TECHNISCHE UNIVERSITÄT WIEN

Operations Research and Control Systems

SWM ORCOS

Differential Terror Queue Games

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

Research Report 2016-02

March, 2016

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

Differential Terror Queue Games

S. Wrzaczek, E.H. Kaplan, J.P. Caulkins, A. Seidl, G. Feichtinger

March 2, 2016

Abstract

We present models of differential terror queue games, wherein terrorists seek to determine optimal attack rates over time while simultaneously the government develops optimal counter terror staffing levels. The number of successful and interdicted terror attacks are determined via an underlying fluid terror queue model. Different information states and commitment abilities derive from different assumptions regarding what the players in the game can and cannot deduce about this underlying model. We consider three different possibilities: open-loop, where both the terrorists and the government have full information but must commit in advance to all future decisions when the game begins; a mixed structure whereby the terrorists are able to observe the total number of terror plots they have launched but are unaware of how many have already been detected by the government, while the government is *only* aware of terror plots they have discovered and unaware of undetected but in-progress terror plots; and the closed-loop structure whereby both players have full information over time. We characterize the optimal controls for both the terrorists and the government in terms of the associated state and costate variables, and deduce the costate equations that must be solved numerically to yield solutions to the game for the different information structures. Using recently assembled data describing both terror attack and staffing levels, we compare the differential game models to each other as well as to the optimal control model of Seidl et al [15]. The paper concludes with a discussion of the lessons learned from the entire modeling exercise.

Keywords: some keywords.

1 Introduction

The threat of terrorism remains one of the world's most pressing problems, and a growing academic literature seeking to understand how both terrorist organizations and governments make decisions has sprouted in response. Enders and Sandler [4] provide a very readable account of empirical and game-theoretic work in this area, while new articles and special journal issues continue to appear; for examples see the January 2015 *Oxford Economic Papers*, the December 2011 issue of *Public Choice*, the May 2011 issue of *Journal of Peace Research*, and the April 2010 issue of *Journal of Conflict Resolution*, among others. Among those articles focusing on the application of game theory to terrorism and counterterrorism, most focus on one- or two-period games at a high level of abstraction.

More recently, Kaplan and colleagues initiated a series of studies that focus on basic operational questions such as estimating the number of undetected terror plots (Kaplan [9]), determining optimal terrorist attack rates and government counterterror agent staffing levels in equilibrium (Kaplan [11], Kaplan [12]), optimal terrorist- and government- force allocation across different insurgent strongholds (Kaplan et al. [8]), and developing optimal time-dependent counterterror staffing plans (Seidl et al. [15]). The present paper builds on this recent work (and most directly upon Seidl et al. [15]) to model differential terror queue games, wherein terrorists seek to determine optimal attack rates over time while simultaneously the government develops optimal counterterror staffing levels. As is often the case with differential game models (Dockner et al. [3], Long [14]), different assumptions governing information structure and commitment ability can lead to different results, and these differences are especially important to recognize in the terrorism setting.

The paper begins with a preliminary analysis of a terrorist organization operating in the absence of government counteractions. The resulting model applies to the decisions of a newly emerging group of terrorists whose existence is as of yet unknown, and provides the other side of the coin examined by Seidl et al. [15] who presumed that only the government acted strategically against a fixed terrorist threat. We will show it is optimal for terrorists to conduct attacks at a constant rate in this scenario, somewhat justifying what was taken as a simplifying assumption in Seidl et al's earlier work.

The paper then defines the differential game that is at the heart of this paper. Both the terrorists and the government are presumed rational in the economic sense. The terrorists seek to maximize the present value of the benefits of successful attacks minus the costs of all attacks initiated by selecting an optimal trajectory of attack rates over time, while the government seeks to minimize the present value of the cost of terror attacks plus the staffing and related costs from fielding counterterror agents by selecting an optimal time-dependent agent staffing plan. The number of successful and interdicted terror attacks are determined via the fluid version of the terror queue model first introduced in Kaplan [9] and most recently utilized in Seidl et al. [15].

We compare the impact of three different information structures which relate to the assumption that each player makes about his opponent's decision rule. In all three cases the players have full information about the form of the game and choose their strategy as best response to their opponent's strategy. In the open-loop structure, the players assume that their opponent is not able to observe the number of known and unknown terror plots, i.e. the state variables, and therefore see their opponent's decision rule only as a function of time. As such the players neglect the impact of a change of the state variable on their opponent's strategy in their determination of the optimal strategy. We contrast the results to closed-loop strategies, where the players are aware that their opponent actively responds to a change of the state variables, and to mixed strategies, where players assume that that their opponent's strategy depends on only one of the state variables and time. The contribution of this analysis is to provide a deeper understanding about the impact of different information structures which are commonly used in differential games, see e.g. Basar and Olsder [2] and Lambertini [13].

From an analysis of the Hamiltonian and first-order conditions for these games, we are able to characterize the optimal controls for both the terrorists and the government in terms of the associated state and costate variables. For the different information structures we deduce the costate equations which must be solved numerically to yield solutions to the game. Using recently assembled data describing both terror attack and staffing levels, we present numerical calculations to compare the differential game models to each other as well as to the optimal control model of Seidl et al in order to see the impact of the various assumptions involved. We include sensitivity analyses to illustrate how the model solutions depend upon the values of key model parameters. The paper concludes with a discussion of the lessons learned from the entire modeling exercise.

2 The Unobservable Terrorist

As already mentioned in the introduction we start with a model of a terrorist group, who can act without any counteractions of the government. We consider the number of terror plots (unknown to the government)

$X(t)$ as a state variable. The total number of terror plots increases by the number of newly initiated terror plots $\alpha(t)$, which is the control variable of the terrorists, and decreased by the number of successsul attacks $\mu X(t)$, where μ is considered to be exogenous.[1] The dynamic state equation thus equals

$$\dot{X} = \alpha - \mu X, X(0) = X_0, \tag{1}$$

where X_0 is the initial number of terror plots at the beginning of the planning horizon.

The terrorists seek to maximize its intertemporal utility of the successful terror attacks while taking into account the costs of initiating a new terror plot (which are assumed to be quadratic). The utility and the costs are weighted by the parameters b and k respectively. Therefore the terrorists maximize the following objective function (subject to (1)) with respect to their optimal control variable $\alpha(t)$,

$$V = \max_{\alpha(t)} \int_0^\infty e^{-r_2 t}\Big(b\mu X(t) - \frac{k}{2}\alpha^2(t)\Big)\, dt, \tag{2}$$

where r_2 denotes the subjective discount rate of the terrorists.

To derive the optimal solution for the control variable we formulate the Hamiltonian of the terrorists, i.e.

$$\mathcal{H} = b\mu X - \frac{k}{2}\alpha^2 + \eta(\alpha - \mu X), \tag{3}$$

where $\eta(t)$ denotes the costate variable. Consequently the first order condition for the control reads,

$$\mathcal{H}_\alpha = -k\alpha + \eta \qquad \Longrightarrow \qquad \alpha = \frac{\eta}{k}. \tag{4}$$

The Legendre Clebsch condition, see e.g. [6], is fulfilled. For the costate equation we obtain

$$\dot{\eta} = (r_2 + \mu)\eta - b\mu. \tag{5}$$

Taking (1) and (5) into account we are able to derive the unique steady state of the system. We obtain

$$\begin{aligned} \hat{X} &= \frac{b}{k(r_2+\mu)} \\ \hat{\eta} &= \frac{b\mu}{r_2+\mu}. \end{aligned} \tag{6}$$

Since the canonical system is linear in the state and costate, we are able to express the optimal trajectories for the control and the state analytically, i.e.

$$\begin{aligned} \alpha(t) &= \frac{b\mu}{k(r_2+\mu)} \\ X(t) &= \frac{e^{-\mu t}\big(\mu X_0 + (e^{\mu t}-1)\frac{b\mu}{k(r_2+\mu)}\big)}{\mu}. \end{aligned} \tag{7}$$

Since the optimal number of newly initiated terror plots is a function of exogenous parameters only, it is constant over time. This relates to the fact that α is an absolute value in the state equation and in the objective, meaning that the efficiency of the control does not depend on the value of the state. Furthermore the time horizon is infinite, such that the terrorists have no aim to push the state variable into any direction

[1] In the subsequent time argument t is ommitted for notational convenience unless necessary.

related to the expected salvage value.

From the above analysis we can conclude the behavior of the terrorists if they assume that they are not or cannot be observed by the government. Independently of the state and the time they always initiate the same number of terror plots. In Seidl et al. [15] we used an optimal control model to derive the optimal strategy of the government given a constant arrival rate of new terror plots, equal to the above derived rate. Having the above model of the unobservable terrorist in mind, the result of the optimal control model can be seen as the optimal behavior of the government if the terrorists are assumed to be unobservable (assumed by both the government and the terrorists). In section 5 we compare the result of this optimal control model with the results of the differential game where both players are assumed to act strategically. This shows the effect of a wrong assumption concerning the terrorists observability.

3 A Differential Game Between the Government and Terrorists

3.1 The Basic Model

We consider a two-state diffusion model where X denotes undetected and Y denotes detected terror plots. The terrorists' control variable α ($\alpha \geq 0$) denotes the arrival rate of the undetected plots, i.e. the number of newly initiated terror plots. The control variable of the government f denotes the number of (teams of) undercover agents deployed by the government to infiltrate terror plots. The agents can be divided into two groups, interdictors (one for each detected plot) and detectors. Analogously to a service point in a queuing model there is a 1-to-1 relation between interdicting agents and detected terror plots. Thus $f - Y$ agents are seeking to detect further plots. The dynamical systems reads

$$\dot{X} = \alpha - \mu X - \delta(f - Y)X, \tag{8}$$

$$\dot{Y} = \delta(f - Y)X - \rho Y, \tag{9}$$

where X and Y are the states, α and f are the controls, μ denotes the rate of successful undetected terror plots, δ the rate that detecting agents reveal planned terror plots, and ρ the interdiction rate. The 1-to-1 relation between interdictors and detected terror plots implies the following constraint

$$Y \leq f, \tag{10}$$

which simply states that the number of detected terror plots cannot exceed the total number of government agents.

The governments objective is to minimize the damage by successful terror plots and the costs of counter terror efforts, which are assumed to be quadratic, compare Seidl et al. [15]. The objective function reads

$$\max_{f} \int_0^\infty e^{-r_1 t} \left[-c\mu X - \frac{c_f}{2} f^2 \right] dt, \tag{11}$$

where c are the estimated costs by one successful terror attack.

The terrorists on the other hand seek to maximize their utility from successful terror attacks minus the costs of planning an attack, i.e.

$$\max_{\alpha} \int_0^\infty e^{-r_2 t} \left[b\mu X - \frac{k}{2} \alpha^2 \right] dt. \tag{12}$$

Note that the government and the terrorists need not value the damage/utility of a successful terror attack equally, i.e. $c \neq b$. To the government, the cost of a terror attack is linked to the number of casualties and their economic value, in addition to the cost of property damage and perhaps even the cost of apparent loss of government control in the wake of a successful attack. On the other hand, while the loss of terror operatives in an attack (e.g. suicide bombers) is factored into account, it is known that successful attacks spark recruitment to terror organizations while contributing to their goals of disrupting the government and perhaps even replacing it (Jacobson and Kaplan [7]; Feinstein and Kaplan [5]).

In this paper we assume that the government and the terrorists both choose their instruments optimally, implying that the model is a differential game. It is known that in differential games there is not only one solution, but that the solution depends on the assumed information structure and commitment ability of the players (see Dockner et al. [3]). Very often it is not really clear which solution is the correct one, that is, which set of assumptions describe reality best. However, given that the solutions are different, it is important to recognize the differences since it is possible to learn something about the value of additional information or commitment ability. In the current paper we consider the following three cases[2]

open-loop: The information structure is symmetric. Neither player is able to observe the state variables instantaneously. Therefore both the terrorists and the government commit to a strategy over time at the start of the game where they assume that their opponent does not adapt its optimal controls according to changes in the state variables[3].

mixed open- and closed-loop: The information structure is asymmetric, the terrorists assume that the government can observe the number of detected terror plots Y but not the number of undetected terror plots X, while the government assumes that the terrorists are able to observe the total number of terror plots $X + Y$, but do not know which have already been detected (so they do not know X or Y individually).

closed-loop: The information structure is symmetric and both the government and the terrorists have full information regarding the state variables due to a combination of observable data and the ability to deduce the value of unobserved states from model logic and said data.

The different information structures above reflect different possible assumptions governing the information players have at hand, or perhaps their flexibility, or even their sophistication. In the open-loop case, inflexibility in suddenly changing plans (e.g. the training required to dispatch counterterror agents, or the difficulty in suddenly increasing the planned rate of terror attacks absent advanced places) combined with literal belief in equations (8)-(9) might lead both players to simply solve their associated optimal control problems as functions of time only. They do not resort to considerations of subgame perfection that might result from asking what to do should the other player deviate from the open-loop strategy (as such deviations would be viewed as exercising infeasible flexibility as just described).

The mixed structure comes about from assuming that, realistically, the government can only observe terror plots it has already detected, while the terrorists do not know which of their plots are or are not known to the government. Also, the flexibility concerns associated with closed-loop play are not assumed to

[2]For a more detailed discussion on the different information structures we refer to Basar and Olsder [2], Dockner et al. [3], and van Long (2015).

[3]Thus, players do not deduce how their opponent actually determines its decision rule despite having full information about the game. This is probably the main reasons why in literature closed-loop or feedback solutions are seen as more reasonable than open-loop strategies.

matter in this structure; both players are free to consider deviations from play that might come about for any number of reasons.

Finally, the closed-loop structure might make the most sense from a strictly mathematical point of view, for given that the underlying terror queue model is deterministic, it is possible to deduce the values of unobserved states from observable data. For example, consider the government. It observes the number of detected plots $Y(t)$ and cannot observe undetected plots, but it can observe successful terror attacks, and the latter occur at rate $\mu X(t)$. As the government knows the value of μ, it can easily deduce $X(t)$. Similarly, while the terrorists do not observe the number of detected plots $Y(t)$ directly, they do know the total number of plots in play, $X(t) + Y(t)$, and like the government can observed successful attacks and hence deduce $X(t)$ by subtraction. If both the government and the terrorists are sufficiently sophisticated to solve optimal control problems to produced equilibrium strategies in a differential game, then surely they are sophisticated enough to deduce full information even if only partial information is directly observable.

In the following section we derive the optimality conditions for each of the three cases described above.

4 Analysis

In order to capture the asymmetric information structure (mixed open- and closed-loop) we have to transform the system. We introduce a new state

$$Z := X + Y, \tag{13}$$

and work in the (Y, Z)-space. Consequently, we can omit X unless necessary otherwise. The resulting dynamics are

$$\dot{Y} = \delta(f - Y)(Z - Y) - \rho Y, \tag{14}$$

$$\dot{Z} = \alpha - \mu(Z - Y) - \rho Y. \tag{15}$$

The Hamiltonians and the first order conditions for the control are equal for all three cases. The Hamiltonian of the government (also called player 1) is

$$\mathcal{H}^G = -c\mu(Z - Y) - \frac{c_f}{2} f^2 + \lambda_Y \big(\delta(f - Y)(Z - Y) - \rho Y\big) + \lambda_Z \big(\alpha - \mu(Z - Y) - \rho Y\big), \tag{16}$$

where λ_Y and λ_Z denote the government's costates for Y and Z respectively. The resulting first order condition is

$$\frac{\partial \mathcal{H}^G}{\partial f} = -c_f f + \lambda_Y \delta(Z - Y) = 0 \quad \Longrightarrow \quad f = \frac{1}{c_f} \lambda_Y \delta(Z - Y). \tag{17}$$

The Hamiltonian of the terrorists (also called player 2) is

$$\mathcal{H}^T = b\mu(Z - Y) - \frac{k}{2} \alpha^2 + \eta_Y \big(\delta(f - Y)(Z - Y) - \rho Y\big) + \eta_Z \big(\alpha - \mu(Z - Y) - \rho Y\big), \tag{18}$$

with η_Y and η_Z being the terrorists' costates. The terrorists control instrument α is determined by

$$\frac{\partial \mathcal{H}^T}{\partial \alpha} = -k\alpha + \eta_Z = 0 \quad \Longrightarrow \quad \alpha = \frac{\eta_Z}{k}. \tag{19}$$

Open-loop Case

As a benchmark case we consider the scenario where the optimal controls of both players are assumed to be functions of time only, i.e. $f(t) = \phi^G(t)$ and $\alpha(t) = \phi^T(t)$. Thus, the costate equations of the government are

$$\dot{\lambda}_Y = (r_1 + \rho)\lambda_Y - c\mu + \lambda_Y \delta((Z - Y) + (f - Y)) - (\mu - \rho)\lambda_Z, \tag{20}$$
$$\dot{\lambda}_Z = (r_1 + \mu)\lambda_Z + c\mu - \lambda_Y \delta(f - Y), \tag{21}$$

and the costate equations of the terrorists are

$$\dot{\eta}_Y = (r_2 + \rho)\eta_Y + b\mu + \eta_Y \delta((f - Y) + (Z - Y)) - (\mu - \rho)\eta_Z, \tag{22}$$
$$\dot{\eta}_Z = (r_2 + \mu)\eta_Z - b\mu - \eta_Y \delta(f - Y). \tag{23}$$

As a result a 6-dimensional canonical system consisting of equations (14), (15), (20)-(23) has to be considered.

Mixed open- and closed-loop case

In this case the terrorists include into their considerations how the government adapts its strategy according to the number of known terror plots Y, but not according to the total number of terror plots Z, which the government is not able to observe. On the other hand, the government takes into account that the terrorists is aware of the total number of terror plots and deduces how the terrorists adapt their strategy when Z changes. The government does not, however, consider the impact of a change of the number of known terror plots on their opponents strategy as they assume that terrorists have no knowledge about X. Thus, the government assumes that the terrorists control is a function of $Z(t)$ and time, whereas the terrorists assume that the governments control is a function of $Y(t)$ and time, i.e.

$$f(t) = \phi^G(Y(t), t), \quad \alpha(t) = \phi^T(Z(t), t). \tag{24}$$

The expressions above add additional terms in the costate equations, which are

$$\dot{\lambda}_Y = (r_1 + \rho)\lambda_Y - c\mu + \lambda_Y \delta((Z - Y) + (f - Y)) - (\mu - \rho)\lambda_Z, \tag{25}$$
$$\dot{\lambda}_Z = (r_1 + \mu)\lambda_Z + c\mu - \lambda_Y \delta(f - Y) - \lambda_Z \frac{\partial \phi^T}{\partial Z}, \tag{26}$$
$$\dot{\eta}_Y = (r_2 + \rho)\eta_Y + b\mu + \eta_Y \delta((f - Y) + (Z - Y)) - (\mu - \rho)\eta_Z - \eta_Y \delta(Z - Y)\frac{\partial \phi^G}{\partial Y}, \tag{27}$$
$$\dot{\eta}_Z = (r_2 + \mu)\eta_Z - b\mu - \eta_Y \delta(f - Y). \tag{28}$$

with $\frac{\partial \phi^G}{\partial Y} = -\frac{1}{c_f}\lambda_Y \delta$ and $\frac{\partial \phi^T}{\partial Z} = 0$. The 6-dimensional canonical system consisting of equations (14), (15), (25)-(28) has to be considered.

Closed-loop Case

In the closed-loop case both players take into account how their opponent adapts its strategy when the state variables change. In mathematical terms that means that the optimal strategies have to be assumed as the following functions

$$f(t) = \phi^G(Y(t), Z(t), t), \quad \alpha(t) = \phi^T(Y(t), Z(t), t). \tag{29}$$

Applying these to the adjoint equations we get

$$\dot{\lambda}_Y = (r_1+\rho)\lambda_Y - c\mu + \lambda_Y\delta((Z-Y)+(f-Y)) - (\mu-\rho)\lambda_Z - \lambda_Z\frac{\partial\phi^T}{\partial Y}, \tag{30}$$

$$\dot{\lambda}_Z = (r_1+\mu)\lambda_Z + c\mu - \lambda_Y\delta(f-Y) - \lambda_Z\frac{\partial\phi^T}{\partial Z}, \tag{31}$$

$$\dot{\eta}_Y = (r_2+\rho)\eta_Y + b\mu + \eta_Y\delta((f-Y)+(Z-Y)) - (\mu-\rho)\eta_Z - \eta_Y\delta(Z-Y)\frac{\partial\phi^G}{\partial Y}, \tag{32}$$

$$\dot{\eta}_Z = (r_2+\mu)\eta_Z - b\mu - \eta_Y\delta(f-Y) - \eta_Y\delta(Z-Y)\frac{\partial\phi^G}{\partial Z}. \tag{33}$$

with $\frac{\partial\phi^G}{\partial Y} = -\frac{1}{c_f}\lambda_Y\delta$, $\frac{\partial\phi^G}{\partial Z} = \frac{1}{c_f}\lambda_Y\delta$ and $\frac{\partial\phi^T}{\partial Y} = \frac{\partial\phi^T}{\partial Z} = 0$. We have to consider the 6-dimensional canonical system consisting of equations (14), (15), (30)-(33).

When we look at the adjoint systems above ((20)-(23) for the open-loop case, (25)-(28) for the mixed open- and closed-loop case and (30)-(33) for the closed-loop case) the implications of the assumed information structure are straight-forward. The open-loop case can be seen as the foundation of all three cases, since its terms also appear in the other two cases. In the mixed open- and closed-loop case the government's and terrorists' optimal strategies are assumed to depend on Y and Z, respectively. This implies that the marginal effect of these state variables on the decision rule weighted by the effect of the control variable on the Hamiltonian is added to the adjoint system. Thus, it incorporates the future effects (aggregated over time) related to the opponent's adaptions of its decision rule corresponding to changes in Y and Z, respectively. Moving to the closed-loop case we observe that additionally the marginal effect of the second state on the controls is added to the adjoint equations. The intuition behind it is analogous to that given before.

Taking now a closer look on the actual decision rules of the players, (17) and (19), it can be observed that f depends immediately on Y and Z (therefore $\frac{\partial\phi^G}{\partial Y} \neq 0$ and $\frac{\partial\phi^G}{\partial Z} \neq 0$), whereas there is no marginal effect on α (therefore $\frac{\partial\phi^T}{\partial Y} = \frac{\partial\phi^T}{\partial Z} = 0$). The reason for this lies in the way both controls enter in the dynamical system and the objective.

The terrorists' control α denotes the number of newly initiated terror plots and is an absolute rate. There is no interaction with the other states in the dynamical system nor in the objective, meaning that a change of Y or Z has no immediate effect on the efficiency of this control. This implies that the adjoint equations for λ_Y and λ_Z are the same no matter what information structure is used. The interpretation of this is that when the government takes into account that terrorists could adapt their strategy as direct response to a change in the state variable, it can deduce from the derivation of the actual decision rule of the terrorists that it actually is not optimal to do so.

It is different for the governments control f (the number of undercover agents). Unknown terror plots can be detected with rate δ, and (of course) the more unknown plots the easier they can be detected. Thus in the dynamical system there is a multiplicative interaction term between the control f and the states. This means that the optimal f should directly be adapted corresponding to the value of the states and implies the additional term in the adjoint equations.

Remark: Note that in general the optimal strategies (of all three above derived information scenarios) of both players also depend on the initial states $(X(0), Y(0)) = (X_0, Y_0)$. In contrast the feedback Nash equilibrium does not depend on the initial states, but only on the current state and on the time. As a result every feedback Nash equilibrium is strongly time-consistent. Due to the structure of our model it is not possible to derive the feedback Nash equilibrium. For a profound discussion on that issue we refer to section

6 of Basar and Olsder [2].

In all three cases it is not possible to derive analytically a steady state. Thus we have to resort on numerical calculations.

5 Numerical Results

In the following we present numerical calculations of the above differential game. The model includes nine (positive) parameters: μ (rate of successful undetected terror plots), δ (rate of revealing undetected terror plots), ρ (interdiction rate), c (costs of one successful terror attack for the government), c_f (weight of quadratic costs of counter terror measures), W (damage of one successful terror attack for the terrorists), k (costs of planning new terror plots), r_1 (discount rate of the government) and r_2 (discount rate of the terrorists). To get a better intuition for the results of our model the results are presented in the (X, Y)-space.

Since our model is an extension to the optimal control model presented in Seidl et al. [15] we use the same parameters, such that we are able to compare the results of the differential game with that of the optimal control model. The parameters are

$$\mu = 1; \quad \delta = 0.002; \quad \rho = 4; \quad c = 5 * 10^7; \quad c_f = 150; \quad r_1 = 0.04. \tag{34}$$

The parameters of Seidl et al. [15] are taken from Kaplan [9] and several other papers (see Kaplan [10], Strom et al. [16], Kaplan [12], Appelbaum [1], Viscusi and Aldy [17]).

The parameters b, k and r_2 only occur in the differential game. Inserting the given parameter values into (7) gives $\alpha = \frac{b}{1.04k}$ and using, like Seidl et al. [15], $\alpha = 40$ as a reference point, gives us as parameters when assuming the utility of a terror plot to the terrorists is a tenth of the damage it causes to the government. For them we assume

$$b = 5 * 10^6; \quad k = 1.2 * 10^5; \quad r_2 = 0.04. \tag{35}$$

The discount rate r_2 of the terrorists is set equally to that of the government. This is within the range of discount rates used in the optimal control and differential game literature. Furtheron Seidl et al. [15] states that the results of the optimal control model are quite robust to small changes in the discount rate. W and k are difficult to estimate and no data are available. W can hardly be measured by pure monetary terms, here also some ideological factors have to be included.

With these parameter values we get the following unique steady state values

open-loop: $\hat{f} = 1002.40, \ \hat{\alpha} = 13.73, \ \hat{X} = 4.57, \ \hat{Y} = 2.29,$

mixed open- and closed-loop: $\hat{f} = 1097.40, \ \hat{\alpha} = 16.97, \ \hat{X} = 5.32, \ \hat{Y} = 2.91,$

closed-loop: $\hat{f} = 849.55, \ \hat{\alpha} = 9.41, \ \hat{X} = 3.49, \ \hat{Y} = 1.48.$

Comparing them with the steady state values of the optimal control model, which are

optimal control: $\hat{f} = 1562.16, \ \hat{X} = 9.73, \ \hat{Y} = 7.56,$

it can be seen that all controls and states in all three cases are smaller than in case of the optimal control model. So in view of the government that means that considering the terrorists as optimizing agent (thus including their best response in the own optimization) is something positive since that reduces the steady state values of the states (i.e. lower number of terror plots)[4].

Furthermore we observe crucial differences between the solutions due to the information structures. The intermediate case with a mixture of open- and closed-loop information has the greatest values for both controls and both states, followed by the open-loop case and the closed-loop case with the lowest values. The reason for the difference can be explained as follows:

The big difference between open-loop and mixed strategy is that, for the mixed case, when calculating the shadow price of Y, η_y, the terrorists take into account to which extent the government adjusts its strategy when Y changes and its impact on the value of the Hamiltonian. The derivative of ϕ^G with respect to Y is strictly negative, and since η_y is positive[5] $\dot{\eta}_y$ increases. In the mixed case, therefore, the terrorists conclude that when the number of known terror plots Y increases, f decreases (when Z remains constant). However, when determining the shadow price of Z, the terrorists do not directly take into account that an increase of Z leads to an increase of f. Thus, the higher evaluation of known terror plots leads to a higher evaluation of the total number of terror plots and thus the terrorists have a higher incentive in the mixed case to make more terror plots. Therefore, terrorists choose a higher α than in the open-loop case which results in a higher number of terror plots (known and unknown). This leads to a higher f, which in turn the terrorists need to compensate by a higher α.

Now in closed-loop terrorists do not only take into account that a change of Y affects f and consequentially their benefits from Y, but also how a change of Z affects f and therefore the shadow price of Z. The derivative of ϕ_G with respect to Z is positive, i.e. when Z increases it is optimal to increase f. Thus, terrorists now take into account that it is not only beneficial to have a high number of terror plots Z (i.e. a higher number of successful terror plots), but also to a stronger extent that there is a negative impact due to a stronger response by the government. This negative impact dominates the attractiveness of known terror plots to occupy terror agents.

So information plays a crucial role in the optimal behavior and in the resulting optimal state values. Starting from the open-loop case, there is an increasing effect if both players can only observe one state (government: the known terror plots, terrorists: all terror plots) and an additional negative effect if both players are able to observe both states. The negative effect dominates the positive one.

5.1 Phase portrait analysis

In Figure 1 and Figure 2 we plot the transitional paths of the model. For both figures the grey line represents the solution of the optimal control model (note that in that case the control of the terrorists is set as a constant), the black line the open-loop, the dashed line the mixed open- and closed-loop, and the dotted line the closed-loop case. In Figure 1 we assume that a new terror group has formed. Thus the initial values are $(X_0, Y_0) = (0, 0)$.

[Figure 1 about here.]

We see that if the number of initial terror plots is low, the terrorists will initiate a comparatively high number of terror plots first. Since it is difficult for the government to detect terror plots when their number is low, they initially will not do much, because costs of terror detection would exceed the utility from it.

[4] Note, however, that this is also strongly related to the used parameters. If the costs of terror plots are lower or the utility is higher, results are the opposite.

[5] The reason for the sign of η_y is that for the terrorists a high Y is beneficial since then interdiction prevents agents from detecting new plots.

Thus, the number of terror plots grows. When there are more terror plots the government will increase its efforts to detect them, not only because it is easier to detect them, but also because more damage is caused by successful attacks. The terrorists will initiate less terror plots as due to the higher control efforts of the government planning new plots is less efficient. The number of terror plots grows until the number of newly initiated terror plots equals the number of successful attacks plus the number of detected plots.

In Figure 2 we assume an already established terror group and a government which only starts with counter-actions against this group. To be consistent with the plots of Seidl et al. [15] we assume $(X_0, Y_0) = (40, 0)$.

[Figure 2 about here.]

Due to the high damage caused by terror attacks, and the high efficiency of detection, the government deploys a high number of agents for detection. As the number of unknown terror plots gets lower, the terrorists have to initiate more and more terror plots to be able to launch more successful attacks. Thus, the number of unknown terror plots X decreases and the number of unknown terror plots Y increases. Therefore, the government can decrease its control efforts. Note, however, that after some time the number of unknown terror plots becomes so low that Y then decreases until its steady state value is reached. I.e. Y decreases once the terrorists once the terrorists cannot draw any advantage from terror plots initiated before the government's involvement into counter-terror actions.

5.2 Sensitivity analysis

Our results are mostly derived numerically, therefore it is natural to look at the robustness of the results if some of the key parameters are changed. In Figure 3 we plot the steady state values of three different solutions of the differential game and vary the interdiction rate ρ in the interval $[0.5, 10]$. The other parameters are the same as introduced above. The left panel of Figure 3 shows the steady state values of X and the right one that of Y.

The stability of the steady states remains the same. Similar to the optimal control model (see Seidl et al. [15]) we find that there are more terror plots, both known and unknown, when the interdiction rate is low. This is reasonable since a high interdiction rate means that the known terror plots are handled quickly.

We can see that the smaller ρ the greater is the difference between the information structures used. Interestingly, the value of X does not change crucially for the open- and the closed-loop case, but for the mixed open- and closed-loop case the change is considerable.

Before we reasoned that terrorists consider plots as more valuable in the mixed case as they include the impact of the government's adaptions of the optimal strategy related to a change in the number of known plots, but neglect reaction to a change in the total number of plots. Therefore it makes sense that when the interdiction rate ρ is low, the terrorists will evaluate a known plot as even higher, because for this lower interdiction rate the chances for a successful terror plot are even higher. In the case of closed-loop, however, terrorists do less because they are more aware of the government's reaction to their actions.

[Figure 3 about here.]

The efficiency rate δ is a parameter that is very hard to observe in reality. In Figure 4 we plot the steady state values and vary δ in the interval $[0.02, 0.4]$ (again the rest of the parameters stay the same as above). The left panel shows the steady state values of X and the right one that of Y.

A higher efficiency rate δ means that more terror plots can be detected and interdicted, therefore it makes sense that both state variables decrease in δ. The impact of the used information structure again makes the biggest difference when δ is small. Similarly to the impact of ρ, one can argue that when the efficiency of the detection agents is low and the terrorists do not properly incorporate the government's reaction to a change in the total number of terror plots, terrorists put relatively more efforts into initiating terror plots than otherwise.

[Figure 4 about here.]

6 Conclusions

The present paper provided two major contributions. First, we extend the terror queue model first introduced in Kaplan [9] to account for the fact that not only the government can choose its strategy optimal, but also the terrorists. We find that if the government is not able to observe terrorists by any means and therefore detect any terror plots, it is optimal for the terrorists to initiate terror plots at a constant rate, which provides a motivation for the assumption of a constant attack rate made in Seidl et al. [15].

To analyze the strategic interactions between the government and the terrorists we set up a differential game. Comparing the results of the optimal control problem of the government and the differential game, we find that the government's optimal strategy only differs quantitatively, but not qualitatively. We see how the optimal number of newly initiated terror plots, i.e. the strategy of the terrorists, is linked to the number of known and unknown terror plots. We find that it is optimal to initiate more terror plots if the number of existing terror plots is high than if it is low.

The second main contribution of the paper is to provide a better understanding about the impact of the choice of commonly used assumptions on the underlying information structure in differential games. While the results of the present model do not differ qualitatively, the underlying information structure has a considerable impact on the extent to which the two players apply their control instruments and consequentially on the long-run development of the number of terror plots.

We see that if terrorists think that the government cannot know (or derive) the total number of terror plots (despite having full information) and only take the adaption of the government's strategy into account when the number of known terror plots changes, terrorists will initiate more terror plots than when they see the government's actions as just a function of time or when they incorporate the government's reaction to the change of the total number of plots. We argue that terrorists then evaluate terror plots as more valuable and initiate more plots, which in turn leads to a stronger response of the government in terms of counter-terror actions. When terrorists incorporate the governement's adaptions to a change of both, the number of known and the total number of terror plots, they will attack less as they know that by planning terror plots they strongly provoke counter-actions.

The presented model is purely deterministic and the players have in principle complete information about all the parameters. They know about the initial number of terror plots and can deduce their development by their knowledge about the state equations and about their opponents objective. While the solution of the present model is an important first step to understand the optimal strategies of the government and the terrorists, it provides many possibilities for extensions. An interesting task would be to study the impact of imperfect information. One could consider what happens if the government errs about the initial number of terror plots or if the terrorists wrongly estimate the efficiency of the government's agents. Furthermore, one could include include uncertainties with respect to the detection of terror plots.

References

[1] Appelbaum, B. (2011) *As US agencies put more value on a life.* Business fret. New York Times, p. A1.

[2] Baser and Olsder. (1982) *Dynamic noncooperative game theory*, Academic Press.

[3] Dockner, E., Jorgensen, S., Long, N.V., Soger, G. (2000) *Differential games in economics and management science*. Cambridge University Press.

[4] Enders, W., Sandler, T. (2012) *The Political Economy of Terrorism.* Cambridge University Press, New York.

[5] Feinstein, J. S., & Kaplan, E. H. (2010). Analysis of a strategic terror organization. Journal of Conflict Resolution, 54(2), 281-302.

[6] Grass, D., Caulkins, J.P., Feichtinger, G., Tragler, G., Behrens, D. (2008) *Optimal control of nonlinear processes: with applications in drugs, corruption and terror.* Springer.

[7] Jacobson, D., Kaplan, E. H. (2007) Suicide bombings and targeted killings in (counter-) terror games. Journal of Conflict Resolution, 51(5), 772-792.

[8] Kaplan, E.H., Kress, M., Szechtman, R. (2010) *Confronting entrenched insurgents.* Operations Research 58(2), pp. 329-341.

[9] Kaplan, E.H. (2010) *Terror queues.* Operations Research 58(4-part-1), pp. 773-784.

[10] Kaplan, E.H. (2012) *Estimating the duration of Jihadi terror plots in the United States.* Studies in Conflict & Terrorism, 35(12), pp. 880-894.

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

[12] Kaplan, E.H. (2015) *Socially efficient detection of terror plots.* Oxford Economic Papers 67(1), pp. 104-115.

[13] Cellini, R., Lambertini, L. (2005). R&D incentives and market structure: dynamic analysis. Journal of Optimization Theory and Applications, 126(1), 85-96.

[14] Long, N.V. (2010) *A Survey of Dynamic Games in Economics. Surveys of Theories in Economics and Business Administration, Vol. 1.* World Scientific.

[15] Seidl, A., Kaplan, E.H., Caulkins, J.P., Wrzaczek, S., Feichtinger, G. (2016) *Optimal control of a terror queue.* European Journal of Operational Research, 248(1), 246–256.

[16] Strom, K., Pope, M., Hollywood, J., Weintraub, G., Daye, C., Gemeinhardt, D. (2010) *Building on clues: Examining successes and failures in detecting US terrorist plots 1999-2009.* Institute for Homeland Security Solutions, Research Triangle Park, North Carolina.

[17] Viscusi, W.K., 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), pp. 5-76.

7 Figures

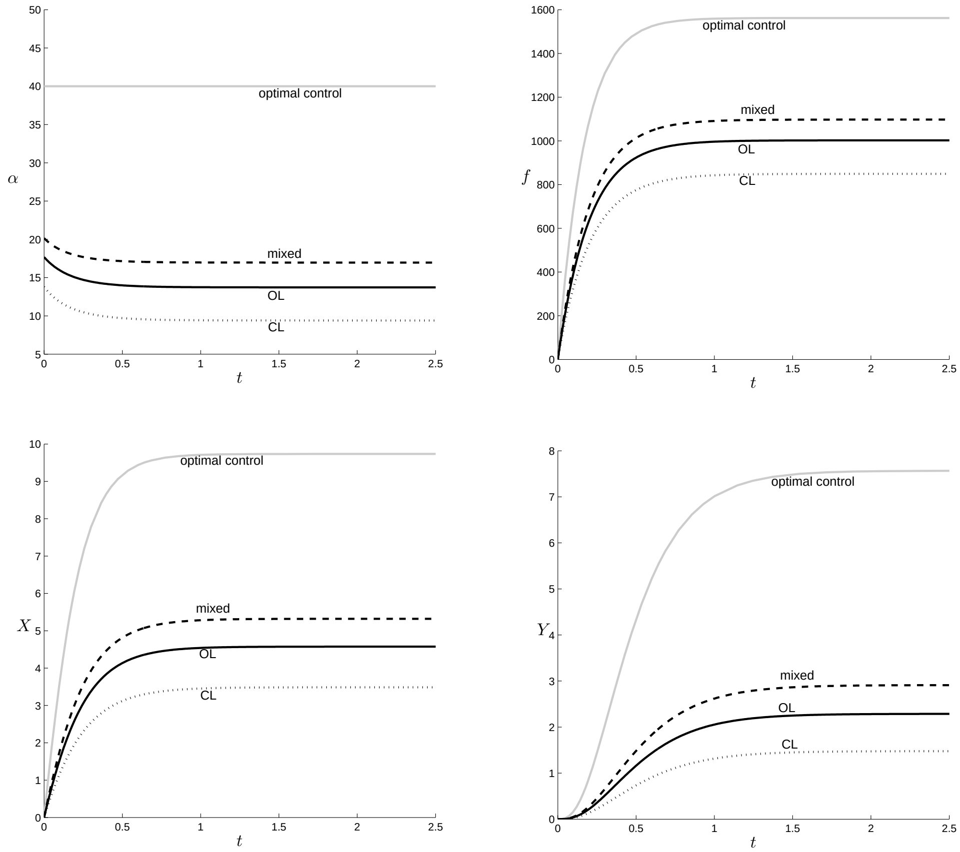

Figure 1: Optimal paths for the controls and the state starting at $(X_0, Y_0) = (0, 0)$

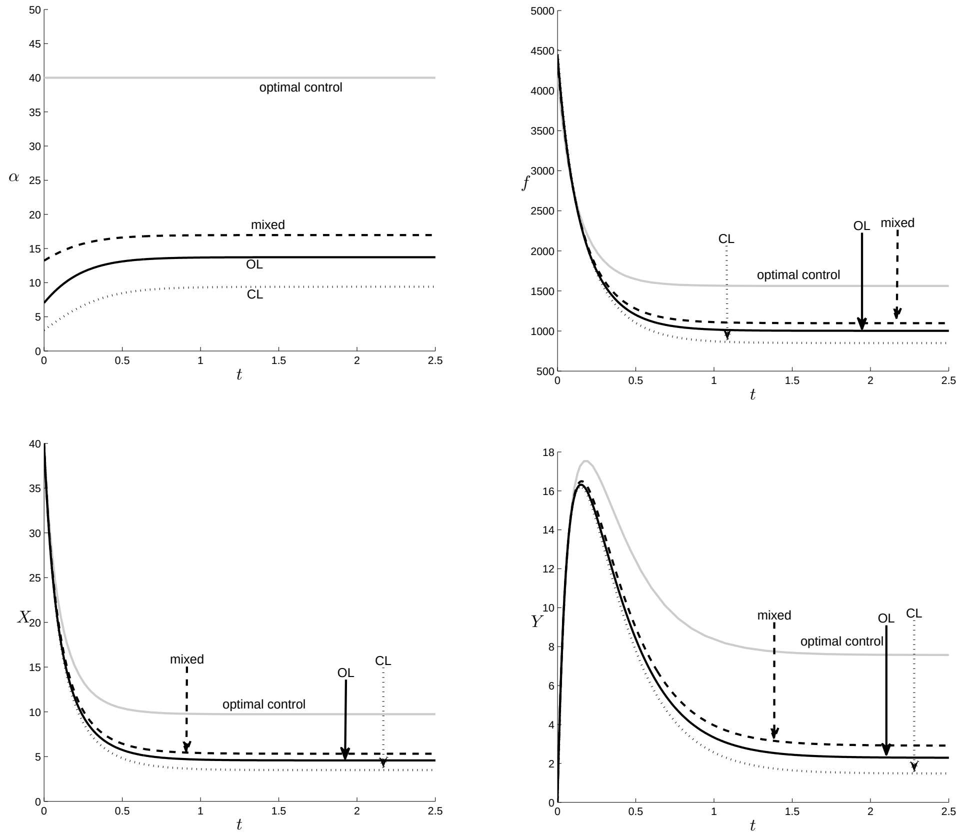

Figure 2: Optimal paths for the controls and the state starting at $(X_0, Y_0) = (40, 0)$

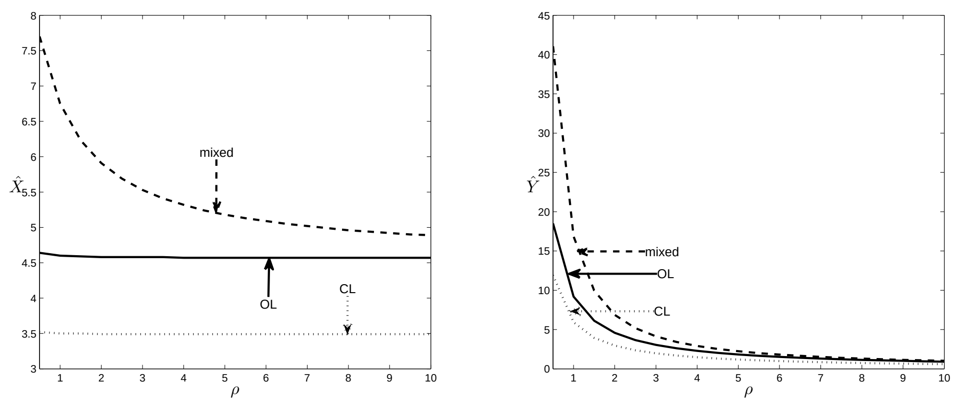

Figure 3: Steady state values of the differential game depending on ρ

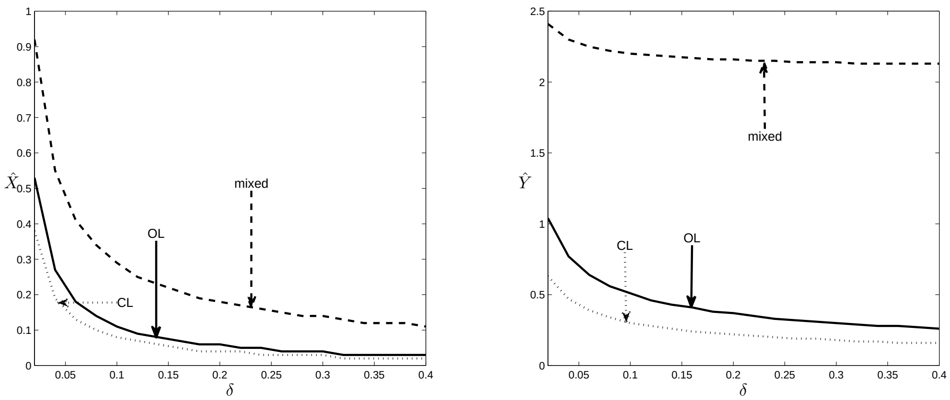

Figure 4: Steady state values of the differential game depending on δ