Pink Slips From the Underground: Changes in Terror Leadership; Supplemental Appendix

Margaret J. Foster* David A. Siegel †

*PhD Candidate, Duke University; E-mail: margaret.foster@duke.edu

†Associate Professor, Department of Political Science, Duke University; E-mail: david.siegel@duke.edu.

We do three things in this supplemental appendix. First, we present a more complex version of the model in the text that incorporates: (i) the decisions of constituency and backers; (ii) different choices of the commander to start a splinter group or to inform on the group; and (iii) parameters for the strength of group rivals and the severity of the commander's infraction that produced the backers' ire. To avoid too much repetition, only new model components are justified here.

Second, we solve for the equilibrium and comparative statics of this more complex model. This serves additionally as a solution to the stripped-down version of the model in the text, since that model is a strict simplification of this one. Third, we briefly discuss results that go beyond those presented in the text.

Full Model

Actors and Actions

We have four actors: a group's *leader* (L), offending *commander* (C), *backers* (B), and the commander's *constituency* (O). The group's *rivals* are a combination of state entities and other anti-state groups that collectively compete against the group and/or desire the group's end. As they are not our focus we do not model them as strategic actors, but they do affect the payoffs of the actors.

The leader acts first and has three classes of punishment at his disposal: kill the commander (K), fire him (F), or discipline him internally (N). We described these in the text.

The commander acts second. Under each of K and N there is no further action for the commander to take. Under F, however, the now-fired commander has three options. One is as described in the text: withdrawal (W). The other two stem from disaggregating the splintering option (S in the text) into starting a splinter group (S here) and informing on the group (I). The splinter group will compete with the original group for constituency support. Given the visibility of this action, we assume that it is common knowledge to all actors. Informing amounts to providing useful intelligence to the group's rivals. For backers and constituency this is indistinguishable from withdrawal: the commander leaves the group and does not continue in the conflict. Thus costs induced by backers and constituency cannot be conditioned on withdrawal versus informing. The leader learns of this action by process of elimination.[1]

The constituency acts third. It observes the behavior of the leader and commander as described above and chooses its optimal level of continued support, s_O, for the group going forward. The backers act last. They observe the level of punishment the leader doles out and choose their optimal level of support, s_B, for the group.

Figure 1 illustrates the reduced form game tree for this model; we have left out the last two stages for readability, instead using standard backward induction to replace s_O and s_B with their equilibrium values. As in the text, we introduce the model's new parameters—two, in this case—before This covers all payoffs.[2]

discussing the model's new and/or altered payoffs.

Payoffs

Beyond the four parameters in the text—β_1, α, σ, and δ—this model introduces two additional ones. The first is the seriousness of the commander's infraction, β_2. More serious infractions are those that backers view as more necessary to punish harshly. This parameter will feed into decisions made by the backers. The second is γ, which is an index representing the strength of the group's rivals, relative to that of the group itself. Stronger rivals have multiple effects in the model, including increasing the group's need for support from all sources and altering the relative advantages to the leader and commander of the S and

[1]We might consider a variant in which the leader could kill an informing commander as revenge. If this could be done secretly the option to do so would eliminate any equilibrium in which I is chosen, leading to weakly less firing in equilibrium. However, our model assumes that the constituency learns when the commander is killed, implying that killing the commander after he informs is strictly worse for the leader than simply killing the commander in the first place and so sparing him rival-induced costs. Thus, in equilibrium, the leader would only kill the commander after he informed if the leader could surreptitiously do so. Post-hoc punishments, though, do not always succeed and failed assassination adds an aura of ineffectiveness.

[2]To accommodate limited space in the figure we have left out functional dependencies. These are: $s^\star_B(\phi, \beta_2), s^\star_O(\phi(1+\delta), \sigma\delta), p(\beta_1, \beta_2), u_N(\sigma(1+\delta)), u_S(\sigma, \delta), u_I(\sigma, \delta), R_N(\sigma, \delta), R_S(\sigma, \delta),$ and $R_I(\sigma, \delta)$.

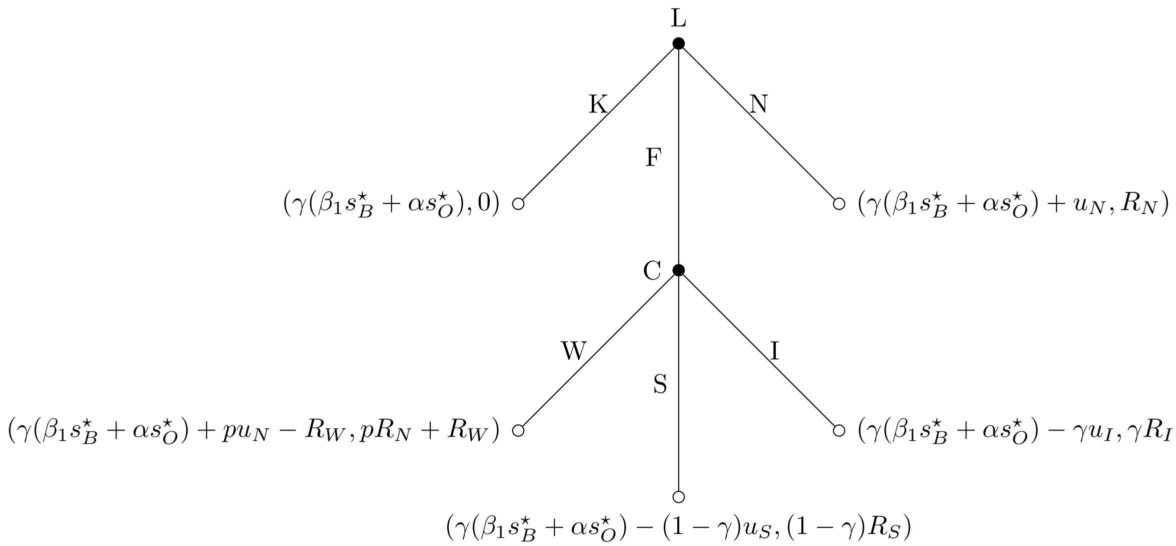

Figure 1: Reduced-Form Game Tree

I options. Note that we do not incorporate rivals' strength in the decision-making of the constituency in the model, though we do discuss what the effect of doing so would entail in the text. Table 1 lists all six parameters, along with their substantive meanings, for easy reference.

Param	Substantive Meaning
β_1	Group's Dependence on Backers / Strength of Backers
β_2	Seriousness of Commander's Infraction
α	Group's Dependence on Constituency / Strength of Constituency
δ	Group's Degree of Decentralization / Connection between Constituency and Commander
σ	Commander's Ability
γ	Rivals' Strength

Table 1: Exogenous Model Parameters

We begin with the backers' payoffs. We assume that the backers' utility, $U_B(s_B; \phi, \beta_2)$, depends on three things: the backers' support, s_B, the level of punishment the commander suffers, $\phi \in \{\phi_N, \phi_F, \phi_K\}$ with this list ordered, and the seriousness of the commander's infraction, β_2. We also assume it is globally concave in s_B. Increasing support increases the backers' utility, conditional on the extent to which their interests align with those of the group; greater punishments and less serious infractions signal more alignment. Under these assumptions, U_B has a unique interior maximum of U_B, denoted $s_B^\star(\phi, \beta_2) = \text{argmax}_{s_B} U_B$. Because the leader-commander interaction, rather than backer or constituency choice, is our focus, we save some time by directly making assumptions on this maximum. These are: $\frac{\partial s_B^\star}{\partial \phi} \geq 0$, $\frac{\partial s_B^\star}{\partial \beta_2} \leq 0$, and $\frac{\partial^2 s_B^\star}{\partial \phi \partial \beta_2} \geq 0$.[3] In words, the equilibrium level of support the backers grants to the group is increasing in the level of punishment and decreasing in the severity of the offense, and the marginal benefit to increasing punishment is itself increasing in the severity of the offense. The first two assumptions follow directly from our discussion, the third from assuming that punishment is more necessary the worse the offense.

We similarly assume that the constituency's utility is $U_O(s_O; \nu, \sigma, \delta)$ and depends on four things: the constituency's support for the group, s_O; the level of punishment the commander suffers, $\phi \in \{\phi_N, \phi_F, \phi_K\}$; the commander's ability, σ; and the group's degree of *de*centralization, δ. We assume that U_O is globally concave in support and we define its unique, interior maximum as $s_O^\star(\phi, \sigma, \delta) = \text{argmax}_{s_O} U_O$, and again choose to make assumptions directly on this. Support is decreasing in the level of punishment, and moreso the more the the constituency values the commander. To represent this, we assume that punishment enters the optimal support function as $\phi(1 + \delta)$. More able commanders lead to more support, but only in decentralized groups. To capture this, we assume that the commander's ability enters the optimal support function as $\sigma\delta$. Together, these assumptions imply that we can write constituency support as $s_O^\star(\phi(1 + \delta), \sigma\delta)$, where $\frac{\partial s_O^\star(a,b)}{\partial a} \leq 0$, $\frac{\partial s_O^\star(a,b)}{\partial b} \geq 0$ and $\frac{\partial^2 s_O^\star(a,b)}{\partial a \partial b} \leq 0$. The first two assumptions follow directly from our discussion, the third from assuming that punishment is viewed more negatively the more capable is the commander.

The leader's payoffs are largely described in the text; here we discuss only new aspects of them. In the stripped-down model, the utility the leader obtains from both backers and constituency depends on the group's attachment to each. In this model, greater dependency

[3]These assumptions follow from some not terribly illuminating, but viable, assumptions on mixed-third-derivatives of U_B. More to the point, they match the substance of the interaction. Group strength and the benefits it provides to backers, including but not limited to pressure on the state to change policy or surrender territory, public or club goods, or status within the group, affects the slope of U_B (and also U_O, below), which we largely elide, and so we do not consider it as a separate parameter.

on both actors, and so greater utility from their support, also follows from stronger rivals. To accommodate this we assume that γ multiplies both $s_B^\star$ and $s_O^\star$ in the leader's utility function.

The leader's additional payoff, beyond support, arising from the continued employ of the commander (choice N) remains $u_N(\sigma(1+\delta))$. Recall that $\frac{\partial u_N}{\partial \sigma} \geq 0$ and $\frac{\partial^2 u_N}{\partial \sigma \partial \delta} \geq 0$. The payoff under W to the leader remains almost the same as well, and is $p(\beta_1, \beta_2)u_N(\sigma(1+\delta)) - R_W$. The sole difference is that now the probability of circumstances changing sufficiently to allow the commander to return to the group depends not only on the strength of the backers, but also on the seriousness of the infractions. We assume that $\frac{\partial p}{\partial \beta_1} \leq 0$, $\frac{\partial p}{\partial \beta_2} \leq 0$, and $\frac{\partial^2 p}{\partial \beta_1 \partial \beta_2} \geq 0$, where the last assumption suggests that more serious infractions and greater backer stability are substitutes in their role in reducing the chance of a commander returning. Recall that R_W is chosen by the leader at the time of the decision to fire.

The leader's payoffs under outcomes S and I vary slightly from that from S in our stripped-down model, in that they depend on rivals' strength as well. As in the text, greater values of σ and δ increase the cost to the leader in a complementary manner. We thus define $u_S(\sigma, \delta)$ and $u_I(\sigma, \delta)$ with all first and mixed-partial derivatives of each weakly positive. The degree to which these costs affect the leader, however, depends on the strength of the rivals, γ. Stronger rivals mean less effective splinter groups that are less risky for the leader, since the new splinter group is likely ill prepared to deal with rivals at inception. In contrast, stronger rivals mean any information passed along will have greater consequences, increasing the cost of outcome I to the leader. To capture each, u_I enters the leader's utility multiplied by γ, and u_S enters multiplied by $(1-\gamma)$.

The commander's payoffs under K, N, and W are identical to those in the model in the text, with the exception that now p depends on both β_1 and β_2, as detailed above. Recall that the payoff for employment by the group is $R_N(\sigma, \delta)$, with both first derivatives as well as the mixed-partial positive. The new payoffs for S and I are similar to that given in the model in the text for its S option: $R_S(\sigma, \delta)$ and $R_I(\sigma, \delta)$, with the same dependencies as with u_S and u_I. As with the leader, the actual benefits to the commander of S and I are also a function of rivals' strengths. Thus, payoffs to the commander under S and I are, respectively, $(1-\gamma)R_S(\sigma,\delta)$ and $\gamma R_I(\sigma,\delta)$.

Equilibrium Behavior and Comparative Statics

We now offer our solution for the full model presented above, but note after presenting each of its equilibrium and comparative statics what changes there are in the stripped-down version in the text.

Equilibrium

This is a four stage, sequential, complete information game. We solve for the subgame perfect equilibria via backward induction. The five terminal nodes of the reduced-form game represented by Figure 1 are potential equilibrium outcomes.

Though the game we presented has four stages, we already determined the constituency's and the backers' equilibrium levels of support, $s^{\star}_O(\phi(1+\delta),\sigma\delta)$ and $s^{\star}_B(\phi,\beta_2)$, in the course of detailing our model. Thus we begin our equilibrium analysis in stage 2.

In stage 2, the commander (C) only has a decision to make in the history in which the leader (L) has chosen to fire (F) him. C's decision in this case is straightforward: Choose W whenever $R_W \geq \max\{(1-\gamma)R_S - pR_N, \gamma R_I - pR_N\}$, I whenever $R_W < \gamma R_I - pR_N$ and $\gamma R_I \geq (1-\gamma)R_S$, and S whenever $R_W < (1-\gamma)R_S - pR_N$ and $\gamma R_I < (1-\gamma)R_S$. Note that we've chosen to eliminate indifference by breaking utility ties in favor of comparatively less "risky" options: W over I over S, under the presumption that the first involves an immediate payoff, the second a delayed payoff potentially affected by the internal politics of the rivals, and the third a payoff dependent on the eventual success of the splinter group.

These conditions depend on the endogenous choice of R_W. In equilibrium, L will choose R_W to be the minimum to make C accept, conditional on L's preferring W to other outcomes. This implies that $R^{\star}_W = \hat{R_W} \equiv \max\{0,(1-\gamma)R_S - pR_N, \gamma R_I - pR_N\}$ whenever L would want C to accept this, which happens whenever $pu_N - \hat{R_W} \geq -\gamma u_I$ if $\gamma R_I \geq (1-\gamma)R_S$ or $pu_N - \hat{R_W} \geq -(1-\gamma)u_S$ if $\gamma R_I < (1-\gamma)R_S$, and $R^{*}_W = 0$ otherwise.

We can split these conditions into five cases to determine what happens upon a firing.

0. $\hat{R_W} = 0$, which implies $(1-\gamma)R_S - pR_N \leq 0, \gamma R_I - pR_N \leq 0$. In this case W is always chosen by C.
1. $\gamma R_I \geq (1-\gamma)R_S, \gamma R_I - pR_N > 0$, so that $\hat{R_W} = \gamma R_I - pR_N$, and $R_W^* = \hat{R_W}$, so that $pu_N - (\gamma R_I - pR_N) \geq -\gamma u_I$. In this case W is always chosen by C.
2. $\gamma R_I < (1-\gamma)R_S, (1-\gamma)R_S - pR_N > 0$, so that $\hat{R_W} = (1-\gamma)R_S - pR_N$, and $R_W^* = \hat{R_W}$, so that $pu_N - ((1-\gamma)R_S - pR_N) \geq -(1-\gamma)u_S$. In this case W is always chosen by C.
3. $\gamma R_I \geq (1-\gamma)R_S, \gamma R_I - pR_N > 0$, so that $\hat{R_W} = \gamma R_I - pR_N$, and $R_W^* = 0$, so that $pu_N - (\gamma R_I - pR_N) < -\gamma u_I$. In this case I is always chosen by C.
4. $\gamma R_I < (1-\gamma)R_S, (1-\gamma)R_S - pR_N > 0$, so that $\hat{R_W} = (1-\gamma)R_S - pR_N$, and $R_W^* = 0$, so that $pu_N - ((1-\gamma)R_S - pR_N) < -(1-\gamma)u_S$. In this case S is always chosen by C.

In stage 1, the leader's choice is dependent on what will happen in stage 2; thus we need to determine L's action in each of the five cases above. This leads to three sets of conditions that determine when K, N, and F are played, where we've again assumed less risky actions are taken in the case of indifference, with K less risky than N, which is less risky than F. In each case the conditions from the previous list are additionally assumed. To make our notation clearer we define four important utility differences: $b^1(\beta_2) = s_B^\star(\phi_K, \beta_2) - s_B^\star(\phi_F, \beta_2)$, $b^2(\beta_2) = s_B^\star(\phi_F, \beta_2) - s_B^\star(\phi_N, \beta_2)$, $o^1(\sigma, \delta) = s_O^\star(\phi_N(1+\delta), \sigma\delta) - s_O^\star(\phi_F(1+\delta), \sigma\delta)$, and $o^2(\sigma, \delta) = s_O^\star(\phi_F(1+\delta), \sigma\delta) - s_O^\star(\phi_K(1+\delta), \sigma\delta)$. In other words, b^1 and o^1 represent the decline in support from, respectively, the backers' and constituency's best outcomes to their second best outcomes, and b^2 and o^2 represent the decline in support from second to third best outcomes. Prior assumptions on support imply that $\frac{\partial b^i}{\partial \beta_2} \geq 0$, $\frac{\partial o^i}{\partial \nu} \geq 0$, and $\frac{\partial^2 o^i}{\partial \sigma \partial \delta} \geq 0$ for $i \in \{1, 2\}$, $\nu \in \{\sigma, \delta\}$.

First set: L chooses K. This requires $\gamma(\beta_1(b^1 + b^2) - \alpha(o^1 + o^2)) - u_N \geq 0$ and:

0. $\gamma(\beta_1 b^1 - \alpha o^2) - pu_N \geq 0$.
1. $\gamma(\beta_1 b^1 - \alpha o^2) - pu_N + \gamma R_I - pR_N \geq 0$.
2. $\gamma(\beta_1 b^1 - \alpha o^2) - pu_N + (1-\gamma)R_S - pR_N \geq 0$.

3. $\gamma(\beta_1 b^1 - \alpha o^2) + \gamma u_I \geq 0$.
4. $\gamma(\beta_1 b^1 - \alpha o^2) + (1-\gamma)u_S \geq 0$.

Second set: L chooses N. This requires $\gamma(\beta_1(b^1 + b^2) - \alpha(o^1 + o^2)) - u_N < 0$ and:

0. $\gamma(\alpha o^1 - \beta_1 b^2) + (1-p)u_N \geq 0$.
1. $\gamma(\alpha o^1 - \beta_1 b^2) + (1-p)u_N + \gamma R_I - pR_N \geq 0$.
2. $\gamma(\alpha o^1 - \beta_1 b^2) + (1-p)u_N + (1-\gamma)R_S - pR_N \geq 0$.
3. $\gamma(\alpha o^1 - \beta_1 b^2) + u_N + \gamma u_I \geq 0$.
4. $\gamma(\alpha o^1 - \beta_1 b^2) + u_N + (1-\gamma)u_S \geq 0$.

Third set: L chooses F. This requires:

0. $\gamma(\beta_1 b^1 - \alpha o^2) - pu_N < 0$ and $\gamma(\alpha o^1 - \beta_1 b^2) + (1-p)u_N < 0$.
1. $\gamma(\beta_1 b^1 - \alpha o^2) - pu_N + \gamma R_I - pR_N < 0$ and $\gamma(\alpha o^1 - \beta_1 b^2) + (1-p)u_N + \gamma R_I - pR_N < 0$.
2. $\gamma(\beta_1 b^1 - \alpha o^2) - pu_N + (1-\gamma)R_S - pR_N < 0$ and $\gamma(\alpha o^1 - \beta_1 b^2) + (1-p)u_N + (1-\gamma)R_S - pR_N < 0$.
3. $\gamma(\beta_1 b^1 - \alpha o^2) + \gamma u_I < 0$ and $\gamma(\alpha o^1 - \beta_1 b^2) + u_N + \gamma u_I < 0$.
4. $\gamma(\beta_1 b^1 - \alpha o^2) + (1-\gamma)u_S < 0$ and $\gamma(\alpha o^1 - \beta_1 b^2) + u_N + (1-\gamma)u_S < 0$.

These conditions specify the unique equilibrium of the game for any set of parameter values. Each of the five possible outcomes of the game, (i) K; (ii) N; (iii) F, W; (iv) F, I; and (v) F, S, occurs in equilibrium in a subspace of the overall parameter space.

In the stripped-down version of the model we present in the text, there are two relevant differences. The first is that we need to eliminate γ, $(1-\gamma)$, and β_2 from all terms in our derivation. The second is that the I option is no longer present, which eliminates all terms with an I subscript. This has a few consequences. First, we instead have that $R_W^{\star} = \hat{R_W} \equiv \max\{0, R_S - pR_N\}$ whenever L would want C to accept this, which happens whenever $pu_N - \hat{R_W} \geq -u_S$, and $R_W^{*} = 0$ otherwise. Second, the first and the third of the five post-firing cases no longer are relevant. This also eliminates the first and third sets of conditions in each of the three sets of conditions corresponding to L's choices. Third, for the remaining three cases, all conditions involving R_I are eliminated.

Comparative Statics

We now compute comparative statics for each of the six parameters in the model: $\beta_1, \beta_2, \alpha, \delta, \sigma$, and γ. We briefly discuss the effect of these parameters on the commander's decisions before turning to our true interest: the leader's choice of K, F, or N.

Commander's Choice

The commander effectively has a nested decision to make. If a non-zero payout R_W is offered or if both other options are worse than the chance of returning to the group despite no payout (i.e., $\hat{R_W} = 0$), then he chooses W. Otherwise, he chooses I if $\gamma R_I \geq (1-\gamma)R_S$ and S if not. There are three parameters that affect the latter decision: γ, σ, and δ. As γ increases, C is (weakly) more likely to choose I than S in equilibrium.[4] As $\nu \in \{\sigma, \delta\}$ increases, I becomes more likely whenever $\frac{\partial}{\partial \nu}(\gamma R_I - (1-\gamma)R_S) \geq 0$. In other words, if commander ability or group decentralization increases payouts to I faster than to S.

It is not necessary to offer any payout when $\hat{R_W} \equiv \max\{0, (1-\gamma)R_S - pR_N, \gamma R_I - pR_N\} = 0$. This is more likely to occur for large p, and so for small β_1 and β_2. It is also more likely to occur when R_N is large compared to R_S and R_I, so when the group offers the best source of employment. This happens when $\nu \in \{\sigma, \delta\}$ increases if $\frac{\partial}{\partial \nu}(\max\{\gamma R_I, (1-\gamma)R_S\} - pR_N) \leq 0$.

When $\hat{R_W} > 0$, $R_W^{\star} > 0$ if $\hat{R_W} \leq pu_N + \gamma u_I$ or if $\hat{R_W} \leq pu_N + (1-\gamma)u_S$, depending on what C would choose if no payoff were offered. If I would be chosen, this becomes $\gamma R_I - pR_N \leq pu_N + \gamma u_I$, or $\gamma(R_I - u_I) \leq p(u_N + R_N)$. Similarly, if S would be chosen, we have: $(1-\gamma)(R_S - u_S) \leq p(u_N + R_N)$. Both are more likely the higher is p, and so the lower are β_1 and β_2. Both are also more likely the larger is the worth to L of C's being in the group, as well as the larger is the cost to L of I or S, but less likely the larger is the value C gets for being in the group. This happens when $\nu \in \{\sigma, \delta\}$ increases if $\frac{\partial}{\partial \nu}(\gamma(R_I - u_I) - p(u_N + R_N)) \leq 0$ and $\frac{\partial}{\partial \nu}((1-\gamma)(R_S - u_S) - p(u_N + R_N)) \leq 0$.

[4] To simplify discussion we often leave off the "(weakly)" in what follows. Including weakly is necessitated by the discrete nature of the first and second stage choices in our game. There are more complex solutions to this problem, including quantal response equilibrium, noise in the utility functions, or distributions over parameters, but as all yield the same insights we stick with the easier one.

Leader's Choice

The first set of conditions above specify when the leader chooses K. By inspection, all six conditions are more easily satisfied as β_1 and β_2 increase, and as α decreases. K is (weakly) more likely to be chosen by L as $\nu \in \{\sigma, \delta\}$ decreases if: (i) L would chose N over F; (ii) L would choose F over N and $\hat{R_W} = 0$; or (iii) L would choose F over N and $\hat{R_W} > 0$ and $-\gamma\alpha\frac{\partial o^2}{\partial \nu} + Z \leq 0$, where $Z = -p\frac{\partial(u_N+R_N)}{\partial \nu} + \gamma\frac{\partial R_I}{\partial \nu}$ if W would be chosen over I and I over S, $Z = -p\frac{\partial(u_N+R_N)}{\partial \nu} + (1-\gamma)\frac{\partial R_S}{\partial \nu}$ if W would be chosen over S and S over I, $Z = \gamma\frac{\partial u_I}{\partial \nu}$ if I would be chosen over W and S, and $Z = (1-\gamma)\frac{\partial u_S}{\partial \nu}$ if S would be chosen over W and I. K is (weakly) more likely to be chosen by L as $\nu \in \{\sigma, \delta\}$ increases if (iii) holds save that its inequality is reversed.

K is (weakly) more likely to be chosen by L as γ increases if $\beta_1(b^1 + b^2) - \alpha(o^1 + o^2) \geq 0$ and $\beta_1 b^1 - \alpha o^2 + Z \geq 0$ where $Z = 0$ if $\hat{R_W} = 0$ and if $\hat{R_W} > 0$, $Z = R_I$ if W would be chosen over I and I over S, $Z = -R_s$ if W would be chosen over S and S over I, $Z = u_I$ if I would be chosen over W and S, and $Z = -u_S$ if S would be chosen over W and I. K is (weakly) more likely to be chosen by L as γ decreases if all inequalities are reversed. γ has an indeterminate effect on the choice of K otherwise.

The second set of conditions above specify when the leader chooses N. By inspection, all six conditions are more easily satisfied as α increases. N is (weakly) more likely to be chosen by L as β_1 decreases if: (i) L would choose K over F; (ii) L would choose F over K and I would be chosen over W and S or S over W and I; or (iii) L would choose F over K and W would be chosen over I or S and $\gamma b^2 + \frac{\partial p}{\partial \beta_1}Z \geq 0$ where $Z = u_N$ if $\hat{R_W} = 0$ and $Z = u_N + R_N$ if $\hat{R_W} > 0$. N is (weakly) more likely to be chosen by L as β_1 increases if (iii) holds save that its inequality is reversed. Similarly, N is (weakly) more likely to be chosen by L as β_2 decreases if: (i) L would choose K over F; (ii) L would choose F over K and I would be chosen over W and S or S over W and I; or (iii) L would choose F over K and W would be chosen over I or S and $\gamma\beta_1\frac{\partial b^2}{\partial \beta_2} + \frac{\partial p}{\partial \beta_2}Z \geq 0$ where $Z = u_N$ if $\hat{R_W} = 0$ and $Z = u_N + R_N$ if $\hat{R_W} > 0$. N is (weakly) more likely to be chosen by L as β_2 increases if (iii) holds save that its inequality is reversed.

L is (weakly) more likely to choose N as $\nu \in \{\sigma, \delta\}$ increases if (i) L would choose K over F; (ii) L would choose F over K and I would be chosen over W and S or S over W and I or $\hat{R_W} = 0$; or (iii) L would choose F over K and W would be chosen with $\hat{R_W} > 0$ and $\gamma\alpha\frac{\partial o^1}{\partial \nu} + (1-p)\frac{\partial u_N}{\partial \nu} - p\frac{\partial R_N}{\partial \nu} + Z \geq 0$, where $Z = \gamma\frac{\partial R_I}{\partial \nu}$ if I would be chosen over W and S, and $Z = (1-\gamma)\frac{\partial R_S}{\partial \nu}$ if S would be chosen over W and I. L is (weakly) more likely to choose N as $\nu \in \{\sigma, \delta\}$ decreases if (iii) holds save that its inequality is reversed.

L is (weakly) more likely to choose N as γ increases if $-\beta_1(b^1 + b^2) + \alpha(o^1 + o^2) \geq 0$ and $-\beta_1 b^2 + \alpha o^1 + Z \geq 0$ where $Z = 0$ if $\hat{R_W} = 0$ and if $\hat{R_W} > 0$, $Z = R_I$ if W would be chosen over I and I over S, $Z = -R_s$ if W would be chosen over S and S over I, $Z = u_I$ if I would be chosen over W and S, and $Z = -u_S$ if S would be chosen over W and I. N is (weakly) more likely to be chosen by L as γ decreases if the inequalities are reversed. γ has an indeterminate effect on the choice of N otherwise.

The third set of inequalities above specify when the leader chooses F. As F, unlike K and N, is the choice taken when L must balance multiple costs and benefits, it is generally not the case that F will become uniformly more likely as one parameter increases or decreases. But it can happen over restricted parameter ranges. We consider each parameter in turn. As increasing α always makes the first of each pair of inequalities easier to satisfy but the second more difficult, it is not the case that increasing α has a uniform effect. Rather, increasing α will weakly increase the likelihood of F if doing so does not alter the fact that the second inequality holds. Similarly, decreasing α will weakly increase the likelihood of F if this does not alter the fact that the first inequality holds. As this behavior is solely an artifact of assumed discontinuities in choice behavior, we do not focus on it further and simply state that the comparative statics on α are not determined.

Next consider β_1. Decreasing it weakly increases the likelihood that the first inequality holds in all cases. However, it weakly decreases the likelihood that the second inequality holds in cases 3 and 4, and so no uniform effect exists there. In cases 0-2, decreasing β_1 increases the likelihood that the second inequality holds whenever $-\frac{\partial p}{\partial \beta_1} Z \geq \gamma b^2$, where $Z = u_N$ if $\hat{R_W} = 0$ and $Z = u_N + R_N$ if $\hat{R_W} > 0$ and W would be chosen over S or I. Thus, in this region of the parameter space, F is more likely as β_1 decreases. The same logic holds for β_2, replacing the condition with $-\frac{\partial p}{\partial \beta_2} Z \geq \gamma \beta_1 \frac{\partial b^2}{\partial \beta_2}$.

Now consider $\nu \in \{\sigma, \delta\}$. Increases in ν have opposite effects on the two inequalities when $\hat{R_W} = 0$ (case 0) and so no uniform effect exists there. When I or S would be chosen over W (cases 3 and 4), the second inequality is always more difficult to satisfy as ν increases. For these cases decreases in ν lead to more F if the first inequality is also more difficult to satisfy as ν increases, which leads to the condition: $Z \geq \gamma\alpha \frac{\partial o^2}{\partial \nu}$ where $Z = \gamma \frac{\partial u_I}{\partial \nu}$ if I would be chosen over W and S and $Z = (1-\gamma)\frac{\partial u_S}{\partial \nu}$ if S would be chosen over W and I. When W would be chosen over S or I and $\hat{R_W} > 0$ (cases 1 and 2), there are two possible conditions under which ν has a uniform effect. If increasing ν makes the first inequality easier to satisfy, which happens when $\gamma\alpha\frac{\partial o^2}{\partial \nu} + p\frac{\partial (u_N + R_N)}{\partial \nu} \geq Z$, then increasing ν leads weakly to more firing whenever $\gamma\alpha\frac{\partial o^1}{\partial \nu} + (1-p)\frac{\partial u_N}{\partial \nu} + Z \leq p\frac{\partial R_N}{\partial \nu}$, where for both conditions $Z = \gamma\frac{\partial R_I}{\partial \nu}$ if I would be chosen over S and $Z = (1-\gamma)\frac{\partial R_S}{\partial \nu}$ if S would be chosen over I. If the first condition fails to hold, so increasing ν makes the first condition harder to satisfy, then decreasing ν leads to weakly more firing whenever the second condition also fails to hold. Otherwise there is no uniform effect of varying ν on firing.

Finally, consider γ and let $Z = 0$ if $\hat{R_W} = 0$ and if $\hat{R_W} > 0$, $Z = R_I$ if W would be chosen over I and I over S, $Z = -R_s$ if W would be chosen over S and S over I, $Z = u_I$ if I would be chosen over W and S, and $Z = -u_S$ if S would be chosen over W and I. L becomes weakly more likely to fire as γ increases whenever $\beta_1 b^1 - \alpha o^2 + Z \leq 0$ and $\alpha o^1 - \beta_1 b^2 + Z \leq 0$ and weakly more likely to fire as γ decreases when both inequalities are reversed. γ has an indeterminate effect on the choice of F otherwise.

In the stripped-down version of the model we present in the text, there are three relevant differences in the comparative statics. One, the definition of $\hat{R_W}$ is now as given above at the end of the equilibrium analysis section, without any R_I or $(1-\gamma)$ terms. Two, all comparative statics relating to the choice of I are not relevant and may be ignored. Thus the only relevant comparison is between the W and S choices, and many conditions are simplified, as discussed above. Three, all comparative statics involving γ and β_2 are not relevant and may be ignored.

Additional Results

To conclude this appendix, we mirror the text's discussion of determinants of the leader's choice, focusing on the effects of the two parameters left out of the stripped-down model: β_2 and γ.

We begin again with the choice of killing the commander. More serious infractions, as one would expect, make this more likely. The strength of rivals acts to enhance the influence of the more important source of group support. If the benefit of backer support outweighs the loss of constituency support, then stronger rivals generally enhance the effect, making killing more likely.[5]

Next we consider internal discipline. Rivals have a similar effect on internal discipline as on killing. If the benefits of constituency support outweigh the cost of the loss of backer support, then stronger rivals make internal discipline more likely.[6] Decreasing the seriousness of the infraction has exactly the same effect on the likelihood of internal discipline as does decreasing the strength of, and the group's dependence on, the backers. As we described it in the text, we do not repeat it here.

Finally, we consider firing. As noted in the text, and as seen above, firing betrays complex relationships with most parameters. However, it does relate cleanly to one parameter: rivals' strength. Increasing rivals' strength makes firing more likely whenever the costs in loss of backer and constituency support upon firing are smaller than the respective costs upon internal discipline and killing, and either (i) the commander would choose to form a splinter group rather than inform, or (ii) the reverse is true but the cost to the leader of being

[5]The only exception to this occurs when the commander would choose to start a splinter group rather than inform and the leader has chosen to kill in order to avoid the high cost of a splinter group or the high withdrawal payout to prevent a splinter group.

[6]The only exception to this occurs when the commander would choose to start a splinter group rather than inform and the leader has chosen to employ internal discipline in order to avoid the high cost of a splinter group or the high withdrawal payout to prevent a splinter group.

informed on is relatively low. If the opposite of these conditions holds, decreasing rivals' strength leads to more firing. Thus, as long as firing is not too costly, strong rivals can enhance the need to keep both backers and constituency in the fold after the commander's infraction.