Optimizing Counter-Terror Operations: Should One Fight Fire with "Fire" or "Water"?

Jonathan P. Caulkins* Dieter Grass† Gustav Feichtinger† and Gernot Tragler†

Abstract

This paper deals dynamically with the question of how recruitment to terror organizations is influenced by counter-terror operations. This is done within a optimal control model, where the key state is the (relative) number of terrorists and the key controls are two types of counter-terror tactics, one ("water") that does not one ("fire") that does provoke recruitment of new terrorists. The model is nonlinear and does not admit analytical solutions, but an efficient numerical implementation of Pontryagin's Minimum Principle allows for solution with base case parameters and considerable sensitivity analysis. Generally this model yields two different steady states, one where the terror-organization is nearly eradicated and one with a high number of terrorists. Whereas water strategies are used at almost any time, it can be optimal not to use fire strategies if the number of terrorists is below a certain threshold.

1 Introduction

One concern expressed about aggressive counter-terror operations is that they might turn public opinion against the counter-terror forces, thereby helping terrorist organizations recruit new personnel. For example, Heymann (2003) notes that "recruitment to the Irish Republican Army (IRA) increased sharply during some periods of overly vigorous British action against suspects". Likewise Kaplan et al. (2005) develop an empirical model that finds that Israeli "killings of terror suspects by Israel sparks estimated recruitment to the terror stock that increases rather than decreases the rate of suicide bombings."

It does not seem credible that the best counter-terror strategy is to do nothing proactive. On the other hand, it also does not seem prudent to ignore entirely the possibility that offensive counter-terror operations might be a double-edged

*H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, U.S.A.

†Department of Operations Research and Systems Theory, Vienna University of Technology, Argentinierstrasse 8, A-1040 Vienna, Austria

sword, stimulating new recruitment even as they kill current terrorists. Presumably there may be some optimal intensity of counter-terror efforts that balances their obvious benefits with potential negative effects on recruitment and/or some optimal mix of tactics that are more or less prone to stimulate new recruitment.

This paper introduces a stylized optimal dynamic control model that frames this trade-off in mathematical notation and explores the implications of the model solutions. The work is clearly exploratory inasmuch as functional forms and parameters are not derived from empirical data, but the mathematical framing lends precision to underlying concepts that may clarify some associated issues.

2 Model

There are a wide range of counter-terror tactics that might be arrayed along a continuum in terms of their likelihood of stimulating recruitment, but the model abstracts these into two discrete categories. "Fire" strategies anger potential recruits to terrorist organizations enough to make it easier for the terrorists to recruit more personnel. They might include killing terrorists via aerial bombing of residential neighborhoods, aggressively searching all people passing through a checkpoint or road block, or other tactics that involve significant collateral damage or inconvenience to innocent third parties.

"Water" strategies "play by the rules" in the eyes of the general population. They are intelligence-driven arrests or "surgical" operations against individuals who are known with very high confidence to be guilty and which do not harm innocent parties. The presumption is that water strategies are more costly and/or are inherently limited, e. g., by lack of "actionable targeting intelligence".

The extent to which the decision maker employs fire and water strategies will be denoted by the (control) variables $v(t)$ and $u(t)$, respectively. Since in dynamic optimization all variables are functions of time, we follow the custom of omitting explicit mention of the time argument in the sequel, and so refer to the fire and water controls as simply v and u.

The strength of the terrorist organizations is reflected in the (single) state variable, x. That variable might most simply be thought of as the number of active terrorists, but more generally could reflect the organization's total resources including financial capital, weapons, technological know-how, etc., as in Keohane and Zeckhauser (2003).

The size and strength of the terrorist organization evolves over time in response to various inflows and outflows. Keohane and Zeckhauser (2003) and Kaplan et al. (2005) model baseline recruitment as a constant. We likewise include such a constant recruitment term, τ. However, we presume that τ is small, accounting for a minority of new recruits. It is common in modeling of drug and criminal enterprises to treat recruitment (initiation) as an increasing function of the current state Caulkins (2005), and other authors (e.g., Castillo-Chavez & Song, 2003) have applied this reasoning to recruitment of terrorists. The more terrorists there are, the more new terrorists the organization can recruit per unit time because recruiting occurs through personal interaction, as in diffusion models of word-of-mouth advertising. Organizational growth is not exponential and unbounded because there are other limiting factors that can reduce per capita recruitment as organizational size increases. On the other hand, a central premise of this paper is that aggressive use of fire strategies (denoted v) increases recruitment. So this second inflow can be denoted by $I(x, v)$, where $I_x > 0, I_{xx} < 0$ and $I_v > 0$. We do not assume that terrorism can be blamed entirely on, let along justified by, over-aggressive counter-terror operations, so $I(0, x) > 0$. In other words, recruitment of new terrorists occurs even when no "fire" tactics are employed.

We distinguish three outflows from the stock of terrorists. The first represents the rate at which people leave the organization, die in suicide attacks, or fall victim to routine or background enforcement efforts that are in addition to those the decision maker in this problem is actively managing. This flow is a function only of x and is represented by $O_1(x)$, with $O_1x > 0$. The second outflow reflects the consequences of water operations and so is given by $O_2(u, x)$. The final is outflow is due to fire operations and is denoted by $O_3(v, x)$.

Over the (potentially infinite) planning horizon, the decision maker seeks to minimize the sum of both some function of x plus the amount spent on both fire (v) and water (u) counter-terror operations. It is standard to presume that outcomes are discounted at some constant rate r per unit time, so the full model can be written:

$$\min_{u,v\geq 0} \int_0^{\infty} e^{-rt} \left(C_1(x) + C\left(u, v\right)\right) dt$$
$$\text{s.t.} \quad \dot{x} = \tau + I(v, x) - O_1(x) - O_2(u, x) - O_3(v, x)$$
$$x(0) = x_0$$

Relatively little of interest can be derived without specifying functional forms for the costs ($C_1(x)$, $C(u, v)$, Outflow rates ($O_1(x)$, $O_2(u, x)$ $O_3(v, x)$), and initiation function ($I(v, x)$). We know of no data or literature that would support any specific functional forms, so we opt for simple and transparent approach of adopting very simple forms and focus on qualitative behavior of the solutions. Again, the goal is to lend precision to a set of concepts, not to compute specific, quantitative policy prescriptions.

The costs of terrorism are assumed to be linear in the number of terrorists. That is $C_1(x) = cx$. Kaplan et al. (2005) model the log of the number of terror attacks as being linear in the stock of terrorists, which would suggest a cost term that is exponential in x. We also analyzed that case and obtained results that are qualitatively similar, but make the linear form our base model for three reasons. The linear form makes intuitive sense, it is easier to analyze, and, most importantly, its simplicity makes clear that the interesting threshold results obtained below stem from the systems dynamics, not the choice of functional form in the objective.

Turning to the control costs, one could imagine that there might be some adverse interaction such that $C_uv > 0$. For example, using fire tactics could make it more difficult (expensive) to use water tactics. We will see below that even without such interactions, the optimal solution does not always use positive levels of both controls. To avoid the appearance that when fire tactics are eschewed it stems from such a cost interaction, we model the control cost function as being separable, $C(u, v) = C_2(u) + C_3(v)$. Exploring interactions might, however, be a fruitful area for further research.

Specifically, the costs of employing the water and fire strategies are presumed to be concave for the usual diminishing returns arguments and so are modeled as being quadratic. That is $C_2(u) = u^2$ and $C_3(v) = v^2$.

Following the tradition of drug and crime models, recruitment is modeled by a power function in x. Counter-terror operations are assumed to accelerate recruitment by a percentage that is proportional to the intensity of the fire operations, with proportionality constant ρ. Hence, function $I(v, x)$ is taken to be $(1 + \rho v)kx^{\alpha}$. Note: eliminating fire operations reduces "word of mouth" recruitment to $I(v, x) = kx^{\alpha}$, but it does not eliminate such recruitment.

The "natural" outflow $O_1(x)$ is presumed to occur at a constant per capita rate $O_1(x) = \mu x$. The outflows $O_2(x, u), O_3(x, v)$ stemming from counter-terror operations are modeled as products of the controls and the number of terrorists. That is $O_2(x, u) = \beta(u)f_1(x)$ and $O_3(x, v) = \nu(v)f_2(x)$. The outflow due to fire strategies $O_2(x, v)$ is modeled as being linear in x, because the methods are perceived to be "shotgun" or "undirected" methods and hence $f_2(x) = x$. The more targets there are (i.e., the larger x is), the more terrorists will be hit. In contrast, the outflow due to water strategies is presumed to be concave in x because actionable intelligence is limited and heterogeneous. There may also be diminishing returns if more specialized skills are needed, so there are a limited number of units that can conduct water operations. In particular, this is modeled by a power function so $f_1(x) = x^{\theta}$.

Finally the functional forms of $\beta(u)$ and $\nu(v)$ have to be specified. They should be concave, and we use the same functional form for both so that any differences in the way the optimal solution employs fire and water strategies stem from their differences in character, not from rather arbitrary decisions concerning functional forms. It turns out that a logarithmic function is convenient for the analysis. The one difference between the water and fire functions is that water operations are more expensive, so for any given level of spending z, $\beta(z) < \nu(z)$. We accommodate that difference simply by pre-multiplying the water function by a constant β that is smaller than the corresponding constant (ν) for fire operations. These two constants reflect the "efficiency" of the two types of operations and below we conduct considerable sensitivity analysis with respect to both constants.

In sum, the model we wish to investigate can be written as

$$
\begin{aligned}
&\min_{u,v\geq 0} \int_0^\infty e^{-rt}\left(cx+u^2+v^2\right)dt \\
&\text{s.t.}\quad \dot{x}=\tau+(1+\rho v)kx^\alpha-\mu\, x-\beta\,\ln(1+u)\,x^\theta-\nu\,\ln(1+v)\,x \\
&\text{and}\quad x(0)=x_0.
\end{aligned}
\tag{FOW}
$$

3 Analysis

3.1 Optimality Conditions

The Fire or Water (FOW) model can be solved by applying Pontryagin's minimum principle (see, e.g., Feichtinger & Hartl, 1986; Leonard & Long, 1992). Therefore we consider the current value Hamiltonian H

$$
H=\lambda_0\left(cx+u^2+v^2\right)+\lambda\left(\tau+(1+\rho v)\,kx^\alpha-\mu\, x-\beta\,\ln(1+u)\,x^\theta-\nu\,\ln(1+v)\,x\right), \tag{1}
$$

where $\lambda_0\in\mathbb{R}$ and λ denotes the costate variable in current value terms.

First we have to exclude the degenerate case $\lambda_0=0$. To do so we first assume $\lambda_0=0$ and then derive a contradiction. If $\lambda_0=0$ we find

$$
H_u=-\frac{\beta x^\theta}{1+u}<0.
$$

Hence H will reach its minimum only for $u=\infty$, which is not a feasible solution. So $\lambda_0>0$ and can without loss of generality be normalized to 1 and hence omitted in the following.

Applying the standard methods of optimal control theory we derive the necessary optimality condition

$$
u^*=\arg\min_{u\geq 0} H,
$$

For an inner minimum $H_u=0$ has to hold, which implies

$$
H_u=2u-\frac{\lambda\beta x^\theta}{1+u},
$$

and hence we get

$$
u^*=\frac{1}{2}\left(\sqrt{1+2\lambda\beta x^\theta}-1\right). \tag{2}
$$

Since

$$
v^*=\arg\min_{v\geq 0} H,
$$

we have to solve $H_v=0$ for an inner minimum v^* with

$$
H_v=2v+\lambda\left(\rho kx^\alpha-\frac{\nu x}{1+v}\right),
$$

which yields

$$v^* = \frac{1}{4}\left(-\rho k x^\alpha - 2 + \sqrt{\left(\lambda\rho k x^\alpha - 2\right)^2 + 8\lambda\nu x}\right). \tag{3}$$

As the derivative of H with respect to the state x equals

$$H_x = c + \lambda\left(k\left(1+\rho v^*\right)\alpha x^{\alpha-1} - \mu - \beta\ln\left(1+u^*\right) - \nu\ln\left(1+v^*\right)\right) \tag{4}$$

the canonical dynamical system with inactive control constraint (SIC) can be written as

$$\dot{x} = \tau + \left(1+\rho v^*\right)k x^\alpha - \mu x - \beta\ln\left(1+u^*\right)x^\theta - \nu\ln\left(1+v^*\right)x \tag{5}$$

$$\dot{\lambda} = \lambda\left(r - k\left(1+\rho v^*\right)\alpha x^{\alpha-1} + \mu + \beta\ln\left(1+u^*\right) + \nu\ln\left(1+v^*\right)\right) - c, \tag{6}$$

where u^* and v^* are given by Eq. 2 and Eq. 3 respectively.

The state x denotes the number of terrorists and the costate λ is its shadow-price, meaning the increase in the objective function due to an infinitesimal increase in x. That both quantities are always positive can be derived from Eqs. 5-6. According to Eq. 5, $\dot{x} = \tau > 0$, whenever $x = 0$. Hence x remains positive for every initial state $x_0 > 0$. A similar argument ensures the positivity of λ. It can be seen from Eq. 6 that $\dot{\lambda} = -c < 0$ if $\lambda = 0$. Hence if there exists a t_0 with $\lambda(t_0) \leq 0$, then $\lambda(t)$ would remain negative thereafter. But in the long run $\lim_{t\to\infty}\lambda(t) > 0$, so there must not exist such a t_0. I.e., $\lambda(t)$ is positive for every t.

What remains to be shown is the minimum property of Eqs. 2-3. But as the Hesse matrix of H in respect to u and v is of the form

Positivity of x and λ leads to three further results. First, from Eq. 2 it follows that u^* is positive since $1 + 2\lambda\beta x^\theta > 1$ holds for $x, \lambda > 0$. Second, $\lambda > 0$ implies that increasing the stock of terrorists is always bad, so one would never invest in a control strategy that stimulates recruitment of more terrorists than it eliminates, regardless of the intervention's cost, so long as costs are not negative. Third, the Hessian matrix of H with respect to u and v

$$D^2 H = \begin{bmatrix} 2 + \frac{\lambda\beta x^\theta}{(1+u)^2} & 0 \\ 0 & 2 + \frac{\lambda\nu x}{(1+v)^2} \end{bmatrix}, \tag{7}$$

is positive definite as long as x and λ are positive, so it satisfies the Legendre-Clebsh condition, which is a sufficient minimum condition for the controls. However, the Hamiltonian H is not concave with respect to x, so the common sufficient condition for the existence of an optimal solution does not hold. Therefore from a rigorous mathematical point of view, every solution calculated to be optimal below can only be denoted as a candidate $(\hat{x}, \hat{\lambda})$ for an optimal solution with the possible optimal controls $(\hat{u}, \hat{v})$.

Nevertheless we can calculate the total costs

$$J(\hat{x}, \hat{u}, \hat{v}) = \int_0^\infty e^{-rt}\left(c\hat{x} + \hat{u}^2 + \hat{v}^2\right)dt$$

for each of these candidates numerically. If multiple candidates exist the costs are compared to find the minimum. See e. g. Fig. 1 where this is displayed for the base case (defined in Sec. 3.3). The detailed analysis of the base case can be found in Sec. 4.1.

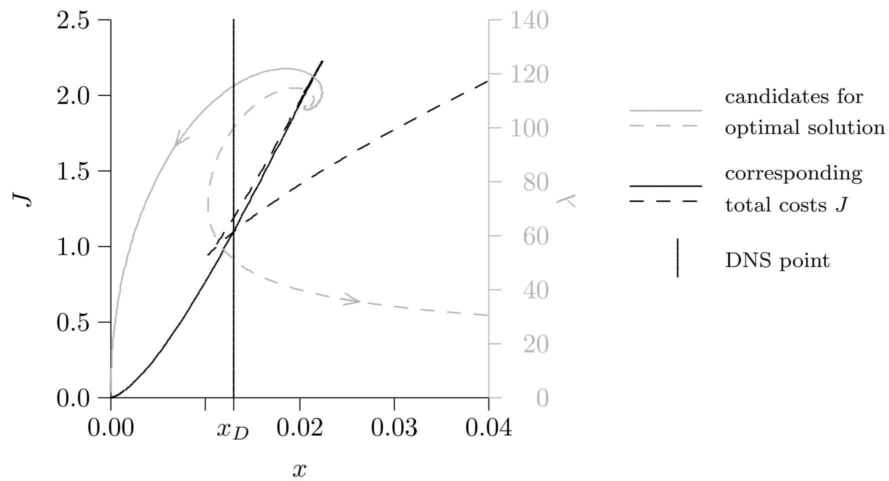

Figure 1: The two candidates for an optimal solution together with the corresponding total costs J are depicted. In x_D the cost functions intersect, which characterizes a Dechert-Nishimura-Skiba (DNS) point.

The water control u has been shown to be positive, but positivity cannot be demonstrated for the fire control v, so we have to include explicitly a non-negativity constraint on v. This means that the Lagrangian

$$L = H + lv \tag{8}$$

has to be considered. The necessary optimality conditions under this control constraint can be stated as (see, e.g., Feichtinger & Hartl, 1986; Leonard & Long, 1992):

$$L_v = 0,$$

which implies an explicit formula for the Lagrange multiplier

$$l = -H_v. \tag{9}$$

Furthermore the following non-positive and slackness conditions have to hold

$$l \leq 0, \quad lv = 0.$$

Using Eqs. 5-6 the dynamical system with active control constraint (SAC), that is $v^* = 0$, is given by

$$\dot{x} = \tau + kx^{\alpha} - \mu x - \beta \ln(1 + u^*)\, x^{\theta} \tag{10}$$

$$\dot{\lambda} = \lambda \left(r - k\alpha x^{\alpha-1} - \mu - \beta \ln(1 + u^*)\right) - c. \tag{11}$$

3.2 Steady States

Due to the nonlinearity of the system dynamics (Eqs. 5-6 and Eqs. 10-11) the steady states can only be calculated numerically. All the calculations are done within the MATLAB® environment, using especially the MATCONT toolbox (see Doedel, Govaerts, and Kuznetsov (2003)). To distinguish between steady states when the control constraint is inactive (SIC) as opposed to active (SAC), a superscript i is used for SIC and superscript a is used for SAC. Likewise we distinguish between steady states with a high and low number of terrorists (x) by the subscripts l for low and h for high. Unstable nodes and foci are denoted by the subscript u. Furthermore the subscript s denotes a switching point, where the control constraint becomes active.

3.3 Base Case Parameters

For the discount rate we use a typical value of $r = 0.05$. The outflow rate μ is assumed to be 5%, which is a typical rate assumed for high-rate criminal offenders (e.g., Greenwood et al. (1994)). That is substantially higher than the natural, peacetime mortality rate for young adults due to arrest, death by suicide attack, etc., but reflects a presumption that in the absence of active counter-terror measures terrorist organizations do not suffer high attrition.

For reasons discussed above, we assume that the constant inflow rate term τ is small. The specific value matters little and we pick a value of $\tau = 10^{-5}$. Having τ be positive, not zero, is a mathematical convenience; it creates a single well-defined low level equilibrium instead of having the entire vertical axis ($x = 0$) in the state-costate space be a steady state. Increasing or decreasing τ by a factor of 10 or even 100 has minimal effect on the analysis.

Parameter k is chosen such that the high steady state x_h^i is normalized to 1 in the absence of counter-terror operations (and assuming that $\tau = 0$ is negligible). That is, x will be measured as a percentage of steady-state size of the terrorist organization when its opponents do not mount counter-terror operations. Since the uncontrolled dynamic is given by

$$\dot{x} = kx^{\alpha} - \mu x$$

the high steady state $x_h^i = \sqrt[\alpha-1]{\frac{\mu}{k}}$, so normalizing x implies that $k = \mu = 0.05$.

In principle the exponent α in the initiation term could take on any value between 0 and 1, but we prefer larger values (emphasizing the potential for existing terrorists to recruit new terrorists). In the absence of relevant empirical data, we choose a basecase value of $\alpha = 0.75$, but explore the influence of this parameter over its entire range in Sec. 5.2.

The efficiency parameters β and ν cannot be derived directly from evaluations of actual interventions. Instead, we choose values for them and the associated parameter θ such that the fire and water interventions reflect the spirit of the concepts we are trying to investigate. In particular, we imagine that the following statements are true. (1) When terrorist organizations have grown to their maximum size (normalized here to $x = 1$) fire strategies are substantially (specifically ten times) more powerful than are water strategies, so $\nu/\beta = 10$. (2) A significant investment in terror-control operations (on the order of $10B per year) can kill or otherwise remove a number of terrorists each year that is equal to 10% of the organization's maximum size. Hence, ν and β should be on the order of $0.1/\ln(1+\sqrt{10})$[1]. (3) When terrorist organizations are small, water strategies are relatively more effective, but the breakeven point where fire and water strategies are equally cost-effective is much less than half the maximum organizational size. Hence, we choose θ such that

$$\nu/\beta x^{1-\theta} = 1 \tag{12}$$

for a fairly small x, specifically $\theta = 0.1$, for which equality holds when $x \approx 1/13$. Given that $\theta = 0.1$, choosing $\beta = 0.01$ and $\nu = 0.1$ meet our other criteria.

The parameter ρ measures the influence of fire interventions on Recruitment. It is set to 1 in the base case since when fire tactics are employed in the optimal solution, v is usually in the vicinity of 0.2 or 0.3. This means that those counter-terror operations stimulate a 20% to 30% increase in recruitment. That is large enough to make the chief dynamic investigated here meaningful, without being so large as to "rig" the results against using fire tactics. As a point of comparison, Kaplan et al. (2005) empirically-calibrated hit-dependent recruitment model suggests that hits could have accounted for nearly all of the increases in the stock of terrorists confronting Israel. The impact on the optimal solution of changing ρ will be analyzed in Sec. 4.4.

The costs c per terrorist is assumed to be 1 in the base case, and the consequences for the optimal strategies on changing these costs are analyzed in Sec. 5.1. With that value and the observed intensities of counter-terror operations we get costs of terror and terror-control that are of sensible orders of magnitude relative to each other. In the low-level equilibrium when terrorists have been nearly annihilated counter-terror operations generate more costs than do terror attacks themselves, but at the high-level equilibrium the opposite is true and terror attacks generate societal costs that are roughly twenty times the amount invested in (offensive) counter-terror operations.

r	α	β	c	k	μ	ν	ρ	θ	τ
0.05	0.75	0.01	1	0.05	0.05	0.1	1	0.1	10^{-5}

Table 1: The specified parameter values for the base case.

[1]The $10B per year figure is an average over five years spending on operations in Afghanistan. The 10% figure is based on estimates that Al-Qaeda's pre-war strength in Afghanistan was 3,000 - 4,000 fighters, and a guess that US operations eliminate about 300-400 Al-Qaeda personnel per year.

4 Optimal Strategies

4.1 Base case

Numerical calculations reveal there is one high-level optimal steady state $x_h^i \approx 0.62$ and one low-level optimal steady state $x_l^a \approx 8 \times 10^{-7}$. The high steady state x_h^i is derived from the SIC version of the model, so fire strategies are employed $v > 0$. The low steady state x_l^a is derived from the SAC version in which the control constraint is active ($v = 0$) so only water control is used to keep the number of terrorists at the low steady state.

Calculating optimal paths leading to the high and low steady states reveals two facts. First there exists a switching point $x_s \approx 0.06$ to the left of which it is suboptimal to use fire interventions. Further to the left there is a second state (at $x_D \approx 0.013$) from which there exist two strategies that are optimal inasmuch as both yield the same over-all objective function value. This state is a so called Dechert, Nishimura, Skiba or DNS point referring to Skiba (1978) and Dechert and Nishimura (1983), where the decision maker is indifferent about which strategy to chose. Moving to the left from x_D by employing a high level of water control $u_D \approx 0.29$ leads to the low steady state and virtual extinction of the terrorist organization (see Fig. 2d). Starting from the same initial point but using only a moderate level of water control allows the stock of terrorists to grow, eventually leading in the long run to the high steady state (see Fig. 2b).

It might seem odd that if it may not be optimal to all but eradicate a terrorist organization when it is possible to do so. However, if the terrorist organization is already large when control begins, eradication can take enormous effort. While it is true that eradication would lead to fewer terrorists and, hence, fewer terror-related losses, at some point those gains can be offset by the increased cost of counter-terror efforts. The DNS point is exactly that point where the gains of eradication are offset by the additional control costs.

If control begins to the left of the DNS point, before the terrorist organization has grown too large, then it is optimal to take actions (specifically water control) to nearly eradicate the terrorists. If the initial number of terrorists exceeds this state x_D, it is still optimal to conduct counter-terror operations, but only to slow not ultimately to prevent growth of the terrorist organization to x_h^i. When the terror organizations are only modestly larger than x_D, only water controls are used, but once x grows beyond x_s fire strategies are also used. Indeed, eventually fire is used more extensively than water since the θ exponent makes fire control more efficient when the number of terrorists is large.

The next few subsections explore the sensitivity of this optimal strategy to three of the most important parameters: the efficiency of fire interventions ν, the efficiency of water interventions β, and the extent to which using fire strategies stimulates greater recruitment of terrorists ρ. The general structure of this solution turns out to be fairly robust. The details depend on the various parameter values, of course, but in quite sensible and predictable ways.

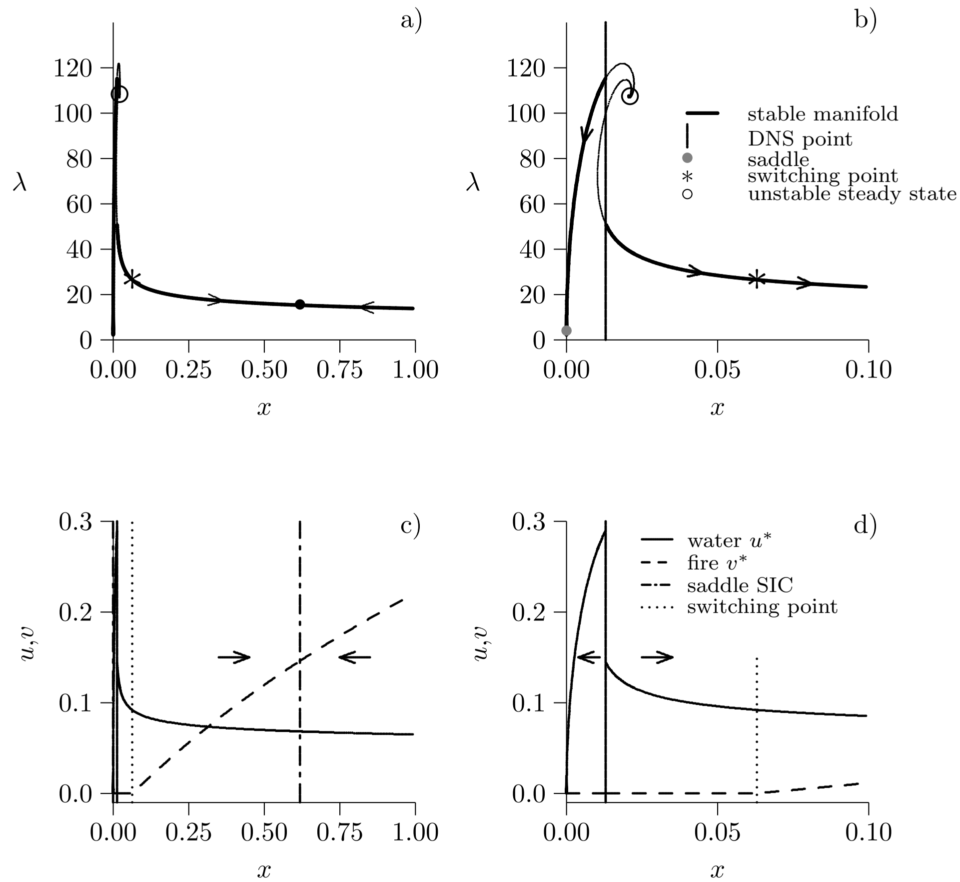

Figure 2: The optimal strategy for the base case is depicted. While in a) and b) the optimal solution is displayed in the state-costate space, in c) and d) the corresponding usage of fire and water interventions are shown. Furthermore b) and d) reveal the details of the solutions behavior near the origin. There the DNS point x_D can be seen and the gap of the water intervention at this point of indifference.

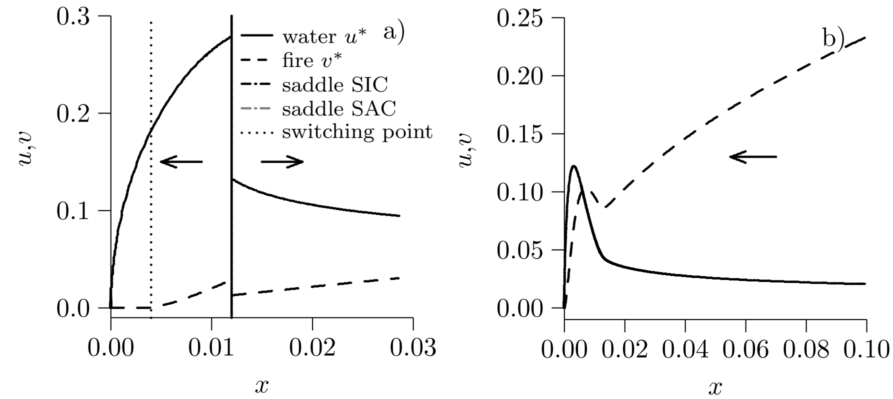

Figure 3: The optimal strategies for large efficiency values ν. Panel a) shows that with $\nu = 0.2$ the DNS point x_D is to the right of the switching point x_s. Panel b) shows there is a unique, optimal (low level) long run steady state x_l^a for $\nu = 1.2$.

4.2 Optimal strategy for increasing fire efficiency ν

In the base case, fire control was not used at the DNS point or any other $x < x_s \approx 0.06$, so those results would not be affected at all if fire control were less effective (smaller ν). However, reducing ν shifts to the right both the switching point x_s (where fire is first used) and the high-level equilibrium; when fire is weak, the range of circumstances for which fire is appropriate and the maximum number of terrorists that it is optimal to eradicate both shrink. If ν is small enough (roughly around $\nu = 0.05$), the switching point moves to the right of the long-run equilibrium, implying that fire control should never be used, unless for some reason the terrorist organization were initially larger than its maximum peacetime size.

The solution can change more substantially if fire interventions become more effective than in the base case $\nu > 0.1$. Increasing ν naturally shifts to the left both the high steady state and the switching point x_s. When ν is increased to approximately 0.16, the switching point moves to the left of the DNS threshold, so it is optimal to use fire as well as water when starting at that critical point from which it is optimal either to eradicate or to accommodate terrorists in the long run. (With more effective fire control, one does not have to accommodate to as many terrorists, since the high-level equilibrium shifts significantly to the left.) This is displayed for $\nu = 0.2$ by Fig. 3a).

If ν is increased much more, to be about twelve times the base case value, the optimal strategy does change dramatically. The high steady state of the SIC is no longer an optimal solution (see Fig. 3b); fire control together with water control are now sufficiently efficient that it is always optimal to drive the number of terrorists down to the low steady state x_l^a. In such circumstances, it would be optimal to use more fire than water initially, but once the stock of terrorists has been sufficiently reduced, water is used more than fire in the end game.

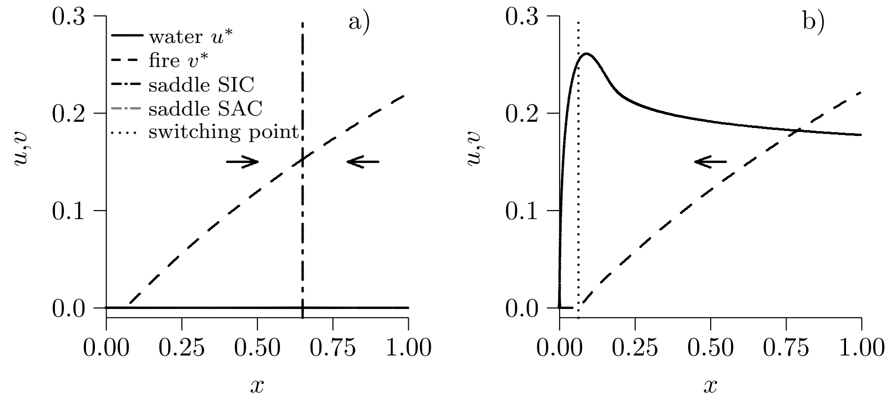

Figure 4: The optimal strategies for increasing β are depicted. Sub-figure a) shows the "extreme" case $\beta = 0$ where a high number of terrorists is optimal in the long run. Contrary to that sub-figure b) represents the case of a high water efficiency with $\beta = 0.03$, where nearly annihilating the terrorists is optimal in the long run.

4.3 Optimal strategy for increasing water efficiency β

Next consider the effects of varying the efficiency of water control β. In the extreme, if water control has no effect $\beta = 0$ then it is of course not optimal to use water interventions. This means the high steady state $x_h^i \approx 0.65$ is always optimal, as can be derived from Eq. 2 and seen in Fig. 4a). There is still a switching point $x_s \approx 0.06$ above which fire control is used, but it only slows not prevents growth toward the high level equilibrium x_h^i.

Water control's efficiency does not need to be very much above zero (specifically $\beta > 3.5 \times 10^{-4}$) in order for the low steady state to emerge as optimal provided that the initial number of terrorists is small enough. (I.e., for the solution to be structurally the same optimal strategy as described for the base case.) For example if β is $\beta = 0.001$ the low and high optimal long run steady states are $x_l^a \approx 1.1 \times 10^{-6}$ and $x_h^i \approx 0.65$, respectively. The switching point will be close to its base value $x_s \approx 0.06$, but the DNS point is very small if β is small $x_D = 5 \times 10^{-4}$ because with low water efficiency, it would be quite expensive to all but eradicate terrorists unless there were initially very few of them.

Increasing β shifts the DNS point to the right and the high-level steady state to the left. Increasing β beyond about double its base case value $\beta \approx 0.02$ moves the DNS point to the right of the switching point, which remains constant at $x_s \approx 0.06$. Finally if β is more than three times higher than in the base case, water becomes so efficient that the DNS point disappears and the low-level steady state (effectively annihilating the terrorists) becomes the globally optimal strategy (see Fig. 4b).

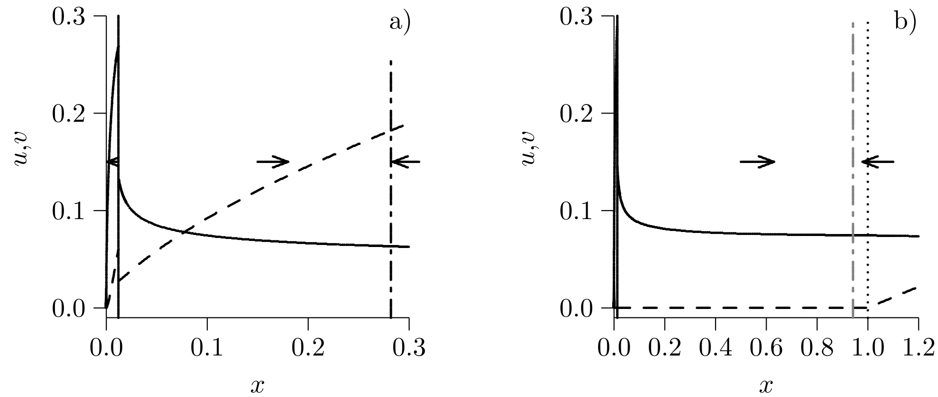

Figure 5: The optimal strategies for no influence of the fire interventions on recruitment a) $\rho = 0$ and doubled influence b) $\rho = 2$ are depicted.

4.4 Optimal strategy for varying influence of fire interventions ρ on recruitment

If there were no backlash from fire interventions on recruitment $\rho = 0$, then common sense suggests and Eq. 3 confirms it is always optimal to use fire interventions. As can be seen in Fig. 5a), there would still be two optimal long run states: a low steady state with $x_l^i \approx 8 \times 10^{-7}$ and a high steady state with $x_h^i \approx 0.28$. Furthermore, the two steady states are still separated by a DNS point $x_D \approx 0.012$. So the main changes are that fire control is always used and the high level steady state has less than half as many terrorists as in the base case. Introducing backlash by increasing ρ from zero naturally reduces the amount of fire control used, so the high long-run number of terrorists increases. The DNS point does as well, but much less dramatically. While increasing ρ from zero to one leads to a 120% increase in the high long-run number of terrorists, (compare Fig. 5a) and Fig. 2a)), the DNS point only increases by 6%.

If ρ increases to above about 0.06 (6% of the base case value of 1.0), fire control is no longer used near the low-level steady state. In other words, if there is even a modest backlash effect, then fire strategies should be curtailed before the terrorists have been fully subdued. If the backlash parameter ρ is bigger than about 2 (double its base case value), the switching point moves to the right of even the high optimal steady state, and fire controls are essentially always counterproductive (see Fig. 5b).

5 Optimal Strategies for other parameters

5.1 Optimal strategy for different costs c

In the base case it is assumed that $c = 1$. This assumption is more or less arbitrarily we next consider the influence of the costs c on the optimal strategy.

Increasing c by a factor of five does not alter the structure of the solution. The low level base case and switching point are essentially unaffected. The DNS point is about four times larger $x_D \approx 0.05$, and the high level steady state is about two thirds as large $x_h^i \approx 0.4$. (See Fig. 6a). That makes sense. If terrorists are more costly, one is willing to work harder to eradicate them even if there initial numbers are a bit larger, and even if it is optimal to accommodate an ongoing terrorism presence, the level tolerated is reduced. To achieve these changes, effort expended on both fire and control strategies is more than doubled. (Compare Fig. 2c and Fig. 6a).

If the costs are increased above $c \approx 13.7$ it is no longer ever optimal to let the long run number of terrorists be high. Instead, it is always optimal to approach the low steady state. The switching point remains at $x_s \approx 0.06$, but as could be expected, the intensity of fire and particularly water controls increases further as compared to the base case.

5.2 Optimal strategy for different α

Parameter α models the extent to which terrorism recruitment is "contagious" in the sense of recruitment being an increasing function of the current number of terrorists. With smaller levels of α, the inflow of new terrorists is closer to a constant and can be large even when the stock of existing terrorists is relatively small. With large levels of α, recruitment grows more nearly linearly in the current number of terrorists and so can get quite small as the number of terrorists is driven down.

Not surprisingly then, if α is small enough, it is never optimal to drive terrorism down to the low level steady state because recruitment would remain high and, hence, so would the ongoing level of terror control effort needed to preserve that low-level steady state. As numerical calculations reveal the low level steady state becomes a relevant optimal solution not until $\alpha > 0.3$.

On the other hand, if α is large enough, it is relatively easy to keep terrorism from springing back once it is eradicated, so it is always optimal to eradicate terrorism, not matter what the starting number of terrorists is. What is perhaps somewhat more surprising is that it can be optimal (e.g., for $\alpha = 1$ to use both fire and water control at all times; there is no switching point.

5.3 Optimal strategy for different θ

Parameter θ models the interaction between the current number of terrorists and the effectiveness of water interventions. As is explained in Sec. 3.3, θ, β and ν are intimately connected via Eq. 12. Keeping ν and the calculated $x = 1/13$ for the base case fixed, we vary θ and β simultaneously under assumption of Eq. 12.

The solution is not affected much at all by reductions in the values of θ and β. Even for $\theta = 0$ and hence $\beta \approx 0.0077$ the primary difference is just that the DNS point shifts right to $x_D \approx 0.02$. That makes sense. The small θ means

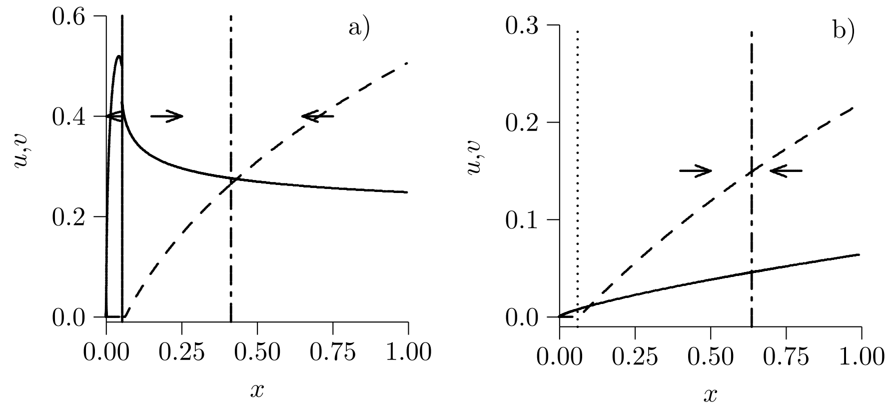

Figure 6: The optimal strategies for a) different costs per terrorist $c = 5$ and b) $\theta = 1$ and $\beta = 0.1$ are depicted.

difficulties locating targets when there are very few terrorists are minor, and that expands the range of initial values for which one should eradicate terror.

On the other hand, increasing θ and hence β reduces the ability to find targets, making it less likely that eradication is optimal. Indeed, if θ is more than about 0.42 it is never optimal to eradicate the terrorists. Fig. 6b) illustrates that for the extreme case of $\theta = 1$, the only long run optimal number of terrorists is $x_h^i \approx 0.29$.

6 Conclusions

A number of authors have raised the possibility that counter-terror operations might stimulate the flow of newly recruited terrorists at the same time that they are removing some of the existing terrorists. Not all counter-terror tactics are equally likely to provoke increased recruitment, so we introduce a simple stylized model with two types of counter-terror operations: fire tactics that stimulate increased recruitment and water tactics that do not. Water tactics are intrinsically more limited in their ability to eliminate large numbers of terrorists, e. g., because they depend on specific intelligence to target strikes in ways that generate minimal collateral damage.

The basic structure of the model solution is quite robust with respect to parameter values, which is helpful since estimating parameter values is very difficult. Generally there are two different steady states. One has very few terrorists and, hence, little recruitment of new terrorists. The other has many more terrorists, sometimes approaching the steady state number when there are no counter-terror operations. We found three possible policy prescriptions, depending on the parameter values: (1) always drive the number of terrorists down to the low-level steady state, (2) always allow the number of terrorists to approach the high-level steady state, albeit using controls to slow the growth, or (3) drive down the number of terrorists if they are not too numerous initially, but otherwise allow them to grow to the higher level equilibrium.

The third case pertains with our base case parameters, and the state value dividing the regions for which it is optimal to approach the low- and high-level equilibria is a so-called DNS (Dechert, Nishimura, Skiba) point.

In terms of tactics, it is essentially always optimal to use the targeted (water) tactics, but fire tactics should only be employed if the number of terrorists is above a critical "switching point". Hence, for many parameter sets, eradicating terrorists should involve greater reliance on fire at first, with water tactics being used predominantly if not exclusively after the initial assault has brought down the number of terrorists. For the base case parameters that switching point is to the right of the DNS point, so moving away from the DNS point only involves using water controls in great volume to drive terrorists down to the low-level steady state, less if one is merely delaying growth to the high-level steady state.

The location of that switching point - and, hence, the circumstances under which fire controls should be used - is moderately sensitive to changes in parameter values. If it moves sufficiently far to the left, it can always be optimal to use fire tactics, although the intensity of fire tactics still seems to always be increasing in the number of terrorists. Conversely, if it moves sufficiently far to the right it may not be optimal to use fire tactics unless the number of terrorists is very large.

Analyzing this simple model may not seem very practical given the complexities of managing counter-terror operations, but precisely because of the model's simplicity, a fortiori a number of important conclusions emerge.

First, if indeed, as some suppose, counter-terror operations can stimulate recruitment of new terrorists, that can have important implications for how the terror war should be fought.

Second, even if some counter-terror operations do stimulate substantial new recruitment and other available tactics do not, that does not necessarily imply that tactics which can stimulate recruitment should never be used. Their benefits may exceed their downside if they are sufficiently effective.

Third, if some counter-terror tactics are more likely than others to be perceived of as outrageous or otherwise to stimulate relatively more new terrorist recruitment, then the optimal mix of terror control operations should vary, perhaps dynamically over time with the changing state of the terrorists.

Fourth, if recruitment is concave in the number of terrorists and the effectiveness of counter-terror tactics proportional to the number of targets, then the intensity with which those recruitment-stimulating tactics should be employed will generally be increasing in the current number of terrorists. So if the decision is made to eradicate the terrorists, the extent to which fire tactics are used should decrease over time. Water strategies should always be used and are relied on most heavily when the number of terrorists is small.

References

Castillo-Chavez, C., & Song, B. (2003). Models of the transmission dynamics of fanatic behaviors. In H. Banks & C. Castillo-Chavez (Eds.), *Bioterrorism: Mathematical modeling applications in homeland security* (Vol. 28, p. 155-172). Philadelphia: SIAM.

Caulkins, J. P. (2005). Models pertaining to how drug policy should vary over the course of a drug epidemic. In *Substance use: Individual behavior, social interactions, markets, and politics* (Vol. 16, p. 407-439). Elsevier.

Dechert, W. D., & Nishimura, K. (1983). A complete characterization of optimal growth paths in an aggregated model with a non-concave production function. *Journal of Economic Theory*, *31*, 332–354.

Doedel, E. J., Govaerts, W., & Kuznetsov, Y. A. (2003). MATCONT: A Matlab package for numerical bifurcation analysis of ODEs. *ACM Transactions on Mathematical Software*, *29*(2), 141-164.

Feichtinger, G., & Hartl, R. F. (1986). *Optimale Kontrolle ökonomischer Prozesse*. Walter de Gruyter.

Greenwood, P. W., Rydell, C. P., Abrahamse, A. F., Caulkins, J. P., Chiesa, J., Model, K. E., et al. (1994). *Three strikes and you're out: Estimated benefits and costs of california's new mandatory sentencing law* (Tech. Rep. No. MR-509-RC). Santa Monica, CA: RAND.

Heymann, P. (2003). Dealing with terrorism after september 11, 2001: An overview. In A. Kowitt & R. Pangi (Eds.), *Countering terrorism: Dimensions of preparedness* (p. 57-72). Cambridge, MA: MIT Press.

Kaplan, E. H., Mintz, A., Mishal, S., & Samban, C. (2005). What happened to suicide bombings in israel? Insights from a terror stock model. *Studies in Conflict and Terrorism*, *28*(3), 225-235.

Keohane, N. O., & Zeckhauser, R. J. (2003). The ecology of terror defense. *Journal of Risk and Uncertainty*, *26*(2-3), 201-29. (available at http://ideas.repec.org/a/kap/jrisku/v26y2003i2-3p201-29.html)

Leonard, D., & Long, N. V. (1992). *Optimal control theory and static optimization in economics*. Cambridge University Press.

Skiba, A. K. (1978). Optimal growth with a convex-concave production function. *Econometrica*, *46*, 527–539.