TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

Operations Research and Control Systems

SWM ORCOS

Optimal Control of a Terror Queue

Andrea Seidl, Edward H. Kaplan, Jonathan P. Caulkins, Stefan Wrzaczek, Gustav Feichtinger

Research Report 2015-12

September, 2015

Operations Research and Control Systems
Institute of Statistics and Mathematical Methods in Economics
Vienna University of Technology

Research Unit ORCOS
Wiedner Hauptstraße 8 / E105-4
1040 Vienna, Austria
E-mail: orcos@tuwien.ac.at

Optimal Control of a Terror Queue

Andrea Seidl Edward H. Kaplan Jonathan P. Caulkins
Stefan Wrzaczek Gustav Feichtinger

Abstract

The task of covert intelligence agents is to detect and interdict terror plots. Kaplan (2010) treats terror plots as customers and intelligence agents as servers in a queuing model. We extend Kaplan's insight to a dynamic model that analyzes the inter-temporal trade-off between damage caused by terror attacks and prevention costs to address the question of how many agents to optimally assign to such counter-terror measures. We compare scenarios which differ with respect to the extent of the initial terror threat and study the qualitative robustness of the optimal solution. We show that in general, the optimal number of agents is not simply proportional to the number of undetected plots. We also show that while it is sensible to deploy many agents when terrorists are moderately efficient in their ability to mount attacks, relatively few agents should be deployed if terrorists are inefficient (giving agents many opportunities for detection), or if terrorists are highly efficient (in which case agents become relatively ineffective). Furthermore, we analyze the implications of a policy that constraints the number of successful terror attacks to never increase. We find that the inclusion of a constraint preventing one of the state variables to grow leads to a continuum of steady states, some which are much more costly to society than the more forward-looking optimal policy that temporarily allows the number of terror attacks to increase.

1 Introduction

Since the terror attacks of September 11, 2001, a growing literature has emerged on the efficient design of counterterror polices. Using the terrorists' resources as state variable, Keohane and Zeckhauser (2003) studied the conditions under which a terrorist organization might be eradicated. Dealing with the question of how to optimally prosecute terrorists, Caulkins et al. (2008) compare what they call 'fire' and 'water strategies'. Further studies on the eradication of terrorists have been given by Kress and Szechtman (2009); Caulkins et al. (2009), Kaplan et al. (2010). More recently game-theoretic aspects have been considered by Behrens et al. (2007); Zhuang and Bier (2007); Feichtinger and Novak (2008); Zhuang et al. (2010); Crettez and Hayek (2014); see also Grass et al. (2008).

A weakness of most of these papers is the assumption that the government can observe the terrorists' state variable's value; that is, the model's solution procedures tacitly assume that government knows the size and strength of the terrorists. However, terrorists are not like conventional armies that can be observed with satellites or other forms of reconnaissance. They challenge authorities in no small part because they operate in small cells that are difficult to detect.

Kaplan (2010) provides a way around this problem. He introduced a new approach to estimate the numbers of terror threats in a given area that have not yet been detected. Then adding the number that have already been detected, which of course is known to the authorities, gives the size of the terrorists' state variable.

Kaplan's innovation was interpreting terror plots as customers and intelligence agents as servers in a queuing model to predict the rate at which such threats can be detected and interdicted. New terror plots (or cases) arrive with a Poisson process, are detected by intelligence agents and enter "service" – meaning the intelligence agents infiltrate and destroy the plot, thereby removing it from the queue. Sometimes terror plots evade detection and complete a terror attack. In terms of a queuing model this is comparable to customers who renege from the queue before being served. The queuing analogy is, however, not exact. Unlike customers in conventional queues, terror plots are not visible upon arrival, but must be discovered before service, i.e. counter measures, can begin. Thus, waiting customers (i.e. undetected plots) and available agents may coexist.

Kaplan (2010) developed a Markovian queuing model for the detection and interdiction of randomly arriving terror plots. The birth-and-death process approach well-known in queuing theory (see, e.g., Hillier and Lieberman 2001) gives rise to balance equations for the state probabilities of a two-dimensional Markovian queuing process. While this equilibrium analysis is restricted to the steady state, the transient case is covered by a deterministic fluid approximation (see, e.g., Newell 1971) that results in a system of two non-linear, ordinary differential equations. The key term describes the rate at which unknown plots are detected and become known.

Originally, the terror queue ansatz was purely descriptive. The aim was to understand the infiltration and interdiction of ongoing terror plots by intelligence agents. In a follow-up papers, however, Kaplan (2013, 2015) started to include optimization with respect to staffing.[1] In particular, he calculated the number of agents that maximizes the benefits-minus-costs of preventing attacks. Moreover, by presuming that terrorists are smart and will infer the staffing level of counter terrorism agencies by observing the fraction of attacks interdicted, Kaplan (2013, 2015) even investigates a simple terror queue staffing game.

Some might find it odd to balance the heroic benefits of saving lives by preventing murderous terrorist attacks with grubby considerations of the costs of control efforts. Such a seemingly cold-hearted calculus can be defended on at least two grounds. First, there are many domains of public policy that balance saving lives against dollars, implicitly if not explicitly, even in such mundane areas as deciding how much to spend on highway improvements that prevent fatal traffic crashes. Second, zealous terror control could require many counterterror agents and their attendant costs, not to mention the accompanying inconvenience, intrusion and loss of privacy such control might entail.

Kaplan's cost-benefit analysis was static and restricted to the long-run steady case. Although such an equilibrium analysis provides important insights, a more complete analysis has to consider the intertemporal structure of the terror plots, their detection and interdiction.

The present paper addresses these dynamic aspects by applying optimal control theory. The tool we will use to study the intertemporal staffing problem is Pontryagin's maximum principle. It provides a useful method for understanding the qualitative behavior of the system. Such an approach is particularly useful when, as here, reliable data

[1] Operations researchers have long studied staffing and manpower planning problems, see Bartholomew (1973).

are scarce, thwarting more empirical or statistical approaches.[2]

We show that the optimal strategy for the government depends on both the number of known and unknown terror plots; as those state variables change dynamically over time, so too should the government strategy evolve. Some results are predictable. When the damage caused by terror is not particularly large, it is optimal from a monetary perspective to accept certain casualties rather than to inefficiently search for further terror plots. Likewise, when there are many terror plots, the government should do much to prevent their success. And the division of agents between the two complementary tasks of detection and interdiction depends on how efficiently these agents act in these different roles. Other results could not be anticipated so easily a priori. For example we see that if terrorists are more successful, the government may reduce not increase the long-run number of intelligence agents. Furthermore we study the impact of the interdiction rate on the optimal long run solution and find that if detection agents are less efficient, more of them are required to successfully prevent terror attacks. We also see that the long-run number of terror plots increases when terrorists are more active in the sense that the inflow to known terror plots is higher.

A strategy that involves – even temporarily – an increase in the number of terror attacks over time might not be acceptable politically, even if it were the optimal way to reduce the total number of attacks over time. Thus, we compare the base case outcome to a scenario in which the decision maker insists that the number of terror plots does not ever rise. This is interesting also from a methodological point of view. Such a constraint means that one of the state equations must always be non-positive. As a result, we can find a manifold of steady states in an area of the state space where otherwise the number of unknown terror plots would increase if this constraint is not taken into account. We find that it can be costly to impose such a constraint, particularly when the initial number of unknown terror plots is low, because it is rather difficult to detect them then. When there is an intermediate initial number of unknown terror plots, imposing the constraint lowers the incentive of a decision maker to put effort into terror detection. While it certainly would be beneficial to temporarily lower the number of attacks, keeping them low would require more agents than would keeping terror plots at an intermediate level.

The paper is structured as follows. Section 2 explains the model in detail. In Section 3 we derive the necessary conditions for optimality for the basic model, and study the numerical results with respect to their robustness. Section 4 analyzes the implications of a constraint that the number of terror attacks must never increase. Section 5 concludes.

2 The Model

The model is a two-state diffusion model, compare e.g. Rogers (2003), that describes the dynamics of the number of undetected, $X(t)$, and detected terror plots, $Y(t)$, respectively. The control variable, $f(t)$, denotes the number of undercover intelligence agents deployed by the government to detect and infiltrate terror plots. (In the sequel we omit the time argument t unless necessary.) The agents, f, are divided between interdictors (of whom there are Y, one for each detected plot) and detectors (everyone else, namely $f - Y$). Interdiction in this model is thus 1-to-1; as in a queuing model, for every customer in

[2]Note that this situation may be compared with those in another field of "deviant behavior", i.e. in the dynamics and control of illicit drug consumption. Tragler et al. (2001) and Behrens et al. (2000) are good examples for the usefulness of optimal control methods in the "economics of crime" (more general economics of deviant behavior).

service, there is a busy server. Based on Kaplan (2010) the system dynamics are given by

$$\dot{X} = \alpha - \mu X - \delta(f - Y)X, \tag{1}$$
$$\dot{Y} = \delta(f - Y)X - \rho Y, \tag{2}$$

where α is the arrival rate of new terror plots. The rate at which undetected terror plots lead to successful terror attacks is denoted by μ. Parameter δ governs the efficiency of intelligence agents with respect to successful plot detection and ρ is the interdiction rate. These two parameters, i.e. δ and ρ, can be thought of as being technology-dependent and thus not controllable, so it is not possible in a simple manner to have them change continuously over time. For example, increasing δ might come about from having better code-breaking algorithms that catch more communications, or better surveillance technologies. Those rates do not share the same flexibility as staffing, and it is also not always obvious how a new technology translates into actual detection/interdiction, so we assume them to be fixed. We have to take the control constraint

$$f \geq Y \geq 0, \tag{3}$$

into consideration so that the flow from X to Y is non-negative. This effectively means that agents are not allowed to drop or ignore a plot that they already know about.

The optimization criterion is to minimize the discounted stream of damages created by successful terror attacks plus the costs of counterterror efforts. The objective is

$$V = \min_f \int_0^\infty e^{-rt} \left[c\mu X + \frac{c_f}{2} f^2 \right] dt, \tag{4}$$

where c denotes the monetary damage of a successful terror plot and c_f scales the quadratic costs of staffing.

Detected plots rarely manage to execute successful attacks both because the terror interdiction rate ρ is so large that detection usually leads to capture fairly quickly and because the undercover agents may literally "take control" of the plots – by infiltrating and becoming part of the plots. Thus, we assume that successful terror attacks only result from unknown terror plots. Accordingly, the number of busy agents Y only enters the objective function indirectly via the optimal total number of agents f.

3 Results

We first analytically derive the necessary optimality conditions and then use numerical methods to find the optimal solution for different initial state values and to study the impact of parameter changes on the solution.

3.1 Analytical Results

The Hamiltonian is

$$\mathcal{H} = c\mu X + \frac{c_f}{2} f^2 + \lambda_1 \left(\alpha - \mu X - \delta(f - Y)X \right) + \lambda_2 \left(\delta(f - Y)X - \rho Y \right),$$

where λ_i, $i = 1, 2$ are the costate variables measuring the marginal costs of the state variables X and Y, respectively. We expect $\lambda_i \geq 0$ as terror plots are harmful, and show below that this is indeed the case. The Lagrangian is

$$\mathcal{L} = \mathcal{H} + \omega\,(f - Y)\,,$$

where ω denotes the Lagrange multiplier associated with constraint (3). Applying Pontryagin's Maximum Principle (see, e.g., Grass et al. 2008) we find

$$\mathcal{L}_f = c_f f + (\lambda_2 - \lambda_1)\,\delta X + \omega = 0. \qquad \Rightarrow \qquad f = \frac{(\lambda_1 - \lambda_2)\,\delta X - \omega}{c_f}. \tag{5}$$

Note that in the interior of the admissible region (i.e. $f \geq Y$ and $\omega = 0$), the above condition can also be written as $c_f f = (\lambda_1 - \lambda_2)\,\delta X$. This means that marginal cost of adding the fth agent equals the marginal benefit of detecting plots due to adding the fth agent, where the marginal benefit is the detection rate added by the last agent times the benefit of shifting plots from undetected to detected status.

The optimal control is non-negative if $\lambda_1 \geq \lambda_2$, i.e. when known terror plots are preferable to unknown terror plots. The Legendre-Clebsch condition is fulfilled as $\mathcal{L}_{ff} = c_f > 0$. The complementary slackness condition $\omega\,(f - Y) = 0$ has to hold. If the constraint $f \geq Y$ is binding, the Lagrange multiplier is

$$\omega = -c_f Y - (\lambda_2 - \lambda_1)\,\delta X.$$

Furthermore, it must be that $\omega \leq 0$.

The costate equations are

$$\dot{\lambda}_1 = (r + \mu)\,\lambda_1 + (\lambda_1 - \lambda_2)\,\delta(f - Y) - c\mu, \tag{6}$$
$$\dot{\lambda}_2 = (r + \rho)\,\lambda_2 - (\lambda_1 - \lambda_2)\,\delta X + \omega. \tag{7}$$

In a steady state $(\hat{X}, \hat{Y}, \hat{\lambda}_1, \hat{\lambda}_2, \hat{f}, \hat{\omega})$ it must hold that

$$\alpha = \mu\hat{X} + \rho\hat{Y},$$

which means that the rate at which new terror plots arrive equals the attack rate plus the interdiction rate. This implies that the steady state number of detected terror plots increases with the arrival rate of terror plots, and decreases with the number of successful terror attacks. Furthermore, for an interior control $\lambda_1 \geq \lambda_2$. Considering (7), this implies that in this case the steady state value of λ_2, and consequently that of λ_1, is positive.

If the constraint is active, i.e. $f = Y$, there is one steady state at

$$\hat{X} = \frac{\alpha}{\mu}, \hat{Y} = 0, \hat{\lambda}_1 = \frac{c\mu}{r + \mu}, \hat{\lambda}_2 = 0, \hat{f} = 0, \hat{\omega} = \frac{\alpha c\delta}{r + \mu}.$$

Due to the positivity of the parameters, the Lagrange multiplier $\hat{\omega}$ is always positive. As such this steady state can never be admissible, meaning that due to the harmfulness of terror attacks it is always better to make some terror detection and interdiction efforts as opposed to doing absolutely nothing.

3.2 Numerical Results

Since most of our results cannot be derived analytically, we resort to numerical calculations. We require values for seven parameters: α (the terror plot arrival rate), μ (the plot completion rate), ρ (the interdiction rate for detected plots), δ (the detection rate for undetected plots), c (the cost of a successful terror attack), c_f (the constant scaling of the quadratic costs of counterterror agents), and r (the social discount rate).

Beginning with α, the Global Terrorism Database (GTD) reports values for the number of terror attacks around the world in addition to deaths and injuries (National Consortium for the Study of Terrorism and Responses to Terrorism (START) 2013). Between 2002-2011, this database reports a total of 80 terror attacks in the United States, or 8 per year. However, these figures ignore plots that were detected and interdicted. Kaplan (2012) reports that 28 of 35 Jihadi plots in the United States were detected between September 11, 2001 and June 30, 2011, while Strom et al. (2010) estimated that 80% of terror plots in the United States occurring between 1999-2009 were detected and stopped prior to attack. This suggests that observed attacks represent only 20% of new plots, thus we set α to $8/0.2 = 40$ newly arriving plots per year.

Following Kaplan (2010), the parameters μ and ρ are assigned the values 1/yr and 4/yr respectively. These values imply that the average time from the inception of a plot until a terror attack equals one year (and in steady state the annual number of successful attacks will equal the number of undetected terror plots $\hat{X}$), while the average time to interdict detected plots equals 3 months.

Kaplan (2015) notes that there are about 2,000 FBI special agents devoted to counterterrorism, and shows that this staffing level is consistent with 80% plot detection if the ratio δ/μ approximately equals 0.002. Having set $\mu = 1$ above, we assign $\delta = 0.002$ to be consistent.

To estimate the cost of a terror attack, Kaplan (2015) uses the GTD's reported numbers of deaths and injuries from terror attacks in Europe and in Israel between 2002-2011 to arrive at an estimated 5.6 statistical deaths per successful terror attack. The US Department of Homeland Security (Appelbaum 2011), based on earlier work by Viscusi and Aldy (2003), values lives at about \$10 million, so for round numbers, we set the cost of a terror attack to \$50 million; note that this is restricted to human casualties and does not consider property or other infrastructure costs and is thus quite conservative.

To estimate c_f, FBI special agents are paid between \$60,000 to \$70,000 per year.[3] Accounting for benefits and other expenses as well as the fact that embedded agents are supported by headquarters or field office staff, we increase the average cost per agent to \$150,000. With quadratic staffing costs and presuming 2,000 agents, an average cost of \$150,000 per agent results when $c_f = 150$.

Finally, we set the social discount rate $r = 0.04$, though later we will show that our results are robust to the specific discount rate assumed. To summarize, the parameter values we employ in our numerical examples are:

$$\alpha = 40, \quad \mu = 1, \quad \rho = 4, \quad \delta = 0.002, \quad c = 50\,000\,000, \quad c_f = 150, \quad r = 0.04.$$

For the numerical calculations we use the Matlab toolbox OCMat[4], see Grass et al. (2008) for a detailed description of the underlying numerical methods.

[3] https://www.fbijobs.gov/113.asp, accessed August 14, 2014.

[4] Available at http://orcos.tuwien.ac.at/research/ocmat_software/ , accessed August 21, 2014

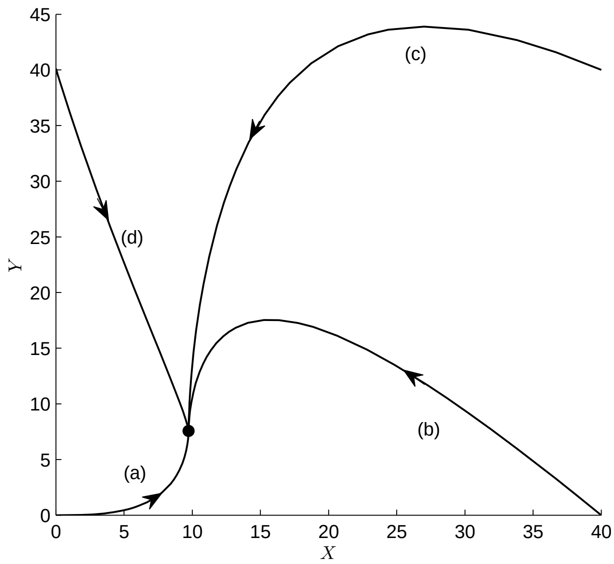

Figure 1: Phase portrait showing the steady state and four exemplary solution paths approaching it.

With these parameter values, there is one steady state, which is a saddle point, at $\hat{X} = 9.73$, $\hat{Y} = 7.56$, $\hat{\lambda}_1 = 12\,094\,015.92$, $\hat{\lambda}_2 = 58\,000.91$, $\hat{f} = 1\,562.16$, $\hat{\omega} = 0$ as can be seen in the phase portrait in Figure 1. The costate variables are positive, meaning that an additional terror plot – known or unknown – is harmful. Obviously, unknown terror plots are harmful as some will succeed and cause damage. Known terror plots require agents' attention and, thus, make detection of new terror plots more costly. However, it makes sense that the steady state value of the first costate is larger than of the second, meaning unknown terror plots are more harmful than known terror plots.

Let us now consider scenarios which differ in the initial numbers of known and unknown terror plots.

- **A new terror threat**: The advent of an entirely new terror threat corresponds to the initial state values $X(0) = 0$ and $Y(0) = 0$ (solution path (a) in Figure 1). The left panel of Figure 2 shows that initially it is not optimal to have many agents, but as the number of (unknown) terror plots, X, increases it is optimal to significantly increase detection efforts. The number of known terror plots Y increases slowly at first because there are few plots to detect. However, as both X and f become larger, so does the flow into Y. Because the inflow to X is constant, this larger outflow leads to a lower growth of X. Thus, it is not necessary after some time to assign more agents to detection. When X and f get closer to their steady state levels the inflow and thus the increase of Y becomes smaller, until it reaches its steady state level. Note that the relatively high value of the parameter ρ implies that agents interdict known plots fairly quickly. As a result, in the long run significantly more intelligence agents work on detecting terror plots than on interdicting them. This can be seen in the right panel of Figure 2. It depicts the share of interdiction agents, which is Y/f, and the share of detection agents, which is $1 - Y/f$.
- **A new governmental organization**: When a government first moves beyond

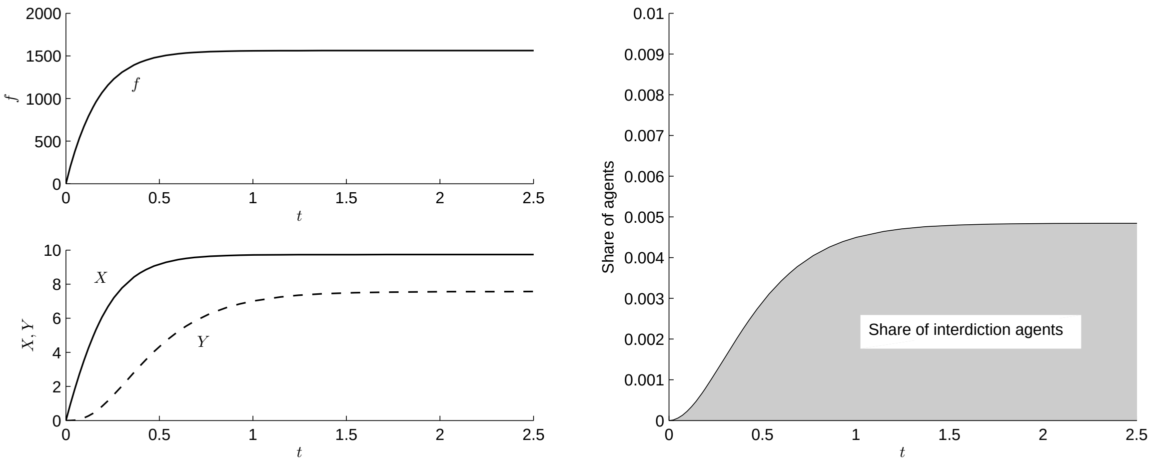

Figure 2: Time path for $X_0 = 0$, $Y_0 = 0$ (left panel) and corresponding share of detection and interdiction agents (right panel)

general homeland security by establishing a new organization that actively seeks to detect and infiltrate plots, this might correspond to the initial state values $X(0) = 40(= \frac{\alpha}{\mu})$ and $Y(0) = 0$ (solution path (b) in Figure 1), since those are the steady state values of the system without governmental intervention with respect to terror detection and interdiction. The scenario is depicted in the left panel of Figure 3. Initially it is optimal to assign many agents to detect plots, consequently X, the number of unknown terror plots decreases and Y, the number of known terror plots, increases. When X falls, it becomes optimal to assign fewer and fewer agents to terror detection as the damage caused by unknown terror plots decreases. Furthermore, it becomes more difficult for agents to detect terror plots when there are not so many of them and this also lowers the incentive to do terror detection. The decline of X also means that the inflow to Y becomes smaller, so that after some time the number of known terror plots also starts to decrease, until it reaches its steady state level.

- **Strategical improvements**: It might also occur that a governmental organization already engaged in terror detection wants to begin optimizing its strategy. For instance, there might initially be something of a knee-jerk response to a new terror threat, but if it proves impossible to root out that threat entirely, society might resign itself to finding a more balanced and sustainable strategy for the long haul. Thus we consider what happens if the initial size of Y is high. If X is also initially high (solution path (c) in Figure 1), it is optimal to assign many agents for terror detection, which leads to an increase of Y. Just as in the scenario described earlier, as X decreases it becomes optimal to do less terror detection. Due to the decrease of X as well as of f, the inflow to Y decreases and Y falls until it reaches its steady state value. If X is initially low (solution path (d) in Figure 1), it would be better to withdraw some agents from their terror prevention tasks. However, we do not allow this by assumption, and thus initially we have $f = Y$. As there is no flow from X to Y then, the number of known plots decreases and the number of unknown plots increases due to the constant inflow. After some time, X becomes so high that it is

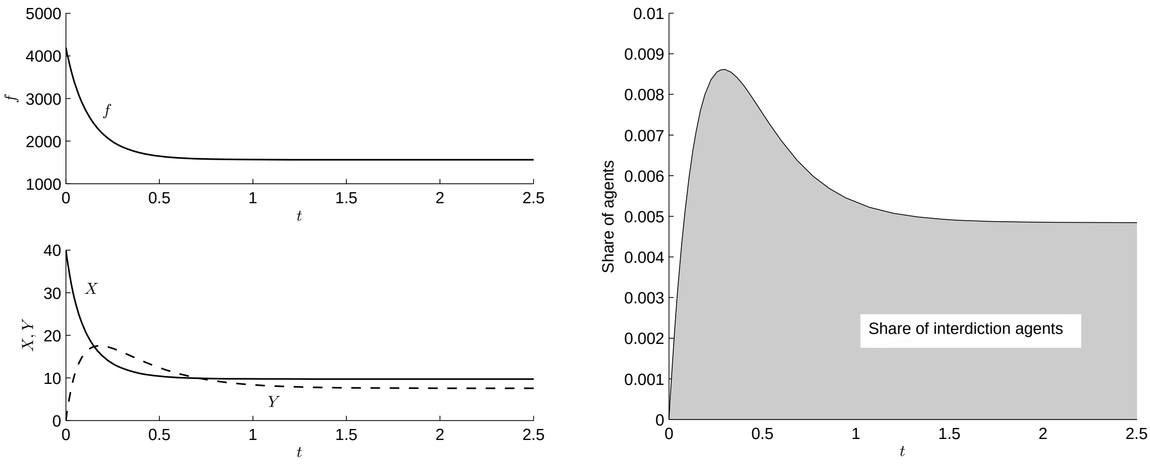

Figure 3: Time path for $X_0 = 40$, $Y_0 = 0$ (left panel) and corresponding share of detection and interdiction agents (right panel)

optimal to increase f, however, the outflow from Y is still so large that the number of terror plots decreases until it reaches its steady state value.

Note that if the initial number of unknown terror plots X is high, we can have an overshooting of Y, i.e., at certain times the optimal solution trajectory may have the number of known terror plots exceeding its steady state value even though the initial state value of Y was below it. The reason for this is that when X is large, so is the flow from X to Y. This happens at such a high rate that undercover agents are not able to handle the terror plots quickly enough to make Y decrease. Thus, the magnitude of X as well as the strong efforts to decrease the number of unknown terror plots imply that at some point Y will exceed its steady state value. However, when X decreases, so does the flow to Y. Consequently, the number of known terror plots will also decrease until it reaches its steady state level.

Figure 4 depicts the number of agents per unknown terror plot (which due to $\mu = 1$ equals the number of successful attacks) over time for the first two scenarios described above. One can see that in case of a small initial number of unknown terror plots one would use more agents per plot than if the number of plots is high. This relates on the one hand to it being easier to detect terror plots when their number is large, and on the other hand to a lower damage caused by terror when X is small.

3.3 Sensitivity Analysis

Since most of our results are derived numerically, it makes sense to ask how robust the solutions are with respect to parameter changes.

In our numerical example, the rate ρ at which undercover agents eliminate terror plots that have already been detected was rather high ($\rho = 4$). Figure 5 shows the phase portrait for a lower value of ρ ($\rho = 0.5$). The numbers of both detected and undetected terror plots are higher than before. However while the number of unknown terror plots is only slightly bigger (i.e. $\hat{X} = 9.86$), the number of known terror plots is 7.98 times as high (i.e. $\hat{Y} = 60.28$). The steady state number of agents is slightly larger than for a high

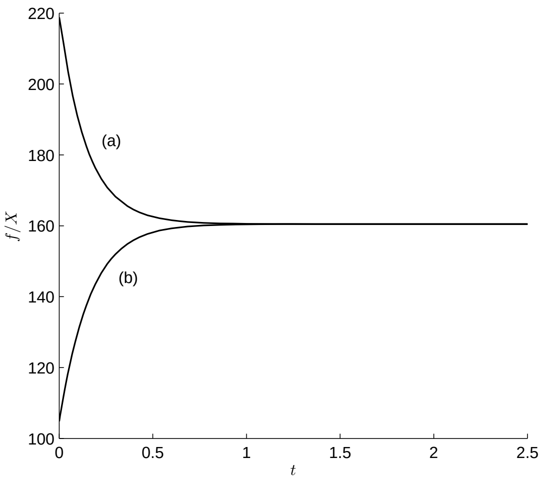

Figure 4: Agent per unknown terror plot over time

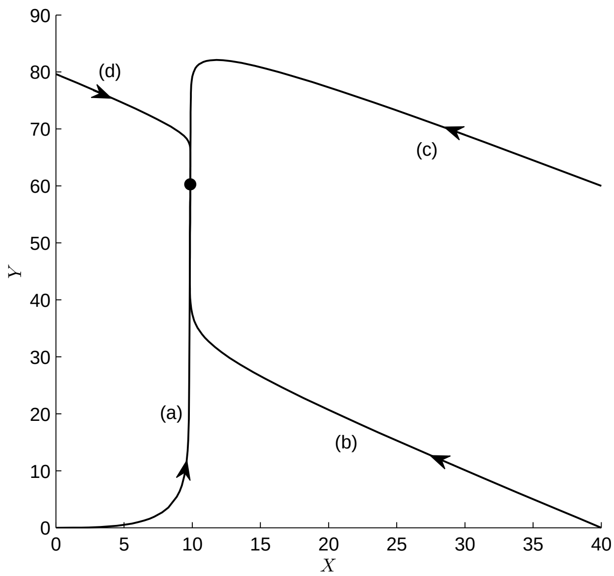

Figure 5: Phase portrait for $\rho = 0.5$ showing the steady state and four different solution paths approaching it.

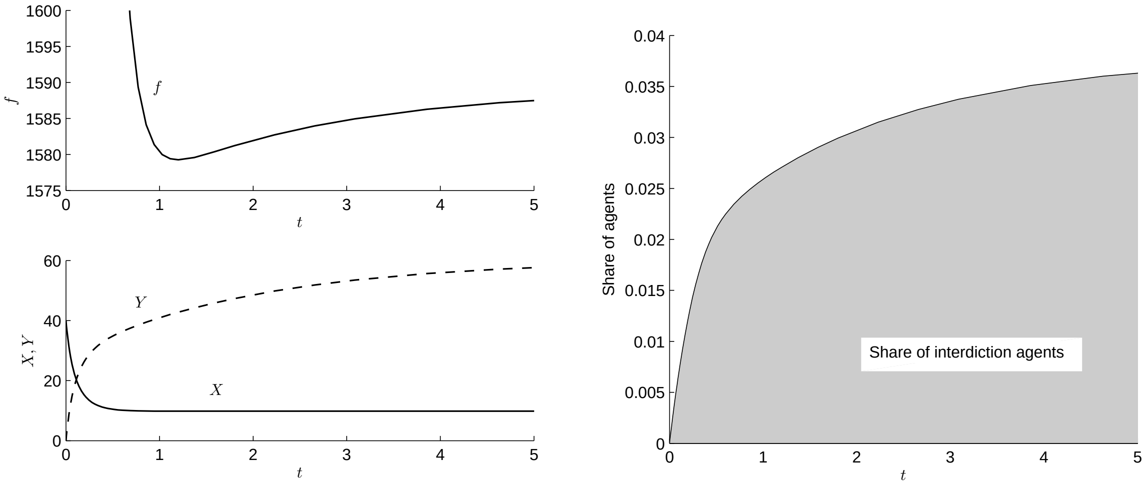

Figure 6: Time path for $\rho = 0.5$ with $X_0 = 40$ and $Y_0 = 0$ (left panel) and corresponding share of intelligence and interdiction agents (right panel)

ρ ($\hat{f} = 1589.11$). Note that f is chosen in a manner such that the steady state detection rate does not differ much for the two different values of ρ. This explains why only $\hat{Y}$ is affected to such an extent by the parameter change. In queuing terms, when the service time increases, there will be more customers in service (and more servers are needed).

The time path for a high initial number of unknown terror plots and no initially known terror plots (corresponding to solution path (b) in Figure 5.) is depicted in Figure 6. While X decreases over time as in the comparable scenario described earlier (Figure 3), that does not hold for f or Y. Due to the longer duration needed to successfully interdict terror plots, the number of known terror plots increases until it reaches its steady state level. Consequently, if agents are occupied with interdiction for a longer time period, one needs more agents in order to decrease X by terror detection. Thus, first when X decreases one can reduce the number of intelligence agents, but later on one needs to slightly increase f again as more and more agents are occupied with interdiction. The right panel of Figure 6 confirms that in this scenario it is optimal to assign more agents to interdiction tasks than for a high ρ.

Note that in this case the control does not follow X over time as before. This illustrates that if one merely adopted the strategy of relating terror control efforts to the estimated number of unknown terror plots that would be myopic and lead to higher costs than in this dynamic framework.

In Figure 7 one can see how the steady state values change if ρ changes. As already explained, a lower ρ leads to a higher $\hat{Y}$ as it takes a longer time to completely interdict the plots. The steady state number of unknown terror plots decreases in ρ. Note that the ratio of unknown terror plots to known terror plots in the steady state increases in ρ. This is because the size of the detection rate has a stronger impact on the size of $\hat{Y}$ than on the size of $\hat{X}$.

If the inflow of unknown terror plots α increases, the steady state number of both unknown and detected terror plots rise even though the optimal level of control efforts also increases. Note that for small α, $\hat{X}$ slightly exceeds $\hat{Y}$. The number of unknown terror plots and as such the damage caused by them is so small then that it does not pay

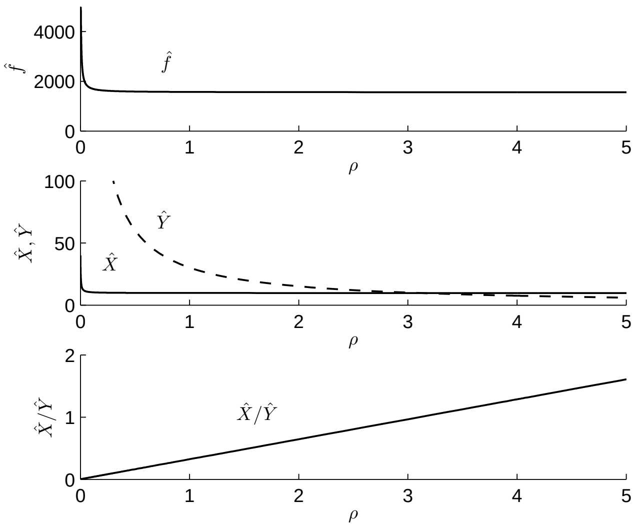

Figure 7: Bifurcation diagram showing how the steady state changes for different values of ρ

off to invest much into terror detection.

Figure 8 shows the dependency of the steady state on parameter μ. A higher μ means that undetected plots mature into successful terror attacks more quickly, so agents have less time to detect a terror plot before it succeeds. This has two implications: On the one side it is necessary to increase detection efforts to compensate for the higher success rate of terror plots. On the other hand X decreases faster. Thus, it is optimal to increase the steady state control efforts only until μ reaches a certain value. Beyond this value terror plots are executed too fast to be handled efficiently, and thus the steady state value of f becomes lower again. By contrast both $\hat{X}$ and $\hat{Y}$ are strictly decreasing in μ, as there is a higher outflow from X and a lower inflow to Y. Due to the increased control efforts, the decrease of $\hat{X}$ is stronger than that of $\hat{Y}$. Consequently, the ratio of known to unknown terror plots increases. However, when $\hat{f}$ smaller again, $\hat{X}$ does not decrease as much anymore due to which the ratio $\hat{X}/\hat{Y}$ even starts to increase again.

Parameter δ reflects the efficiency of undercover agents with respect to detecting terror plots. If this parameter increases, it requires less effort to detect a terror plot, thus $\hat{Y}$ increases and $\hat{X}$ decreases so much that it is possible to decrease $\hat{f}$ without negatively affecting the number of known terror plots.

A change of the discounting rate r does not have much impact on the steady state, although an increase of r, i.e. a more myopic decision maker, leads to a slightly lower $\hat{f}$ and a slightly higher$\hat{X}$.

Of course costs of terror and of terror prevention also affect the steady state values. Higher costs c caused by successful terror attacks mean that more efforts should be devoted to combating terror. Thus the optimal number of undercover agents in the steady state, $\hat{f}$, becomes so much higher if c increases that the steady state number of unknown terror plots $\hat{X}$ decreases and the number of detected terror plots $\hat{Y}$ increases.

Increasing the costs of employing undercover agents, c_f, has the opposite effect: it becomes less efficient to assign agents to terror detection and interdiction, so $\hat{f}$ decreases. This leads to an increase of $\hat{X}$ and a decrease of $\hat{Y}$.

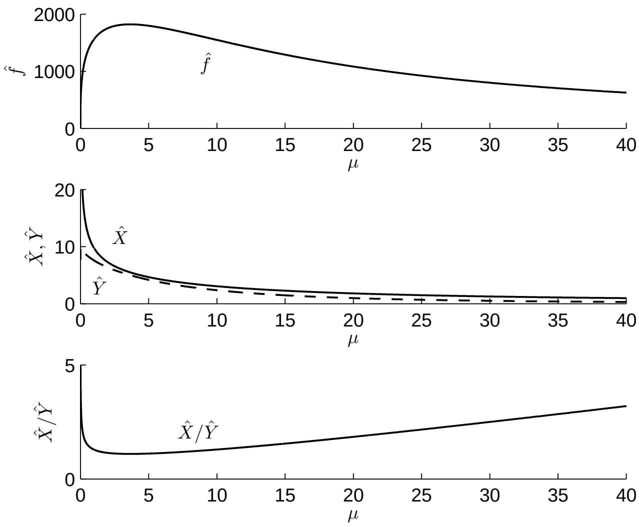

Figure 8: Bifurcation diagram showing how the steady state changes for different values of μ

4 Requiring the Number of Terror Attacks to Always Decrease

Figure 1 shows that it is not always optimal to make the number of unknown terror plots and, consequently the number of successful attacks, decrease for every initial state value. In some cases it was optimal to let terror attacks increase at least for a time. This follows from the pure mathematical optimization problem. For a decision maker, however, allowing terror attacks to increase might be infeasible politically, even if doing so might lead to a better outcome in aggregate, summing over all time. Thus, in the subsequent section we add a constraint that does not allow the rate of terror attacks to grow over time. As only unknown terror plots lead to successful terror attacks in our model this condition reads as

$$\dot{X} \leq 0. \tag{8}$$

4.1 Analytical Results

Rewriting condition (8) as

$$\alpha - \mu X - \delta(f - Y)X \leq 0,$$

we see that if we want to prevent X from increasing, we have to choose the control in a manner such that $\dot{X} = 0$, which gives

$$f = \frac{\alpha - \mu X + \delta Y X}{\delta X}. \tag{9}$$

In order to be admissible it has to hold that $f \geq Y$ which translates to

$$X \leq \frac{\alpha}{\mu}.$$

For $X > \frac{\alpha}{\mu}$ it is not possible to choose the control in a manner such that both constraints, namely $f \geq Y$ and $\dot{X} \leq 0$, remain active at the same time (if they become active at all in this particular area of the state space). Note, however, that the constraint $\dot{X} \leq 0$ cannot be active for $X > \frac{\alpha}{\mu}$, because $\alpha - \mu X$ is negative then, and since the term $\delta(f - Y)X$ is always non-negative, the constraint is automatically fulfilled.

If the control constraint $\dot{X} \leq 0$ becomes active for $X < \frac{\alpha}{\mu}$, the control will always be chosen in a manner such that $f \geq Y$. Therefore, both constraints cannot be active at the same time and we refer to the derivation of the necessary optimality conditions for the constraint $f \geq Y$ to the previous section. The Lagrangian is now

$$\mathcal{L} = \mathcal{H} - \sigma\left(\alpha - \mu X - \delta(f - Y)X\right),$$

and we determine for $X \leq \frac{\alpha}{\mu}$ the Lagrange multiplier σ to be

$$\sigma = \frac{-c_f f + (\lambda_1 - \lambda_2)\,\delta X}{\delta X},$$

in case the constraint on the state equation becomes active. It has to hold that $\sigma \leq 0$ and $\sigma\dot{X} = 0$. The costate equations become

$$\begin{aligned}
\dot{\lambda}_1 &= (r + \mu)\,\lambda_1 + (\lambda_1 - \lambda_2)\,\delta(f - Y) - c\mu - \sigma\left(\mu + \delta(f - Y)\right), \\
\dot{\lambda}_2 &= (r + \rho)\,\lambda_2 - (\lambda_1 - \lambda_2)\,\delta X + \sigma\delta X.
\end{aligned}$$

When the constraint is active, the state equation's value is zero, and the value of X does not change. Let us denote by $\bar{X}$ this number of unknown terror attacks which does not change unless the Lagrange multiplier σ becomes positive, which happens if it becomes optimal to decrease the number of unknown terror plots by detection. If the constraint is active with $\sigma \leq 0$ inserting the optimal control $f = \frac{\alpha - \mu\bar{X} + \delta Y\bar{X}}{\delta\bar{X}}$ into the state equations, we find

$$\begin{aligned}
\dot{X} &= 0, && (10) \\
\dot{Y} &= \alpha - \mu\bar{X} - \rho Y. && (11)
\end{aligned}$$

Note that since neither state depends on the costates, the calculation of the costates is only required for the calculation of σ, which we need to find a possible point where the constraint becomes inactive. Since X is constant at $\bar{X}$ when the constraint is active, we can find a *continuum of steady states* by setting $\dot{Y} = 0$ which yields

$$\hat{Y} = \frac{\alpha - \mu\bar{X}}{\rho}.$$

Let us assume that the constraint becomes active at $t = 0$, the initial number of unknown terror plots is X_0, and the initial number of known terror plots is given by Y_0. Then we have $\bar{X} = X_0$ and we find that

$$Y(t) = Y_0 e^{-\rho t} + \frac{\alpha - \mu\bar{X}}{\rho}\left(1 - e^{-\rho t}\right).$$

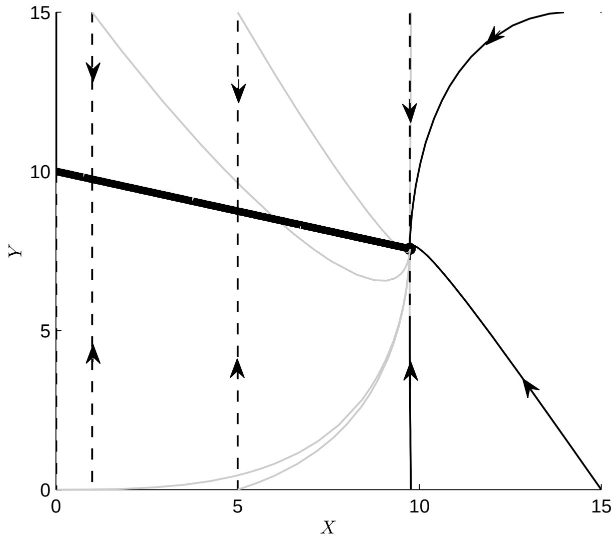

Figure 9: Phase portrait with constraint $\dot{X} \leq 0$. The bold solid line depicts a continuum of steady states, the thin lines trajectories approaching them. On the dashed parts the constraint is active. The gray lines show the solution paths for the unconstrained case for comparison.

We have in case of an active constraint

$$\dot{\lambda}_2 = (r+\rho)\,\lambda_2 - (\lambda_1 - \lambda_2)\,\delta X + \sigma\delta X = (r+\rho)\,\lambda_2 - c_f f.$$

Thus,

$$\hat{Y} = \frac{\alpha - \mu\bar{X}}{\rho}, \quad \hat{f} = \frac{\alpha - \mu\bar{X} + \delta\hat{Y}\bar{X}}{\delta\bar{X}}, \quad \hat{\lambda}_2 = \frac{c_f\hat{f}}{r+\rho}, \quad \hat{\lambda}_1 = \frac{1}{r}\left[\mu\left(c - \hat{\lambda}_2\right) - \frac{c_f\hat{f}}{\delta\bar{X}}\left(\mu + \delta\left(\hat{f} - \hat{Y}\right)\right)\right]$$

4.2 Numerical Results

The phase portrait for the problem with the political constraint can be seen in Figure 9. We have a manifold of steady states (depicted by the bold, solid line) where one can end up depending on the initial state values. If the initial value of X_0 is small, then the constraint $\dot{X} = 0$ remains active along the entire solution path. (In the phase portrait, parts of the solution paths with active constraints are drawn with a dashed line, and the gray lines show the optimal solution path if no constraint is imposed on the number of unknown terror plots).

From (9) we see that the optimal control, i.e. the optimal number of intelligence agents, increases with the number of known terror plots, Y, for the active constraint. The intuition behind this is that in order to keep the number of unknown terror plots constant, these unknown terror plots must be detected and transformed into known plots. This can only be achieved if there is a higher flow from X to Y than if this constraint is not taken into account. Consequently, as the number of interdiction agents corresponds to the number of known terror plots, the total number of agents must be greater or equal to Y and thus, f must increase in Y.

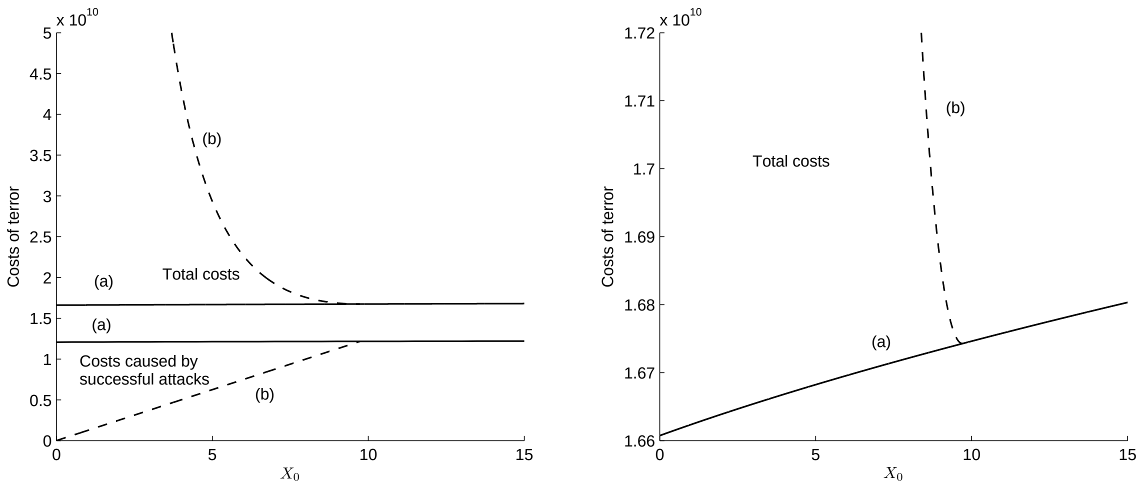

Figure 10: Costs of terror without (a) and with constraint $\dot{X} \leq 0$ (b) depending on the initial number of terror plots X_0 for $Y_0 = 0$ with a zooming in the right panel.

If starting with a low number of (known) terror plots one would assign initially only a small number of agents to terror detection. On the other hand to keep $\dot{X}$ constant at zero, detection measures lead to a flow from X to Y and as such Y increases until its steady state value is reached. If Y is initially rather large, the constant inflow to $\dot{X}$ and the resulting flow from X to Y are not strong enough to keep Y growing or constant, so it falls until it reaches its steady state value.

If the initial number of unknown terror plots X_0 is very large, the political constraint becomes active only very close to the steady state. For some intermediate values of X_0 and a low initial number of known terror plots, Y_0, the constraint is not active in the beginning. However, after some time X becomes slightly smaller than its (unconstrained) steady state value. As it is not allowed for X to increase again, the constraint becomes active, and X remains constant and Y increases until it reaches its steady state level.

Figure 10 shows the costs arising through terror. On the one hand we can see the total costs due to terror (i.e. the costs arising directly because of successful terror attacks and the costs of employing agents to detect and interdict terror plots) if one enforces the constraint that X must not increase (the dashed line (b)), on the other hand we can see the losses if one does not impose the constraint (solid line (a)). Obviously, the enforcement of the constraint leads to higher costs as terror prevention is costly.

The overall societal cost caused by terror increases with the initial number of unknown terror plots X_0 if there is no constraint (solid line). In marked contrast, if politicians require that the number of successful terror plots must not increase, then the costs of terror and terror prevention depend comparatively to an enormous extent on the size of the initial threat for small initial state values and, indeed, are actually higher – not lower. The reason for this surprising result is that if X_0 is very small, then the constraint that X never be allowed to increase becomes more restrictive. Note, however, as one can see in Figure 10, the costs of successful terror attacks are lower when imposing the constraint, it is the significantly higher number of agents required to keep the number of terror attacks low which leads to the higher costs.

Enforcing the constraint that the number of successful terror attacks must not increase can force the hiring of more agents even if that leads to higher control costs. Figure 11

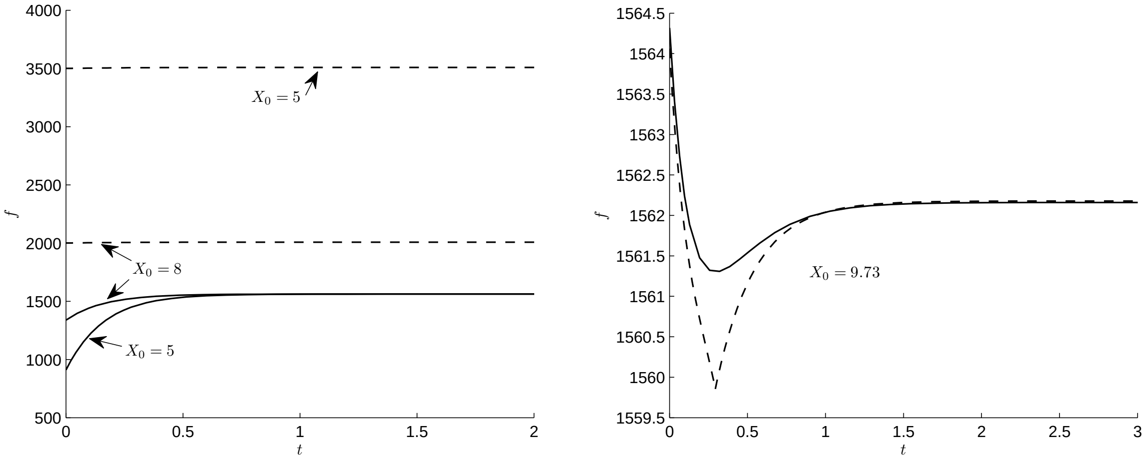

Figure 11: Optimal number of agents over time with (dashed line) and without constraint $\dot{X} \leq 0$ (solid line) for $Y_0 = 0$.

shows the optimal number agents with and without the political constraint for the initial state values $(X_0, Y_0) = (5, 0)$ and $(8, 0)$. If the constraint is active and one starts without any known terror plots, one would increase f over time. The lower the initial number of unknown terror plots the higher is the number of agents required to keep the number of successful terror attacks low as detection is more difficult. Note, however, if X_0 is very close to the steady state value of the unconstrained case, it is initially optimal for a small Y_0 to use more agents if we do not impose the constraint as one can see in the right panel of Figure 11. Due to the relatively high number of unknown terror plots, terror detection is rather efficient then. Thus, in the unconstrained case, a higher application of the control leads to a decrease of X. This sounds beneficial, so one might wonder why not to do the same for the constrained case. The reason is, as one could see before, it is much more costly to prevent X to increase when it is low, so – from a purely monetary perspective – it does not pay off to decrease successful terror attacks and then have to keep them low. However, in the long run the number of agents is slightly higher for the case where we impose the constraint than for the case where we do not.

Note that also other types of constraints might be interesting. E.g. instead of requiring the number of attacks never to increase, one could also consider the pure state constraint that the number of attacks must be below a certain threshold, i.e. $\mu X(t) \leq \chi$ for $t \geq \tau$. If the initial number of terror attacks μX_0 is above this threshold, it makes sense to let the number of attacks decrease before imposing the constraint (otherwise there would not be an admissible solution to this problem). Consequently, $\tau \geq 0$ is the time (exogenously) given to reach the threshold χ. Alternatively, one could just require the number of terror attacks to never to rise above its initial level. Then we impose $\chi = \mu X_0$.

In both cases, if $X(t)$ is below the threshold χ/μ, $t < \tau$, one can afford to let the number of terror attacks increase until it reaches the critical threshold. Once it is reached at time $t = t_\chi$, if $t_\chi \geq \tau$ one would enforce the constraint for $t \geq t_\chi$. Of course, this requires more efforts for a small χ than for a large one (in case of an active constraint we have $f(t) = \frac{\mu(\alpha-\chi)+\delta\chi Y(t)}{\delta\chi}$). If $t_\chi < \tau$ it might be optimal to have the number of terror attacks increase above the critical threshold for some time. Whether such a strategy is optimal depends on the costs of decreasing the number of unknown terror plots until time

τ. In case $\chi = \mu X_0$ we again have a continuum of steady states with active constraint.

5 Conclusion

The present paper considered how a government should optimally employ undercover agents to fight terror. We saw that the optimal strategy depends on the initial state values. In particular, if there are initially no terror plots, it is not necessary to do much in the beginning as the costs of terror are not very large. As the number of terror plots increases, it becomes optimal to increase the number of active agents. On the other hand if there are initially many terror plots, it is optimal to assign many agents to terror prevention first, and do less as the number of terror plots falls.

As most of our results were derived numerically, we studied how a change of parameter values affects the long run steady state. We studied the dependency of the steady state strategy on the interdiction rate ρ. One could say that if agents operate more and more efficiently, the number of terror plots in the steady state becomes smaller. The change in the steady state level of counter-terror efforts is quite interesting with respect to changes in the terror success rate μ. First f is increasing and then decreasing rather than exhibiting a monotonic relationship. A reason for this is that when terrorists act more efficiently, a government must do more to reduce the damage due to successful attacks. However, when terrorists act more efficiently, the chances of detecting a terror plot before its execution become smaller and smaller, and thus it becomes less efficient to even try to detect the attacks.

The basic model seeks to balance the costs of terror attacks with the cost of terror control. So if the number of terror plots is initially very low, it is optimal to do relatively little. Thus, pursuing the optimal policy that is optimal for an all-powerful social planner may involve "allowing" the number of terror attacks to increase. However, such a policy is certainly not politically convenient or perhaps even feasible for a democratically elected government that must be responsive to public opinion. Thus, we also consider scenarios that impose the constraint that the number of successful terror attacks must not increase. This constraint is not only of practical interest from a policy perspective, but also mathematically interesting as one state is kept constant once the constraint becomes active. This leads to a continuum of steady states. In particular, if the initial number of unknown terror plots is low, one would keep this number fixed, allowing only the number of known terror plots to increase or decrease depending on its initial size. Naturally imposing any constraint increases costs, and the impact on the objective value is higher, the lower is the initial number of (undetected) terrorists, because then the costs of detecting terror plots exceeds the utility of preventing terror attacks the most. What is perhaps less obvious is that it can lead to the paradox of lower overall social costs when the initial scale of the terrorist problem is large. However, the difference in the objective value compared to the unconstrained solution becomes smaller as the strategy becomes closer to the strategy when not imposing the political constraint. For a high number of initial terror plots, the unconstrained optimal strategy would have had that the number of terror attacks decreases in any event.

There are many possibilities for extensions. In the present paper it was assumed that the arrival rate of terror plots is constant. However, terrorists might adapt their behavior in response to governmental actions, and a forward thinking government should take this behavior into consideration. Thus, it would make sense to consider a differential game between the government and the terrorists. In this context it would be interesting to see how asymmetries with respect to available information affects the optimal strategies. In particular, the government might be able to observe directly the number of known terror plots but only have some rough idea about the total number of terror plots. Or, the terrorists might know how many terror plots there are, but not know how many of these plots are known to the government.

Given the very real possibility of players having imperfect information, it would be interesting to study what happens if a government errs about the number of terror plots or the terrorists err in their assumptions about what proportion of their plots have already been detected.

Acknowledgments

This research was supported by the Austrian Science Fund (FWF) under Grants P25979-N25 and P25275-G11. Edward H. Kaplan was supported by the Yale School of Management's faculty research funds.

References

References

Appelbaum, B. (2011). As U.S. agencies put more value on a life, businesses fret. New York Times, p. A1. http://www.nytimes.com/2011/02/17/business/economy/17regulation.html?pagewanted=all (last accessed May 26, 2015).

Bartholomew, D. J. (1973). *Stochastic models for social processes.* John Wiley & Sons, London, 2^{nd} edition edition.

Behrens, D., Caulkins, J., Tragler, G., and Feichtinger, G. (2000). Optimal control of drug epidemics: prevent and treat – but not at the same time? *Management Science*, 46(3):333–347.

Behrens, D. A., Caulkins, J. P., Feichtinger, G., and Tragler, G. (2007). Incentive Stackelberg strategies for a dynamic game on terrorism. In Jørgensen, S., Quincampoix, M., and Thomas, V. L., editors, *Advances in Dynamic Game Theory*, volume 9 of *Annals of the International Society of Dynamic Games*, pages 459–486, Boston. Birkhäuser.

Caulkins, J., Grass, D., Feichtinger, G., and Tragler, G. (2008). Optimizing counter-terror operations: Should one fight fire with "fire" or "water"? *Computers & Operations Research*, 35(6):1874–1885.

Caulkins, J. P., Feichtinger, G., Grass, D., and Tragler, G. (2009). Optimal control of terrorism and global reputation. *Operations Research Letters*, 37(6):387–391.

Crettez, B. and Hayek, N. (2014). Terrorists' eradication versus perpetual terror war. *Journal of Optimization Theory and Applications*, 160(2):679–702.

Feichtinger, G. and Novak, A. (2008). Terror and counterterror operations: Differential game with cyclical Nash solution. *Journal of Optimization Theory and Applications*, 139(3):541–556.

Grass, D., Caulkins, J. P., Feichtinger, G., Tragler, G., and Behrens, D. A. (2008). *Optimal Control of Nonlinear Processes: With Applications in Drugs, Corruption and Terror*. Springer, Heidelberg.

Hillier, F. S. and Lieberman, G. J. (2001). *Introduction to Operations Research*. McGraw-Hill, New York, 5th edition edition.

Kaplan, E., Kress, M., and Szechtman, R. (2010). Confronting entrenched insurgents. *Operations Research*, 58(2):329–341.

Kaplan, E. H. (2010). Terror queues. *Operations Research*, 58(4-part-1):773–784.

Kaplan, E. H. (2012). Estimating the duration of Jihadi terror plots in the United States. *Studies in Conflict & Terrorism*, 35(12):880–894.

Kaplan, E. H. (2013). Staffing models for covert counterterrorism agencies. *Socio-Economic Planning Sciences*, 47(1):2–8.

Kaplan, E. H. (2015). Socially efficient detection of terror plots. *Oxford Economic Papers*, 67(1):104–115.

Keohane, N. O. and Zeckhauser, R. J. (2003). The ecology of terror defense. *Journal of Risk and Uncertainty*, 26(2–3):201–29.

Kress, M. and Szechtman, R. (2009). Why defeating insurgencies is hard: The effect of intelligence in counterinsurgency operations–a best-case scenario. *Operations Research*, 57(3):578–585.

National Consortium for the Study of Terrorism and Responses to Terrorism (START) (2013). Global terrorism database [data file]. Retrieved from http://www.start.umd.edu/gtd (last accessed May 26, 2015).

Newell, G. F. (1971). *Applications of Queueing Theory*. Chapman and Hall, London.

Rogers, E. (2003). *Diffusion of Innovations*. Free Press, 5th edition edition.

Strom, K., J. Hollywood, M. P., Weintraub, G., Daye, C., and Gemeinhardt, D. (2010). Building on clues: Examining successes and failures in detecting U.S. terrorist plots, 1999–2009. Institute for Homeland Security Solutions, Research Triangle Park, North Carolina. http://sites.duke.edu/ihss/files/2011/12/Building_on_Clues_Strom.pdf (last accessed August 14, 2014).

Tragler, G., Caulkins, J. P., and Feichtinger, G. (2001). Optimal dynamic allocation of treatment and enforcement in illicit drug control. *Operations Research*, 49(3):352–362.

Viscusi, W. K. and Aldy, J. E. (2003). The value of a statistical life: a critical review of market estimates throughout the world. *Journal of Risk and Uncertainty*, 27(1):5–76.

Zhuang, J. and Bier, V. M. (2007). Balancing terrorism and natural disasters-defensive strategy with endogenous attacker effort. *Operations Research*, 55(5):976–991.

Zhuang, J., Bier, V. M., and Alagoz, O. (2010). Modeling secrecy and deception in a multiple-period attacker–defender signaling game. *European Journal of Operational Research*, 203(2):409–418.