TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

Operations Research and Control Systems

Optimizing counter-terror operations

J.P. Caulkins
G. Feichtinger, D. Grass, G. Tragler

Research Report 2010-17
December, 2010

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

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

Optimizing counter-terror operations

Jonathan P. Caulkins[a]
Gustav Feichtinger[b], Dieter Grass[b], Gernot Tragler[b]

[a]Qatar Campus and H. John Heinz III School of Public Policy and Management,
Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA
E-mail address: Caulkins@cmu.edu (J.P. Caulkins).
[b]Institute for Mathematical Methods in Economics, Vienna University of Technology
Department of Operations Research and Control Systems,
Argentinierstrasse 8, A-1040 Vienna, Austria
E-mail address: gustav@eos.tuwien.ac.at (G.Feichtinger),
orcos@eos.tuwien.ac.at web: http://orcos.tuwien.ac.at

Abstract: This paper deals with the question of how optimal control theory can be applied to design counter-terror measures. Three illustrative examples are presented. In the first model the key controls are two types of counter-terror tactics, one ("water") that does not and one ("fire") that does provoke recruitment of new terrorists. Generally, this model yields two different steady states, one where the terror organization is nearly eradicated and one with a high number of terrorists. The second model extends the approach by including public opinion as a catalysator for counter-terror operations. Again multiple equilibra and some sort of tipping behavior (path-dependence of optimal solutions) occur. The third model analyzes a dynamic game between a terror organization and a target country. Essentially, stationary Nash and Stackelberg equilibra are compared.

Keywords: Optimal control; Terrorism; Counter-terrorism; Social interactions; Pontryagin´s minimum principle; Differential games;

1. INTRODUCTION

In 1968 Gary S. Becker [1] created a new field in economic research – the economics of crime. One of his main questions was: How much offending should be tolerated by our society? While Becker's analysis was essentially *static*, more realistic analysis of 'deviant' behavior use an *intertemporal* (dynamic) approach. Important examples of deviant behavior occur in illicit drug consumption, corruption, violence, and terrorism. The present paper offers three illustrations of how dynamic optimization has been applied to the control of terrorism. In particular, the models use optimal control theory to study the impact of the optimal mix of counter-terror measures.

In what follows the strength of a terroristic organization is reflected in a state variable, x. That variable might most simply be thought of as the number of active terrorists, but more generally could measure the organization's total resources including financial capital, weapons, and technological know-how; compare Keohane and Zeckhauser [7].

The stock of a terrorist organization changes over time by inflows and outflows. Of course these flows can depend on the actions (counter-measures) of the target country, but there is also an epidemiological structure of the inflows, with incidence rates depending on the prevalence, i.e. stock(s) of the system. For example, new terrorists may be primarily recruited by existing terrorists, or simply be more likely to join the organization if the organization appears to be strong. It turns out that those *(social) interactions* imply substantial nonlinearities which may generate *multiple equilibria* and *path-dependence* of the solution paths.

2. MODELS

Model #1: Should one Fight Fire with "Fire" or "Water"? [2]

One concern expressed about aggressive counter-terror operations is that they might turn public opinion against the counter-terror forces, thereby helping terrorist organizations recruit new personnel. For example, Heymann [5] notes that *"recruitment to the Irish Republican Army (IRA) increased sharply during some periods of overly vigorous British action against suspects"*. Likewise Kaplan et al. [6] develop an empirical model that finds that *"killings of terror suspects by Israel sparks estimated recruitment to the terror stock that increases rather than decrease the rate of suicide bombings"*.

It does not seem credible that the best counter-terror strategy is to do nothing. On the other hand, it also does not seem prudent to ignore entirely the possibility that offensive counter-terror operations might be a double-edged sword, stimulating new recruitment even as they kill current terrorists. Presumably, there may be some optimal intensity of counter-terror efforts that balances their obvious benefits with potential negative effects on recruitment and/or some optimal mix of tactics that are more or less prone to stimulate new recruitment.

This paper deals dynamically with this question within an optimal control model, where the key state is the number of terrorists and the key controls are two types of counter-terror tactics, one (*"water"*) that does not and one (*"fire"*) that does provoke additional recruitment of new terrorists. The model is nonlinear and does not admit analytical solutions, but an efficient numerical implementation of Pontryagin's minimum principle allows for solution with base case parameters and considerable sensitivity analysis. Generally, this model yields two different steady states, one where the terror organization is nearly eradicated and one with a high number of terrorists. It turns out that the basic structure of the model solution is quite robust with respect to parameter values, which is helpful since estimating parameter values is very difficult.

We found three possible policy prescriptions, depending on the parameter values:

(1) always drive the number of terrorists down to the low-level steady state, (2) always allow the number of terrorists to approach the high-level steady state, albeit using controls to slow the growth, or (3) drive down the number of terrorists if they are not too numerous initially, but otherwise allow them to grow to the higher level equilibrium.

The third case pertains using our base case parameters, and the state value dividing the regions for which it is optimal to approach the low- and high-level equilibria is a so-called Skiba point, see, e.g. Grass et al [4]. In terms of tactics, it is always essentially optimal to use the targeted (water) tactics, but fire tactics should only be employed if the number of terrorists is above a critical "switching point".

Hence, for many parameter sets, eradicating terrorists should involve greater reliance on fire at first, with water tactics being used predominantly if not exclusively after the initial assault has brought down the number of terrorists. For the base case parameters that switching point is to the right of the Skiba point, so moving away from the Skiba point only involves using water controls in great volume to drive terrorists down to the low-level steady state, and with lesser intensity if one is merely delaying growth to the high-level steady state.

The location of that switching point—and, hence, the circumstances under which fire controls should be used—is moderately sensitive to changes in parameter values. If it moves sufficiently far to the left, it can always be optimal to use fire tactics, although the intensity of fire tactics still seems to be increasing in the number of terrorists. This makes sense; the bigger and more powerful the terrorist threat, the more aggressive one should be in countering it. Conversely, if the switching point moves sufficiently far to the right it may not be optimal to use fire tactics unless the number of terrorists is very large. For more details see Caulkins et al. [2].

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

1st, if indeed counter-terror operations can stimulate recruitment of new terrorists that can have important implications for how the terror war should be fought; even in this bare bones model including that feedback has dramatic effects on the optimal solution.

2nd, even if some counter-terror operations do stimulate substantial new recruitment and other available tactics do not, that does not necessarily imply that tactics which can stimulate recruitment should never be used. Their benefits may exceed their downside if they are sufficiently effective.

3rd, if some counter-terror tactics are more likely than others to be perceived as outrageous or otherwise to stimulate relatively more new terrorist recruitment, then the optimal mix of terror-control operations should vary, perhaps dynamically over time with the changing state of the terrorists.

4th, if recruitment is concave in the number of terrorists and the effectiveness of counter-terror tactics proportional to the number of targets, then the intensity with which those recruitment-stimulating tactics should be employed will generally be increasing in the current number of terrorists. So if the decision is made to eradicate the terrorists, the extent to which fire tactics are used should decrease over time. Water strategies should always be used and are relied on most heavily when the number of terrorists is small.

Model #2: Public Opinion as Catalyst for Counter-Terror [3]

There is also a conventional wisdom that winning the "hearts and minds" of non-combatants is an important strategy in non-conventional conflict. We capture this idea by introducing a two-state optimal control model that explicitly includes a state variable representing the level of public sympathy for the counter-terror forces.

The key innovation in this model is to presume that the outflow from the stock of terrorist is increasing (and concave) in the level of public sympathy for those operations, as well as in the level of counter-terror efforts. The reason for this is that public support encourages the civilian population, within which the terrorists are embedded, to provide information or otherwise assist the counter-terror forces, or at least to refrain from actively helping the terrorists. The analysis yielded interesting results, both mathematically and substantially. From an optimal control perspective, we showed that the long-run optimal outcome can depend on the initial conditions in ways that are not monotonic with respect to the initial value of either state.

That is, we found a so-called *Skiba curve* separating different regions in state space, for which it is optimal to drive the system to steady states with either a lower or a higher number of terrorists. There are places in the state space where a slight increase in the initial number of terrorists can tip the optimal strategy, from approaching the lower-level to approaching the higher-level of terrorists.

But there are other places where the same slight increase can tip one in the other direction. The same odd trait can hold for the initial level of public sympathy. There can be places where a bit more initial sympathy can tip one from approaching the high to approaching the lower level equilibrium, and other places where the same shift tips one in the opposite direction.

It has long been recognized that one state models can exhibit so-called "weak" Skiba points, where one has to stay put if starting exactly at this point, but slight deviations from this initial position can place one on a trajectory leading to different long run steady states. In our model, we found an analogous one dimensional "weak" Skiba curve. If one starts exactly on this curve, the optimal path leads to an intermediate steady state, whereas a slight deviation in initial conditions leads to a small or large steady state.

Indeed, we examine an intermediate case where one part of the curve contains "weak" Skiba points whereas the other part contains "strong" Skiba points.

From an applications perspective, we do not for a minute think that this model, as currently formulated, provides much practical guidance for fighting terrorists. It is far too stylized. However, even these preliminary results make one thing perfectly clear. If there is any merit to the conventional wisdom that public sympathy can *catalyze* the effectiveness of counter-terror operations, then it may be important for counter-terror models to incorporate that fact explicitly. Even in this most elementary model, doing so can have dramatic implications for the results and policy prescriptions.

A principal limitation of the present analysis is the lack of validation of functional forms relating to the public sympathy variable, with respect to both how terrorists' actions generate sympathy for counter-terror forces and how zealous counterterrorism efforts might erode that sympathy. Hence, a fruitful next step would be empirical analysis of longitudinal data on terror attacks, counter-terror operations, and public opinion.

There are databases on terror attacks (e.g., the RAND Worldwide Terrorism Incident Database). Content analysis of media coverage might be a sufficient proxy for public opinion. (e.g., all the major parties in Iraq have radio stations and publish newspapers). Data on the tempo of counter-terror operations is often classified, but might be made available to scholars working closely with the military of the nation(s) involved.For a more detailed analysis of the reputation model see Caulkins et al [3].

Model #3: A Post September 11th Game on Terrorism [4]

This model is also based upon the description of the evolution of a stock of terrorists over time discussed in the introduction. Thus all arguments given there carry over. There is, however, one significant difference.

Whereas the terrorism models in the previous sections have only one decision maker, here we consider the strategic interaction between two decision-makers. We shall refer to the decision-maker on the terrorist side as ITO, an abbreviation for some International Terror Organization.

The ITO's opponent is assumed to be the government of a country that is at high risk of being a target of a future terror attack; for ease of exposition referred to as "the target country." Both the ITO and the target country have access to (at least) one instrument for altering the status quo. The ITO can choose the rate at which it commits terror attacks; the target country can choose the rate at which it counterattacks. In times of "war", attacks by both sides reduce the stock of terrorists in contrast to "peacetime," when neither side is making attacks.

In a post September 11th scenario, it seems plausible that the mere existence of terrorists is perceived as being a perpetual threat to any likely target country. Thus, at any instant of time, both the ITO and the target country have to decide whether or not to launch an attack in the "terror war." The state variable of the resulting (modeled) conflict corresponds to the number of terrorists (denoted by x).

The ITO's attacks (denoted by v) reduce the number of ITO members. It is obvious that suicide attacks reduce the stock but even non-suicide attacks will tend to reduce the size of the ITO because every attack exposes the organization to law enforcement risks and its operatives to an increased likelihood of imprisonment or death.

Terror-control activities pursued with intensity u (=control variable of the target country) kill, capture, or incapacitate terrorists; that is their purpose. The greater the intensity of terror-control interventions, the more resources there are that can be devoted to investigating the aftermath of an attack. These interventions may in turn provide information about other ITO operatives, so that subsequent attackers can be effectively stopped. Therefore, the greater the target country's intensity of terror-control measures, the higher is the probability of finding and arresting ITO operatives.

The target country (player 1) has three policy goals: to minimize the attacks of ITO, v, to minimize the power of the ITO, x, and to minimize the costs of terror-control, u. Thus, the objective of player 1 is a weighted sum of three components. Note that sometimes the costs for counter-terror measures can be neglected because they can be much smaller than the other terms.

ITO (player 2) wants become powerful (large x) and likes successful attacks (v), but terror attacks are costly (convex in v). In addition, ITO may have political objectives aimed at inducing "excessive" counter-attacks by player 1 eliciting high values of u as an indirect way of stirring up sentiments against the target country. This could, for instance, smooth the way for changes in local regimes. For a more detailed model formulation compare Grass et al [5].

As a reminder of some basic definitions, a **Nash equilibrium** is one in which none of the *simultaneous* acting players has an incentive to depart from her/his/its own Nash strategy as no other players do. A Leader-Follower Game and associated **Stackelberg equilibrium** pertains when one of the players, the leader, has the opportunity to move first, i.e. *before* the opponent, the follower, makes his/her decision. So there is a *hierarchical* (sequential) situation.

We determine the Stackelberg equilibrium with the ITO as leader. We consider the open-loop Stackelberg game with the ITO (player 2) as leader and the target country (player 1) as follower. We assume that the cost of terror control is perceived to be negligible. The Stackelberg solution of the game requires the following three steps:

Step 1: The ITO announces its strategy v.

Step 2: The target county solves, for this given v, the same control problem as in the Nash case.

Step 3: Using the follower's reaction function, the leader solves the optimal control problem, where the adjoint variable of the player 1 acts as second state variable of the player 2. In our analysis the stationary Nash and Stackelberg solutions are computed and compared, Remarkably, since the game is state-separable, any open-loop Nash solution qualifies as a Markov perfect equilibrium.

The qualitative properties of the stationary Stackelberg solution are essentially the same as the ones of the stationary Nash solutions, but not entirely. We study the two solution's similarities and differences. It turns out that

$$\hat{u}_S > \hat{u}_N, \text{ and } \quad \hat{v}_S > \hat{v}_N \qquad (*)$$

Thus, the war on terror will be more intense if the ITO has a first mover advantage (and a political objective) and announces the number of terror attacks to be carried out to the target country (compared to the case where both players act at the same time). Thus the ITO being the Stackelberg-leader, leads to a higher number of terror attacks and more intensive counter-terror interventions, and, thus, a more intensive conflict.

Since both the Stackelberg leader and follower increase their efforts in contrast with the players in the Nash equilibrium, it is interesting to analyze the effect of (*) on the equilibrium. The difference between the equilibrium values of $\hat{x}_S$ and $\hat{x}_N$ is positive, and therefore we can conclude that the ITO being the leader yields an increase in the size of the ITO,

$$\hat{x}_S > \hat{x}_N.$$

We conclude the comparison of the results of the two solution concepts with a comparison of their objective function values.

It turns out that the difference between the objective function values, is positive; the ITO is better off leading when playing the Stackelberg strategy than when playing the Nash strategy.

Inasmuch as being a Stackelberg leader involves committing upfront to a course of action, whereas the Nash concept involves simultaneous and hence more interactive decision-making, we have uncovered a logical reason why an ITO might want to cultivate a reputation for being fanatically devoted to a course of action. Simply put, the ITO achieves a better outcome (objective function value) if it can convince the target country to play a Stackelberg game in which the ITO's fanatical (superficially 'irrational') favor gives it the first-mover advantage.

These results provide an informative picture of possible 'war on terror' scenarios and the impact of the many possible types of interaction on the results. This is remarkable because the model is stunningly simple (while allowing for analytical calculations). With these pieces of information decision-makers can gain an understanding of the interaction between terrorists and likely target countries (both striking back and attacking) and insight into the consequences of policy-making. Better informed policy-makers can then make better policy.

3. CONCLUSIONS

There is an emerging literature on quantitative analysis of terrorism and counter-terror oprations, e.g. [7], [5], [6] and [8]

In the present contribution we sketched three intertemporal optimization models to design efficient counter terror operations. Although they are highly stylized cartoons of the reality, they may deliver structural insights into the dynamics and control of terrorism and counter-terrorism. Since the data situation is notoriously poor, Pontryagin´s mimimum principle provides an appropriate tool to analyze models of this type. The results obtained by these optimality conditions mainly rely on *qualitative* properties of their function involved.

There are many applications of dynamic optimization to socio-economic problems. Real-world applications, however, involve more than one decision-maker. To exemplify the situation in the economics of crime (more generally: economics of "deviant" behavior): decision making for illicit drug problems, corruption, violence, and terrorism is confronted with rational or bounded rational behavior of offenders. Drug users interact with drug dealers, and authorities, bribers with bribees, terrorists react to counter terror operations etc. The adequate framework to deal with interactions of decision-makers including strategic competition is game theory.

Since the dynamics of the involved processes is crucial, differential games provide the appropriate toolkit to analyze optimal control in a competitive (non-cooperative or cooperative) way.

4. REFERENCES

[1] Becker, G.S. (1968) Crime and punishment: An economic approach, *The Journal of Political Economy*, 76(2):169-217.

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

[3] Caulkins J.P., Feichtinger G., Grass D., Tragler G. (2009) Optimal control of terrorism and global reputation: A case study with novel treshold behavior, *Operations Research Letters*, 37(6):387 - 391

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

[5] Heymann P. (2003) Dealing with terrorism after September 11, 2001: an overview. In: Kowitt A, Pangi R, editors. *Countering terrorism: dimensions of preparedness*. Cambridge, MA: MIT Press; p. 57–72

[6] Kaplan E.H., Mintz A., Mishal S., Samban C. (2005) What happened to suicide bombings in Israel? Insights from a terror stock model. *Studies in Conflict and Terrorism*; 28(3):225–35

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

[8] Memon N., Farley J.D., Hicks D.L., Rosenorn T. (2009) *Mathematical Methods in Counterterorism*. Springer Wien New York