
Politics and the Suboptimal Provision of Counterterror

Ethan Bueno de Mesquita

Abstract I present a model of interactions between voters, a government, and a terrorist organization. The model focuses on a previously unexplored conceptualization of counterterrorism as divided into tactic-specific observable and general unobservable tactics. When there is divergence between voters and government preferences, strategic substitution among different modes of attack by terrorists and agency problems between the voters and government create a situation in which the politically optimal counterterrorism strategy pursued by the government in response to electoral and institutional incentives is quite different from the security maximizing counterterrorism strategy. In particular, in response to electoral pressure, the government allocates resources to observable counterterror in excess of the social optimum. This problem is particularly severe when governments put great weight on rent-seeking or care less about counterterror than do voters and when terrorists have a large set of tactics from which to choose. Voters can decrease the magnitude of the agency problem by increasing the benefits of reelection by, for example, slackening requirements for nonsecurity related public goods.

In pursuit of security from terrorism, governments engage in a broad array of counterterror tactics. Since the attacks of 11 September 2001, the United States has employed methods including enhanced airport and airplane security, a tightening of border controls, attempts to freeze terrorist assets, bombings of terrorist training camps, covert operations to infiltrate terrorist organizations and funding networks, and the invasions of Afghanistan and Iraq, among others. Other countries facing serious terrorist threats use a similarly large set of counterterror tactics. For instance, Israel has engaged in border closings, curfews, targeted assassinations, military incursions, home demolitions, and the construction of a security barrier.

Different counterterror strategies offer different costs and benefits. As such, in trying to understand the politics of counterterrorism, it is often fruitful to disaggregate the general category of counterterrorism into stylized subcategories of tac-

I have received valuable comments and advice from Bruce Bueno de Mesquita, Randy Calvert, Martin Cripps, James Fearon, Amanda Friedenberg, Robert Powell, Matthew Stephenson, and Barbara Walter. I thank the Weidenbaum Center on the Economy, Government, and Public Policy at Washington University for financial support.

International Organization 61, Winter 2007, pp. 9–36

© 2007 by The IO Foundation.

DOI: 10.1017/S0020818307070087

tics that vary in their strategic implications. For instance, Sandler and Siqueira draw a distinction between proactive and defensive counterterrorism to demonstrate that proactive counterterrorism, which is a public good among countries, may be undersupplied.¹ Rosendorff and Sandler maintain this same distinction but, building on Sandler and Lapan's argument that one country's proactive counterterrorism efforts may impose negative externalities on other countries by causing terrorists to divert attacks toward foreign soft targets, they show that in some circumstances proactive measures may be oversupplied from the global perspective.² Bueno de Mesquita and Dickson consider a different division of counterterrorist tactics: indiscriminate tactics that are relatively inexpensive but impose negative economic externalities on a population of terrorist sympathizers, and discriminate tactics that are relatively expensive but impose fewer externalities.³ Building on models of the effects of counterterrorism on mobilization, they explore how the choice of different types of counterterrorism tactics can affect the level of terrorist violence by altering the opportunity costs of mobilization and perceptions of government type.⁴ Powell analyzes the optimal division of government resources between site-specific defenses and general counterterrorism expenditures (for example, intelligence) under a variety of assumptions regarding the nature of the terrorist threat.⁵

In the spirit of those papers, this article analyzes the politics of counterterrorism by considering a new stylized division of the set of counterterrorism strategies. I model counterterrorism tactics as differentiated on two dimensions: those that are both tactic specific and publicly observable and those that are general and not publicly observable. Tactic-specific observable counterterrorism includes a wide array of security measures—airport security, securing of ports—that guard against specific types of attacks and are readily verified by voters and terrorists alike. General, unobservable counterterrorism is made up of those strategies that governments undertake in secret, such that they might never be observed by voters and such that it is not possible for terrorists to avoid their consequences by switching tactics. These strategies might include the infiltration of terrorist cells by intelligence agents, secret operations to freeze funding networks or disrupt the weapons supply chain, covert bombings of terrorist training camps, and other such methods.

Observable, tactic-specific counterterrorism has an important liability—the terrorists, on observing an increase in a particular government counterterrorism program, can switch tactics, pursuing attacks less affected by the government's efforts. General counterterrorism, not directed at a specific tactic, does not suffer from this liability. A variety of empirical studies find that tactic-specific counterterrorism

1. Sandler and Siqueira 2006.

2. See Rosendorff and Sandler 2004; and Sandler and Lapan 1988.

3. Bueno de Mesquita and Dickson 2007.

4. For models of the effect of counterterrorism on mobilization, see Lichbach 1987; de Figueiredo and Weingast 2001; Rosendorff and Sandler 2004; and Bueno de Mesquita 2005.

5. Powell Forthcoming.

ism generates such substitution effects.⁶ For instance, Enders and Sandler demonstrate that when the United States installed metal detectors in airports in the 1970s, hijackings decreased but other forms of terrorism increased.⁷

From a political perspective, however, unobservable tactics also have an important disadvantage—the voters do not observe them. Politicians engage in security spending for at least two reasons: in order to increase security and to serve their political interests. Unobservable operations have the liability, then, that politicians receive no credit for them from voters, other than the credit the politicians receive for the absence of attacks.

To explore the consequences of this distinction between tactic-specific observable and general unobservable counterterrorism, I present a model of the interactions between a voter, a government, and a terrorist organization. The terrorists have resources to devote to violence. They divide these resources amongst a variety of tactics in response to the counterterrorism spending of the government. The government, for its part, has a counterterrorism budget that it must decide how to allocate. The government can spend resources on a variety of publicly observable counterterrorism tactics, each designed to thwart a particular type of terrorist attack (for example, airport security). The government can also devote resources to general unobservable operations designed to disrupt the terrorists' access to resources such as personnel, financing, weapons, and information. Finally, the government can divert resources supposedly allocated to unobservable counterterrorism to other, non-counterterrorism-related projects. The voter observes three things: how the government divides the counterterrorism budget between observable and unobservable spending, the amount of spending devoted to each of the government's observable counterterrorism programs, and the success or failure of the terrorist campaign. The voter does not observe the amount of the unobservable budget that is devoted to counterterrorism. The voter then decides whether or not to reelect the government.

The model yields several implications. First, despite the inefficiency of tactic-specific publicly observable counterterrorism, unless the government's and voter's preferences are perfectly aligned (that is, they weight the importance of counterterrorism versus other priorities equally and the government has no preference for rent seeking), the government will always allocate resources to observable counterterrorism in excess of the social optimum. This inefficient distribution of counterterrorism resources is because, not in spite, of the voter. Voters, concerned about government diversion of unobservable spending away from counterterrorism, force the government, through electoral incentives, to overspend on observable counterterrorism. Doing otherwise would leave the voters even worse off.

Second, as the number of potential terrorist tactics increases, tactic-specific spending becomes less effective relative to general counterterrorism spending. This is because

6. See Enders and Sandler 1993 and 2002; Enders, Sandler, and Cauley 1990; and Im, Cauley, and Sandler 1987.

7. Enders and Sandler 1993.

substitution effects allow the terrorists to evade tactic-specific counterterror, forcing the government to spread its observable resources thin.

Third, unobservable counterterror spending is undertaken by the government to secure reelection. Hence, if voters can increase the benefits of holding office (for example, by slackening requirements for non-security-related public goods provision), they can decrease the moral hazard problem, which increases government spending on covert operations and allows the voters to decrease their demand for observable counterterror.

Definitional Issues

There is substantial discussion as to the proper definition of terrorism.⁸ For the purposes of this article, I do not need to take a firm position on many of the specific definitional issues (for example, must terrorism be an illegal activity, can terrorism be carried out by the state, is terrorism inherently targeted against civilians, and so on). The important features of the terrorism I study in this model are that it is carried out by an organization that is neither affiliated with nor supported by the government, that the voters find the presence of terrorism costly, and that the terrorist organization be clandestine, so that counterterror tactics beyond simple military conflict are needed to effectively curtail violence. Since these conditions are consistent with most of the definitions of terrorism in the literature, I do not take a strong position on definitional issues beyond these requirements.

It is worth noting that this definitional restriction is consistent with phenomena other than terrorism. In particular, organized crime meets the criteria discussed above. Thus, to the extent that governments and voters also have divergent preferences over expenditure on crime prevention versus rent seeking or other government priorities, the analysis below applies to anticrime policy in much the same way as it does to counterterrorism policy.

Connection to the Public Debate over Counterterror

Two key substantive assumptions underlie this analysis. First, I assume that government and voter preferences may partially diverge, in the sense that the government may engage in rent seeking—using counterterrorism resources for purposes other than enhancing security. Second, I argue that there is an important distinction, both strategically and politically, between observable and unobservable counterterrorism.

8. For several important contributions to this debate, see Laqueur 1977; Gibbs 1989; Crenshaw 1995; and Hoffman 1998.

Recent debates over the proper geographical division of counterterrorism resources demonstrate the problem of rent seeking in counterterrorism policy. The argument has been made that, for reasons of political patronage, counterterrorism resources have been diverted to states that face little or no terrorist threat. For instance, the *New York Times* reports that in 2004, Wyoming received a counterterrorism allocation of \$38 per person, seven times the per capita spending in New York, a decision which New York City Mayor Michael Bloomberg referred to as, “pork barrel politics at its worst.”⁹

Commentators have also expressed concern about the government focusing on publicly observable programs due to political expediency. In a recent *Newsweek* interview, security expert Bruce Scheier articulated this critique, describing certain counterterrorism expenditures as “an enormous waste of money” and explaining: “Politicians tend to prefer security countermeasures that are very visible, to make it look like they’re doing something. So they will tend to pick things that are visible even if they are less effective.”¹⁰

Perhaps the most comprehensive analysis of counterterrorism issues in the United States, from a policy perspective, is provided by the recent *9/11 Commission Report*. This report, which recommends a variety of steps toward improving counterterrorism, highlights the salience of the distinction between observable and unobservable counterterrorism tactics explored in this article.

Not surprisingly, many of the Commission’s recommendations involve observable tactics. For instance, with respect to airport security, the report suggests: “The TSA [Transportation Safety Administration] and the Congress must give priority attention to improving the ability of screening checkpoints to detect explosives on passengers. As a start, each individual selected for special screening should be screened for explosives.”¹¹

The Commission also discusses strategies for strengthening border control and securing nuclear materials, arguing: “The U.S. border security system should be integrated into a larger network of screening points that includes our transportation system and access to vital facilities, such as nuclear reactors.”¹²

In addition to recommendations whose implementation would require concrete, observable policy changes, the 9/11 Commission also suggests strategies that call for more secretive, unobservable actions. For instance, the Commission rejected as ineffective observable attempts to freeze terrorist resources:

After 9/11, the United States took aggressive actions to designate terrorist financiers and freeze their money. . . . These actions appeared to have little effect. . . . Worldwide asset freezes have not been adequately enforced and have

9. *New York Times*, 1 June 2004, B4.

10. *Newsweek*, 17 March 2004. Available at (<http://msnbc.msn.com/id/4549661/>). Accessed 11 October 2006.

11. “The 9/11 Commission Report”, 393. Available at (<http://www.9-11commission.gov/report/911Report.pdf>). Accessed 11 October 2006.

12. *Ibid.*, 387.

been easily circumvented, often within weeks, by simple methods. But trying to starve the terrorists of money is like trying to catch one kind of fish by draining the ocean. A better strategy has evolved since those early months, as the government learned more about how al Qaeda raises, moves, and spends money.¹³

Evidently, the sort of public strategies to block terrorist financing that were widely publicized following the 9/11 attacks were ineffective because terrorists, aware of the government's actions, were able to avoid them. A better technique, the report suggests, is to use secret intelligence about terrorist funding to unobservably track and undermine terrorist organizations themselves. As the report states: "Vigorous efforts to track terrorist financing must remain front and center in U.S. counterterrorism efforts. The government has recognized that information about terrorist money helps us to understand their networks, search them out, and disrupt their operations."¹⁴

Clearly, implementing these types of recommendations requires operations that remain largely out of the public eye. Moreover, for such operations to be maximally effective, the terrorists must remain unaware of them until it is too late to avoid the consequences.

The government, of course, also engages in secret, unobservable counterterrorism actions, such as infiltrating terrorist cells. Indeed, the rise in the salience of counterterrorism has led to an increase in expenditures on such covert tactics. The U.S. budget request for the Special Operations Command increased 34 percent from fiscal year 2003 to fiscal year 2004, to \$6.7 billion from \$5 billion.¹⁵ Not surprisingly given the large sums of government money at stake, there have been struggles over tactical and fiscal control of covert operations within the executive branch. Both Kibbe and Hersh report that, with the support of the 9/11 Commission, the Defense Department moved to wrest control of such operations away from the Central Intelligence Agency (CIA).¹⁶ The existence of such infighting not only highlights the importance of unobservable tactics, but also points to the fact that control over these methods provides decision makers with the ability to advance their particular agendas and engage in rent seeking.

The Model

Consider a game between three players: the government (G), a terrorist organization (T), and a representative voter (V). The government chooses how to allocate its counterterror budget. Then the terrorists choose what types of attacks to invest

13. Ibid., 382.

14. Ibid.

15. Kibbe 2004.

16. See *ibid.*; and *The New Yorker*, 24 and 31 January 2004.

their resources in. Finally, the voter observes spending decisions and terror outcomes and chooses whether to reelect the government. The timing is chosen to reflect the terrorists' and voter's ability to respond strategically to counterterror tactics.

Formally, the model is closely related to the multitask model of Holmström and Milgrom.¹⁷ However, there are important differences. First, in this model the principal (the voter) uses electoral threats, rather than wages, to provide the agent (the government) with incentives. Since electoral incentives are discontinuous, voting offers less leverage to the principal than do wages. Consequently, different types of behavior by the agent can be induced by the two mechanisms. Second, in the standard economic model, the principal simply seeks to induce the agent to produce some product through effort. Here, the government is "producing" counterterror. The production of counterterror is itself an interesting problem, since the terrorists respond strategically to counterterror outputs.

Actions

The Government. The government has a counterterror budget A and has to allocate it between three types of spending: tactic-specific publicly observable counterterror spending (a_{pub}), general unobservable counterterror spending (a_0), and unobservable spending for non-counterterror purposes ($A - a_{pub} - a_0$). The observable spending can be divided among n different types of counterterror, corresponding to different types of terrorist tactics. The unobservable spending is general, directed toward disrupting terrorist financing and infrastructure. Thus, a strategy for the government is an $n + 1$ -tuple, $(a_0, a_1, \dots, a_n) \in [0, A]^{n+1}$, subject to the budget constraint (BC_G)

$$\sum_{i=0}^n a_i \leq A. \tag{BC_G}$$

Notice that $a_{pub} = \sum_{i=1}^n a_i$.

The counterterror technology is as follows. The unobservable spending (a_0) decreases the total resources to which the terrorists have access. It does so according to a function $g(a_0): [0, A] \rightarrow [0, 1]$, where $g(\cdot)$ is increasing and concave in a_0 . Public spending directed at a particular type of attack, i , decreases the effectiveness of terrorist resources directed toward that type of attack in a manner to be described later.

It is worth noting that this formalization conflates two issues: observable versus unobservable counterterror and tactic-specific versus general counterterror. That

17. Holmström and Milgrom 1991.

is, observable counterterror is taken to always be tactic specific, while unobservable counterterror is taken to always be general. While this is descriptive of many policies, it need not be the case. For instance, border security is a general, observable counterterror policy. However, as will become clear, some of the results will be driven primarily by observability while others will be driven primarily by the distinction between general and tactic-specific methods. Throughout the analysis I highlight which aspect of the formalization is at work.

The Terrorists. The terrorists have resources, $(1 - g(a_0))R$, which they divide among the n possible types of attacks. I label a generic type of attack t_i . To avoid notational clutter, define a new function $R(\cdot) \equiv (1 - g(\cdot))R$, where the assumptions on $g(\cdot)$ imply that $R(\cdot)$ is decreasing and convex in a_0 . Label $R(0) = R$. A strategy for the terrorists is an n -tuple $(t_1, t_2, \dots, t_n) \in [0, R]^n$, subject to the budget constraint (BC_T)

$$\sum_{i=1}^n t_i \leq R(a_0). \quad (BC_T)$$

The Voter. The voter chooses whether or not to reelect the government based on observing the division of the budget between observable and unobservable spending, the allocation of the observable spending, and the success or failure of the terrorist campaign.

The Technology of Terrorism

The marginal impact of a terrorist dollar spent on each type of tactic is a function of two factors. First, each terrorist tactic, i , is characterized by a tactic-specific parameter $\theta_i \in R$. This parameter can be thought of as the underlying effectiveness of a tactic. That is, holding counterterror constant, a dollar spent on certain types of attacks may simply be more effective at achieving a terrorist organization's goals than a dollar spent on another type of attack. This *inherent effectiveness* of each tactic $(1, 2, \dots, n)$ is captured by the vector of parameters $(\theta_1, \theta_2, \dots, \theta_n)$. Second, tactic-specific government counterterror (a_i) can potentially dampen the effectiveness of even the most inherently effective tactics. Thus, the total effectiveness of a tactic, i , is given by a function $f(a_i, \theta_i) : [0, A] \times \mathbb{R} \rightarrow \mathbb{R}^+$, where $f(\cdot, \cdot)$ is decreasing and concave in the first argument (a_i) and increasing in the second argument (θ_i). I will refer to the total effectiveness of a tactic ($f(a_i, \theta_i)$), induced by both its inherent effectiveness and the level of counterterror spending on that tactic, as its *induced effectiveness*.

The total impact of the terror campaign is taken to be additively separable across types of attack and is given by $\tau = \sum_{i=1}^n f(a_i, \theta_i)t_i$. The assumption of additive separability implies that there is no interaction between types of terrorist attacks. This is clearly a simplification, since many types of attacks may be complements

with one another. For instance, the anthrax attacks that were perpetrated through the mail in the United States in 2001 probably would not have had as large an impact as they did, had they not followed closely on the heels of the 9/11 attacks. I make this assumption because such effects are not the primary focus of this article and it simplifies the analysis by generating pure substitution by terrorists in response to counterterrorism. Enders and Sandler provide a thorough discussion of substitution versus complementarity in terrorist tactics and an empirical estimation of their relative import in terrorist decision making.¹⁸

The terror campaign can result in one of two outcomes: success for the terrorists (\bar{S}) or failure for the terrorists (\underline{S}). The probability that the terrorists succeed in their goals is $\pi(\tau): \mathbb{R}^+ \rightarrow [0, 1]$, where $\pi(\cdot)$ is increasing and concave in τ —the higher the impact of the terrorist campaign, the more likely the terrorists are to succeed.¹⁹

Payoffs

The Government. The government has preferences over reelection, the level of terrorism, and spending that are common knowledge and are given by

$$U_G = (1 - \alpha)u_G(A - a_{pub} - a_0) + \alpha(\pi(\tau)v(\bar{S}) + (1 - \pi)v(\underline{S})) \\ + \Pr(\text{reelection} | a_{pub}, a_0, S)B.$$

B is the benefit associated with reelection. The function u_G represents the payoff from resources redirected away from counterterror. I assume that u_G is increasing, concave, and satisfies $\lim_{x \rightarrow 0} u'_G(x) = \infty$. The function $v(S)$ represents the social costs of terrorism. Without loss of generality, I assume that $v(\bar{S}) = 0$, which implies $v(\underline{S}) > 0$ —it is better for society (that is, the voter) for the terrorists to fail. The parameter α measures how closely the government's preferences align with the voter's welfare. When $\alpha = 1$, the government has identical preferences to the voter and there is no agency problem. If $\alpha = 0$, then the government's preferences completely diverge from the voter's—the government cares about counterterrorism only insofar as it affects reelection. Intermediate values of α reflect a government that, as with the voter, cares about security for its own sake, but has priorities not perfectly aligned with the voter.

The idea behind this specification is that, absent reelection concerns, the government might prefer to spend some of its counterterrorism budget on things other than security measures. Depending on the government in question, these prefer-

18. Enders and Sandler 1993.

19. The definition of "success" is left nonspecific, since terrorism is used for different purposes in different contexts (for example, extracting concessions, gaining media attention, causing economic disturbance, recruiting supporters, and so on).

ences could take a variety of forms. Some governments might simply be corrupt, misappropriating money either by stealing it or by turning funds supposedly directed toward security into pork or patronage. Other governments might not engage in graft but rather view the importance of counterterror, relative to other programs, differently than the voters do. Such governments might be tempted to redirect some of the unobservable counterterror budget toward these other priorities. The key assumption for the model is that voter and government preferences over counterterrorism spending differ. In the absence of such a difference, as I will show, the government will pursue the voter's optimal counterterrorism policy. I will refer, generally, to the government diverting or redirecting resources to indicate the allocation of resources intended by the voter for counterterrorism to some other purpose.

The Terrorists. The terrorists have preferences over the success of the terror campaign given by

$$U_T = \pi(\tau)u_T(\bar{S}) + (1 - \pi(\tau))u_T(\underline{S}).$$

Without loss of generality, I assume that $u_T(\underline{S}) = 0$, which implies that $u_T(\bar{S}) > 0$.

The Voter. The voter has preferences over the social costs of the terror campaign given by

$$U_V = \pi(\tau)v(\bar{S}) + (1 - \pi(\tau))v(\underline{S}).$$

Recall that $v(\bar{S}) = 0$ and $v(\underline{S}) > 0$.

Equilibrium

Since there is hidden action by the government, the appropriate solution concept is Perfect Bayesian equilibrium (PBE). Notice that at the point where the voter votes, the payoffs are invariant to the voter's action. Thus, all voting rules are sequentially rational. As such, I impose an equilibrium selection criterion, restricting attention to voting rules that are "optimal" in the sense that they induce behavior from the government that the voter likes. More formally, I assume that the voter chooses the voting rule that maximizes the voter's *ex ante* expected utility.²⁰ Practically, this is akin to assuming that the voter votes in a retrospectively opti-

20. This is the standard approach in models of retrospective voting. See, for example, Barro 1973; Ferejohn 1986; and Austen-Smith and Banks 1989. Bueno de Mesquita and Friedenbergh 2005 formalize retrospective voting as an equilibrium selection criterion.

mal manner. Because the voter is indifferent at the point of reelection, such a strategy is credible.

I further restrict attention to those PBE in which the terrorists and voter have “reasonable” beliefs about unobservable effort. In particular, off the equilibrium path, they believe that the government chooses its utility-maximizing level of unobservable counterterrorism. This selection criterion rules out equilibria of the following form. The voter believes (unreasonably) that, for what would be the optimal choice of observable effort, the government will engage in a level of unobservable effort less than what the government would actually engage in. Consequently, the voter gives the government electoral incentives to choose some other level of observable effort. The voter’s beliefs are consistent with equilibrium, because the relevant information set is never reached. However, the equilibrium is only sustainable because the voter’s unreasonable beliefs deter the government from taking an action it, and the voter, would have liked, if the voter had reasonable beliefs. Given this assumption, I suppress notation on beliefs, since beliefs will be identical to optimizing behavior, on and off the path.

The Terrorists

The terrorists allocate their resources to maximize the impact of the terror campaign. Formally, this means that they solve the following problem:

$$\max_{(t_1, \dots, t_n)} \pi \left(\sum_{i=1}^n f(a_i, \theta_i) t_i \right) u_T(\bar{S}),$$

subject to (BC_T) . Since the budget constraint will obviously bind at the optimum, the optimal allocation at an interior solution is characterized by the budget constraint and the following Kuhn-Tucker conditions:

$$\pi'(\tau) u_T(\bar{S}) f(a_i, \theta_i) = \lambda,$$

for all t_i , where λ is the Lagrangian multiplier on the budget constraint. This implies that if more than one type of attack is pursued, then all types of attack that receive positive budget shares must have the same equilibrium level of induced effectiveness $(f(a_i, \theta_i))$. Furthermore, only those types of attacks with the highest induced effectiveness will receive positive budget shares. Define this set of highly effective tactics as $\bar{\mathbf{t}} = \{t_i | f(a_i, \theta_i) \geq f(a_j, \theta_j), \text{ for all } j\}$. Now, the following result can be stated.

Lemma 1. *The terrorists choose a positive level of t_i only if $t_i \in \bar{\mathbf{t}}$. Moreover, the terrorists are indifferent between all $t_i \in \bar{\mathbf{t}}$.*

All proofs are in the Appendix. The intuition is as follows. Different types of terror attacks have different levels of induced effectiveness. The terrorists are free to substitute among these tactics. Consistent with the findings of Enders and Sandler, the terrorists pursue those types of attacks that offer the highest level of impact.²¹ It is important to note, in interpreting this result, that it is with regard to *induced effectiveness*. That is, the lemma does not imply that terrorists will only invest in the most inherently effective tactics. Rather, it implies that, in equilibrium, the terrorists will invest in those tactics that will have the largest marginal impact on the success of the terror campaign, given both the inherent effectiveness of the tactics and the government’s counterterror strategy.

The Government

The Allocation of Observable Spending. The first question to ask about the government is, given a division between observable and unobservable spending, how will it allocate the publicly observable spending? Since the voter knows how observable resources are allocated, the voter will choose a reelection rule that induces the optimal allocation from the government. What, then, is this optimal allocation?

The voter wants the government to minimize the total impact of terrorism (τ). As shown in Lemma 1, the terrorists will only invest in the types of attacks that have the highest induced impact. Thus, the optimal allocation of observable spending should minimize the effectiveness of the chosen forms of terror. The government’s problem, then, is

$$\min_{(a_1, \dots, a_n)} \max_{i=1, \dots, n} f(a_i, \theta_i).$$

The government solves this problem by choosing (a_1, \dots, a_n) to minimize the impact of the most inherently effective tactics. This means that the government will disperse its resources in such a way that all tactics either have the same level of induced effectiveness or receive no budget share. Thus, for any allocation a_{pub} and vector of parameter values $\theta = (\theta_1, \theta_2, \dots, \theta_n)$, the optimal allocation of observable counterterror spending will result in all tactics, i and j , that receive a positive budget share satisfying $f(a_i, \theta_i) = f(a_j, \theta_j) \equiv \hat{f}(a_{pub}, \theta)$, and all other tactics receiving a budget share of 0. Moreover, because the terrorists invest in the most effective tactics, it is clear that those tactics, k , that the government does not invest counterterror resources in will satisfy $f(0, \theta_k) < \hat{f}(a_{pub}, \theta)$. Define a tactic i as *maximally effective with respect to a_{pub} and θ* if $f(0, \theta_i) > \hat{f}(a_{pub}, \theta)$.

Powell, who finds the same result regarding optimal counterterror expenditures in a model developed simultaneously and independently of this one, describes an

21. Enders and Sandler 1993.

algorithm that provides the intuition for the concept of maximal effectiveness defined above.²² The government spends counterterror resources on the most inherently effective tactic until that tactic has the same induced effectiveness as the next most inherently effective tactic. It then spends counterterror resources on these two tactics until they have the same induced effectiveness as the next most effective tactic. It continues this procedure until it runs out of money earmarked for observable counterterror. Those tactics that received positive budget shares under this method qualify as maximally effective.

The above argument gives rise to the following result.

Remark 1. The government chooses its observable counterterror allocation such that $f(a_i, \theta_i) = f(a_j, \theta_j) \equiv \hat{f}(a_{pub}, \bar{\theta})$, for all i, j that receive positive counterterror budget shares a_i . Moreover, the government devotes counterterror resources to, and only to, preventing maximally effective tactics. Finally, the terrorists devote an equal amount of resources to each type of maximally effective tactic.

Remark 1 implies that in equilibrium the set of tactics the terrorists will devote resources to ($\bar{\mathbf{i}}$) will be equal to the set of maximally effective tactics. Denote this set by $\bar{\mathbf{i}}^*(a_{pub}, \bar{\theta}) = \{t_i | f(0, \theta_i) > \hat{f}(a_{pub}, \bar{\theta})\}$ and let $|\bar{\mathbf{i}}^*(a_{pub}, \bar{\theta})|$ be the number of elements in this set. I will drop the functional notation when no confusion will be caused. It follows from Remark 2, that the total impact of terror is

$$\tau = \sum_{\{i | t_i \in \bar{\mathbf{i}}^*\}} f(a_i, \theta_i) t_i = \sum_{\{i | t_i \in \bar{\mathbf{i}}^*\}} \hat{f}(a_{pub}, \bar{\theta}) \frac{R(a_0)}{|\bar{\mathbf{i}}^*|} = \hat{f}(a_{pub}, \bar{\theta}) R(a_0).$$

That is, the optimal allocation of observable counterterrorism resources is such that the total impact of terrorism is simply a function of the parameters (θ, R) and the division of the budget between observable and unobservable spending.

Another implication of the remark is that, since in equilibrium all terror tactics that receive positive budget shares must have the same effectiveness, the government will especially focus its counterterror efforts on inherently effective tactics.

Remark 2. The level of counterterror spending on a particular type of attack (a_i) is weakly increasing in the inherent effectiveness of that tactic (θ_i).

These results also have implications for patterns of terrorist violence and counterterrorism in response to major innovations in terrorist tactics. From time to time terrorists develop wholly new tactics (for example, skyjackings in the 1960s, suicide bombings in the 1980s). When major innovation of this sort occurs, the new tactic may be so effective as to temporarily supplant all other tactics. The model

22. Powell 2005.

suggests, in this case, that the terrorists will direct all resources toward the new tactic until government counterterror technology catches up and decreases the new tactic's inherent effectiveness enough to once again make other tactics maximally effective and, therefore, attractive to the terrorists.

The Allocation of Unobservable Spending. The next question is, given a level of observable and unobservable spending, how much of the unobservable spending will the government spend on counterterror and how much will it divert to other priorities? The government's objective function is

$$B[\text{Pr}(\text{reelect}|\bar{S})\pi(\hat{f}(a_{pub}, \bar{\theta})R(a_0)) + \text{Pr}(\text{reelect}|\underline{S})(1 - \pi(\hat{f}(a_{pub}, \bar{\theta})R(a_0)))] + \alpha(1 - \pi(\hat{f}(a_{pub}, \bar{\theta})R(a_0)))v(\underline{S}) + (1 - \alpha)u_G(A - a_{pub} - a_0).$$

The first-order condition, with respect to unobservable spending (a_0), is

$$\pi'[\hat{f}(a_{pub}, \bar{\theta})R(a_0^*)]\hat{f}(a_{pub}, \bar{\theta})R'(a_0^*)(B[\text{Pr}(\text{reelect}|\bar{S}) - \text{Pr}(\text{reelect}|\underline{S})] - \alpha v(\underline{S})) = (1 - \alpha)u'_G(A - a_{pub} - a_0^*).$$

First note that $\lim_{x \rightarrow 0} u'_G(x) = \infty$ implies that there is an interior solution. Furthermore, the first-order condition makes clear that the level of spending on unobservable counterterror, holding fixed observable counterterror, is strictly increasing in $\text{Pr}(\text{reelect}|\underline{S})$ and strictly decreasing in $\text{Pr}(\text{reelect}|\bar{S})$. That is, the government's unobservable counterterror expenditures are maximized when it is rewarded for policy successes and punished for policy failures. Thus, the voter's optimal reelection rule assigns $\text{Pr}(\text{reelect}|\underline{S}) = 1$ and $\text{Pr}(\text{reelect}|\bar{S}) = 0$. Substituting these in, the first-order condition can be rewritten

$$-\pi'[\hat{f}(a_{pub}, \bar{\theta})R(a_0^*)]\hat{f}(a_{pub}, \bar{\theta})R'(a_0^*)(B + \alpha v(\underline{S})) = (1 - \alpha)u'_G(A - a_{pub} - a_0^*). \tag{1}$$

The intuition behind the first-order condition is clear. The left-hand side represents the marginal benefit of unobservable spending. Increasing the level of unobservable spending directed toward counterterror decreases the resources to which the terrorists have access, which diminishes the impact of terrorism. This benefits the government by increasing both the probability of reelection and social welfare. The marginal costs are on the right-hand side. Money spent on unobservable counterterror cannot be appropriated for other priorities. The optimal amount of spending on unobservable operations balances these concerns. Given this, as long as there is some difference between voter and government preferences, the gov-

ernment will divert some counterterrorism resources. This intuition is formalized in the following remark.

Remark 3. As long as $a_{pub} < A$ and $\alpha < 1$, the government will divert some of the counterterrorism budget.

As mentioned above, increasing counterterrorism spending decreases the risk of successful terrorism, which benefits the government in two ways. First, it increases the probability of reelection. Second, to the extent that the government shares the voter’s preferences for security (α), it also derives a direct benefit from preventing attacks. Thus, as the congruence between government and voter preferences increases, independent of electoral pressures, the government has an increased incentive to invest in unobservable counterterrorism. This intuition is formalized in the following remark.

Remark 4. As α increases, government spending on unobservable counterterrorism (a_0^*) increases.

The Division Between Observable and Unobservable Spending. The voter can force the government to choose a division between observable and unobservable spending by conditioning the voting rule on the amount of observable spending. The government will choose the division that the voter demands as long as doing so yields a higher expected payoff than forgoing reelection and choosing its most preferred allocation. This gives rise to the following incentive compatibility constraint (IC_G):

$$(1 - \pi[\hat{f}(a_{pub}, \bar{\theta})R(a_0^*(a_{pub}))])(B + \alpha v(\underline{S})) + (1 - \alpha)u_G(A - a_{pub} - a_0^*(a_{pub}))$$

$$\geq \max_{\tilde{a}_{pub}, \tilde{a}_0} \{(1 - \pi[\hat{f}(\tilde{a}_{pub}, \bar{\theta})R(\tilde{a}_0)])\alpha v(\underline{S}) + (1 - \alpha)u_G(A - \tilde{a}_{pub} - \tilde{a}_0)\}(IC_G)$$

Let \tilde{a}_{pub} be the largest a_{pub} that satisfies (IC_G).

Given the government’s optimal allocations of observable and unobservable spending, and the incentive compatibility constraint, how much of the budget does the voter want to devote to each activity?

The voter’s problem can be written as follows:

$$\max_{a_{pub}} (1 - \pi[\hat{f}(a_{pub}, \bar{\theta})R(a_0^*(a_{pub}))])v(\underline{S}),$$

subject to (IC_G).

If the incentive compatibility constraint does not bind, then the optimum is given by the first-order condition:

$$-\pi' [\hat{f}(a'_{pub}, \bar{\theta})R(a^*_0(a'_{pub}))] \times \left[\hat{f}_a(a'_{pub}, \bar{\theta})R(a^*_0(a'_{pub})) + \hat{f}(a'_{pub}, \bar{\theta})R'(a^*_0(a'_{pub})) \frac{\partial a^*_0}{\partial a_{pub}} \right] v(\underline{\mathcal{L}}) = 0,$$

where \hat{f}_a is the partial derivative of \hat{f} with respect to a_{pub} . Hence, the optimal division of the budget between observable and unobservable spending is given by:

$$a^*_{pub} = \begin{cases} a'_{pub} & \text{if } a'_{pub} \leq \bar{a}_{pub} \\ \bar{a}_{pub} & \text{if } a'_{pub} > \bar{a}_{pub} \end{cases}$$

The first-order condition implies that at an interior optimum,

$$\hat{f}_a(a^*_{pub}, \bar{\theta})R(a^*_0(a^*_{pub})) = \hat{f}(a^*_{pub}, \bar{\theta})R'(a^*_0(a^*_{pub})) \frac{-\partial a^*_0}{\partial a_{pub}}. \tag{2}$$

The first fact to note from this first-order condition is that, at an interior optimum, the level of unobservable spending on counterterror (a^*_0) is decreasing in the level of observable spending on counterterror.

Remark 5. At the optimum, if the level of publicly observable spending (a_{pub}) increases, the optimal level of unobservable spending (a^*_0) decreases.

Given this, there is a clear intuition underlying the voter’s induced preferences over how to divide the counterterrorism budget between observable and unobservable spending. There are two effects of increasing the level of publicly observable spending that can be seen in the first-order condition. First, it decreases the effectiveness of attempted attacks ($\hat{f}_a(a_{pub}, \bar{\theta})R(a^*_0(a_{pub})) < 0$), which is a marginal benefit for the voter. Second, it changes the equilibrium allocation of unobservable spending ($\partial a^*_0 / \partial a_{pub} < 0$), which changes the resources to which the terrorists have access ($R'(a^*_0(a_{pub})) < 0$). This is a marginal cost from the voter’s perspective. The voter wants the government to balance these marginal benefits and costs, understanding that the government will engage in some diversion of resources used unobservably.

Whether or Not to Reelect

The voter chooses the voting rule to induce the best possible behavior from the government. With respect to the division between observable and unobservable

spending, the voter does this by threatening not to reelect the government unless it chooses the optimal (from the voter’s perspective) incentive-compatible division. As discussed in Remark 3, because of the moral hazard problem, the voter cannot force the government to invest all of the unobservable resources in counterterror. However, given a division between observable and unobservable spending, the voter maximizes the government’s spending on unobservable operations by rewarding counterterror success and sanctioning failure. Thus, the voter’s optimal voting rule is:

$$\Pr(\text{reelect} | a_{pub}, S) = \begin{cases} 1 & \text{if } a_{pub} = a_{pub}^* \text{ and } S = \underline{S} \\ 0 & \text{else.} \end{cases}$$

While this result’s basic idea seems intuitive (counterterrorism failures are likely to hurt a government’s chances of reelection), there certainly exist cases where it does not hold. For example, after 9/11, there seemed to be a “rally” ’round the flag effect” that boosted support for the incumbent government despite the security failure. One should not, however, generalize too much from such examples. Governments in countries that have experienced relatively little terrorism may indeed experience a boost in support following a catastrophic attack. However, this does not describe the bulk of terrorist conflicts, where violence is ongoing and at a lower level. It seems likely that, as more attacks accumulate over time in a society, something like the monotonicity relationship predicted above (more violence implies a lower probability of being reelected) will emerge.²³

Results

Suboptimal Counterterror

The preceding analysis identifies the politically optimal allocation of counterterror resources. How does this political equilibrium compare to the social optimum?

The social optimum would occur if the government had the same preferences as the voters, since this would eliminate the agency problem. Consider, then, a social planner (*SP*) that is a perfect agent. This social planner’s expected utility is given by:

$$U_{SP} = (1 - \pi(\tau))v(\underline{S}).$$

What would the social planner do? First, given a division between observable and unobservable spending, the social planner would allocate the observable portion

23. See Berrebi and Klor 2006, for an estimate of the impact of terrorism on electoral outcomes.

in precisely the same manner that the government in the model above did. Second, given a division between observable and unobservable spending, the social planner would spend all of the unobservable allocation on counterterror, diverting none of it. As such, the social planner’s problem reduces to

$$\max_{a_{pub}} (1 - \pi(\hat{f}(a_{pub}, \bar{\theta})R(A - a_{pub})))v(\underline{\mathcal{L}}).$$

The optimal division between observable and unobservable spending is characterized by the first-order condition

$$-\pi'(\hat{f}(a_{pub}^{SP}, \bar{\theta})R(A - a_{pub}^{SP})) \times [\hat{f}_a(a_{pub}^{SP}, \bar{\theta})R(A - a_{pub}^{SP}) - \hat{f}(a_{pub}^{SP}, \bar{\theta})R'(A - a_{pub}^{SP})]v(\underline{\mathcal{L}}) = 0,$$

which holds only if

$$\hat{f}_a(a_{pub}^{SP}, \bar{\theta})R(A - a_{pub}^{SP}) = \hat{f}(a_{pub}^{SP}, \bar{\theta})R'(A - a_{pub}^{SP}). \tag{3}$$

Comparing equations (2) and (3), the only difference is the presence of the term $-\partial a_0^*/\partial a_{pub}$ in the former. This term provides the essential intuition for the over-spending on observable counterterrorism that occurs in the political equilibrium.

The social planner spends the entire budget on counterterror ($a_0^{SP} + a_{pub}^{SP} = A$). Thus, in effect $\partial a_0^{SP}/\partial a_{pub} = -1$; an additional dollar allocated to publicly observable counterterror by the social planner “costs” exactly a dollar in unobservable counterterror. However, as discussed in Remark 3, in the political equilibrium some of the resources budgeted for unobservable counterterror spending are redirected ($a_0^* + a_{pub}^* < A$). As a result, an additional dollar allocated to publicly observable spending in the political equilibrium costs less than a dollar in unobservable counterterror because only a portion of that dollar was being spent on counterterror; the rest was being diverted. Hence, the marginal cost of publicly observable spending is lower in the political equilibrium than in the social optimum. This drives the voter, in the political equilibrium, to choose a voting rule that induces the government to distort its counterterror spending toward publicly observable counterterror.

Proposition 1. *If the equilibrium amount of publicly observable spending is interior ($a_{pub}^* \leq \bar{a}_{pub}$), then the equilibrium amount of resources devoted to publicly observable counterterror (a_{pub}^*) is greater than the social optimum (a_{pub}^{SP}).*

The voter’s inability to fully monitor the government, coupled with divergence between government and voter preferences, leads the voter to use electoral pressure to force the government to distort the allocation of counterterror resources.

Governments overinvest in publicly observable counterterror not due to irrationality or incompetence; rather, they do so because voters, operating within the constraints of the political-institutional environment, want the government to do so.

This induced overinvestment in observable counterterror means that total counterterror spending is inefficient. That is, holding the level of spending on total counterterror constant, there is always a reallocation of resources that makes both the government and the voter better off. This is because, all else equal, the government and the voter prefer less terrorism. Unconstrained by the voter, the government would invest those resources that it devotes to counterterror efficiently (that is, the marginal benefit of observable spending would equal the marginal benefit of unobservable spending). However, in response to the monitoring problem, the voter demands additional spending on observable tactics. As discussed in Remark 5, the government responds to a demand for increased observable spending with a decrease in unobservable spending. Consequently, the division of resources is inefficient in equilibrium—the marginal benefit of unobservable spending is higher than the marginal benefit of observable spending.

Importantly, despite the inefficiency it induces, the voter is not acting irrationally by constraining the government in this way. The voter imposes a higher level of observable spending because doing so increases the government's total counterterror expenditures, albeit at the price of an inefficient allocation. This increase in total spending more than compensates for the inefficiency, from the voter's perspective.

Proposition 2. *If the equilibrium amount of publicly observable spending is interior ($a_{pub}^* \leq \bar{a}_{pub}$), then, holding total spending fixed, there is a reallocation of observable and unobservable spending that makes the government and voter better off by reducing the impact of terrorism.*

Government Preferences and the Level of Observable Spending

Remark 4 highlighted the fact that, as the level of divergence between government preferences and voter preference decreases, the government becomes increasingly willing to invest resources in unobservable counterterrorism. This suggests that, when the government and voters have similar preferences, the voters should be less concerned about the agency problem and should, therefore, not insist on as much observable spending. This intuition is formalized in the following proposition.

Proposition 3. *The less the government values resources diverted from counterterrorism spending either for rent seeking or other governmental priorities (larger α), the less observable government counterterrorism spending there is and the closer the level of observable spending is to the social optimum.*

Thus, the magnitude of the inefficiency caused by monitoring problems discussed above is determined, in part, by the level of divergence between voter and government preferences.

The Effectiveness of Counterterror

Recall from Remark 1 that, because the terrorists strategically substitute among different tactics in order to maximize their impact, in equilibrium the government allocates tactic-specific observable resources to minimize the effectiveness of the most effective tactics. This leads the government to choose an observable counterterror strategy that makes all maximally effective tactics have the same induced effectiveness. Because of this, the more possible tactics the terrorists have at their disposal (that is, the greater the terrorists' capacity to innovate), the more thinly the government must spread its observable counterterror resources. This gives rise to the following result.²⁴

Proposition 4. As the number of maximally effective terrorist tactics increases, tactic-specific publicly observable counterterror becomes less effective relative to general unobservable counterterror. Moreover, this results in less spending on tactic-specific observable counterterrorism.

If terrorists are reasonably good innovators, so that they have many tactics at their disposal, then substitution effects imply that tactic-specific counterterror is likely to be ineffective at curtailing violence. General operations, alternatively, are unaffected by the terrorists' ability to innovate. Thus, the inefficiency discussed in Proposition 1 can have significant negative security consequences. Of course, the voter responds to this, allowing the government to engage in more unobservable spending, even though this also increases the amount of rent seeking in which the government can engage.

It is important to note that this result is driven not by issues of observability, but by the distinction between tactic-specific and general counterterrorism. An increase in the number of maximally effective tactics decreases the effectiveness of publicly observable counterterror because observable counterterror was assumed to be directed against specific types of attack. Similarly, the reason an increase in the number of maximally effective tactics leads to an increase in unobservable counterterror is because such spending was assumed to be general. As mentioned above, this conflates two issues: observability versus nonobservability and tactic-specific versus general counterterrorism. Proposition 4 shows that when the number of maximally effective tactics increases, the government should substitute counterterrorism spending toward methods that are not tactic specific. Thus, in a model with observable, but non-tactic-specific, counterterrorism, an increase in

24. Faria 2006 presents a similar result.

the number of maximally effective tactics would not necessarily lead to an increase in unobservable counterterrorism spending.

Increased Electoral Incentives

Voters use electoral incentives to induce the government to take actions in line with voter preferences. The strength of these incentives is a function of the benefits associated with reelection.

Recall from equation (1) that the amount of the unobservable counterterror budget that the government spends on counterterrorism balances the marginal benefit associated with increasing the probability of reelection against the marginal cost of spending the money on counterterrorism rather than some other priority. If the voters could increase the benefit associated with reelection, this would raise the marginal benefit of increasing the probability of reelection, which would heighten incentives to invest in unobservable counterterror. This intuition is formalized in the following proposition.

Proposition 5. The amount of the unobservable budget allocated to counterterrorism (a_0^) is increasing in the benefits of holding office (B).*

The fact that the government engages in more unobservable counterterrorism when electoral incentives are large also affects voter incentives. In particular, because the agency problem is mitigated, the voter will demand less observable counterterrorism spending. This intuition is summarized in the following result.

Proposition 6. The level of observable counterterror spending that the voter demands and receives (a_{pub}^) is decreasing in the benefits of holding office (B).*

Voters benefit from increasing the benefits politicians enjoy from being in office. One way voters might increase these benefits is by relaxing the constraints that the government faces in other policy areas. For example, consider a simple, and highly stylized, extension in which, in addition to providing counterterror, the government supplies some other generic public good q at cost $c(q)$. Suppose, further, that the voter has an electoral rule similar to before, but with the added requirement that the voter will not reelect the government unless it produces at least q^* of the public good.²⁵

In this model, the government's benefit from holding office is $B - c(q^*)$. Thus, according to Proposition 5, the voter can increase the government's investment in unobservable counterterrorism by decreasing the amount of other public goods

25. See Persson, Roland, and Tabellini 2000; and Ashworth 2005 for models where this kind of voting rule is derived endogenously. But, also see Padró i Miquel 2005 for a discussion of the problems of using electoral mechanisms to control politicians in multitask settings.

that the government is required to provide in order to be reelected. The voter will, of course, choose q^* strategically, to maximize expected utility. One implication, then, is that if circumstances change such that the amount the voter values security versus other governmental services increases, then the voter can decrease the required provision of public goods other than security and, in so doing, increase government incentives to provide unobservable counterterrorism.

There is an important subtlety to this argument. One would naturally intuit that, following a terrorist attack, voters would reprioritize security issues and there would be a consequent increase in the counterterrorism budget. While this is consistent with the discussion above, the idea here is somewhat different. The implication of the above argument is that, even given an increase in the counterterrorism budget and a resultant decrease in other budgetary spending, the voters should go even further and increase the amount of discretion (or corruption) that the government can exercise over the non-counterterrorism budget by relaxing electoral accountability with respect to non-security-related public goods. That is, for any given non-counterterrorism budget, voters should allow the government to increase the amount of rents it extracts by decreasing the expected output of public goods derived from that budget. In so doing, the voters increase the value to the government of holding office, which decreases the government's incentive to act corruptly with respect to the counterterrorism budget.

Conclusion

I have presented a model that captures aspects of the strategic relationship between voters, governments, and terrorists in democratic societies facing a terrorist threat. The model focuses on a previously unexplored categorization of counterterrorism strategies into observable and unobservable tactics. I demonstrated that, when there is divergence between voter and government preferences, agency problems between the voter and government create a situation in which the government's politically optimal counterterrorism strategy is quite different from the security-maximizing counterterrorism strategy, even holding total counterterror spending constant. In particular, in response to electoral pressure, the government overallocates resources to tactic-specific observable counterterror. This problem is particularly severe when the divergence between government and voter preferences is significant. Moreover, a focus on tactic-specific counterterrorism is particularly ineffective when terrorists have a large set of tactics from which to choose. Voters can decrease the magnitude of the agency problem by increasing the benefits of reelection by, for example, slackening requirements for non-security-related public goods.

The model is clearly quite stylized. Nonetheless, the key intuition does have some implications for public policy. When voters are concerned that there is significant divergence between their and their government's preferences over counterterror, the model suggests that counterterror allocations will be severely distorted. In order to overcome the moral hazard problem that leads to government overin-

vestment in observable counterterrorism, voters need some mechanism by which to monitor government spending on unobservable counterterrorism. The problem, of course, is that such spending must be kept secret to be effective and to avoid substitution by the terrorists.

One possible institutional response that could mitigate the agency problem while preserving secrecy is a nonpartisan, nonelected monitoring body (such as the 9/11 Commission in the United States). Such an organization, of course, could not fully observe covert spending. However, if the monitoring body were able to uncover even a noisy signal of unobservable spending and communicate it to the voters, then the voters could condition their voting decisions on this information and increase incentives for the government to invest unobservable resources in counterterrorism. This would have two benefits: first, it would increase the amount of unobservable spending directed toward the counterterrorism uses for which it was intended; and, second, in so doing it would diminish the voters' incentives to use electoral pressure to induce the government to overspend on observable counterterrorism.

Of course, such a prescription would have implementation problems and is only a partial solution. How can a commission be created that is in fact independent, given that it requires government cooperation to do its work effectively? How can information be credibly transmitted to voters without violating secrecy requirements? In systems with separation of powers, legislative oversight, multiple bureaucracies responsible for similar tasks, and other complications, who exactly should be held to account for policy failures? These issues are critical, and any workable policy would have to address them in one form or another. The contribution of this analysis, thus, is not a precise solution. Rather, the hope is that identifying a causal mechanism, and elucidating its implications, brings us one step closer to a more nuanced understanding of the politics of counterterrorism.

Appendix

Proof of Lemma 1. It follows from the Kuhn-Tucker conditions that if t_i and t_j receive positive budget shares, then $f(a_i, \theta_i) = f(a_j, \theta_j)$. Consider the possibility that some type of attack $t_k \notin \bar{\mathbf{t}}$ receives positive budget share. This implies that $\pi'(\tau)u_T(\bar{S})f(a_k, \theta_k) = \lambda$. By the definition of $\bar{\mathbf{t}}$, there is some $t_i \in \bar{\mathbf{t}}$ where $\pi'(\tau)u_T(\bar{S})f(a_i, \theta_i) > \lambda$ that does not receive positive budget share. This is clearly suboptimal, so we have a contradiction, so all $t_k \notin \bar{\mathbf{t}}$ must receive no budget share. Finally, from $f(a_i, \theta_i) = f(a_j, \theta_j)$ for all $t_i, t_j \in \bar{\mathbf{t}}$ it follows that the terrorists are indifferent among the elements of $\bar{\mathbf{t}}$.

Proof of Remark 1. I first establish that $f(a_i, \theta_i) = f(a_j, \theta_j)$ for all tactics that receive positive counterterrorism budget share. Suppose there is an optimal allocation where tactics i and j each receive positive budget share with $f(a_i, \theta_i) < f(a_j, \theta_j)$. By Lemma 1, the terrorists will not invest in tactic i . Thus, the government could reduce a_i , reallocate it to tactics the terrorists do invest in, and improve its expected payoff. Thus, the allocation was not optimal.

Second, I establish that the government does not devote resources to tactics that are not maximally effective. Suppose the government devoted resources to a non-maximally effective tactic, k . By the argument above, we have $f(a_k, \theta_k) = \hat{f}_a(a_{pub}, \bar{\theta})$. But by the definition of maximal effectiveness, we have $f(a_k, \theta_k) < \hat{f}_a(a_{pub}, \bar{\theta})$, a contradiction.

Third, I establish that the government invests resources in all maximally effective tactics. Suppose there is an optimal allocation (a_1, \dots, a_n) that puts no budget share on some maximally effective tactic i . By the definition of maximal effectiveness, that tactic will have effectiveness $f(0, \theta_i) > \hat{f}_a(a_{pub}, \bar{\theta})$. By Lemma 1, this implies that the terrorists will not invest in any of the tactics receiving positive counterterror budget shares. Thus, the government could improve its expected utility by reallocating resources to those tactics that the terrorists are investing positive budget share in, so the allocation was not optimal.

Finally, I show that the terrorists invest equally in all maximally effective tactics. Suppose the terrorists invest less in some maximally effective tactic i than in others. One can identify two maximally effective tactics i and j that have the same effectiveness $\hat{f}_a(a_{pub}, \bar{\theta})$ but where $t_i < t_j$. The government's expected utility would be improved by reallocating from i to j . But then, $f(a_i, \theta_i) > \hat{f}_a(a_{pub}, \bar{\theta})$, so the terrorists reallocate to t_i . Thus, there is no equilibrium where the terrorists do not equally invest in all maximally effective tactics.

Proof of Remark 3. The first-order condition is

$$-\pi'[\hat{f}(a_{pub}, \bar{\theta})R(a_0^*)]\hat{f}(a_{pub}, \bar{\theta})R'(a_0^*)(B + \alpha v(\underline{S})) = u'_G(A - a_{pub} - a_0^*).$$

Since the left-hand side is finite, and $\lim_{x \rightarrow 0} u'_G(x) = \infty$, we have $a_0^* < A - a_{pub}$, thus the government diverts some of the budget.

Proof of Remark 4. The government's objective function is

$$U_G = (1 - \pi(\hat{f}(a_{pub}, \bar{\theta})R(a_0))) (B + \alpha v(\underline{S})) + (1 - \alpha)u_G(A - a_{pub} - a_0).$$

Taking the cross-partial with respect to a_0 and α yields:

$$\frac{\partial^2 U_G}{\partial a_0 \partial \alpha} = -\pi'(\hat{f}(a_{pub}, \bar{\theta})R(a_0))R'(a_0) + u'_G(A - a_{pub} - a_0) > 0,$$

where the inequality follows from $\pi'(\cdot) > 0, R'(\cdot) < 0$, and $u'_G(\cdot) > 0$. Thus, the objective function has strictly increasing marginal returns in a_0 and α , so Theorem 3 of Edlin and Shannon implies that a_0^* is increasing in α .²⁶

Proof of Remark 5. The first-order condition implies that

$$\frac{\partial a_0^*}{\partial a_{pub}} = -\frac{\hat{f}_a(a_{pub}^*, \bar{\theta})R(a_0^*(a_{pub}^*))}{\hat{f}(a_{pub}^*, \bar{\theta})R'(a_0^*(a_{pub}^*))} < 0, \tag{4}$$

26. Edlin and Shannon 1998.

where the inequality follows from the fact that $\hat{f}_a < 0$ and $R' < 0$.

Proof of Proposition 1. I will use the following claim.

Claim 1. $\partial a_0^*/\partial a_{pub} > -1$

At an interior solution, the equilibrium level of observable spending is given by equation (2):

$$\hat{f}_a(a_{pub}^*, \bar{\theta})R(a_0^*(a_{pub}^*)) = \hat{f}(a_{pub}^*, \bar{\theta})R'(a_0^*(a_{pub}^*)) \frac{-\partial a_0^*}{\partial a_{pub}}$$

From equation (3) the social optimum is:

$$\hat{f}_a(a_{pub}^{SP}, \bar{\theta})R(A - a_{pub}^{SP}) = \hat{f}(a_{pub}^{SP}, \bar{\theta})R'(A - a_{pub}^{SP})$$

The left-hand sides are the same in both equations. The only difference on the right-hand side is the presence of the term $-\partial a_0^*/\partial a_{pub}$. By the claim, $-\partial a_0^*/\partial a_{pub} < 1$, so the right-hand side of the former is point-wise less than the right-hand side of the latter, which implies that $a_{pub}^* > a_{pub}^{SP}$. All that remains is to prove the claim.

Proof of Claim 1. From equation (4),

$$\frac{\hat{f}_a(a_{pub}^*, \bar{\theta})R(a_0^*(a_{pub}^*))}{\hat{f}(a_{pub}^*, \bar{\theta})R'(a_0^*(a_{pub}^*))} = -\frac{\partial a_0^*}{\partial a_{pub}}$$

From equation (3),

$$\frac{\hat{f}_a(a_{pub}^{SP}, \bar{\theta})R(A - a_{pub}^{SP})}{\hat{f}(a_{pub}^{SP}, \bar{\theta})R'(A - a_{pub}^{SP})} = 1$$

By the definition of a social optimum, for any a_{pub} , $a_0^{SP} = A - a_{pub}$. From Remark 3, for any a_{pub} , $a_0^* < A - a_{pub}$. This implies that for any a_{pub} , $R(a_0^*) < R(a_0^{SP})$ and $R'(a_0^*) > R'(a_0^{SP})$. Thus, for any a_{pub} we have

$$-\frac{\partial a_0^*}{\partial a_{pub}} = \frac{\hat{f}_a(a_{pub}, \bar{\theta})R(a_0^*(a_{pub}))}{\hat{f}(a_{pub}, \bar{\theta})R'(a_0^*(a_{pub}))} < \frac{\hat{f}_a(a_{pub}, \bar{\theta})R(a_0^{SP}(a_{pub}))}{\hat{f}(a_{pub}, \bar{\theta})R'(a_0^{SP}(a_{pub}))} = 1,$$

which establishes the claim.

Proof of Proposition 2. Both the voter's and the government's objective are decreasing in $\tau = \hat{f}(a_{pub}, \bar{\theta})R(a_0)$, holding total spending fixed. Thus, it suffices to show that in equilibrium, τ can be decreased without altering spending. Let total spending be $a = a_{pub} + a_0$.

Then, we solve

$$\min_{a_{pub}, a_0} \hat{f}(a_{pub}, \bar{\theta})R(a_0) + \lambda(a - a_{pub} - a_0),$$

where λ is the Lagrangian multiplier. Thus, the first-order conditions imply that at the minimum $[\hat{f}_a(a_{pub}, \bar{\theta})R(a_0)]/[\hat{f}(a_{pub}, \bar{\theta})R'(a_0)] = 1$. But, equation (4) and claim (1) show that this is not true in equilibrium.

Proof of Proposition 3. I will use the following claim.

Claim 2. If $\alpha > \alpha'$, then $\partial a_0(\alpha')/\partial a_{pub} > \partial a_0^(\alpha)/\partial a_{pub}$.*

At an interior solution, the equilibrium level of observable spending is given by equation (2):

$$\hat{f}(a_{pub}, \bar{\theta})R(a_0^*) = \hat{f}(a_{pub}, \bar{\theta})R'(a_0^*) \frac{-\partial a_0^*(\alpha)}{\partial a_{pub}}.$$

By claim (2), the right-hand side of this first-order condition is point-wise greater under α' than α , which implies that $a_{pub}^*(\alpha) < a_{pub}^*(\alpha')$. Moreover, it is clear that $a_{pub}^*(1) = a_{pub}^{SP}$, so as α approaches 1, the level of public spending converges monotonically to the social optimum. All that remains is to prove the claim.

Proof of Claim 2. From equation (4),

$$\frac{\hat{f}_a(a_{pub}^*, \bar{\theta})R(a_0^*(\alpha, a_{pub}^*))}{\hat{f}(a_{pub}^*, \bar{\theta})R'(a_0^*(\alpha, a_{pub}^*))} = -\frac{\partial a_0^*(\alpha)}{\partial a_{pub}}.$$

By Remark 4, $a_0^*(\alpha) > a_0^*(\alpha')$. Thus, for any a_{pub} , $R(a_0^*(\alpha)) < R(a_0^*(\alpha'))$ and $R'(a_0^*(\alpha)) > R'(a_0^*(\alpha'))$. This implies that

$$-\frac{\partial a_0^*(\alpha')}{\partial a_{pub}} = \frac{\hat{f}_a(a_{pub}^*, \bar{\theta})R(a_0^*(\alpha', a_{pub}^*))}{\hat{f}(a_{pub}^*, \bar{\theta})R'(a_0^*(\alpha', a_{pub}^*))} < \frac{\hat{f}_a(a_{pub}^*, \bar{\theta})R(a_0^*(\alpha, a_{pub}^*))}{\hat{f}(a_{pub}^*, \bar{\theta})R'(a_0^*(\alpha, a_{pub}^*))} = -\frac{\partial a_0^*(\alpha)}{\partial a_{pub}},$$

which establishes the claim.

Proof of Proposition 4. Consider a vector of parameters $\bar{\theta}$, with m maximally effective tactics $(\theta_1, \dots, \theta_m)$ and the corresponding optimal allocation of public spending (a_1, \dots, a_m) , $\sum_{i=1}^m a_i = \hat{a}_{pub}$. From Lemma 1, $f(a_i, \theta_i) = f(a_j, \theta_j) \equiv \hat{f}(\hat{a}_{pub}, \bar{\theta})$ for all $i, j \in \{1, 2, \dots, m\}$. Now consider adding another maximally effective tactic $m + 1$, with parameter θ_{m+1} , so that, by the definition of maximal effectiveness, $f(0, \theta_{m+1}) > \hat{f}(\hat{a}_{pub}, \bar{\theta})$. To maintain the same level of effectiveness of publicly observable counterterror, the government would have

to allocate resources to tactic $m + 1$ such that $f(a_{m+1}, \theta_{m+1}) = \hat{f}(\hat{a}_{pub}, \bar{\theta})$, which, for any positive a_{m+1} would violate the budget constraint $\sum_{i=1}^{m+1} a_i = \hat{a}_{pub}$. Thus, publicly observable counterterror spending becomes less effective as the number of tactics increases. The effectiveness of unobservable counterterror spending ($R'(a_0)$) is constant in the number of tactics.

The second result follows from the observation that as the number of maximally effective tactics increases, \hat{f}_a decreases. Thus, the left-hand side of equation (2) decreases, decreasing the optimal a_{pub}^* .

Proof of Proposition 5. Taking the cross-partial of equation (1) with respect to B yields:

$$\frac{\partial^2 U_G}{\partial a_0 \partial B} = -\pi' [\hat{f}(a_{pub}, \bar{\theta})R(a_0)] \hat{f}(a_{pub}, \bar{\theta})R'(a_0) > 0.$$

Thus, U_G has strictly increasing marginal returns in a_0 and B , which, by Theorem 3 of Edlin and Shannon, implies that the optimal a_0^* is strictly increasing in B .²⁷

Proof of Proposition 6. I will make use of the following claim.

Claim 3. Let $B > B'$. Then $\partial a_0^(B')/\partial a_{pub} > \partial a_0^*(B)/\partial a_{pub}$.*

Given the claim, the proof is the same as the proof of Proposition 3. Proposition 5 implies that $a_0^*(B) > a_0^*(B')$. Given this, the proof of the claim follows from the same argument as the proof of claim (2).

References

- Ashworth, Scott. 2005. Reputational Dynamics and Political Careers. *Journal of Law, Economics and Organization* 21 (2):441–66.
- Austen-Smith, David, and Jeffrey Banks. 1989. Electoral Accountability and Incumbency. In *Models of Strategic Choice in Politics*, edited by Peter C. Ordeshook, 121–48. Ann Arbor: University of Michigan Press.
- Barro, Robert. 1973. The Control of Politicians: An Economic Model. *Public Choice* 14 (1):19–42.
- Berrebi, Claude, and Esteban Klor. 2006. On Terrorism and Electoral Outcomes: Theory and Evidence from the Israeli-Palestinian Conflict. *Journal of Conflict Resolution* 50 (6):899–925.
- Bueno de Mesquita, Ethan. 2005. The Quality of Terror. *American Journal of Political Science* 49 (3):515–30.
- Bueno de Mesquita, Ethan, and Eric Dickson. 2007. The Propaganda of the Deed: Terrorism, Counterterrorism, and Mobilization. *American Journal of Political Science* 51 (2), forthcoming.
- Bueno de Mesquita, Ethan, and Amanda Friedenberg. 2005. Optimal Retrospective Voting. Unpublished manuscript, Washington University, St. Louis, Mo.
- Crenshaw, Martha. 1995. Thoughts on Relating Terrorism to Historical Contexts. In *Terrorism in Context*, edited by Martha Crenshaw, 3–26. University Park: Pennsylvania State University Press.

27. Ibid.

- de Figueiredo, Rui J. P., Jr., and Barry R. Weingast. 2001. Vicious Cycles: Endogenous Political Extremism and Political Violence. Institute of Governmental Studies Working Paper 2001-9. University of California, Berkeley.
- Edlin, Aaron S., and Chris Shannon. 1998. Strict Monotonicity in Comparative Statics. *Journal of Economic Theory* 81 (1):201–19.
- Enders, Walter, and Todd Sandler. 1993. The Effectiveness of Anti-Terrorism Policies: A Vector-Autoregression-Intervention Analysis. *American Political Science Review* 87 (4):829–44.
- . 2002. Patterns of Transnational Terrorism, 1970–99: Alternative Time Series Estimates. *International Studies Quarterly* 46 (2):145–65.
- Enders, Walter, Todd Sandler, and Jon Cauley. 1990. UN Conventions, Technology, and Retaliation in the Fight Against Terrorism: An Econometric Evaluation. *Terrorism and Political Violence* 2 (1):83–105.
- Faria, Joao Ricard. 2006. Terrorist Innovations and Anti-terrorist Policies. *Terrorism and Political Violence* 18 (1):47–56.
- Ferejohn, John. 1986. Incumbent Performance and Electoral Control. *Public Choice* 50 (1–3):5–26.
- Gibbs, Jack. 1989. Conceptualization of Terrorism. *American Sociological Review* 54 (3):329–40.
- Hoffman, Bruce. 1998. *Inside Terrorism*. New York: Columbia University Press.
- Holmström, Bengt, and Paul Milgrom. 1991. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics and Organization* 7 (Special Issue):24–52.
- Im, Eric Iksoo, Jon Cauley, and Todd Sandler. 1987. Cycles and Substitutions in Terrorist Activities: A Spectral Approach. *Kyklos* 40 (2):238–55.
- Kibbe, Jennifer D. 2004. The Rise of the Shadow Warriors. *Foreign Affairs* 83 (3):102–15.
- Laqueur, Walter. 1977. *Terrorism*. London: Weidenfeld and Nicolson.
- Lichbach, Mark Irving. 1987. Deterrence or Escalation?: The Puzzle of Aggregate Studies of Repression and Dissent. *Journal of Conflict Resolution* 31 (2):266–97.
- Padró i Miquel, Gerard. 2005. Multitask Political Agency. Unpublished manuscript, Stanford University, Stanford, Calif.
- Persson, Torsten, Gerard Roland, and Guido Tabellini. 2000. Comparative Politics and Public Finance. *Journal of Political Economy* 108 (6):1121–61.
- Powell, Robert. Forthcoming. Defending Against Terrorist Attacks with Limited Resources. *American Political Sciences Review*.
- Rosendorff, Peter, and Todd Sandler. 2004. Too Much of a Good Thing? The Proactive Response Dilemma. *Journal of Conflict Resolution* 48 (5):657–71.
- Sandler, Todd, and Harvey E. Lapan. 1988. The Calculus of Dissent: An Analysis of Terrorists' Choice of Targets. *Synthese* 76 (2):245–61.
- Sandler, Todd, and Kevin Siqueira. 2006. Global Terrorism: Deterrence versus Preemption. *Canadian Journal of Economics* 39 (4):1370–87.