Terrorism: A Model of Intimidation*

Sambuddha Ghosh† Boston University

Gabriele Gratton‡ University of New South Wales

Caixia Shen§ Shanghai University of Finance and Economics

November 2013

Abstract

This paper uses a reputational model to study the strategic core of terrorism: intimidation. The model consists of a negotiation phase and a conflict phase. Failure of negotiations leads to the conflict phase, which we model as war-of-attrition with two-sided incomplete information. For any level of discounting there is a unique equilibrium, where conflict can be prolonged but wasteful. Our results suggest that religious terrorists fight shorter but more frequent wars, while secular terrorists fight longer wars but less frequently. If the government appears tougher, the terrorist starts a conflict with lower probability. Yet, conditional on a first attack, the government's continuation payoff is independent of its perceived toughness. Our model provides insights into why and how negotiations fail and conflicts begin.

* PRELIMINARY *

Keywords: terrorism, reputation, intimidation, repeated games, two-sided asymmetric information. ["terrorism, reputation, intimidation, costly war of attrition"]

JEL Classification Numbers: D74, D82.

*We are grateful to Gaurav Aryal, Mauro Caselli, Anton Kolotilin, Juan Ortner, Carlos Pimienta, and seminar participants at the Australian National University and the 2013 Australasian Economic Theory Workshop.

†Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215 USA.

‡School of Economics, Australian School of Business, University of New South Wales, Sydney 2052, Australia.

§School of International Business Administration, Shanghai University of Finance and Economics, 777 Guoding Road, Shanghai, China 200433.

1 Introduction

Terrorism is a multifaceted phenomenon and no two definitions are identical. Common characteristics are that terrorism involves premeditated use of violence to *intimidate* an opponent to attain a political goal.[1] Much has been written on the violent tactics and the political objectives of terrorist groups.[2] In contrast, we know surprisingly little about the *strategic* core of terrorism: intimidation. This paper is a first step towards reducing this gap.

We build a fully dynamic game-theoretic model of intimidation driven by reputation. This model sheds light on the conditions under which a terrorist conflict starts and its duration. We believe these insights match some anecdotal evidence on the duration of attacks by some major terrorist groups. We review this evidence in Section 3.5.

Before we present our model, let us start with the following one-shot game—a terrorist threatens to attack if the government does not give up a contested resouce. Should the government concede? If the attack is costly for the terrorist, then the threat is not credible:There is no reason for the terrorist to incur the cost of an attack even if the government does not concede. Therefore, it is optimal for the government to not concede. Let us refer to the above one-shot game as the GT model. In GTGT (or TGTG), the unique equilibrium is that there is no attack. If GT (or TG) is repeated finitely many times, there is a unique equilibrium where T never attacks and G never concedes.

Thus, an interesting model requires some form of incomplete information—the government must believe that the terrorist is committed to carrying out an attack if her demands are not met.[3] In fact terrorist groups employ this strategy—first cause harm by launching an attack and then threaten more attacks until a concession. By launching an initial attack, the terrorist can build a reputation for being committed to attacking unless the demands are met. Our model captures this phenomenon.

We develop a model with infinitely many periods that are classified into two phases—the negotiation phase and the conflict phase. The first phase covers only period 0; while the second one could extend over periods $1, 2, \ldots$. It seems more natural to not impose an exogenous terminal period.

In each period of the conflict phase, a terrorist and a government play a two-stage extensive-form game. In the first stage, the terrorist decides whether or not to attack. The terrorist can be of two types—*fundamentalist* or *normal*. With probability μ_0, the terrorist is a fundamentalist who is committed to attack until the government gives up; with probability $1 - \mu_0$, the terrorist is *normal* and incurs a cost in each period she attacks. In the second stage, a government decides whether to concede to the terrorist's request or not. The government can be of two types—*tough* or *normal*. With probability π_0, the government is *tough* and never concedes; with probability

[1] We review some common definitions of terrorism in Section 1.1.

[2] See, among others, Enders and Sandler (2012), Krueger (2007), and Richardson (2006).

[3] Indeed, Gould and Klor (2010) show that in some measure terrorism works, meaning that terrorist attacks induce the victim population to be more willing to concede.

$1-\pi_0$, the government is *normal* and suffers a loss from each terrorist attack. Each normal type maximizes the discounted sum of future utilities.

Although infinitely repeated games are often marked by a plethora of equilibria, we show that, for any level of discounting, there is a unique perfect Bayesian equilibrium of the game of conflict and therefore of the entire game. Uniqueness allows us make a variety of comparative predictions. We introduce a notion of what it means for one player to be relatively more committed than the other; the precise nature of the unique equilibrium depends on who is more committed, the terrorist or the government.

The conflict game gives us three major predictions regarding the likelihood and length of the conflict once negotiation fails. First, when the prior probability that the terrorist is fundamentalist is larger, then the terrorist begins to attack with higher probability. Terrorist groups with strong religious beliefs are perceived as more likely to be committed to fight indefinitely for their goal. Our results suggest that these groups are more likely to start a terrorist conflict or break a truce.

Second, if the government is more likely to be tough, then the terrorist begins to attack with lower probability. An immediate implication is that a government wants to be perceived as tough. Yet, the perceived advantage of showing commitment should not be overstated: Once an armed conflict begins, the expected payoff for the government is independent of the probability of being tough. Intuitively, when the government is perceived to be tough, it expects a rational terrorist to attack with low probability.

Third, if a first attack is carried out by the terrorist, then the length of the conflict depends on the player most likely to be committed. When this player is more likely to be committed, then the conflict is shorter (unless, of course, both players are actually of the committed type). In our interpretation—in the plausible scenario where the terrorist is more likely to be committed than the government—religious terrorists should be expected to fight short (and frequent) wars, while secular terrorist fight longer (though rare) wars. In Section 3.5 we explain why this prediction seems broadly consistent with some observations.

Our equilibrium exhibits intimidation, that is after an attack, the government must update its beliefs and give higher probability to the terrorist being a fundamentalist. <WHY ONLY 1 sentence?>

The negotiation phase gives us predictions regarding when and why negotiations fail. Intelligence can do much to diminish the risk of terrorist's attacks, but it is unthinkable for a government to stop all attacks.[4] The best course of action is therefore to prevent the organization of terrorist groups by offering sufficient concessions to alienate support for the violent part of opposition.[5] We assume that the government is uncertain about the preferences of the

[4] Indeed, surveys indicate that some three quarters of Americans believe occasional acts of terrorism to be part of the future and around 45% of them thinks there is not much more the government can do to prevent them (see Pew Research Center for the People & the Press, April 23, 2013, http://www.people-press.org/2013/04/23/most-expect-occasional-acts-of-terrorism-in-the-future/).

[5] Popular support is crucial for terrorist groups. For example, Bueno de Mesquita and Dickson (2007) show

opposition at the time it makes an offer. Our results predict that negotiations systematically fail only when the terrorist is relatively more likely to be a fundamentalist but the government expects to curb the terrorist's support from the opposition when it offers partial concessions. In equilibrium, the government knows that in case of failure an armed conflict will ensue with probability 1. Yet, it fails to make a sufficiently large offer to prevent the formation of a violent terrorist group. Crucially, the government would have conceded more had it known with certainty the preferences of the opposition group. Following our previous interpretation, our results suggest that negotiations are more likely to fail when the terrorist group has strong religious beliefs.

Often more than two parties in the group are in potential conflict between each other.[6] Drawing on Siqueira (2005) and Enders and Sandler (2012), we extend the model to introduce a political wing on the side of the terrorist. The opposition group is formed by a political wing and a terrorist wing. During the negotiation stage, the government offers a partial concession to the political wing. If the political wing refuses the offer, then negotiations break down, a terrorist is born and we enter the conflict phase. If it accepts the offer, then he decides whether or not to repress the terrorist wing that might be born. Repression carries a cost for the political wing. If the repression succeeds, then conflict is avoided and the political wing gets the concession offered by the government. If repression fails then the terrorist is born and the conflict phase starts. We say that the political wing is *weak* if the cost of repression is high. Otherwise it is *strong*. Crucially, the cost of repression is the political wing's private information.

In the realistic situation where the terrorist is more likely to be committed than the government is, the latter can only lose from engaging in conflict and, once conflict starts, its continuation payoff is the same as being attacked a single time and then be forced to concede with probability 1. (Note we are not claiming this is what happens in equilibrium, but rather that the payoffs are as if this were the case.) We predict that mutually undesirable war will arise if (i) there is a significant probability that the political wing is strong, but it is actually weak, and (ii) the terrorist is not very likely to be fundamentalist. In this case, the government makes a small concession, gambling that this might be enough. But when this is not enough, armed conflict starts with probability 1. Crucially, the government would have conceded more had it known the political wing's type with certainty. Both our models of negotiation predict that negotiations fail with positive probability when the government even makes significant concessions.

The following section presents a brief review of the related literature. Section 2 introduces the baseline model. Sections 3 and 4 study the conflict and negotiation phases respectively. Section 5 extends the model to a scenario where the opposition group is made up of a political and a military wing. Section 6 concludes.

---

how terrorists choose to attack in order to create support for their cause; Berman and Laitin (2008) put forward the idea that even volunteering for suicide attacks can be explained by the crucial role that religious terrorists play in providing public goods to their communities.

[6] "Examples include the ETA and Batasuna; the PLO and Fatah; the IRA and Sinn Féin; and the Islamic Resistance and Hezbollah" (Enders and Sandler, 2012).

1.1 Related Literature

The related literature is vast, as it covers a variety of fields, from politics to sociology, passing through defense economics. We shall therefore not attempt to provide a comprehensive summary of the literature.[7] The definition of terrorism is also subject of intense debate, as the objectives and backgrounds of the faction that perpetuates the violence might enter the definition. Admittedly, we define terrorism quite narrowly, focusing on the strategic aspect of it. Yet, almost all definitions of terrorism contain this aspect. For example, Igor Primoratz (2012) defines terrorism as "coercive intimidation" and "intimidation with a purpose... meant to make others do things they would otherwise not do" (p. 10). Enders and Sandler (2012) argue that a conservative definition of terrorism is "the premeditated use or threat to use violence by individuals or dissident groups to obtain a political or social objective through intimidation of a large audience beyond that of the immediate victims" (p. 4). Krueger (2007) defines the goal of terrorism as "spreading fear" and clarifies that terrorism is a "tactic" (p. 14-15) as opposed to a state of mind of the terrorist.

Within the game theoretic literature, Lapan and Sandler (1988) were among the first to model terrorism as a game between players who are irrational with some probability. However, in their model, absent a concession, the probability of being a commitment type jumps up to an arbitrary and exogenously given quantity. These probilities are endogenous in our model. Hodler and Rohner (2012) carries out a Bayesian analysis of the probability of the players being committed but they have only two periods; this in turn means that they predict attacks only when the probability of the terrorist being committed is very large, which does not sit well with the view that they are often rational calculating agents.There is no negotiation stage in these papers. Also, our model is infinitely repeated and so avoids last-period or "end-game" effects.

From a technical stand-point, the closest to our work is that of Abreu and Gul (2000), which provides a thorough analysis of reputational effects in bargaining. However, the model is a continuous-time model, and results available for the discrete case need the time interval between periods to vanish; we believe that this assumption is not suitable for the analysis of terrorism, where the drama is played out on a much slower and more extended scale. Our uniqueness results hold for any discount factor.

One of our innovation is to model a negotiation phase followed by a conflict phase. Adding a negotiation stage makes this fit case-studies of terrorism much better as they rarely (if ever!). start with a conflict. Part of the challenge is to understand when a struggle by a dissident group against a government will be resolved peacefully and when it will slide into a conflict that is beneficial to neither side, but could still drag on and consume resources, including lives. In particular, Section 5 extends our model of negotiation to include a political wing of the terrorist group. A somewhat similar idea is in Siqueira (2005) who studies the interaction of militant and

[7] Richardson (2006) provides a comprehensive treatment of the literature within politics. Bueno de Mesquita (2008) provides a selective overview of the recent literature on terrorism in the context of political economy.

political wings of a terrorist group.

Intimidation is not the only description that fits the strategy of terrorism, especially international terrorism. Baliga and Sjöström (2012) is a model of persuasion where an extremist can induce other players to take hawkish actions. Thus, terrorism is communication aimed at inducing countries to enter in conflict with each other. Bueno de Mesquita (2005) considers an infinitely repeated game of complete information in which the government makes concession to secure the help of moderates to stamp out the extremists. However since the extremists are left in control, violence is higher until one party is defeated. There is no role for reputation. To get around the plethora of equilibria, the author restricts attention to efficient Markovian equilibria. Our model is complementary: it not only captures the reputational side of terrorism and intimidation, but also derives a unique equilibrium.

Following the seminal works Schelling (1956, 1960, 1966) and Crawford (1982), studies of counterterrorism have advocated that the government commit to more hawkish bargaining positions.This leaves open the question:How this reputation for hawkishness built? Our model delivers reputation endogenously as the result of the government resisting the attacks of the terrorist without making concessions. Yet, our model highlights how terrorism itself is a strategy aimed at gaining reputation and that the advantages of more hawkish positions should not be overstated.

2 The Baseline Model

There are infinitely many periods $t \in \mathbb{Z}^+ = \{0\} \cup \mathbb{N}$. In the baseline model there are two players—a government G and a terrorist T. The game starts with G in possession of a resource over which T has a claim.

There are two *phases* in the game: the first, *negotiation*, lasts for just one period, $t = 0$; the next, *conflict*, can last though periods $t \in \mathbb{N}$. We now describe these phases in more detail, but it is convenient to do so in reverse order.

Conflict Phase We describe the second phase first. In each period $t \in \mathbb{N}$ the extensive-form stage-game is played. The stage-game itself has two steps and the first has the terrorist T deciding whether or not to attack. T can be of two types—with probability $\mu_0 > 0$, T is a fundamentalist who is committed to attacking until the government gives up; otherwise, T is *normal* and pays a cost $b > 0$ in each period she attacks, and rationally chooses its decision to attack. In the second stage, a government G decides whether to concede to T's demand or not. If the government does not concede, it enjoys a rent $\gamma > 0$ in the current period from retaining control over the resource; if it concedes, the terrorist enjoys a rent $\tau > 0$ in the current period and in each subsequent periods. The government can be of two types: With probability π_0, the government is *tough* and never concedes; otherwise,it is *normal*, and suffers a loss $L > 0$ from each terror attack. Both normal types, the government and the terrorist, maximize the sum of future utilities discounted by the factor $\delta > 0$. Our solution concept is the perfect Bayesian equilibrium (henceforth *equilibrium*), since subgame-perfection does not have any bite in such models of incomplete information.

At each period $t+1$, at stage 1 the state of the game is defined by a vector $(\mu_t, \pi_t, \theta_{t+1})$, where μ_t is G's belief about the probability that T is the fundamentalist type, π_t is T's beliefs that G is the tough type, and $\theta_t \in \{W, NW\}$ is W if G has not conceded yet and NW otherwise; at stage 2, the state vector is $(\mu_{t+1}, \pi_t, \theta_{t1})$, where G's beliefs about T's type have been updated from μ_t to μ_{t+1} in light of T's action at stage 1. By assumption, $\theta_1 = W$. We assume that NW is an absorbing state such that the expected present discounted value of NW is 0 for the government and $U^T \equiv \tau / (1-\delta)$ for the terrorist. Also, if $\mu_t = 0$ (i.e. if G believes that the terrorist is of normal type with probability 1) then the present discounted value is $U^G = \gamma / (1-\delta)$ for the government and 0 for the terrorist.

The extensive form and the parameters of the game, including the probabilities μ_0 and π_0, are common knowledge. Since the game has perfect monitoring of actions, this also means that subsequent beliefs are common knowledge. We start with two parametric assumptions that make the game interesting.

Assumption 1. *$\delta L > \gamma$; $b < U^T = \tau / (1-\delta)$.*

The first of the two assumptions says that the loss from one attack in the next period exceeds the enjoyment of the contested resource in the current period. The second says that the cost of attacking is strictly less than the discounted value of getting the contested resource forever, starting from the current period; it immediately follows from this that T strictly prefers to attack if she knows that one attack will result in the government conceding immediately.

Negotiation Phase We now describe the first phase, which is an extensive-form game with four steps, listed in chronological order below.

1. At the beginning, G makes a take-it-or-leave-it offer $\beta^G \in [0,1]$ to T, when the terrorist has not yet built attacking capabilities.

2. If G offers β^G, the terrorist is develops attacking capabilities with probability $f(\beta^G)$, where $f : [0,1] \to (0,1)$ is a decreasing function. We argue that this assumption captures the terrorist's need for a large enough mismatch between the expectations and the offer to aggregate enough consensus and recruit soldiers. So even if the government's offer were spurned, the terrorist must work harder to launch an armed struggle if the the initial offer is more generous.

3. After observing if she can attack or not, T chooses to accept or reject the offer. If T accepts the offer, the game ends and G receives $(1-\beta^G)U^G$ and T receives $\beta^G U^T$. If T rejects the offer, then types of G and T are drawn and the game of conflict is played as above.

Note that the entire extensive-form negotiation game described above lasts for a "single" period $t = 0$. Point 3 above reflects our view that the tough government is unable for internal or political reasons to concede in the face on an attack; such a type will be known (or "drawn") only when an attack signals the start of a conflict. Likewise, the terrorist group is just a group of dissidents at the moment of negotiation; only when negatiations fail and the group has the choice to arm itself, will it be clear how fundamentalist a leader assumes charge of the group.

Another parametric assumption, which is generically satisfied, allows us to obtain a unique equilibrium of the two-sided incomplete information game comprising the negotiation and conflict phases. While it is stated in terms of the primitives of the model, we defer the statement to a point where its interpretation will be evident.

3 Conflict

We characterize what is generically the unique equilibrium of the conflict phase, dropping the state θ_t from our notation and assuming that all strategies are defined conditional on $\theta_t = W$.

A (behavior) strategy[8] for T is a sequence of mappings $\sigma^T_{t+1} : [0,1]^2 \to [0,1]$, $t \in \mathbb{N}$, where $\sigma^T_{t+1}(\mu_t, \pi_t)$ is the probability that the normal type of T concedes in period $t+1$ conditional on the public beliefs and history. A (behavior) strategy for G is a sequence of mappings $\sigma^G_{t+1} : [0,1]^2 \to [0,1]$, one for each $t \in \mathbb{N}$, where $\sigma^G_{t+1}(\mu_{t+1}, \pi_t)$ is the probability with which G gives up in period $t+1$ when the public beliefs are μ_{t+1} and π_t. Note that the strategy of player T in period $t+1$ depends on the beliefs at the end of period t, as is standard; in contrast, G observes T's move at $t+1$ and updates her belief about T's type to μ_{t+1} before acting in period $t+1$; this is because each period is thought of as a length of time during which an extensive-form game is played out.

3.1 Beliefs

First, it is obvious that if G believes T to be a fundamentalist with sufficiently high probability, then G would concede immediately. Similarly, if T believes G to be tough with sufficiently high probability, then T would concede immediately. This follows from two observations: (i) each player concedes if it/she believes the other to be committed with probability 1; (ii) payoffs, and hence optimal strategies, are continuous in beliefs. In order to prove such results formally we introduce some notation.

Since the committed types never concede, the total probabilities of concession by G and T respectively are obtained by multiplying the normal type's probability with the probability that

[8] Properly speaking, we are restricting attention to Markovian formulations, where the strategy depends only on the public beliefs. While this is often done in applied settings for simplification, it is without loss of generality here since our results apply to all equilibria, not just Markovian ones.

the (respective) normal type concedes:

$$\begin{aligned}\bar{\sigma}^G_{t+1}(\mu_{t+1},\pi_t) &= (1-\pi_t)\,\sigma^G_{t+1}(\mu_{t+1},\pi_t)\,;\\ \bar{\sigma}^T_{t+1}(\mu_t,\pi_t) &= (1-\mu_t)\,\sigma^T_{t+1}(\mu_t,\pi_t)\,.\end{aligned}$$

If T has not conceded until period $t+1$, the updated belief μ_{t+1} that T is the commitment type is derived by Bayes' rule from μ_t and σ^T_{t+1}:

$$\mu_{t+1} = \frac{\mu_t}{\mu_t + (1-\mu_t)\left(1-\sigma^T_{t+1}(\mu_t,\pi_t)\right)} = \frac{\mu_t}{1-\bar{\sigma}^T_{t+1}(\mu_t,\pi_t)}; \tag{1}$$

obviously, $\mu_{t+1} = 0$ at any history where T concedes. Similarly, if G has refused to concede throughout we have

$$\pi_{t+1} = \frac{\pi_t}{\pi_t + (1-\pi_t)\left(1-\sigma^G_{t+1}(\mu_{t+1},\pi_t)\right)} = \frac{\pi_t}{1-\bar{\sigma}^G_{t+1}(\mu_{t+1},\pi_t)}, \tag{2}$$

and $\pi_{t+1} = 0$ as soon as G concedes. The following two quantities are in $(0,1)$ thanks to Assumption 1:

$$\bar{\mu} := \frac{U^G + \gamma/\delta}{U^G + L} = \frac{\gamma}{\delta\left[\gamma + (1-\delta)L\right]}, \text{ and} \tag{3}$$

$$\bar{\pi} := 1 - \frac{b(1-\delta)}{\tau}. \tag{4}$$

Lemma 1. *In any equilibrium,*

(i) if $\pi_t > \bar{\pi}$, then T concedes with probability 1 at time $t+1$;

(ii) if $\sigma^G_{t+1}(\mu_t,\pi_t) = 1$, then T strictly prefers to attack if $\pi_t < \bar{\pi}$ and is just indifferent at $\bar{\pi}$; and

(iii) if $\pi_t < \bar{\pi}$ and $\mu_t > \bar{\mu}$, then T attacks at $t+1$ with probability 1 and G concedes if and only if T attacks.

Proof. T's payoff from choosing $\sigma^T_{t+1}(\mu_{t+1},\pi_t) = 0$, i.e. attacking for sure at $t+1$, is maximum if after an attack at $t+1$ the normal type of G concedes for sure, *and* if T concedes at $t+2$ if G does not concede at $t+1$. Note that when G concedes in period $t+1$, the terrorist gets a flow payoff of τ from the resource starting with the current period. When G is the committment type, it doesn't concede at $t+1$ and T gets none of the resource because she concedes at $t+2$. Hence T's *maximum* payoff from attacking at $t+1$ is

$$(1-\delta)(-b) + (1-\pi_t)\tau + \pi_t \cdot 0. \tag{5}$$

Expression (5) is zero at $\bar{\pi}$, and negative above it. This proves part (i).

For part (ii) note that $\sigma^G_{t+1}(\mu_t,\pi_t) = 1$ implies that if G does not concede in period $t+1$ then player T must put probability 1 on the crazy type, i.e. $\pi_t = 1$ and hence will find it optimal to concede at $t+2$ by step (i) above. Then the payoff of attacking is given by (5) and the payoff of conceding is 0; hence $\bar{\pi}$ is again the point of indifference.

Part (iii) follows from a similar argument to part (i). □

If it is very likely that G is tough, then a normal T would not attack and G would concede any time there is an attack. It is worth noting that this holds independently of the probability T is fundamentalist, i.e. the probability of a tough G trumps that of a fundamentalist T; this is a second-mover advantage. If G is not very likely to be tough but T is very likely to be a fundamentalist, then T attacks in period 1 and G responds by conceding immediately; if G does not concede, then T will never attack again.

We now define two notions of how committed the players are: the first is an absolute notion, while the second is a relative notion. Both compare the priors μ_0 and π_0 to the threshold values identified in Lemma 1. For simplicity, we make the following genericity assumption, which rules out a very small set (of zero Lebesgue measure) of priors but gives us uniqueness.

Assumption 2. *The quantities* $\ln \bar{\pi} / \ln \pi_0$ *and* $\ln \bar{\mu} / \ln \mu_0$ *are not integers.*

Definition 1. A conflict is of commitment order $n \in \mathbb{N} \cup \{0\}$ if n is the largest non-negative integer such that $\mu_0 < \bar{\mu}^n$ and $\pi_0 < \bar{\pi}^n$.

Note that the commitment order is 0 if both $\bar{\mu} < \mu_0 < 1$ and $\bar{\pi} < \pi_0 < 1$; otherwise the commitment order is non-zero.

Definition 2. In a conflict of order n, the terrorist is *relatively more likely* to be committed if $\bar{\mu}^{n+1} < \mu_0 < \bar{\mu}^n$ and $\pi_0 < \bar{\pi}^{n+1}$; the government is *at least as likely* to be committed if $\mu_0 < \bar{\mu}^n$ and $\bar{\pi}^{n+1} < \pi_0 < \bar{\pi}^n$.

3.2 Recursively Evaluating Expected Payoffs

Conditional on $\theta_t = W$, the payoff of T in the game starting at time $t+1$ can be expressed as a function of the public beliefs that each player is crazy. Let $V^i_{t+1} : [0,1]^2 \to \mathbb{R}$ be the value function of player $i \in \{T, G\}$. Then the value function for T satisfies the following recursive equation, which is elaborated on below:

$$V^T_{t+1}(\mu_t, \pi_t) = \left(1 - \sigma^T_{t+1}(\mu_t, \pi_t)\right) [\bar{\sigma}^G_{t+1}(\mu_{t+1}, \pi_t) U^T + \\ + \left(1 - \bar{\sigma}^G_{t+1}(\mu_{t+1}, \pi_t)\right) \delta V^T_{t+2}(\mu_{t+1}, \pi_{t+1}) - b]. \quad (6)$$

If T does not attack at the start of period $t+1$, then there is no cost or gain and the payoff is zero. If T attacks at a cost of b, which happens with probability $1 - \sigma^T_{t+1}(\mu_t, \pi_t)$, she gets the entire surplus if G concedes, and if G does not concede she gets the discounted value of payoffs from the next period, keeping in mind that both T and G are more likely to be commitment types at the start of period $t+2$.

Similarly, the value function for G from the start of period $t+1$, after T has already attacked, and hence $\mu_t > 0$, is recursively given by

$$V_{t+1}^{G}\left(\mu_{t+1},\pi_t\right) = \left(1-\sigma_{t+1}^{G}\left(\mu_{t+1},\pi_t\right)\right)\left[\gamma + \bar{\sigma}_{t+2}^{T}\left(\mu_t,\pi_t\right)\delta U^{G} + \left(1-\bar{\sigma}_{t+2}^{T}\left(\mu_t,\pi_t\right)\right)\delta V_{t+2}^{G}\left(\mu_{t+2},\pi_{t+1}\right)\right] - L. \quad (7)$$

It is worth noting that the value function calculates the payoff of G from the start of period $t+1$ conditional on an attack having already occurred; hence the payoff L is subtracted. The first term is the product of two factors, the probability that G does not concede and the payoff conditional on not conceding; this payoff comprises three components—the current utility derived from the resource, the future value in case T gives up at the start of the next period, and the continuation value in case T attacks one more time.

Let us now focus, setting aside for now the justification, on an equilibrium where T and G are indifferent. For T to be indifferent we need $V_{t+1}^{T}\left(\mu_t,\pi_t\right) = 0$, which is the payoff of T from conceding. Similarly, G is indifferent if and only if $V_{t+1}^{G}\left(\mu_{t+1},\pi_t\right) = -L$, the payoff G gets if it concedes.[9] Suppose that T is indifferent for two consecutive periods. Setting $V_{t+1}^{T}\left(\mu_t,\pi_t\right) = V_{t+2}^{T}\left(\mu_{t+1},\pi_{t+1}\right) = 0$ in equation (6), and assuming $\sigma_{t+1}^{T}\left(\mu_t,\pi_t\right) \neq 1$, we get

$$\bar{\sigma}_{t+1}^{G}\left(\mu_{t+1},\pi_t\right) = \frac{b}{U^T} = \frac{b(1-\delta)}{\tau} = 1-\bar{\pi}.$$

This corresponds to a strategy for G such that

$$\sigma_{t+1}^{G}\left(\mu_{t+1},\pi_t\right) = \frac{1-\bar{\pi}}{1-\pi_t} =: \tilde{\sigma}_{t+1}^{G}(\pi_t) \in [0,1] \text{ if } \pi_t \leq \bar{\pi}. \quad (8)$$

Similarly, suppose that G is indifferent for two consecutive periods (or is indifferent at time t and concedes at time $t+1$). Putting $V_{t+2}^{G}\left(\mu_{t+2},\pi_{t+1}\right) = -L = V_{t+1}^{G}\left(\mu_{t+1},\pi_t\right)$ into equation (7) and assuming $\sigma_{t+1}^{G}\left(\mu_{t+1},\pi_t\right) \neq 1$, we must have

$$\bar{\sigma}_{t+2}^{T}\left(\mu_{t+1},\pi_{t+1}\right) = 1 - \frac{U^G + \gamma/\delta}{U^G + L} = 1 - \frac{\gamma}{\delta\{\gamma + (1-\delta)L\}} = 1-\bar{\mu}.$$

This corresponds to a strategy for T such that

$$\sigma_{t+2}^{T}\left(\mu_{t+1},\pi_{t+1}\right) = \frac{1-\bar{\mu}}{1-\mu_{t+1}} =: \tilde{\sigma}_{t+1}^{T}(\mu_{t+1}) \in [0,1] \text{ if } \mu_{t+1} \leq \bar{\mu}. \quad (9)$$

[9] This payoff is not 0 but $-L$ for the same reason that the term $-L$ appears in equation (7)—when it is G's turn to decide if it wants to concede or prolong the fight, T has already attacked and the loss will be experienced by G in the current period regardless of her choice of move.

The following lemma is a summary of the above discussion.

Lemma 2. *If T is indifferent between conceding and fighting at times t and $t+1$ in any equilibrium, then G's equilibrium concession probability and the public beliefs about G are*

$$\tilde{\sigma}_t^G(\pi_{t-1}) := \frac{1-\bar{\pi}}{1-\pi_{t-1}}; \text{ and } \pi_t = \frac{\pi_{t-1}}{\bar{\pi}}$$

respectively. Similarly, if G is indifferent between conceding and fighting at times t and $t+1$ in any equilibrium, then T's probability of conceding and the public beliefs about his type are

$$\tilde{\sigma}_{t+1}^T(\mu_t) := \frac{1-\bar{\mu}}{1-\mu_t}, \text{ and } \mu_{t+1} = \frac{\mu_t}{\bar{\mu}}.$$

Remark 1. T's mixing probability at $t=1$ need not equal $\tilde{\sigma}_1^T$; G's mixing at $t=1$ can be different from $\tilde{\sigma}_1^G$ only if T strictly prefers to attack at $t=1$.

It remains to show that there is an equilibrium of the two-sided incomplete information game of conflict where player $i \in \{T, G\}$ concedes with the probabilities in Lemma 2 above, except possibly at $t=1$. In what follows we shall show that the two indifference conditions (8) and (9) capture key features of all equilibrium strategy profiles, so that beliefs must evolve according to the above lemma, except possibly at $t=1$ and until they hit $\bar{\mu}$ or $\bar{\pi}$.

3.3 Equilibrium of the Conflict Phase

We can now characterize the unique equilibrium of the conflict phase, assuming that the terrorist group has developed the potential to attack. For ease of exposition and interpretation we state two separate results, starting with one where the government is at least as likely to be committed. In both propositions the last two items are about the evolution of beliefs, whereas the ones before are about equilibrium actions.

Proposition 1. *In a conflict of order n, if the government is at least as likely to be committed, then there exists a unique equilibrium in which*

(i) *an armed conflict begins with probability $\mu_0/\bar{\mu}^n < 1$;*

(ii) *both the (normal) government and the terrorist concede with probabilities strictly less than 1 for n periods and with probability 1 from $n+1$ onwards;*

(iii) *if T attacks at time $t \geq 2$ and $\mu_{t-1} \leq \bar{\mu}$, we have $\mu_t = \mu_{t-1}/\bar{\mu}$; and*

(iv) *if G resists at t and $\pi_{t-1} \leq \bar{\pi}$, we have $\pi_t = \pi_{t-1}/\bar{\pi}$.*

Proof. All omitted proofs are in appendix. □

Figure 1: How Beliefs Evolve 1

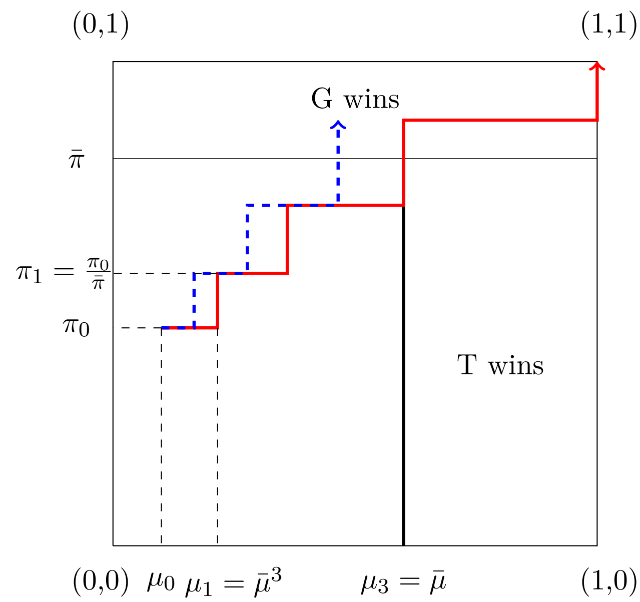

For intuition, consider first how a conflict of order n plays out. In equilibrium, any new attack increases the belief that the terrorist is a fundamentalist. Similarly, any time the government does not concede, this increases the belief that the government is tough. Thus, the conflict cannot last for more than n periods if at least one player is normal. To see this point, recall that Lemma 2 shows that each act of non-concession increases the probability of a fundamentalist T and a tough G by factors of $1/\bar{\mu}$ and $1/\bar{\pi}$ respectively. By definition of "order or a conflict" we know that $\pi_0/\bar{\pi}^n > \bar{\pi}$, i.e. after n periods of non-concession the probability of the government being committed goes above $\bar{\pi}$. Lemma 1 then shows that conflict ends at period $n+1$ with T conceding. Note that at stage 1 of time $n+1$, the terrorist knows that the government will concede after even a single attack. Yet, since she believes the government is tough with probability greater than $\bar{\pi}$, she concedes independently of what she expects a normal government to do. Now we proceed backwards.

What should the government do in stage 2 of period n? It knows that the terrorist will immediately concede if it does not. Yet, it believes the terrorist to be fundamentalist with probability exactly $\bar{\mu}$ and is therefore indifferent between conceding and not. To see this, suppose first that the government holds beliefs $\mu_n < \bar{\mu}$. Then it would concede with probability 0. But if this were the equilibrium strategy, then we would have $\pi_n = \pi_{n-1} < \bar{\pi}$. Second, if $\mu_n > \bar{\mu}$, then the government would concede with probability 1. But if this were the equilibrium strategy, the terrorist would attack with probability 1 in period n and we would have $\mu_n = \mu_{n-1} < \bar{\mu}$. Thus, beliefs must be such that $\mu_n = \bar{\mu}$ and $\pi_{n-1} < \bar{\pi}$, which implies that in period 1 the terrorist must "gain some reputation". The only way to do so is by attacking with probability less than

1.

The solid line in Figure 1 depicts the evolution of beliefs in a conflict of order 3. The dashed line represents the evolution of beliefs if it is common knowledge that the terrorist attacks with greater probability in period 1. In this case, π_3 would be greater than $\bar{\pi}$. Thus, in period 3 the government knows that the terrorist would concede with probability 1 if the government has not conceded yet. Since $\mu_3 < \bar{\mu}$, then the government strictly prefers not to concede.

The total probability of an attack in period 1 is therefore $\mu_0/\bar{\mu}^n$. From then on, G resists with probability $\bar{\pi}$ and T attacks with probability $\bar{\mu}$. Thus, the unconditional probability of an attack in period $t \in [1, n]$ is given by

$$\frac{\mu_0}{\bar{\mu}^n} \left(\bar{\mu}\bar{\pi}\right)^{t-1}.$$

We now turn to the case when the terrorist is relatively more likely to be committed.

Proposition 2. *In a conflict of order n, if the terrorist is more likely to be committed, then there exists a unique equilibrium in which*

(i) *an armed conflict begins with probability* 1*;*

(ii) *G does not concede with probability $\pi_0/\bar{\pi}^n$ in period* 1*;*

(iii) *both the government and the terrorist concede with probability strictly less than* 1 *until stage* 1 *of period $n+1$, and with probability* 1 *from stage 2 of period $n+1$;*

(iv) *if T attacks at time $t \geq 2$ and $\mu_{t-1} \leq \bar{\mu}$, we have $\mu_t = \mu_{t-1}/\bar{\mu}$; and*

(v) *if G resists at time $t \geq 2$ and $\pi_{t-1} \leq \bar{\pi}$, we have $\pi_t = \pi_{t-1}/\bar{\pi}$.*

The basic intuition is not different from Proposition 1. The main difference is that in period 1 it is now the government who needs to gain sufficient reputation for being tough to play on a leveled field with the terrorist. Thus, the government concedes with a higher probability than the one that would make the terrorist indifferent between attacking and not attacking. It follows that in period 1 the terrorist attacks with probability 1.

After the first period beliefs are such that $\mu_1 = \mu_0$ and $\pi_1 = \bar{\pi}^n$. From then on, both players keep playing the (mixed) strategies in Lemma 2 so that at stage 2 of period $n+1$, beliefs are $\mu_{n+1} > \bar{\mu}$ and $\pi_n = \bar{\pi}$ and the government concedes with probability 1.

The solid line in Figure 2 depicts the evolution of beliefs in a conflict of order 2. The dashed line represents the evolution of beliefs if it is common knowledge that the terrorist attacks with smaller probability in period 1. In this case, μ_2 would be greater than $\bar{\mu}$. Thus, in period 2 the terrorist knows that the government would concede with probability 1 after the attack. Since $\pi_1 < \bar{\pi}$, then the terrorist strictly prefers to attack.

The total probability of an attack in period 1 is 1. In the same period, G resists with probability $\pi_0/\bar{\pi}^n$. From then on, G resists with probability $\bar{\pi}$ and T attacks with probability

Figure 2: How Beliefs Evolve 2

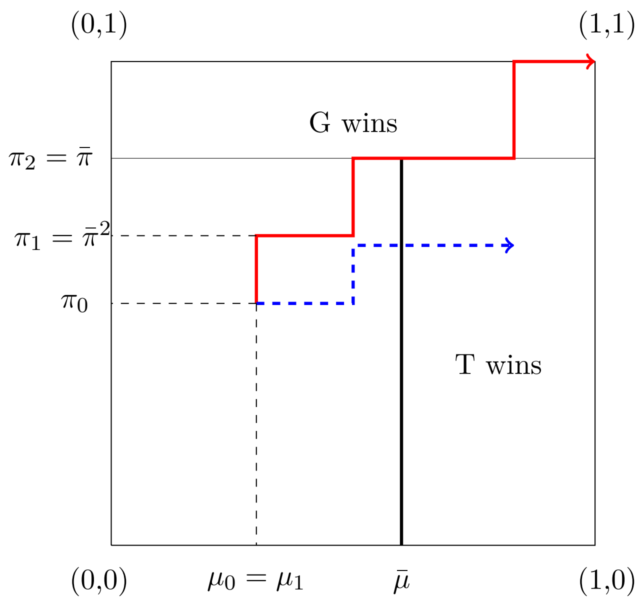

$\bar{\mu}$. Thus, the unconditional probability of an attack in period $t \in [2, n+1]$ is given by

$$\frac{\pi_0}{\bar{\pi}^n}\bar{\mu}^{t-1}\bar{\pi}^{t-2}.$$

3.4 Comparative Statics and Some Predictions

We derive some comparative statics. This is well-defined thanks to our proof of uniqueness.

The following corollary says that the probability of an armed conflict is increasing in the relative likelihood of the terrorist being committed. Thus, our model predicts that terrorist groups which are perceived to be fundamentalist are more likely to recur to terrorist strategies and start a conflict. In this sense, a truce with Hamas, for example, should not be expected to last long (see Section 3.5).

Figure 3 depicts the total probability of first attack as a function of μ_0 for two different values of π_0, $\pi_0' \in (\bar{\pi}^{h+1}, \bar{\pi}^h)$ and $\pi_0'' \in (\bar{\pi}^{l+1}, \bar{\pi}^l)$. In both cases, the probability of first attack is strictly increasing for low values of μ_0 and equals 1 for higher values.

Corollary 1. *Fix the likelihood π_0 that the government is committed. The probability that the terrorist starts an armed conflict is increasing in the likelihood of the terrorist being committed μ_0.*

For a government, an image of toughness can pay: if the government is at least as likely to be committed as the terrorist, then the probability of a conflict is strictly less than 1. In this case, the probability of a conflict is $\mu_0/\bar{\mu}^n$, where n is the largest natural number such that $\pi_0 \leq \bar{\pi}^n$.

Figure 3: Probability of First Attack

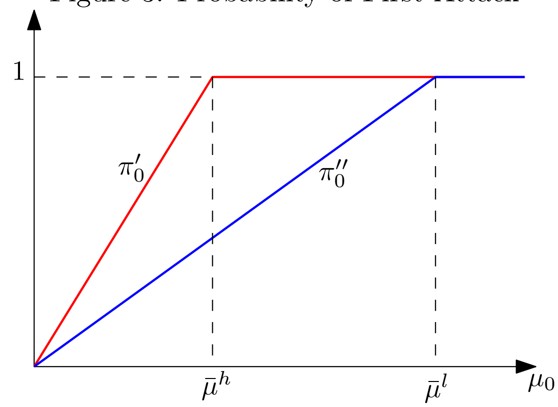

Thus, if π_0 increases, the probability of an armed conflict decreases. Figure 3 depicts the total probability of first attack as a function of μ_0 for $n = h$ (blue line) and $n = l < h$ (red line). Increasing π_0 from π_0' to $\pi_0'' > \pi_0'$ moves the line representing the probability of first attack to the right.

Corollary 2. *Let the government be at least as likely to be committed as the terrorist. Then, the probability that the terrorist starts an armed conflict is decreasing in π_0.*

The advantage of being perceived as tough should not be overstated. Recall that after the first attack, the expected payoff for the government is $-L$, independently of π_0. Indeed, in equilibrium, the government is indifferent between conceding and resisting whenever it plays.

The next result concerns the length of the conflict. We first show that the maximum length of a conflict depends on the absolute likelihood that the terrorist and the government are committed. That is, on the order of commitment n. If the terrorist is more likely to be committed, then there is still an attack with positive probability in period $n+1$. Otherwise, attacks must end with period n. Thus, our model predicts that conflicts between players believed to be fundamentalist are shorter (see Section 3.5) in the sense that they are more likely to end abruptly after few periods. This said, conditional on the conflict protracting after period 1, the probability that it goes on to $t \leq n$ is independent of which player is more likely to be committed or how much the two are likely to be committed. Between period 2 and n, the survival probability of the conflict depends only on the threshold values $\bar{\pi}$ and $\bar{\mu}$. Two conflicts, one with a large and one with small order$n$ are empirically identical from period 2 until they end. Figure 4 depicts the probability of an attack in period t conditional on the conflict protracting at the end of period 1.

they are indistinguishable from conflicts facing players likely to be normal.

Corollary 3. *Unless both players are of the committed type, the maximum length of a conflict is determined by the order of conflict n. If the terrorist is more likely to be committed, there is never an attack after period $n+1$. If the government is at least as likely to be committed as the*

Figure 4: Probability of conflict after period 1 ($n = 90$)

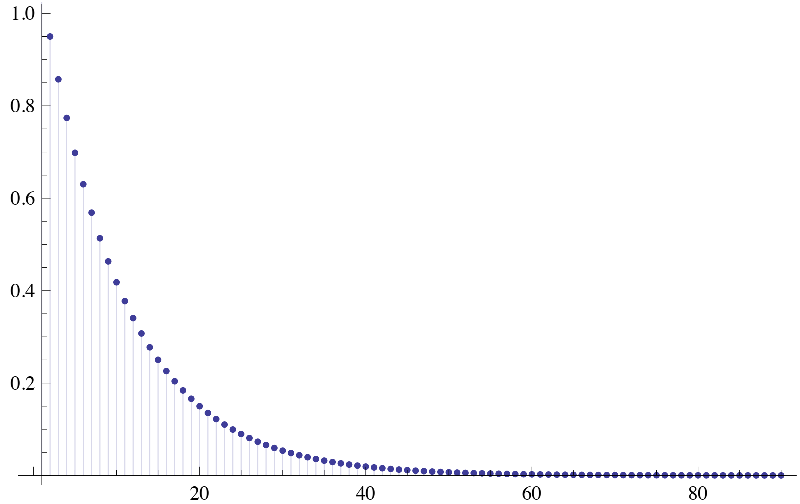

terrorist, there is never an attack after period n*. Conditional on there being a conflict at the end of period* 1*, the probability of an attack in period* $t \in [1, n]$ *is*

$$\bar{\mu}^{t-1}\bar{\pi}^{t-2}.$$

3.5 Commitment and the Length of Conflicts

Terrorist conflicts can last from a few months to many decades. Thus, understanding what determines this variance is of key importance. Figure 5 depicts the number of attacks by year-quarter for four major terrorist groups from 1970 to 2012.[10] All four terrorist groups are ultimately motivated by a separatist goal. Hamas and Lashkar-e-Taiba (LeT) are Islamic terrorist groups originating respectively in Palestine and Afghanistan (active in Pakistan and India). The (Provisional) Irish Republican Army (IRA) and Basque Fatherland and Freedom (ETA) are separatist military groups respectively in Ireland and the Basque Country in Spain and France. Both Hamas and LeT show short burst of conflicts followed by periods of little or no activity. The time-series of Hamas's attacks, in particular, shows clearly three brief periods of conflict corresponding to the First Intifada (from Hamas foundation in 1988 to 1993), the Second Intifada (September 2001 to 2005) and the conflict culminated with the Gaza War (June 2006 to January 2009). In stark contrast, both the IRA and ETA show no such a pattern. The Troubles in Ireland and the UK which begun in 1969 have lasted almost uninterrupted until the late 90s, when the Good Friday Agreements where signed (April 1998). Similarly, ETA's attacks continued uninterrupted from the assassination of of Admiral Luis Carrero Blanco (Franco's chosen successor) in 1973 to the *Spirit of Ermua* massive demonstrations against ETA in 1997.

[10] Our calculations on data from National Consortium for the Study of Terrorism and Responses to Terrorism (START), 2012, Global Terrorism Database [globalterrorismdb_1012dist]. Retrieved from http://www.start.umd.edu/gtd.

Figure 5: Number of Attacks by Quarter of Major Terrorist Groups

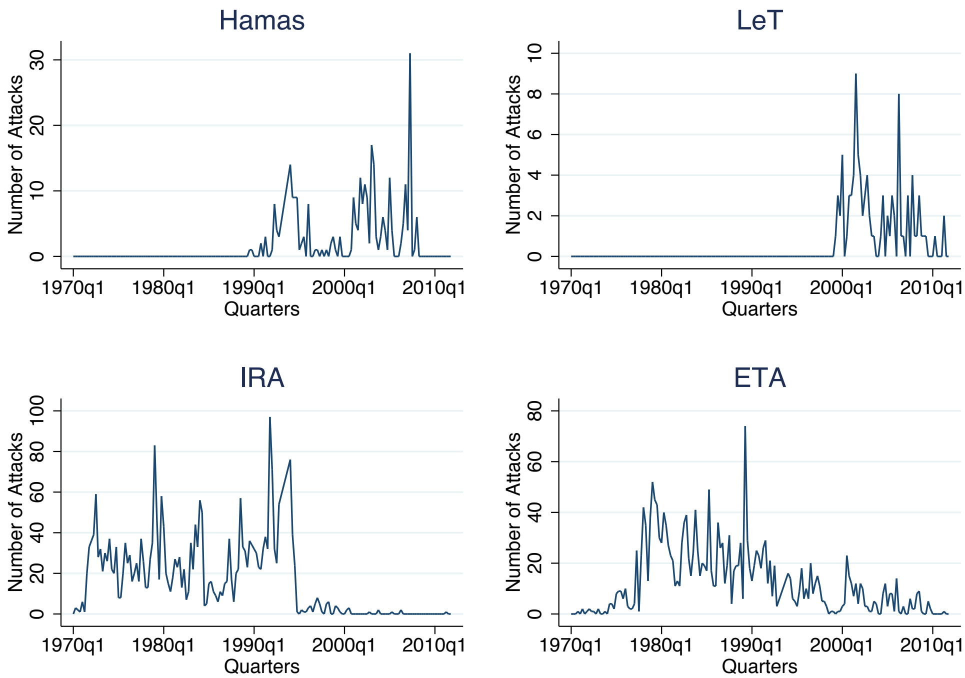

A second conflict started with a car bomb on January 21, 2000.

Our model suggests that shorter and frequent conflicts are typical of organizations which are perceived as committed to an ultimate goal and reject compromise. This result is consistent with the pattern in Figure 5. A common distinction in the literature is the one between religious and secular terrorism.[11] Religious terrorist groups are believed to be fiercely committed to an ultimate goal and reject compromise. For example, Taheri (1987) notes that a key difference between Islamic and secular terrorism is that the first "is clearly conceived and committed as a form of Holy War which can only end when total victory has been achieved" (p. 7) (see also Hoffman, 1995). In the case of Hamas, its charter states that the land of Palestine is "endowed in perpetuity for all generations of Muslims" and it rejects all peace initiatives as "contrary to the beliefs of Islamic resistance." Similarly, LeT's pamphlet *Why Are We Waging Jihad* extends the goal of the organization to the establishment of Islamic rule over the entire Indian subcontinent. On the contrary, both the ETA and the IRA are secular groups whose leadership is commonly held to be political in nature. Military and public officials have been the targets of the bulk attacks by both the IRA and the ETA; in contrast, Hamas and LeT commonly attack civilians, holding them responsible for the occupation of Islamic land.

[11] A notable exception is Pape (2003), who shows that suicide bombings is employed by both religious and secular terrorists.

4 Negotiation in the Baseline Model

We characterize the set of equilibria in the negotiation stage.

First, we show that if the government is at least as likely as the terrorist to be committed, then a conflict is always avoided.

Proposition 3. *If the government is at least as likely as the terrorist to be committed, then in the unique equilibrium the government offers* 0 *and the terrorist accepts. The conflict phase is avoided with probability* 1*.*

In equilibrium, the terrorist's expected payoff of a conflict is 0. To see this, recall from Proposition 1 that in period 1 the terrorist is indifferent between attacking and conceding. Thus, an offer of 0 is sufficient to induce the terrorist to avoid the conflict.

We now turn to the case when the terrorist is more likely than the government to be committed. We show that if the support for the terrorist is sufficiently low when the government makes a small concession, then conflict arises with strictly positive probability.

Let $\bar{\beta} \in [0,1]$ be the offer that would be accepted by a terrorist group with attacking capabilities.[12] Recall that $f(\beta)$ is the probability that the terrorist is capable of building attacking capabilities if the government makes an offer $\beta^G = \beta$.

Proposition 4. *If the terrorist is more likely than the government to be committed, let*

$$\beta^{\dagger} := \frac{\left(\bar{\beta} - \beta\right) U^G}{(1-\beta) U^G + L}.$$

One of two situations below prevails:

1. *if $f(\beta) > \beta^{\dagger}$ for all $\beta < \bar{\beta}$, there exists a unique equilibrium where the government offers $\beta^* = \bar{\beta}$ and the terrorist accepts, and the conflict phase is avoided with probability* 1*;*

2. *if there exists $\beta < \bar{\beta}$ such that $f(\beta) < \beta^{\dagger}$, in equilibrium the government offers $\beta^* < \bar{\beta}$ and the terrorist accepts β^* only if it does not develop attacking capabilities, so that the conflict phase is avoided with probability $1 - f(\beta^*) < 1$.*

The government knows that by offering $\bar{\beta}$, the conflict will be avoided with probability 1. The government thus compares this outcome with what it expects to happen if it offers $\beta < \bar{\beta}$. If the terrorist's consensus is sufficiently large for all such offers, then the government chooses to concede $\bar{\beta}$. Otherwise, it chooses to gamble and make an offer that although not accepted, it significantly reduces the consensus for the terrorist. Thus, negotiations fail not because the

[12] In Appendix B we show that this is

$$\bar{\beta} \equiv 1 - \left(\frac{\pi_0}{\bar{\pi}^n} \frac{1-\bar{\pi}^n}{1-\pi_0} + \frac{b}{U^T} \right).$$

parts are committed to fight, but because the government does not believe that the terrorist will be able to organize into a military force. Crucially, the government would have conceded more had she known with certainty the preferences of the opposition group.

Comparing Propositions 3 and 4, we predict that negotiations systematically fail only when the terrorist is relatively more likely to be a fundamentalist but the government expects to curb the terrorist's support from the opposition when it offers partial concessions. Following our previous interpretation, our results suggest that negotiations are more likely to fail when the terrorist group has strong religious beliefs.

5 Political Wings: a Model of Negotiation and Repression

While the model discussed thus far considers terrorism as more than just conflict, it presents the dissidents as a unified militant group. More often than not, a more political faction and a more militant faction co-exist and represent broadly the same aspirations. They are collaborators and competitors at the same time. This section extends the model to include a political wing P alongside the terrorist group T. We now describe this three-person game.

Conflict Phase Conditional on this phase being triggered, the game is unchanged.

Negotiation Phase This is our point of divergence from the earlier model. At the beginning of period $t = 0$, G can commit to a take-it-or-leave-it offer $\beta^G \in [0, 1]$ to the political wing P, rather than to the terrorist wing T. However this offer is *conditional* on P successfully repressing T. If the political wing accepts the offer, it attempts to repress T. If T is the normal type the attempt is successful, but a fundamentalist ("crazy") militant wing cannot be repressed successfully and the conflict phase starts. The decision to repress must be taken before P knows the type of T.

We view repression as a process of internal conflict that comes at a cost $\phi(\bar{\beta}^T - \beta^G)$, where β^G is the actual offer, and $\bar{\beta}^T = 1 - b\,(1 - \delta)\,/\tau$ is the offer that would to convince a rational T to never attack, independently from initial beliefs π_0 and μ_0; our model of the cost captures the very intuitive feature that if G's initial offer is worse it will take more for P to force T to not launch into an armed struggle. These costs are consistent with more than one concrete interpretation. The parameter $\phi \in \{\phi_S, \phi_W\}$, which we interpret as P's type, is P's private information; with probability $p > 0$, P is *strong* and has low cost of repression $\phi_S \geq 0$; with probability $(1 - p)$, he is *weak* and has high cost of repression $\phi_W > \phi_S$. If P successfully represses the terrorist, then its payoff is the conditional offer that was made by the government minus the actual cost of repression:

$$\beta^G - \phi\left(\bar{\beta}^T - \beta^G\right).$$

Otherwise, if repression is either not attempted or fails, P's payoff is 0.

This completes the description of the negotiation phase and the model.

We assume that G makes the lowest possible offer whenever indifferent and that P always attempts a repression if he is indifferent.

We begin with noting that (each type of) the political wing follows a simple threshold strategy.

Lemma 3. *A weak political wing attempts a repression if and only if*

$$\beta^G \geq \bar{\beta}_W := \frac{\phi_W}{\phi_W + (1-\mu_0)}\bar{\beta}^T.$$

A strong political wing attempts a repression if and only if

$$\beta^G \geq \bar{\beta}_S := \frac{\phi_S}{\phi_S + (1-\mu_0)}\bar{\beta}^T \in (0, \bar{\beta}_W).$$

When the government is at least as likely as the terrorist to be committed, then it knows that even if the political wing does not accept an offer, then T will attack with probability strictly less than 1 (Proposition 1). We show that if $\bar{\beta}_S$ is sufficiently large, then the government offers 0 whenever μ_0 is sufficiently low and 1 otherwise. Intuitively, if the cost of a deal with the political wing is too large, then G prefers to deal directly with the terrorists. If the likelihood that the terrorist is a fundamentalist is very large, then the government decides to avoid a conflict at all costs. Otherwise, it offers nothing to the political wing.

Proposition 5. *Define $\hat{\mu}_n \equiv \bar{\mu}^n \frac{U^G}{U^G+L}$. If the government is at least as likely as the terrorist to be committed and*

$$\bar{\beta}_L > \frac{1-\bar{\mu}^n}{1-\hat{\mu}_n}$$

then the government offers 0 if $\mu_0 \leq \hat{\mu}_n$ and 1 otherwise.

When the $\bar{\beta}_S$ is sufficiently small, the government might be willing to offer a partial concession. If the likelihood that the terrorist is a fundamentalist is very large, then it decides to avoid the conflict at all costs and it offers $\beta^G = 1$. If the likelihood that the terrorist is a fundamentalist is very small, then it offers nothing. Yet, for intermediate values, the government offers a partial concession. When the offer needed to induce a weak political wing to attempt a repression is too high, then the government offers $\beta^G = \bar{\beta}_S$ and a repression is attempted only by strong political wings. If, on the contrary, the offer needed to induce a weak political wing to attempt a recession is sufficiently small, then the government offers $\beta^G = \bar{\beta}_L$ if the likelihood that the terrorist is a fundamentalist is sufficiently large.

Proposition 6. *Let*

$$\hat{\mu}_n := \bar{\mu}^n \frac{U^G}{U^G+L},\ \beta^\dagger := \min\left\{\frac{1-\bar{\mu}^n}{1-\hat{\mu}_n}, \frac{\bar{\beta}_L p \hat{\mu}_n \frac{U^G}{L} + (1-\bar{\mu}^n)(1-p)}{\hat{\mu}_n \frac{U^G}{L} + \bar{\mu}^{n(1-p)}}\right\},$$

$$\mu^{\dagger 1} := \frac{\hat{\mu}_n}{\hat{\mu}_n + \frac{1-\bar{\mu}^n}{\bar{\beta}_L}} < \mu^{\dagger 2} := \frac{\hat{\mu}_n}{\hat{\mu}_n + \frac{(1-\bar{\mu}^n)(1-p)}{\bar{\beta}_H - p\bar{\beta}_L}} < \mu^{\dagger 3} := \frac{\hat{\mu}_n}{\hat{\mu}_n \left(1 + \frac{U^G}{L\left(1-\bar{\beta}_H\right)}\right)}.$$

If the terrorist is more likely than the government to be committed, then

1. *if* $\bar{\beta}_L > \frac{1-\bar{\mu}^n}{1-\hat{\mu}_n}$, *G offers* 0 *for* $\mu_0 \leq \hat{\mu}_n$ *and everything for* $\mu_0 > \hat{\mu}_n$*;*
2. *if* $\bar{\beta}_L < \frac{1-\bar{\mu}^n}{1-\hat{\mu}_n}$, *G offers* 0 *if* $\mu_0 \leq \frac{\hat{\mu}_n}{\hat{\mu}_n + \frac{1-\bar{\mu}^n}{\bar{\beta}_L}}$ *and*
   (a) *if* $\bar{\beta}_H \geq \beta^\dagger$, *G offers* $\bar{\beta}_L$ *if* $\mu_0 \leq \hat{\mu}_n \frac{1-\hat{\beta}_L p}{1-\left(1-\bar{\mu}^n \bar{\beta}_L\right)p}$ *and everything otherwise;*
   (b) *if* $\bar{\beta}_H \leq \beta^\dagger$, *G offers* $\bar{\beta}_L$ *if* $\mu_0 \in \left[\mu^{\dagger 1}, \mu^{\dagger 2}\right)$, *offers* $\bar{\beta}_H > \bar{\beta}_L$ *if* $\mu_0 \in \left[\mu^{\dagger 2}, \mu^{\dagger 3}\right)$ *and everything otherwise.*

6 Conclusions

This paper provides a game-theoretic framework for studying terrorism as violence with the purpose of intimidation, with ultimate political goals. It proposes the use of a reputational model to study the phenomenon: players are very likely rational actors, who exploit the slight uncertainty about their type to build a reputation for not conceding until their demands are met. It shows how such a model can be used to derive realistic results regarding the probability and length of an armed conflict. In particular, dissident groups that are more likely to be committed to fighting until their final victory are more likely to engage in terrorism, but the conflicts they start are likely to be short.

The strategy of religious terrorists is commonly believed to differ from that of secular ones. We provide anecdotal evidence of how the first is associated with frequent short conflicts, while the latter gives rise to longer conflicts followed by lasting periods of peace. Associating religious terrorism with a higher prior probability of being committed, the result mentioned above provides a novel explanation of how and why the strategy of religious terrorism differs from secular terrorism.

References

[1] **Abreu, Dilip and Faruk Gul (2000).** Bargaining and Reputation, *Econometrica,* 68(1): 85–117.

[2] **Baliga, Sandeep and Thomas Sjöstrom (2012).** The Strategy of Manipulating Conflict, *The American Economic Review*, 102(6): 2897–2922.

[3] **Berman, Eli and David D. Laitin (2008).** Religion, Terrorism and Public Goods: Testing the Club Model, *Journal of Public Economics*, 92: 1942–1967.

[4] **Bueno de Mesquita, Ethan (2005).** Conciliation, Counterterrorism, and Patterns of Terrorist violence, *Industrial Organization*, 59: 149–176.

[5] **Bueno de Mesquita, Ethan (2008).** The Political Economy of Terrorism: A Selective Overview of Recent Work, *The Political Economist*, 10(1): 1–12.

[6] **Bueno de Mesquita, Ethan and Eric S. Dickson (2007).** The Propaganda of the Deed: Terrorism, Counterterrorism, and Mobilization, *American Jounral of Political Science*, 51(2): 364–381.

[7] **Crawford, Vincent P. (1982).** A Theory of Disagreement in Bargaining, *Econometrica*, 50(3): 607–637.

[8] **Enders, Walter and Todd Sandler (2012).** *The Political Economy of Terrorism*, 2nd Ed. Cambridge: Cambridge University Press.

[9] **Gould, Eric D. and Esteban F. Klor (2010).** Does Terrorism Work?, *The Quarterly Journal of Economics*, 125(4): 1459–1510.

[10] **Hodler, Roland and Dominic Rohner (2012).** Electoral Terms and Terrorism, *Public Choice*, 150(1): 181–153.

[11] **Krueger, Alan B. (2007).** *What Makes a Terrorist*, Princeton University Press.

[12] **Lapan, Harvey E. and Todd Sandler (1988).** To bargain or not to bargain: That is the question, *The American Economic Review*, 78(2): 16–21.

[13] **Lapan, Harvey E. and Todd Sandler (1993).** Terrorism and signalling, *European Journal of Political Economy*, 9(3): 383–397.

[14] **Hoffman, Bruce (2006).** *Inside Terrorism*, Columbia University Press; Revised & enlarged edition.

[15] **Hoffman, Bruce (1995).** "Holy Terror": The Implications of Terrorism Motivated by Religious Imperatives, *Studies in Conflict & Terrorism*, 18(4): 271–284.

[16] **Pape, Robert A. (2003).** The Strategic Logic of Suicide Terrorism, *American Political Science Review*, 97(3): 343–361.

[17] **Primoratz, Igor (2012).** *Terrorism: A Philosophical Investigation*, Wiley.

[18] **Richardson, Louise (2006).** *What Terrorists Want: Understanding the Enemy, Containing the Threat*, Random House.

[19] **Siqueira, Kevin (2005).** Political and Militant Wings within Dissident Movements and Organizations, *Journal of Conflict Resolution*, 49(2): 218–236.

[20] **Schelling, Thomas C. (1956).** An Essay on Bargaining, *American Economic Review*, 6: 281–306.

[21] **Schelling, Thomas C. (1960).** *The Strategy of Conflict*, Cambridge, MA: Harvard University Press.

[22] **Schelling, Thomas C. (1966).** *Arms and Influence*, New Haven, CT: Yale University Press.

[23] **Taheri, Amir (1987).** *Holy Terror: The Inside Story of Islamic Terrorism*, Sphere Books, London.

[24] Yared, Pierrre (2010). *A Dynamic Theory of War and Peace.*

Omitted Proofs

A Proofs of Section 3

We characterize the unique equilibrium of the conflict phase. We prove Propositions 1 and 2 as corollaries of two theorems that will be stated later on in the section. While the structure of the game is potentially very complex, we find that all equilibria are marked by a few intuitive properties. We begin with a few preliminary results. Recall that Assumptions 1 and 2 are in force. Every strategy profile, equilibrium or otherwise, can be represented as a sequence of probabilities as formalised in the definition below.

Definition 3. The *concession sequence* $\langle c_k \rangle_{k \in \mathbb{N}}$ of any strategy profile is a sequence of numbers in $[0, 1]$, where each odd (even) term is the probability that T (respectively, G) concedes at that time conditional on the game continuing to that point. A concession sequence arising from an equilibrium profile is an called an *equilibrium concession sequence.*

This construct helps us prove uniqueness. Lemma 4 shows that in equilibrium a zero probability of concession is possible only at the initial history; all other histories are marked by a strictly positive concession probability.

Lemma 4. *In any equilibrium concession sequence, all terms (except possibly the first) must be strictly positive.*

Proof. STEP 1. The proof is based on the key idea that if the string $(c_k, 0, c_{k+2})$ appears in an equilibrium concession sequence and $c_{k+2} > 0$, then $c_k = 1$: If the opponent is not conceding in the interim the value of concession can only go down because there is positive cost to fighting; therefore concession should have been strictly better at the step before.

STEP 2. We now show that, along any concession sequence, adjacent terms cannot be 0. Let $c_k = 0 = c_{k+1}$; if $c_{k+2} > 0$, it would contradict Step 1. Induction implies that if two adjacent terms of the concession sequence are 0, all subsequent terms are 0 too. But since there is a positive probability of the commitment type, it cannot be an equilibrium to never concede, knowing that your opponent will not. Therefore, no equilibrium concession sequence contains adjacent 0s.

STEP 3. Suppose $c_k = 0$ for some $k > 1$. By Step 2, we must have $c_{k+1} > 0$; from Step 1 it means that $c_{k-1} = 1$. If the player who is supposed to concede with probability 1 does not do so, his reputation immediately jumps to 1 and his opponent must concedes immediately thereafter with certainty, i.e. $c_k = 1$ — a contradiction! □

Remark 2. Lemma 4 does not apply to the first term, i.e. $k = 1$, because there is no prior term, i.e. there is no c_0.

Lemma 1 in Section 3 says that if one of the two players is of the commitment type with sufficiently high (prior) probability, then the other player must concede immediately. The next lemma says that beliefs cannot jump from below the cutoff points $\bar{\mu}$ and $\bar{\pi}$ to above these without touching at least one of the two cutoff points. In other words if normal types finish conceding at time n, it must be the case that at least one reputation is exactly at its cutoff i.e. either $\pi_n = \bar{\pi}$ or $\mu_n = \bar{\mu}$.

Lemma 5. *In equilibrium (i) $\pi_t < \bar{\pi}$ and $\mu_{t+1} < \bar{\mu}$ implies $\pi_{t+1} \leq \bar{\pi}$; (ii) $\mu_t < \bar{\mu}$ and $\pi_t < \bar{\pi}$ implies $\mu_{t+1} \leq \bar{\mu}$.*

Proof. Suppose not. Let $\pi_t < \bar{\pi}$, $\mu_{t+1} < \bar{\mu}$ but $\pi_{t+1} > \bar{\pi}$. Lemma 1 implies that the normal type of T will concede w.p. 1 at time $t+2$ if G does not concede at $t+1$. So if G does not concede at time $t+1$ it gets a continuation payoff of γ from $t+2$ onwards if T is the normal type; since T is normal with probability $1-\mu_{t+1}$, G's payoff from $t+1$ (the current period) onwards is

$$(1-\delta)\{-L+\gamma\} + \delta\{(1-\mu_{t+1})\gamma + \mu_{t+1}(-L(1-\delta) + 0 \cdot \delta)\}.$$

G strictly prefers to not concede if the above exceeds the payoff $-(1-\delta)L$ from conceding immediately at $t+1$:

$$(1-\delta)\gamma + \delta\{(1-\mu_{t+1})\gamma - L(1-\delta)\mu_{t+1}\} > 0. \tag{10}$$

Inequality (10) reduces to

$$\mu_{t+1} < \frac{\gamma}{\delta\{L(1-\delta)+\gamma\}} \equiv \bar{\mu},$$

which is true by assumption. Therefore G strictly prefers to fight at $t+1$, i.e. $\sigma^G_{t+1}(\mu_{t+1}, \pi_t) = 0$ — which contradicts Lemma 4, implying that $\pi_t < \bar{\pi}$ and $\mu_{t+1} < \bar{\mu}$ cannot lead to $\pi_{t+1} > \bar{\pi}$.

Let $\mu_t < \bar{\mu}$ and $\pi_t < \bar{\pi}$, but $\mu_{t+1} > \bar{\mu}$. By a similar logic T strictly prefers to fight at $t+1$ if

$$(1-\delta)(-b) + (1-\pi_t)\tau + \pi_t \cdot 0 > 0.$$

The net utility for T to fight at period t is $-b$ and with probability $1-\pi_t$ the government will concede and T will get τ forever. The expression above reduces to $\pi_t < 1 - \frac{b(1-\delta)}{\tau} \equiv \bar{\pi}$. So T strictly prefers to fight at $t+1$, i.e. $\sigma_{t+1}^{T}(\mu_t, \pi_t) = 0$; this contradicts Lemma 4. $\square$

Note that all but the first move by T have a strictly positive probability of concession given by equations (8) and (9).

The next lemma shows that along the equilibrium path, provided no one concedes, both reputations grow according to equations (1) and (2) from period 2 onwards until a time t when either $\mu_t = \bar{\mu}$ or $\pi_t = \bar{\pi}$.

Lemma 6. *For any period $t \geq 2$, if $\pi_{t-1} \leq \bar{\pi}$ and $\mu_t \leq \bar{\mu}$, then G plays $\tilde{\sigma}_t^G(\pi_{t-1})$ and T plays $\tilde{\sigma}_t^T(\mu_t)$.*

Proof. We show the result for G. The result for T follows a symmetric argument.

Procede by contradiction. If $\sigma_t^G(\pi_t) \neq \sigma_t^{G*}(\pi_t)$, by Lemma 2, T is not indifferent at either t or at $t+1$. There are two possibilities. First, she strictly prefers to concede. But then G would concede with probability 0 in the previous period, contradicting Lemma 4. Second, she strictly prefers to fight. But then by Lemma 4 she is T in period 0 and $t = 1 < 2$. $\square$

The lemmas above are useful in proving Theorems 1 and 2, which apply, respectively, to the cases where G is at least as committed as T and where T is more committed than G.

Theorem 1. *If there exists $n \in \mathbb{N}$ such that*

$$\mu_0 < \bar{\mu}^n \text{ and } \bar{\pi}^{n+1} < \pi_0 < \bar{\pi}^n \tag{11}$$

there is a unique equilibrium σ^* where
(i) in period 1, T concedes with probability

$$\sigma_1^{T*}(\mu_0, \pi_0) = 1 - \frac{\mu_0}{\bar{\mu}^n}\frac{1-\bar{\mu}^n}{1-\mu_0}; \tag{12}$$

(ii) after that both players concede with probabilities

$$\sigma_{t+1}^{G*}(\mu_{t+1}, \pi_t) = \begin{cases} \frac{1-\bar{\pi}}{1-\pi_t} & \text{if } \pi_t \leq \bar{\pi} \\ 1 & \text{if } \pi_t \geq \bar{\pi} \end{cases}; \quad \sigma_{t+1}^{T*}(\mu_t, \pi_t) = \begin{cases} \frac{1-\bar{\mu}}{1-\mu_t} & \text{if } \mu_t \leq \bar{\mu} \\ 1 & \text{if } \mu_t \geq \bar{\mu} \end{cases};$$

(iii) the path of beliefs absent concession is

$$\left(\mu_0, \pi_0; \bar{\mu}^n, \frac{\pi_0}{\bar{\pi}}; \bar{\mu}^{n-1}, \frac{\pi_0}{\bar{\pi}^2}; \cdots; \bar{\mu}, \frac{\pi_0}{\bar{\pi}^n}; 1, 1; 1, 1; \cdots\right).$$

Proof. Existence. It can be checked that Lemmas 1 and 2 imply that the above is an equilibrium. In particular, σ_1^{T*} and Bayes' rule imply that the equilibrium belief about T's type after non-concession at $t = 1$ is given by $\bar{\mu}^n$.

Uniqueness. If $\mu_0 \geq \bar{\mu}$, then Lemma 1 implies that the above is the only equilibrium; similarly for the case $\pi_0 \geq \bar{\pi}$. Therefore let $(\mu_0, \pi_0) < (\bar{\mu}, \bar{\pi})$, so that $n \geq 1$. If normal types follow $\tilde{\sigma}$ defined in equations (8) and (9) up to and including time n, there will be a jump since $\pi_0/\bar{\pi}^n > \bar{\pi}$; but jumps are ruled out by Lemma 5. By Lemmas 2 and 6, the only freedom we have is in choosing different strategies for $t = 1$.

Case 1: $\sigma_1^T < \sigma_1^{T*}$. Suppose that $\sigma_1^T < \sigma_1^{T*}$. The inequality $\sigma_1^T < \sigma_1^{T*}$ implies that T's reputation increases at a slower rate such that $\mu_n < \bar{\mu}$. If $\sigma_1^G < \tilde{\sigma}_1^G$, then T prefers to concede immediately ($\sigma_1^T = 1$) since T is just indifferent at $\tilde{\sigma}_1^G$; this contradiction implies that $\sigma_1^G \geq \tilde{\sigma}_1^G$, which in turn gives $\pi_1 \geq \pi_0/\bar{\pi}$ and therefore $\pi_n > \bar{\pi}$ i.e. there exists $m \leq n$ such that belief profile is (μ_m, π_m) with $\mu_m < \bar{\mu}$ and $\pi_m > \bar{\pi}$, contradicting Lemma 5.Therefore, $\sigma_1^T \geq \sigma_1^{T*}$ is the only possibility in equilibrium.

Case 2: $\sigma_1^T > \sigma_1^{T*}$. Suppose that $\sigma_1^T > \sigma_1^{T*}$. Now $\mu_1 > \mu_0/\bar{\mu}$, $\mu_2 > \mu_0/\bar{\mu}^2$, etc. Since Proposition 1 implies that G's reputation is growing as the same rate $1/\bar{\pi}$ it follows from inequality (11) and $\bar{\mu}^{n+1} < \mu_0$ that $\mu_n > \bar{\mu}$, i.e. a jump occurs by time n. Therefore, $\sigma_1^T \leq \sigma_1^{T*}$ is the only possibility in equilibrium.

Last, since T must be indifferent at $t = 0$ to play σ_1^{T*}, then $\sigma_1^G = \tilde{\sigma}_1^G = \sigma_1^{G*}$. □

The next theorem considers the case where T is more committed than G.

Theorem 2. *If there exists $n \in \mathbb{N}$ such that*

$$\bar{\mu}^{n+1} < \mu_0 < \bar{\mu}^n \text{ and } \pi_0 < \bar{\pi}^{n+1}, \tag{13}$$

then there is a unique equilibrium σ^* where
(i) in period 1, T attacks with probability 1 and G concedes with probability

$$\sigma_1^{G*}(\mu_1, \pi_0) = 1 - \frac{\pi_0}{\bar{\pi}^n} \frac{1 - \bar{\pi}^n}{1 - \pi_0}; \tag{14}$$

(ii) subsequently, both players mix with probabilities

$$\sigma_{t+1}^{G*}(\mu_{t+1}, \pi_t) = \begin{cases} \frac{1-\bar{\pi}}{1-\pi_t} & \text{if } \pi_t \leq \bar{\pi} \\ 1 & \text{if } \pi_t \geq \bar{\pi} \end{cases}; \quad \sigma_{t+1}^{T*}(\mu_t, \pi_t) = \begin{cases} \frac{1-\bar{\mu}}{1-\mu_t} & \text{if } \mu_t \leq \bar{\mu} \\ 1 & \text{if } \mu_t \geq \bar{\mu} \end{cases};$$

(iii) starting with the second stage of period $n + 1$ (i.e. G's move at $n + 1$) all normal types concede with probability 1 if the game has not ended, and the path of beliefs absent concession is

$$\left(\mu_0, \pi_0; \mu_0, \bar{\pi}^n; \frac{\mu_0}{\bar{\mu}}, \bar{\pi}^{n-1}; \cdots ; \frac{\mu_0}{\bar{\mu}^{n-1}}, \bar{\pi}; \frac{\mu_0}{\bar{\mu}^n}, 1; 1, 1; \cdots \right).$$

Proof. Existence. It can be checked that Lemmas 1 and 2 imply that the above is an equilibrium. In particular, σ_1^{T*} and Bayes' rule imply that the equilibrium belief about T's type after non-concession at $t = 1$ is given by $\bar{\mu}^n$.

Uniqueness. If $\mu_0 \geq \bar{\mu}$, then Lemma 1 implies that the above is the only equilibrium; similarly for the case $\pi_0 \geq \bar{\pi}$. Therefore let $(\mu_0, \pi_0) < (\bar{\mu}, \bar{\pi})$, so that $n \geq 1$. If normal types follow $\tilde{\sigma}$ defined in equations (8) and (9) up to and including time n, there will be a jump since $\pi_0/\bar{\pi}^n > \bar{\pi}$; but jumps are ruled out by Lemma 5. By Lemmas 2 and 6, the only freedom we have is in choosing different strategies for $t = 1$.

By contradiction, suppose that T concedes with positive probability in period 1. This implies she expects G to conceding with probability at least $\tilde{\sigma}_1^G$. But this implies that there is $m \leq n$ such that beliefs are (μ_{m+1}, π_m) with $\mu_{m+1} > \bar{\mu}$ and $\pi_m < \bar{\pi}$, contradicting Lemma 5.

Last, since T cannot concede with probability less than 0, we have that $\mu_{n+1} > \bar{\mu}$. Thus, by Lemma 5, G must concede in period 1 with probability exactly σ_t^{G*}. □

B Proofs of Section 4

Proof of Proposition 3. From Proposition 1, if the government is at least as likely as the terrorist to be committed, then G's expected payoff of a conflict is $U^G - (U^G + L)\mu_0/\bar{\mu}^n$; T's expected payoff of a conflict is 0.

First, we consider T's play. Accepting all offers (i) weakly dominates all other strategies, and (ii) strictly dominates all strategies which do not accept some strictly positive offer. To see this, let $\beta^G \geq 0$ be G's offer. A terrorist that does not develop attacking capabilities accepts all offers. A terrorist that develops attacking capabilities accepts an offer if

$$\begin{aligned} \beta^G U^T &\geq 0 = \text{ expected payoff of conflict} \\ \iff \beta^G &\geq 0. \end{aligned}$$

Second, we show that T accepting all offers and G offering $\beta^G = 0$ is an equilibrium. To see this, let $\beta^G = 0$. From the previous step, T has a best response in accepting all offers. G's best response is to offer $\beta^G = 0$, since G's expected payoff of a conflict is strictly less than U^G.

Last, we show that the equilibrium is unique. To see this, let T play the only other undominated strategy: accept with probability 1 if and only if the offer is strictly positive; accept with probability $\xi \in (0, 1)$ if $\beta^G = 0$. Then, G's best response is $\inf\left\{\beta^G \in \mathbb{R}_+ : \beta^G > 0\right\}$. □

Proof of Proposition 4. From Proposition 2, is more likely than the government to be committed, then G's expected payoff of a conflict is $-L$; T's expected payoff of a conflict is

$$U^T \left(1 - \frac{\pi_0}{\bar{\pi}^n} \frac{1 - \bar{\pi}^n}{1 - \pi_0}\right) - b.$$

Let $\bar{\beta} \in [0,1]$ be the offer that would be accepted by a terrorist group with attacking capabilities. That is,

$$\bar{\beta} \equiv 1 - \left(\frac{\pi_0}{\bar{\pi}^n}\frac{1-\bar{\pi}^n}{1-\pi_0} + \frac{b}{U^T}\right),\ \beta^{\dagger} = \frac{\left(\bar{\beta}-\beta\right)U^G}{\left(1-\beta\right)U^G + L}.$$

The expected payoff for G of an offer β^G is therefore given by:

$$\begin{cases} \left(1-f\left(\beta^G\right)\right)\left(1-\beta^G\right)U^G - f\left(\beta^G\right)L & \text{if } \beta^G < \bar{\beta}; \\ \left(1-\beta^G\right)U^G & \text{otherwise.} \end{cases}$$

We divide the remaining of the proof in two cases.

Case 1: If there exists $\beta < \bar{\beta}$ such that $f(\beta) < \beta^{\dagger}$, then G's best response is to offer

$$\beta^* = \max_{\beta < \bar{\beta}} \left[\left(1-f\left(\beta^G\right)\right)\left(1-\beta^G\right)U^G - f\left(\beta^G\right)L\right].$$

Thus, in equilibrium, G offers $\beta^* < \bar{\beta}$, and a terrorist group with attacking capabilities accept (i) all offers $\beta^G > \bar{\beta}$ with probability 1, and (ii) offer $\beta^G = \bar{\beta}$ with probability $\xi \in [0,1]$. Notice that the equilibrium is essentially unique and a conflict is not avoided with probability $f(\beta^*)$.

Case 2: If $f(\beta) > \beta^{\dagger}$ for all $\beta < \bar{\beta}$ then G's best response is (i) offer $\bar{\beta}$ if T accepts all offers $\beta^G \geq \bar{\beta}$ with probability 1, or (ii) $\inf\left\{\beta^G \in \mathbb{R}_+ : \beta^G > \bar{\beta}\right\}$ if T accepts all offers $\beta^G > \bar{\beta}$ with probability 1 and offer $\beta^G = \bar{\beta}$ with probability $\xi \in (0,1)$. Thus, in the unique equilibrium, G offers $\bar{\beta}$ and T accepts all offers $\beta^G \geq \bar{\beta}$. A conflict is avoided with probability 1.

□

C Proofs of Section 5

Proof of Lemma 3. The expected payoff for P of type ϕ of attempting a repression is $(1-\mu_0)\beta^G - \phi\left(\bar{\beta}^T - \beta^G\right)$. If he does not repress, he gets 0. Thus, he attempts a repression if and only if $\beta^G \geq \frac{\phi}{\phi + (1-\mu_0)}\bar{\beta}^T$. □

Proof of Proposition 5. Let $n \in \mathbb{N}$ be such that $\bar{\pi}^{n+1} < \pi_0 < \bar{\pi}^n$ and $\mu_0 > \bar{\mu}^n$. The expected payoff for G is given by

$$\begin{aligned}
\text{if } \beta^G = 1, &\quad \text{then} \quad 0; \\
\text{if } \beta^G \in \left[\bar{\beta}_W, 1\right), &\quad \text{then} \quad \left(1 - \bar{\beta}_W\right) U^G - \mu_0 L; \\
\text{if } \beta^G \in \left[\bar{\beta}_S, \bar{\beta}_W\right), &\quad \text{then} \quad p \left[\left(1 - \bar{\beta}_S\right) U^G - \mu_0 L\right] - (1-p) L; \\
\text{if } \beta^G \in \left[0, \bar{\beta}_S\right), &\quad \text{then} \quad -L.
\end{aligned}$$

It is easy to see that $\beta^G < \bar{\beta}_S$ is dominated by $\beta^G = 1$. Also, we have to concentrate only on $\bar{\beta}_S$, $\bar{\beta}_W$, and 1 since the lower bound of each of the four intervals dominates any other value in the interval.

By comparing the expected payoffs above, we conclude that G prefers $\bar{\beta}_W$ to 1 if and only if

$$\mu_0 \leq \frac{\left(1 - \bar{\beta}_W\right) U^G}{\left(1 - \bar{\beta}_W\right) U^G + L}.$$

Similarly, we can show that G prefers $\bar{\beta}_S$ to 1 if and only if

$$\mu_0 \leq \frac{\left(1 - \bar{\beta}_S\right) U^G - \frac{1-p}{p} L}{\left(1 - \bar{\beta}_W\right) U^G + L}$$

and $\bar{\beta}_W$ to $\bar{\beta}_S$ if and only if

$$p \geq \frac{\left(1 - \bar{\beta}_W\right) U^G + L}{\left(1 - \bar{\beta}_S\right) U^G + L}.$$

Comparing these three thresholds gives the proposition. $\square$

Proof of Proposition 6. Let $n \in \mathbb{N}$ be such that $\bar{\pi}^{n+1} < \pi_0 < \bar{\pi}^n$ and $\mu_0 < \bar{\mu}^n$. Define $\hat{\mu}_n \equiv \bar{\mu}^n \frac{U^G}{U^G + S}$. The expected payoff for G is given by

$$\begin{aligned}
\text{if } \beta^G = 1, &\quad \text{then} \quad 0; \\
\text{if } \beta^G \in \left[\bar{\beta}_W, 1\right), &\quad \text{then} \quad \left(1 - \bar{\beta}_W\right) U^G - \mu_0 L; \\
\text{if } \beta^G \in \left[\bar{\beta}_S, \bar{\beta}_W\right), &\quad \text{then} \quad p \left[\left(1 - \bar{\beta}_S\right) U^G - \mu_0 L\right] - (1-p) \left[U^G - \frac{\mu_0}{\bar{\mu}^n} \left(U^G + L\right)\right]; \\
\text{if } \beta^G \in \left[0, \bar{\beta}_S\right), &\quad \text{then} \quad U^G - \frac{\mu_0}{\bar{\mu}^n} \left(U^G + L\right).
\end{aligned}$$

It is easy to see that we have to concentrate only on 0, $\bar{\beta}_S$, $\bar{\beta}_W$, and 1 since the lower bound of each of the four intervals dominates any other value in the interval.

By comparing the expected payoffs above, we conclude that G prefers everything to 0 if and only if

$$\mu_0 > \hat{\mu}_n.$$

Similarly, $\bar{\beta}_S$is better than 0 if and only if

$$\mu_0 > \frac{\hat{\mu}_n}{\hat{\mu}_n + \frac{1-\bar{\mu}^n}{\bar{\beta}_S}};$$

and $\bar{\beta}_W$ is better than 0 if and only if

$$\mu_0 > \frac{\hat{\mu}_n}{\hat{\mu}_n + \frac{1-\bar{\mu}^n}{\bar{\beta}_W}}.$$

We can immediately notice that if $\bar{\beta}_W$ is better than 0, then $\bar{\beta}_S$ is also better than 0.

Also, $\bar{\beta}_W$ is better than $\bar{\beta}_S$ if and only if

$$\mu_0 > \frac{\hat{\mu}_n}{\hat{\mu}_n + \frac{(1-\bar{\mu}^n)(1-p)}{\bar{\beta}_W - p\bar{\beta}_S}};$$

$\bar{\beta}_W$ is better than 1 if and only if

$$\mu_0 \leq \frac{\hat{\mu}_n}{\hat{\mu}_n \left(1 + \frac{U^G}{S\left(1-\bar{\beta}_W\right)}\right)};$$

and finally $\bar{\beta}_S$ is better than 1 if and only if

$$\mu_0 \leq \hat{\mu}_n \frac{1 - \hat{\beta}_S p}{1 - \left(1 - \bar{\mu}^n \bar{\beta}_S\right) p}.$$

Define

$$\beta^{\ddagger} = \frac{1 - \bar{\mu}^n}{1 - \bar{\mu}^n \frac{U^G}{U^G + L}}.$$

Hence, if both $\bar{\beta}_W$ and $\bar{\beta}_S$ are greater than $\beta^{\ddagger}$, then 1 is better than both of them whenever 0 is not optimal. This shows the first part of the proposition. Instead, if only $\bar{\beta}_S$ is less or equal to $\beta^{\ddagger}$, then there exist intermediate values of μ_0 for which $\bar{\beta}_S$ is the optimal choice. This is between the thresholds for $\bar{\beta}_S$ better than 0 and for 1 better than $\bar{\beta}_S$. Finally, $\bar{\beta}_W$ is optimal if and only if it is less than $\beta^{\ddagger}$ and it is better than $\bar{\beta}_S$ and 1. For this to be true, we need that the threshold for $\bar{\beta}_W$ better than $\bar{\beta}_S$ is smaller than the threshold for $\bar{\beta}_W$ better than 1. This implies that

$$\bar{\beta}_W < \frac{\bar{\beta}_S p \hat{\mu}_n \frac{U^G}{S} + (1 - \bar{\mu}^n)(1 - p)}{\hat{\mu}_n \frac{U^G}{S} + \bar{\mu}^{n(1-p)}} = \beta^{\dagger}.$$

This proves the second part of the proposition. $\square$