

Discussion Papers
Department of Economics
University of Copenhagen

No. 07-14

Terrorism, Anti-Terrorism,
and the Copycat Effect

Thomas Jensen

Stu­di­estræde 6, DK-1455 Copenhagen K., Denmark
Tel. +45 35 32 30 82 - Fax +45 35 32 30 00
<http://www.econ.ku.dk>

ISSN: 0902-6452 (print) ISSN: 1601-2461 (online)

Terrorism, Anti-Terrorism, and the Copycat Effect

Thomas Jensen*
Department of Economics
University of Copenhagen

July 15, 2007

Abstract

In this paper we formulate and analyze a simple dynamic model of the interaction between terrorists and authorities. Our primary aim is to analyze how the introduction of a so called copycat effect influences behavior and outcomes. We first show that our simple model of terrorist cells implies that an increase in anti-terrorism makes it more likely that cells will plan small rather than large attacks. Furthermore, we see that an increase in anti-terrorism can make a terrorist attack more likely. Analyzing the problem of optimal anti-terrorism we see that the introduction of a copycat effect rationalizes an increase in the level of anti-terrorism after a large attack. Using this result we show how the copycat effect changes the dynamic pattern of terrorism attacks and what the long run consequences are.

Keywords: Terrorist Cells, Optimal Anti-Terrorism, Copycat Effect, Dynamic Pattern of Terrorism.

JEL Classification: D74, H56.

*Department of Economics, University of Copenhagen, Studiestraede 6, DK-1455 Copenhagen K, Denmark. Email: Thomas.Jensen@econ.ku.dk.

1 Introduction

In this paper we formulate and analyze a simple dynamic model of the interaction between terrorists and authorities. The primary aim of the paper is to analyze how the introduction of a so called copycat effect influences behavior and outcomes. We say that a copycat effect exists if terrorist cells are more likely to be formed after a period with a high level of terrorism (for example a large attack) than after a period with low terrorist activity. There are several good reasons to expect that a copycat effect exists. Media attention to terrorism is higher when there has been a lot of terrorist activity in the recent past. Therefore the possibility of becoming a terrorist is more salient for potential terrorists. Furthermore, the increased media attention means that even relatively minor terrorist acts get a lot of publicity. Therefore it becomes more attractive to form a terrorist cell and thus it is likely that more cells are formed.

In our model a terrorist cell lives for one period only and its sole decision is whether to plan a small or a large attack. Planning a large attack is more risky because it requires more planning and therefore involves a higher risk of being rolled up by the authorities. The difference in risk between the two types of attacks is increasing in the authorities' spending on anti-terrorism. Therefore it follows that if the authorities increase the level of anti-terrorism then a cell is more likely to plan a small attack. The effect of an increase in anti-terrorism on the probability that a cell will be successful in making an attack (small or large) is ambiguous. It can be the case that increased spending on anti-terrorism makes a terrorist attack more likely.

In each period of time a terrorist cell is formed with some probability. The aim of the authorities is to minimize the sum of (discounted) expected damage from terrorism and anti-terrorism costs over all periods by choosing the level of anti-terrorism in each period. The horizon is infinite. We solve for optimal anti-terrorism in two cases. First we consider a benchmark case where the probability of a cell being formed is the same in all periods. Then we move on to a case where a copycat effect is in play. More specifically we assume that the probability of a cell being formed is higher if there was a large attack in the previous period. We show that the authorities choose a higher level of anti-terrorism after a large attack. Using that result we see that if a cell is formed then the probability of a small attack is highest and the probability of a large attack and the expected damage is lowest after a large attack. Finally, we compare long run distributions for the benchmark case and the copycat case. In the long run the copycat effect implies more anti-terrorism, more small attacks and a higher per period sum of terrorism damage and anti-terrorism costs. On the other hand it implies less large attacks and less damage from terrorism.

A substantial number of papers have studied economic and game theoretic

models of the interaction between terrorists and authorities. For a review see for example Sandler and Enders (2004). Among the specific problems that have been studied are terrorists choice of targets (see e.g. Sandler and Lapan (1988)), hostage taking (Lapan and Sandler (1988)), substitutions by terrorists after policy changes (Enders and Sandler (1993)), the choice between proactive and defensive counterterrorism measures (Rosendorff and Sandler (2004)), and the effect of concessions to terrorists (Bueno de Mesquita (2005)). We are not aware of papers studying the implications of copycat effects. Furthermore, our model is distinct from most of the literature because it is dynamic (with infinite horizon), although it should be noted that the dynamic structure is very simple because terrorist cells live for one period only. Another dynamic model is Faria (2003).

The paper is organized as follows. In Section 2 we set up the model. Then we consider the behavior of the terrorist cells in Section 3 and the problem of optimal anti-terrorism in the two cases in Section 4. Finally, in Section 5, we discuss our results and some ideas for further research.

2 The Model

In each period of time $t = 0, 1, 2, \dots$ a terrorist cell is formed with probability $\lambda_t \in (0, 1]$. A cell lives for one period only and its only decision is whether to plan a small or a large attack (it can only make one attack). If the cell formed in period t succeeds in making a small attack then the damage is $D > 0$. If the cell succeeds in making a large attack then the damage is

$$D(1 + \varepsilon_t),$$

where ε_t is drawn from a probability distribution on $[0, \infty)$ with cumulative distribution function F . The realization of ε_t is known to the cell when it makes its decision. We assume that $F(0) = 0$ and that $F(\varepsilon) < 1$ for all ε . Furthermore we assume that F is differentiable on $[0, \infty)$ such that it has the density function $f = F'$.

The authorities choose a level $a \in [0, \infty)$ of anti-terrorism in each period. The level of anti-terrorism in some period t , a_t , decides how likely it is that a cell formed in period t is rolled up before it attacks *if the cell plans a large attack*. If the cell plans a small attack then the probability that it is rolled up is zero. While this is hardly realistic it is a simple way of modelling that a cell preparing a small attack is less likely to be rolled up because a small attack requires less planning. Formally, if the cell decides to plan a large attack then the probability that is rolled up before the attack is $p(a_t)$, where $p : [0, \infty) \rightarrow [0, 1]$ is a differentiable and strictly increasing function satisfying $p(0) = 0$. The level of anti-terrorism is

known to the cell when it makes its decision. We assume that the cell maximizes expected damage. Therefore the cell plans a small attack if

$$D > (1 - p(a_t))D(1 + \varepsilon_t)$$

and a large attack if we have the opposite inequality. If the cell is indifferent then we assume that it plans a small attack.

The aim of the authorities is to minimize the sum of discounted expected damages and anti-terrorism costs by choosing the level of anti-terrorism in each period. The discounting rate of the authorities is $\delta \in (0, 1)$. The cost of anti-terrorism is given by a differentiable and strictly increasing function $c : [0, \infty) \rightarrow [0, \infty)$ with $c(0) = 0$.

The timing of events and decisions in period t is as described in the list below. It is important to note that when the authorities decide on the level of anti-terrorism they know λ_t and F but they do not know whether a cell will be formed and what the realized value of ε_t will be in that case. r_t denotes the damage from terrorism in period t .

Timing of events and decisions in period t :

1. The authorities decide on a_t and pays the cost $c(a_t)$;
2. A new cell is formed with probability λ_t ;
3. If a cell was formed then the value of ε_t is realized and the cell decides on what kind of attack to plan;
4. If a cell was formed and planned a small attack then it launches its attack ($\Rightarrow r_t = D$);
5. If a cell was formed and planned a large attack then it is rolled up with probability $p(a_t)$ ($\Rightarrow r_t = 0$);
6. If a cell was formed, planned a large attack and was not rolled up then it launches its attack ($\Rightarrow r_t = D(1 + \varepsilon_t)$).

In the following we first take a closer look at the behavior of the terrorist cells and its consequences. Then we move on to the problem of optimal anti-terrorism. We will focus on two cases. First, we consider a benchmark case where in each period the probability of a new cell being formed does not depend on actions or events in previous periods. Secondly, we consider a case where a copycat effect is in play, i.e. it is more likely that a cell is formed in the current period if there was a large terrorist attack in the previous period than if there was not. Finally, we compare the copycat case to the benchmark case.

3 The Behavior of the Terrorist Cells

Suppose that in some period t the authorities choose the level of anti-terrorism a . Furthermore suppose that a cell is formed. As noted above the cell plans a small attack if (and only if)

$$D \geq (1 - p(a))D(1 + \varepsilon_t).$$

Thus the probability that the cell launches a small attack is

$$\Pr(D \geq (1 - p(a))D(1 + \varepsilon_t)) = \Pr(\varepsilon_t \leq \varepsilon^*) = F(\varepsilon^*),$$

where

$$\varepsilon^* = \frac{1}{(1 - p(a))} - 1 = \frac{p(a)}{(1 - p(a))}.$$

Note that

$$\frac{\partial \varepsilon^*}{\partial a} = \frac{p'(a)}{(1 - p(a))^2} > 0$$

which implies

$$\frac{\partial F(\varepsilon^*)}{\partial a} = \frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*) > 0.$$

Thus we see that an increase in the level of anti-terrorism makes it more likely that the cell will make a small attack. The probability of the cell successfully launching a large attack is

$$(1 - p(a))(1 - F(\varepsilon^*)).$$

Since $p' > 0$ and $\frac{\partial F(\varepsilon^*)}{\partial a} > 0$ this expression is decreasing in a , so an increase in a makes a large attack less likely. Adding the two probabilities above we get the probability that the cell launches some kind of attack (i.e. the probability that it is not rolled up):

$$P(a) = F(\varepsilon^*) + (1 - p(a))(1 - F(\varepsilon^*)) = 1 - p(a)(1 - F(\varepsilon^*)).$$

We see that

$$\frac{\partial P}{\partial a} = -p'(a)(1 - F(\varepsilon^*)) + p(a)\frac{\partial \varepsilon^*}{\partial a}f(\varepsilon^*)$$

The first term arises from a 's effect on the probability of the cell being rolled up. This term is obviously negative. The second term arises from a 's effect on the cell's decision about what kind of attack to plan. An increase in a makes it more likely that the cell will plan a small attack which decreases the probability that it is rolled up. Thus this term is positive. Generally we cannot say which of the two effects dominates, it depends on the functions p and F and the value of a . Below we show by an example that an increase in the level of anti-terrorism can make a terrorist attack more likely. That is an interesting observation.

Consider the following simple example:

$$\begin{aligned} p(a) &= \frac{a}{1+a}, \\ F(\varepsilon) &= 1 - \exp(-\varepsilon). \end{aligned}$$

Then we have

$$\varepsilon^* = \frac{a}{(1+a)(1 - \frac{a}{1+a})} = a$$

and

$$f(\varepsilon) = \exp(-\varepsilon).$$

Thus the probability of the cell making an attack is

$$P(a) = 1 - \frac{a}{1+a} \exp(-a)$$

and hence we have

$$\frac{\partial P}{\partial a} = \frac{\exp(-a)}{(1+a)^2} (a(1+a) - 1).$$

Loosely speaking, for small levels of a an increase in the level of anti-terrorism makes a terrorist attack less likely and for large levels of a we have the opposite effect. More precisely, $\frac{\partial P}{\partial a}$ is negative for a 's below the positive root of $a(1+a) - 1$ and positive for a 's above this root.

4 Optimal Anti-Terrorism

In this section we consider the problem of the authorities. First, we consider our benchmark case where, for each t , λ_t (the probability of a cell being born in period t) does not depend on what has happened in earlier periods. In that case the authorities' problem is just a sequence of independent static problems which are easy to solve. Secondly, we introduce a simple type of copycat effect which makes the authorities problem truly dynamic. We solve this problem by dynamic programming. Finally, we compare the two cases.

4.1 The Benchmark Case

In this case the level of anti-terrorism chosen in some period t does not influence the problem of the authorities in future periods. Thus in each period the authorities simply choose the level of anti-terrorism that minimizes the sum of the expected damage from the cell possibly formed in that period and the cost of anti-terrorism.

Consider the authorities problem in period t . For simplicity we suppress subscript t 's such that we write λ , a and ε instead of λ_t , a_t and ε_t . The expected damage from a cell formed in period t is

$$\Delta(a) = F(\varepsilon^*)D + (1 - F(\varepsilon^*))(1 - p(a))D(1 + E[\varepsilon|\varepsilon > \varepsilon^*]).$$

Note that

$$E[\varepsilon|\varepsilon > \varepsilon^*] = \frac{\int_{\varepsilon^*}^{\infty} \varepsilon f(\varepsilon) d\varepsilon}{1 - F(\varepsilon^*)}.$$

We can write the problem of the authorities as

$$\min_{a \in [0, \infty)} \pi(a),$$

where

$$\pi(a) = \lambda\Delta(a) + c(a).$$

Since p , c and F are differentiable so is π . The first order condition for an interior solution is

$$\pi'(a) = 0,$$

which can also be written

$$c'(a) = -\lambda\Delta'(a).$$

This is a simple "marginal cost equals marginal benefit" equation. The left hand side is of course the marginal cost of anti-terrorism. The right hand side is minus the marginal effect of anti-terrorism on the expected damage from terrorism. By differentiating Δ and collecting terms (see the Appendix for details) we see that the condition can be written as

$$c'(a) = \lambda p'(a)(1 - F(\varepsilon^*))D(1 + E[\varepsilon|\varepsilon > \varepsilon^*]).$$

Under some additional assumptions on the functions p and c the solution to the authorities problem is unique, interior and the only solution to the first order condition.

Theorem 4.1 *Suppose p and c are twice differentiable and that we have the following conditions:*

1. $p'(0) > 0$, $c'(0) = 0$ and $\lim_{a \rightarrow \infty} c'(a) = \infty$;
2. $p'' \leq 0$ and $c'' > 0$.

Then there is a unique solution to the authorities problem and it is interior and the only solution to the first order condition.

Proof. See the Appendix.

If the conditions in the theorem are satisfied then we let \bar{a} denote the unique solution to the authorities problem. By using the Implicit Function Theorem on the first order condition it is easily seen that \bar{a} is a differentiable function of λ and that

$$\frac{\partial \bar{a}}{\partial \lambda} > 0$$

(simply note that $\frac{\partial \pi'}{\partial \lambda} < 0$ and, by the proof of Theorem 4.1, $\pi''(a) > 0$). So if the probability of a cell being formed is increased then the authorities will choose a higher level of anti-terrorism. This immediately implies that the probability of a small attack,

$$\lambda F(\varepsilon^*(\bar{a})),$$

is increasing in λ (remember that ε^* increases with the level of anti-terrorism). We cannot generally say in which direction the probability of a large attack,

$$\lambda(1 - F(\varepsilon^*(\bar{a})))(1 - p(\bar{a})),$$

changes when λ increases. The same is true for the sum of the two probabilities and for the expected damage, $\lambda \Delta(\bar{a})$. Note, however, that conditional on a cell being formed we have that the probability of a large attack and the expected damage from terrorism are decreasing in λ . Conditioning on a cell being formed does not change the result that the probability of a small attack is increasing in λ or the fact that we cannot determine in which direction the probability of some kind of attack changes when λ increases.

As an example of functions satisfying the conditions in Theorem 4.1 we can take the example of p from the previous section and $c(a) = a^2$. If we again let $F(\varepsilon) = 1 - \exp(-\varepsilon)$ then the first order condition for optimal anti-terrorism becomes

$$2a - \lambda \frac{\exp(-a)}{(1+a)^2} (2+a)D = 0.$$

4.2 Introducing a Copycat Effect

Now we introduce a simple type of copycat effect. More specifically we assume that λ_t is higher if there was a large terrorist attack in period $t-1$ than if there was not. To model the copycat effect define the variable x at time t as

$$x_t = \left\{ \begin{array}{l} s \text{ if } r_{t-1} \leq D \\ l \text{ if } r_{t-1} > D \end{array} \right\}.$$

So $x_t = s$ (for small) if the damage from terrorism was at most D in period $t - 1$. If the damage was higher than D in period $t - 1$ then $x_t = l$ (for large). We then let λ_t depend on x_t and assume that

$$\lambda_t(s) < \lambda_t(l),$$

which reflects the copycat effect. We furthermore assume that $\lambda_t(s)$ and $\lambda_t(l)$ does not depend on t . Thus we can write

$$\lambda_t(s) = \lambda^s < \lambda^l = \lambda_t(l) \text{ for all } t.$$

Having modelled the copycat effect it is obvious that the level of anti-terrorism chosen in period t influences the probability that a cell is born in period $t + 1$ because it influences r_t and thus x_{t+1} . Therefore the authorities must solve a truly dynamic problem in order to find their optimal level of anti-terrorism in each period. We solve the problem by dynamic programming.

The Bellman equation for the dynamic programming problem can be written as

$$V(x) = \inf_{a \in [0, \infty)} [\pi(a, x) + \delta(P(s, x, a)V(s) + P(l, x, a)V(l))],$$

where

$$\pi(a, x) = \lambda^x \Delta(a) + c(a)$$

and

$$P(x', x, a) = \Pr(x_{t+1} = x' | x_t = x, a_t = a) \text{ for all } x, x' \in \{s, l\}.$$

Note that the Bellman equation is a little non-standard because the transition probabilities depend on a . Writing the transition probabilities in detail we get

$$P(l, x, a) = \lambda^x(1 - p(a))(1 - F(\varepsilon^*))$$

and

$$\begin{aligned} P(s, x, a) &= 1 - P(l, x, a) \\ &= 1 - \lambda^x(1 - p(a))(1 - F(\varepsilon^*)) \end{aligned}$$

for all $x \in \{s, l\}$, $a \in [0, \infty)$. By plugging in the transition probabilities the Bellman equation becomes

$$V(x) = \inf_{a \in [0, \infty)} [\pi(a, x) + \delta \lambda^x(1 - p(a))(1 - F(\varepsilon^*))(V(l) - V(s)) + \delta V(s)].$$

Lemma 4.2 *There exists a unique solution \bar{V} to the Bellman equation above. It satisfies $\bar{V}(s) < \bar{V}(l)$.*

Proof. See the Appendix.

Now, for each $x \in \{s, l\}$, consider the problem

$$\min_{a \in [0, \infty)} [\pi(a, x) + \delta(P(s, x, a)\bar{V}(s) + P(l, x, a)\bar{V}(l))].$$

Pairs of solutions to these two minimization problems are solutions to the dynamic programming problem of the authorities. In the theorem below we present an existence and uniqueness result. For simplicity we define

$$g(a, x) = \pi(a, x) + \delta(P(s, x, a)\bar{V}(s) + P(l, x, a)\bar{V}(l)).$$

Theorem 4.3 *Suppose the assumptions from Theorem 4.1 are satisfied, that F is twice differentiable, and that*

$$\frac{\partial^2}{\partial a^2} F(\varepsilon^*) = \frac{\partial}{\partial a} \left(\frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*) \right) \leq 0.$$

Then, for each $x \in \{s, l\}$, there is a unique solution to the problem considered above and it is interior and the only solution to the first order condition

$$\frac{\partial}{\partial a} g(a, x) = 0$$

Furthermore, letting $\bar{a}_+(x)$ denote the optimal level of terrorism in state x we have

$$\bar{a}_+(s) < \bar{a}_+(l).$$

Proof. See the Appendix.

The result that authorities choose a higher level of anti-terrorism when $x = l$ than when $x = s$ is perhaps not surprising but it does have some interesting implications. Consider the probability of a small attack as a function of x . This probability is

$$\lambda^x F(\varepsilon^*(\bar{a}_+(x)))$$

and thus it is highest when $x = l$. So if there was a large attack in the previous period then there is a higher probability of a small attack than if there was not. This is also true if we instead consider the probability of a small attack conditional on a cell being born, which is of course equal to $F(\varepsilon^*(\bar{a}_+(x)))$. The probability of a large attack as a function of x is

$$\lambda^x (1 - p(\bar{a}_+(x)))(1 - F(\varepsilon^*(\bar{a}_+(x)))).$$

We cannot generally say whether this function is highest when $x = s$ or when $x = l$. However, note that the probability of a large attack conditional on a cell being formed is highest when there was not a large attack in the previous period. With respect to the expected damage from terrorism, $\lambda^x \Delta(\bar{a}_+(x))$, we again cannot say whether it is highest when $x = s$ or when $x = l$. But conditional on a cell being born it is highest when there was not a large attack in the previous period. Finally, note that the per period sum of damage from terrorism and anti-terrorism costs,

$$\pi(a, x) = \lambda^x \Delta(\bar{a}_+(x)) + c(\bar{a}_+(x)),$$

is highest when $x = l$. This follows from the observation that $\pi(a, s)$ is increasing for $a \geq \bar{a}_+(s)$, which follows easily from Theorem 4.4 in the following section.

Now we will consider the problem of finding the long run distribution of x when the authorities behave optimally. Define

$$Q(x', x) = P(x', x, \bar{a}_+(x)) \text{ for all } x', x \in \{s, l\}$$

and note that

$$Q(x', x) \in (0, 1) \text{ for all } x', x \in \{s, l\}.$$

These transition probabilities defines a map Q^* on the set of probability distributions on $\{s, l\}$ into itself by

$$(Q^*\beta)(x') = Q(x', s)\beta(s) + Q(x', l)\beta(l), \quad x' \in \{s, l\}.$$

Define $\bar{\beta}$ by

$$\bar{\beta}(s) = \frac{Q(s, l)}{1 - Q(s, s) + Q(s, l)}.$$

It is easily seen that $\bar{\beta}$ is the unique fixed point for Q^* and that $(Q^*)^t \beta \rightarrow \bar{\beta}$ for any β . Hence $\bar{\beta}$ is the unique stationary distribution of x and for any distribution of x_0 the distribution of x_t converges to $\bar{\beta}$. Therefore we conclude that the long run distribution of x is given by $\bar{\beta}$. Of course $\bar{\beta}$ then determines the long run distribution of the level of anti-terrorism and all functions thereof.

Finally we return briefly to the example of p , F and c considered earlier. Since F is twice differentiable and

$$\frac{\partial^2}{\partial a^2} F(\varepsilon^*) = \frac{\partial^2}{\partial a^2} (1 - \exp(-a)) = -\exp(-a) < 0$$

we have that the conditions in Theorem 4.3 are satisfied. For each x the first order condition for optimal anti-terrorism is

$$2a - \lambda^x \frac{\exp(-a)}{(1+a)^2} (2+a)D - \delta \lambda^x \left[\frac{\exp(-a)}{(1+a)^2} (2+a) \right] (\bar{V}(l) - \bar{V}(s)) = 0,$$

which can be simplified to

$$2a - \lambda^x \frac{\exp(-a)}{(1+a)^2} (2+a)(D + \delta(\bar{V}(l) - \bar{V}(s))) = 0.$$

By comparing with the first order condition from the benchmark case we see that in each state the authorities behave as if they were in a case where there is no copycat effect and D is replaced by $D + \delta(\bar{V}(l) - \bar{V}(s))$.

4.3 Comparing the Two Cases

Consider the authorities' problem with and without the copycat effect in some period t . Suppose that $\lambda = \lambda^{x^t}$, i.e. that the probability of a new cell being formed is the same whether or not there is a copycat effect. Then the following result shows that the authorities will choose a strictly higher level of anti-terrorism if the copycat effect is present. Note that we still assume that the assumptions in Theorem 4.3 (which include the assumptions in Theorem 4.1) are satisfied.

Theorem 4.4 *If $\lambda = \lambda^x$ then*

$$\bar{a} < \bar{a}_+(x).$$

Proof. By the first order conditions for the two cases we have that

$$\lambda \Delta'(\bar{a}) + c'(\bar{a}) = 0$$

and

$$\lambda^x \Delta'(\bar{a}_+(x)) + c'(\bar{a}_+(x)) > 0.$$

Since $\lambda = \lambda^x$ and $\lambda \Delta(a) + c(a)$ is convex it follows that $\bar{a} < \bar{a}_+(x)$. \square

The intuition behind this result is the following. With the copycat effect the authorities do not only consider the sum of expected damage and anti-terrorism costs in the present period, they also take into account that raising the anti-terrorism level makes it less likely that a cell will be formed in the following period. Thus the marginal benefit from anti-terrorism is higher with the copycat effect and therefore a higher level is chosen.

From the result it follows easily that if we are in a period with $\lambda = \lambda^{x^t}$ then the probability of a small attack is higher with the copycat effect than without it. On the other hand the probability of a large attack and the expected damage from terrorism is lower with the copycat effect (note, however, that the sum of damages and costs is higher with the copycat effect).

Ultimately we want to compare long run distributions for the two cases. The problem with this is how to choose the parameters λ , λ^s and λ^l in order to get

meaningful comparisons. We do it the following way. Fix λ^s and λ^l and suppose that in each period the probability of a cell being formed in the benchmark case is

$$\lambda = \left\{ \begin{array}{l} \lambda^s \text{ with probability } \bar{\beta}(s) \\ \lambda^l \text{ with probability } \bar{\beta}(l) \end{array} \right\},$$

where $\bar{\beta}$ is the long run distribution of x from the copycat case. The realization of λ is known to the authorities when they choose the level of anti-terrorism. We define λ this way to ensure that the long run distributions of the probability of a cell being formed are the same in the benchmark and the copycat case. Thus any difference between the two cases does not arise because of differences in these long run distributions.

We say that a variable (e.g. the level of anti-terrorism or the expected damage from terrorism) is higher in the long run with (without) the copycat effect if the long run distribution of the variable with (without) the effect strictly first order stochastically dominates the long run distribution without (with) the effect. Note that this implies that the long run average of the variable is higher with (without) the effect. With this definition we have the following results.

Theorem 4.5 *Assuming λ is distributed as described above the following statements hold.*

1. *The level of anti-terrorism is higher in the long run with the copycat effect.*
2. *The probability of a small attack ($r = D$) is higher in the long run with the copycat effect.*
3. *The probability of a large attack ($r > D$) is higher in the long run without the copycat effect*
4. *The expected damage from terrorism is higher in the long run without the copycat effect.*
5. *The sum of expected damage and anti-terrorism costs is higher in the long run with the copycat effect.*

Proof. When λ is distributed as described above then any variable v depending on the level of anti-terrorism is higher in the long run with (without) the copycat effect if and only if

$$\begin{array}{l} v(\bar{a}_+(x)) \geq v(\bar{a}(\lambda^x)) \text{ for each } x \in \{s, l\} \\ (\leq) \end{array}$$

with strict inequality for at least one x . Using that observation all conclusions follow easily from Theorem 4.4. \square

It is worth noting that we cannot generally say whether the long run probability of some kind of attack ($r > 0$) is highest with or without the copycat effect.

5 Discussion

Using our simple model of terrorist cells we saw that an increase in the level of anti-terrorism makes it more likely that a cell will make a small attack and less likely that it will make a large attack. The probability that a cell makes some kind of attack (which is equal to the probability that it is not rolled up by the authorities) can change in either direction. This is an interesting observation - spending more on anti-terrorism may increase the probability of a terrorist attack. Suppose that there has just been a terrorist attack and that the authorities increase the level of anti-terrorism only to try to calm down the public. This effort can have the effect that another terrorist attack becomes more likely! Note, however, that an increase in the level of anti-terrorism always decreases the expected damage made by a terrorist cell.

By analyzing the problem of optimal anti-terrorism we saw that the existence of a copycat effect offers a rational choice explanation of why authorities increase the level of anti-terrorism after a large attack. Therefore, when a copycat effect exists a terrorist cell formed after a large attack is more likely to make a small attack and less likely to make a large attack. This implies that after a large attack there is a larger probability of a small attack. But, because of the increased likelihood of a cell being formed, it does not necessarily imply that there is a smaller probability of a large attack. By the same argument we have that while the expected damage made by a terrorist cell is smaller after a large attack, the a priori expected damage from terrorism may be higher.

In our comparison of the copycat case and the benchmark case we saw that the long run distribution of several variables differs systematically in the two cases. With the copycat effect there is more anti-terrorism, more small attacks and a higher sum of damages and costs while there is less large attacks and less damage. Note that the benchmark case is the better one for the authorities because the sum of damages and costs are lower.

The way we define the copycat effect is evidently stylized. Instead of assuming that λ_t is a piecewise constant function of r_{t-1} with a jump at D it would be more desirable to assume only that it is some increasing function of r_{t-1} . That would, however, also make the model more technically challenging to analyze. Our intuition tells us that a model with a more realistic assumption on λ_t 's dependence

on r_{t-1} would give results that are qualitatively similar to ours. Still, it would be nice to see the analysis of such a model carried out.

The copycat effect is introduced exogenously into the model. As we have mentioned earlier there are good reasons for assuming that a copycat effect exists. Nevertheless, it would be desirable to have a model where the copycat effect follows endogenously from the dynamic interaction between terrorists and authorities (and perhaps the public and the media). This is an interesting direction for further research.

A different way of rationalizing that authorities increase the anti-terrorism level after a large attack is to assume that such an attack reveals information that the authorities use to update beliefs. For example, it could be information about the number of existing cells, the probability that a cell is formed during some period of time, or the striking capabilities of existing or new cells. Modelling this is another possible direction for further research.

6 References

1. Bueno de Mesquita, E. (2005). Conciliation, Counterterrorism, and Patterns of Terrorist Violence, *International Organization* 59: 145-176.
2. Enders, W. and Sandler, T. (1993). The Effectiveness of Anti-Terrorism Policies: Vector-Autoregression Intervention Analysis, *American Political Science Review* 87: 829-844.
3. Faria, J. R. (2003). Terror Cycles, *Studies in Nonlinear Dynamics and Econometrics* 7(1), Article 3.
4. Lapan, H. E. and Sandler, T. (1988). To Bargain or Not to Bargain: That is the Question, *American Economic Association Papers and Proceedings* 78: 16-20.
5. Rosendorff, P. and Sandler, T. (2004). Too Much of a Good Thing? The Proactive Response Dilemma, *Journal of Conflict Resolution* 48: 657-671.
6. Sandler, T. and Enders, W. (2004). An Economic Perspective on Transnational Terrorism, *European Journal of Political Economy* 20: 301-316.
7. Sandler, T and Lapan, H. E. (1988). The Calculus of Dissent: An Analysis of Terrorists' Choice of Target, *Synthese* 76: 245-261.
8. Stokey, N. L. and Lucas, Jr., R. E. (with Prescott, E. C.) (1989). Recursive Methods in Economic Dynamics. Harvard University Press.

7 Appendix

Proof of $\Delta'(a) = -p'(a)(1 - F(\varepsilon^))D(1 + E[\varepsilon|\varepsilon > \varepsilon^*])$.*

First write $\Delta(a)$ as

$$\Delta(a) = F(\varepsilon^*)D + (1 - F(\varepsilon^*))(1 - p(a))D + (1 - p(a))D \int_{\varepsilon^*}^{\infty} \varepsilon f(\varepsilon) d\varepsilon.$$

By differentiation we get

$$\begin{aligned} \Delta'(a) &= \frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*)D - \frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*)(1 - p(a))D - p'(a)(1 - F(\varepsilon^*))D \\ &\quad - p'(a)D \int_{\varepsilon^*}^{\infty} \varepsilon f(\varepsilon) d\varepsilon - (1 - p(a))D \frac{\partial \varepsilon^*}{\partial a} \varepsilon^* f(\varepsilon^*). \end{aligned}$$

By collecting terms we then get

$$\begin{aligned} \Delta'(a) &= -p'(a)(1 - F(\varepsilon^*))D(1 + E[\varepsilon|\varepsilon > \varepsilon^*]) \\ &\quad + \frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*)D(1 - (1 - p(a))(1 + \varepsilon^*)). \end{aligned}$$

Since $\varepsilon^* = \frac{1}{(1-p(a))} - 1$ it follows that the last term is equal to zero and thus we are done. \square

Proof of Theorem 4.1.

First note that it suffices to show that

$$\pi'(0) < 0, \quad \lim_{a \rightarrow \infty} \pi'(a) = \infty \quad \text{and} \quad \pi'' > 0.$$

We know that

$$\pi'(a) = c'(a) - \lambda p'(a)(1 - F(\varepsilon^*))D(1 + E[\varepsilon|\varepsilon > \varepsilon^*]).$$

Thus we have

$$\begin{aligned} \pi'(0) &= c'(0) - \lambda p'(0)(1 - F(0))D(1 + E[\varepsilon|\varepsilon > 0]) \\ &= c'(0) - \lambda p'(0)D(1 + E[\varepsilon]). \end{aligned}$$

And then it follows from the two first assumptions in 1. that $\pi'(0) < 0$. Now rewrite π' as

$$\pi'(a) = c'(a) - \lambda p'(a)(1 - F(\varepsilon^*))D - \lambda p'(a)D \int_{\varepsilon^*}^{\infty} \varepsilon f(\varepsilon) d\varepsilon.$$

By that expression and the assumption that $p'' \leq 0$ we see that

$$\begin{aligned}\pi'(a) &\geq c'(a) - \lambda p'(a)D(1 + E[\varepsilon]) \\ &\geq c'(a) - \lambda p'(0)D(1 + E[\varepsilon])\end{aligned}$$

Using that inequality it follows from the last assumption in 1. that $\lim_{a \rightarrow \infty} \pi'(a) = \infty$. By differentiating π' we get

$$\begin{aligned}\pi''(a) &= c''(a) - \lambda p''(a)(1 - F(\varepsilon^*))D(1 + E[\varepsilon | \varepsilon > \varepsilon^*]) \\ &\quad + \lambda p'(a) \frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*)D(1 + \varepsilon^*).\end{aligned}$$

By our assumptions the first term is strictly positive and each of the last two terms are non-negative. Thus we have $\pi'' > 0$. \square

Proof of Lemma 4.2.

Define the map T from the set of real functions on $\{s, l\}$ (which can be identified with \mathbb{R}^2) into itself by

$$(Tf)(x) = \inf_{a \in [0, \infty)} [\pi(a, x) + \delta(P(s, x, a)f(s) + P(l, x, a)f(l))].$$

It is easily checked that T satisfies Blackwells sufficient conditions for a contraction (see e.g. Stokey and Lucas (1989), Theorem 3.3, p. 54). And then it follows by Banach's Fixed Point Theorem / The Contraction Mapping Theorem (see e.g. Stokey and Lucas (1989), Theorem 3.2, p. 50) that there exists a unique \bar{V} such that

$$T\bar{V} = \bar{V}.$$

Furthermore we have $T^n f \rightarrow \bar{V}$ for all f .

To show $\bar{V}(s) < \bar{V}(l)$ it suffices to show that, for any f ,

$$f(s) \leq f(l) \Rightarrow (Tf)(s) < (Tf)(l).$$

Because then we can pick such an f to get

$$\bar{V}(s) = \lim_n (T^n f)(s) \leq \lim_n (T^n f)(l) = \bar{V}(l)$$

and thus

$$\bar{V}(s) = (T\bar{V})(s) < (T\bar{V})(l) = \bar{V}(l).$$

Suppose $f(s) \leq f(l)$. Then we have

$$(Tf)(l) = \inf_{a \in [0, \infty)} [\pi(a, s) + \delta(P(s, s, a)f(s) + P(l, s, a)f(l)) \\ + (\lambda^l - \lambda^s)\Delta(a) + \delta(\lambda^l - \lambda^s)(1 - p(a))(1 - F(\varepsilon^*))(f(l) - f(s))].$$

From this equation we see that

$$(Tf)(l) \geq \inf_{a \in [0, \infty)} [\pi(a, s) + \delta(P(s, s, a)f(s) + P(l, s, a)f(l))] + (\lambda^l - \lambda^s) \inf_{a \in [0, \infty)} [\Delta(a)] \\ = (Tf)(s) + (\lambda^l - \lambda^s) \inf_{a \in [0, \infty)} [\Delta(a)]$$

and since $\inf_{a \in [0, \infty)} [\Delta(a)] = D$ it then follows that

$$(Tf)(l) > (Tf)(s).$$

□

Proof of Theorem 4.3.

First note that to prove the first statement of the theorem it suffices to show that, for each $x \in \{s, l\}$,

$$\frac{\partial}{\partial a}g(a, x)|_{a=0} < 0, \quad \lim_{a \rightarrow \infty} \frac{\partial}{\partial a}g(a, x) = \infty \quad \text{and} \quad \frac{\partial^2}{\partial a^2}g(a, x) > 0.$$

By differentiation we get (after collecting some terms)

$$\frac{\partial}{\partial a}g(a, x) = \frac{\partial}{\partial a}\pi(a, x) \\ + \delta\lambda^x[-p'(a)(1 - F(\varepsilon^*)) - (1 - p(a))\frac{\partial\varepsilon^*}{\partial a}f(\varepsilon^*)](\bar{V}(l) - \bar{V}(s)).$$

By plugging in $a = 0$ and using Theorem 4.1 we get

$$\frac{\partial}{\partial a}g(a, x)|_{a=0} < \frac{\partial}{\partial a}\pi(a, x)|_{a=0} < 0.$$

Since $p'' \leq 0$ and $\frac{\partial}{\partial a}(\frac{\partial\varepsilon^*}{\partial a}f(\varepsilon^*)) \leq 0$ it follows that the term

$$p'(a)(1 - F(\varepsilon^*)) + (1 - p(a))\frac{\partial\varepsilon^*}{\partial a}f(\varepsilon^*)$$

is bounded. By Theorem 4.1 we have $\lim_{a \rightarrow \infty} \frac{\partial}{\partial a} \pi(a, x) = \infty$ and thus we can conclude that

$$\lim_{a \rightarrow \infty} \frac{\partial}{\partial a} g(a, x) = \infty.$$

By differentiation of $\frac{\partial}{\partial a} g(a, x)$ we get

$$\begin{aligned} \frac{\partial^2}{\partial^2 a} g(a, x) &= \frac{\partial^2}{\partial^2 a} \pi(a, x) \\ &+ \delta \lambda^x [-p''(a)(1 - F(\varepsilon^*)) + 2p'(a) \frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*) - (1 - p(a)) \frac{\partial}{\partial a} (\frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*))] (\bar{V}(l) - \bar{V}(s)). \end{aligned}$$

By our assumptions the term in the square brackets is non-negative and by Theorem 4.1 we have $\frac{\partial^2}{\partial^2 a} \pi(a, x) > 0$. Thus we see that

$$\frac{\partial^2}{\partial a^2} g(a, x) > 0.$$

To prove the last statement of the theorem note that

$$\begin{aligned} \frac{\partial}{\partial a} g(a, l) &= \frac{\partial}{\partial a} g(a, s) \\ &+ (\lambda^l - \lambda^s) [\Delta'(a) - \delta(\bar{V}(l) - \bar{V}(s))(p'(a)(1 - F(\varepsilon^*)) + (1 - p(a)) \frac{\partial \varepsilon^*}{\partial a} f(\varepsilon^*))]. \end{aligned}$$

Since the term in the square brackets is negative we have

$$\frac{\partial}{\partial a} g(a, l) < \frac{\partial}{\partial a} g(a, s).$$

Therefore

$$\frac{\partial}{\partial a} g(a, l)|_{a=\bar{a}_+(s)} < \frac{\partial}{\partial a} g(a, s)|_{a=\bar{a}_+(s)} = 0.$$

And then it easily follows by the convexity of $g(a, l)$ that

$$\bar{a}_+(s) < \bar{a}_+(l).$$

□