Appendix A

Proof of Proposition 2

For part i, writing out the probability of protest and simplifying gives:

$$Pr(\overline{\mu}(s_\omega) \geq k/\overline{\beta}) = \Phi\left((\mu_\omega - k/\overline{\beta})\frac{\alpha_0}{\rho\alpha_s\sqrt{\rho^2\alpha_0^{-1} + \alpha_s^{-1}}}\right).$$

Since $\frac{\alpha_0}{\rho\alpha_s\sqrt{\rho^2\alpha_0^{-1}+\alpha_s^{-1}}}$ is decreasing in ρ, this probability is increasing in ρ if and only if $\mu_\omega < k/\overline{\beta}$.

For part ii, rewriting the protest condition substituting $\mu_\omega + \nu_\omega$ for ω:

$$\lambda\mu_\omega + (1-\lambda)(\mu_\omega + \nu_\omega + \epsilon_\omega/\rho) \geq k/\overline{\beta}$$
$$\epsilon_\omega \geq \frac{\rho}{1-\lambda}(k/\overline{\beta} - \mu_\omega) - \rho\nu_\omega \equiv \overline{\epsilon}_\omega.$$

So, the probability of protest is increasing in ρ when $\overline{\epsilon}_\omega$ is increasing in ρ. Taking the appropriate derivative:

$$\frac{\partial\overline{\epsilon}_\omega}{\partial\rho} = \left(1 - \frac{\alpha_0}{\rho^2\alpha_s}\right)(k/\overline{\beta} - \mu_\omega) - \nu_\omega.$$

Rearranging this gives the desired result. ∎

Results with More General Distributions of θ and ω.

The results relating the tactical decision hold as long as the signal conveys *any* information about θ in the following sense. Suppose θ and s_t are drawn from a joint distribution, and again write the conditional mean of θ given s_t as $\overline{\mu}_\theta$. Then the optimal tactical choice is still $\overline{\mu}_\theta$, and the expect cost when choosing this tactic is:

$$\mathbb{E}[k(\overline{\mu}_\theta - \theta)^2] = k\mathbb{V}_\theta[\theta|s_t]$$

If $\mathbb{V}_\theta[\theta|s_t]$ is less than $\mathbb{V}_\theta[\theta]$ for all s_t (as with the normal distribution), then there is always more protest with a tactical signal than without a tactical signal. However, this property is not true for all distributions. For example, suppose θ is binary, taking on value 0 or 1. If the prior places a very high likelihood on 0 (or 1) and the signal indicates that 1 is more likely than in the prior, the variance of the posterior distribution will go up.

A property that always holds is that the *average* cost of protest over the potential signals always decreases upon observing s_t, which is just a consequence of the law of total variance:

$$k\mathbb{E}_{s_t}[\mathbb{V}_\theta[\theta|s_t]] = k(\mathbb{V}_\theta[\theta] - \mathbb{V}_{s_t}[\mathbb{E}_\theta[\theta|s_t]])$$

When $\mathbb{V}_{s_t}[\mathbb{E}_\theta[\theta|s_t]]$ is high – which means the signal has a large effect on the posterior belief of θ – then in general the expected cost of protest will be low. So, if we label $\mathbb{V}_{s_t}[\mathbb{E}_\theta[\theta|s_t]]$ the informativeness of the signal, then more informative signals will (on average) leads to a lower expected cost of protest.

The results regarding the signal of the regime's popularity are also much more general. For an prior on ω and noise term ϵ_ω, the law of iterated expectations implies that:

$$\mathbb{E}[\mathbb{E}[\omega|s_\omega]] = \mathbb{E}[\omega]$$

that is, the average posterior belief about ω must be ω. So, for there to be realizations of s_ω that increase the belief about ω, there must be realizations that decrease ω. Another property that always holds is that as $\rho \to 0$, the citizen learns nothing about ω from s_ω, and hence protests if and only if $\mu_\omega > k/\overline{\beta}$. So, there will always be a region of the parameter space where protest is more likely for $\rho > 0$ than for $\rho = 0$ and always a region where protest is less likely for $\rho > 0$ than for $\rho = 0$.

Derivation of Equilibrium Condition in Section 2

I search for symmetric Perfect Bayesian Equilibria that are monotone in the individual level of dis-satisfaction with the regime. This means that (1) all citizens use the same strategy, (2) citizens protest if and only if they sufficiently dislike the regime, and (3) citizens' beliefs about the regime strength and level of protest are consistent with Bayes' rule and the other citizens' strategy profile. Since $\overline{s}$ is a sufficient statistic for the public signals, I also restrict attention to equilibria where citizens only condition their behavior on this average rather than the individual signals. The *ex ante* distribution of $\overline{s}$ is normal with mean ω and precision $n\alpha_\omega \equiv \alpha_s$, i.e., the more publicly observed signals, the more precise the information gleaned from the signals.

A citizen with regime sentiment's belief about ω is normal with mean:

$$\overline{\overline{\mu}}_\omega(\omega_i) = \frac{\alpha_0\mu_\omega + \alpha_s\overline{s} + \alpha_\omega\omega_i}{\alpha_0 + \alpha_s + \alpha_\omega}$$

and precision $\alpha_0+\alpha_s+\alpha_\omega$. The only way the public signals enter this term (and hence the expected payoff to protest for a fixed strategy) is through the $\overline{s}$ term, justifying the focus on strategies that only condition on the average public signal.

From the perspective of a citizen with ω_i, the probability of citizen j protesting is:

$$Pr(\omega_j > \hat{\omega}(\overline{s})) = Pr(\omega + \nu_j > \hat{\omega}(\overline{s})) = \Phi\left(\tilde{\alpha}^{1/2}\left(\alpha_s(\omega_i) - \hat{\omega}(\overline{s})\right)\right),$$

where $\tilde{\alpha} = \frac{(\alpha_0+\alpha_s+\alpha_\omega)\alpha_\omega}{\alpha_0+\alpha_s+2\alpha_\omega}$. The equilibrium condition is that the marginal citizen is indifferent between protesting and not, or:

$$\hat{\omega}(\overline{s}) + v_A\Phi\left(\tilde{\alpha}^{1/2}\frac{\alpha_0+\alpha_s}{\alpha_0+\alpha_s+\alpha_\omega}\left(\overline{\mu}_\omega - \hat{\omega}(\overline{s})\right)\right) = k/\overline{\beta},$$

where $\overline{\mu}_\omega = \frac{\alpha_0\mu_\omega+\alpha_s\overline{s}}{\alpha_0+\alpha_s}$ is the public mean belief about ω.

Proof of Proposition 4

The *ex ante* expected size of protest is given by

$$\mathbb{E}[A] = \int_{\overline{s}=-\infty}^{\infty} Pr(\omega_i > \hat{\omega}(\overline{s})|\overline{s})\phi(\overline{s}; \mu_\omega, \alpha_s)d\overline{s},$$

where $\phi(\cdot; \mu, \alpha)$ is the probability density function of a random variable with mean μ and precision α. This is decreasing in $\hat{\omega}(\overline{s})$. Implicitly differentiating equation 6 implies the level of protest is decreasing in $k/\overline{\beta}$, proving parts i-ii.

The level of protest is a complicated function of α_s, but extensive numerical analysis indicates that a result similar to proposition 2 applies, i.e., the level of protest is increasing in α_s when protest is *ex ante* unlikely and decreasing in α_s otherwise.

Tactical Coordination Model

The actors in the model are a continuum mass 1. For this model, θ represents the average preferred tactic. At the outset, citizens share a common prior on θ that this is normally distributed with mean normalized to 0 and precision β_0.

Citizens have an individual preferred tactic x_i, given by:

$$x_i = \theta + \nu_i,$$

where ν_i is normally distributed with mean 0 and precision β_x.

Citizens also observe a public signal of θ given by:

$$s_\theta = \theta + \epsilon_s,$$

where ϵ_s is normally distributed with mean 0 and precision β_s.

All of the primitive random variables: ϵ_0, ϵ_s, ν_i and θ are independent.

Let $a_i \in \{0,1\}$ be the protest decision for citizen i and $t_i \in \mathbb{R}$ be the tactic chosen by citizen i if protesting. The citizen payoff is:

$$u(a_i, t_i; x_i) = a_i \left[\omega + v_A A - g(R) - k[r(t_i - x_i)^2 + (1-r)(t_i - \theta)^2]\right]. \tag{9}$$

As before, if a citizen does not protest ($a_i = 0$), she gets a payoff of 0. The A term is the proportion of citizens that protest and R is the range of tactics chosen by the protesters. Assume $v_A \geq 0$ and g is a continuous and increasing function with $g(0) = 0$ and $\lim_{R\to\infty} g(R) = \overline{g}$ for some finite $\overline{g} > 0$. A high v_A means the protest size has a large impact on the benefit of joining, and the g function captures how much the effectiveness of protest decreases as participants choose a less cohesive set of tactics. For the tactical coordination model, assume $\omega \in \mathbb{R}$ is common knowledge.[22]

The k term scales the general cost of protesting, which is comprised of two components. First, the $(t_i - x_i)^2$ term means the cost of protesting with tactic t_i is increasing in the distance between citizen i's chosen and preferred tactic. Second, the $(t_i - \theta)^2$ term indicates the cost is increasing in the distance between citizen i's chosen tactic and the average preferred tactic. A direct interpretation of this term is that θ represents an objectively optimal tactic, and citizens want to choose effective tactics. The term also has the effect of inducing citizens to choose a tactic closer to what they think others will do rather than simply picking their ideal tactic, and hence indirectly captures the incentive to coordinate with others. The $r \in (0,1)$ parameter scales the relative importance of the two components of the cost.

The coordination incentive would more directly be captured by a term increasing in the distance between citizen i's chosen tactic and the average *chosen* tactic of others as in related models (e.g., Morris and Shin, 2002; Dewan and Myatt, 2008) where, to use the terminology here, all citizens participate in the protest by assumption. Unfortunately this formulation greatly complicates the analysis when participation is endogenous, since the expectation of the preferred tactic of par-

[22]Or equivalently, that the citizens share a common prior on the regime's unpopularity with mean ω.

ticipants is generally non-linear in x_i unless all citizens participate. As a result, there is no linear equilibrium. The payoff structure used here rescues this linearity, making the model substantially more tractable.

After observing x_i and s_θ, the citizens simultaneously choose whether or not to protest, followed by the protesters simultaneously choosing a tactic. I solve for Perfect Bayesian equilibria that are symmetric in the sense that all citizens use the same strategy, and in a stronger sense formalized below.

Regardless of the strategies of others, the optimal protest tactic for citizen i (if joining) is a weighted average of her preferred tactic and her expectation about the average preferred tactic of others: $t_i^* = rx_i + (1-r)\overline{\overline{\mu}}_i$, where $\overline{\overline{\mu}}_i = \mathbb{E}[\theta|s, x_i]$. So, the expected cost of protest when choosing the optimal tactic is:

$$
\begin{aligned}
&\mathbb{E}[k(r(t_i^* - x_i)^2) + (1-r)(t_i^* - \theta)^2] \\
&\quad = k(r(x_i - (rx_i + (1-r)\overline{\overline{\mu}}_i))^2 + (1-r)(\overline{\overline{\mu}}_i - rx_i + (1-r)\overline{\overline{\mu}}_i)^2 + (1-r)\mathbb{E}[(\theta - \overline{\overline{\mu}}_i)^2]) \\
&\quad = k(r(1-r)^2(x_i - \overline{\overline{\mu}}_i) + (1-r)r^2(x_i - \overline{\overline{\mu}}_i) + (1-r)(\beta_0 + \beta_s + \beta_x)^{-1}) \\
&= k(1-r)\left[(\beta_0 + \beta_s + \beta_x)^{-1} + r\left(\frac{\beta_0 + \beta_s}{\beta_0 + \beta_s + \beta_x}\right)^2 d_i^2\right] \equiv K(d_i) \qquad (10)
\end{aligned}
$$

where $d_i \equiv |x_i - \overline{\mu}_\theta|$ and $\overline{\mu}_\theta \equiv \frac{\beta_s s_\theta}{\beta_0 + \beta_s}$ is the *public mean belief* about the average preferred tactic conditional on the public signal s_θ but not the individual preferred tactic. So, d_i is a measure of how unusual or extreme citizen i believes her preferred tactic to be. The expected cost of protest when choosing the optimal tactic is increasing in d_i. Motivated by this symmetry, I restrict attention to equilibria of the form "protest if and only if $d_i < \hat{d}$" for some $\hat{d} \geq 0$.

In an equilibrium of this form, the effect of increasing the amount of public information about the preferred tactics on the protest size is determined by how increasing β_s affects the cost to the *marginal citizen*, i.e., the citizen observing exactly $d_i = \hat{d}$. The expected cost of protest for the marginal citizen has the following properties:

Proposition 6. *The cost of protest to the marginal citizen is decreasing in β_s if and only if the expected size of protest is less than $2\Phi(1/\sqrt{2r}) - 1$.*

Proof Differentiating equation 10:

$$\frac{\partial K}{\partial \beta_s} = k(1-r)\left[-(\beta_0+\beta_s+\beta_x)^{-2} + 2r\frac{\beta_x}{(\beta_0+\beta_s+\beta_x)^2}\left(\frac{\beta_0+\beta_s}{\beta_0+\beta_s+\beta_x}\right)d_i^2\right],$$

which is positive for the marginal citizen (i.e., $d_i = \hat{d}$) if and only if:

$$\hat{d} > \sqrt{\frac{(\beta_0+\beta_s+\beta_x)}{2r\beta_x(\beta_0+\beta_s)}} \equiv \hat{d}^*.$$

Given $\hat{d}$, the probability a given citizen protests is:

$$Pr(d_i < \hat{d}) = Pr(-\hat{d} < x_i - \overline{\mu}_\theta < \hat{d}),$$

where

$$x_i - \overline{\mu}_\theta = \theta + \nu_i - \frac{\beta_s}{\beta_0+\beta_s}(\theta+\epsilon_s) = \nu_i + \frac{\beta_0}{\beta_0+\beta_s}\theta - \frac{\beta_s}{\beta_0+\beta_s}\epsilon_s,$$

which is normally distributed with mean 0 and precision. $\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_s+\beta_x}$. So the cost of protest is increasing in β_s if and only if the the expected size of protest is above:

$$\Phi\left(\sqrt{\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_s+\beta_x}}\hat{d}^*\right) - \Phi\left(-\sqrt{\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_s+\beta_x}}\hat{d}^*\right)$$
$$= 2\Phi\left(\sqrt{\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_s+\beta_x}}\sqrt{\frac{\beta_0+\beta_s+\beta_x}{2r\beta_x(\beta_0+\beta_s)}}\right) - 1 = 2\Phi(1/\sqrt{2r}) - 1 \quad \blacksquare$$

As $r \to 0$, this critical level approaches 1, meaning the cost to the marginal citizen is always decreasing in β_s. The critical threshold is decreasing in r, but even as $r \to 1$ it only goes down to $2\Phi(1/\sqrt{(}2)) - 1 \approx 0.52$. Further, as shown below, the value of protest is increasing in β_s through other channels, so the payoff to protest for the marginal citizen may be increasing in β_s even if the expected size of protest is above this threshold.

Completing the derivation of the equilibrium requires determining the marginal citizen's belief about the protest size (A) and the range of chosen tactics (R). Recall a citizen observing x_i's belief about θ is normal with mean $\overline{\overline{\mu}}(x_i)$ and precision $\beta_0 + \beta_x + \beta_s$, so the expected protest level for such a citizen is given by:

$$\begin{aligned}
\mathbb{E}[A|x_i;\hat{d}] &= \Phi\left(\beta_A^{1/2}(\overline{\mu} + \hat{d} - \overline{\overline{\mu}}(x_i))\right) - \Phi\left(\beta_A^{1/2}(\overline{\mu} - \hat{d} - \overline{\overline{\mu}}(x_i))\right) \\
&= \Phi\left(\beta_A^{1/2}\left(\frac{\beta_x(\overline{\mu} - x_i)}{\beta_0 + \beta_x + \beta_s} + \hat{d}\right)\right) - \Phi\left(\beta_A^{1/2}\left(\frac{\beta_x(\overline{\mu} - x_i)}{\beta_0 + \beta_x + \beta_s} - \hat{d}\right)\right),
\end{aligned}$$

where $\beta_A = \frac{\beta_x(\beta_0+\beta_s+\beta_x)}{\beta_0+\beta_s+2\beta_x}$. By the symmetry of the normal distribution, this can be written as a function of d_i, and for a citizen observing exactly $d_i = \hat{d}$ the expected protest size is:

$$\begin{aligned}
\mathbb{E}[A|d_i = \hat{d};\hat{d}] &= \Phi\left(\beta_A^{1/2}\left(\hat{d}\frac{\beta_0 + \beta_s + 2\beta_x}{\beta_0 + \beta_x + \beta_s}\right)\right) - \Phi\left(\beta_A^{1/2}\left(-\hat{d}\frac{\beta_0 + \beta_s}{\beta_0 + \beta_x + \beta_s}\right)\right) \\
&= \Phi\left(\hat{d}\sqrt{\frac{\beta_x(\beta_0 + \beta_s + 2\beta_x)}{\beta_0 + \beta_s + \beta_x}}\right) - \Phi\left(-\hat{d}\sqrt{\frac{\beta_x(\beta_0 + \beta_s)^2}{(\beta_0 + \beta_s + \beta_x)(\beta_0 + \beta_s + 2\beta_x)}}\right).
\end{aligned}$$

Both terms are increasing in $\hat{d}$, so the expected protest from the perspective of the marginal citizen is increasing in $\hat{d}$.

For the range term, it is common knowledge that citizens between $x_i \in (\overline{\mu} - \hat{d}, \overline{\mu} + \hat{d})$ will protest and common knowledge what tactic they will select. So, the lowest tactic chosen by a protester is:

$$\underline{t} = r(\overline{\mu} - \hat{d}) + (1 - r)\frac{\beta_s s + \beta_x(\overline{\mu} - \hat{d})}{\beta_0 + \beta_s + \beta_x},$$

and the highest tactic $\bar{t}$ is the same as above but with $\overline{\mu}+\hat{d}$ replacing the $\overline{\mu}-\hat{d}$ terms, so the range is given by:

$$R=\bar{t}-\underline{t}=2\left(r+\frac{(1-r)\beta_x}{\beta_0+\beta_s+\beta_x}\right)\hat{d}.$$

This expression is increasing in $\hat{d}$ and decreasing in β_s (for a fixed $\hat{d}$). Interestingly, protest participation is a strategic *substitute* through this channel, as more participants means a wider range of tactics, rendering protest less appealing. (See Myatt (2015) for an example where strategic substitutability can have counterintuitive effects of protest behavior.)

Summarizing, an equilibrium $\hat{d}$ solves:

$$U_1^*(\hat{d};\hat{d})\equiv\omega+b_A\mathbb{E}[A|d_i=\hat{d};\hat{d}]-g\left(2\left(r+\frac{(1-r)\beta_x}{\beta_0+\beta_s+\beta_x}\right)\hat{d}\right)-K(\hat{d})=0. \tag{11}$$

As $\hat{d}\to 0$, the $\mathbb{E}[A|d_i=\hat{d};\hat{d}]$ and $g(R)$ terms drop out, and $K(\hat{d})$ approaches $\frac{(1-r)}{\beta_0+\beta_s+\beta_y}$. As $\hat{d}\to\infty$, the $K(\hat{d})$ term approaches $-\infty$ while the other terms are finite. So, if $\omega>\frac{(1-r)}{\beta_0+\beta_s+\beta_y}$ there must be at least one finite $\hat{d}$ that meets this equilibrium condition. Under some additional restrictions the equilibrium is unique:

Lemma 7. *i. If $\omega>\frac{(1-r)}{\beta_0+\beta_s+\beta_y}$, then there is no equilibrium with no protest and at least one $\hat{d}>0$ meeting equation 11.*
ii. If b_A and $-\underline{g'}$ are sufficiently small, this intersection is unique.

Proof Part i is demonstrated above. If the condition in part ii is met, then equation 11 is always decreasing and hence is equal to zero for a unique $\hat{d}$. ∎

If the first condition does not hold there is always an equilibrium with no protest, though in this case (and in general) there maybe multiple intersections and hence multiple equilibria. When this is the case I restrict attention to the equilibrium with the highest level of protest.

Taking comparative statics on the size of protest:

Proposition 8. *The expected size of protest is:*

i) decreasing in k,

ii) increasing in ω, and

iii) increasing in β_s if $\mathbb{E}[A|\hat{d}] < 2\Phi\left(1/\sqrt{2r}\right)$ (and sometimes increasing even if this does not hold)

Proof Recall the expected size of protest is

$$\mathbb{E}[A|\hat{d}] = 2\Phi\left(\sqrt{\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_s+\beta_x}}\hat{d}\right) - 1$$

For part i, the sign of $\frac{\partial \mathbb{E}[A|\hat{d}]}{k}$ is equal to the sign of $\frac{\partial \hat{d}}{\partial k}$. Implicitly differentiating gives:

$$\frac{\partial \hat{d}}{\partial k} = \frac{-\frac{\partial U_1^*}{\partial k}}{\frac{\partial U_1^*}{\partial \hat{d}}}.$$

The denominator is positive and the denominator is negative in the largest (or unique) $\hat{d}$ such that $U_1^*(d) = 0$, so the expression is negative. Part ii follows from an analogous calculation.

For part iii:

$$\frac{\partial \mathbb{E}[A|\hat{d}]}{\partial \beta_s} = 2\phi\left(\sqrt{\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_x+\beta_s}}\hat{d}\right)\left[\hat{d}\frac{1}{2}\left(\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_x+\beta_s}\right)^{-1/2}\left(\frac{\beta_x}{\beta_0+\beta_s+\beta_x}\right)^2 + \sqrt{\frac{\beta_x(\beta_0+\beta_s)}{\beta_0+\beta_x+\beta_s}}\frac{\partial \hat{d}}{\partial \beta_s}\right].$$

All but the $\frac{\partial \hat{d}}{\partial \beta_s}$ terms are guaranteed to be positive. This term is given by:

$$\frac{\partial \hat{d}}{\partial \beta_s} = \frac{-\frac{\partial U_1^*}{\partial \beta_s}}{\frac{\partial U_1^*}{\partial \hat{d}}}.$$

Again the denominator is negative at an equilibrium $\hat{d}$. The numerator is given by the sign of

$\frac{\partial \mathbb{E}[A|\hat{d},d_i=\hat{d}]}{\partial \beta_s}$. It is useful to write the expected size of protest as:

$$\mathbb{E}[A|\hat{d}, d_i = \hat{d}] = \Phi\left(\underbrace{\hat{d}\beta_A^{1/2}\left(1+\frac{\beta_x}{\beta_0+\beta_s+\beta_x}\right)}_{I}\right) + \Phi\left(\underbrace{\hat{d}\beta_A^{1/2}\left(1-\frac{\beta_x}{\beta_0+\beta_s+\beta_x}\right)}_{II}\right) - 1.$$

So, this derivative can be written:

$$\begin{aligned}\frac{\partial \mathbb{E}[A|\hat{d}, d_i = \hat{d}]}{\partial \beta_s} =& \phi(I)\hat{d}\left[\frac{\partial\sqrt{\beta_A}}{\partial \beta_s}\left(1+\frac{\beta_x}{\beta_0+\beta_s+\beta_x}\right)+\sqrt{\beta_A}\frac{-\beta_x}{(\beta_0+\beta_s+\beta_x)^2}\right] \\ &+ \phi(II)\hat{d}\left[\frac{\partial\sqrt{\beta_A}}{\partial \beta_s}\left(1-\frac{\beta_x}{\beta_0+\beta_s+\beta_x}\right)+\sqrt{\beta_A}\frac{\beta_x}{(\beta_0+\beta_s+\beta_x)^2}\right].\end{aligned}$$

The only negative term is the second in the upper square bracket, but since $0 < II < I$, $\phi(II) > \phi(I)$, the corresponding second term in the bottom square bracket is larger, and hence the expression is positive, and hence $\frac{\partial \hat{d}}{\partial \beta_s} > 0$.

Next, $R(\hat{d})$ is decreasing in β_s and hence $g(R(\hat{d}))$ is increasing in β_s. As shown above, $K(\hat{d})$ is increasing in β_s if if $\mathbb{E}[A|\hat{d}] < 2\Phi\left(1/\sqrt{2r}\right)$, so this is a sufficient but not necessary condition for the protest level to be increasing in β_s.

Figure A1 illustrates this result. As the general cost of protest (k) increases, the expected size of protest drops and eventually reaches zero. This makes protest more cohesive – measured by a linearly decreasing function of tactics chosen[23] – as only those with typical preferred tactics participate. Conversely, as the cost of protest decreases, protests become larger and less cohesive.

[23] In particular, the cohesion is defined as $1 - .5R$ where R is the range of chosen tactics, which is selected to make the size and cohesion on a comparable scale. A substantively similar picture would result from any decreasing function.

Figure A1: Effect of increasing the cost of participation (left panel) and precision of public information about tacts (right panel) on the size (black curve) and cohesion (grey curve) of protests.

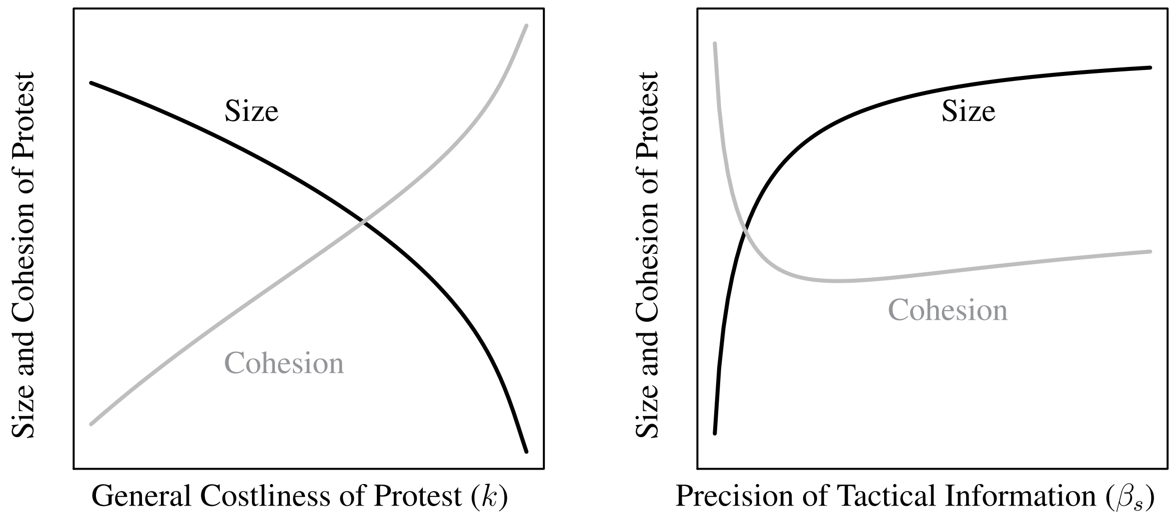

Biased Signals

With the bias term, the distribution of s_θ given θ is now normal with mean $\theta + m_\theta$ and precision:

$$\frac{\beta_s \gamma_\theta}{\beta_s + \gamma_\theta} \equiv \beta'_s.$$

By standard Bayesian updating, her belief about θ is normal with mean:

$$\overline{\mu}'_\theta = (s_\theta - m_\theta)\frac{\beta'_s}{\beta'_s + \beta_0}$$

and precision:

$$\overline{\beta}' = \beta_0 + \beta'_s.$$

When choosing tactic $\overline{\mu}_\theta$, the expected cost of protest is $k/\overline{\beta}'$.

Similarly, rearranging the formula for s_ω gives:

$$(s_\omega - m_\omega)/\rho = \omega + (\epsilon_\omega + \nu_b)/\rho,$$

where $\nu_b = b_\omega - m_\omega$, i.e., the "error" in the estimation of the bias. Since $(\epsilon_\omega + \nu_b)/\rho$ is a random variable with mean 0 and precision:

$$\rho^{-2}\frac{\alpha_\omega \gamma_\omega}{\alpha_\omega + \gamma_\omega} \equiv \alpha'.$$

So, $(s_\omega - m_\omega)/\rho$ is an unbiased signal of ω the belief about ω upon observing s_ω is now normal with mean:

$$\overline{\mu}'_\omega = \lambda' \mu_\omega + (1-\lambda')(s_\omega - m_\omega)/\rho,$$

where $\lambda' = \frac{\alpha_0}{\alpha_0 + \alpha'}$. So, the probability of protest is now $Pr(\overline{\mu}'_\omega > k/\overline{\beta}')$, and the *ex ante* distribution of $\overline{\mu}'_\omega$ is normal with mean μ_ω and precision $(1-\lambda)^{-2}\alpha'$.

To prove proposition 5, note $k/\overline{\beta}'$ is decreasing in β_0, β_s, and γ_θ, proving part i. The probability of protest can be written as $\Phi\left((1-\lambda)(\alpha')^{1/2}(k/\overline{\beta}' - \mu_\omega)\right)$, which is increasing in α' if and only if $k/\overline{\beta}' < \mu_\omega$, and α' is increasing in γ_ω and ρ, proving part ii. The probability of protest is not a function of m_θ and m_ω, proving part iii.

Endogenous Bias

Next, consider the case where the bias terms in s_ω is endogenously chosen by the incumbent regime (or some other actor), denoted I. To ease interpretation, suppose a higher level of bias leads to a more pro-regime (i.e., lower) s_ω:

$$s_\omega = \rho\omega - b_\omega + \epsilon_\omega.$$

Let the incumbent payoff be:

$$u_I = -af(t, \theta) - c_\omega(b_\omega),$$

where f is strictly positive and c_ω is increasing and convex in b_ω. That is, when the citizen protests, the incumbent payoff is reduced by $f(t, \theta)$, and c_ω represents the cost of manipulating the signal of its unpopularity.

As is typical in a "career concerns" style model, the main condition for a pure strategy equilibrium is that if the citizen expects the incumbent to choose a bias level b^*_ω and behaves optimally given this belief, it is in fact optimal for the regime to choose exactly $b_\omega = b^*_\omega$. If the citizen expects the bias to be b^*_ω, then the mean of her posterior belief about ω upon observing s_ω is:

$$\overline{\mu}_\omega = \lambda\mu_\omega + (1 - \lambda)(s_\omega + b^*_\omega)/\rho,$$

where λ is as defined in the main text. Note that the higher the expectation of bias, the higher the belief about the regime's unpopularity. Rearranging, she protests if and only if $\overline{\mu}_\omega > k/\overline{\beta}$, or:

$$s_\omega > \left(k/\overline{\beta} - \lambda\mu_\omega\right)\frac{\rho}{(1-\lambda)} - b^*_\omega \equiv \hat{s}_\omega.$$

So, the probability of protest when choosing bias level b_ω when the citizen expects bias level b^*_ω is:

$$\begin{aligned} &Pr\left(\rho\omega - b_\omega + \epsilon_\omega > \left(k/\overline{\beta} - \lambda\mu_\omega\right)\frac{\rho}{(1-\lambda)} - b^*_\omega\right) \\ &= Pr\left(\rho\omega + \epsilon_\omega > \left(k/\overline{\beta} - \lambda\mu_\omega\right)\frac{\rho}{(1-\lambda)} - b^*_\omega + b_\omega\right) \\ &\Phi\left(\sqrt{\rho^2\alpha_0 + \alpha_s}\left(\rho\mu_\omega - \left(k/\overline{\beta} - \lambda\mu_\omega\right)\frac{\rho}{(1-\lambda)} + b^*_\omega - b_\omega\right)\right). \end{aligned}$$

So, the expected payoff for choosing bias level b_ω when the citizen expects b_ω^* is:

$$u_i(b_\omega; b_\omega^*) = f(t,\theta)\Phi\left(\sqrt{\rho^2\alpha_0 + \alpha_s}\left(\rho\mu_\omega - \left(k/\overline{\beta} - \lambda\mu_\omega\right)\frac{\rho}{(1-\lambda)} + b_\omega^* - b_\omega\right)\right) - c(b_\omega),$$

giving first order condition:

$$\sqrt{\rho^2\alpha_0 + \alpha_s}\mathbb{E}[f(t,\theta)]\phi\left(\sqrt{\rho^2\alpha_0 + \alpha_s}\left(\rho\mu_\omega - \left(k/\overline{\beta} - \lambda\mu_\omega\right)\frac{\rho}{(1-\lambda)}\right)\right) = c'(b_\omega^*). \tag{12}$$

The LHS is not a function of b_ω, so whether this is met for any or more than one b_ω depends on the shape of the c function.

If c meets the Inada conditions (i.e., c is increasing and convex with $c'(0) = 0$ and $\lim_{b_\omega \to \bar{b}} c'(b_\omega) = \infty$) then there is a unique solution to equation 12.

When there is a unique pure strategy equilibrium to the model, the only difference in computing the comparative statics on the probability of protest with respect to the exogenous parameters (e.g., ρ and β_s) is that these also may affect the equilibrium level of bias chosen. However, as long as there is a pure strategy, the equilibrium bias choice does not affect the probability of protest, as the citizen adjusts for the bias in equilibrium. So, all of the comparative statics in proposition 1 still hold.

If the incumbent plays a mixed strategy, it becomes much more difficult to characterize the equilibrium strategies, but given the analysis of the model with uncertainty over the bias – which the mixed strategy case endogenizes – the main results seem unlikely to change.

Nonstandard Beliefs

To model noise in the posterior belief about the regime popularity, suppose the citizen's conditional belief about ω given s_ω is:

$$\overline{\mu}'_\omega(s_\omega) = \lambda\mu_\omega + (1-\lambda)(s_\omega/\rho) + \gamma$$

where γ is normally distributed with mean μ_γ and precision α_γ. Assuming the citizen behaves optimally given this belief – i.e., the citizen is unaware that she forms her belief incorrectly, she protests if $\overline{\mu}'_\omega(s_\omega) > k/\overline{\beta}$. The *ex ante* distribution of this belief is normal with mean $\mu_\omega + \mu_\gamma$ and precision $\frac{\tau_\mu\alpha_\gamma}{\tau_\mu+\alpha_\gamma}$ where τ_μ is defined in section 1. So the probability of protest is

$$\Phi\left(\frac{\tau_\mu\alpha_\gamma}{\tau_\mu+\alpha_\gamma}(k/\overline{\beta} - (\mu_\omega + \mu_\gamma))\right)$$

This is the same probability of protest derived in section 1 with $\frac{\tau_\mu\alpha_\gamma}{\tau_\mu}$ replacing τ_μ and $\mu_\omega + \mu_\gamma$ replacing μ_ω. So, the conclusions in the main model remain unchanged. In particular, increasing β_s increases the probability of protest and increasing ρ increases the probability of protest when $k/\overline{\beta} > \mu_\omega + \mu_\gamma$.

The citizen having the wrong belief about the effect of ρ is complex since ρ affects both the normalization of s_ω and the weight (i.e, λ). To simplify, suppose the misperception of ρ only affects the normalization. In particular, let:

$$\overline{\mu}^B_\omega(s_\omega) = \lambda\mu_\omega + (1-\lambda)(s_\omega/\rho'(\rho))$$

where $\rho'(\rho)$ is the adjusting parameter used by the citizen, which may be a function of the true ease of airing grievances. E.g., if $\rho'(\rho) = \rho$ the citizen forms the correct Bayesian belief, and if ρ' is constant in ρ then changes in the true ease of criticism has no effect on how the citizen adjusts.

Given the incorrect adjustment, the posterior belief as a function of the primitive random variables is:

$$\overline{\mu}_{\omega}^{B}(s_{\omega}) = \lambda\mu_{\omega} + (1-\lambda)((\rho\omega + \epsilon_s)/\hat{\rho}(\rho))$$

which is normally distributed with mean

$$\left(\lambda + (1-\lambda)\frac{\rho}{\hat{\rho}(\rho)}\right)\mu_{\omega}$$

and precision

$$\left(\frac{1-\lambda}{\hat{\rho}(\rho)}\right)^2 \frac{\rho^2\alpha_{\omega}\alpha_s}{\rho^2\alpha_{\omega}+\alpha_s}$$

If $\mu_{\omega} > 0$, the mean of the distribution of $\overline{\mu}_{\omega}^{B}(s_{\omega})$ is increasing in ρ if:

$$\frac{\partial}{\partial\rho}\left[\frac{\rho}{\hat{\rho}(\rho)}\right] = \frac{\hat{\rho}(\rho) - \rho\hat{\rho}'(\rho)}{\hat{\rho}(\rho)^2} > 0$$
$$\hat{\rho}(\rho) > \rho\hat{\rho}'(\rho)$$

If this inequality holds, say the citizen under-adjusts for the effect of ρ. So, when the citizen under-adjusts, there will be more protest through this channel.

However, increasing ρ also affects the precision of the posterior mean (even ignoring the effect of the changing the weight λ). The first term of the precision is decreasing in ρ (as long as $\hat{\rho}(\rho)$ is increasing in ρ) and the second is increasing, so the sign of that derivative is indeterminate. As in the main text, if this precision is increasing in ρ, this will lead to more protest when protest is *ex ante* likely and less protest when protest is less likely (setting aside the effect on the mean.)