - (1-\gamma)m}{\sigma_\gamma}\right)\frac{\theta}{\sigma_\gamma}\left(\gamma - \frac{dT^*(\hat{t};p)}{d\hat{t}}\right). \tag{8}$$

$$\frac{dG^{\theta}(\hat{s};r)}{d\hat{s}} = \phi\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)\frac{\left(\lambda\hat{s}+(1-\lambda)m+\sigma_\lambda\frac{\phi\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)}{1-\Phi\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)}\right)}{\sigma_\lambda}\left(\lambda-\frac{d\theta^*(\hat{s};r)}{d\hat{s}}\right)$$
$$+\left(1-\Phi\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)\right)\left(\lambda+\sigma_\lambda\frac{d\frac{\phi}{1-\Phi}}{d\hat{s}}\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)\right). \tag{9}$$

5.1 Three Substantive Effects of Increased Stringency

Consider the effects of increasing the stringency of the cutoff rule in Γ^T (i.e., decreasing $\hat{t}$). Making the cutoff rule more stringent has two competing effects on the function G^T, captured by the term $(\gamma - \frac{dT^*(\hat{t};p)}{d\hat{t}})$ in Equation 8. First, when the cutoff rule is more stringent, a player whose signal equaled the cutoff rule received a better signal and so believes the state is more favorable to regime change. Call this the *beliefs effect* of increased stringency. Second, when the cutoff rule is more stringent, conditional on a state of the world (i.e. a true T), fewer people participate. Hence, when the cutoff rule is more stringent, the true state of the world must be more favorable (i.e., must be lower) in order for regime change to be achieved. Call this the *critical-threshold effect* of increased stringency. The beliefs effect (represented by γ) tends to make G^T increasing and the critical-threshold effect (represented by $-\frac{dT^*(\hat{t};p)}{d\hat{t}}$) tends to make it decreasing.

The function G^θ exhibits the same two effects. (In Equation 9, the beliefs effect is represented by λ and the critical-threshold effect is represented by $-\frac{d\theta^*(\hat{s};r)}{d\hat{s}}$.) In addition, there is a third effect on G^θ, represented by the second line of Equation 9. When the cutoff rule is more stringent, a player whose signal equaled the cutoff rule received a better signal and so believes the payoff from successful regime change is higher. (Notice that $\theta^*(\hat{s};r)-\lambda\hat{s}$ is increasing in $\hat{s}$ and, because the normal distribution has the monotone hazard rate property, $\frac{\phi}{1-\Phi}$ is increasing.) Call this the *expected payoff effect* of increased stringency. This effect doesn't exist for G^T because the state is not about the payoff from success in Γ^T. Since we are trying to understand why G^θ is *decreasing* for stringent enough rules, and the expected payoff effect tends to make G^θ increasing in stringency, we can safely ignore this third effect in trying to understand the differences between G^T and G^θ.[9]

The question, then, is the following: Why, for highly stringent rules, does the beliefs effect dominate in G^T but the critical-threshold effect dominate in G^θ? I develop intuitions

[9] Moreover, substituting from Equation 1 into Equation 9, the expected payoff effect becomes negligible

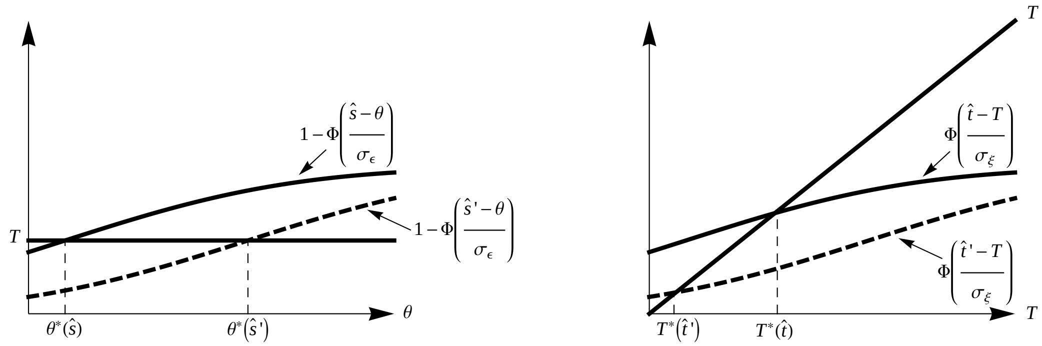

Figure 7: Changing $\hat{t}$ has less of an effect on $T^*(\hat{t})$ than changing $\hat{s}$ has on $\theta * (\hat{s})$.

to answer this by considering the effects one at a time.

5.2 The Beliefs Effect

The beliefs effects in Γ^θ and Γ^T are represented by λ and γ, respectively. This reflects the fact that the more informative is the signal in either game, the larger is the beliefs effect. These magnitudes are unaffected by the stringency of the cutoff rule. It will be important that both λ and γ are strictly less than 1.

5.3 The Critical-Threshold Effect

Recall the definitions of the critical thresholds themselves: $\theta^*(\hat{s}; r)$ is the minimal θ that leads to regime change in $\Gamma^\theta(r)$, given a cutoff rule $\hat{s}$. Similarly, $T^*(\hat{t}; p)$ is the maximal T that leads to regime change in $\Gamma^T(p)$, given a cutoff rule $\hat{t}$. These two thresholds are defined in Equations 1 and 5, respectively, and are represented graphically (for two values of $\hat{s}$ and $\hat{t}$) in Figure 7.

It will be useful to develop intuitions in three steps. First, I will discuss the fact that the critical-threshold effect is larger in the game Γ^θ than in the game Γ^T. Then I will show that the critical-threshold effect is, in fact, so large in Γ^θ that it is always larger than the beliefs effect. This implies that the only reason G^θ is ever increasing is because of the

---

as stringency increases. To see this, note

$$\lim_{\hat{s}\to\infty}\left(1-\Phi\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)\right)\left(\lambda+\sigma_\lambda\frac{d\frac{\phi}{1-\Phi}}{d\hat{s}}\left(\frac{\theta^*(\hat{s};r)-\lambda\hat{s}-(1-\lambda)m}{\sigma_\lambda}\right)\right)=0.$$

See the appendix for a proof of this claim.

expected payoffs effect. Third, I will discuss why, for stringent enough cutoff rules, the critical-threshold effect is in fact smaller than the beliefs effect in the game Γ^T.

5.3.1 The Critical-Threshold Effect is Larger in Γ^θ than in Γ^T

It is clear, from Figure 7, that changing $\hat{t}$ has a smaller impact on T^* than changing $\hat{s}$ has on θ^*. One can see this formally by implicitly differentiating Equations 5 and 1:

$$\frac{dT^*(\hat{t};p)}{d\hat{t}} = \frac{\frac{1}{\sigma_\xi}\phi\left(\frac{\hat{t}-T^*(\hat{t};p)}{\sigma_\xi}\right)}{\frac{1}{\sigma_\xi}\phi\left(\frac{\hat{t}-T^*(\hat{t};p)}{\sigma_\xi}\right)+1} < 1 \tag{10}$$

and

$$\frac{d\theta^*(\hat{s};r)}{d\hat{s}} = \frac{\phi\left(\frac{\hat{s}-\theta^*(\hat{s};r)}{\sigma_\epsilon}\right)}{\phi\left(\frac{\hat{s}-\theta^*(\hat{s};r)}{\sigma_\epsilon}\right)} = 1.$$

Substantively, why is the critical-threshold effect larger in the game Γ^θ than in the game Γ^T?

In both games, when the cutoff rule is made more stringent, participation decreases. Hence, the state of the world must become more favorable in order to achieve regime change. Making the state of the world more favorable in the game Γ^θ means a higher realization of θ. Such a change has only one effect and it is strategic—when the state of the world is better, more people receive a signal that crosses the cutoff rule, hence more people's strategy calls on them to participate. This is why $\frac{d\theta^*}{d\hat{s}} = 1$. There is a one-for-one trade-off, in terms of achieving regime change, between making the cutoff rule more stringent (which decreases participation) and making the state more favorable (which increases participation). However, things are different in the game Γ^T.

Making the state of the world more favorable in the game Γ^T means a lower realization of T. Such a change has two effects—one strategic and one mechanical. The strategic effect is just as in the game Γ^θ. When the state of the world is more favorable, more people receive a good enough signal to cross the cutoff rule, increasing participation. The second effect is mechanical and does not have an analogue in the game Γ^θ. When the state of the world (T) is more favorable, fewer people need to participate in order to achieve regime change. Because of this second effect, there is a less than one-for-one trade-off, in terms of achieving regime change, between making the cutoff rule more stringent (thereby reducing participation) and improving the state of the world (thereby increasing participation *and* making it easier to achieve regime change). That is, for any given incremental increase in the stringency of $\hat{t}$, a decrease in T^* that is of a smaller size than the increase in $\hat{t}$ will continue to assure regime change. Hence, the critical-threshold effect is smaller in the game Γ^T than in the game Γ^θ.

5.3.2 The Critical-Threshold Effect is Larger than the Beliefs Effect in Γ^θ

The fact that there is a one-for-one trade-off between the stringency of the cutoff rule and the critical threshold is crucial for understanding multiplicity in the game Γ^θ. In particular, the fact that $\frac{d\theta^*(\hat{s};r)}{d\hat{s}} = 1$ implies that it is impossible for the beliefs effect (represented by $\lambda = \frac{\sigma_\theta}{\sigma_\theta + \sigma_\epsilon} < 1$) to be greater than the critical-threshold effect. Taking into account only these two effects, increasing the stringency of the cutoff rule, therefore, always makes the player whose signal is at the cutoff rule worse off (because the strategic decrease in participation more than compensates for the increased beliefs about the state of the world sustaining regime change). That is, taken together, in the game Γ^θ, the net of the beliefs effect and the critical-threshold effect is for G^θ to be decreasing in stringency. Hence, the only reason that G^θ is increasing anywhere is because of the expected payoffs effect. But, as shown in Footnote 9, as stringency increases, the expected payoff effect becomes negligible, so eventually G^θ becomes decreasing in stringency.

The above argument, of course, does not hold in the game Γ^T. There, it is possible for $\gamma = \frac{\sigma_T}{\sigma_T + \sigma_\xi}$ to be greater than $\frac{dT^*(\hat{t};p)}{d\hat{t}} < 1$. Indeed, as we will see, for high enough levels of stringency this must be the case.

5.3.3 The Critical-Threshold Effect in Γ^T Becomes Negligible as Stringency Increases

We have seen that the critical-threshold effect need not necessarily be larger than the beliefs effect in the game Γ^T. It remains to be shown that, for sufficiently stringent rules, it indeed is not, and to develop an intuition for why.

Let's start by showing that the critical-threshold effect is indeed smaller than the beliefs effect for sufficiently stringent cutoff rules in the game Γ^T. To see, notice from Equation 5 that for all $p \in \mathcal{P}$,

$$T^*(\hat{t}; p) \in (0, 1)$$

and

$$\lim_{\hat{t} \to -\infty} T^*(\hat{t}; p) = 0.$$

Given this, it is clear from Equation 10, that

$$\lim_{\hat{t}\to-\infty} \frac{dT^*(\hat{t};p)}{d\hat{t}} = 0.$$

In Γ^T, for very high levels of stringency, the critical-threshold effect becomes negligible. Why is this?

First, notice that, regardless of the stringency of the cutoff rule, changing the state of the world always has a mechanical effect on the likelihood of regime change. That is, regardless of the cutoff rule, decreasing T directly makes it easier to achieve regime change since less participation is required. This mechanical effect is represented by the 1 in the denominator of $\frac{dT^*(\hat{t};p)}{d\hat{t}}$ in Equation 10.

The same is not true for the strategic effect. As the cutoff rule becomes very stringent, there is almost no density of population members who received signals near the cutoff rule. Hence, a small change in the cutoff rule has almost no negative effect on participation for very stringent cutoff rules. And, for the exact same reason, for very stringent cutoffs, an improvement in the state of the world has almost no positive effect on participation. In the derivative $\frac{dT^*(\hat{t};p)}{d\hat{t}}$, these facts can be see in the term $\frac{1}{\sigma_\xi}\phi\left(\frac{\hat{t}-T^*(\hat{t};p)}{\sigma_\xi}\right)$. This term represents the measure of marginal participants (i.e., those who would stop participating due to a marginal increase in stringency or who would start participating due to a marginal improvement in the state of the world). It appears in both the numerator and the denominator because the measure of marginal participants has implications for the effect of a change in $\hat{t}$ and for the effect of a change in T. As $\hat{t}$ goes to minus infinity (i.e., as the cutoff rule becomes very stringent), this term clearly goes to zero. That is, for very stringent cutoff rules, there are essentially no marginal participants.

The arguments above show the following. For very stringent cutoff rules, participation is essentially unaffected by a small change in stringency or by a small change in the state of the world. This is because there are essentially no marginal participants when the cutoff rule is very stringent. However, regardless of stringency, improving the state of the world mechanically increases the probability of regime change by lowering the required level of participation. Hence, as the cutoff rule becomes very stringent, an incremental increase in stringency requires essentially no improvement in the state of the world to continue to sustain regime change. And this is why the critical-threshold effect becomes negligible when the cutoff rule becomes very stringent.

Importantly, the fact that $\frac{dT^*(\hat{t};p)}{d\hat{t}}$ goes to zero (i.e., that the critical-threshold effect becomes negligible) is not driven by special features of the normal distribution. The fact that the effect of stringency on participation becomes negligible as the rule become stringent follows from the density of population signals going to zero in its tails. And that, of course, is a feature of any density with full support on the real line, since the density must integrate to one.

5.3.4 Two Intuitive Conditions and their Relationship to Limit Dominance

Taken together, what do these arguments suggest drives the difference between Γ^T and Γ^θ with respect to existence of a zero participation equilibrium? In Γ^T, the uncertainty is over a parameter whose realization has both a strategic and a mechanical effect, whereas in Γ^θ the uncertainty is over a parameter than has only a strategic effect. This fact, as we have seen, implies that it is impossible for the beliefs effect to be larger than the critical-threshold effect in Γ^θ but not so in Γ^T.

This fact is closely related to two-sided limit dominance. In particular, notice that it is precisely because T has a mechanical effect on regime change that it can produce two-sided limit dominance. In a complete information version of Γ^T, if $T < 0$, participation is a dominant strategy because the regime will fall even if only one player participates. This is entirely due to the mechanical effect. When T is negative, mechanically, regime change will occur even if no one participates. Similarly, for $T > 1$, not participating is a dominant strategy. Again, this is entirely due to the mechanical effect. When T is bigger than 1 (which is the total mass of the population) the regime will not fall even if everyone participates.

The game Γ^θ, by way of contrast, does not have two-sided limit dominance even with full support because θ has no mechanical effect on the probability of the regime falling. Because there is no mechanical effect, no matter how high θ is, if a player expects no one else will participate, it is a best response not to participate.

The intuitions developed above, thus, suggest that absence of a zero-participation equilibrium in Γ^T is driven by the assumption that the prior puts positive weight on the possibility that the regime will fall even with zero participation. But, has already been discussed, this assumption sits uncomfortably with the assumption on payoffs—if a single, measure zero person participates when $T < 0$, she derives the benefit θ when the regime falls, even though she was irrelevant to causing that event. Hence, there is an important sense in which the assumptions underlying the non-existence of a zero-participation equilibrium in Γ^T are substantively less well motivated than the assumptions underlying Γ^θ. This suggests another reason why equilibrium multiplicity, and thus arguments like Schelling's (1960) and Weingast's (1997), should remain an important feature of our understanding of mass uprisings and revolution.

6 Conclusion

For half a century, our understanding of the strategic logic of revolutionary threats and mass uprisings has been based, in no small part, on the argument that such settings have multiple equilibria. Yet a recent game theoretic literature makes use of a modeling technology—global games of regime change—that typically yield a unique equilibrium. This raises the question of whether natural models of regime change are inconsistent with standard analyses dating back at least to Schelling (1960).

I argue that such uniqueness is not a natural feature of such models. Most importantly, I provide a substantively plausible model of regime change—in which the citizens are uncertain of the payoff from regime change rather than regime strength—that has the advantages of the global games models (most notably, a locally smooth equilibrium correspondence), yet yields multiple equilibria. Indeed, I argue that this model is more substantively plausible for an application to mass uprisings that is the standard global game of regime change. The standard global game assumes that there are states of the world in which: (i) the regime will fall no matter what and (ii) in such a state of the world, a single individual who mobilizes receives a discontinuous benefit when the regime falls, even though in truth she was irrelevant to this outcome. The model with uncertainty over the payoff to regime change does not make this assumption. And, as I demonstrate, this assumption is essential to the non-existence of a zero-participation equilibrium in the standard global game of regime change.

All told, then, I show there is no tension between arguments about mass uprisings that depend on equilibrium multiplicity and a modeling approach that introduces uncertainty in order to generate smoothness. As such, the analysis provided here highlights the importance of theorists introducing uncertainty into models of regime change in a substantively plausible way that does not, inadvertently, rule out deep arguments about the strategic logic of such phenomena.

Appendix

Notation

The following notation will be useful:

- $\alpha \equiv \frac{(1-\lambda)}{\sigma_\lambda}$
- $\beta \equiv \frac{(1-\lambda)m+\sigma_\epsilon \Phi^{-1}(1-T)}{\sigma_\lambda}$

Now, substituting for $\overline{m}_i$ and $\theta^*(\hat{s};r)$ we have

$$G^\theta(\hat{s};r) \equiv (1-\Phi(\alpha\hat{s}-\beta))\left(\lambda\hat{s}+(1-\lambda)m+\sigma_\lambda\frac{\phi(\alpha\hat{s}-\beta)}{1-\Phi(\alpha\hat{s}-\beta)}\right).$$

Proofs of Numbered Results

Proof of Lemma 3.1.

(i) Substituting for $\theta^*(\hat{s};r)$ and $\overline{m}_i$, and slightly rearranging, the limit can be rewritten as

$$\lim_{s_i\to-\infty}\left(1-\Phi\left(\frac{\hat{s}-\Phi^{-1}(1-T)\,\sigma_\epsilon-\lambda s_i-(1-\lambda)m}{\sigma_\lambda}\right)\right)(\lambda s_i+(1-\lambda)m) + \lim_{s_i\to-\infty}\sigma_\lambda\phi\left(\frac{\hat{s}-\Phi^{-1}(1-T)\,\sigma_\epsilon-\lambda s_i-(1-\lambda)m}{\sigma_\lambda}\right).$$

The second term clearly goes to zero. Thus, all that remains is to show that the first term goes to zero. By simple rearrangement, the first term can be rewritten:

$$\lim_{s_i\to-\infty}\frac{1-\Phi\left(\frac{\hat{s}-\Phi^{-1}(1-T)\sigma_\epsilon-\lambda s_i-(1-\lambda)m}{\sigma_\lambda}\right)}{\frac{1}{\lambda s_i+(1-\lambda)m}}.$$

Using l'Hopital's rule and the definition of the normal PDF, this equals:

$$-\lim_{s_i\to-\infty}\frac{(\lambda s_i+(1-\lambda)m)^2}{e^{\frac{(\hat{s}-\Phi^{-1}(1-T)\sigma_\epsilon-\lambda s_i-(1-\lambda)m)^2}{2}}\sigma_\lambda\sqrt{2\pi}}.$$

Again using l'Hopital's rule, this equals:

$$-\lim_{s_i\to-\infty}\frac{-2(\lambda s_i+(1-\lambda)m)}{e^{\frac{(\hat{s}-\Phi^{-1}(1-T)\sigma_\epsilon-\lambda s_i-(1-\lambda)m)^2}{2}}(\hat{s}-\Phi^{-1}(1-T)\,\sigma_\epsilon-\lambda s_i-(1-\lambda)m)\,\sigma_\lambda\sqrt{2\pi}}.$$

Again using l'Hopital's rule this equals:

$$-\lim_{s_i\to-\infty}\frac{2}{\sigma_\lambda\sqrt{2\pi}e^{\frac{(\hat{s}-\Phi^{-1}(1-T)\sigma_\epsilon-\lambda s_i-(1-\lambda)m)^2}{2}}\left(1+(\hat{s}-\Phi^{-1}(1-T)\,\sigma_\epsilon-\lambda s_i-(1-\lambda)m)^2\right)}.$$

Now the numerator is constant and the denominator goes to infinity, establishing the result.

(ii) Substituting for $\theta^*(\hat{s};r)$ and $\overline{m}_i$, and slightly rearranging, the limit can be rewritten as

$$\lim_{s_i\to\infty}\left(1-\Phi\left(\frac{\hat{s}-\Phi^{-1}\left(1-T\right)\sigma_\epsilon-\lambda s_i-(1-\lambda)m}{\sigma_\lambda}\right)\right)(\lambda s_i+(1-\lambda)m)$$
$$+\lim_{s_i\to\infty}\sigma_\lambda\phi\left(\frac{\hat{s}-\Phi^{-1}\left(1-T\right)\sigma_\epsilon-\lambda s_i-(1-\lambda)m}{\sigma_\lambda}\right).$$

The first term clearly goes to infinity and the second term clearly goes to zero.

(iii) Differentiating with respect to s_i, we have that at a critical point the following first-order condition must hold:

$$\frac{dg(s_i^*,\hat{s};r)}{ds_i}=\phi\left(\frac{\hat{s}-\Phi^{-1}\left(1-T\right)\sigma_\epsilon-\lambda s_i^*-(1-\lambda)m}{\sigma_\lambda}\right)\frac{\lambda}{\sigma_\lambda}(\lambda s_i^*+(1-\lambda)m)$$
$$+\left(1-\Phi\left(\frac{\hat{s}-\Phi^{-1}\left(1-T\right)\sigma_\epsilon-\lambda s_i^*-(1-\lambda)m}{\sigma_\lambda}\right)\right)\lambda$$
$$-\sigma_\lambda\phi'\left(\frac{\hat{s}-\Phi^{-1}\left(1-T\right)\sigma_\epsilon-\lambda s_i^*-(1-\lambda)m}{\sigma_\lambda}\right)\left(\frac{\lambda}{\sigma_\lambda}\right)=0.$$

For notational convenience, let $f(s_i)=\frac{\hat{s}-\Phi^{-1}(1-T)\sigma_\epsilon-\lambda s_i-(1-\lambda)m}{\sigma_\lambda}$. Rearranging, the first order condition holds if and only if:

$$\frac{1-\Phi(f(s_i^*))}{\phi(f(s_i^*))}-\frac{\phi'(f(s_i^*))}{\phi(f(s_i^*))}=-\frac{\lambda s_i^*+(1-\lambda)m}{\sigma_\lambda}.$$

Using the fact that $\phi'(x)=-x\phi(x)$ (note that the chain rule has already been applied), this can again be rewritten:

$$\frac{1-\Phi(f(s_i^*))}{\phi(f(s_i^*))}+f(s_i^*)=-\frac{\lambda s_i^*+(1-\lambda)m}{\sigma_\lambda}.$$

Substituting for $f(s_i^*)$ and rearranging, this can be rewritten:

$$\frac{1-\Phi\left(\frac{\hat{s}-\Phi^{-1}(1-T)\sigma_\epsilon-\lambda s_i^*-(1-\lambda)m}{\sigma_\lambda}\right)}{\phi\left(\frac{\hat{s}-\Phi^{-1}(1-T)\sigma_\epsilon-\lambda s_i^*-(1-\lambda)m}{\sigma_\lambda}\right)}+\frac{\hat{s}-\Phi^{-1}\left(1-T\right)\sigma_\epsilon-(1-\lambda)m}{\sigma_\lambda}=-\frac{(1-\lambda)m}{\sigma_\lambda}.$$

Since the normal distribution has the monotone hazard rate property, the left-hand side is increasing in s_i^* and the right-hand side is constant. Thus, $g(\cdot, \hat{s}; r)$ can have at most one critical point.

Given that $g(\cdot, \hat{s}; r)$ has at most one critical point, it follows from the first two points of this lemma that, if it has a critical point, it is a minimum and that $g(s_i^*, \hat{s}; r) < 0$. Hence, $g(\cdot, \hat{s}; r)$ is increasing everywhere to the right of s_i^* and, since $\lim_{s_i \to \infty} g(s_i, \hat{s}; r) = \infty$, it eventually crosses k.

■

Proof of Lemma 3.2. Necessity follows from the argument in the text. For sufficiency, consider a profile where all players employ such a cutoff rule. Consider a player with type $s_i < \hat{s}$. Lemma 3.1 establishes that $g(s_i, \hat{s}; r) < k$ for all such players, so they have no profitable deviation to participating. Similarly, consider a player with type $s_i > \hat{s}$. Lemma 3.1 establishes that $g(s_i, \hat{s}; r) > k$ for all such players, so they have no profitable deviation to not participating. ■

Proof of Lemma 3.3.

(i) $G^\theta(\hat{s}; r)$ can be rewritten $(1 - \Phi(\alpha\hat{s} - \beta))(\lambda\hat{s} + (1-\lambda)m) + \sigma_\lambda \phi(\alpha\hat{s} - \beta)$. Given this, we can write

$$\lim_{\hat{s}\to\infty} G^\theta(\hat{s}; r) = \lim_{\hat{s}\to\infty} \frac{(1 - \Phi(\alpha\hat{s} - \beta))}{\frac{1}{\lambda\hat{s} + (1-\lambda)m}} + \lim_{\hat{s}\to\infty} \phi(\alpha\hat{s} - \beta).$$

It is straightforward that the second term equals 0. Thus, consider the first term in isolation:

$$\begin{aligned}
\lim_{\hat{s}\to\infty} \frac{(1 - \Phi(\alpha\hat{s} - \beta))}{\frac{1}{\lambda\hat{s} + (1-\lambda)m}} &= \lim_{\hat{s}\to\infty} \frac{\alpha(\lambda\hat{s} + (1-\lambda)m)^2}{e^{\frac{(\alpha\hat{s}-\beta)^2}{2}} \lambda\sqrt{2\pi}} \\
&= \lim_{\hat{s}\to\infty} \frac{2\alpha\lambda(\lambda\hat{s} + (1-\lambda)m)}{(\alpha\hat{s} - \beta)\alpha e^{\frac{(\alpha\hat{s}-\beta)^2}{2}} \lambda\sqrt{2\pi}} \\
&= \lim_{\hat{s}\to\infty} \frac{2\lambda}{\alpha e^{\frac{(\alpha\hat{s}-\beta)^2}{2}} \sqrt{2\pi} + (\alpha\hat{s} - \beta)^2 \alpha e^{\frac{(\alpha\hat{s}-\beta)^2}{2}} \sqrt{2\pi}} \\
&= 0,
\end{aligned}$$

where, in order, the equalities follow from (1) l'Hopital's rule and the definition of the PDF of the standard normal, (2) l'Hopital's rule, (3) l'Hopital's rule, and (4) the observation that the numerator of the limit is a positive constant in s and the denominator of the limit goes to infinity. Hence, the whole limit goes to 0.

(ii) Using the same rewriting as the previous point,

$$\lim_{\hat{s}\to-\infty} G^{\theta}(\hat{s};r) = \lim_{\hat{s}\to-\infty}(1-\Phi(\alpha\hat{s}-\beta))(\lambda\hat{s}+(1-\lambda)m) + \lim_{\hat{s}\to-\infty}\phi(\alpha\hat{s}-\beta).$$

Again, it is straightforward that the second term equals 0. The first term equals $-\infty$, since $1-\Phi(\alpha\hat{s}-\beta)$ clearly goes to 1 and $\lambda\hat{s}+(1-\lambda)m$ goes to $-\infty$.

(iii) Given the first two points of this lemma, establishing the following two steps suffices:

(a) There exists a $\hat{s}$ such that $G^{\theta}(\hat{s};r) > 0$.

(b) $G^{\theta}(\hat{s};r)$ has at most one critical point.

The first step will establish that $G^{\theta}(\hat{s};r)$ has a maximum. The second point will establish that it has no minima. Taken together, these establish single peakedness. I take them in order.

(a) Consider a $\hat{s} > -\frac{(1-\lambda)m}{\lambda}$. Recall, we can write

$$G^{\theta}(\hat{s};r) = (1-\Phi(\alpha\hat{s}-\beta))(\lambda\hat{s}+(1-\lambda)m) + \sigma_{\lambda}\phi(\alpha\hat{s}-\beta).$$

The first term is positive since $(1-\Phi(\alpha\hat{s}-\beta)) > 0$ and $\hat{s} > -\frac{(1-\lambda)m}{\lambda}$. The second term is positive for all $\hat{s}$. Hence, for $\hat{s} > -\frac{(1-\lambda)m}{\lambda}$, $G^{\theta}(\hat{s};r) > 0$.

(b) Differentiating, we have that at a critical point the following first order condition holds:

$$\frac{dG^{\theta}(\hat{s}^{*};r)}{d\hat{s}} = -\phi\left(\alpha\hat{s}^{*}-\beta\right)\alpha(\lambda\hat{s}^{*}+(1-\lambda)m)+\left(1-\Phi\left(\alpha\hat{s}^{*}-\beta\right)\right)\lambda+\sigma_{\lambda}\alpha\phi'\left(\alpha\hat{s}^{*}-\beta\right) = 0.$$

Rearranging, this holds if and only if

$$\frac{\lambda\left(1-\Phi\left(\alpha\hat{s}^{*}-\beta\right)\right)}{\alpha\phi\left(\alpha\hat{s}^{*}-\beta\right)} + \frac{\sigma_{\lambda}\alpha\phi'\left(\alpha\hat{s}^{*}-\beta\right)}{\alpha\phi\left(\alpha\hat{s}^{*}-\beta\right)} = \lambda\hat{s}^{*}+(1-\lambda)m.$$

Using the fact that $\phi'(x) = -x\phi(x)$ and canceling, this can be rewritten

$$\frac{\lambda\left(1-\Phi\left(\alpha\hat{s}^{*}-\beta\right)\right)}{\alpha\phi\left(\alpha\hat{s}^{*}-\beta\right)} - \sigma_{\lambda}\left(\alpha\hat{s}^{*}-\beta\right) = \lambda\hat{s}^{*}+(1-\lambda)m.$$

Since the normal distribution has the monotone hazard rate property, $\frac{1-\Phi(\alpha\hat{s}^*-\beta)}{\phi(\alpha\hat{s}^*-\beta)}$ is decreasing in $\hat{s}^*$. Thus, the entire left-hand side is decreasing in $\hat{s}^*$ while the right-hand side is increasing in $\hat{s}^*$, so there can be at most one $\hat{s}^*$ satisfying the first-order conditions.

■

Proof of Theorem 3.1.

I begin with the first claim. To see that, for all $r \in \mathcal{R}$, there is a Bayesian Equilibrium with no participation, consider a strategy profile with $a_i = 0$ for all s_i. The probability of regime change is zero. If a player were to deviate to participating, the probability of regime change would still be zero, since all individuals are measure zero. Thus, the payoff to deviating is $-k < 0$.

Now turn to finite cutoff equilibria.

Definition 6.1 *Let* $s^*(r) = \arg\max_s G^\theta(s; r)$.

Lemma 3.3 establishes that $s^*(r)$ is unique. Further, it is clear that $s^*(r)$ is constant in k, so $G^\theta(s^*(r_{-k}, k); (r_{-k}, k))$ is constant in k.

First consider the case of no finite cutoff equilibria. Fix an r_{-k}. There are two cases:

(i) First, suppose that, for all $k \geq 0$, $G^\theta(s^*(r); (r_{-k}, k)) < k$. Then, by Lemma 3.2 there are no finite cutoff equilibria, establishing the existence of an open set $O_-(r_{-k})$.

(ii) Next, suppose there exists a finite $\overline{k}(r_{-k})$ such that $G^\theta(s^*(r); (r_{-k}, \overline{k}(r_{-k}))) = \overline{k}(r_{-k})$. Then, since $G^\theta(s^*(r); (r_{-k}, \overline{k}(r_{-k})))$ is constant in k, there are no finite cutoff equilibria for any $k < \overline{k}(r_{-k})$ establishing the existence of an open set $O_-(r_{-k})$.

Next consider the case of at least one finite cutoff equilibria. By hypothesis, there exists a finite $\overline{k}(r_{-k})$ such that $G^\theta(s^*(r); (r_{-k}, \overline{k}(r_{-k}))) = \overline{k}(r_{-k})$. (Otherwise no finite cutoff equilibrium would exists for $r = (r_{-k}, k)$). Since G^θ is continuous on $\mathcal{R}$, single peaked, and satisfies $\lim_{\hat{s}\to-\infty} G^\theta(\hat{s}; r) = -\infty$ and $\lim_{\hat{s}\to\infty} G^\theta(\hat{s}; r) = 0$, it follows that, for $\hat{s} \in (-\infty, \infty)$, $G^\theta(\hat{s}; r)$ takes all values in $(0, G^\theta(s^*(r); r))$ twice and takes the value $G^\theta(s^*(r); r)$ exactly once. This implies that for any $k > \overline{k}(r_{-k})$, $G^\theta(\hat{s}; r)$ crosses k twice, establishing the existence of an open set $O_+(r_{-k})$ with two finite cutoff equilibria. There is a unique finite cutoff equilibrium if and only if $G^\theta(s^*(r); (r_{-k}, k)) = k$. Since $G^\theta(s^*(r); (r_{-k}, k))$ is constant in k, there is exactly one such k. ■

Proof of Theorem 3.2.

A finite cutoff rule $\hat{s}$ is a Bayesian Equilibrium of $\Gamma^{\theta}(r)$ if and only if it satisfies $G^{\theta}(\hat{s};r) = k$ (Lemma 3.2). Thus, if there is a finite cutoff rule that is an equilibrium, it must be that $G^{\theta}(s^*(r);r) \geq k$.

Since G^{θ} is continuous on $\mathcal{R}$, single peaked, and satisfies $\lim_{\hat{s}\to-\infty} G^{\theta}(\hat{s};r) = -\infty$ and $\lim_{\hat{s}\to\infty} G^{\theta}(\hat{s};r) = 0$, it follows that, for $\hat{s} \in (-\infty,\infty)$, $G^{\theta}(\hat{s};r)$ takes all values in $(0, G^{\theta}(s^*(r);r))$ twice.

If $G^{\theta}(s^*(r);r) < k$, then there are no equilibria in finite cutoff strategies. To see that this is possible, fix all other parameter values and let $k \to \infty$.

If $G^{\theta}(s^*(r);r) \geq k$, then there are equilibria in finite cutoff strategies. To see that this is possible, fix parameter values such that $G^{\theta}(s^*(r);r) > 0$ and let $k \to 0$. Moreover, if $G^{\theta}(s^*(r);r) > k$, then $G^{\theta}(\hat{s};r)$ takes the value k twice for $\hat{s} \in (-\infty,\infty)$. Each instance where $G^{\theta}(\hat{s};r) = k$ is a finite cutoff equilibrium, by Lemma 3.2.

All that remains is to show that the set $\{r \in \mathcal{R} | G^{\theta}(s^*(r);r) = k\}$ is non-generic. Label this set D. I make use of the following claim.

Claim 6.1 *$G^{\theta}(s^*(r);r) - k$ is continuous on $\mathcal{R}$ in each element of r, and strictly monotone on $\mathcal{R}$ in k.*

Given the claim it is straightforward to establish non-genericity by showing that the set D is closed with empty interior.

The set D is closed because it is the continuous pre-image of the closed set $\{0\}$.

The set D has empty interior. To see this, notice that if the interior were non-empty, then D would contain some point $r' \in \mathcal{R}$ and some ϵ-ball around r'. That ball would contain two points with the 5^{th} element (i.e., the value of k) strictly ordered by the standard order on the real line, contradicting monotonicity in k.

Since D is closed and has empty interior, it is meagre.

All that remains is to prove the claim.

Proof of Claim. Rewrite $G^{\theta}(s^*(r);r) - k$ as:

$$(1 - \Phi(\alpha s^*(r) - \beta))(\lambda s^*(r) + (1-\lambda)m) + \sigma_{\lambda}\phi(\alpha s^*(r) - \beta) - k.$$

Continuity is immediate from the Theorem of the Maximum (Mas-Collel, Whiston and Green (1995), Theorem M.K.6). Monotonicity in k follows from the fact that α, β, and $s^*(r)$ are constant in k, so $G^{\theta}(s^*(r);r) - k$ is monotonically decreasing in k. ■ ■

Proof of Lemma 4.1. Necessity follows from the argument in the text. For sufficiency, consider a player with type $t_i < \hat{t}$. Such a player's payoff from participating is strictly positive, so there is no profitable deviation to not participating. Consider a player with type $t_i > \hat{t}$. Such a player's payoff to participation is strictly negative, so there is no profitable deviation to participating. A player with type $\hat{t}$ is indifferent by construction. ■

Proof of Lemma 4.2. I will make use of the following two facts: $\lim_{x\to 1} \Phi^{-1}(x) = \infty$ and $\lim_{x\to 0} \Phi^{-1}(x) = -\infty$.

(i) It is immediate from Equation 5 that for any $p \in \mathcal{P}$, $\lim_{\hat{t}\to\infty} T^*(\hat{t};p) = 1$. Using this fact and the fact that $\lim_{x\to 1} \Phi^{-1}(x) = \infty$, we have that for any $p \in \mathcal{P}$,

$$\lim_{\hat{t}\to\infty} G^T(\hat{t};p) = \lim_{\hat{t}\to\infty} \Phi\left(\frac{\sigma_\xi}{\sigma_T\sqrt{\sigma_T^2+\sigma_\xi^2}}(T^*(\hat{t};p)-m) - \frac{\sigma_T}{\sqrt{\sigma_T^2+\sigma_\xi^2}}\Phi^{-1}(T^*(\hat{t};p))\right)\theta = 0.$$

(ii) It is immediate from Equation 5 that for any $p \in \mathcal{P}$, $\lim_{\hat{t}\to-\infty} T^*(\hat{t};p) = 0$. Using this fact and the fact that $\lim_{x\to 0} \Phi^{-1}(x) = -\infty$, we have that for any $p \in \mathcal{P}$,

$$\lim_{\hat{t}\to-\infty} G^T(\hat{t};p) = \lim_{\hat{t}\to-\infty} \Phi\left(\frac{\sigma_\xi}{\sigma_T\sqrt{\sigma_T^2+\sigma_\xi^2}}(T^*(\hat{t};p)-m) - \frac{\sigma_T}{\sqrt{\sigma_T^2+\sigma_\xi^2}}\Phi^{-1}(T^*(\hat{t};p))\right)\theta = \theta.$$

(iii) The third point follows from the fact that, for any $p \in \mathcal{P}$ and any finite t, $\Phi\left(\frac{\sigma_\xi}{\sigma_T\sqrt{\sigma_T^2+\sigma_\xi^2}}(T^*(\hat{t};p)-m) - \frac{\sigma_T}{\sqrt{\sigma_T^2+\sigma_\xi^2}}\Phi^{-1}(T^*(\hat{t};p))\right)$ is strictly between 0 and 1.

■

Proof of Proposition 4.1. The proof follows Angeletos, Hellwig and Pavan (2007).

Differentiating $G^T(\hat{t};p)$ with respect to $\hat{t}$ yields

$$\frac{dG^T(t;p)}{d\hat{t}} = \phi\left(\frac{(1-\gamma)(T^*(\hat{t})-m) - \gamma\Phi^{-1}(T^*(\hat{t}))\sigma_\xi}{\sigma_\gamma}\right)\frac{\theta}{\sigma_\gamma}\frac{dT^*}{d\hat{t}}\left[(1-\gamma) - \frac{\gamma\sigma_\xi}{\phi(\Phi^{-1}(T^*(\hat{t})))}\right].$$

Implicitly differentiating Equation 5 shows that $\frac{dT^*(\hat{t};p)}{d\hat{t}} = \frac{\phi\left(t - \frac{T^*(\hat{t};p)}{\sigma_\xi}\right)}{\phi\left(\hat{t} - \frac{T^*(\hat{t};p)}{\sigma_\xi}\right) + \sigma_\xi} > 0$, so the derivative of $G^T(\cdot;p)$ is negative if and only if $(1-\gamma) - \frac{\gamma\sigma_\xi}{\phi(\Phi^{-1}(T^*(\hat{t};p)))} < 0$. Substituting for γ and using the fact that $\min_x \frac{1}{\phi(\Phi^{-1}(x))} = \sqrt{2\pi}$, the right-hand side of Equation 7 is monotonically decreasing if $\sigma_\xi < \sigma_T^2\sqrt{2\pi}$. When this condition holds, Lemma 4.1 implies that the game has at most one cutoff rule consistent with equilibrium and points 1 and 2 of Lemma 4.2 establish that it has at least one, since $\theta > k$ for all $p \in \mathcal{P}$. ■

Proof of Lemma 4.3. From the proof of Proposition 4.1, if $p \in \mathcal{P}^*$, then $G^T(\cdot;p)$ is decreasing on its entire domain, in which case the set $G^T(\cdot;p)$ has no critical points.

I argue that the only other alternative is $G^T(\cdot;p)$ having exactly two critical points, and that there exist parameter values such that this is true, in five steps.

(i) **There are parameter values for which $G^T(\cdot;p)$ is increasing on at least part of the domain of its first argument.** Recall from the proof of Proposition 4.1 that the derivative of $G^T(\cdot;p)$ is negative if and only if $(1-\gamma) - \frac{\gamma\sigma_\xi}{\phi(\Phi^{-1}(T^*(\hat{t};p)))} < 0$, which, for a fixed $\hat{t}$, clearly does not hold for small enough values of γ.

(ii) **For all parameter values, G^T is decreasing on at least part of the domain of its first argument.** For all p, $G^T(\hat{t};p)$ converges to 0 as $\hat{t} \to \infty$ and to θ as $\hat{t} \to -\infty$ (by points 1 and 2 of Lemma 4.2). Since, for all $p \in \mathcal{P}$, $\theta > 0$, $G^T(\cdot;p)$ is decreasing on part of the domain of its first argument.

(iii) **There exist parameter values where the set G^T has at least one critical value.** Follows directly from points 1 and 2.

(iv) **If $G^T(\cdot;p)$ has any critical points it has at least two**. By Lemma 4.2 point 1, for all $p \in \mathcal{P}$, $G^T(\hat{t};p)$ converges θ as $\hat{t} \to -\infty$. By Lemma 4.2 point 3, for all $p \in \mathcal{P}$ and all $\hat{t} \in \mathbb{R}$, $G^T(\hat{t};p) \in (0,\theta)$. This implies that as $T \to -\infty$, $G^T(\hat{t},p)$ converges to its upper bound. Thus, if $G^T(\hat{t};P)$ has any critical points, the smallest one must be a local minimizer.

   By Lemma 4.2 point 2, for all $p \in \mathcal{P}$, $G^T(\hat{t};p)$ converges 0 as $\hat{t} \to \infty$. By Lemma 4.2 point 3, for all $p \in \mathcal{P}$ and all $\hat{t} \in \mathbb{R}$, $G^T(t;p) \in (0,\theta)$. This implies that as $T \to \infty$, $G^T(\hat{t};p)$ converges to its lower bound. Thus, if $G^T(\hat{t};p)$ has any critical points, the largest one must be a local maximizer.

   Taken together, these two points imply the claim.

(v) **For all p, $G^T(\cdot;p)$ has at most two critical points.** From the proof of Proposition

4.1, $\frac{dG^T(\hat{t};p)}{d\hat{t}} = 0$ if and only if:

$$\phi\left(\Phi^{-1}\left(T^*\left(\hat{t};p\right)\right)\right) = \frac{\gamma\sigma_\xi}{1-\gamma}.$$

The right-hand side is constant in $\hat{t}$. From the proof of Proposition 4.1, $T^*(\hat{t};p)$ is increasing in $\hat{t}$ for all $p \in \mathcal{P}$ and clearly Φ^{-1} is strictly increasing. Thus, for any $p \in \mathcal{P}$, $\phi\left(\Phi^{-1}\left(T^*\left(\hat{t};p\right)\right)\right)$ can equal any given constant at most twice.

Take together, points 3–5 complete the proof. ■

Proof of Theorem 4.1. Fix a $p \in \mathcal{P}$. First notice that there can be no infinite cutoff equilibria. For $\hat{t} \to \infty$, the equilibrium would call for all players to participate. But since $\lim_{\hat{t}\to\infty} G^T(\hat{t};p) = -k$, players with signals that are arbitrarily large have payoffs strictly less than 0, so they will not participate. Similarly, for $\hat{t} \to -\infty$, the equilibrium would call for all players not to participate. But since $\lim_{\hat{t}\to-\infty} G^T(\hat{t};p) = \theta$, players with signals that are arbitrarily low have payoffs strictly higher than k, so they will participate.

Now turn to finite cutoff strategies.

Fix a $p_{-k} \in \mathcal{P}_{-k}$. Notice that $\frac{dG^T(\hat{t};(p_{-k},k))}{d\hat{t}} = \frac{dG^T(\hat{t};(p_{-k},k'))}{d\hat{t}}$ for all $k, k' \in \mathbb{R}_+$. Thus, there are only two cases to consider.

First, fix a $p_{-k} \in \mathcal{P}_{-k}$ such that the set $\left\{\hat{t}|\frac{dG^T(\hat{t};(p_{-k},k))}{d\hat{t}} = 0\right\} = \emptyset$. By Lemma 4.2, $\lim_{\hat{t}\to\infty} G^T(\hat{t};(p_{-k},k)) = 0$ and $\lim_{\hat{t}\to-\infty} G^T(\hat{t};(p_{-k},k)) = \theta$. This implies that, for all k, $G^T(t;(p_{-k},k))$ is monotonically decreasing in its first argument. Thus, for any $k \in (0,\theta)$, $G^T(\hat{t};(p_{-k},k))$ crosses k exactly once. By Lemma 4.1, this is the only finite cutoff equilibrium.

Next, fix a $p_{-k} \in \mathcal{P}_{-k}$ such that the set $\left\{\hat{t}|\frac{dG^T(\hat{t};(p_{-k},k))}{d\hat{t}} = 0\right\}$ is non-empty. By Lemma 4.3 the function $G^T(\cdot;(p_{-k},k))$ has exactly two critical points. Define them as follows:

$$\bar{t}(p_{-k},k) = \arg \max_{\hat{t} \in \left\{\hat{t}|\frac{dG^T(\hat{t};(p_{-k},k))}{d\hat{t}}=0\right\}} G^T(\hat{t};(p_{-k},k))$$

and

$$\underline{t}(p_{-k},k) = \arg \min_{\hat{t} \in \left\{\hat{t}|\frac{dG^T(\hat{t};(p_{-k},k))}{d\hat{t}}=0\right\}} G^T(\hat{t};(p_{-k},k)).$$

By the third bullet point of Lemma 4.2, both $G^T(\bar{t}(p_{-k},k);(p_{-k},k))$ and $G^T(\underline{t}(p_{-k},k);(p_{-k},k))$ are contained in $(0,\theta)$. Notice, further, that $\bar{t}$ and $\underline{t}$ are not a function of k.

Given that $G^T(\hat{t};(p_{-k},k))$ converges to 0 as $\hat{t} \to \infty$ and to θ as $\hat{t} \to -\infty$, it must be that $\underline{t}(p_{-k},k) < \bar{t}(p_{-k},k)$ and that $G^T(\hat{t};(p_{-k},k))$ is decreasing for $\hat{t} \in (-\infty, \underline{t}(p_{-k},k)) \cup (\bar{t}(p_{-k},k),\infty)$. This implies that, for $\hat{t} \in \mathbb{R}$, $G^T(\hat{t};(p_{-k},k))$ takes values in $(G^T(\underline{t}(p_{-k},k);(p_{-k},k))$, $G^T(\bar{t}(p_{-k},k);(p_{-k},k)))$ exactly three times, takes values in $\{G^T(\underline{t}(p_{-k},k);(p_{-k},k)),G^T(\bar{t}(p_{-k},k);(p_{-k},k))\}$ exactly twice, and takes values in $(0,(G^T(\underline{t}(p_{-k},k);(p_{-k},k))) \cup (G^T(\bar{t}(p_{-k},k);(p_{-k},k)),\theta)$ exactly once. Thus, for $k \in (0,(G^T(\underline{t}(p_{-k},k);(p_{-k},k))) \cup (G^T(\bar{t}(p_{-k},k);(p_{-k},k)),\theta)$ there is a unique cutoff equilibrium. To see that this set is not empty, notice that $\lim_{k\to 0} G^T(\underline{t}(p_{-k},k);(p_{-k},k)) > k$ and that for all k, $G^T(\underline{t}(p_{-k},k);(p_{-k},k)) < G^T(\bar{t}(p_{-k},k);(p_{-k},k)) < \theta$. Since the union of open sets is an open set, this establishes the result. ■

Proof of Theorem 4.2. Consider a cutoff rule $\hat{t}$. Write the expected payoff of participating to a player of type t_i, given the cutoff rule $\hat{t}$, as:

$$U(a_i = 1;(t_i,\hat{t})) = \Phi\left(\frac{T^*(\hat{t};p) - \gamma t_i - (1-\gamma)m}{\sigma_\gamma}\right)\theta - k.$$

Clearly, $U(a_i = 1;(t_i,\hat{t}))$ is continuous and decreasing in t_i.

Now, consider the possibility of a negative, infinite cutoff equilibrium (i.e., one with no participation). From Lemma 4.2 point 2 and Lemma 4.3, for any $p \in \mathcal{P}$, there is a finite $\underline{t}$ such that $G^T(\cdot;p)$ is monotonically decreasing in $\hat{t}$ for any $\hat{t} < \underline{t}$. Moreover, as $\hat{t} \to -\infty$, $G^T(\hat{t};p) \to \theta > k$. Thus, there is a finite $\underline{\underline{t}}$ such that, for any $\hat{t} < \underline{\underline{t}}$, $G^T(\hat{t};p) > k$. Consider any cutoff rule $\hat{t} < \underline{\underline{t}}$. We have that $U(a_i = 1;(\hat{t},\hat{t})) = G^T(\hat{t};p) > k$. Then, by continuity, there is a $t_i > \hat{t}$ such that $U(a_i = 1;(t_i,\hat{t})) > 0$, so a type t_i player has an incentive to participate. Hence, the cutoff rule $\hat{t}$ was not an equilibrium. This implies that no $\hat{t} < \underline{\underline{t}}$ is an equilibrium, so there is no negative infinite cutoff equilibrium.

Next consider the possibility of a positive infinite cutoff equilibrium (i.e., one with full participation). From Lemma 4.2 point 1 and Lemma 4.3, for any $p \in \mathcal{P}$, there is a finite $\bar{t}$ such that $G^T(\hat{t};p)$ is monotonically decreasing in $\hat{t}$ for any $\hat{t} > \underline{t}$. Moreover, as $\hat{t} \to \infty$, $G^T(\hat{t};p) \to 0 < k$. Thus, there is a finite $\bar{\bar{t}}$ such that, for any $\hat{t} > \bar{\bar{t}}$, $G^T(\hat{t};p) < k$. Consider any cutoff rule $\hat{t} > \bar{\bar{t}}$. We have that $U(a_i = 1;(\hat{t},\hat{t})) = G^T(\hat{t};p) - k < 0$. Then, by continuity, there is a $t_i < \hat{t}$ such that $U(a_i = 1;(t_i,\hat{t})) < 0$, so a player of type t_i does not want to participate. Hence, the cutoff rule $\hat{t}$ was not an equilibrium. This implies that no $\hat{t} > \underline{\underline{t}}$ is an equilibrium, so there is no positive infinite cutoff equilibrium. ■

Proof of Claim in Footnote 9. The limit in the statement of the proposition can be rewritten:

$$\lim_{\hat{s}\to\infty}\left(1-\Phi\left(\alpha\hat{s}-\beta\right)\right)\left(\lambda+\sigma_\lambda\frac{d\frac{\phi}{1-\Phi}}{d\hat{s}}\left(\alpha\hat{s}-\beta\right)\right).$$

Further, note that $\frac{d\frac{\phi}{1-\Phi}}{d\hat{s}}(\alpha\hat{s}-\beta)=\frac{\phi'(\alpha\hat{s}-\beta)\alpha}{1-\Phi(\alpha\hat{s}-\beta)}+\frac{\phi(\alpha\hat{s}-\beta)^2\alpha}{(1-\Phi(\alpha\hat{s}-\beta))^2}$, so we can again rewrite the limit as follows:

$$\lim_{\hat{s}\to\infty}\left(1-\Phi\left(\alpha\hat{s}-\beta\right)\right)\lambda+\lim_{\hat{s}\to\infty}\sigma_\lambda\alpha\phi'\left(\alpha\hat{s}-\beta\right)+\lim_{\hat{s}\to\infty}\sigma_\lambda\frac{\phi\left(\alpha\hat{s}-\beta\right)^2}{1-\Phi\left(\alpha\hat{s}-\beta\right)}.$$

The first term clearly goes to zero.

Next consider the second term. Using the definition of the normal pdf and differentiating, it can be rewritten:

$$-\frac{\sigma_\lambda\alpha^2}{\sqrt{2\pi}}\lim_{\hat{s}\to\infty}\frac{\alpha\hat{s}-\beta}{e^{\frac{(\alpha\hat{s}-\beta)^2}{2}}}.$$

By l'Hopital's rule, this can be rewritten

$$-\frac{\sigma_\lambda\alpha^2}{\sqrt{2\pi}}\lim_{\hat{s}\to\infty}\frac{\alpha}{e^{\frac{(\alpha\hat{s}-\beta)^2}{2}}(\alpha\hat{s}-\beta)\alpha}=0.$$

Finally, consider the third term. By l'Hopitals rule, it can be rewritten:

$$\sigma_\lambda\lim_{\hat{s}\to\infty}\frac{2\phi\left(\alpha\hat{s}-\beta\right)\phi'\left(\alpha\hat{s}-\beta\right)\alpha}{-\phi\left(\alpha\hat{s}-\beta\right)\alpha}.$$

Rearranging, this can be rewritten

$$-2\sigma_\lambda\lim_{\hat{s}\to\infty}\phi'\left(\alpha\hat{s}-\beta\right).$$

Now an argument identical to that for the second term shows this third term goes to zero.

■

References

Angeletos, George-Marios, Christian Hellwig and Alessandro Pavan. 2006. "Signaling in a Global Game: Coordination and Policy Traps." *Journal of Political Economy* 114(3):452–484.

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.

Baliga, Sandeep and Tomas Sjöström. 2004. "Arms Races and Negotiations." *Review of Economic Studies* 71:351–369.

Baliga, Sandeep and Tomas Sjöström. 2012. "The Strategy of Manipulating Conflict." *American Economic Review* 102(6):2897–2922.

Boix, Carles and Milan W Svolik. 2013. "The Foundations of Limited Authoritarian Government: Institutions, Commitment, and Power-Sharing in Dictatorships." *Journal of Politics* 75(2):300–316.

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

Carlsson, Hans and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61(5):989–1018.

Chan, Kenneth S. and Y. Stephen Chiu. 2002. "The role of (non-)transparency in a currency crisis model." *European Economic Review* 46(2):397–416.

Chassang, Sylvain and Gerard Padro i Miguel. 2010. "Conflict and Deterrence Under Strategic Risk." *Quarterly Journal of Economics* 125(4):1821–1858.

Corsetti, Giancarlo, Bernardo Guimaraes and Nouriel Roubini. 2006. "International lending of last resort and moral hazard: A model of IMF's catalytic finance." *Journal of Monetary Economics* 53(3):441–471.

Dewan, Torun and David P. Myatt. 2007. "Leading the Party: Coordination, Direction and Communication." *American Political Science Review* 101(4):825–843.

Dewan, Torun and David P. Myatt. 2008. "The Qualities of Leadership: Direction, Communication, and Obfuscation." *American Political Science Review* 102(3):351–368.

Echenique, Federico and Aaron Edlin. 2004. "Mixed equilibria are unstable in games of strategic complements." *Journal of Economic Theory* 118(1):61–79.

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

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(04):645–668.

Fearon, James D. 2011. "Self Enforcing Democracy." *The Quarterly Journal of Economics* 126(4):1661–1708.

Frankel, David M., Stephen Morris and Ady Pauzner. 2003. "Equilibrium Selection in Global Games with Strategic Complementarities." *Journal of Economic Theory* 108(1):1–44.

Goldstein, Itay and Ady Pauzner. 2005. "Demand Deposit Contracts and the Probability of Bank Runs." *Journal of Finance* 60(3):1293–1327.

Greene, William H. 2003. *Econometric Analysis, Fifth Edition.* Prentice Hall.

Guimaraes, Bernardo and Stephen Morris. 2007. "Risk and Wealth in a Model of Self-Fulfilling Currency Attacks." *Journal of Monetary Economics* 54:2205–2230.

Hardin, Russell. 1996. *One for All: The Logic of Group Conflict.* Princeton, New Jersey: Princeton University Press.

Hellwig, Christian. 2002. "Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games." *Journal of Economic Theory* 107:191–222.

Kuran, Timor. 1989. "Sparks and Prairie Fires: A Theory of Unanticipates Political Revolution." *Public Choice* 61(1):41–74.

Little, Andrew T. 2012. "Elections, fraud, and election monitoring in the shadow of revolution." *Quarterly Journal of Political Science* 7(3):249–283.

Mas-Collel, Andreu, Michael D. Whiston and Jerry R. Green. 1995. *Microeconomic Theory.* Oxford: Oxford University Press.

Meirowitz, Adam and Joshua A Tucker. 2013. "People Power or a One-Shot Deal? A Dynamic Model of Protest." *American Journal of Political Science* 57(2):478–490.

Morris, Stephen and Hyun Song Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88(3):587–597.

Morris, Stephen and Hyun Song Shin. 2000. "Rethinking Multiple Equilibria in Macroeconomic Modeling." *NBER Macroeconomics Annual* 15:139–161.

Morris, Stephen and Hyun Song Shin. 2003. Global Games: Theory and Applications. In *Advances in Economics and Econometrics: Theory and Applications, 8th World Congress of the Econometric Society*, ed. Mathias Dewatripont, Lars Peter Hansen and Stephen J. Turnovsky. Cambridge, UK: Cambridge University Press.

Morris, Stephen and Hyun Song Shin. 2004. "Coordination Risk and the Price of Debt." *European Economic Review* 48:133–153.

Persson, Torsten and Guido Tabellini. 2009. "Democratic Capital: The Nexus of Political and Economic Change." *American Economic Journal: Macroeconomics* 1(2):88–126.

Rochet, Jean-Charles and Xavier Vives. 2004. "Coordination Failures and the Lender of Last Resort: Was Bagehot Right After All?" *Journal of the European Economic Association* 2(6):1116–1147.

Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Shadmehr, Mehdi and Dan Bernhardt. 2011*a*. "Collective Action with Uncertain Payoffs: Coordination, Public Signals and Punishment Dilemmas." *American Political Science Review* 105(4):829–51.

Shadmehr, Mehdi and Dan Bernhardt. 2011*b*. "Coordination Games with Information Aggregation.". Paper Presented at the 69th Annual Meetings of the Midwest Political Science Association.

Smith, Alastair and Scott Tyson. 2014. "Two-Sided Coordination: Competition in Collective Actions." Paper Presented at the 72nd Annual Meetings of the Midwest Political Science Association.

Weingast, Barry R. 1997. "Political Foundations of Democracy and the Rule of Law." *American Political Science Review* 91:245–263.