European Journal of Operational Research 248 (2016) 246–256

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Optimal control of a terror queue

Andrea Seidl[a,*], Edward H. Kaplan[b], Jonathan P. Caulkins[c], Stefan Wrzaczek[a,d], Gustav Feichtinger[a,e]

[a] Department for Operations Research and Control Systems, Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Wiedner Hauptstraße 8, 1040 Vienna, Austria
[b] Yale School of Management, Yale School of Public Health, and Yale School of Engineering and Applied Science, 165 Whitney Avenue, New Haven, CT 6511, USA
[c] Carnegie Mellon University, H. John Heinz III College, 5000 Forbes Avenue, Pittsburgh, PA 15213-3890, USA
[d] Department of Business Administration, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria
[e] Wittgenstein Centre for Demography and Global Human Capital (IIASA, VID/ÖAW, WU), Vienna Institute of Demography/Austrian Academy of Sciences, Wohllebengasse 12-14, 1040 Vienna, Austria

A R T I C L E I N F O

Article history:
Received 19 January 2015
Accepted 6 July 2015
Available online 11 July 2015

Keywords:
Counterterrorism
Optimal control
Queues
Intelligence
Staffing models

A B S T R A C T

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, Grass, Feichtinger, and Tragler (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, Feichtinger, Grass, and Tragler (2009), and Kaplan, Kress, and Szechtman (2010). More recently game-theoretic aspects have been considered by Behrens, Caulkins, Feichtinger, and Tragler (2007), Zhuang and Bier (2007), Feichtinger and Novak (2008), Zhuang, Bier, and Alagoz (2010), Crettez and Hayek (2014); see also Grass, Caulkins, Feichtinger, Tragler, and Behrens (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

* Corresponding author. Tel.: +43 1 58801 105471; fax: +43 1 58801 9 105471.
E-mail addresses: andrea.seidl@tuwien.ac.at (A. Seidl), edward.kaplan@yale.edu (E.H. Kaplan), caulkins@andrew.cmu.edu (J.P. Caulkins), stefan.wrzaczek@tuwien.ac.at (S. Wrzaczek), gustav@eos.tuwien.ac.at (G. Feichtinger).

http://dx.doi.org/10.1016/j.ejor.2015.07.010
0377-2217/

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 & 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 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 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 (teams 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 (terror plot) in service, there is a busy server, which in case of this model is either a single agent or a team of agents. 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

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

[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, Caulkins, and Feichtinger (2001) and Behrens, Caulkins, Tragler, and Feichtinger (2000) are good examples for the usefulness of optimal control methods in the "economics of crime" (more general economics of deviant behavior).