# Models of Terrorist Organization *

Jacob N. Shapiro†

This version August 1, 2013.

## 1 Introduction

This working paper presents several models of the organizational difficulties facing terrorist organizations. My goal here is not to provide a single best model of terrorist organization. Rather, it is to some different ways of modeling the strategic situation underlying the managerial challenges outlined in *The Terrorist's Dilemma*. Other approaches to doing so are certainly possible and probably more elegant. These materials are intended to be useful to students or more skilled theorists interested in improving on the analysis in *The Terrorist's Dilemma*.

The basic setup throughout is quite simple. There are two actors, a terrorist boss, B, and an operative, T. B and T have different preferences over targets, basically different beliefs about what is politically optimal, but B's preferences are unknown to T and T's preferences are unknown to B. The group has to conduct a series of two attacks. B can send a signal to T about his preferences, but pays some security cost for doing so. After T conducts the first attack, B can also punish him at some additional security cost and hire a new T for the second attack.

There are many different ways to represent this basic interactions. Section 2 presents a full analysis of the model in chapter 5 of Shapiro (2013). In this model there are two targets, B is always right about which one is the best for political reasons from his perspective, and we want to know how the difference in T's preferences affects how B manages the interaction. Section 3 examines the same model with two targets using a more traditional information structure, one in which B and T each have separate views about what is politically optimal

*The author thanks Jon Bendor, James Fearon, Adam Meirowitz, Kris Ramsay, Peter Schram, and David Siegel for critical feedback at various points. All errors and omissions are my own.

†Assistant Professor, Woodrow Wilson School and Department of Politics, Princeton University, Princeton, NJ 08544; Email: jns@princeton.edu, URL: http://www.princeton.edu/~jns

(why they have such views is not modeled) but the probability they agree on one target is different than the probability they agree on the other.[1] The focus is of section 2 is on understanding when different canonical management approaches can emerge in equilibrium, including: (1) what we can think of as the normal organizational outcome in which B tells T what to attack in the first period and T follows orders; (2) management by exception in which B only pays the cost to tell T where to attack if a specific situation occurs (i.e. he learns that he likes target 2 instead of target 1 or vice versa); and (3) an analogue to the leaderless resistance model advocated by Beam (1992) in which B provides broad ideological guidance but does not interact with his operatives in any other way. Section 4 studies the same problem with a continuous target set and signaling space using the models from my 2008 working paper "Bureaucracy and Control in Terrorist Organizations." The models in section 4 are my preferred way of thinking about terrorist organization.

The key point across all three models is that as long as the level of preference divergence between B and T is not too large, and as long as the costs of communicating and punishing when needed are modest, then leaders can exercise control and ensure their preferred target is hit in the first period. When one of those parameters gets too large, then cooperation within the group breaks down.

Finally, an important caveat to this analysis is that I do little to address equilibrium selection and do not deal with mixed strategies in sections 2 and 3. A fuller formal analysis would do both.

# 2 Model from *The Terrorist's Dilemma*

We can start by outlining the conditions for what can call normal organizational behavior to occur in equilibrium. In most organizations, loosely speaking, leaders tell their employees what they would like to happen, employees carry out those requests, and leaders remain uncertain whether employees took the desired action because they wanted to, or because of the threat of disciplinary action for not doing so. In the terrorist context, this amounts to leaders telling their operatives which targets to attack and the operatives following those orders. We can think of this as the high control equilibrium.

This kind of high-control equilibrium is possible in this section's model of the interaction between a terrorist boss, B, and his agent, T, when three conditions are met:

[1] As chapters 2, 3, and 8 of *The Terrorist's Dilemma* show, preference divergence within terrorist organizations stems from many sources and is almost always present.

1. T must prefer doing what B wants rather than attacking her preferred target and being sanctioned for doing so, what we will call the Motivational Constraint (MC).

2. It must be in B's best interests to sanction T if she takes actions that reveal she does not share B's preferences about targets, but not otherwise, what we will call the Credibility Constraint (CC).[2]

3. B must find that gains from communicating his preferred target outweigh the security costs of doing so given beliefs about T's preferences and the assumption that MC and CC are met (SC).

In practical terms, for these conditions to be met it must be the case that: discrimination between targets is important enough to leaders that they are willing to pay the security costs of communicating; there is sufficient uncertainty about targets to make communicating valuable; and preference divergence is not so large that operatives would rather face potential punishment than do what their leaders want.

## 2.1 Setup

The basic setup is quite simple. There are two actors, a terrorist boss, B, and an operative, T. There are two targets. B and T have different preferences over targets, basically different beliefs about what is politically optimal, but B's preferences are unknown to T and T's preferences are unknown to B. The group has to conduct a series of two attacks.[3]

### 2.1.1 Sequence of Play

There are two targets, $\theta \in \{1, 2\}$. Before the interaction begins B collects information through some process and calculates (without loss of generality) that hitting target 1 will best serve the group's political goals, $\theta_B = 1$. This choice is unobserved by T who also collects information through her own process. The probability that she also decides that target 1 is the best option is $1 - \rho$, with $\rho \in [0, \frac{1}{2}]$.[4]

[2] This is analogous to the time-consistency issue in Shapiro and Siegel (2012).

[3] Note that we are not concerned here with a principal-agent problem in which the principal and the agent each get separate information and try to come to some sort of best choice despite different preferences as in Crawford and Sobel (1982), or Gilligan and Krehbiel (1987). Rather, this is a model in which B is always right from B's perspective, and we want to know the difference in T's preferences affects how B manages the interaction.

[4] We can think of these preferences as being drawn from two different, but correlated distributions. Suppose, for example, that the information the players get about the world is drawn from a bivariate normal distribution with $\mu_B = 0, \sigma_B = 1$, $\mu_T = 0, \sigma_T = 2$, and correlation $\hat{\rho} = .75$ so that there is more

This probability, which is common knowledge, captures the theoretical concept of uncertainty over targets given information about the world. Since B's preferences are unknown to T, and T's preferences are unknown to B, uncertainty peaks at $\rho = \frac{1}{2}$. Anything more than that and they know that it's more likely than not that the other party does not share their preferences. Throughout this section uncertainty refers to the chance that B and T's preferences are not aligned.

After preferences are realized, the following stage game is repeated twice:

**Communication Phase.** B decides to communicate about what target it wants hit or not, paying cost $c$ if he does, and send a message $m \in \{1, 2\}$. Here $c$ captures the marginal risk B incurs from government counterterrorism when he communicates; think of the increased risk from a courier being intercepted, for example.[5]

**Attack Phase.** T chooses where to attack, at target 1 or target 2. Denote this choice $a \in \{1, 2\}$.

**Punishment Phase.** Having observed T's choice of attacks, B updates his prior beliefs about the probability that T's preferred target is target 1, to $\rho_1$ and can decide to sanction T or not, paying cost $s$ if he does. We can think of $s$ as the unit cost of punishment given prior investment in a security apparatus. Here $s$ captures factors such as the probability that the punished agent changes sides or retaliates against the group. T suffers disutility $r$ from being sanctioned and is not rehired.

**Final Phase.** If B sanctions, then the act of doing so provides information to potential recruits about his preferences, so that the new T is more likely to shares B's preferences and therefore $\rho_2 < \rho$. We can think of this as the pool of T's learning something about the type of targets B likes given information about the world, and so those who do not agree stay away, rendering the remaining potential recruits more likely to share B's preferences.[6] In the next round, B's preferences remain the same, as do T's if she is

variance in the information T gets but the information is correlated. Assume that both players prefer target 1 if their information, $x_i \leq 0$. With this underlying probability structure $\rho$ in the model is just $Prob(x_B < 0, x_T > 0 | x_B < 0) = Prob(x_B > 0, x_T < 0 | x_B > 0) \approx .33$. Notice too that in this setting the symmetry of the probability model for the preferences means that $\rho$ does not depend on the signal. When that is not symmetric, then T's expectation about B's preferences, call it $\rho_T$, will depend on the signal it gets. The next section analyzes a more standard model.

[5]Security costs that do not depend on whether B communicates are captured in the discount factor, $\delta$, which will be discussed under payoffs.

[6]An alternative setup would be to give B some positive utility from sanctioning when he is not obeyed, to have all subsequent operatives drawn from the same pool. An even simpler model could drop the second

rehired. If a new T is hired, that T's learning from B's actions in the first round are captured in the fact that $\rho_2 < \rho$, uncertainty is reduced.

Figure 1 illustrates the first stage of this interaction, assuming $\theta_B = 1$ and suppressing information sets at the last node. For readability the figure does not show the branch where $\theta_B = 2$, but note that T knows his own preference but is uncertain as to B's, as illustrated by the information sets. Note also that some of the information sets in the final move for B may not be there in equilibrium.

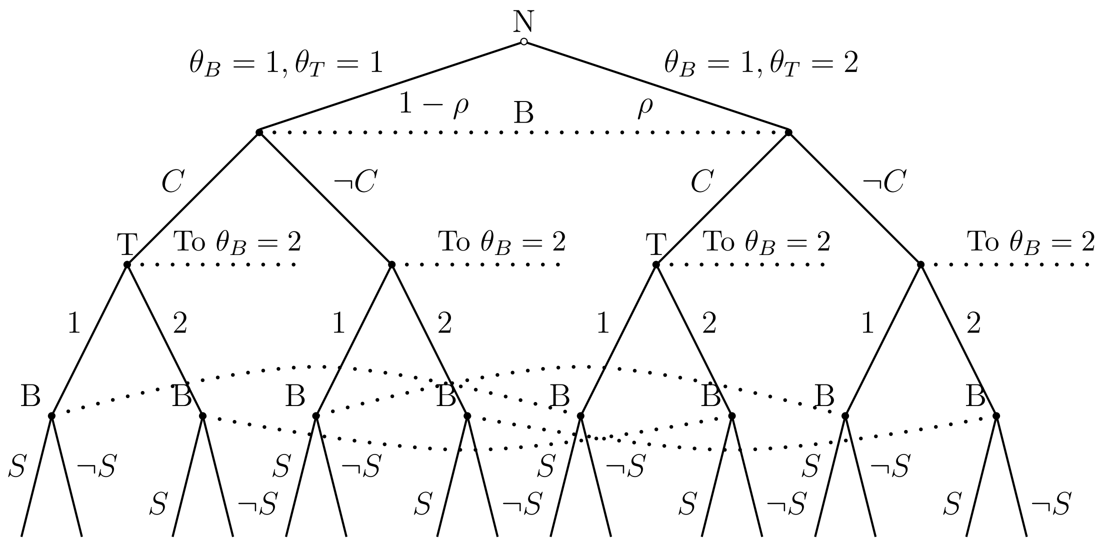


Figure 1: Extended form of security-control tradeoff game.

### 2.1.2 Payoffs

In each stage B gets utility $b > 0$ if T attacks his preferred target and 0 otherwise. T gets $t > 0$ for attacking her preferred target and 0 otherwise. Thus $\rho$ capturing the probability of preference divergence and $t$ captures its intensity when it occurs. Payoffs in the second period are discounted by $\delta \in [0, 1]$, which can be thought of as capturing the portion of the risk operatives face from government counterterrorism efforts regardless of whether B communicates.

Preference divergence is captured in $t$, which represents how much T values doing what it feels best, regardless of B's desired target. The key effect of $t$ is to determine whether MC can be met: whether T will find the cost of being sanctioned, $r$, sufficiently high that she wants to attack where B says even if that is her least preferred target.

period and simply examine the 1-period interaction under two conditions: (1) when punishment is credible through some un-modeled value B gets by maintaining his reputation; and (2) when that value is less than the cost of sanctioning.

Discrimination is captured in $b$, the value to B of having an attack on its desired target. As we will see, $b$ determines whether it is worth it to B to take the actions required to achieve a high level of control (communicating its preferences and sanctioning T if she does not attack where told).

As noted, uncertainty is captured in $\rho$, the probability that B and T do not share preferences over targets.

## 2.2 Analysis

This section analyzes the model. We first derive the three constraints that divide up the parameter space and then describe some of the sustainable equilibrium in different parts of that space.

### 2.2.1 Constraints

**Motivation Constraint: MC**

B's ability to influence T's actions in the first period depends on the threat of being fired motivating T who do not share B's preferences from attacking her favored target in the first period. Whether the threat does so is a function of three factors: how likely is it that T will get to carry out the second attack given the operational environment, which is captured by $\delta$; how much T values hitting her preferred target, which is captured in $t$; and how effectively B can punish T, which is captured in $r$. If being sanctioned is not that bad for T (given the inherent risks of being a terrorist) and the value of hitting her preferred target is large enough that $r < (1-\delta)t$, then there is no incentive to do what B wants and T has a dominant strategy of attacking her preferred target. If, however, the net benefit of hitting her preferred target and getting sanctioned is less for T than the benefit of hitting the target B wants now and then hitting her preferred target in period 2, if

$$r \geq (1-\delta)t, \tag{MC}$$

then T's best response is to listen to B if it communicates honestly (assuming of course that she will be hired again).

**Credibility Constraint: SC**

Suppose MC binds so that T has an incentive to avoid being sanctioned. B's strategy about sanctioning must hinge on whether the improvement in the outcome from sanctioning given what he learned from T's actions, $\delta b\rho_1(\cdot) - \delta b\rho_2$, exceeds the cost of sanctioning T. CC thus involves two constraints. First, B must want to fire T if she takes actions that reveal her to have different preferences. Second, B must want to retain her when T does what she is told. If, after all, T is going to be fired in any case, then the T who does not agree with B about targets has no incentive to obey him in the first period. More formally, it must be the case that the cost of sanctioning divided by the benefits of hitting the good target, $\frac{s}{\delta b}$ is less than the reputational benefit (given beliefs) if T does not hit B's preferred target, but more than the reputational benefit (given beliefs) if T obeys: $\rho_1(C,O)-\rho_2 \leq \frac{s}{\delta b} < \rho_1(C,\neg O)-\rho_2$. Clearly what B believes after T's attack, $\rho_1(\cdot)$, depends on his communication strategy and T's best response to that, but if communications motivate all types of T to follow B's orders (if the T pool on attacking where told), then we have a simple statement of CC:

$$\rho - \rho_2 \leq \frac{s}{\delta b} < 1 - \rho_2. \tag{CC}$$

We will refer to the situation where the left hand inequality is violated so that B always wants to fire as $\neg CC_{AF}$ and to that where the right hand inequality is violated so that B will never want to fire as $\neg CC_{NF}$.

**Security Constraint: SC**

Solving for when $U_B(C) \geq U_B(NC)$ defines $c^*$, the largest cost of communications for which B finds it worthwhile to guarantee his desired outcome in the present period by communicating. This $c^*$ is the security constraint, SC. To pin down SC, note that if B's communication does not influence T's decision at the attack phase, then it can never be the case that the inequality is met since communicating would be a pure cost. Assuming that T's best response to communications is to attack where told, then rearranging the inequality a bit yields the condition that it is a best response for B to communicate when

$$c \leq c^* = \rho(s + b(1 - \delta(1 - \rho_2))). \tag{SC}$$

These three constraints divide the parameter space up into twelve blocks (2x3x2) which can be aggregated into four regions.

1. When MC, CC, and SC are all met then what we will call the High Control equilibrium is possible. In this equilibrium B communicates his preferences in the first round, T attacks where she is told, B rehires T, and T attacks her preferred target in the second round. Several other perverse equilibrium are possible as well.

2. When either SC is not met, MC is not met, or both are not met, then what we will call

Trust but Punish is possible. In this equilibrium B does not communicate, T attacks her preferred target, B then fires if T attacks his least preferred target and rehires T otherwise, and whichever T is working in the second period attacks her preferred target.

3. When CC is not met because B never wants to fire T even if he knows for sure she is the 'wrong' type, then what we will call the Independent Agents equilibrium is possible. In this equilibrium B never communicates, T attacks her preferred target in the first round, B rehires, and in the second period T attacks her preferred target.

4. When CC is not met because B wants to fire T if her actions are uninformative about her type, then if either MC or SC are not met there is another Independent Agents equilibrium in which T attacks her preferred target in the first round, B does not rehire if T attacks his least preferred target, and whichever T is working in the second period T attacks her preferred target. If CC is not met because B's best response is to fire when the T's pool, then there are multiple equilibrium depending on the parameter values.

The following subsections discuss each of these regions in turn.

### 2.2.2 High-Control

**Proposition 2.1**: If MC, SC, and CC are met, then there is a 'high-control' equilibrium in which B communicates his preferences honestly, T attacks where told if she receives communications, B rehires T following the attack, and then T attacks her preferred target in the last period. This equilibrium is sustained by the following off-the-path beliefs: if T receives no communication then she simply uses her priors from play on the path and assumes that she and B share preferences with probability $1 - \rho$; if B sees an attack it did not order he assumes that he and T share preferences with probability 0.[7]

To shows that this is an equilibrium we need to verify that there are no profitable one-period deviations given off-the-path beliefs and that all actions are best-responses given beliefs on the path. Start with the final phase. Sanctioning is not a credible threat in

[7]This equilibrium is sustainable for any off-the-path beliefs for B in which seeing an attack he did not order weakly increases his estimate of the probability that he and T have different preferences over targets. If, perversely, B's belief upon seeing an attack he did not order decreases his estimate of the probability that they have different preferences then this equilibrium cannot be sustained as B will never want to fire a T who deviates.

the final period, B cannot impact any future actions by T since the game ends after the last sanctioning stage. Since B will not sanction, T's best response is to take her preferred action in the final attack stage. T is therefore guaranteed a payoff of $\delta t$ for being rehired. Knowing his utility is the same whether he communicates or not in the final stage, B will never communicate in the second period. B's value for the second period is thus $\delta(1-\rho_1(\cdot))b$ if he sticks with the same T after period 1 and $\delta(1-\rho_2)b$ if he sanctions T and hires someone new.

Suppose that B were to deviate at the punishment phase. If B fires a T who listened he pays cost $s$ and gets a new T, yielding expected utility $-s+(1-\rho_2)\delta b$ which is clearly less than B gets in expectation from re-using the T who listened, $(1-\rho)\delta b$, under the assumption that CC is met. There is likewise no profitable deviation for T at the attack phase is she if asked to hit her least preferred target. If she attacks her preferred target, then she gets $t$ with probability 1 and is fired, paying cost $r$. If she attacks where told, she is guaranteed a payoff of $\delta t$ in the final period. Clearly this cannot be a profitable one-period deviation if T is asked to attack her preferred target, and if MC is met it cannot be a profitable one period deviation since $t-r \leq \delta t$.

Finally, we need to check whether there is a profitable deviation for B at the communication phase. If B does not communicate T's expected utility from attacking her preferred target given B's punishment strategy (fire if the less preferred target is attacked) and her off-the-path beliefs is $t+(1-\rho)\delta t-\rho r$ and from attacking the other target it is $\rho\delta t-(1-\rho)r$. Setting up the inequality for when T will attack her preferred target and rearranging yields $t(1+\delta) \geq -r$ which is clearly true. Thus if B fails to communicate, T's best response given off-the-path beliefs is to hit her preferred target. With that in hand we can compare B's expected utility from his equilibrium strategy, $U_B(C_H) = b-c+(1-\rho)\delta b$, with what he gets if he does not communicate, $U_B(NC) = b(1+\delta-\rho-\delta\rho\rho_2)-\rho s$, or if he communicates dishonestly, $U_B(C_D) = -c-s+(1-\rho_2)\delta b$. Not communicating is not a profitable deviation since $U_B(C_H)-U_B(NC) = \rho(s+b(1-\delta(1-\rho_2)))-c$ which is positive as long as SC is met. Communicating dishonestly is also not a profitable deviation since $U_B(C_H)-U_B(C_D) = b(1-\delta(\rho-\rho_2))+s$ which is strictly greater than zero.

### 2.2.3 Examples of other equilibrium when MC, SC, and CC are met

As one would expect in an interaction of this type, a number of other equilibrium can be sustained by different off-the-path beliefs even when MC, SC, and CC are all met. There is, for example, a no-communications equilibrium in which B does not communicate, T attacks where she prefers, B fires the T who attacks where he does not like and rehires the other T, and then whichever T is employed in the last period attacks her preferred target. This equilibrium is sustainable, for example, if T's off-the-path belief is that if it receives any signal then she and B share preferences with probability zero.

As before there can be no profitable deviations in the final phase. Suppose that B were to deviate at the punishment phase by rehiring the T who attacked the target he did not like. Since T's equilibrium strategy up to this point is to attack her preferred target, B's posteriors about her type correctly identify her preferences and if B rehires he gets nothing in the final stage. That payoff of zero is clearly less than the payoff to firing, $(1-\rho_2)\delta b - s$, if CC is met. Firing after seeing an attack on his preferred target also cannot be a profitable deviation as doing so yields $(1-\rho_2)\delta b - s$ which is clearly less than the $\delta b$ he gets from rehiring the operative who has shown that she prefers the same target. Turning to T the fact that $\rho < \frac{1}{2}$ means there can be no profitable deviation for T at the attack phase. Given B's equilibrium strategy attacking her least preferred target yields $-r(1-\rho)+\rho(\delta t)$ which is clearly inferior to the $(1-\rho)(t+\delta t) - r\rho$ she gets for attacking her preferred target. Finally, we need to check whether B would want to communicate. The difference between B's utility when he communicates and when he does not is $-c-s(1-\rho)-b(1+\delta(\rho^2+\rho_2)-2\rho(1+\delta\rho_2))$ which is clearly less than zero as all the constituent terms are negative under the assumption that $0 < \rho_2 < \rho < \frac{1}{2}$.

Other equilibrium can be constructed, but all require some form of intuitively unappealing off-the-path beliefs or dishonest signaling strategy.[8]

### 2.2.4 Trust But Punish

**Proposition 2.2**: If CC is met but either SC is not met, or MC is not met, or neither is met, then there is a 'trust but punish' equilibrium in which B does not communicate, T attacks her preferred target, B fires the T who attacks his least preferred target and rehires the other T, and then whichever T is employed in the last period attacks her preferred target. This equilibrium is sustainable for any off-the-path beliefs by T about B's type if he signals.

Consider first the case where SC is not met but MC is, so that if T knew what B wanted she would find it worth doing. As before there can be no profitable deviations in the final phase. Suppose that B were to deviate at the punishment phase by rehiring the T who

[8]For example the equilibrium where B signals the opposite of its type and T responds as though B is always signaling opposite to type is sustainable with off-the-path beliefs that the two parties share preferences with probability $\rho$.

attacked the target he did not like. Since T's equilibrium strategy up to this point is to attack her preferred target, B's posteriors about her type correctly identify her preferences and if B rehires he gets nothing in the final stage which is clearly less than the payoff to firing $(1-\rho_2)\delta b - s$ if CC is met. Firing after seeing an attack on his preferred target also cannot be a profitable deviation as doing so yields $(1-\rho_2)\delta b - s$ which is clearly less than the $\delta b$ he gets from rehiring the operative who has shown that she prefers the same target.

Given the lack of a signal from B in equilibrium there can be no profitable deviation by T at attack stage; once again attacking her least preferred target yields $-r(1-\rho)+\rho(\delta t)$ which is clearly inferior to the $(1-\rho)(t+\delta t)-r\rho$ she gets for attacking her preferred target. Finally, consider deviations by B at communication stage. For any deviation to be profitable it must be that the off-the-path beliefs for T and T's best response function lead to an outcome that is worth the cost of communicating. That clearly cannot be the case as when SC is not met even being able to guarantee his preferred outcome in the next period does not make it worthwhile for B to communicate relative to T's equilibrium strategy of attacking her preferred target.

Now consider the case where MC is not met but SC is, so that if B could get T to do what he wanted in the first period that would be worth the cost of communicating. As when SC is not met there can be no profitable deviations in the final phase or the punishment phase. Now, however, the threat of punishment does not outweigh the value of hitting her preferred target in the present period, so T has a dominant strategy to attack her preferred target regardless of her beliefs about B's preferences. Given that, there can be no profitable deviation for B; communicating is a pure cost.

### 2.2.5 Independent Agents

**Proposition 2.3**: If CC is not met because $s > (1-\rho_2)\delta b$, so that B will never find it worthwhile to fire T even if he learns she likes the wrong target from his perspective (what we called $\neg CC_{NF}$), then there is an 'independent agents' equilibrium in which B does not communicate, T attacks her preferred target, B rehires T, and T attacks her preferred target. This equilibrium is sustainable for any off-the-path beliefs by T about B's type if he signals.

As before there is never any reason for the new T to condition her behavior in the final period and therefore never any reason for B to communicate in that period. Now, consider a 1-period deviation by B at the punishment phase. If B sees an attack he liked, it can clearly not be a profitable deviation to fire T as he trades of $\delta b$ against $-s+(1-\rho_2)\delta b$. If B sees an attack he did not like, then he knows for sure T does not share his preferences, but it is not worth it to deviate as paying the cost of punishment exceeds the value of getting a better operative in the final round, $s > (1-\rho_2)\delta b$. Given that B will not condition hiring devisions on T's actions, there can be no profitable deviation by T at the attack stage. Note that this is true regardless of what T learns about B's type from its actions at the communications stage. There can therefore be no profitable deviation by B at communication stage.

**Proposition 2.4**: When CC is not met because $s \leq (\rho - \rho_2)\delta b$, so that B wants to fire if T's actions are uninformative about her type (what we called $\neg CC_{NF}$), and either MC is not met or SC is not met, then another Independent Agents equilibrium exists. In this Independent Agents equilibrium T attacks her preferred target in the first round, B does not rehire if T attacks his least preferred target but rehires otherwise, and whichever T is working in the second period T attacks her preferred target. This equilibrium is sustainable for any off-the-path beliefs by T about B's type if he signals.

As before there is never any reason for the new T to condition her behavior in the final period and therefore never any reason for B to communicate in that period. Now, consider a 1-period deviation by B at the punishment phase. If B sees an attack he liked, it can clearly not be a profitable deviation to fire T as he trades of $\delta b$ against $-s + (1-\rho_2)\delta b$. If B sees an attack he did not like, then he knows for sure T does not share his preferences, but it is not worth it to deviate as failing to punish guarantees a payoff of zero in the next round which is less than $(1-\rho_2)\delta b - s$ in this part of the parameter space. In equilibrium T receives no signal and so her best bet is to attack her preferred target (recall that $\rho < \frac{1}{2}$ so B and T are more likely than not to share preferences) and so there is no profitable deviation at the attack stage. If either MC or SC is not met then there is no profitable deviation for B: if MC is not met then his signal doesn't create any incentives for T making signaling a pure cost; and if SC is not met then even guaranteeing his preferred attack is not worth the cost of signaling.

**Proposition 2.5**: When CC is not met because $s \leq (\rho - \rho_2)\delta b$ ($\neg CC_{AF}$) so that B wants to fire if T's actions are uninformative about her type, but both MC and SC are met, then no equilibrium exists in which B signals honestly and T's off-the-path beliefs if he sees no signal are that they share target preference with probability greater than $\frac{1}{2}$. There are parameter values for which no equilibrium exists in which B does not signal and T's off-the-path beliefs are that B signals honestly as well as parameter values for which an equilibrium in which B does not signal can be supported by the same off-the-path beliefs.

To see the first part of the proposition first consider the equilibrium in which B signals honestly at the communication phase, T attacks the target that maximizes her expected utility looking forward, B fires if T attacks his least preferred target, and then strategies in the final period are as above. Once again, no deviation is profitable before the attack phase for the reasons outlined above. If B signals, T can guarantee herself a chance to hit her preferred target in the next round which because MC is met means she while attack where B requests. Suppose there is a a one-period deviation by B at communication stage in which B is silent. Given B's strategy at the attack phase T gets $t + (1-p)\delta t - pr$ for attacking her most preferred target and $p\delta t - (1-p)r$ for attacking her least preferred target, where $p$ is her off-the-path belief about the probability she and B prefer the same target. Taking the difference and simplifying shows that attacking her most preferred target is the best response given B's equilibrium strategy whenever $r(1-2p) + t(1+\delta(1-2p)) > 0$ which is always the case for $p \leq \frac{1}{2}$. Since T's best response is to attack her most preferred target if B is silent, the payoff following this one period deviation is $(1-\rho)(b+\delta d) + \rho(-s + (1-\rho_2)\delta b)$, while the payoff from signaling is $b - c + (1-\rho)\delta b$. Taking the difference between these utilities shows that deviating to no-signal is always weakly better when $c \leq \rho(s + b(1 - \delta(1-\rho_2)))$, which has to be true given that SC is met.

Now consider the equilibrium in which B signals honestly at the communication phase, T attacks the target that maximizes her expected utility looking forward, B fires regardless of where T attacks, and then strategies in the final period are as above. This cannot be an equilibrium as B would obviously like to retain the T who attacks his preferred target. For the same reason, it cannot be an equilibrium for B to never fire T after signaling honestly.

To see the second part of the proposition consider the equilibrium in which B's strategy is to not signal at the communication phase, T attacks the target that maximizes his expected utility looking forward, B fires if T attacks his least preferred target, and then strategies in the final period are as above. Once again, no deviation is profitable before the communications phase. However since both MC and SC are met B would find getting his preferred attack worth more than the cost of communicating and T would find the threat of punishment bad enough to make attacking B's preferred target in the first period a best response if she knew what it was. Now a one-period deviation by B at the communications phase where he signals honestly can be profitable. Knowing that B's equilibrium strategy is to fire if he sees his least preferred target attacked, T's best response given off-the-path beliefs that B is signaling honestly is to attack B's preferred target. Returning then to his equilibrium strategy B would rehire T and T would attack her preferred target in the final phase. That deviation earns B $b + \rho\delta b - s$ compared to $\rho(b + \delta b) + (1 - \rho)((1 - \rho_2)\delta b - c)$. Taking the difference and re-arranging yields a net benefit to deviating of

$$b(1 - (1 - \rho_2)\delta) + c - \frac{s}{1 - \rho}. \tag{1}$$

There are a range of values for which that net benefit in (1) is positive when MC and SC are both met. Trivially, for example, when $s$ is vanishingly small, the inequality can always be satisfied. Notice, though, that signaling would not be sustainable as an equilibrium if $\neg CC_{AF}$ is true given what we can think of sensible off-the-path beliefs by T. The issue is that B would not be best responding at the punishment phase by rehiring if the T pool (which would mean that her actions are uninformative as to her type) and therefore T would have no incentive to respond to the communication.

Finally, to show the third part of the proposition some algebra shows that there are parameter values for which that deviation is not profitable (equation 1 is less than 0) even though both MC and SC are met, in which case the equilibrium can be sustained.

# 3 Terrorist management with two targets and a traditional information structure

In this section we analyze the same sequence of play as in section 2 but in a more traditional principal-agent setup in which the principal and the agent each get separate information but the probability they agree that target 1 is the best place to attack is different than the probability that they agree target 2 is the best one. Defining preference divergence as the value T places on hitting her preferred target whether or not B agrees, then the core insight from this setup is that there a range of different management options are available depending on how costly communications are and how much preference divergence there is. When preference divergence is too high, no cooperation is possible. When preference divergence is in the middle range and communications are not too costly, then the high control equilibrium is possible. If preference divergence is low and communications are not too costly, then B a management by exception equilibrium is possible in which B tells T where to attack only if his preferences are unusual. If preference divergence is low and are communications expensive (i.e risky) then a leaderless resistance equilibrium is possible in which B establishes some shared priors on his preferences but does not signal directly and T attacks where she B will want her too.

## 3.1 Sequence of Play

The core setup is almost the same as in section 2. There are two targets, $\theta \in \{1, 2\}$ and the players' types are the target they prefer. Before the interaction begins B collects information and determines with probability $p > \frac{1}{2}$ that hitting target 1 will best serve the group's political goal. T also collects information and with probability $q$ she too decides that target 1 is the best option. We then have the following distribution over preferences: $Prob(\theta_B = 1, \theta_t = 1) = pq$, $Prob(\theta_B = 1, \theta_t = 0) = p(1-q)$, $Prob(\theta_B = 0, \theta_t = 1) = (1-q)p$, $Prob(\theta_B = 0, \theta_t = 0) = 0$. These probabilities are common knowledge.

Notice right away that this information structure obscures some of the core theoretical concepts in *The Terrorist's Dilemma*. Divergence in underlying preferences is captured in $t$, as before. In this setup, however, there is no clear distinction between uncertainty by T about what will be politically optimal, which we can think of as being captured by how close $p$ is to $\frac{1}{2}$, and uncertainty reflected in the difference between the probabilities that B and T like target 1, $p - q$.

After the players' types are determined actions proceed as before: B decides sends a message $m \in \{ns, 1, 2\}$ where $ns$ stands for the decision to send no signal, paying cost $c$ if he sends 1 or 2; T decides where to attack striking a target $a \in \{1, 2\}$, B decides whether to fire T or not, paying cost $c$ if he does so and imposing cost $r$ on T. The process is then repeated with the stipulation that B has the same preferences in the second period as does any T who is rehired; and if a new T is hired she too prefers target 1 with probability $q$. B receives payoff $b$ for getting his preferred target hit and zero otherwise. T gets $t$ for attacking her preferred target and zero otherwise. Utilities in the second period are discounted by common discount factor $\delta \in (0, 1)$ which captures the probability the players survive to the second period.

We will look for Sub-game Perfect Nash Equilibrium (SPNE) in pure strategies, focusing on the following:

**High control.** In this equilibrium the B separate by sending honest messages about their type and the T pool on attacking where told in the first period.

**Management by exception.** We will look at two management by exception (MBE) equilibrium. In MBE1 B signals only if he has unusual preferences, if $\theta_B = 2$, and the T pool on attacking target 1 if there is no signal and on attacking target 2 if B tells them to. In MBE2 B signals if $\theta_B = 1$, and the T pool on attacking target 1 if B tells them to and on attacking target 2 if B sends no signal. Note that since B is more likely to prefer target 1, MBE1 yields higher *ex ante* expected utility.

**Leaderless resistance.** In this equilibrium the B pool on not signaling and the T pool by attacking target 1 in the first period.

**Independent Agents.** In this equilibrium both types of B pool on not signaling and the T separate by attacking their preferred target in the first period.

The first equilibrium represents the 'normal' organizational outcome. The second two can be thought of as more efficient versions of the first as the outcomes in terms of attacks are the same but off-the-path beliefs no longer generate incentives for both types of B to signal (which is inefficient in a world with only two targets). The fourth equilibrium represents the model of terrorism where the leadership provides broad guidance and the operatives make decisions for themselves without direct communication. The fifth is what happens when the central leadership has no means to enforce discipline.

## Preliminaries

**Lemma 3.1**: In the final period B's best response at the communications phase is to not signal, T's best response at the attack phase is to strike where told, and B's best response at the punishment phase is to do nothing.

The lemma follows straightforwardly from the players' utility functions. B pays cost $c$ for punishing in the final phase but gains no benefit, so his best response at the punishment phase is to do nothing. Knowing this, T's best response is to attack her preferred target as not doing so yields utility 0, and doing so yields utility $t$. Since communicating cannot influence T's actions at this stage, B receives the same utility from T's actions in expectation regardless of what he does but pays cost $c$ for communicating, and thus his best response at the communications phase in the final period is not to signal.[9]

**Lemma 3.2, Upper Motivational Constraint (UMC)**: T would be willing to attack her least preferred target in the first period to avoid being fired whenever preference divergence is not too large given the value of the future and cost of being punished.

Formally, the value to T of attacking her preferred target in the present period is $t$. The value of being rehired given lemma 1 is $\delta t$. The cost of being fired is $r$. When $r > (1 - \delta)t$

[9]This need not be the case in a different model where preferences where determined by noisy signals of some underlying state of the world which in turn influenced the players' utilities from hitting different targets, but that would be a different kind of model altogether.

T's best response is to take an action that will allow it to attack again in the next period, even if that means attacking her least preferred target. If UMC is not met, then T obviously has a dominant strategy to attack her preferred target in the first period. Note that we can also write UMC as a constraint on the level of preference divergence,

$$t < \frac{r}{1-\delta}, \tag{UMC}$$

which will be useful in a moment.

**Lemma 3.3, Lower Motivation Constraint (LMC)**: Assume that B's best response at the punishment phase is to fire any T who does not attack his preferred target. If $\theta_T = 1$ then T's best response at the attack stage if B does not signal is to set $a = 1$ for any off-the-path beliefs in which B prefers target 1 with probability greater than $\frac{1}{2}$. If $\theta_T = 2$ then T's best response at the attack stages is to hit her preferred target if preference divergence ($t$) is large enough, otherwise it is to hit target 1.

Define $\tilde{p} \equiv Prob(\theta_B = 1|m = ns)$ and assume that $\tilde{p} > \frac{1}{2}$. The first part of the lemma follows from the fact that if $\theta_T = 1$ then the difference between T's utility for attacking target 1 and that for attacking target 2, $E(U_T(a = 1) - U_T(a = 2)|\theta_T = 1)$ is $(2\tilde{p} - 1) + \frac{t}{\delta t + r}$ which is clearly greater than zero for any $\tilde{p} > \frac{1}{2}$. The second part of the lemma follows from the fact that if $\theta_T = 2$ then $E(U_T(a = 1) - U_T(a = 2)|\theta_T = 2) = (2\tilde{p} - 1)(\delta t + r) - t$. That difference is negative whenever

$$t \geq \frac{r}{1-\delta(2\tilde{p}-1)}, \tag{LMC}$$

and thus if $\theta_T = 2$ and $t$ is large enough T will attack target 2 instead of her preferred target if B sends $m = nc$ and his best response at the punishment phase is as assumed.

Note that for $\tilde{p} \in (\frac{1}{2}, 1)$ there is always a gap between UMC and LMC and thus there can always be $t$ where both are met.

**Lemma 3.4, Credibility Constraint (CC)**: B will be willing to fire T if he knows she does not share his preferences over targets whenever the probability of getting his preferred type in the next period is high enough compared to the cost of firing weighted by the value of having his preferred target hit.

Suppose MC binds so that T has an incentive to avoid being sanctioned for sure. B's strategy about sanctioning must hinge on whether the improvement in the outcome from sanctioning given what he learned from T's actions. Given Lemma 3.1, B's best response is to fire T if

$$\delta s < b\rho - \delta b\rho(\cdot). \tag{CC2}$$

Here $\rho(\cdot)$ is the posterior probability that T shares B's his preferences (given either Bayes' Rule or off-the-path beliefs) and $\rho$ is the probability a new T will do so.

For the HC, MBE1, and MBE2 equilibrium in which the T pool in the first period in equilibrium, it turns out that off-the-path beliefs that are critical. To see why, note first that it can never be worthwhile for B to fire a T he knows to share his type (if $\rho(\cdot) = 1$) as the inequality in CC2 cannot be met. This is because there is no reputational gain as in the previous setup. If, on the other hand, we set $\rho(\cdot) = 0$ to mimic the condition where B has learned T does not like the same target then this condition is met when $s < \frac{q}{\delta b}$ if $\theta_B = 1$ and when $s < \frac{1-q}{\delta b}$ if $\theta_B = 2$. Whichever B is less likely to get his preferred target in the next period will want to refrain from firing for a lower cost of sanctioning. While this may seem counterintuitive at first, it occurs because all previous costs are sunk by the punishment phase, and so a B who is sure his current T has different preferences is simply trading the cost of sanctioning against the expected value of getting a new T who might share his preferences with probability $q$.

**Lemma 3.5, Security Constraint (SC)**: B will be willing to communicate it's type whenever the gain in the expected value of the game going forward exceeds the cost of communicating.

Solving for when $U_B(C) \geq U_B(NC)$ defines the largest cost of communications for which B finds it worthwhile to guarantee his desired outcome in the present period by communicating. That value is the security constraint, SC. The conditions for SC to be pinned down are less straightforward than for UMC, LMC, or CC and so we will leave addressing them to our analysis of the equilibrium.

## 3.2 High Control (HC)

Formally the high control equilibrium is a messaging strategy for B, an attack strategy for T, a disciplinary strategy for B, and then a final period strategies for both players such that all players are best responding on and off-the-path given the other ' equilibrium strategies, beliefs formed by Bayes Rule on the path, and off-the-path beliefs. We will consider the following strategy profile: in the first period in B sends an honest message about his type, T attacks where told, T has off-the-path beliefs if she receives no signal that $Prob(\theta_B = 1) > \frac{1}{2}$ (exactly what they are in the range may matter as described below), and B has off-the-path beliefs that T prefers the target she attacks, and B rehires the T who attacks where told.

**Proposition 3.1**: If LMC, UMC, and CC are all met then the high control equilibrium can be sustained as long as the cost of communicating is not too high and the probability that T prefers target 1 is not too low.

On the path the T pool on attacking where told. B therefore learns nothing about T's type from her first period attack and thus faces the same lottery in the next stage whether he fires or not. Since he pays a cost $c$ for firing there can be no profitable deviation for B at the punishment phase. If MC is met there can be no profitable deviation for T at the attack phase as long as B's off-the-path belief that he and T prefer the same target are sufficiently small given the costs of signaling and the value of getting the right attack. To see this note that B's best response at the punishment phase is to fire T as long as $\rho(\cdot) < \rho - \frac{s}{\delta b}$, which is just CC. If that condition is met then the T who is asked to attack her least preferred target gets $t - r$ for doing so (given B's best response) which is less than the value of attacking her preferred target in the next period so long as MC is met.

Now consider the payoffs to B from a one-period deviation at the communications phase. If B does not communicate T's off-the-path beliefs are that $Prob(\theta_B = 1) > \frac{1}{2}$ and we have assumed that both LMC and UMC are met. Lemma 3 shows that T's best response is therefore to attack her preferred target: given her uncertainty about B's preferences she likes the gamble of hitting her preferred target better than that of hitting target 1. With that best response off-the-path by T, B has to tradeoff his expected utility from staying on the equilibrium path against the potential gains from not signaling and having T attack her preferred target. That tradeoff is negative when $c$, the cost of communications, is small and we can define $c_1$ ($c_2$) as the largest $c$ for which a B who prefers target 1 (target 2) has no incentive to make a single-period deviation from the high-control equilibrium. Formally,

$$\begin{aligned} c_1 &= (1-q)(s + b(1 - q\delta)) \\ c_2 &= q(b+s) + b(1 - q - q^2)\delta. \end{aligned} \tag{SC}$$

Three facts are immediately apparent. First, $c_1$ is weakly positive so that for any $q < 1$ there exists some $c > 0$ s.t. $c < c_1$, meaning that for cheap enough communications it can always be the case that B does not want to deviate if $\theta_B = 1$. Second, $c_1 \leq c_2$ if $q \geq \frac{1}{2}$. In other words, the type of B who is more likely to share T's preferences has a lower threshold cost at which they will want to deviate and not communicate in the first period, which makes intuitive sense. Third, $c_1$ is decreasing in $q$ and $c_2$ is increasing in $q$, so that the highest cost of communicating each type will tolerate before deviating is increasing in the probability that T has the opposite preferences. Pulling these facts together shows that the high-control equilibrium can clearly be sustained if $q > \frac{1}{2}$ and $c < c_1$.

To see the second part of the proposition note that it is possible that $c_2$ can be negative for low enough $q$, meaning that even if the risks from communicating are zero B will want to deviate if $\theta_B = 2$. If that is the case, the equilibrium cannot be sustained as some types of B would want to deviate at the communication phase. To quantify the point at which the possibility of the high-control equilibrium working breaks down define $\tilde{q}$ as the lowest $q$ for which there is always a $c$ low enough (though it could be zero) that neither type of B's best response is to deviate (assuming that B stays on the path if he is indifferent). Some algebra provides the condition that

$$\tilde{q} = \frac{b(1 - \delta(1 - \sqrt{5 + \frac{(b+s)(b+s+2\delta b)}{\delta^2 b^2}} - s}{2\delta b}. \tag{2}$$

Noting that $\tilde{q} \leq \sqrt{2} - 1$ shows that as long as $q$ is larger than that value there will always be a $c$ small enough that the high-control equilibrium can be sustained if CC, UMC, and LMC are met.

Taking derivatives of $\tilde{q}$ shows that the threshold probability for which the B who prefers target 2 is not willing to deviate is increasing in $b$ and decreasing in $s$. This makes sense intuitively as we would think that the threshold probability of getting his preferred attack required to get a B who likes target 2 to deviate is lower (and thus $\tilde{q}$ higher) when he places a higher value on getting his preferred attack, but should increase as the costs of sanctioning (which he can avoid with certainty by staying on the path) go up.

## 3.3 Management by Exception

There are two ways in which the high-control equilibrium can break down that do not involve so much preference divergence (such large $t$) that T simply wants to attack her preferred target. The first is that the value to T of hitting her preferred target, $t$, could be so small that under HC her best response given off-the-path beliefs if she received no signal would be to play the odds by attacking target 1 instead of by hitting her preferred target. If that were the case, then the B who preferred target 1 would have a profitable deviation at the communications phase. When that is the case what we will call the MBE1 equilibrium can emerge in which B only signals honestly if the unusual event happens and he prefers target 2. The second way HC can break down is that T is sufficiently likely to prefer the less likely target, target 2, that the B who also like that target no longer find it worth the cost of communicating to ensure that T attacks there. When HC breaks down for this reason, then the MBE2 equilibrium can emerge in which B signals honestly if the more likely event happens (he prefers target 1) and remains silent otherwise.

Below we analyze each situation in turn.

### 3.3.1 MBE1

When the value to T of hitting her preferred target, $t$, becomes so small that under HC her best response given off-the-path beliefs if she received no signal would be to play the odds by attacking target 1 instead of by hitting her preferred target, then B will have a profitable deviation at the communication phase if $\theta_B$=1.

To examine this situation suppose that UMC binds but that LMC does not so that if B does not communicate then all types of T prefer to attack target 1 given some off the path belief $\tilde{p} > \frac{1}{2}$ s.t. $t < \frac{r}{1-\delta(2\tilde{p}-1)}$. In this situation an MBE equilibrium is possible in which B sends $m = ns$ if $\theta_B = 1$ and $m = 2$ if $\theta_B = 2$, T's off-the-path beliefs are that B likes target 1 with probability 1 if he sends $m = 1$, T attacks target 1 if B does not signal and target 2 otherwise, B's off-the-path beliefs if T attacks target 2 after he sends no signal are that $q = 0$ and if T attacks target 1 after he sends $m = 2$ that $q = 1$, B rehires T who hit his preferred target and fire T who do not.

**Proposition 3.2**: If $t$ is small enough that UMC is met so that T would like to avoid being fired for sure regardless of his preferences, but LMC is not, so that T who are unsure of B's type will hit target 1 to avoid any risk of being fired, then the MBE equilibrium can be sustained as long as the costs of sanctioning is modest but non-zero and the costs of communicating are small enough.

As before we can prove the proposition by checking for profitable one-period deviations at each phase assuming others will best respond to those given off-the-path beliefs. At the sanctioning stage we need to check that neither B has any incentive to deviate. If $\theta_b = 1$ then B does not signal in equilibrium and T sets $a = 1$ regardless of her preferences. If $\theta_B = 2$ sends $m = 2$ and sees his preferred target hit. In neither case can he update on T's type given her equilibrium strategy and so pays a cost $s$ for sanctioning but gets no gain in his expected utility for the next stage. Deviating at the punishment phase thus cannot make sense for him.

What about best responses off-the-path at the punishment phase? If $\theta_B = 1$ and T attacks target 2 then given equilibrium play and off-the-path beliefs B is certain they do not share preferences and so gets zero expected utility in the next period from rehiring but gets $q\delta b - s$ from firing. Firing yields positive utility as long as $s < q\delta b$ and is thus a best response is $s$ is small. If $\theta_B = 2$ and T attacks target 1 then given equilibrium play and off-the-path beliefs B is certain they do not share preferences and so gets a value of zero in the next period from rehiring but gets $(1-q)\delta b - s$ from firing. Firing is therefore a best response off the path as long as $s < (1-q)\delta b$.

At the attack phase, note that B's equilibrium strategy after not signaling is to fire if $a \neq 1$ and after sending $m = 2$ it is to do so if $a \neq 2$. Given that, as long as UMC is met there can be no deviation profitable deviation for T. To see this first consider the situation where B has not signaled. If $\theta_T = 1$ then T gets $t + \delta t$ for attacking where told and $-r$ for deviating, which clearly cannot be profitable. If $\theta_T = 2$ then deviating yields $t - r$ while staying on the path yields $\delta t$ and that cannot be a profitable deviation either if UMC is met. Now consider the situation where B has sent $m = 2$. Given equilibrium strategies T knows B prefers target 2. If $\theta_T = 1$ then T gets $\delta t$ for attacking where told and $t - r$ for deviating, which clearly cannot be profitable if UMC is met. If $\theta_T = 2$ then deviating yields $-r$ while staying on the path yields $t + \delta t$ and that cannot be a profitable deviation.

Finally, we need to check for deviations at the communications phase. T's off-the-path beliefs if she sees $m = 1$ is that B likes target 1 with certainty. Her best response to a one-period deviation in which B sends $m = 1$ is therefore clearly to attack target 1 whatever her type; remember that when B returns to the path at his next information set he will follow his equilibrium strategy of firing if T does not attack his preferred target. Given that best response off-the-path by T, deviating to $m = 1$ cannot be profitable when $\theta_B = 1$ as B pays cost $c$ but gets the same outcome as from not signaling. Likewise it clearly cannot be optimal for B to send $m = 2$ if $\theta_b = 1$ as doing so yields $-c - s + q\delta b$ which is clearly less than the $b + q\delta b$ he gets from staying on the path. If $\theta_B = 2$ then neither sending $m = 1$ nor not signaling are profitable deviations. Taking either action gets T his less desired attack and he then has to pay the cost of sanctioning (recall he'll return to equilibrium play).

The interesting question in MBE1 is how large the cost of communicating would have to be to make it worthwhile for B to deviate to sending no message if $\theta_B = 2$. Sending a message yields utility $b - c + (1-q)\delta b$ and keeping quite yields $-s + (1-q)\delta b$. A bit of algebra shows that the largest $c$ for which B weakly prefers staying on the MBE1 path is $\hat{c} = b + s$. The cost for which MBE1 can be sustained is increasing in the value to B of getting his preferred target hit and in the cost of sanctioning, which B avoids by staying on the path.

Notice that MBE1 can be sustained everywhere HC can but is also possible when $t$ is too small for HC to work.

### 3.3.2 MBE2

When the probability T likes target 2 gets large enough it can be the case that the B who likes target 2 would prefer to deviate from the high-control equilibrium by not signaling and having T play her off-the-path best response of attacking where she prefers. We can think of this as the situation where B knows he has a problematic agent, expects that one target will be best and tells his agent so, but then changes his mind. He is tempted then to want to give up control, even though he thought it would be worthwhile *ex ante*.

To examine this situation suppose that $q < \sqrt{2} - 1$ and $c > c_2$ so that B would prefer to deviate from HC by not communicating if $\theta_B = 2$. Then the MBE2 equilibrium is still possible. In MBE2 B sends $m = 1$ if $\theta_B = 1$ and $m = ns$ if $\theta_B = 2$, T's off-the-path beliefs are that B likes target 2 with probability 1 if he sends $m = 2$, T attacks target 1 if B signals and target 2 otherwise, B's off-the-path beliefs if T attacks target 2 after he sends a message to attack target 1 are that $q = 0$ and if T attacks target 1 after he sends no signal that $q = 1$, B rehires T who hit his preferred target and fire T who do not.

**Proposition 3.3**: If $t$ is between LMC and UMC then the MBE2 equilibrium can be sustained as long as the costs of sanctioning is small enough and the costs of communicating are modest. This equilibrium can be sustained for higher communications costs than HC.

As before we can prove the proposition by checking for profitable one-period deviations at each phase assuming others will best respond to those deviations given their off-the-path beliefs. At the sanctioning stage we need to check that neither B has any incentive to deviate. If $\theta_b = 1$ then B signals in equilibrium and sees his preferred target hit and if $\theta_B = 2$ B does not signal but still sees his preferred target hit. In neither case can he update on T's type given her equilibrium strategy and so pays a cost $s$ for sanctioning but gets no gain in his expected utility for the next stage. Deviating at the punishment phase thus cannot make sense for him. Thus as long as the cost of sanctioning is low enough, neither type of B has an incentive to deviate at the punishment stage.

What about best responses off-the-path at the punishment phase? If $\theta_B = 1$ and T attacks target 2 then given off-the-path beliefs B knows they do not share preferences and so gets a value of zero in the next period from rehiring but gets $q\delta b - s$ from firing, which yields positive utility as long as $s < q\delta b$. If $\theta_B = 2$ and T attacks target 1 then given off-the-path beliefs B knows they do not share preferences and so gets a value of zero in the next period from rehiring but gets $(1-q)\delta b - s$ from firing which yields positive utility as long as $s < (1-q)\delta b$.

Turning to the attack phase, we start by noting that B's equilibrium strategy after communicating is to fire if $a \neq 1$ and after sending no signal it is to do so if $a \neq 2$. So long as LMC and UMC are met there can be no deviation profitable deviation for T given that play by B. To see this first consider the situation where B has communicated. If $\theta_T = 1$ then T gets $t + \delta t$ for attacking where told and $-r$ for deviating, which clearly cannot be profitable. If $\theta_T = 2$ then deviating yields $t - r$ while staying on the path yields $\delta t$ and that cannot be a profitable deviation either if UMC is met. Now consider the situation where B has not communicated. Given equilibrium strategies T knows B prefers target 2. If $\theta_T = 1$ then T gets $\delta t$ for attacking where told and $t - r$ for deviating, which clearly cannot be profitable if UMC is met. If $\theta_T = 2$ then deviating yields $-r$ while staying on the path yields $t + \delta t$ and that cannot be a profitable deviation.

Finally, we need to check for deviations at the communications phase. Given T's off-the-path beliefs that B likes target 2 with probability 1, her best response to a one-period deviation in which B sends $m = 2$ is clearly to attack target 2 whatever her type; remember that when B returns to the path at his next information set he will follow his equilibrium strategy of firing if T does not attack his preferred target. It clearly cannot be optimal for B to send $m = 2$ if $\theta_b = 1$ as doing so yields $-c - s + q\delta b$ which is clearly less than the $b - c + q\delta b$ he gets from staying on the path. Likewise, if $\theta_B = 2$ then sending $m = 2$ yields the same utility as sending no signal but costs more and so cannot be a profitable deviation. The interesting question is how large the cost of communicating would have to be to make it worthwhile for B to deviate to sending no message if $\theta_B = 1$. A bit of algebra shows that the largest $c$ for which B weakly prefers staying on the MBE2 path is $c^* = b + s$.

Note that the last part of the proposition follows from the fact that $b + s > c_1$ and $b + s > c_2$. Note also that MBE2 can be sustained anywhere that MBE1 can.

### 3.4 Leaderless resistance

Suppose that the costs of communications are high and that the intrinsic value T places on hitting her preferred target is low. Then there can be a leaderless resistance equilibrium in which B does not communicate, the T pool in the first period on hitting the target B is more likely to want, and then B retains T. This is analogous to the situation in which an ideological leader can publish broad guidance to coordinate prior beliefs, but eschews direct contact with operational units. Such leaders can often still impose modest costs on their agents by declaring those who deviate from the party line as apostates or non-believers, and so those subscribing to the cause have some incentive to try to attack where they think leaders prefer. Note that in such settings, sanctioning is very cheap for leaders as it amounts to simply disavowing individuals.

**Proposition 3.4**: If $t$ is small enough, the costs of communicating large enough, and the costs of sanctioning small enough, then there is a leaderless resistance equilibrium in which B does not signal, if he does T's off-the-path beliefs are that he signals honestly, all types of T attack target 1 in the first period, and B has off-the-path beliefs if he sees $a = 2$ that $q = 0$.

As before we can prove the proposition by checking for profitable one-period deviations at each phase assuming others will best respond to those deviations given off-the-path beliefs. At the sanctioning stage we need to check that neither B has any incentive to deviate. Given equilibrium play up to the punishment phase B learns nothing about T's type (they pool on attacking target 1) and thus the net value to deviating by firing T is $-s$ for both types of B and so deviating cannot be a profitable.

What about best responses off-the-path at the punishment phase? If $\theta_B = 1$ and T attacks target 2 then given off-the-path beliefs B knows they do not share preferences and so gets a value of zero in the next period from rehiring but gets $q\delta b - s$ from firing, which yields positive utility as long as $s < q\delta b$. The fact that $s$ is low enough will be critical for sustaining the equilibrium because the off-the-path best response by B when $\theta_B = 1$ is what keeps the T who prefers target 2 from deviating at the attack phase. If $\theta_B = 2$ and T attacks target 2 then given off-the-path beliefs B knows they share preferences and so gets a value of $\delta b$ in the next period from rehiring but gets $(1 - q)\delta b - s$ from firing. The B who prefer target 2 have a best response to a one period deviation by T of retaining.

Turning to the attack phase, we need to ask when the T would pool on attacking target 1 even though if $\theta_B = 2$ then B's best response is to rehire if T deviates by attacking target 2. If $\theta_T = 1$ this is an easy choice. Given B's equilibrium strategy of rehiring if $a = 1$, T gets $t + \delta t$ for staying on the path and $-r$ for deviating. If $\theta_T = 2$ then T can stay on the path and guarantee herself a payoff of $\delta t$ in the next period or she can deviate by striking her preferred target in the first stage, effectively taking a gamble that B actually prefers target 2. That cannot be a profitable deviation for any $t^* < \frac{pr}{1-\delta p}$.[10] Intuitively, the more T discounts the future and the more likely it is that B prefers target 1, the less willing B will be to risk punishment in the present period against the chance of future gains if B does indeed prefer target 2. So, if $t < t^*$ there is no profitable deviation at the attack phase.

Finally, we need to check for deviations at the communications phase. Given T's off-the-path beliefs that B signals honestly if he deviates, her best response to a one-period deviation by B is to attack where B says as doing so guarantees she will be rehired (recall B has a best response to rehire if $a = 1$ when $\theta_B = 1$ and if $a = 2$ when $\theta_B = 2$). B therefore faces a tradeoff between getting $a = 1$ on the path followed by a lottery between attacks in the next period or paying the cost of communication and guaranteeing his preferred attack in the present period. If $\theta_B = 1$ deviating is pure cost and so cannot be profitable. If $\theta_B = 2$ then taking the difference between B's utility on- and off-the-path shows that the B has no profitable deviation if $c > b$, which shows the last part of the proposition.

## 3.5 Independent agents

Finally, we can consider the situation where B's threats do nothing to motivate T because she would rather attack her favored target in the first period and accept the costs of punishment than defer to B.

**Proposition 3.5**: If $t$ is large enough there are independent agent equilibrium in which B does not signal even if the cost of doing so is arbitrarily small, T attacks her preferred target regardless of beliefs off-the-path, and B takes different actions depending on the cost of sanctioning.

There are several cases to consider. We will begin with case (1) where $s < \min\{q\delta b, (1-q)\delta b\}$ so that either type of B would want to fire if he sees an attack he does not like. We will then consider case (2) where $\min\{q\delta b, (1-q)\delta b\} < s < \max\{q\delta b, (1-q)\delta b\}$, so that only one type of B is willing to fire if he sees an attack he does not like. Finally, we will consider case (3) where $s > \max\{q\delta b, (1-q)\delta b\}$ so that neither type of B would want to fire if he sees an attack he does not like. As we will see, the $t$ required to sustain an independent agent equilibrium is decreasing as we move from case (1) to case (3).

Suppose we are in case (1). Then B's best response at the punishment stage is to fire if

[10] Comparing $t^*$ to LMC notice that whether $t^*$ is above or below LMC will depend in general on the parameters, but $t^* < LMC$ whenever LMC is met and $\delta < \frac{1}{2p}$. When $\delta > \frac{1}{2p}$ it can be the case that $t^*$ is met but LMC is not. In that case the equilibrium still works as long as $t < t^*$.

T attacks his less preferred target regardless of his type and rehire otherwise. This follows from the fact that T's choice is completely informative as to his type in this equilibrium.

Given B's best response at the punishment phase, there can be no profitable deviation for either type of T. If $\theta_T = 1$ then the expected utility from staying on the path minus that from deviating is $(p(t + \delta t) + (1 - p)(t - r)) - (p(-r) + (1 - p)(\delta t))$ which a bit of algebra shows is positive for any $p > \frac{1}{2}$. If $\theta_T = 2$ then the expected utility from staying on the path minus that from deviating is $p((t - s) + (1 - p)(t + \delta t)) - (p(\delta t) + (1 - p)s)$ which is positive as long as $t$ is larger than $\hat{t} = \frac{r(2p-1)}{1-\delta(2p-1)}$.

B's best responses off the-path also impose a restriction on $t$. For B not to be tempted to signal for some off-the-path beliefs it must be the case that even the T who knew for sure that attacking her preferred target in the first period would get her fired would still want to hit that target. That condition is just that $t > \frac{r}{1-\delta}$ so that UMC is not met. Notice that $\hat{t} \leq \frac{r}{1-\delta}$ (since $p \in (\frac{1}{2}, 1]$) which tells us that as long as the necessary off-the-path conditions are met then the on-the-path ones will be as well.

Finally, consider B's decision about whether to deviate by sending a signal as to his preferences. Since we have specified that $t > \frac{r}{1-\delta}$ there can be no impact of B's signals on T's actions, and so any signal is pure cost and thus cannot be a profitable deviation.

Now consider case (2) where for one type of B the discounted value of getting a new T does not exceed the costs of sanctioning given the low probability that T shares his preferences, and thus the B would not be willing to fire even if he sees an attack he does not like. Suppose without loss of generality that $q > \frac{1}{2}$ and that $(1 - q)\delta b < s < q\delta b$ so that B's best response at the punishment phase would be to fire a T who attacks target 2 if $\theta_B = 1$ but B will not fire a T who attacks target 1 if $\theta_B = 2$. That means that no T who attacks target 1 will be fired, but that T who attack target 2 actually risk being fired.

We need to check whether this would induce a profitable deviation for T at the attack phase such that even if $\theta_T = 2$ her best response is to attack target 1. If $\theta_T = 1$ the value of staying on the path minus the expected value of deviating is $(t + \delta t) - (p(-r) + (1 - p)\delta t)$ which is clearly positive. If $\theta_T = 2$ the value of staying on the path minus the expected value of deviating is $(p(t - s) + (1 - p)(t + \delta t)) - \delta t$ which is positive whenever $t > \frac{pr}{1-\delta p}$. This constraint will be met below UMC when $p < \frac{1}{2\delta}$ and above it otherwise. In other words when the costs of sanctioning are high enough that one type of B will be unwilling to fire, then a level of preference divergence below UMC can sustain the independent agents equilibrium.

Now consider case (3) where neither type of B would be willing to fire if they knew for certain they were facing the opposite type. In that situation, there can clearly be no profitable deviation for T at the attack stage for any $t$ as he is always rehired and thus no communications can influence her actions.

The key point from analyzing the independent agents equilibrium is that as the cost of sanctioning goes up, we go from a situation where independent agents can only be sustained for very high preference divergence, to one where it can be sustained for lower levels of preference divergence if there is sufficient uncertainty about B's preferences, to one where it can always be sustained because no B will ever find it worthwhile to sanction T.

## 3.6 Summary

In this section we explored a model in which B and T have different preferences over targets but those preferences are not necessarily symmetric so that the probability they both prefer target 1 is different than the probability they both prefer target 2. This setting adds some richness to the simpler model in section 2 but does so at the cost of substantial additional complexity. The core intuition about the conditions under which a high control equilibrium can be maintained is unchanged.

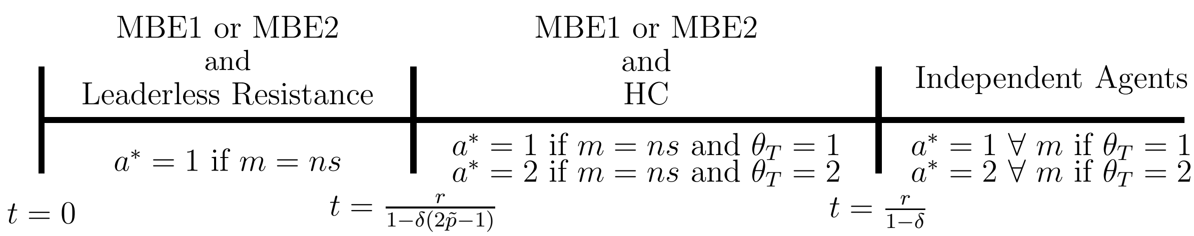


Figure 2: T's best response in the first period ($a^*$) and resulting equilibrium possibilities assuming B will fire if T attacks the target he does not like and that T's off the path beliefs if B does not signal in the HC equilibrium are that he is more likely to prefer target 1, that $\tilde{p} > \frac{1}{2}$.

To see how the equilibrium outcomes play out as preference divergence increases figure 2 plots T's best response function and the resulting equilibrium possibilities as a function of the level of preference divergence ($t$). At certain $t$ the set of equilibrium changes and where these are is a function of the costs of being sanctioned ($r$) and T's off-the-path beliefs about the probability B prefers target 1 if B's equilibrium strategy is to signal honestly ($\tilde{p}$). This is all assuming that the cost of sanctioning T is low enough for B that he would prefer to fire T at the punishment phase if he knows for sure that she does not share his preferences over targets.

As the analysis showed, the range of equilibrium similarly changes as the costs of sanctioning change for B. For low values of $s$ all the equilibrium in figure 2 are possible. As $s$ gets larger it reaches a point where one type of B no longer finds it worth it to sanction. When that happens all equilibrium in which the T pool in the first period can break down more easily as one type of T gains an incentive to deviate and hope that he has a B who prefers the same target and thus would not be willing to punish. At very high levels of $s$, where neither type of B is willing to sanction even if he knows for sure that he has a bad agent, then only the independent agents equilibrium can be sustained.

# 4 Terrorist management with a continuous target set

This section presents two models of terrorist organization using a continuous target space. These models were developed for a 2008 working paper "Bureaucracy and Control in Terrorist Organizations."

Understanding how terrorist groups will organize under different conditions is easiest if we break their organizational challenge down into two discrete questions. First, how much communication should we expect in groups when leaders and operatives have different preferences over targets and receive different information about the political impact of their actions? Second, how should these expectations depend on the ability of leaders to discipline operatives? This section addresses both questions by modeling the situation where a terrorist leader, B, has to delegate attacks to an operational terrorist, T.

We begin with a static model in which nature generates the ideal target, $\theta \in \mathbb{R}$, from some known distribution of potential targets. B observes $\theta$ and T does not. B can send a message, $m$ to T about $\theta$, but doing so increases the risk of being identified by the government and so reduces the probability that the attack succeeds. After receiving B's message, or noting the lack thereof, T decides where to attack. B and T can have different preferences over targets, so that while B's utility is maximized when T attacks $\theta$, T would rather attack $\theta - d$. In this setting, problems of credible signaling are central and lead to a number of interesting predictions about when leaders will accept the security costs entailed in communicating with their operatives.

Of course, the story does not end with communications. Terrorist leaders face the second problem that even if operatives know what leaders want, they may not do it. Historically, efforts to exercise control over terrorist operatives have often entailed sanctioning those who do not behave as leaders would like.[11] Turning to a two-period model that cleanly separates communications from sanctioning allows us to explore how leaders' credibility and reputational concerns impact group organization. Incorporating these elements also helps to illuminate the important distinction between efforts to exercise control—by providing information and possibly sanctioning operatives—and the realized level of control leaders achieve.[12] As we will see, the model makes an explicit argument about how the amount of control leaders need to exercise to get a given closeness of outcomes depends on the extent of preference divergence in a group. This fuller model also provides insight into why actually exercising control can be so problematic in terrorist organizations.

Both models are unusual from an agency theory perspective in that the information advantage is the reverse of that in standard models of delegation.[13] Here the principals are better informed than the agents about the mapping from policies into outcomes.[14] Leaders can share information to help their agents achieve better outcomes, but doing so comes at a cost. This information structure is actually quite common in more mundane organizations. Whether by dint of experience, or because they have more developed social networks, senior leaders in many organizations often know more than their subordinates do about how to respond to the environment. This is especially true in organizations like the military where managers work their way up from the lower ranks and so can drawn on more contacts and a much larger experiential base when deciding what to do given ambiguous information. Delegation is not always about who can do a better job, often it is about the exigencies of the situation, or the limited work capacity of bosses. The models thus have theoretical relevance beyond the immediate setting.

As the proofs of key results build on well-established results or simple algebra, I relegate them to section 4.4.

[11]For many specic examples, see the outstanding discussion of the mechanics of discipline in the PIRA in McIntyre (2003).

[12]By making predictions about both, this model differ from Baccara and Bar-Isaac (2006). Their models deal only with groups' information structures and do not make predictions about outcomes.

[13]For a review of models of delegation, see Bendor *et al.* (2001). These models do not fall into either of their two canonical categories.

[14]This will generally be the pattern in covert political groups where the exigencies of underground organization mean agents are ill-informed about the political impact of their actions (Bell, "Aspects of the Dragonworld", 1989a; Bell, "Revolutionary Dynamics", 1990b).

## 4.1 Different Preferences and Different Information

Play in the the static game proceeds through three steps. First nature chooses the B's optimal target $\theta$ from a symmetric, unimodal distribution $g(\theta)$ with $var(\theta) = b$. B then observes $\theta$ and can send a message $m$ to T. If B sends a message, the probability of the attack succeeding drops from $l$ to $h$. Taking B's decision into account, T attacks a target $a$ in the support of $g(\theta)$. Utilities are as follows:

$$
\begin{aligned}
U_B(\theta, a, nc) &= -u(\beta|a-\theta|)l - f(1-l) \\
U_B(\theta, a, c) &= -u(\beta|a-\theta|)h - f(1-h) \\
U_T(\theta, a, nc) &= -v(|a-(\theta-d)|)l - f(1-l) \\
U_T(\theta, a, c) &= -v(|a-(\theta-d)|)h - f(1-h)
\end{aligned}
$$

Here $u(\cdot)$ and $v(\cdot)$ are increasing functions so that B's expected utility is maximized at $a = \theta$ and T's is maximized at $a = \theta - d$. The intensity of government counterterrorism is captured in the reduced probability of success when B communicates with $h$ being the probability of success when B communicates and $l$ being the probability when B does not. Assume $h < l$ so that telling T about $\theta$ makes it more likely government will disrupt the plot.[15] Let $f$ be the cost of failure which may vary in response to political conditions or in response to leaders' risk aversion. Uncertainty over targets is captured by $var(\theta)$. The value of discrimination is reflected in $\beta \in [0, 1]$ which scales the costs to B from hitting the wrong target relative to the costs to T. The greater is $\beta$, the more costly it is for B to hit the wrong target. Preference divergence is captured by $d$.

This simple setting allows us to address three basic questions. First, how will the level of communication vary in response to changes in basic environmental parameters such as the discrimination in violence required by the group's political goals? Second, how will the level of communication vary in response to government counterterrorism? Third, how will the precision of communication vary in response to changes in the amount of preference divergence within a group?[16]

[15] We can also think of $h$ and $l$ as high and low levels of control. When B exercises low levels of control, T can only rely on her prior beliefs about what targets should be hit but the plot is less vulnerable to government disruption.

[16] All of these variables could be endogenized, though doing so would detract from our central focus. Recent studies of how the strategic interaction between a terrorist group conducting attacks and a government deciding how much to invest in counterterrorism determines the level of government counterterrorism include Bueno de Mesquita and Dickson (2007) and Siqueira and Sandler (2006).

### 4.1.1 Level of Communication

We can begin by analyzing how the level of communication varies in response to changes in key environmental variables.[17] To simplify the exposition, we'll consider absolute value loss utilities and $\theta \sim unif[0, b]$ so that larger $b$ indicate greater uncertainty about what targets will serve the political goals.[18]

A Perfect Bayesian Equilibrium in this interaction is a messaging strategy for B, $m(\theta)$, an attack strategy for T, $\alpha(m)$ and off-the-path beliefs for T. As in any such game, there are a wide variety of possible equilibrium.[19] For $d > 0$ there can be no equilibrium in which B signals honestly. Suppose B did signal honestly, setting $m(\theta) = \theta$. Then T's best response would be to play $\alpha(m) = m - d$. But if T played $\alpha(m) = m - d$, then B's best response would be to set $m(\theta) = \theta + d$, and so on. Drawing on the classic Crawford and Sobel (1982) model of cheap talk, however, we can construct an equilibrium based on cutpoints. B can communicate that $\theta$ is in a given range and the size of these ranges will depend on the extent of the preference divergence.[20]

Consider the following two cut-point messaging strategy:

$$m(\theta) = \begin{cases} r_1 & \text{if } \theta < \theta_1 \\ nc & \text{if } \theta_1 \leq \theta \leq \theta_2 \\ r_2 & \text{if } \theta_2 < \theta \end{cases}$$

As $\theta$ is drawn from a uniform distribution, T's best response function is to hit $d$ below the middle of each segment:[21]

$$\alpha(m) = \begin{cases} \frac{\theta_1}{2} - d & \text{if } m = r_1 \\ \frac{\theta_1 + \theta_2}{2} - d & \text{if } m = nc \\ \frac{\theta_2 + b}{2} - d & \text{if } m = r_2 \end{cases}$$

[17] Technically the results in this model address the probability of communication, not the level of communication. The statement that comparative statics about the probability of communication in a single interaction tell us something about the level of communication expected from a group is a common shorthand in the literature. It implicitly assumes that the interactions modeled here occur repeatedly in a world with stochastic parameters.

[18] Quadratic loss utilities yield the same results but do not admit illustrative algebraic simplifications. For $\theta \sim unif[a, b]$, $var(\theta) = \frac{(b-a)^2}{12}$ which is increasing in b.

[19] For a recent discussion of this issue and one method of solving it, see Olszewski (2006).

[20] A key finding from the Crawford and Sobel paper is that as preference divergence goes to zero, the size of these ranges shrinks to a point and communication becomes completely informative.

[21] For a generic distribution and symmetric loss function T's best response would be to hit $d$ below the expected value of each segment.

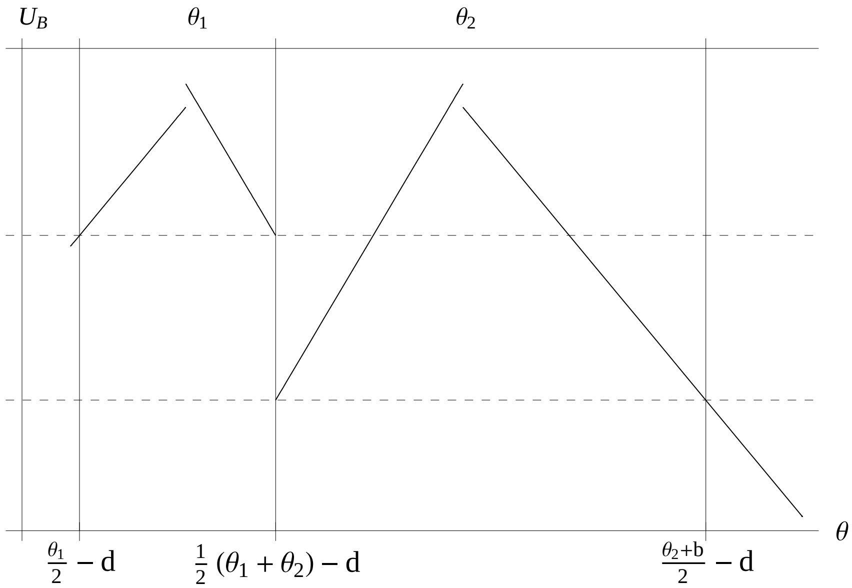


Figure 3: Cutpoint Equilibrium

If T sees any $m \notin \{r_1, nc, r_2\}$ then it believes B is one of the three types, which one does not matter. Since our equilibrium concept places no restrictions on off-the-path beliefs this weak restriction is fine.

At $\theta_1$, the terrorist boss must be indifferent between: (1) paying the security cost of signaling and having T hit $d$ below the mid-point of the lower section; and (2) paying no security cost and having T hit $d$ below the mid-point of the no-communications region. Likewise at $\theta_2$ B must be indifferent between: (1) paying the security cost of signaling and having T hit $d$ below the mid-point of the upper section; and (2) paying no security cost and having T hit $d$ below the mid-point of the no-communications region. Figure 1 illustrates B's utility as a function of $\theta$ and the cutpoints.

These two indifference conditions describe a system of two equations in two unknowns which we can solve for $\theta_1$ and $\theta_2$, providing an explicit expression of the distance between them.[22] Noting that the probability of communication is decreasing in this distance leads to the following propositions:

**Proposition 4.1:** There are at least two cutpoints whenever preference divergence is sufficiently small given the level of uncertainty about the optimal targets. The range of parameters for which a two-cutpoint equilibrium exists increases as the value of discrimination increases and as uncertainty increases. The range shrinks as preference divergence decreases

[22]This approach essentially modifies that presented in Osborne (2003).

and as the relative reduction in success probability for communicating gets smaller.

**Proposition 4.2:** When a two-cutpoint equilibrium exists, the probability of communication is decreasing in the cost of failure and in the relative reduction in the success probability when B communicates.[23] The probability of communication is increasing in the level of discrimination required by political goals. The probability of communication is increasing in the uncertainty of the operational environment provided the costs of failure are not too low. The level of preference divergence has no effect on the distance between the cut-points in a two-cutpoint equilibrium.

Two characteristics of the static game merit greater attention: the impact of counterterrorism and the relationship between preference divergence and the precision of communications.

### 4.1.2 The impact of government counterterrorism

The marginal impact of an increase in government counterterrorism on the probability an attack succeeds depends on both the specific nature of that increase and on whether B communicates. Target hardening, for example, can be expect to have a similar impact on both $l$ and $h$ while increased spending on communications intercepts would have a greater impact on $h$. Whether such changes in government strategy increase or decrease communication depends on the relationship between the costs of uncertainty and the costs of failure, on whether success is more important than precision. If it is, then leaders would rather have a high probability of success against the wrong target ($b < \frac{6f}{\beta}$) Otherwise a lower probability of success against the right target ($b > \frac{6f}{\beta}$) is preferred. Proposition 4.3 formalizes the relationship between government counterterrorism, the level of uncertainty, and the costs of failure for a two-cutpoint equilibrium.

**Proposition 4.3:** When $\frac{h'}{h} < \frac{l'}{l}$ so that the proportional decrease in the success probability when B communicates is larger than the proportional decrease in the no-communication success probability, then the probability of communications is decreasing in counterterrorism if $b < \frac{6f}{\beta}$ and increasing otherwise. When $\frac{h'}{h} > \frac{l'}{l}$ the comparative statics are reversed.

These results demonstrate an interesting relationship between uncertainty over targets – here the range of possible targets – and the impact of government counterterrorism. The intuition behind them is that when $b$ gets larger for a fixed number of cutpoints, the sizes of both the no-communication and communication regions increase. Both changes mean that

[23]Another way to state this is that the size of the no communication region is increasing.

the relative gains to communicating shrink. $b = \frac{6f}{\beta}$ effectively marks the point at which the world is sufficiently certain that the value of communicating outweighs the cost of failure. Once $b$ drops below this point, we find the intuitive relationship that government actions which lower T's success probability proportionally more when T communicates than when he does not make T less likely to communicate. The key point to emphasize is that when the effects of increased government counterterrorism depend on whether B communicates, increased counterterrorism can actually lead to increased communications.

### 4.1.3 Preference Divergence and the Precision of Communication

Despite the fact that there is no relationship between $d$ and the size of the no-communication region in a two cut-point equilibrium, there is a sense in which the level of control gets larger as $d$ gets smaller. If $d$ is sufficiently small, then B can partition the signaling spaces below $\theta_1$ and above $\theta_2$ into more than one region, leading to attacks that are closer to $\theta$ in expectation. Doing so will also decrease the size of the no-communication region. The following proposition formalizes this insight.

**Proposition 4.4:** As $d$ gets small, the number of feasible partitions of the communication regions is weakly increasing and the size of the no-communication region is weakly decreasing. Both the probability of communication and the precision of communication are thus weakly increasing as $d$ gets small.

This result leads naturally to the following remark.

**Remark 4.1:** B's *ex ante* expected utility from the *ex ante* pareto-optimal equilibrium is increasing as $d$ gets small.[24]

There are two reasons for this result. First, B does better for a fixed number of partitions with smaller $d$ because T's attacks will be closer to what B would like. Second, B's utility is increasing in the number of partitions of the communications regions. The intuition for this result is that B pays no additional cost for sending more distinct messages in the signaling regions and receives strictly greater payoffs in those regions by doing so.

This result on precision has strong implications for the two-period game. There B's ability to credible threaten T, and so sustain an honest equilibrium in the first period, depends critically on B's expected utility from the second period.

[24] Crawford and Sobel (1982) prove this result for any fixed interval. See section 4.4 for a proof in this context.

## 4.2 Discipline, Communication, and Control

So far we have discussed how preference divergence between B and T limits B's ability to communicate, thereby lowering the expected level of communication. While the static model provided a rich set of predictions on when and how effectively terrorist leaders will communicate, it left two key questions unanswered: (1) Under what conditions will operatives be motivated by threats; and (2) Under what conditions will leaders' threats be credible?

Terrorists' own observations about command and control suggest there is an interesting interaction between B's ability to punish T for not taking the desired action, the extent to which B will communicate, and the level of control leaders exercise. When asked to explain why the leadership of his group did not do more to restrain wayward operatives, former UVF bomb-maker David Ervine responded "in a military organization the admiral doesn't have to worry about the sailor getting off watch and shooting him. My admiral did have that concern."[25] Ervine's observation suggests that if we want to understand how terrorist organize, it is critical to understand their disciplinary challenge.

To do so, consider a two-period extension of the model presented in section 2.1. After observing T's attack, B can sanction T for poor performance in the first period and thereby motivate better performance. For B's decision about whether to sanction B to depend on T's actions in the first period—and thereby make it possible for the threat of being sanctioned to motivate T—B must be uncertain as to the level of preference divergence. To keep things tractable restrict T to three types: 'good' ($d_g = 0$), 'bad' ($d_b > 0$), and 'acceptable' ($0 < d_a < d_b$). The players' utility functions for each period are the same as in the previous model with second-period payoffs discounted by a common discount factor $\delta$.

Play begins when nature chooses $\mathrm{T}_1$'s type. Both players share the belief that $T_1$ is bad with probability $q$ and good with probability $(1-q)$. The stage game then proceeds as before. Having observed $\mathrm{T}_1$'s choice of attacks, B decides whether to sanction $\mathrm{T}_1$, paying cost $s$ if it does. $\mathrm{T}_1$ suffers disutility $w$ from being sanctioned and is not rehired. If B sanctions it provides information to potential recruits about its preferences, so that B is guaranteed an acceptable agent in the second round.[26] In the next round B's preferences remain the same, and the stage game is repeated.

A Subgame Perfect Bayesian Equilibrium in this repeated interaction has four compo-

[25] Author interview, David Ervine, March 8, 2006.

[26] We can think of this as B establishing a reputation for toughness, or as the pool of T's learning something about B and so those who don't agree leave, rendering the remaining potential recruits more likely to share B's preferences.

nents: (1) a messaging strategy for B in every period conditional on B's prior ($q$) and posterior ($q_1(\cdot)$) beliefs about T, $m(\theta^1, \theta^2, s(a_1), q, q_1(a_1))$, (2) an attack strategy for every T in every period conditional on B's sanctioning strategy, $\alpha_t^{T_i}(m_1, m_2, s(a_1))$; (3) a sanctioning strategy for B after seeing $T_1$'s attack in the first period, $s(a_1)$; (4) and off-the-path beliefs for B and T at both stages if applicable. Importantly, the core results from the previous section carry through when we introduce uncertainty about $T_1$'s type. B simply takes this uncertainty into account when generating it's messaging strategy by integrating over the distribution of possible types.[27]

In this setting B's communication strategy clearly depends on its ability to force bad T to play as though they are good in the first period. B's ability to do so depends on the set of credible sanctioning strategies, which in turn depend on B's expected outcome in the final period. Given this complexity, one way to get at the relationship between the level of communication and B's disciplinary capacity is to focus on the conditions under which an honest equilibrium can be sustained in the first period.

Because sending messages is costly, an honest equilibrium is one in which B sends an honest message, $m_1(\theta, q) = \theta$ when $\theta$ is sufficiently far from $E(\theta)$ that the reduced probability of success from communicating is less costly than having T rely on its prior beliefs. T then attacks exactly where B instructs or at $E(\theta)$ if no message is received—$\alpha(m_1) = m_1$ if T gets a message and $\alpha(nc) = E(\theta)$ otherwise. An honest equilibrium in the first round thus achieves the maximal level of communication in that the no-communications region shrinks to what it would be if B and T had identical preferences over targets. Formally an honest equilibrium entails the following messaging strategy for B in the first period:

$$m_1(\theta) = \begin{cases} \theta & \text{if } \theta < E(\theta) - \frac{(l-h)f}{\beta l} \\ \theta & \text{if } \theta > E(\theta) + \frac{(l-h)f}{\beta l} \\ nc & \text{otherwise} \end{cases}$$

For T's best response function in sustain an honest equilibrium B's sanctioning strategy must guarantee there is no profitable deviation for T. Since the bad T would clearly prefer to deviate from this strategy in any one-shot interaction, two constraints must be met to sustain an honest equilibrium in the first round:

1. Motivational Constraint (MC): $T_1$ must prefer the honest equilibrium over deviating,

[27] B's utility from any message is a strictly concave function of T's type. Because any linear combination of strictly concave functions is a strictly concave function, the method of establishing $n$ cutpoints by solving a system of $n$ indifference conditions in $n$ unkowns still establishes B's equilibrium messaging strategy.

being sanctioned, and losing the discounted value of second-period game.

2. Credibility Constraint (CC): $T_1$ must believe B will sanction if he deviates from the honest equilibrium.

We can rewrite the constraints more formally using some additional notation. $\alpha(m,d)$ indicates the optimal attack strategy given $d$. $m_t(\theta,q)$ is the optimal messaging strategy in period $t$ given B's prior beliefs over T's type. $m_t(\theta,d_i)$ is the optimal messaging strategy if B believes T to be of type $i \in \{g,b,a\}$. $m_t(\theta,q_1(\alpha_1,m_1))$ is B's optimal messaging strategy given posterior beliefs about T's type. The motivational and credibility constraints amount to the following two equations which respectively define $w^*$, the smallest $w$ for which MC holds and $s^*$, the largest $s$ for which CC holds:

$$\begin{aligned} w^* &\geq EU_1^{T_1}(\alpha_1(m_1(\theta,d_g),d_b)) - EU_1^{T_1}(\alpha_1(m_1(\theta,d_g),d_g)) - (1-\delta)EU_2^{T_1}(\alpha(m_2(\theta,q)),d_b), \\ \frac{s^*}{1-\delta} &\leq EU_2^B(m_2(\theta,d_a)) - EU_2^B(m_2(\theta,q_1(m_1,\alpha_1))). \end{aligned}$$

Carefully inspecting these equations makes it immediately clear that the constraints interact to produce the outcomes we care about. We will therefore work through the logic of the results before discussing how changes in the core parameters of the model impact the likelihood of getting different outcomes.

The motivational constraint depends on three factors: B's messaging strategy, the signal T receives in the first round, and B's sanctioning strategy. The intuition for this dependence is clear from considering the situation where B plays as though he faces a good agent in the first round. If B signals, then T knows $\theta$ exactly and so gets the maximum gain in the present period from defecting. If B does not signal—because $\theta$ is close to $E(\theta)$—then defecting from an honest equilibrium in the present period is not as good for T as he cannot perfectly absorb the impact of uncertainty over targets. This illustrates a more general principle, the bad T's value from defecting is greater when B's message is more precise. This principle means that the motivational constraint will be met for lower $w$ in the no-communication region. We can therefore identify three motivational regimes: (1) the motivational constraint is met in both regions; (2) the motivational constraint is only met in the no-communication region; and (3) the motivational constraint is never met. Because the last term in the MC equation captures the value to T of being rehired, we need to understand both B's sanctioning strategy and his second-period messaging strategy to fully analyze how the MC works.

The CC is a bit simpler. It is essentially the value to B of guaranteeing himself an acceptable agent in the next round and so clearly depends on whether T's actions are informative as to her type. If the Ts separate, then it is easier to meet the credibility constraint as the expected value of the one-shot game with a bad T is lower than the expected value of the game where T faces a lottery between a bad T and a good T. Lemma 1 formalizes this intuition.

**Lemma 4.1:** The highest cost of sanctioning, $s^*$, for which the credibility constraint is met is greater when the T play true to type the when the Ts pool.

We can thus identify three credibility regimes: (1) the credibility constraint is never met; (2) the credibility constraint is met only when B knows it is facing a bad T; and (3) the credibility constraint is weak enough the B will fire if it remains uncertain as to T's type.

The three regimes each for two constraints lead to 9 possible equilibrium outcomes in pure strategies. Rather than discuss each regime, we can focus on three key qualitative features of the interaction between MC and CC.[28].The first is that whether or not the MC is met depends critically on the stochastic signal B receives about the world. The second is that communication is maximized—in the sense of being most likely in the first period—when B can effectively threaten T, but when doing so is sufficiently costly that B will only fire if he is certain he is facing a bad T. If firing is too inexpensive for B, then the incentives for bad T to play honestly in the first round disappear and an honest equilibrium cannot be sustained. The third is that realized discipline occurs as the interaction of a random event—nature's choice of the ideal target—and the players' strategic incentives. This means that observing disciplinary actions in group X, but not in group Y, provides limited evidence of greater preference divergence in group X. It may simply be that leaders in group Y wield a sufficient threat that their agents toe the line.

With these results in mind, we turn to closer inspection of the two constraints to see how the chances of being in the different regimes vary with the key independent variables.

### 4.2.1 Motivation

The first condition that must be met for bad T to be willing to follow an honest equilibrium is that deviation not be profitable. This motivational constraint is met when the combined costs of being sanctioned and losing the benefits of the second-round outweigh the present gains from defection. T's first period payoff to defecting from an honest equilibrium clearly

[28]Table A1 in the section 4.4 describes the equilibrium outcomes for each constraint combination under the assumption that B has off-the-path beliefs which ensure that if the Ts pool, they do so on $\alpha(m, d_g)$. This is not a strong restriction. Good T would always benefit by deviating from pooling on $\alpha(m, d_b)$ as long as B's off the path beliefs do not lead B to fire T upon getting an unexpectedly good outcome in the first period.

depends on whether or not B signals. If B sends an honest message, T can attack its own ideal point. If B sends no message, T's utility is maximized by attacking $d$ below the expected value of the no-communication region. We will consider each possibility in turn.

If B communicates honestly, T knows $\theta$ and the MC reduces to

$$w^* \geq ld - (1-\delta)EU_2^{T_1}(\alpha(m_2(\theta, q)), d_b).$$

For all the parameters except for the level of preference divergence, taking comparative statics on this amounts to taking comparative statics on the bad T's continuation value in a pooling equilibrium. Doing so leads to the following proposition.

**Proposition 4.5:** If B communicates in the first period in a two-cutpoint equilibrium, $w^*$ is increasing in the cost of failure, the uncertainty of the operational environment, and the level of preference divergence for the bad T. $w^*$ is decreasing in the level of discrimination required by political goals.[29]

If B does not communicate, then the situation is a bit more complicated. Because B's messaging strategy in an honest equilibrium neatly bounds the no-communication region we can explicitly solve for T's expected utility from defection. Doing so shows that T's utility from defecting in the first period responds differently to changes in the parameters than its utility in the second period. For example, the first-period value to the bad T of playing true to type is decreasing in $f$ while the second period expected utility given the T pooled is increasing in $f$. Proposition 6 formalizes how the consequences of changes in the cost of failure on $w^*$ depend on other parameters.

**Proposition 4.6:** If B does not communicate in the first period in a two-cutpoint equilibrium, $w^*$ is increasing in the uncertainty of the operational environment and the level of preference divergence. $w^*$ is increasing in the cost of failure as long as the discount factor and level of preference divergence are not too large. $w^*$ is decreasing in the level of discrimination required by political goals as long as the discount factor and level of preference divergence are not too large.

Regardless of whether B communicates, the MC is easier to meet as preference divergence and uncertainty drop.

[29] $w^*$ is insensitive to changes in the probability of getting a bad T because with linear utility functions changes in $q$ do not impact T's second-period utility.

### 4.2.2 Credibility

The second condition that must be met for $T_1$ to be willing to follow the honest equilibrium is that he must believe that B will sanction any first-period deviation. Such a threat is only credible if the discounted gains from getting an acceptable agent in the second round exceed the present cost. Notice that first period outcomes have no impact on this constraint. Because no sanctioning is credible in the second round, the discounted gains amount to the difference in B's expected utility if he gets the better agent in the second round. The credibility constraint thus depends on whether the T pool. We can thus identify two constraints

$$\begin{aligned}\frac{s^*_{sep}}{1-\delta} &\leq EU_2^B(m_2(\theta, d_a)) - EU_2^B(m_2(\theta, d_b)),\\ \frac{s^*_{pool}}{1-\delta} &\leq EU_2^B(m_2(\theta, d_a)) - EU_2^B(m_2(\theta, q)).\end{aligned}$$

Using the implicit solutions for a two-cutpoint equilibrium we can identify how the probability of meeting the constraint varies with the parameters and thereby identify the minimum reputational gain, $d_b - d_a$, for which the credibility constraint is met in a two-cutpoint equilibrium. Lemma 1 shows that $s^*_{sep} < s^*_{pool}$ and so we focus on the former. Because B's signaling strategy depends on the type it is facing, the exact difference in expected utilities is quite complex. What we can do is sign the cross-partial derivatives of B's expected utility with respect to preference divergence and the parameters. Where the cross partial is positive, the gains from getting an acceptable agent are increasing in the parameter, meaning the minimum reputation gain required to meet CC is decreasing in the parameter. Taking this approach yields the following proposition.

**Proposition 4.7:** For a two-cutpoint equilibrium, the minimum reputational gain for which B can credibly threaten to sanction T for deviating from an honest equilibrium in the first round is in decreasing in uncertainty. The minimum reputational gain is also decreasing in the level of discrimination required to meet political goals as long as uncertainty is sufficiently large. The greater the level of preference divergence in first period, the smaller the reputational gain required to make sanctioning credible.

These results make intuitive sense. The value to B of having T attack close to $\theta$ is obviously increasing in uncertainty and discrimination. This means the value of getting an acceptable agent in the second period is increasing in uncertainty and discrimination. Getting a slightly better agent yields higher returns the worse the first agent because B's expected utility for the second period is strictly concave in $d$.

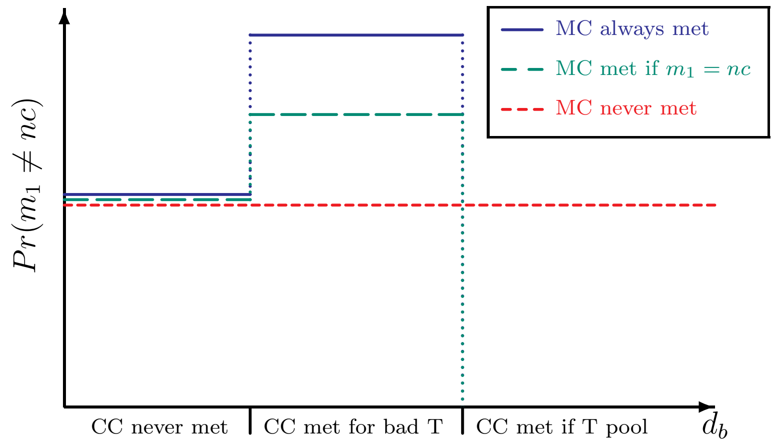

Figure 4: Probability of Communication

Before moving on, we should not that changes in the precision of communication can have outsized impacts on the expected utility of the second round. This means that if $d_b$ is close to the level that would allow a finer partition of the communication regions, then reputational gains that would be too small to meet the CC in a two-cutpoint equilibrium may suffice to make sanctioning credible. As a practical matter, this finding once again suggests the link between latent preference divergence and observed discipline and is quite complex.

### 4.2.3 Outcomes of Two-Period Model

Combining the results from the previous two sub-sections shows how the outcomes of the game in a two-cutpoint equilibrium vary depending on which combination of constraints hold. Figure 4.2.3 illustrates how the probability of communication varies by regime. Figure 4.2.3 shows how the *ex ante* expected level of control in the first period, $-E(|a^* - \theta|)$, does so. In both figures the x-axis measures preference divergence between B and the bad T, $d_b$. Moving right on the x-axis thus equates to increased reputational gain, $d_b - d_a$, for B from firing a bad T. As proposition 7 shows, the larger this gain, the more likely it is that the credibility constraint is met. Each figure has three lines representing the three possible regimes on the MC.[30]

Two features of the outcomes with respect to the level of communication merit further explanation. First, the impact of preference divergence on outcomes is distinctly non-

[30]It is important to keep in mind that an increase in $d_b$ can make it harder to meet the MC. To focus in on the interaction between the constraints, these figures assume the changes in $d_b$ over the range of the figure are insufficient to change the MC regime. Instead the different MC regimes pictured here depend on B's ability to sanction, $w$.

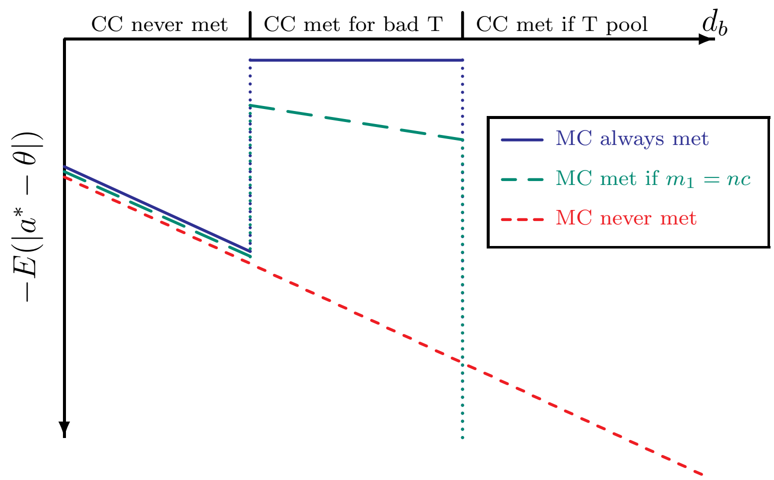

Figure 5: Expected Level of Control

monotonic. Starting from a point at which preferences are sufficiently aligned that CC is never met, the probability of communication is insensitive to the level of preference divergence until it becomes sufficiently large that B would actually want to fire a bad T.[31] At this point B's threat of sanctioning T becomes relevant and the MC can begin to bite. If the MC is met, the probability of communication increases suddenly as B is able to motivate either a pooling equilibrium—if the MC is met regardless of B's message—or a semi-pooling equilibirum—when MC is met only if $m_1 = nc$. As we should expect, communication is maximized when the MC is met regardless of the message from B.

Second, when it becomes too easy for B to punish T, when $d_b$ is large enough that the CC is met even if the T pool, then the there is no informative pure-strategy equilibrium. This is a standard result. If the T pool on $\alpha(m, d_g)$ they are fired, so the bad T would play true to type. But knowing this, B would not fire on seeing a T play $\alpha(m, d_g)$. But then the bad T would want to dissemble if the MC bites. And so on. One way to interpret this result is that being able to threaten helps T, but when B wields too great a threat credible communication becomes impossible and internal communication breaks down.

Turning to the outcomes with respect to control, the key results mirror those for the probability of communication. Starting from low levels of preference divergence, the expected level of control is decreasing until preference divergence becomes large enough that the CC is met for bad T. At that point if the MC is met, B is able to motivate bad T to play as good T, increasing the level of control. If the MC is met regardless of the signal B sends then the level of control becomes insensitive to preference divergence for some time as B can motivate

[31]Recall that both cutpoints shift by the same amount as $d$ changes and so the probability that $m_1 \neq nc$ does not change.

all T to play $\alpha(m, d_g)$ in the first period. However, if the MC is only met when $m_1 = nc$, then the level of control will be decreasing in $d_b$ in this region, albeit at a slower pace than in the case where MC is never met. Once the level of preference divergence with the bad T becomes so large that the CC is met when the T pool, cooperation within the group breaks down if the MC is met. If the MC is never met, then the CC simply does not matter for a linear utility function.

Pulling these results together suggests an interesting paradox: groups that suffer such high levels of preference divergence that the CC will be met even when the T pool may be better off with limited disciplinary capacity. At least when the MC is never met, informative equilibria are possible with such high levels of preference divergence. As in many organizations, leaders in terrorist groups appear to do best when they can punish their agents but cannot do so too easily.

## 4.3 Common Results

Certain common threads run through both models in this section. The first is that the costliness of hitting the wrong targets is a key factor. The greater the discrimination required by a group's political goals, the more control leaders will exert. In the models in sections 4.1 and 4.2 this meant that the terrorist boss was more likely to communicate when hitting the wrong targets was costly. This suggests that giving groups a stake in the political process so that they become concerned with maintaining popular support can have two beneficial effects. First, it can cause them to be more discriminate in their attacks; perhaps by focusing on legitimate military targets, thereby sparing civilians lives. Second, it can push them to accept lower levels of security, rendering government's counterterrorism job easier.

The second is that leaders will exercise more control when there is more uncertainty about how specific targets or methods of attack will support political goals. This result should not be surprising, but put more generally it has real bite. The greater the information advantage of the principals, the more likely they are to exert control.[32]

Finally, we saw that greater preference divergence leads to decreased communication for two reasons. In the static model greater preference divergence makes it harder to credibly signal about the nature of the world. The less B and T agree, the more B will try to mislead T when he does signal. This means the amount of information T can take from the signal

[32] Additionally, in the face of divergent preference, greater uncertainty can facilitate more precise communication. Greater uncertainty does so discontinuously by allowing B to squeeze in more partitions of the type space.

is relatively small. In the dynamic model, greater preference divergence meant B needed to wield a greater threat to compel T to behave as he would like. In general, greater preference divergence can be thought of as raising the price of exerting control. Taken together, these results mean that greater preference divergence will lead to less communication and control.

## 4.4 Proofs

This document provides proofs of the propositions and lemmas in section 4.

**Proof of Proposition 4.1:** The indifference conditions lead to the following system of two equations in two unknowns:

$$-h\beta(\theta_1 - (\frac{\theta_1}{2} - d)) - (1-h)f = -l\beta((\frac{(\theta_1+\theta_2)}{2} - d) - \theta_1) - (1-l)f \tag{3}$$
$$-h\beta(\frac{(\theta_2+b)}{2} - d) - (\theta_2)) - (1-h)f = -l\beta(\theta_2 - (\frac{(\theta_1+\theta_2)}{2} - d) - (1-l)f \tag{4}$$

Solving these yields the following which define the interior solution with two cutpoints.

$$\theta_1 = \frac{2fh(h-l) + (bhl - 2d(h+l)(h+2l))\beta}{h(h+2l)\beta} \tag{5}$$
$$\theta_2 = \frac{2fh(l-h) + ((l+h)(bh - 2d(h+2l))\beta}{h(h+2l)\beta} \tag{6}$$

Jointly solving (3) and (4) and checking for when $0 < \theta_1 < \theta_2 < b$ yields the following restrictions on the existence of a two-cutpoint equilibrium.

$$b > \frac{2f(l-h)}{l\beta} \tag{7}$$
$$d < \frac{h(2f(h-l) + bl\beta}{2(h+l)(h+2l)\beta} \tag{8}$$

The proposition follows immediately from taking comparative statics on the restrictions.
■

**Proof of Proposition 4.2:** Subtracting $\theta_1$ from $\theta_2$ yields:

$$\frac{4f(l-h) + bh\beta}{(h+2l)\beta}. \tag{9}$$

All results except those on $b$ follow directly from taking comparative statics on this difference. The probability of communication is the cumulative density of the no-communication region which is just $\frac{1-(\theta_2-\theta_1)}{b}$. The results on $b$ follow from taking the derivative of that density with respect to $b$ and imposing the constraints for a two-cutpoint equilibrium to exist plus the additional constraint that $f > \frac{(h+2l)\beta}{4(l-h)}$.

■

**Proof of Proposition 4.3:** Let $h$ and $l$ be functions of the level of government counterterrorism, $x$. Define $\theta_{dif} = \theta_2 - \theta_1$. Applying the quotient rule to (15) and noting that the denominator is strictly positive gives us $\frac{\partial(\theta_{dif})}{\partial x} < 0$ so long as $2(h'l - hl')\beta(b\beta - 6f) < 0$. If uncertainty is small relative to the cost of failure but still large enough so that a two-cutpoint equilibrium exists, if $\frac{2f(l-h)}{l\beta} < b < \frac{6f}{\beta}$ then we find the familiar result that communication is increasing in pressure if the elasticity of $h$ is greater than that of $l$:[33]

$$\frac{\partial(\theta_{dif})}{\partial x} < 0 \quad \text{if} \quad \frac{h'}{h} > \frac{l'}{l}; \tag{10}$$

$$\frac{\partial(\theta_{dif})}{\partial x} > 0 \quad \text{if} \quad \frac{h'}{h} < \frac{l'}{l}. \tag{11}$$

$$\tag{12}$$

However, if uncertainty is sufficiently large relative to the cost of failure, if $b > \frac{6f}{\beta}$, then the relationship is reversed so that now communications are decreasing if the elasticity of $h$ is greater than that of $l$.

$$\frac{\partial(\theta_{dif})}{\partial x} > 0 \quad \text{if} \quad \frac{h'}{h} > \frac{l'}{l}; \tag{13}$$

$$\frac{\partial(\theta_{dif})}{\partial x} < 0 \quad \text{if} \quad \frac{h'}{h} < \frac{l'}{l}. \tag{14}$$

■

**Proof of Proposition 4.4:** For each cutpoint B must be indifferent between the mid-point of the region below the cut-point and the mid-point of the region above the cutpoint. Let $\theta_k$ be the highest cut-point below the no communications region and $\theta_{k+1}$ be the lowest cutpoint above that region. Suppose there are $T$ intervals. It must be that $U_B(\alpha(r_t), \theta_t) = U_B(\alpha(r_{t+1}), \theta_t)\ \forall\ t < k$ and $\forall\ t > k$. For $t = k$, we must have $U_B(\alpha(r_t), \theta_t) = U_B(\alpha(nc), \theta_t)$

[33] The elasticities are negative so this means that $l$ is decreasing proportionally faster in $x$.

and $U_B(\alpha(nc), \theta_{t+1}) = U_B(\alpha(r_{t+1}), \theta_{t+1})$. These indifference conditions lead to a system of $T$ equations in $T$ unknowns which can be solved.

However, because B's utility in the no communication interval is shifted up and has a steeper slope than it does in the communication regions, this system of equations cannot be simplified to a second-order difference equation.[34] Instead, consider arbitrary cutoffs $\underline{\theta}$ and $\overline{\theta}$ which bound the no communication region. In each communication region, B's indifference implies that $\theta_t + d = \frac{1}{2}[\frac{1}{2}(\theta_{t-1} + \theta_t) + \frac{1}{2}(\theta_t + \theta_{t+1})]$.[35] Since the intervals must add up to $\underline{\theta}$ in the lower region we get that for $\underline{T}$ intervals in the lower region

$$\underline{T}\theta_1 + 2d\underline{T}(\underline{T} - 1) = \underline{\theta}. \tag{15}$$

So when $d < \frac{\underline{\theta}}{2\underline{T}(\underline{T}-1)}$ there is some $\theta_1 > 0$ such that (13) is satisfied. Clearly for smaller $d$, larger $\underline{T}$ are possible. Notice also that as $\underline{\theta}$ increases, more partitions may become possible. In the two cut-point equilibrium, the lower cut-point is increasing in $b$. So as $b$ increases, at some point more partitions become possible. At the moment this happens, there is a discontinuous increase in $\underline{\theta}$ which must now make B indifferent between the no-communication region and something higher than the mid-point of the original lower communication region.

The upper region is somewhat more complicated as the lower bound on that region is not fixed, but depends on the number of partitions in the lower region. Clearly for any fixed $\overline{\theta}$ the number of possible partitions will be weakly increasing as $d$ gets small.[36] However, the lower bound on the upper communication region depends on the mid-point of the no-communication region. This in turn depends on the number of segments in the lower communication region. In fact, the mid-point of the no-communication region is increasing in the number of segments in the lower communication region. Thus in order to discuss the impact of $d$, we need to be able to say what happens to the upper bound on the no-communication region, $\theta_{k+1}$ as $d$ gets small.

To show that the no-communications region is shrinking as $d$ gets small, first consider the case where there are no further partitions above $\theta_{k+1}$. Then solving the indifference condition for $\theta_{k+1}$ and taking the first derivative of that with respect to $\theta_k$ yields $\frac{l}{h+l} > 0$. We take the derivative with respect to $\theta_k$ because we want to show how the size of the

[34]The slope in the internal region is steeper because $l > h$, meaning the probability of paying the cost for a bad attack is greater. The utility is shifted up because $(1-h)f > (1-l)f$ meaning the probability of paying the failure cost is lower in the no-communication region.

[35]This proof essentially follows the same lines as the discussion of the Crawford and Sobel model in Osborne, 347-8.

[36]Also, as long as there are two segments in the upper communication region, further subdivisions depend only on making B indifferent between these, so that $\theta_{k+2} + d = \frac{1}{2}[\frac{1}{2}(\theta_{k+1} + \theta_{k+2}) + \frac{1}{2}(\theta_{k+2} + \theta_{k+3})]$.

no-communication region changes for any arbitrary lower bound on that region. While this derivative is positive, it is strictly less than 1 so the size of the no-communication region will be shrinking in $d$ if $\theta_{k+1}$ is the highest partition.[37]

Now consider the case where there is a second cutpoint $\theta_{k+2}$ which is also going to increasing in $\theta_k$ as it is influenced by $\theta_{k+1}$. Solving the indifference condition for $\theta_{k+1}$ and taking the derivative with respect to $\theta_k$ yields $\frac{l+h\frac{d\theta_{k+2}}{d\theta_k}}{h+l}$. The indifference condition for $\theta_{k+2}$ gives us that $\frac{d\theta_{k+2}}{d\theta_k} = \frac{1}{2}\frac{d\theta_{k+1}}{d\theta_k}$. Plugging this back in and simplifying yields $\frac{d\theta_{k+1}}{d\theta_k} = \frac{2l}{2l+h}$ which is less than one, showing that for 1 further partition the range of the no communications region is shrinking as $d$ gets small.

For $n$ partitions of the upper sample space we mu still have $\frac{d\theta_{k+1}}{d\theta_k} = \frac{l+h\frac{d\theta_{k+2}}{d\theta_k}}{h+l}$. The indifference condition for further partitions provides that for $T$ partitions we get $\frac{d\theta_{k+2}}{d\theta_k} = \sum_{i=1}^{T-1}\frac{1}{2^n}\frac{d\theta_{k+1}}{d\theta_k} = \frac{2^n-1}{2^n}\frac{d\theta_{k+1}}{d\theta_k}$. Plugging this into the previous equation we get $\frac{d\theta_{k+1}}{d\theta_k} = \frac{l}{l+h(1-\frac{2^n-1}{2^n})}$ which proves the result since $\frac{2^n-1}{2^n} < 1$.

■

**Proof of Remark 4.1:** The statement that B's 's expected utility is increasing as $d$ gets small follows from two facts: (1) for a fixed number of partitions, B's expected utility is increasing as $d$ gets small; and (2) B's expected utility is increasing for changes in $d$ that allow more partitions. Consider each in turn.

For a fixed number of partitions, first note that the size of the no-communication region is unchanging. To see this let $\theta_k$ be the cutpoint bounding the no communication region from below and $\theta_{k+1}$ bound that region from above. Using the indifference conditions to solve for these in terms of $\theta_{k-1}, \theta_{k+2}$, and the parameters yields $\theta_{k+1} - \theta_k = \frac{4f(l-h)+h\beta(\theta_{k+2}-\theta_{k-1})}{\beta(h+2l)}$. Applying the condition from the previous proof that $\theta_t = \frac{1}{2}[\frac{1}{2}(\theta_{t-1}+\theta_t) + \frac{1}{2}(\theta_t+\theta_{t+1})] - d$ for $t < k$ and for $t > k+1$ shows that any two cutpoints within the communications regions increase by the same amount as $d$ gets small. Next note that B's expected utility for a fixed-size no-communication region, $\int_{\theta_k}^{\theta_{k+1}}\frac{1}{\theta_{k+1}-\theta_k}(-l\beta(|(\frac{\theta_k+\theta_{k+1}}{2} - d) - \theta) - (1-l)f)d\theta$ is strictly increasing as $d$ gets small. Finally, note that as $d$ gets small, all the cutpoints shift up by the same amount. The contribution of all fixed-size regions to B's expected utility is clearly increasing. The only complication arises from the fact that the lowest and smallest partitions gets larger as $d$ gets small while the upper and largest partition gets smaller. This

[37]A one-unit increase in the lower bound of the no-communication region yields a change in the upper bound of that region that is less than one unit, so the size of the region shrinks as the lower bound moves up.

change must be utility enhancing since B's expected utility over any portion of a partition is decreasing in the distance between that portion and the midpoint of the partition. Figure 1 provides the graphic intuition for this point. As the cutpoints shift right, B trades the low-utility region on the right for the high utility region on the left.

Now consider the case where the change in $d$ allows for more partitions. Geanakopolos (2006, 1466-1467) proves that in any single player Bayesian Nash game the receiver is made better off in expectation by a finer partition of the signal space. To complete the proof we need only show that B's utility at any point is increasing in T's utility at that point. Using the absolute value loss utility functions given in the paper we have

$$U_B(\cdot) = \beta U_T(\cdot) + d - (1-\beta)(1-h)f \tag{16}$$

which is clearly increasing in $U_T(\cdot)$.

■

Table 4.1: Equilibrium Outcomes in Two-Period Game

| | CC never met | CC met for bad T | CC met if T pool |
|---|---|---|---|
| MC met* in both regions | - B send $m_1(\theta, q)$<br>- Ts separate<br>- B rehires T<br>- B sends $m_2(\theta, d_i)$ | - B send $m_1(\theta, d_g)$<br>- Ts pool on $\alpha(m, d_g)$<br>- B rehires T<br>- B sends $m_2(\theta, q)$ | - No informative pure strategy equilibrium‡ |
| MC met if $m_1 = nc$ | - B sends $m_1(\theta, q)$<br>- Ts separate<br>- B rehires T<br>- B sends $m_2(\theta, d_i)$ | - B sends $m_1(\theta, q)^\dagger$<br>- Ts pool if $m_1 = nc$<br>- B fires bad T, rehires if pool<br>- B sends $m_2(\theta, q)$ if Ts pool<br>- B sends $m_2(\theta, d_i)$ otherwise | - No informative pure strategy equilibrium‡ |
| MC never met | - B sends $m_1(\theta, q)$<br>- Ts separate<br>- B rehires T<br>- B sends $m_2(\theta, d_i)$ | - B sends $m_1(\theta, q)$<br>- Ts separate<br>- B fires bad T<br>- B sends $m_2(\theta, d_i)$ | - B sends $m_1(\theta, q)$<br>- Ts separate<br>- B fires bad T<br>- B sends $m_2(\theta, d_i)$ |

* Here MC being met means it is met only if B fires. For linear utility functions, MC cannot be met without firing. However, the MC can be met for sufficiently concave non-linear loss functions. This can happen because uncertain B signal more informatively than B facing a bad T for certain, meaning the value to a bad T of keeping B uncertain can exceed the present value of playing true to type.
† $m_1(\theta, q)$ here takes into account fact that T pool on $\alpha(m, d_g)$ in the $nc$ region.
‡ Babbling equilibrium can always be sustained with appropriate beliefs.

**Proof of Table 4.1:** Consider each cell of the table in turn and check for profitable deviations from the equilibrium strategy:

- CC never met, MC always met: Even though MC is always met, B will never fire so bad T have no incentive to play as an honest type. By construction then B's first period expected utility is maximized by sending $m_1(\theta, q)$. T's off-the-path beliefs sustain this so long as when T receives any off-the-path signal it plays as though it received some arbitrary on-the-path signal. Because CC is never met, B will never fire and so T has no incentive to play other than true to type. T's actions resolve B's uncertain as to T's type, B does not fire because CC is not met. B sends $m_2(\theta, d_i)$ which is optimal by definition given that all T will play true to type in the final period.

- CC never met, MC met if $m_1 = nc$: By construction B's messaging strategy maximizes his first period expected utility given uncertainty over $d$. T's off-the-path beliefs sustain this so long as when T receives any off-the-path signal it believes it received some arbitrary on-the-path signal. Because CC is never met, firing cannot be a profitable deviation which means the T have no reason not to play true to type. B's second-period messaging strategy contingent on T's behavior is optimal by construction.

- CC never met, MC never met: By construction B's messaging strategy maximizes his first period expected utility given uncertainty over $d$. Since MC is never met, T has no incentive to play other than true to type. Because CC is never met, firing cannot be a profitable deviation. Because T's first-period actions are fully revealing, B's optimal second-period strategy is to send $m_2(\theta, d_i)$.

- CC met for bad T, MC always met: Sending $m_1(\theta, d_g)$ is optimal if both T will play as the good type. Because CC is met for bad T, B will fire if it gets a bad attack. Since MC is always met, T has no incentive to deviate from playing $\alpha(m, d_g)$ in the first round. As the T pool, B receives no further information as to T's type and will not fire since the CC is only met for bad T. By construction B's second period messaging strategy maximizes its expected utility given uncertainty.

- CC met for bad T, MC met if $m_1 = nc$: By construction B's messaging strategy maximizes his first period expected utility given uncertainty over $d$. Solving for the B's optimal strategy given pooling in the no communication region merely involves assuming all T set $\alpha(nc) = E(\theta)$. Since MC is met only if $m_1 = nc$ and since CC is only met for bad T, there is not profitable deviation for any T from pooling on $\alpha(nc) = E(\theta)$ in the no-communication region and from playing true to type elsewhere. If B gets a good attack when it does not communicate, it learns nothing about T's type and so sends $m_2(\theta, q)$. If B sends a signal the T separate. B sends $m_q(\theta, d_g)$ if it got a good attack in the first period and $m_q(\theta, d_a)$ if it got a bad attack and had to sanction T.

- CC met for bad T, MC never met: By construction B's messaging strategy maximizes his first period expected utility given uncertainty over $d$. Since MC is never met, T have no incentive to deviate from playing true to type. Since CC is met for bad T, B will fire upon getting a bad attack. This threat creates no incentives for bad T as MC is never met. B sends $m_2(\theta, d_g)$ if it got a good attack in the first period and $m_2(\theta, d_a)$ if it got a bad attack and had to sanction T. No deviation can be profitable for B in the second round as it knows T's type with certainty.

- CC met if T pool, MC always met: There can be no informative pure strategy equilibrium. Consider a pooling equilibrium. There is no incentive for bad T to play $\alpha(m, d_g)$ if doing so gets him fired. Bad T will therefore deviate to $\alpha(m, d_b)$. If the T separate, though, B would deviate by not firing upon getting a good outcome, in which case bad T would deviate to $\alpha(m, d_g)$.

- CC met if T pool, MC met if $m_1 = nc$: First note that there can be no informative pure strategy equilibrium when $m_1 = nc$ because there is no incentive for bad T to play $\alpha(nc, d_g)$ if doing so gets him fired. Bad T will therefore deviate to $\alpha(nc, d_b)$. But if the T separate when $m_1 = nc$, then B would deviate by not firing upon getting a good outcome after sending $nc$. T would then deviate to $\alpha(nc, d_g)$. As there is no pure-strategy equilibrium when B does not communication, B's expected utility from not communicating is undefined and there can be no informative pure-strategy equilibrium with two or more cutpoints.

- CC met if T pool, MC never met: By construction B's messaging strategy maximizes his first period expected utility given uncertainty over $d$. Since MC is never met, T have no incentive to deviate from playing true to type. Since CC is met for bad T, B will fire upon getting a bad attack. This threat creates no incentives for bad T as MC is never met. B sends $m_q(\theta, d_g)$ if it got a good attack in the first period and $m_q(\theta, d_a)$ if it got a bad attack and had to sanction T. No deviation can be profitable for B in the second round as it knows T's type with certainty.

■

**Proof of Proposition 4.5:** With a linear utility function and uniform distribution of $\theta$, T's expected utility for the second round simplifies to

$$\frac{-4f(h-l)^2 - 3bhl\beta + 4f(-3+2h+l)(h+2l)\beta}{4(h+2l)\beta}. \tag{17}$$

The result follows directly from taking comparative statics under the constraints for a two-cutpoint equilibrium.

■

**Proof of Proposition 4.6:** With a linear utility function and uniform distribution of $\theta$, $w^*$ when B does not communicate simplifies to

$$\frac{-2d^2l^2(h+2l)\beta^2 - 3bfh(h-l)l\beta(-1+\delta) + 4f^2(h-l)(-(h-l)^2 + (-3+2h+l)(h+2l)\beta)(-1+\delta)}{4f(h-l)(h+2l)\beta} \tag{18}$$

The first part of the result follows directly from taking comparative statics under the constraints for a two-cutpoint equilibrium. The conditional results on $w^*$ follow from solving for the values of $d$ and $\delta$ such that $\frac{\partial w^*}{\partial f} > 0$ and $\frac{\partial w^*}{\partial \beta} < 0$.

■

A natural question given Proposition 4.6 is to ask whether the change in $w^*$ is larger in one region or the other because doing so provides leverage on the relative likelihood of different equilibrium outcomes. If, for example, a given reduction in preference divergence lowers the MC when B communicates ($w^*(c)$) more quickly than when B does not ($w^*(nc)$), then conditional on the smaller MC being met, we can say the probability of getting a pooling equilibrium—and thus the highest level of communication in the first period—has increased. Taking comparative statics on the difference between the two MC yields an additional proposition omitted from the text for space reasons.

**Proposition 4.6A:** Conditional on the no-communication MC being met in a two-cutpoint equilibrium, the relative probability that the MC is met whether or not B communicates is decreasing in preference divergence and the cost of failure.

The intuition behind this result is that the first-period value of playing true to type for the bad T is increasing faster in the no-communication region than in the communication region.

**Proof of Proposition 4.6A:** Note that the discounted second period utility drops out when taking the difference between $w^*(nc)$ and $w^*(c)$. Taking comparative statics on

$$ld - \frac{d^2 l^2 \beta}{2f(l-h)} \tag{19}$$

proves the result.

■

**Proof of Lemma 4.1:** To show that it is easier to meet the CC for bad T than if the T pool it is sufficient to show that $EU_2^B(m_2(\theta, q)) - EU_2^B(m_2(\theta, d_b)) > 0$. Solving for the expected utility for any arbitrary level of preference divergence and cutpoints yields

$$\frac{1}{4}\left(4f(-3+2h+l) + \left(-bh - (h-l)(\theta_1 - \theta_2) + 4d^2\left(-\frac{h}{\theta_1} + \frac{l}{\theta_1 - \theta_2} + \frac{h}{-b+\theta_2}\right)\right)\beta\right). \tag{20}$$

$EU_2^B(m_2(\theta, q))$ amounts to the following lottery where the cutpoints account for B's uncertainty: (1) with probability $q$ B gets the value of equation (19) with $d = d_b$; and (2) with probability $(1-q)$ B gets that value with $d = 0$. Because in a two-cutpoint equilibrium the distance between the cutpoints, $\theta_2 - \theta_1$, is invariant to $d$ and $q$, we can eliminate everything but $4d^2(-\frac{h}{\theta_1} + \frac{l}{\theta_1-\theta_2} + \frac{h}{-b+\theta_2}))$ from consideration.[38] Noting that this term is strictly negative for $d > 0$ proves the result.

■

**Proof of Proposition 4.7:** B's expected utility in the second round in a two-cutpoint equilibrium with $\theta$ $unif[0, b]$ and absolute value loss utility is:

$$\begin{aligned}&\frac{1}{4}(4f(-3+2h+l) + \beta(-bh + \frac{(h-l)(4f(-h+l)+bh\beta)}{(h+2l)\beta} - \\ &\quad \frac{4d^2(h+2l)\beta(4f^2h^2(h-l)^2(4h-l) - 4bfh^2(h-l)^3\beta + l(4d^2(h+l)^2(h+2l)^2 - b^2h^2(2h^2+l^2))\beta^2)}{(4f(h-l)-bh\beta)(4f^2h^2(h-l)^2 + 4bfh^2(h-l)l\beta - (-b^2h^2l^2 + 4d^2(h+l)^2(h+2l)^2)\beta^2)})).\end{aligned} \tag{21}$$

The first part of the proof follows from taking the cross partial of (20) with respect to preference divergence and the other parameters. Where the cross partial is positive, the gains from getting an acceptable agent are increasing in the parameter. The last part of the

[38]Note the locations of the cutpoints and thus B's expected utility is not invariant to $d$ and $q$.

proof follows from noting that the second derivative with respect to $d$ on B's second-period expected utility is negative.

■

# 5 Conclusion

This working paper has presented several ways of modeling terrorist organization. The goal was to provide a useful reference and show some of the details behind some of the logic in *The Terrorist's Dilemma*. Other approaches to modeling these phenomena are possible and the models here could surely be presented in a more intuitive way. Two important limitations are worth noting. First, I did little to prove uniqueness in sections 2 and 3 and did not deal there with mixed strategies. Multiple equilibrium are possible in many parts of the parameter space in all three models. A fuller formal discussion of terrorist organization could do more to address equilibrium selection and to analyze where certain outcomes are unique.

# References


Baccara, Mariagiovanna and Bar-Isaac, Heski, “How to Organize Crime,” Working paper, New York University (2006).

Beam, Louis, “Leaderless Resistance,” *The Seditionist*, vol. 12 (1992): 12–13.

Bell, J. Bowyer, “Aspects of the Dragonworld: Covert Communications and the Rebel Ecosystem,” *International Journal of Intelligence and Counterintelligence*, vol. 3, no. 1 (1989): 15–43.

Bell, J. Bowyer, “Revolutionary Dynamics: The Inherent Inefficiency of the Underground,” *Terrorism and Political Violence*, vol. 2 (1990): 193–211.

Bendor, Jonathan, Glazer, Amihai and Hammond, Thomas H., “Theories of Delegation,” *Annual Review of Political Science*, vol. 4 (2001): 235–269.

Bueno de Mesquita, Ethan and Dickson, Eric S., “The Propaganda of the Deed: Terrorism, Counterterrorism, and Mobilization,” *American Journal of Political Science*, vol. 51, no. 2 (2007): 364–381.

Crawford, Vincent P. and Sobel, Joel, “Strategic Information Transmission,” *Econometrica*, vol. 50, no. 6 (1982): 1431–1451.

Geanakopolos, John, “Common Knowledge,” in Aumann, Robert J. and Hart, Segiu eds., *Handbook of Game Theory with Economic Applications*, chap. 40, pp. 1437–1496 (Elsevier, 2006).

Gilligan, Thomas W. and Krehbiel, Keith, “Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures,” *Journal of Law, Economics, and Organization*, vol. 3, no. 2 (1987): 287–335.

McIntyre, Anthony, “Provisional Republicanism: Internal Politics, Inequities and Modes of Repression,” in McGarry, Fearghal ed., *Republicanism in Modern Ireland*, pp. 178–198 (Dublin: University College Dublin Press, 2003).

Olszewski, Wojciech, “Rich language and refinements of cheap-talk equilibria,” *Journal of Economic Theory*, vol. 128 (2006): 164–186.

Osborne, Martin J., *An Introduction to Game Theory* (New York: Oxford University Press, 2003).

Shapiro, Jacob N., *The Terrorist's Dilemma: Managing Violent Covert Organizations* (Princeton University Press, 2013).

Shapiro, Jacob N. and Siegel, David A., "Moral Hazard, Discipline, and the Management of Terrorist Organizations," *World Politics*, vol. 64, no. 1 (2012): 39–78.

Siqueira, Kevin and Sandler, Todd, "Terrorists versus the Government: Strategic Interaction, Support, and Sponsorship," *Journal of Conflict Resolution*, vol. 50, no. 6 (2006): 878–898.