National Responses to Transnational Terrorism: Intelligence and Counterterrorism Provision*

Thomas Jensen†

November 15, 2012

Abstract

Intelligence about transnational terrorist threats is generally gathered by national agencies. I set up and analyze a game theoretic model to study the implications of national intelligence gathering for the provision of domestic (defensive) counterterrorism when two countries are facing a transnational terrorist threat. It is shown that, relative to a benchmark case where all intelligence is commonly known, national intelligence gathering often leads to increased overprovision, although it can be the other way around. By extending the model with a communication stage, I also explore the possibilities for intelligence sharing prior to decisions on counterterrorism provision. If verifiable sharing is a viable option for each country, there exists an equilibrium with full intelligence sharing. On the other hand, if only cheap talk communication is possible then full sharing cannot happen in equilibrium.

Keywords: Transnational Terrorism, Counterterrorism, Intelligence

---

*I gratefully acknowledge funding from The Danish Council for Independent Research | Social Sciences (grant number 09-066752).

†Department of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, Denmark. Email: Thomas.Jensen@econ.ku.dk.

1 Introduction

Counterterrorism policy decisions are generally made in an environment of uncertainty. Intelligence is rarely complete, so policy makers face difficult choices when deciding whether the benefits of particular policies outweigh the costs. When the terrorist threat is transnational, such as the threat from Al-Qaeda and related groups in recent years, a new dimension of uncertainty is added. Since intelligence is generally gathered by national agencies (e.g., Walsh 2009), the authorities in one country may well be uncertain not only about the terrorists' capabilities, but also about the information available to other countries. And since counterterrorism policies are decided on the national level, this type of uncertainty has potentially large implications for the provision of counterterrorism. Because of transnational externalities, the optimal policy for one country will typically depend on the policies of other countries. Therefore, national intelligence gathering leads to a fundamentally different strategic situation for the targeted countries than if all intelligence was commonly known.

The purpose of this paper is to study the implications of national intelligence gathering for the provision of domestic (defensive) counterterrorism[1] by setting up and analyzing a game theoretic model. In the model, two countries receive private signals about the capabilities of a transnational terrorist organization and then independently choose whether to make a costly investment in domestic counterterrorism or not. Investment reduces the likelihood that an attempted terrorist attack in the homeland will be successful. Therefore, if one country invests while the other does not, the terrorists will attack the non-investing country.[2] Otherwise the terrorists are equally likely to attack each country.

To understand the implications of national intelligence gathering, the outcome of the private information game is compared to the outcome of a common intelligence benchmark where both signals are known to both countries before they make their policy decisions. The benchmark model displays the well known result that, due to negative transnational externalities, there will be overprovision of domestic (defensive) counterterrorism (e.g., Sandler and Siquiera 2006). So a fundamental question is whether the introduction of national intelligence gathering leads to increased overprovision or it is the other way around.

I present results showing that national intelligence will often lead to more overprovision than if all intelligence was commonly known. However, it can be the other way

[1] The type of counterterrorism considered can be referred to as domestic and/or defensive. The important characteristic is that it involves actions in the homeland that make terrorist attacks less likely to succeed there without substantially reducing the general capabilities of the transnational terrorist group(s). I refer to it as domestic because it can include measures that would typically be called proactive, for example programs that seek to infiltrate terrorist cells in the homeland. Such measures make it harder for cells to fulfill their missions and thus the terrorist group will become more likely to fund/send cells elsewhere.

[2] For empirical evidence on transference of terrorist attacks, see Enders and Sandler 1993, 2004, 2006, and Sandler and Enders 2004.

around. More precisely, suppose first that the common ex ante expectation about the terrorists' capabilities is sufficiently high that, if no further intelligence was received, the countries would both invest. Then national intelligence gathering leads to an increase in overprovision of domestic counterterrorism relative to the benchmark. On the other hand, if the ex ante expectation about the capabilities of the terrorists is sufficiently low, it can be that the outcome is more efficient when the countries' signals are private.

The model's main contribution lies in the intuition it provides about the implications of national intelligence gathering. Intuitively, why does the realistic assumption that countries facing a common transnational terrorism threat have private information often lead to increased overprovision of domestic counterterrorism? The main reason is that a country's incentive to invest in domestic counterterrorism is higher when it expects the other country to invest. Suppose we are in the benchmark case and that the signals received are such that it is just barely optimal for each country not to invest if the other country does the same. Then, in equilibrium, the countries can coordinate on mutual non-investment. However, if the signals are private then, from the point of view of country i, there is always a possibility that country j will invest. This uncertainty may well imply that it is optimal for country i to invest. So national intelligence gathering leads to uncertainty about the other country's decision, which again leads to investment in situations where there would be mutual non-investment with commonly known signals. The reason why this intuition does not always hold will be carefully explained later.

Even though intelligence is gathered by national agencies, it could be that countries are generally able and willing to share intelligence such that the amount of private information becomes negligible. I explore the possibilities for intelligence sharing by extending the main model with a communication stage taking place before the simultaneous investment decisions. It turns out that results on intelligence sharing are highly dependent on whether it is possible to verifiably share intelligence or not. If intelligence can be shared in a way such that the content can be verified by the receiver, full intelligence sharing can happen in equilibrium and then the outcome of the investment game will be the benchmark outcome. However, the nature of intelligence and/or the relations between targeted countries and their intelligence agencies often makes verifiable sharing either completely impossible or at least a non-viable option (Walsh 2006, 2009). Therefore, I also consider a case where only cheap talk communication is possible. With this assumption, intelligence will not be fully shared. Thus there will be private information in the subsequent counterterrorism provision game.

The reasons behind the results on intelligence sharing are relatively simple. In both cases, the key to the results is the observation that each country always prefers that the other country does not invest. This implies that each country prefers the other country to believe that it has received a low signal about the capabilities of the terrorists. With verifiable sharing, this means that a country receiving a low signal will reveal it, because this will make the other country less likely to invest. Further, this observation will make slightly higher types of the country reveal as well (to distinguish itself from types with even higher signals) and so on. Thus it follows that all signals will be revealed.

When only cheap talk communication is possible, the incentive of each country to minimize the probability that the other country invests means that there is always an incentive to pretend to have received a low signal. This implies that full intelligence sharing is not possible.

This paper builds on a substantial literature about counterterrorism provision. Sandler and Lapan (1988), Arce and Sandler (2005), and Sandler and Siquiera (2006) all study national counterterrorism provision when two countries are facing a transnational terrorist threat. As mentioned above, a central result is that defensive counterterrorism is typically oversupplied because of negative externalities. It is assumed that all countries have the same information, so this paper adds to this strand of literature by introducing the realistic assumption that intelligence is gathered on the national level and exploring the possibilities for intelligence sharing. Rosendorff and Sandler (2004) show that proactive policies can also lead to negative externalities if they increase terrorist recruitment and new attacks are transferred abroad. When only targets within a country is considered, terrorists' substitution between targets is internalized by the authorities and optimal levels of counterterrorism can be reached. Powell (2007) and Bueno de Mesquita (2007) characterizes optimal policies in this setting.

2 The Model

Two countries, $i = 1, 2$, are facing a common terrorist threat from a transnational terrorist organization. Each country can decide whether to make a costly investment in domestic counterterrorism or not.[3] The cost is denoted C. Investment reduces the vulnerability of the country to terrorism. More precisely, the counterterrorism investment reduces the probability of an attempted attack being successful from one to $p \in (0, 1)$. The countries make their decisions simultaneously and independently. Each country's objective is to minimize the sum of expected damages from terrorism and counterterrorism costs.[4]

The terrorists observe the investment decisions of the countries before deciding where to attack. They can launch only one attack and their objective is to maximize expected damage. Thus, if only one country invests in domestic counterterrorism then the terrorists will attack the other country because the attack is then more likely to be sucessful. If neither or both countries invest, the terrorists are indifferent about where to attack. I assume that they will then attack each country with probability one half. The damage from a sucessful attack is denoted D.

[3] The setup also covers implementation of domestic counterterrorism policies that are not directly costly, but leads to substantial distortionary costs for the country.

[4] This type of objective function for the government of a country was first used by Sandler and Lapan (1988).

If the two countries are fully informed about D, the situation can be represented by the simple 2×2 game shown below.

		Country 2	
		Invest (I)	Not Invest (N)
Country 1	Invest (I)	$-\frac{p}{2}D - C, -\frac{p}{2}D - C$	$-C, -D$
	Not Invest (N)	$-D, -C$	$-\frac{1}{2}D, -\frac{1}{2}D$

This game is easily analyzed. The pure strategy Nash equilibria are[5]

$$\begin{aligned} (N, N) \quad & \text{if} \quad D \leq 2C, \\ (I, I) \quad & \text{if} \quad D \geq \frac{2C}{2-p}. \end{aligned}$$

The efficient outcome (lowest total sum of expected damages and costs) is (N, N) if $D \leq \frac{2C}{1-p}$ and (I, I) if $D \geq \frac{2C}{1-p}$. Thus the efficient Nash equilibrium is

$$\begin{aligned} (N, N) \quad & \text{if} \quad D \leq 2C, \\ (I, I) \quad & \text{if} \quad D > 2C. \end{aligned}$$

So we see that, assuming the countries are always able to coordinate on the efficient equilibrium, the outcome of the full information game will be efficient unless $D \in (2C, \frac{2C}{1-p})$. In this region of the parameter space the game is a prisoners dilemma: I strictly dominates N for each country, while the efficient outcome is that neither country invests. This is a representation of the well known result that countries facing a common transnational terrorist threat will often overinvest in defensive counterterrorism because of negative externalities: increased investment in one country makes the terrorists more likely to attack elsewhere.

2.1 National Intelligence

In the real world, countries are unlikely to have full information about the capabilities of transnational terrorist organizations. Intelligence is rarely complete. Furthermore, intelligence is generally gathered by national agencies. Thus, counterterrorism authorities are typically uncertain about both the capabilities of the terrorists and the intelligence available to other countries. Below I will analyze the implications of introducing national gathering of intelligence in the simple model presented above. I will assume that each country receives a private signal about the capabilities of the terrorists, more precisely a signal correlated with D. Thus the intelligence gathering process is completely exogenous. The outcome of the private information game will be compared to a benchmark case where both signals are commonly known to the two countries.

[5] For $D \in [\frac{2C}{2-p}, 2C]$ there is also a mixed strategy Nash equilibrium, which is straightforward to find.

Let $d_i \in [0, \bar{d}]$ denote the signal received by country i. This signal is used to form a belief about the distribution of D. In particular, after receiving the signal d_i, country i's expectation about D is $E[D|d_i]$. I assume that $E[D|d_i]$ is a continuously differentiable and strictly increasing function of d_i. It is also assumed that the expected value of D given both signals, $E[D|d_1, d_2]$, is continuously differentiable, strictly increasing with respect to both d_1 and d_2, and symmetric.

Since the signal of each country is correlated with D, in general the two signals are also correlated. Let $F(\cdot|\cdot)$ denote the cumulative distribution function of one country's signal given the signal of the other country (the situation is assumed to be completely symmetric). I assume that $F(d_i|d_j)$ is continuously differentiable with respect to each signal. Furthermore, it is also assumed that an increase in d_j shifts the conditional distribution of d_i (weakly) to the right. More precisely, if $d' < d''$ then the distribution of d_i conditional on $d_j = d''$ first order stochastically dominates the distribution of d_i conditional on $d_j = d'$:

$$F(d_i|d_j = d'') \leq F(d_i|d_j = d') \text{ for all } d_i.$$

I will also make the following minor technical assumption. Let d' be a signal such $E[D|d_i = d'] > E[D]$ ($E[D|d_i = d'] < E[D]$), where $E[D]$ denotes the ex ante expected value of D. Then it is assumed that

$$\begin{aligned} E[D|d', d'] &> E[D|d_i = d'] \\ (E[D|d', d'] &< E[D|d_i = d']). \end{aligned}$$

In other words, if the reception of a particular signal makes a country update its expected value of D in one direction, then learning that the other country has received the same signal will make the country update its expected value of D further in the same direction.

In the national intelligence game, each country can condition its decision whether to invest or not on its signal. We write the strategy of country i as $s_i(d_i)$. Since investment is more attractive the more damage the terrorists are capable of causing, it is natural to look for (Bayesian Nash) equilibria where each country invests if and only if its signal is above some cutoff value. Furthermore, since the game is symmetric, we will restrict attention to symmetric cutoff equilibria. Thus, an equilibrium is given by a cutoff signal x such that each country will invest if and only if its signal is above x:

$$s_i^*(d_i) = \left\{ \begin{array}{c} N \text{ if } d_i \leq x \\ I \text{ if } d_i > x \end{array} \right\}, i = 1, 2. \tag{1}$$

In equilibrium it should be optimal for each country to use this strategy given that the other country is using it.

3 Analysis of the National Intelligence Game

Before analyzing the national intelligence game, I will first analyze the natural benchmark, namely the game where the signals are commonly known by the two countries.

In this game each country can condition its decision on both signals, so the strategy of country i is written $s_i(d_1, d_2)$. This game is closely related to the full information game because the two countries have exactly the same information available. Indeed, it is easy to see that the efficient (Bayesian Nash) equilibrium is given by the strategies

$$s_i^*(d_1, d_2) = \left\{ \begin{array}{c} N \text{ if } E[D|d_1, d_2] \leq 2C \\ I \text{ if } E[D|d_1, d_2] > 2C \end{array} \right\}, i = 1, 2. \tag{2}$$

In other words, the game is essentially equivalent to the full information game with D replaced by $E[D|d_1, d_2]$. So the natural benchmark to which the outcome of the national intelligence game should be compared is that each country will invest precisely if the expected value of D given all available information in the game is higher than $2C$.

To analyze the national intelligence game, assume that country j is using the strategy given by the cutoff signal x. That is, country j invests if and only if $d_j > x$. Then country i's sum of expected terrorism damages and counterterrorism costs if it invests is

$$C + \frac{p}{2}(1 - F(x|d_i))E[D|d_i, d_j > x].$$

If country i chooses not to invest then the expected damage is

$$\begin{aligned} &\frac{1}{2}F(x|d_i)E[D|d_i, d_j \leq x] + (1 - F(x|d_i))E[D|d_i, d_j > x] \\ &\qquad = \frac{1}{2}E[D|d_i] + \frac{1}{2}(1 - F(x|d_i))E[D|d_i, d_j > x]. \end{aligned}$$

Thus it is optimal for country i not to invest precisely if

$$\frac{1}{2}E[D|d_i] + \frac{1}{2}(1 - F(x|d_i))E[D|d_i, d_j > x] \leq C + \frac{p}{2}(1 - F(x|d_i))E[D|d_i, d_j > x],$$

which is equivalent to

$$E[D|d_i] + (1 - p)(1 - F(x|d_i))E[D|d_i, d_j > x] \leq 2C. \tag{3}$$

Each of the terms on the left hand side of the inequality are increasing and continuous in d_i (this follows easily from earlier assumptions). Thus it is a best response for country i to use the cutoff strategy given by x precisely if

$$E[D|d_i = x] + (1 - p)(1 - F(x|x))E[D|d_i = x, d_j > x] = 2C. \tag{4}$$

So if x satisfies this equation then we have an equilibrium. If $E[D|d_i = 0] < \frac{2C}{2-p}$ and $E[D|d_i = \bar{d}] > 2C$ then there exists a solution $x > 0$ to the equation. In the following it is assumed that these assumptions hold. While we do not necessarily have uniqueness of equilibrium, it is easy to establish that a highest cutoff equilibrium exists and that this is the efficient equilibrium. All these results are proved in the appendix.

From equation (4) it immediately follows that $E[D|d_i = x] < 2C$. Recall that in the benchmark outcome each country invests precisely if $E[D|d_1, d_2] > 2C$. Thus, in any equilibrium of the national intelligence game each country will choose to invest for lower expectations about the capabilities of the terrorists than in the common intelligence benchmark. The intuition is simple and illuminating. In the common intelligence game the countries can coordinate on not investing as long as the common expectation about D is such that N is the best response to N in the full information game. In the national intelligence game each country is uncertain about the intelligence gathered by the other country. Thus there is always a positive probability that the other country's expected value of D will make it invest. This means that, even when country i's expected D is such that the countries could coordinate on not investing if that expectation had been common for the two countries, it can be optimal for country i to invest. Below I formulate this observation as a proposition.

Proposition 1 *In any equilibrium of the national intelligence (private information) game, each country will invest in domestic counterterrorism for lower expected capabilities of the terrorists than in the efficient equilibrium of the common intelligence benchmark game.*

While this result is interesting, it does not imply that national intelligence gathering necessarily worsens the overprovision of domestic counterterrorism relative to the benchmark. Suppose both countries have received the cutoff signal. Then each country's expected value of D is $E[D|x] < 2C$. However, it could be that the expected value of D given the information that both signals is equal to x, $E[D|x, x]$, is higher than $2C$. If this is the case then there exist signal pairs such that national intelligence gathering leads to efficient non-investment while the countries would have both invested if all intelligence was commonly known. In the following I will analyze when national intelligence gathering does in fact lead to more inefficient provision of domestic counterterrorism.

Clearly, national intelligence gathering can lead to outcomes where one country invests while the other does not. This is a rather trivial consequence of introducing private information, so when comparing the games with private and public information I will primarily focus on the part of the space of possible signal pairs where, in the national intelligence equilibrium, the countries make the same decision.

The comparison of the outcomes of the national intelligence game and the benchmark game depends critically on whether $E[D|x, x] < 2C$ or $E[D|x, x] > 2C$. Examples of these two cases are shown in figure 1 and 2. Remember that in the efficient equilibrium of the benchmark game the countries both invest precisely if the expected D given both signals is above $2C$. The efficient outcome is for both countries to invest if $E[D|d_1, d_2] > \frac{2C}{1-p}$ and for none of them to invest otherwise.

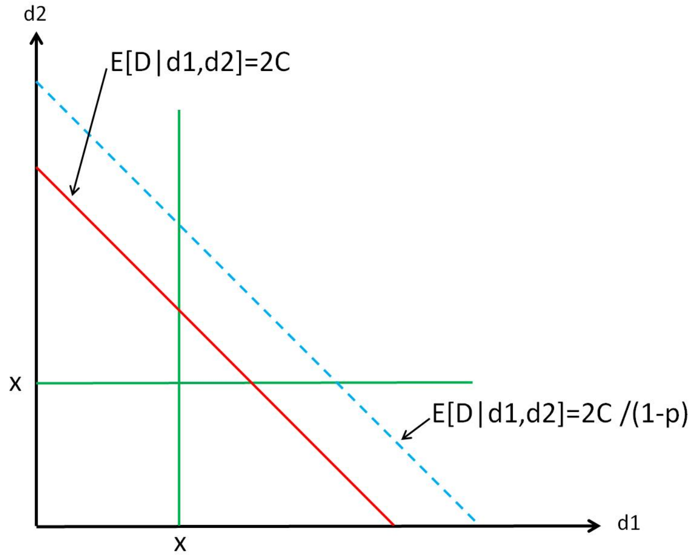

Figure 1: $E[D|x,x] < 2C$

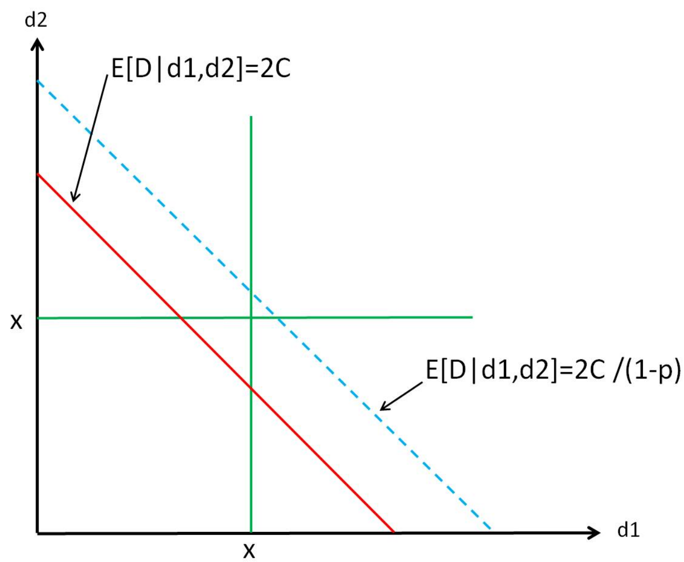

Figure 2: $E[D|x,x] > 2C$

If $E[D|x,x] < 2C$ then the set of signal pairs such that neither country invests when intelligence is gathered nationally is contained in the corresponding set for the benchmark game. Thus it is reasonable to say that national gathering of intelligence generally makes the overprovision of domestic counterterrorism worse relative to a world where all intelligence is commonly known. However, note that for the benchmark outcome to be at least as good as the national intelligence outcome for *all* realized signal-pairs we need $p \leq \frac{1}{2}$. The problem is that when one country invests while the other does not in the national intelligence game and $E[D|d_1, d_2] \in (2C, \frac{2C}{1-p})$, then (I, I) is not necessarily better than only one country investing. This is only the case if $p \leq \frac{1}{2}$.[6]

Now suppose $E[D|x, x] > 2C$. Then, if we are in the national intelligence game and both countries have received the cutoff signal, each country's expected value of D is below $2C$, but the expected value given all available information in the game is above $2C$. This implies that there exist signal pairs such that national intelligence gathering leads to efficient non-investment by the two countries, while the outcome with commonly known intelligence would be mutual investment. More specifically, this is the case for signal pairs with $d_1, d_2 \leq x$ and $E[D|d_1, d_2] \in (2C, \frac{2C}{1-p})$.

If $E[D|x, x] > \frac{2C}{1-p}$ then national intelligence leads to underprovision of domestic counterterrorism since (I, I) is the efficient outcome when the expected value of D is above $\frac{2C}{1-p}$. So while the national intelligence outcome is still better than the benchmark for signals $d_1, d_2 \leq x$ with $E[D|d_1, d_2] \in (2C, \frac{2C}{1-p})$, the benchmark is better for signals $d_1, d_2 \leq x$ with $E[D|d_1, d_2] > \frac{2C}{1-p}$.

Finally, it should be noted that if we take into consideration signal pairs where $d_i < x < d_j$ such that one country invests while the other does not, then there are always (no matter if $E[D|x, x]$ is above or below $2C$) regions where the benchmark is better than the national intelligence outcome. For example, if $d_1 < x$, $d_2 > x$, and $E[D|d_1, d_2] > \frac{2C}{1-p}$, then the (I, I) outcome results in a lower total sum of expected damages and costs than (N, I). So even though national gathering of intelligence in some sense mitigates the overprovision problem when $E[D|x, x] \in (2C, \frac{2C}{1-p})$, we never have that the national intelligence outcome is at least as good as the benchmark for *all* possible pairs of signals. Thus, the "national intelligence is better than the benchmark" result when $E[D|x, x] \in (2C, \frac{2C}{1-p})$ is clearly weaker than the opposite result when $E[D|x, x] < 2C$ and $p \leq \frac{1}{2}$.

The conclusions reached above are collected in the following proposition.

Proposition 2 *Let x be an equilibrium of the national intelligence game. Then the following statements hold.*

1. *If $E[D|x, x] < 2C$ then, for all signal pairs such that the national intelligence outcome is (N, N) or (I, I), the benchmark outcome results in a weakly lower total sum of counterterrorism costs and terrorism damages (strictly lower for some signal pairs). If $p \leq \frac{1}{2}$ this holds for all signal pairs.*

[6] To see this, note that (I, I) is better than (I, N) (or (N, I)) precisely if

$$2C + pE[D|d_1, d_2] \leq C + E[D|d_1, d_2],$$

which is equivalent to

$$E[D|d_1, d_2] \geq \frac{C}{1-p}.$$

So, when $E[D|x, x] < 2C$, the benchmark outcome is better than the national intelligence outcome for all signal pairs precisely if $\frac{C}{1-p} \leq 2C$, which is equivalent to $p \leq \frac{1}{2}$.

2. *If $E[D|x,x] \in (2C, \frac{2C}{1-p})$ then, for all signal pairs such that the national intelligence outcome is (N,N) or (I,I), the benchmark outcome results in a weakly higher total sum of counterterrorism costs and terrorism damages (strictly higher for some signal pairs). However, there exist signal pairs such that the national intelligence outcome is (N,I) or (I,N) and the benchmark outcome results in a strictly lower total sum of costs and damages.*

3. *If $E[D|x,x] > \frac{2C}{1-p}$ then there are signal pairs such that the national intelligence outcome is (N,N) and the benchmark outcome results in a strictly lower total sum of counterterrorism costs and terrorism damages. However, there are also signal pairs with the same national intelligence outcome such that the benchmark outcome results in a strictly higher total sum of costs and damages.*

While the proposition reveals how the comparison of the national intelligence outcome and the common intelligence benchmark depends on the expected capabilities of the terrorists conditional on both countries receiving the cutoff signal, it does not directly reveal how the comparison depends on the primitives of the model. For example, under which conditions does it hold that $E[D|x,x] < 2C$?

Since $E[D|x] < 2C$, it immediately follows that we are in case one of Proposition 2 if $E[D|x,x] \leq E[D|x]$. By an assumption made earlier, this is the case if $E[D|x] < E[D]$. So if $E[D] \geq 2C$ then it follows that $E[D|x,x] < 2C$. Thus we have the following important result.

Proposition 3 *Suppose the common ex ante expected value of D is above the level where it makes investment a strictly dominating strategy for each country if no further information is received ($E[D] \geq 2C$). Then national intelligence gathering makes the overprovision of domestic counterterrorism worse relative to the common intelligence benchmark.*

If $E[D] < 2C$ then it is possible that that we are in case two or three of Proposition 2. Below I present an example of the general model studied so far. The example shows that it is indeed possible for national intelligence gathering to lead to more efficient outcomes than if all intelligence was commonly available. However, for this to happen $E[D]$ has to be substantially below $2C$, especially when p is relatively small (i.e., when investment in counterterrorism is quite effective in reducing a country's vulnerability to a terrorist attack).

3.1 An Example

Suppose the two countries collect independent intelligence about different parts of the transnational terrorist organization's activities. For example, it could be that each country primarily collects intelligence in a specific geographical region because their intelligence network is most developed there for historical or other reasons. Further, assume that each country learns everything about the particular activities it surveils and that no activities escape scrutiny, such that the true capabilities of the entire terrorist organization is simply the sum of the signals received by the two countries:

$$D = d_1 + d_2.$$

Thus, I basically assume away the inherent noisiness of intelligence. A country is only uncertain about the capabilities of the terrorists because it does not have information about the activities surveilled by the other country. The signals d_1 and d_2 are assumed to be independent and uniformly distributed on $[0, 1]$. While this is clearly a stylized example, I believe that it serves the purpose of illustrating the mechanisms at play in the model, for example how national intelligence gathering can, under some circumstances, make the provision of domestic counterterrorism more efficient.

In this example, the benchmark outcome is (N, N) if $d_1 + d_2 \leq 2C$ and (I, I) otherwise. Note that the benchmark game is completely identical to the full information game because the true value of D is known when both signals are known.

In the game with national intelligence gathering, the equilibrium equation (4) becomes

$$x + \frac{1}{2} + (1-p)(1-x)(x + \frac{1+x}{2}) = 2C.$$

This is simply a second order equation in x:

$$3(1-p)x^2 - 2(2-p)x - (2-p) + 4C = 0. \tag{5}$$

The assumptions made earlier that $E[D|d_i = 0] < \frac{2C}{2-p}$ and $E[D|d_i = \bar{d}] > 2C$ become

$$\frac{1}{2} < \frac{2C}{2-p} \text{ and } 1 + \frac{1}{2} > 2C,$$

which is equivalent to

$$\frac{2-p}{4} < C < \frac{3}{4}.$$

Here these assumptions imply that there is a unique equilibrium $x \in (0, 1)$.[7] This equilibrium is found by solving equation (5) for the smallest root (with the assumptions made, the polynomial on the left hand side is positive at $x = 0$ and negative at $x = 1$). The exact expression for x can be found in the appendix.

Having found the equilibrium x in the national intelligence game, we are now ready to analyze the example with respect to when national intelligence gathering makes overprovision of domestic intelligence worse than in the benchmark model and when it is the other way around. From the general theory we know that if $E[D] \geq 2C$ then national

[7] If $C \leq \frac{2-p}{4}$ then we still have a unique equilibrium: $x = 0$. If $C \geq \frac{3}{4}$ then $x = 1$ is an equilibrium. And while there can be other equilibria, $x = 1$ is the efficient one.

intelligence gathering does lead to worse outcomes. Here $E[D] = E[d_1] + E[d_2] = 1$, so the condition $E[D] \geq 2C$ corresponds to $C \leq \frac{1}{2}$. Therefore, we are primarily interested in analyzing what happens when $\frac{1}{2} < C < \frac{3}{4}$.

First, consider the special case $p = \frac{1}{2}$. From the formula in the appendix we get that the equilibrium is

$$x = 1 - \sqrt{2 - \frac{8}{3}C}.$$

So $E[D|x, x] = x + x = 2 - 2\sqrt{2 - \frac{3}{8}C}$. This means that we are in case one of Proposition 2 precisely if

$$2 - 2\sqrt{2 - \frac{8}{3}C} < 2C,$$

i.e., if

$$C < -\frac{1}{3} + \frac{1}{6}\sqrt{40} \approx .72.$$

So, efficiency wise, national intelligence gathering is worse than the benchmark for values of C such that the ex ante expected value of D (which is equal to one) is substantially above $2C$. But if, approximately, $C \in (.72, .75)$ then we are in case two of Proposition 2 (it is easy to check that we are not in case three).

In figure 3 I have plotted $E[D|x, x] - 2C$ as a function of C for different values of p ($p = .1, .25, .5, .75$, and $.9$). So for each value of p it is easy to see when $E[D|x, x]$ is below $2C$, such that we can use Proposition 2 to conclude that national intelligence gathering makes overprovision of domestic counterterrorism worse than in the benchmark model. We see that the smaller p is, the larger is the region where $E[D|x, x] < 2C$ (this observation is formally proved in the appendix). Furthermore, for p sufficiently small, the benchmark outcome is better than the national intelligence outcome for all C's with $\frac{2-p}{4} < C < \frac{3}{4}$.

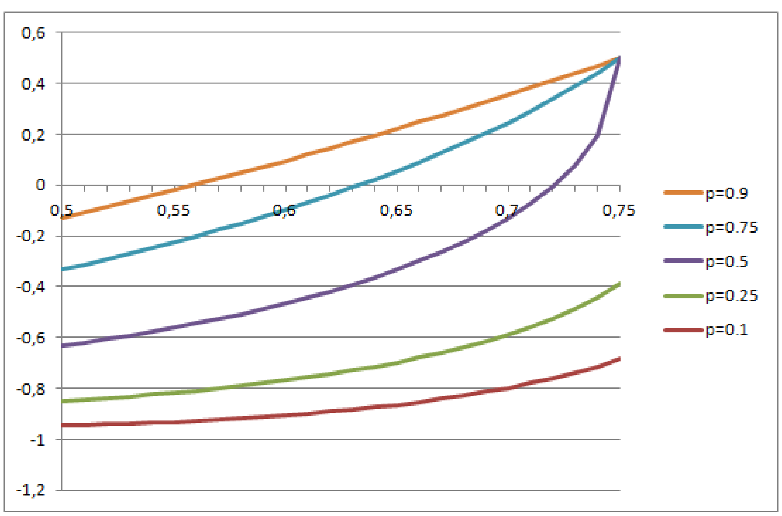

Figure 3: $E[D|x, x] - 2C$ as a function of C for different values of p

To sum up, this example has illustrated that, relative to the benchmark with commonly known intelligence, it is possible for national intelligence gathering to mitigate the overprovision of domestic intelligence due to negative externalities. If the ex ante expectation about the capabilities of the terrorists is very low relative to $2C$, then it is only if the observed intelligence greatly increases the expected value of D that a country will choose to invest. Such a strong adjustment of expectations can happen for lower signals if we are in the benchmark case where two signals are received than if each country only observes its own signal. And this effect can dominate the effect deriving from the fact that the adjustment of expectation needed to make a country invest is lower in the national intelligence game (because of uncertainty about the information of the other country, see Proposition 1 and the discussion before it). When this is the case, the national intelligence outcome is more efficient than the benchmark (if we only consider the signal pairs where the countries make the same decision in the national intelligence game).

We also saw that the smaller p is, the larger is the set of investment costs for which the benchmark is more efficient than the national intelligence outcome. The intuition behind this result is worth discussing. For all $p < 1$, the reduction in expected domestic damage for a country if it invests is higher when the other country invests than when it does not. However, the difference in reduction is larger the smaller p is. This means that the smaller p is, the more it matters that, with national intelligence gathering, there is always a positive probability that the other country will invest. And this implies that national intelligence gathering leads to increased overprovision of domestic intelligence for a larger set of C's when p is smaller.

Finally, it is worth emphasizing that if investment in counterterrorism is just reasonably effective in reducing a country's vulnerability to a terrorist attack, national intelligence gathering leads to increased overprovision unless the ex ante expectation of D is very low relative to $2C$. So, in times of a substantial general terrorist threat such as the threat from Al-Qaeda and related groups in recent years, the relevant implication of the model is that, because intelligence is gathered on the national level, overprovision of domestic counterterrorism will be even worse than what is to be expected solely because of the negative transnational externalities.

4 Intelligence sharing

Until now we have assumed that national gathering of intelligence implies that the intelligence collected by each country is private information when the countries decide whether to invest in domestic counterterrorism or not. However, it could be that the countries share their intelligence. If so, the observation that intelligence is gathered by national agencies is clearly less relevant for counterterrorism provision.

In this section I explore whether intelligence will actually be shared by the individual countries. I do so by extending the game studied above with a communication stage taking place before the simultaneous investment decisions. In this stage, each country can either send a message to the other country or not. The possible messages are the possible signals that the countries can receive. If each country always sends a message equal to its received signal then all intelligence become commonly known and the subsequent investment game will then be completely equivalent to the benchmark game. If each country never sends a message (or, for example, always sends the message $m_i = 0$), no information is revealed and the countries will then play the national intelligence game in the second stage.

I will distinguish between two cases with respect to the possibilities for intelligence sharing. In the first case it is assumed that the messages are pure cheap talk. Thus the sender country can send any message at no cost and the receiver country has no way of verifying that a message is correct. The nature of intelligence implies that it is often not possible to directly communicate the content of a particular piece of intelligence in a verifiable way. For example, information that one country has obtained by interviewing informants can typically not be shared with another country in a way such that this country can be sure that the content has not been manipulated. Furthermore, and perhaps more importantly, intelligence agencies generally have strong incentives not to reveal their sources and other details about their activities (e.g., Walsh 2009). Thus, even when some piece of intelligence can in principle be shared verifiably, doing so may well not be seen as a viable option because it would reveal, for example, the sources from which the intelligence was obtained.

The second case corresponds to a situation where the content of a message is verifiable. I assume that each country can either send the true message $m_i = d_i$ (and country j will then know the signal of country i with certainty) or it can choose not to send a message.[8] This case could be relevant when the two countries are close allies and there is a large degree of trust between their intelligence agencies.[9] With generally shared interests and a high level of trust, it is likely that authorities are less reluctant to reveal the highly classified information necessary for verifiable intelligence sharing because it is unlikely to be passed on and to be used against the interest of the sharing country. Furthermore, it can be argued that even when the intelligence collected is of a kind that makes verifiable sharing difficult, large potential reputation costs from sending a false message to a trusted partner implies that this could still be the relevant case, at least as an approximation. Thus, to give an example, the cheap talk case is likely to be more relevant if the two countries are the US and Pakistan while the verifiable messages case is more relevant if they are the US and the UK.

It turns out that results on intelligence sharing are highly dependent on whether

[8] The possibility of sending unverifiable cheap talk messages could have been included in this case, but this would not change the results.

[9] Even when two countries are close allies, it does not necessarily imply that there is a high level of mutual trust between their intelligence agencies. For example, the "Curveball" case related to the 2003 invasion of Iraq illustrates that there was a large level of distrust between the intelligence agencies of Germany and the US (Drogin 2007).

verifiable sharing is possible or not. If verifiable sharing is possible, full truthful revelation can happen in equilibrium. In the cheap talk case, full intelligence sharing will not happen and thus there will be private information in the subsequent counterterrorism investment game. The analysis of the two cases is below.

4.1 Cheap Talk Communication

Suppose there exists a (Perfect Bayesian) equilibrium with full revelation of each country's signal. With full revelation, the investment game in the second stage is equivalent to the benchmark game studied earlier. For simplicity, assume that each country always sends the true message, i.e., $m_i(d_i) = d_i$ for both i.[10] Consider country 1 and pick a realized signal d' such that $E[D|d_1 = d'_1, d_2 = 0] < 2C$. Thus, in equilibrium, the range of d_2's such that country 2 does not invest is higher when $d_1 = d'$ than when $d_1 = d'' > d'$ (I assume that the countries play the efficient benchmark equilibrium in stage two). Now consider the decision of country 1 in the message stage when it has received the signal $d_1 = d''$. In equilibrium, country 1 sends the true message $m_1 = d''$. However, suppose the country deviates and instead sends the message $m_1 = d'$. Then, in the investment stage, country 2 will believe that country 1 has received the signal d' and will therefore invest for a strictly smaller range of realized d_2's. And since it is always better for country 1 that country 2 does not invest, this means that it is better for country 1 to send the false message $m_1 = d'$ rather than the true message $m_1 = d''$. In other words, a country can, by pretending to have received a smaller signal than it really has, make the other country less likely to invest, which results in a lower sum of damages and costs for the country under consideration. This is why there will not be full revelation in equilibrium.

The argument above clearly illustrates the barriers to intelligence sharing when information cannot be shared verifiably, either because of the nature of the intelligence or the lack of trust and common interests between targeted countries. Each country always has an incentive to make the other country believe that it has low expectations about the capabilities of the terrorists because this will make the other country less likely to invest and thus more likely to become the target of the terrorists.

Finally, it is easy to see that there exist equilibria with no revelation of information, such that the second stage investment game is equivalent to the national intelligence game. Suppose each country never sends a message and then, in the investment stage, plays the equilibrium strategy with the highest cutoff signal from the national intelligence game. Then we just have to check that it is never profitable for a country to actually send some message in stage one. To ensure that this is the case, we can just choose the out of equilibrium belief of each country after seeing some message to be that the sender country has received the highest possible signal $\bar{d}$ (and is thus maxi-

[10]This is of course only one possible equilibrium with full revelation, but the argument below holds generally.

mally likely to invest). Then it is obviously better to send no message, which reveals no information to the other country.

4.2 Verifiable Sharing

When messages are verifiable, we get the opposite results of the ones from the cheap talk case. There exists an equilibrium with full revelation such that the outcome will simply be the benchmark outcome. And, on the other hand, there does not exist an equilibrium where no information is revealed.

I first show that there exists an equilibrium with full revelation. Let $m_i(d_i) = d_i$. Then we just have to check that a country will never find it optimal not to send a message rather than verifiably share its signal. This of course depends on the belief of the other country if it receives no message, which is an out of equilibrium belief. Let this belief be that the signal of the country sending no message is the highest possible signal $\bar{d}$. With this out of equilibrium belief, neither country can ever profitably deviate from revealing its signal, because such a deviation will maximize the probability that the other country invests. So when intelligence can be shared verifiably, it is possible to reach the benchmark outcome even though intelligence is gathered at the national level.

To see why no information revelation is not possible in equilibrium, assume that we have such an equilibrium. Then the outcome will be identical to the outcome of the national intelligence game. But this means that if a country has received a very low signal, for example $d_i = 0$, then it has an incentive to deviate from sending no message to verifiable sharing of that signal. Because this will make the other country invest for a strictly smaller range of d_j's, which makes country i better off.

5 Conclusion

I have explored the implications of introducing national intelligence gathering in a game theoretic model of domestic counterterrorism provision where two countries are facing a common transnational terrorist threat. The analysis revealed that the private information environment following from national intelligence collection often makes the overprovision of domestic counterterrorism worse than in the common intelligence benchmark. Loosely speaking, unless the ex ante expectation about the capabilities of the terrorists is quite low, there is more overprovision in the national intelligence case. So this is clearly the relevant implication of the model in times of a substantial and well known general threat from transnational terrorist groups, for example the period since the existence, intentions, and general capabilities of Al-Qaeda and related groups became known in the 1990s. On the other hand, in times where the ex ante belief is that transnational terrorist threats are very minor, it can be that national gathering of intelligence leads to more efficient counterterrorism provision.

I also studied the possibilities for sharing of nationally gathered intelligence. It was shown that full sharing of intelligence cannot happen in equilibrium if the only possibility for communication between the countries is through cheap talk messages. On the other hand, if verifiable sharing is a viable option for each country then it is possible to reach the common intelligence benchmark by intelligence sharing.

References

Arce M., Daniel G. and Todd Sandler. 2005. “Counterterrorism: A Game-Theoretic Analysis,” *Journal of Conflict Resolution* 49(2): 183-200.

Bueno de Mesquita, Ethan. 2007. “Politics and the Suboptimal Provision of Counterterror,” *International Organization* 61(1): 9–36.

Drogin, Bob. 2007. *Curveball: Spies, Lies, and the Con Man Who Caused a War.* Random House.

Enders, Walter and Todd Sandler. 1993. “The Effectiveness of Anti-Terrorism Policies: Vector-Autoregression-Intervention Analysis.” *American Political Science Review* 87(4): 829–844.

Enders, Walter and Todd Sandler. 2004. "What Do We Know About the Substitution Effect in Transnational Terrorism?" In *Researching Terrorism Trends, Achievements, Failures*, ed. Andrew Silke and G. Ilardi. Frank Cass.

Enders, Walter and Todd Sandler. 2006. "Distribution of Transnational Terrorism among Countries by Income Classes and Geography after 9/11." *International Studies Quarterly* 50(2): 367–393.

Powell, Robert. 2007. “Allocating Defensive Resources with Private Information about Vulnerability.” *American Political Science Review* 101(4): 799–809.

Rosendorff, Peter and Todd Sandler. 2004. “Too Much of a Good Thing? The Proactive Response Dilemma.” *Journal of Conflict Resolution* 48(4): 657–671.

Sandler, Todd and Walter Enders. 2004. "An Economic Perspective on Transnational Terrorism." *European Journal of Political Economy* 20(2): 301-316.

Sandler, Todd and Harvey E. Lapan. 1988. “The Calculus of Dissent: An Analysis of Terrorists’ Choice of Targets.” *Synthese* 76(2): 245–261.

Sandler, Todd and Kevin Siqueira. 2006. “Global Terrorism: Deterrence versus Preemption.” *Canadian Journal of Economics* 39(4): 1370–1387.

Walsh, James Igoe. 2006. "Institutions Are Not Enough: Intelligence Sharing in the European Union." *Journal of Common Market Studies* 44(3): 625-643.

Walsh, James Igoe. 2009. *The International Politics of Intelligence Sharing.* Columbia University Press.

Appendix

Claims made immediately after equation (4).

1. There exists an equilibrium $x > 0$.

   Proof. Let $G : [0, \bar{d}) \to \mathbb{R}$ be the function defined by

   $G(d) = E[D|d_i = d] + (1-p)(1-F(d|d))E[D|d_i = d, d_j > d] - 2C$ for all $d \in [0, \bar{d})$.

   From our assumptions it follows that G is continuous and that

   $$G(0) = E[D|d_i = 0] + (1-p)E[D|d_i = 0] - 2C < (2-p)\frac{2C}{2-p} - 2C < 0$$

   and

   $$G(d) > E[D|d_i = d] - 2C > 0 \text{ for } d \text{ sufficiently close to } \bar{d}.$$

   From these observations it follows that there must exist an $x \in (0, \bar{d})$ with $G(x) = 0$, which means that equation (4) is satisfied. □

2. There exists a highest cutoff equilibrium and this is the efficient equilibrium.

   Proof. By continuity of G (see claim 1 above), the set of equilibrium cutoff signals,$\{x|G(x) = 0\}$, is closed and bounded in $\mathbb{R}$. Thus there exists a highest cutoff equilibrium $x^H \in (0, \bar{d})$.

   To see that x^H is the efficient symmetric cutoff equilibrium, let $u_i(x_i, x_j)$ denote the ex ante expected utility of country i when it uses the cutoff strategy given by x_i and country j uses the cutoff strategy given by x_j. It is easy to see that country i is always better off when country j is less likely to invest, i.e., when x_j is higher. Thus, if $x \neq x^H$ is a symmetric cutoff equilibrium, then we have

   $$u_i(x, x) < u_i(x, x^H).$$

   Furthermore, since x^H is a symmetric equilibrium cutoff, we also have

   $$u_i(x, x^H) \leq u_i(x^H, x^H).$$

   Thus, each country i receives a higher ex ante expected utility in the x^H equilibrium than in any other symmetric cutoff equilibrium. □

The unique equilibrium in the example in Section 3.1.

Solve the second order equation (5) to get that the smallest root (and thus the equilibrium cutoff signal) is

$$x = \frac{(2-p) - \sqrt{(2-p)^2 - 3(1-p)(-2+p+4C)}}{3(1-p)}.$$

Claim made in the example in Section 3.1.
The smaller p is, the larger is the set of C's such that $E[D|x,x] < 2C$.

Proof. First note that $E[D|x,x] = x + x = 2x$. Thus $E[D|x,x] < 2C$ is equivalent to $x < C$. Therefore it suffices to show that $\frac{\partial x}{\partial p} > 0$ for all $p \in (0,1)$ and $C \in (\frac{2-p}{4}, \frac{3}{4})$. From the expression for x from above we get

$$\frac{\partial x}{\partial p} = 3\frac{1 + \frac{1}{2}(1-p)(13-8p-12C)A^{-\frac{1}{2}} - A^{\frac{1}{2}}}{9(1-p)^2},$$

where $A = (2-p)^2 - 3(1-p)(-2+p+4C)$. So it follows that $\frac{\partial x}{\partial p} > 0$ is equivalent to

$$A^{\frac{1}{2}} + \frac{1}{2}(1-p)(13-8p-12C) - A > 0.$$

By further calculations this is equivalent to

$$120(1-p)^2 C(\frac{5}{6} - C) > 9(1-p)^2,$$

which is easily seen to hold for all $p \in (0,1)$ and $C \in (\frac{2-p}{4}, \frac{3}{4})$. $\square$