

Terrorism Works, for Its Supporters

Online Appendix

Andrew J. Coe, Peter Schram, and Heesun Yoo

April 2, 2024

Contents

I	Proofs for Results in the Paper	2
1	Proof of the Conditions in Assumption 1 and 2	2
2	Proof of Proposition 1	4
II	Model Extensions	8
3	Uncertainty about Interests	8
4	Uncertainty about Actions	12
5	Decisive Government Victory	13

Part I

Proofs for Results in the Paper

1 Proof of the Conditions in Assumption 1 and 2

We first prove that the condition stated in Assumption 1 ensures that an unsupported V will not prefer to attack if G sets x in every period, even if $x = 0$. Obviously if V attacks and wins, then in any equilibrium it will set its ideal policy $x = \hat{x}_V$ in every period. Comparing the continuation values of attacking and not, we have:

$$\begin{aligned} -c_V + P(0)\frac{-|\hat{x}_V - \hat{x}_V|}{1 - \delta} + [1 - P(0)]\frac{-|x - \hat{x}_V|}{1 - \delta} &< \frac{-|x - \hat{x}_V|}{1 - \delta} \\ \Leftrightarrow P(0)\frac{|x - \hat{x}_V|}{1 - \delta} &< c_V \end{aligned}$$

Clearly the left side falls in x , so if the inequality is satisfied at $x = 0$, it will be satisfied for any value of $x \in [0, \hat{x}_V]$. Thus the condition stated in the assumption guarantees V will not attack without support, even if G makes no policy concession.

Next we show that the condition stated in Assumption 2 ensures that S will not support V until it wins if G sets x in every period, even if $x = 0$. The continuation value of supporting V at the level f until it wins, assuming that f suffices to motivate V to attack, is:

$$\begin{aligned} CV_S^{strong} &= -f - c_S + P(f)\frac{-|\hat{x}_V - \hat{x}_S|}{1 - \delta} + [1 - P(f)](-|x - \hat{x}_S| + \delta CV_S^{strong}) \\ \Leftrightarrow CV_S^{strong} &= \frac{1}{1 - [1 - P(f)]\delta} \left(-f - c_S + P(f)\frac{-|\hat{x}_V - \hat{x}_S|}{1 - \delta} - [1 - P(f)]|x - \hat{x}_S| \right) \end{aligned}$$

Setting this to be less than the value of never supporting V , we obtain:

$$\begin{aligned} \frac{1}{1 - [1 - P(f)]\delta} \left(-f - c_S + P(f) \frac{-|\hat{x}_V - \hat{x}_S|}{1 - \delta} - [1 - P(f)] |x - \hat{x}_S| \right) &< \frac{-|x - \hat{x}_S|}{1 - \delta} \\ \Leftrightarrow -f - c_S + P(f) \frac{-|\hat{x}_V - \hat{x}_S|}{1 - \delta} - [1 - P(f)] |x - \hat{x}_S| &< -|x - \hat{x}_S| - \frac{P(f)\delta}{1 - \delta} |x - \hat{x}_S| \\ \Leftrightarrow \frac{P(f)}{1 - \delta} (|x - \hat{x}_S| - |\hat{x}_V - \hat{x}_S|) &< f + c_S \end{aligned}$$

If $x > \hat{x}_S$, the quantity in parentheses is negative and the inequality is trivially satisfied. If $x \leq \hat{x}_S$, it is equivalent to the condition stated in the assumption for S .

If S is to obtain this continuation value, it must be that f suffices to motivate V to attack, given that G is setting x and V is receiving f and attacking in every round until it wins. V 's continuation value of attacking with support f in every round until it wins is:

$$\begin{aligned} CV_V^{strong} &= -c_V + \zeta f + P(f) \frac{-|\hat{x}_V - \hat{x}_V|}{1 - \delta} + [1 - P(f)] (-|x - \hat{x}_V| + \delta CV_V^{strong}) \\ \Leftrightarrow CV_V^{strong} &= \frac{1}{1 - [1 - P(f)]\delta} (-c_V + \zeta f - [1 - P(f)] |x - \hat{x}_V|) \end{aligned}$$

Comparing this value to that of deviating to not attacking for one period, we have:

$$\begin{aligned} CV^{strong} &\geq -|x - \hat{x}_V| + \delta CV_V^{strong} \\ \Leftrightarrow -c_V + \zeta f - [1 - P(f)] |x - \hat{x}_V| &\geq -\frac{1 - [1 - P(f)]\delta}{1 - \delta} |x - \hat{x}_V| \\ \Leftrightarrow \zeta f + \frac{P(f)}{1 - \delta} |x - \hat{x}_V| &\geq c_V \end{aligned}$$

The left side rises continuously and unboundedly in f , and Assumption 1 guarantees the inequality is not satisfied at $f = 0$, so there is a unique $\underline{f}_x > 0$ that renders V just willing to attack. Because the left side falls in x , we know that \underline{f}_x is minimized at $x = 0$, so any f that motivates V must satisfy the condition for f^* stated in the assumption.

2 Proof of Proposition 1

First observe that our assumptions guarantee the existence of a baseline “no-concession” equilibrium in which G sets $x = 0$ in every period, regardless of what S and V do; S provides no support to V , regardless of what G and V do; and V attacks if and only if G sets x and S offers support of at least $\underline{f}_x^{baseline} > 0$ from S . G 's strategy is clearly a best response. If G sets $x = 0$, Assumption 2 guarantees S will not support V , and Assumption 1 guarantees that V will not attack, in response. If G sets $x > 0$, it will not change the behavior of S or V but will lead to a worse policy for G , so there is no profitable deviation for G . Given G 's strategy, Assumption 2 guarantees S will prefer never to support V . Given G 's and S 's strategies, Assumption 1 ensures that V will not attack without support. V 's cutpoint $\underline{f}_x^{baseline}$ is determined by the least f that motivates V to attack:

$$\begin{aligned} -c_V + \zeta f + [1 - P(f)] \left(-|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |0 - \hat{x}_V| \right) &\geq -|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |0 - \hat{x}_V| \\ \Leftrightarrow \zeta f + P(f) \left(|x - \hat{x}_V| + \frac{\delta}{1 - \delta} \hat{x}_V \right) &\geq c_V \end{aligned}$$

The left side rises continuously in f , is negative at $f = 0$ by Assumption 1, and is unbounded in f , so that $\underline{f}_x^{baseline}$ is well-defined and must be positive.

We will use this baseline equilibrium as a reversion point to construct a second, “terror threat” equilibrium, the existence of which is determined by the conditions stated in Proposition 1. We begin by defining the strategy profile in this terror threat equilibrium. In any period in which the players have not yet reverted to the baseline equilibrium, G sets $x = x^*$. If G sets $x \geq x^*$, S provides no support to V ; otherwise, S provides support of f_x^* , to be defined momentarily. If G sets x and S provides support $f \geq \min \left\{ f_x^*, \underline{f}_x^{terror} \right\}$, V attacks; otherwise it does not. The players revert to the baseline equilibrium if and only if, in this or any previous period:

- G set $x < x^*$ and either S provided $f < f_x^*$ or S provided $f \geq f_x^*$ and V did not attack.
- G set $x \geq x^*$ and either S provided $f > 0$ or V attacked.

Informally, in this strategy profile, G makes a policy concession x^* which should lead to S not supporting and V not attacking. If G reneges on this concession, setting $x < x^*$, S will punish G for one period by providing support to V and thereby inducing V to attack. If at any point when G is supposed to be punished, S doesn't provide the specified support, or S does but V doesn't attack, then from there on out, G won't make any policy concession, S won't support V , and V won't attack. For this to be in equilibrium, the support S provides in a punishment period must be enough to motivate V to attack, S must be willing to provide that support in order to get G to return to making the expected policy concession, and the punishment must be severe enough to deter G from deviating to a lower policy.

Once the players revert to the baseline, their behavior is in equilibrium, so consider a period in which reversion has not occurred. Starting from the last choice, V must decide whether to attack or not, given that G set x and S provided support f . There are three cases to consider:

1. $x \geq x^*$, $f = 0$: If V attacks and does not win, the players will revert to the baseline equilibrium; otherwise they will not. V prefers attacking if and only if:

$$\begin{aligned}
-c_V + [1 - P(0)] \left(-|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |0 - \hat{x}_V| \right) &\geq -|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |x^* - \hat{x}_V| \\
\Leftrightarrow P(0) \left(-x + \frac{\hat{x}_V}{1 - \delta} \right) &\geq c_V + \frac{\delta}{1 - \delta} x^*
\end{aligned}$$

Assumption 1 implies this inequality is not satisfied for any x , so V always prefers not to attack.

2. $x < x^*$, $f < f_x^*$ or $x \geq x^*$, $f > 0$: Whether V attacks and does not win or does not attack, the players will revert to the baseline equilibrium in the next period. V prefers

attacking if and only if:

$$\begin{aligned} -c_V + \zeta f + [1 - P(f)] \left(-|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |0 - \hat{x}_V| \right) &\geq -|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |0 - \hat{x}_V| \\ \Leftrightarrow \zeta f + P(f) \frac{\hat{x}_V}{1 - \delta} - P(f)x &\geq c_V \end{aligned}$$

By Assumption 1, this inequality is not satisfied at $f = 0$, but the left side increases continuously and unboundedly in f , so there is some least f for which it is satisfied, which defines \underline{f}_x^{terror} and proves it is positive.

3. $x < x^*$, $f \geq f_x^*$: If V attacks and does not win, the players will not revert to the baseline equilibrium in the next period; if V does not attack they will. V prefers attacking if and only if:

$$\begin{aligned} -c_V + \zeta f + [1 - P(f)] \left(-|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |x^* - \hat{x}_V| \right) &\geq -|x - \hat{x}_V| - \frac{\delta}{1 - \delta} |0 - \hat{x}_V| \\ \Leftrightarrow \zeta f + P(f) \frac{\hat{x}_V}{1 - \delta} + [1 - P(f)] \frac{\delta x^*}{1 - \delta} - P(f)x &\geq c_V \end{aligned}$$

Since $\hat{x}_V > x^*$, x , the left side rises in f . So, if the inequality is satisfied at $f = f_x^*$, then it is also satisfied for any higher f . It follows that V 's specified strategy is a best response if and only if the inequality is satisfied at $f = f_x^*$, which is the first condition stated in the proposition.

Now back up to S 's choice of how much support to provide, given that G set x . Observe that it cannot be a best response for S to provide a positive level of support that does not induce V to attack—this costs f , does not affect V 's behavior or this period's policy, and leads to reversion and a policy permanently less favorable to S . So S 's best response must be either no support, or a level sufficient to induce V to attack. There are two cases to consider:

1. $x \geq x^*$: The level of support necessary to induce V to attack is determined by the inequality in V 's case 2. Notice that any f that satisfied that condition would also satisfy the condition for V in Assumption 2; we will use this fact shortly. If S offers positive support, then so long as V does not attack and win, the players will revert to the baseline equilibrium in the next period. If S offers no support, V will not attack and the players will not revert. S would prefer to offer no support over offering enough to induce attack if and only if:

$$\begin{aligned}
-|x - \hat{x}_S| - \frac{\delta}{1 - \delta} |x^* - \hat{x}_S| &\geq -f - c_S + P(f) \frac{-|\hat{x}_V - \hat{x}_S|}{1 - \delta} + [1 - P(f)] \left(-|x - \hat{x}_S| - \frac{\delta |0 - \hat{x}_S|}{1 - \delta} \right) \\
&\Leftrightarrow f + c_S \geq P(f) \frac{2\hat{x}_S - \hat{x}_V}{1 - \delta} + P(f) [|x - \hat{x}_S| - \hat{x}_S] - \frac{\delta x^*}{1 - \delta}
\end{aligned}$$

If $x \leq \hat{x}_S$, then the term in brackets is negative, and by Assumption 2 the inequality must be satisfied. If instead $x > \hat{x}_S$, then the term in brackets is positive but cannot be greater than $\hat{x}_V - 2\hat{x}_S$. This implies that $\hat{x}_V > 2\hat{x}_S$, which implies that the whole right side of the inequality is negative, so that it is satisfied for any f . Thus, in this case it is always a best response for S to offer no support.

2. $x < x^*$: The level of support necessary to induce V to attack is determined by the inequality in V 's case 3. If S supports at f_x^* or higher and V attacks but does not win, then the players will not revert in the next period. If S supports at any lower level and V does not attack and win, then the players will revert. S prefers offering enough to induce attack over offering no support if and only if:

$$\begin{aligned}
-f - c_S + P(f) \frac{-|\hat{x}_V - \hat{x}_S|}{1 - \delta} + [1 - P(f)] \left(-|x - \hat{x}_S| - \frac{\delta |x^* - \hat{x}_S|}{1 - \delta} \right) &\geq -|x - \hat{x}_S| - \frac{\delta |0 - \hat{x}_S|}{1 - \delta} \\
&\Leftrightarrow P(f) \frac{2\hat{x}_S - \hat{x}_V}{1 - \delta} + [1 - P(f)] \frac{\delta x^*}{1 - \delta} - P(f)x \geq f + c_S
\end{aligned}$$

S 's prescribed strategy is a best response if and only if this inequality is satisfied at

$f = f_x^*$, which is the second condition stated in the proposition, and V is willing to attack in this situation at that level of support, which is the first condition stated in the proposition.

Finally, back up to G 's decision of what policy to set. Clearly it cannot be a best response for G to set $x > x^*$, since this will not change S 's or V 's behavior relative to setting $x = x^*$, will not affect the future policy, and will result in a worse policy for G in this period. So G 's only viable options are to set $x = x^*$ or $x < x^*$. The former will lead to S offering no support and V not attacking; the latter will lead to S supporting at f^* and V attacking. G prefers x^* if and only if:

$$\begin{aligned} \frac{-x^*}{1-\delta} &\geq -c_G + P(f_x^*) \frac{-\hat{x}_V}{1-\delta} + [1 - P(f_x^*)] \left(-x - \frac{\delta x^*}{1-\delta} \right) \\ c_G + P(f_x^*) \frac{\hat{x}_V - x^*}{1-\delta} &\geq [1 - P(f_x^*)] (x^* - x) \end{aligned}$$

This is the third condition stated in the proposition, and completes the proof.

Part II

Model Extensions

3 Uncertainty about Interests

As we note in the paper, because our model assumes all the players' interests and actions are common knowledge, the equilibrium path features policy concessions from the government, but neither support to the terrorists nor terrorist attacks. The government perfectly anticipates the supporters' and terrorists' reactions to its choice of policy, and chooses one that will not induce support or attacks. We argued in the paper that this setup could still

be used to study the question of whether supporters could use the threat of supporting terrorism to extract a concession from the government, because we could simply analyze when this threat was credible off the equilibrium path and see how that off-path threat coerced the government into making a policy concession on the path.

Still, one might reasonably be concerned about applying predictions from a model to empirical cases, when those predictions concern behavior that does not occur on the equilibrium path in the model. We therefore analyzed an incomplete information version of our model in order to show that this would lead to support for terrorists and terrorist attacks occurring with positive probability on the equilibrium path, and to confirm that our qualitative conclusions about how this threat works would survive this more realistic setup.

Because of length restrictions, we will only describe the changes to the model setup and the equilibrium we analyzed, and then explain the intuitions for why support and attacks can happen on the path in that equilibrium and why our conclusions still come through. The formal definition of the equilibrium, propositions characterizing its existence, and proofs of those propositions are available from the authors on request.

The extended model introduces uncertainty about S 's resolve: that is, how much S cares about the issue at stake, relative to the costs of supporting and suffering terrorist attacks. In particular, we assume that the cost of supporting V at the level f costs θf , rather than just f as in the model in the main paper. The larger θ is, the more it costs S to provide whatever level of support is necessary to motivate V to attack, and the more heavily this cost weighs in S 's utility against the policy change that this support and attack might extract from G policy x . (We believe that instead treating S 's cost of suffering the terrorist campaign c_S as the uncertain parameter would have very similar effects, precisely because it also simply amplifies the total cost of supporting and suffering a terrorist campaign relative to whatever policy change might occur.)

At the beginning of the game, Nature chooses $\theta \in \{\underline{\theta}, \bar{\theta}\}$, with $0 < \underline{\theta} < \bar{\theta}$, so that $\underline{\theta}$ is

the “tough” type of S , since she bears a smaller marginal cost of providing support, and $\bar{\theta}$ is the “weak” type. The probability of $\underline{\theta}$ is α and of $\bar{\theta}$ is $1 - \alpha$. This choice is observed by S and V , but not by G . The game then proceeds exactly as described in the main paper.

We still impose Assumption 1 on V and Assumption 2 on both types of S , were they playing a complete information game where G knew their type. We distinguish the two types by assuming that $\underline{\theta}$ is such that this type of S meets the conditions of Proposition 1 of the main paper. Thus, if S is the tough and G is certain of this, then the terror threat equilibrium exists. By contrast, $\bar{\theta}$ is such that this type of S cannot meet the conditions of Proposition 1, so that if S is the weak type and G knows this, then only the baseline, no-concession equilibrium exists. This ensures a game in which, at least initially, G does not know whether S 's threat to support V and induce it to attack if G does not make an expected policy concession is credible.

We analyze the conditions for the following strategy profile to be in Perfect Bayesian Equilibrium. If, in some period, G has learned with certainty that he faces the tough type of S , then the players employ the strategy profile of the terror threat equilibrium, with G setting policy $x^* > 0$, S not supporting V , and V not attacking on the equilibrium path. Off-path, no matter what happens, G continues to be certain he faces the tough type. By construction, this is a Perfect Bayesian Equilibrium for such periods. If, in some period, G has learned with certainty that he faces the weak type of S , then the players follow the strategy profile of the baseline, no-concession equilibrium, with G setting policy 0, S not supporting, and V not attacking on the path. Off-path, no matter what happens, G continues to be certain he faces the weak type. This too is clearly in equilibrium. In any period where G does not know with certainty which type he faces, he chooses between two options: the risky policy of 0 or the safe policy of x^* , depending upon his belief about S 's type. If he is confident enough that S is tough, he sets x^* ; otherwise he chooses 0. If G chooses x^* , both types of S provide no support to V , V does not attack, and G does not update his belief about S . If

G chooses 0, the tough type of S responds as in the terror threat equilibrium by providing a level of support f_0^* that induces V to attack, while the weak type does not support V and V does not attack. G thereby infers with certainty which type he faces. If G observes an off-path level of support (i.e., $f \in (0, f_x^*)$), he believes with certainty that S is weak.

The details arise from characterizing the cutpoint for G 's belief that determines whether he selects the risky or the safe policy, from ensuring that the weak type of S does not wish to imitate the strong type when G offers the risky policy, and from characterizing optimal off-path behavior when an uncertain G sets a policy other than 0 or x^* .

Consider what happens when this profile is in equilibrium. If G starts the game skeptical enough that S is tough, he screens S by setting $x = 0$. If S is the tough type, both support for terrorism and terrorist attack will happen on the equilibrium path, G will infer that S 's threat to support terrorism is credible, and from then on G will make a policy concession in order to avoid this threat, just like in the model in the main paper. If S is the weak type, no support or attack will occur, G will infer that S 's threat to support terrorism is not credible, and no concession will be forthcoming, also like the full-information model.

If G starts the game worried enough that S is tough, he will make a policy concession from the start, and no support or attack will occur on the equilibrium path. But it will still be true that it is S 's (possibly bluffed) threat to support terrorism and induce attacks that coerces G to make the safe policy choice.

Observations 1, 2, and 3—that this threat is easier for S to wield than the weapon of the strong; that the threat may still be credible if, or might require that, V be extreme; and that when used this threat resembles the stylized facts of a terrorism campaign—still hold under this incomplete-information setup. The authors have worked out a numerical example illustrating Observation 2 that is analogous to the ones in the main paper.

4 Uncertainty about Actions

Our model in the main paper also assumes that the level of support f that S provides to V is common knowledge. It seems reasonable that S would know how much support she is providing V , and V would know how much it is receiving. However, given the government's counter-terrorism efforts, supporters of terrorism would surely need to conceal their identities and the precise nature, channel, and degree of the support they contribute. Thus it seems likely that in the real world, the government would be quite uncertain about f . Here, we explain why we think this is not so worrisome, because a much more plausible assumption about what G can observe ends up being good enough for our results.

To understand the potential problem, consider the role G 's observation of f plays in the terror threat equilibrium. First, f enters into G 's continuation values, but only through its effect on $P(f)$. Second, f enters into G 's strategy, because his actions are conditioned on whether S provided at least the specified level of support in the punishment phase and whether S provided any support outside of the punishment phase. Thus it might be concerning that, if G can't observe f , he may not be able to tell which action he is supposed to play or what the consequences will be for his continuation value.

However, it turns out that it suffices for G to be able to observe only whether an attack occurred and, if one occurred, how serious of a danger it posed to G remaining in power—that is, $P(f)$. Empirically, a government should obviously be able to observe whether it is being attacked in a terrorist campaign, and it seems very plausible that a government would be able to assess the degree to which the campaign, even if it had not yet been victorious, placed the government in jeopardy. In our model, this is akin to the government being able to observe the value of $P(f)$, even if it does not necessarily observe f , and even if it can only form the roughest estimate of f from its observation of $P(f)$.

Reconsidering the role of f in the terror threat equilibrium, if G observes $P(f)$, then

he will know his own continuation values. In equilibrium (on or off the path), S only ever provides positive support to V if it will induce V to attack, enabling the government to observe $P(f)$. Because $P(f)$ is increasing in f , there is a one-to-one correspondence between f and $P(f)$. This implies that, at any point in the game tree where G 's strategy is conditioned on f and an attack occurring, we could write an equivalent strategy that is instead conditioned on the observed value of $P(f)$. The only points where G 's strategy is conditioned on f but an attack does not occur, are those where S offered less than the specified support to V so that it did not attack. But then G can observe the absence of attack, infer that S must have deviated from equilibrium, and react accordingly.

Summarizing, our analysis does not actually depend on the government observing the level of support provided to a terrorist organization. It is enough if a government can tell when it is being attacked and how dangerous an ongoing terrorist campaign is.

5 Decisive Government Victory

In the model in the main paper, the terrorist organization can (possibly with vanishing probability) win a decisive victory, usurping the government's power to set policy, but the government cannot vanquish the terrorist organization so long as supporters continue contributing to it. We explain in the paper why we think this is empirically plausible in many cases, but it is also plausible that in some cases a government might be able to inflict such a crushing defeat on the terrorists themselves, or successfully identify and arrest enough of their supporters or interdict enough of their channels for receiving support, to shut the terrorist organization down permanently. Here, we summarize our analysis of an extension of the model in which we incorporate this possibility. In short, we find that the qualitative conclusions still hold, though the conditions for the existence of the terror threat equilibrium become more restrictive.

We modify the model setup so that, if V attacks in some period, then three outcomes are possible. As before, with probability $P(f)$, where f is the level of support V received from S in that round, V wins and sets policy from then on. But now, with probability R , the government wins and sets policy from now on, and S and V have no further actions. With probability $1 - P(f) - R$, neither side wins, the government's choice of policy from the beginning of the period is implemented, and the game repeats.

It turns out that this changes very little of the analysis. In any subgame, V is more hesitant to attack, because although it still offers the same chance of victory for V , it now also risks defeat and G 's ideal policy permanently implemented. This in turn raises the level of support S must provide to motivate V , rendering it more costly for S to induce an attack. Holding the required level of support constant, S is also more reluctant to induce an attack, because now she risks not only V 's victory (whose ideal policy is too high for her) but also G 's victory (whose ideal policy is too low for her). Finally, G becomes more difficult for the S 's threat of supporting V to deter, because attacks now offer the possibility for G of permanently ending any policy concession to S .

Thus, all three conditions in Proposition 1 are modified to become more stringent, because the costs of punishing G for S and V now include risking G 's victory, and the benefit of deviating from the policy concession for G now includes the possibility of his victory. As the probability of G 's victory R rises, these conditions become increasingly stringent, and for high enough values, they cannot be satisfied. When this happens, S has no ability to coerce a concession from G , and the only equilibrium is the baseline, no-concession one.