Analysis of a Strategic Terror Organization[1]

Jonathan S. Feinstein[2], Edward H. Kaplan[3]

July 2009
Revised September, October 2009

Mailing address:

Profs. Jonathan S. Feinstein and Edward H. Kaplan
Yale School of Management
135 Prospect Street
New Haven, CT 06511

Feinstein phone / e-mail: 203-432-5975 / jonathan.feinstein@yale.edu
Kaplan phone / e-mail: 203-432-6031 / edward.kaplan@yale.edu

Keywords: terrorism; expected utility; capabilities; counterterrorism

Note: the keywords above were selected given the choices available on the JCR website. If allowed to be more specific, our keywords would be: terrorism; terrorist decision making; evolution of terror organizations; timing and scale of terror attacks; mathematical model; backwards induction

[1] We thank the Daniel Rose Technion-Yale Initiative in Homeland Security and Counterterror Operations Research for support.

[2] John G. Searle Professor of Economics and Management, Yale School of Management (jonathan.feinstein@yale.edu)

[3] William N. and Marie A. Beach Professor of Management Sciences, Yale School of Management; Professor of Public Health, Yale School of Public Health; and Professor of Engineering, Yale School of Engineering and Applied Science (edward.kaplan@yale.edu)

Analysis of a Strategic Terror Organization

July 2009; Revised September, October 2009

Abstract

We model a terrorist organization's choice over the scale and planning horizon of terror attacks and the consequences for the organization's evolution. The organization can engage in short-term attacks planned and executed in a single period, characterized by a low fixed cost and relatively high marginal cost; and longer term attacks planned and executed over two periods, having a high fixed cost but relatively low marginal cost. Longer term attacks require more resources and cause more damage if successful. Successful attacks increase the organization's size; in addition the organization has a natural growth rate. Attacks can fail due to failed execution or counter-terror interdiction. In a two period version of this model we analyze the terror organization's attack decisions. We use simulations to characterize optimal strategies and explore their implications for the growth of the organization. We identify a set of strategic regimes and our results show that they always occur in a fixed order as a function of the organization's initial strength.

Introduction

In this paper we present a model of a strategic terror organization that plans and pursues attacks to maximize its growth, through a mechanism whereby successful attacks attract new members, increasing the organization's strength. Our particular focus is on terror organizations that act strategically and in pursuit of long-term objectives, formulating attacks over multiple periods, and evolving based on attack outcomes. In particular, we use formal decision-theoretic modeling to explore the nature of the strategies such organizations will pursue. We also explore the impact of counter-terror enforcement in deterring a terror organization from launching attacks and growing, though we do not treat the government as a strategic player given our focus on the evolution of terror organizations.[1]

In our model, a terrorist organization has access to two kinds of attack technologies. One is a smaller, simpler technology, characterized by a relatively low fixed cost of resources, but a relatively high marginal cost of increasing the scale of an attack. An example would be an attack with conventional firearms with modest planning. The other is a larger, more complex, and more destructive technology, characterized by a high fixed cost but a low marginal cost of increasing the scale of an attack. Examples include larger organized attacks like 9-11 and a large-scale biological attack (e.g. Kaplan, Craft and Wein 2002; Wein, Craft and Kaplan 2003). The Homeland Security Council's list of possible terror attack scenarios provides examples of both small and large attacks that correspond to what we have in mind (Homeland Security Council 2004). A small attack can be planned and executed in one period, whereas a large attack requires two periods from inception to execution. The fixed cost of an attack can be thought of as including the training needed to engage in an attack, as well as the basic resources required for planning an attack (regardless of size). The variable cost is the cost required for a given scale attack, and increases as the scale increases – it will include weapons, transportation, and delivery. Rosendorff and Sandler (2004) present a related one period model in which a terror organization can engage in either a normal or "spectacular" attack. Successful attacks benefit the organization by increasing its membership (and more generally its strength). Two additional important parameters in our model are the terror organization's "natural" rate of growth (or, possibly, decline), and the government's probability of detecting and thwarting an attack.

While the two attack technologies in this paper are stylized, we believe they capture an essential aspect of strategic terror: the choice between smaller and simpler vs. larger, more complex attacks requiring more planning. Understanding how terror organizations evaluate this choice and what factors drive them to be willing to undertake a large-scale attack is important for evaluating terror risk and formulating counter-terrorism policy. Only an organization with sufficient resources will choose to pursue a large attack, and one important question is how large an organization will be before undertaking such an attack. We show that in the context of our model a terror organization will in fact choose to pursue a large-scale attack only when it has grown to a size well in excess of the absolute minimum needed to pursue such an attack - at smaller sizes the organization prefers to engage in a series of smaller attacks, building its capability. While we focus on growing organizations, we have also explored the case of an organization that is declining and find that such an organization will be willing to "go for it all" and attempt a large-scale attack at a size threshold that is lower than the corresponding threshold for an organization that is growing.

We take as given the terror organization's initial size and analyze how its strategy in terms of attack planning depends on its size, and how its strategy and attack outcomes determine its evolution. Through extensive simulations we characterize the terror organization's strategy as a set of regimes, each regime corresponding to a particular pattern of attacks. We describe these regimes and exhibit them for a few representative cases. An interesting finding from the simulations is that the terror organization's objective, which in our context is the expected size of the organization at the end of the model's timeline, is convex in the initial size of the organization, so that as the organization grows it becomes an increasing threat at an increasing rate. From the viewpoint of counterterrorism this result highlights how important it is to keep terror organizations small and not let them reach the point on their growth curve where they can begin to grow rapidly and launch more frequent and larger attacks.

We also study the impact of the other key parameters in the model, the terror organization's natural rate of growth and the government's probability of thwarting an attack. We find that the degree of convexity of the value function itself depends on the interactions of certain parameters. Thus the convexity is greater when government enforcement is weaker, so that the government must be especially vigilant in crafting a strong counter-terrorism detection force when facing a relatively larger terror organization. Further, there is an interaction between the degree of government enforcement and the natural rate of growth of the terror organization: when the natural rate of growth is relatively high the government must again be especially vigilant, else the terror organization may (in expectation) grow rapidly.

We select parameter values for our model simulations from a qualitative reading of a series of terror attacks, many associated with Al Qaeda, over the past decade, including 9-11, the London 7/07/05 bombings, the Madrid 3/11/04 train bombings, the Istanbul truck bombings of 2003, and the series of terror attacks in Indonesia over 2003-05.[2] We evaluate all costs and scale effects in terms of human resource requirements to plan and execute an attack. More generally, both money and manpower are required to launch an attack, but the recent history points to the human resource requirements as crucial, thus we focus on these.

Much of the literature on terror organizations emphasizes the importance of viewing these organizations as networks (for examples see Sageman (2004, 2008) and Enders and Jindapoon (2010 [this issue]); Johnson et. al. (2006) provide an interesting description of a situation in which terror cells coalesce in an environment without higher level organizational structure). We think of the organization we study as a network, with its size measured by the number of affiliated members. But more work could be done extending our model to identify how the network structure of the organization relates to identification of terror attack opportunities and the implementation of an attack. For example, the organization may be structured in layers – leaders, operational managers, and foot soldiers who carry out tasks and attacks. In the context of our model, the production of an attack would require human resources from all levels, and in turn the organization's growth could be modeled in terms of the growth and training of individuals at the different levels.

The remainder of the paper is organized as follows. In the next section we present our model of a terror organization, its attack decisions, the factors that determine success or failure of an attack, including counter-terrorism detection, and our specification of how the organization evolves. In the third section, we outline how the model is solved and present the basic structure of the optimal strategies for the terror organization. The fourth section contains additional simulations and the final section concluding remarks.

Model

The organization is characterized by its overall size, or strength, which we denote by q. Size is a single summary measure that is meant to capture several different dimensions of the organization. The most basic dimension is the number of active members who can play a role in planning and executing attacks. We will describe much of our model in these terms. In particular, we focus on the human resource requirements for an attack and assume that all other costs can be converted to the "currency" of human resources. This approach is consistent with the literature on terror organizations, which emphasizes that human resources are the greatest constraint on such organizations, limiting the number and scale of attacks that may be engaged in at any point in time; the monetary cost of attacks is generally fairly small, and the technology of attacks, at least to this point, has been relatively simple (Enders and Sandler 2006; Financial Action Task Force 2008; National Research Council 2007). In a broader interpretation, q can be thought of as capturing the overall capability of the terror organization or "terror stock" (Keohane and Zeckhauser 2003; Kaplan et al. 2005). Capability depends on membership size, but may also depend on how well the organization functions, for example the degree to which orders are followed and the skill of the organization's leaders.

We study the terror organization over two periods, focusing on the attacks it attempts to carry out and the evolution of q. At the beginning of the first period the organization's q is q_0, which we take as given. The organization's objective is to maximize its expected value of q at the end of period 2, its planning horizon. Its focus is thus on its own growth as an organization. This can be viewed as an interim goal: if the organization grows and becomes stronger, it will be able to achieve its larger objectives. Clearly there are other objectives the organization is also likely to value, such as the damage it inflicts on society or specific individuals or groups and the achievement of specific political goals. If we assume that damage the organization inflicts through its attacks is proportional to the benefits it realizes from these attacks, then this damage is essentially incorporated in q. However, political or other objectives are likely to be related in a more complicated way to the organization's strength and attacks and are not fully incorporated in our model, though we believe it should be possible to enrich the model so that the terror organization's objective includes more terms or the organization has multiple objectives.

Terror Attacks

We distinguish between two different kinds of attacks. One is a small scale, short-term attack that can be planned and executed in a single period. The technology to produce this kind of attack requires a modest fixed cost and a significant marginal cost that scales with the size of the attack. In particular, to launch a small scale attack of scale s requires a fixed cost of K_s and an additional outlay $c_s s$. In our formulation, we think of all costs as denominated in the currency of human resources, for example man-years. Thus K_s is the number of man-years required to plan and execute, regardless of the scale of the attack, and $c_s s$ is the number of additional man-years required to plan and execute an attack of size s. If a small scale attack is successful, the organization reaps benefits $b_s s$ from the attack. We think of these benefits as increasing the organization's q, through increasing recruitment of new members and possibly through other channels as well. Examples of small scale attacks are attacks with conventional weapons like automatic rifles on standard targets and car bombings.

The other kind of attack is a large scale attack that is planned and executed over two periods. The technology to produce this kind of attack is characterized by a large fixed cost and a small marginal cost. In particular, the fixed cost required for this kind of attack is K_b, with K_b much larger than K_s. The marginal cost is $c_b s_b$. If a large scale attack is successful, the organization reaps benefits $b_b s_b$ from the attack. An example of this kind of attack would be a multi-pronged attack on a city or military installation. A large scale attack, as we describe, generally requires more resources than a small scale attack (in the relevant range in which s for a small attack is not too large), but yields higher benefits if successful. We assume also that b_b/s_b is significantly larger than b_s/s_s. We assume that if an attack fails it generates zero benefits. It is possible to generalize our model to allow for the possibility that an attack fails in such a way as to actually reduce the organization's strength even further, reducing q. We think this is more likely for a large-scale attack and discuss this possibility further below, but we do not model this outcome formally.

We assume that once human resources are devoted to an attack, they are lost to the organization and not available for any subsequent attacks. This may happen either because the perpetrators die - suicide missions, or because they become known to the authorities and are therefore no longer effective. Implicitly, we are thus focusing on human resource costs at the operational level of individuals "on the ground." Typically, the organization's leadership is available over time for successive attacks.

A terror attack is not guaranteed to succeed. It may fail, either because the organization is unable to carry out the attack successfully, or because the authorities thwart the attack. The ability of the organization to carry out an attack successfully depends on the scale of the attack as compared with the scale (q) of the organization. Specifically, the probability that the organization successfully carries out an attack of size s is given by

$$1 - \frac{c_s s}{q - K_s}$$

for a small attack, and

$$1 - \frac{c_b s}{q - K_b}$$

for a large attack.[3] These formulas link the probability of success to the ratio of the size of attack, multiplied by the cost parameter c which converts size to human resources, to the resources left after the fixed cost of the attack is covered. As we show below, and intuitively, this factor tends to lead the organization to choose attacks that are of medium size compared with its available resources: it does not wish to plan and attempt to carry out an attack that is close to exhausting all available resources, because such an attack strains the capabilities of the organization to the point where the probability it can successfully carry out the attack is small.

Even if the organization is able to plan and carry out an attack, the attack can still fail if the authorities thwart the attack. As our focus on this paper is not on deriving optimal government enforcement strategies, we choose a simple specification for this possibility. Specifically, we let $1-\theta$ denote the probability that an attack is thwarted. We assume this probability is the same for small and large attacks, and that whether the authorities thwart any given attack is independent of whether they thwart any other attacks. More generally, we could allow $1-\theta$ to be different for large and small attacks. Intuitively, the probability of thwarting a large, "spectacular" attack may be higher, especially as such an attack requires two periods for planning and execution. However, modeling this distinction does not provide sufficient additional insights, beyond confirming the evident intuition that a large attack is relatively less attractive if the probability it will be thwarted is greater, relative to the added complexity of solving the model. Thus we choose not to explore this issue further.

Combining the two possible sources of attack failure generates the overall probability an attack is successful. Thus, for a small scale attack the probability of success is

$$(1-\frac{c_s s}{q-K_s})\theta\,.$$

Decisions, Outcomes, and Evolution of q

At the beginning of the first period the terror organization decides first whether to engage in a large scale attack, and, if so, of what scale. If the organization decides to engage in a large scale attack of size s_b, then the resources required for this attack are subtracted from q_0, defining

$$q'_0 = q_0 - K_b - c_b s_b.^{4}$$

We require q'_0 to be greater than zero. In turn, this implies that the organization can only engage in a large scale attack if $q_0 > K_b$, and can only engage in a large scale attack of size s_b if

$$q_0 - K_b - c_b s_b > 0.$$

Thus, the initial size of the organization limits its strategic options. Following the discussion above, the probability an attack of size s_b is successful is

$$(1 - \frac{c_b s_b}{q_0 - K_b})\theta.$$

After the organization makes its decision about the large scale attack, it decides whether or not to engage in a small attack during period 1, and, if so, of what scale. The expressions for the small attack are based on the organization's capability available after allowing for any resources devoted to a large scale attack. For convenience, we denote this available capability as q'_0, but note that this is equal to q_0 if there is no large scale attack undertaken. Thus when a large attack is undertaken, any small attacks must be planned and executed using other individuals (and resources) in the organization. A small attack is feasible only if

$$q_0' > K_s$$

and a small attack of scale s_1 is feasible only if

$$q_0' - K_s - c_s s_1 > 0.$$

The probability a small attack of size s_1 is successful is

$$(1 - \frac{c_s s_1}{q_0' - K_s})\theta.$$

Once the organization has made its decisions about the large attack and first period small attack, it spends the duration of the period engaged in planning whichever attacks it has decided upon, and, if it has decided to launch a small attack, attempts to carry out this attack. We do not model these processes, though the model could be enriched to model the staging of an attack.

At the end of period one the outcome of the small attack, if there is one, is learned. If the attack succeeds it generates benefits $b_s s_1$ which in our model accrue to the organization through increasing q. These “benefits” are presumably costs to society as a whole, which do not enter directly in our model.

At the close of period one q is updated. Recall that the initial q_0 has been diminished by whatever resources have been devoted to planning and executing attacks. Two factors may now increase q. One is a successful attack: if a small attack succeeds q is increased to

$$q_0' + b_s s_1$$

where q_0' is equal to

$$q_0 - K_b 1_{\{s_b > 0\}} - K_s 1_{\{s_1 > 0\}} - c_b s_b - c_s s_1$$

($1_{\{A\}}$ equals one if the event A is true and zero otherwise). As noted above, if an attack fails there is no further impact on q.

The second factor that generates an increase in q is an exogenous growth rate r, which can also be thought of as a recruitment rate (Faria and Arce 2005). When r is high, the organization has a high natural rate of growth, and this can influence its strategic choices, for example about whether to engage in a first period small attack or wait until the second period to launch a small attack.[5] We model exogenous growth as a multiplier on the value of q. Applying this multiplier generates the value q_1 entering period 2:

$$q_1 = (1 + r)(q_0' + 1_{\{\text{period 1 small attack succeeds}\}} b_s s_1).$$

Throughout the body of the paper we focus on the case in which $r \geq 0$. However, all of our analysis continues to apply when r is negative as long as it is greater than -1, so that $1 + r > 0$. The case of r being positive is a growing organization, which perhaps poses a greater long-term threat. Such an organization is more likely to be at an earlier stage of its life-cycle. The case of r negative is a declining organization. Such an organization may have been more powerful earlier, and may still be large (a large q) but is typically at a later point in its life-cycle. Note that r and q are conceptually and empirically distinct. An organization might be small (small q) but be growing quickly (high r), or might be large (large q) but declining (negative r). Later we discuss, for our basic results, an interesting difference in strategic choices made by these two kinds of organizations.

If the organization has engaged in a first period attack, q_1 has two possible values, and its realization is known at the end of the first period, while if the organization has not engaged in a first period attack q_1 has one possible value.

In period 2 the organization has only a single decision, whether to engage in a small attack or not. If the organization decided to undertake a large scale attack in period 1, the planning and execution of this attack continues into period 2, but there are no further decisions about this attack.

The organization's decision about whether or not to launch a period two small attack is made in light of its strength as of the beginning of the period, q_1. The organization can undertake an attack only if it has sufficient strength: q_1 must exceed K_s and for an attack of size s_2, q_1 must exceed $K_s + c_s s_2$.

If the organization engaged in a period one small attack, then in general its decision about a period two attack depends on the outcome of the period one attack. In particular, q_1 is higher when a period one attack succeeds than when it fails. In fact it is easy to show that if the organization engages in a period one attack that fails and chooses to engage in a period two attack, then it will also choose to engage in a period two attack if the period one attack succeeds. If the organization chooses to undertake an attack of size s_2, the probability the attack succeeds equals

$$(1 - \frac{c_s s_2}{q_1 - K_s})\theta.$$

Once the organization has made its decision about a period 2 small attack, it spends the duration of the period planning and executing its active attacks. At the end of period 2 the outcomes of these attacks are realized. Each attack can either succeed or fail and if there is more than one attack underway (large and small) their outcomes are independent. If a large scale attack succeeds, benefits $b_b s_b$ are generated for the terror organization. If a small attack has been engaged in and succeeds, benefits $b_s s_2$ are generated. As noted earlier, we assume that an attack that fails has no further impact on q. It is possible to extend our model to allow for a failure to have a negative impact on q. This seems more plausible for a large attack: a spectacular attack can go spectacularly wrong, and this might cost the organization in terms of its popularity. For example, the November 9, 2005 bombings at three hotels in Amman, Jordan that killed many Jordanian citizens (including wedding guests among other victims) produced a backlash against Al Qaeda in Iraq and its leader Abu Musab al-Zarqawi who claimed responsibility for the attack (Fattah 2005). While such a case is interesting, it adds complexity to our model and we do not formally pursue it in this paper (but see our discussion around Figure 1). Finally, q_1 is updated to its terminal value q_2:

$$q_2 = q_1 - 1_{\{s_2>0\}} K_s - c_s s_2 + 1_{\{\text{period 2 small attack succeeds}\}} b_s s_2 + 1_{\{\text{big attack succeeds}\}} b_b s_b.$$

The terror organization's objective is to maximize the expected value of q_2.

Model Solution: Structure

We solve the model and display results as a function of q_0, the strength of the terror organization at the beginning of period 1. We show that the q_0 dimension divides into intervals with each interval corresponding to a given strategy. Thus, our model predicts which kinds of attacks the organization will engage in as a function of q_0. We also compute the value function for the organization - the expected value of q_2 - as a function of q_0. The intervals in the q_0 dimension and the value function depend on the full set of parameters in the model, including the costs of large and small attacks, the benefits to the organization of successful attacks, government detection (θ), and r.

We solve the model using rollback. First we solve for the organization's optimal decision in period 2 regarding a small attack (whether to launch an attack, and if so its scale). Then we solve for the organization's optimal decisions in period 1: working backwards within this period we first solve for the organization's decision about a small attack, then for its decision concerning a large scale attack.

The state variable in the model is q. The value of q entering a period is what guides the organization's choices in that period. Thus the period 2 decision concerning a small attack is conditioned on q_1 (note that because the value function is linear in outcomes, the organization's decision in period 2 does not directly depend on whether or not it is engaged in a large scale attack, even though such an attack is ongoing in period 2). Period 1 decisions are conditioned on q_0: the large scale attack decision is made directly based on q_0, while the small attack decision is made based on q_0'.

Period 2

In period 2 the organization decides whether to engage in a small attack and, if so its scale s_2. Because the objective function for the organization – maximization of the expected value of q_2 – is linear in the different attack outcomes, the maximization problem associated with this decision is simply:

$$\max_{s_2} \; [\theta b_s s_2(1 - \frac{c_s s_2}{q_1 - K_s}) - c_s s_2 - K_s, 0]$$

subject to the constraint:

$$c_s s_2 + K_s \leq q_1.$$

The first-order condition for this problem when $s_2 > 0$ is

$$\theta b_s - c_s - \frac{2\theta b_s c_s}{q_1 - K_s} = 0.$$

We define the parameter

$$\gamma = 1 - \frac{c_s}{\theta b_s}.$$

The first-order condition above can then be written as

$$s_2 = \frac{\gamma}{2c_s}(q_1 - K_s).$$

To determine the range of $q_1$for which $s_2 > 0$, we compare the value of the objective function for this choice of s_2 to 0. The objective function at the optimum interior value of s_2 is

$$(q_1 - K_s)(\theta b_s \frac{\gamma}{2c_s}(1 - \gamma/2) - \gamma/2) - K_s.$$

Comparing this expression to zero, s_2 is chosen as its positive solution, given above, whenever

$$q_1 > K_s(1 + \frac{2c_s}{\gamma\theta b_s(1 - \gamma/2) - \gamma c_s}).$$

Thus, the solution to the organization's decision problem in period 2 is

$$s_2^* = \begin{cases} \frac{\gamma}{2c_s}(q_1 - K_s) & \text{if} \quad q_1 > K_s(1 + \frac{2c_s}{\gamma\theta b_s(1-\gamma/2)-\gamma c_s}) \\ \\ 0 & \text{otherwise} \end{cases}.$$

This solution is intuitive. The optimal scale of an attack depends on two factors. One is γ, which essentially is a cost-benefit ratio: θb_s is the expected benefit of an attack and c_s is the cost. If this ratio is close to one, γ is close to zero and the optimal attack is very small. As this ratio increases, γ rises towards one and the optimal scale attack rises. The other factor is the residual organizational capacity available after the fixed cost of the attack is taken into account: $q_1 - K_s$. Recall that the chance of launching a successful attack depends on the ratio of the attack scale to available organizational capability. The implication of this assumption, worked through the mathematics of the first-order condition, is that the optimal attack scale increases linearly as available capacity rises. Because there is a fixed cost to engage in an attack it will only pay to engage in an attack if the scale can be made large enough. This is the logic behind the threshold condition on $q_1$which determines whether the organization chooses to engage in an attack.

If the organization chooses to engage in a small attack in period 1, then q_1 has two possible values, depending on whether this attack succeeds or fails. In circumstances where the organization chooses not to launch a small attack in period one, it may or may not find it optimal to launch a small attack in period two, again depending on q_1. The key parameter in this circumstance is r. If r is high enough, the organization may find it optimal not to engage in a small attack in period one (it may or may not engage in a big attack) and then launch a small attack in period two.

Period 1

In period 1 the terror organization first decides whether to engage in a large scale attack, and then decides whether to engage in a small attack. In this section we discuss the decision about the small attack, and the structure of the solution over both periods when only small attacks are under consideration. We add the large attack in the next section. Figure 1 depicts the organization's small attacks decision problem over both periods.

[Figure 1 near here]

The organization's decision about whether or not to launch a small attack in period 1 depends on q_0'. In making its decision about a period 1 attack, the organization must take into account the impact such an attack will have on period 2 decisions and outcomes since its objective is to maximize the expected value of q_2. In particular, the period 1 attack may either succeed or fail and each outcome is associated with a value for q_1, which in turn impacts the period 2 small attack decision. Because a successful attack increases q it is easy to show that if the organization engages in an attack in period 1 that succeeds, it will always choose to launch a small attack in period 2. But if the organization launches a period 1 small attack that fails, it may or may not choose to launch a period 2 small attack.[6] Thus in analyzing the period 1 small attack decision, we can assume that a successful attack is always followed by a period 2 small attack, but for a period 1 attack that fails we must explicitly model whether or not a period 2 small attack is launched. Recall from the previous section that the organization's period 2 decision about whether or not to launch a small attack depends on q_1: whenever q_1 exceeds a threshold it engages in an attack. Given q_0', we can compute the value of s_1 such that, if the period 1 attack fails, the resulting value for q_1 is on the threshold – this is the value s_1^t that solves:

$$(q_0' - K_s - c_s s_1^t)(1+r) = K_s(1 + \frac{2c_s}{\gamma\theta b_s(1-\gamma/2) - \gamma c_s}).$$

It is easy to see that s_1^t is increasing in q_0'.

There are thus two cases to analyze: (i) there is a second period small attack only if the first period small attack succeeds; (ii) there is a second period small attack regardless of the outcome of the first period small attack. The viable s_1 values for case (i) are bounded below by s_1^t and the viable s_1 values for case (ii) are bounded above by s_1^t. We solve for the optimal s_1 for each, take the maximum (maximizing $E(q_2)$) and then compare to not launching an attack to determine the optimal decision. Case (ii) turns out to be essentially identical to the second period small attack decision problem. Substituting in for the optimal second period attack s_2, rolling back, and computing the first-order condition, the formula for the optimal level of first period small attack is

$$s_1^* = \frac{\gamma}{2c_s}(q_0' - K_s),$$

identical to the formula for s_2^* with q_0' substituted for q_1.

Case (i) requires a separate set of calculations. Again substituting the optimal value for s_2 and its expected value when the first period attack is a success, and setting $s_2 = 0$ when the first period attack fails, the equation to determine s_1 is:

$$\max_{s_1} -cs_1 + \theta(1 - \frac{s_1 c_s}{q_0' - K_s})[bs_1 + (\theta b\frac{\gamma}{2c_s}(1 - \gamma/2) - \gamma/2)(b - c_s)s_1].$$

From this we derive the first-order condition for the optimal value of s_1, s_1^*:

$$s_1^* = \frac{1}{2c_s}[\frac{\gamma + (\theta\frac{\gamma}{2c_s}(1 - \gamma/2) - \gamma/2b)(b - c)}{1 + (\theta\frac{\gamma}{2c_s}(1 - \gamma/2) - \gamma/2b)(b - c_s)}](q_0' - K_s)$$

This formula is similar to the formulas for s_2 and s_1 for case (ii) in that s_1^* is increasing in $q_0' - K_s$ and includes the term $\frac{1}{2c_s}$; however, it includes an additional somewhat complex fraction. By inspection the s_1^* for this second equation is strictly larger than the expression for s_1^* for case (ii) as long as γ is less than 1. Since the value associated with launching a small attack in period

2 is increasing in q_1, it follows that when s_1 is smaller, a failure is less costly and therefore a second period attack is more attractive.

Defining

$$A = \theta b \frac{\gamma}{2c_s}(1 - \gamma/2) - \gamma/2,$$

for case (ii) the value function depends on q_0' as:

$$(q_0' - K_s)\frac{\gamma}{2c_s}[\theta b(1 - \gamma/2)(1 + A) - c_s(1 + A)](1 + r)$$

and for case (i) as:

$$(q_0' - K_s)\frac{1}{2c_s}\frac{\gamma + A/b}{1 + A/b}[\theta b(1 - 1/2(\frac{\gamma + A/b}{1 + A/b})(1 + A)) - c_s(1 + \theta(1 - 1/2(\frac{\gamma + A/b}{1 + A/b})A))](1 + r).$$

Since both functions are linear in q_0', it follows that they can cross at most once; thus there can be at most one switch from one case to the other.

The Four Regimes

Consider the case where the terror organization can only mount small attacks, in which case $q_0' = q_0$. Analyzing the first and second period small attack decisions together, there are 4 solution regimes: (i) launch no attacks; (ii) do not launch an attack in period one but launch an attack in period two; (iii) launch an attack in period one, then launch an attack in period two only if the period one attack succeeds; and (iv) launch an attack in both periods. These four regimes do not always occur – for some parameter values, only a subset occurs. However, extensive simulations indicate that the regimes always occur in the same order as a function of q_0. Regime (i) is first, followed by (ii), (iii), (iv). When a subset of regimes occurs, they always occur following this ordering.

Above we showed that the value functions associated with strategies (iii) and (iv) are linear in q_0' ($= q_0$ when only small attacks are allowed). It is easy to show that the value functions associated with strategies (i) and (ii) are also linear in q_0. Thus for any given set of parameters the relative slopes of the value functions associated with the different strategies can be computed and compared, and it is easy to see that the strategies are ordered from least to greatest slope $-$ with the proviso that some may not be optimal for any q_0 values. The overall value function is the upper envelope of these 4 functions, and it follows from the fact that they are ordered from least to greatest slope that the overall value function is piecewise linear and convex (or linear in special cases).

[Figure 2 near here]

As a function of q_0, Figure 2 depicts the solution for one representative set of parameters: $c_s = 1$, $b_s = 5$, $K_s = 1$, $\theta = 0.6$, and $r = 0.5$ (all costs and benefits are measured in person years). We chose these parameter values, and the later values for a large-scale attack, based on our reading of the evidence associated with the string of well known terror attacks of the past decade, as noted in the introduction. Because costs and benefits are measured in person-years, we are assuming in effect that a small attack requires one year of fixed cost planning, which might for example be accomplished by 3 individuals working for 4 months. It also requires one full year of human resources for implementation for each single unit of scale (linked for example to intended casualties), perhaps several bombers and several tactical specialists working for a few months. Benefits are five times this operational cost – five units added to q for each one unit of scale. For this simulation, the probability of not being detected and stopped by counter-terror authorities is sixty percent, and the natural rate of growth of the organization is 50 percent per period. Also graphed in the figure is the value function, $V(q_0)$ or $E(q_2)$.

In the figure there are four regimes corresponding to the four strategies of the terror organization. The first regime extends from zero to $q_0 = 2.67$. In this regime the terror organization engages in no attacks and its growth is due entirely to r. The second regime extends from $q_0 = 2.67$ to $q_0 = 3.54$. In this regime the terror organization engages in only a second period attack (with scale shown by the lightly-dashed line in this interval). The third regime begins at $q_0 = 3.54$ and extends to $q_0 = 5.08$. In this regime the terror organization engages in a first period attack (the unmarked solid line) and engages in a second period attack only if the first period attack succeeds (the heavily-dashed line). The fourth regime extends from $q_0 = 5.08$ up. In this regime the terror organization engages in a first period and a second period attack regardless of whether it fails or succeeds with its first period attack. The value function (solid line with $\times$'s) is piecewise linear and convex, as noted above. Thus the terrorist organization value function, which is its expected strength at the end of period

2, is convex in its initial strength q_0.

There is an option value effect in our model: if a period 1 attack succeeds, a period 2 attack can be launched of relatively large scale, which generates the chance for a large q_2. Even so, there are values of q_0 for which the organization chooses not to launch a small attack in period one when its available resources q_0 are above K_s. This is due to the technology for attack success: the probability of success depends on

$$1 - \frac{c_s s}{q_0 - K_s},$$

so that if $q_0 - K_s$ is small, then even for a small attack the chance of success is very low. Further, there is a range of q_0 for which the organization chooses not to launch a period one small attack, but does launch an attack in period two.

The Large-Scale Attack Decision

The terror organization's decision about whether to undertake a large-scale attack is made at the beginning of the first period. Recall that the results of this attack – success or failure – are not realized until the end of period two. Hence there are no further decisions by the organization that are conditioned on the outcome of such an attack.[7]

We fold the analysis of the organization's decision about whether or not to engage in a large-scale attack into our analysis of the small scale attacks (see Figure 1 for the small attacks decision tree). If the organization chooses to undertake a large-scale attack of magnitude s_b, this reduces available organizational resources for a small attack. In particular, q_0 is reduced to q_0' where

$$q_0' = q_0 - K_b - c_b s_b.$$

Given q_0', we can use the previous analysis to determine the organization's optimal strategy with respect to small scale attacks. This decomposition is possible because the value function, $E(q_2)$, is linear in the outcomes of the different attacks: the overall value function may be written as the sum of two parts,

$$V(q_0'(s_b)) + \theta[1 - \frac{s_b c_b}{q_0 - K_b}] b_b s_b,$$

adding the expected benefit from a successful large-scale attack to the previous expression.[8] To determine whether the organization will find it in its interest to engage in a large-scale attack, and the magnitude it will plan for such an attack if it does choose to undertake an attack, we maximize this overall value function by computing the optimal s_b if it does undertake an attack and comparing the expected consequences to the value if it chooses not to undertake a large-scale attack.

[Figure 3 near here]

Figure 3 depicts the optimal strategy for the terror organization including its decision about launching a large-scale attack. For this simulation we use the same parameters as before except that we set K_s to 1/2. Parameters for the large attack are: $K_b = 3$, $c_b = 1$, and $b_b = 40$. Finally we set $\theta = 0.6$ and $r = .1$. Compared with a small attack, the large attack requires six times as much fixed cost investment, the equivalent of for example 3 individuals working full time for a full year (or 6 individuals part-time). It requires the same operational cost in terms of scale. The benefit for a given scale is 8 times larger.

The figure shows that for small initial terror organization size, a large attack is not engaged in and the previous analysis of small attacks applies. However, for these parameters only 3 of the regimes occur. Above $q_0 = 3.92$, the terror organization undertakes a large attack. The three regimes associated with the small attacks then repeat. In particular, for $q_0 = 3.92$ to $q_0 = 6.09$, the terror organization undertakes only a large attack. From $q_0 = 6.09$ to $q_0 = 9.08$, in addition to a large attack the organization engages in a period one small attack and undertakes a small attack in period two if the period one attack is successful. Finally, for q_0 above 9.08, the terror organization undertakes a full set of attacks – a large attack and small attacks in both periods. Note that the q_0 intervals for the three regimes in which a large attack is undertaken are wider than the corresponding intervals for the similar regimes (in terms of small attacks) when a large attack is not undertaken.

A large-scale attack requires K_b of human resources as a fixed cost, hence for q_0 below this level the organization cannot undertake a large-scale attack. Above K_b a large attack becomes possible. An interesting question is, what is the minimum value of q_0 for which the organization first finds it optimal to engage in a large-scale attack? Our simulations reveal that this value of q_0 is in fact well above K_b. For example, in the simulation results shown in Figure 3, the fixed cost to engage in a large scale attack is 3, but the organization chooses not to engage in such an attack until its initial size reaches 4. In other simulations, with different parameter values, we find an even greater cushion between the fixed cost level and the minimum initial size at which the terror organization undertakes a large-scale attack. Thus it is not optimal for the organization, in our model, to undertake a large-scale attack as soon as it has the minimum required resources to do so. Rather, it must have grown to a point well beyond this threshold before it finds it optimal to engage in such an attack. This result derives from two effects. One is the technology of attacks: an attack for which its scale s is very close to the maximum possible scale $s_{\max}$ has a low probability of success. The other is the tradeoff between large and small attacks: When the organization chooses to undertake a large-scale attack it foregoes having the resources it devotes to such an attack available for small scale attacks.

Though our focus is on the case in which r is positive, we have also explored the case of a negative r, corresponding to a declining terror organization. Interestingly, we find that in such cases the minimum value of q_0 for which the organization is willing to engage in a large attack is lower: For $r = -.1$ the threshold falls to 3.7 (from 3.9 for $r = .1$ shown in Figure 3) and for

$r = -.5$ the threshold falls further to 3.34. This is intuitive: A declining organization will not reap as great benefits from a small period one attack that succeeds, and is more likely to "go for it all" attempting a large attack that, if it succeeds, may rescue the organization from decline.

Model Solution: Additional Simulation Results

Figures 4 and 5 show additional simulation results. Figure 4 depicts expected final terror organization size for various values of r, the natural rate of growth of the organization. The graph is interesting and not initially intuitive. The graph shows that the terror organization's expected size is greater for larger r for all initial size values, as expected (note that the lines for the different r values do not cross). What is surprising is the way the expected final size lines converge for intermediate initial size values, roughly 4 to 5.5, and then diverge. The reason is that when r is very low, the terror organization finds it in its own interest to launch a large attack for even a relatively small initial size. It has no other way to grow and thus takes the risk of failure for the chance of success. For large values of r, the organization is not willing to engage in a large-scale attack for these modest initial size levels, but only when initial size is greater. The reason is that it will grow naturally even if does not launch an attack – the opportunity cost of expending resources on an attack with a relatively high chance of failure is so high as to discourage undertaking such an attack. Thus there is an interesting interaction between r and initial terror organization size. This insight also picks up on the regimes depicted in Figure 3. The region in Figure 4 in which the lines for different r values converge is exactly the regime in Figure 3 for which only a large attack is undertaken. There are no recruitment gains for large attacks at the end of period 1 (since the large attack does not happen until the end of period 2), hence the trade-off between undertaking a large attack vs. a small period one attack is shifted more in favor of a small attack as r rises. Further, when a large attack is undertaken, recruitment can only be based on the q "left over" after period 1, which is proportionately less than when a small attack is undertaken in period one with no big attack. These effects lead the curves to converge, as shown in Figure 4.

[Figure 4 near here]

Figure 5 explores another interesting interaction, between r and θ (focusing just on the case in which r is positive). The figure shows a positive, superlinear relationship as both θ and r rise: In the "far" corner of the contour box, when both θ and r rise towards one, the expected final terror organization size rises rapidly. This result highlights how the government's counter-terrorism policy should depend upon other parameters, in this case r. When r is low the government has more leniency in terms of the need for a strong detection ability, whereas for large r if the government fails to have a good counter-terrorism detection force the terror organization may (in expectation) grow rapidly.

[Figure 5 near here]

Conclusion

We have presented a stylized model of a strategic terror organization operating over two periods, with access to two attack technologies. Our findings highlight how a terror organization's decision to engage in an attack depends upon its current size, other attacks it has in the works, and parameters in its environment. Our results show that a terror organization's attack strategy falls into well-defined regimes, and that these regimes are ordered in a regular way, which should be helpful in guiding counter-terrorism policy formulation. Our results also show that a terror organization leverages its size at an increasing rate, so that counter-terrorism will be most effective through keeping a terror organization small.

In future work we plan to extend our model to a multi-period setting so we can study the longer run dynamics and life-cycle of a terror organization. We also intend to model the government's counter-terrorism policy, thus developing a game-theoretic analysis. It will be useful to see how the government's enforcement responds to the terror organization's size, which is stochastic depending on attack outcomes. We also envision more rigorous empirical research addressing the life-cycle of selected terrorist organizations, including the relationship between the results of terror attacks, the organization and its evolution.

We are only beginning to develop more rigorous, sophisticated models of terror organizations that can themselves be quite sophisticated. Since such terror organizations undoubtedly pose a greater threat to society, this approach seems important.

Footnotes

[1] See Arce and Sandler (2005) and Sandler and Siqueira (2009) for surveys of game theory models of terrorism and counterterrorism.

[2] See Financial Action Task Force (2008) and Silber and Bhatt (2007) for descriptions of these attacks; see International Crisis Group (2006) for information specifically about the Indonesian bombings; see Wright (2004) for a review of Al-Qaeda's historical development.

[3] An alternative way to derive this formula is to define the maximum feasible scale for an attack, given q. This is given by $K + cs = q$. Solving for the largest value of s that satisfies this equality and labeling this $s_{\max}$ yields

$$s_{\max} = \frac{q - K}{c}.$$

Our formulation can then be seen as $1 - s/s_{\max}$.

[4] More generally, the organization may recover a fraction of the resources spent on the attack after the attack is completed, which would then be added back into q next period.

[5] The value of r is also important for government counter-terrorism policy.

[6] One way to incorporate the idea that an attack may fail so badly as to further reduce q is to add a third outcome to each q node in the figure, corresponding to the case of "Fail, negative impact on q" (our current fail scenario would then be "Fail with 0 further impact on q"). In principle the basic structure of our analysis will still go through in this case.

[7] In terms of Figure 1, the decision about the large attack is an initial decision

node at the start of the tree. If a large attack is undertaken the outcome is represented by a random event node at the back of the tree.

[8] As noted earlier in the main text, we could extend the model to allow for the possibility of a "Fail with negative impact on q" node into our model. We note also the model should in principle be able to be extended for an arbitrary number of periods and solved via dynamic programming arguments such as those found in Jacobson and Kaplan (2007).

References

Arce, Daniel G. and Todd Sandler. 2005. Counterterrorism: A game-theoretic analysis. *Journal of Conflict Resolution* 49 (2): 183-200.

Faria, João R. and Daniel G. Arce. 2005. Terror support and recruitment. *Defence and Peace Economics* 16 (4): 263-273.

Enders, Walter and Paan Jindapon. 2010. Network externalities and terrorist network structure. *Journal of Conflict Resolution* 54 (2): this issue.

Enders, Walter and Todd Sandler. 2006. *The Political Economy of Terrorism.* Cambridge, UK: Cambridge University Press, 2006.

Fattah, Hassan M. 2005. Bombing in Jordan: reaction; angry Jordanians mourn dead, and a bruised image. *New York Times*, November 11.

Financial Action Task Force. 2008. *Terrorist Financing, 2008.* Paris: Organisation for Economic Cooperation and Development (available at http://www.fatf-gafi.org/dataoecd/28/43/40285899.pdf).

Homeland Security Council. 2004. *Planning Scenarios Created for Use in National, Federal, State, and Local Homeland Security Preparedness Activities.* Washington, DC: White House Homeland Security Council (available at http://www.globalsecurity.org/security/ops/ter-scen.htm).

International Crisis Group. 2006. *Terrorism in Indonesia: Noordin's Networks*. Asia Report No. 114.

Jacobson, Daniel and Edward H. Kaplan. 2007. Suicide bombings and targeted killings in (counter) terror games. *Journal of Conflict Resolution* 51 (5): 772-792.

Johnson, Neil F., Mike Spagat, Jorge A. Restrepo, Oscar Becerra, Juan C. Bohorquez, Nicolas Suarez, Elvira M. Restrepo and Roberto Zarama. 2006. Universal Patterns Underlying Ongoing Wars and Terrorism. *Physics and Society* (available at arXiv:physics/0605035v1).

Kaplan, Edward H., David L. Craft and Lawrence M. Wein. 2002. Emergency response to a smallpox attack: the case for mass vaccination. *Proceedings of the National Academy of Sciences of the United States of America* 99 (16): 10935-10940.

Kaplan, Edward H., Alex Mintz, Shaul Mishal and Claudio Samban. 2005. What happened to suicide bombings in Israel? Insights from a terror stock model. *Studies in Conflict and Terrorism* 28 (3): 225-235.

Keohane, Nathaniel O. and Richard J. Zeckhauser. 2003. The Ecology of Terror Defense. *Journal of Risk and Uncertainty* 26 (2/3): 201-229.

National Research Council. 2007. *Countering the Threat of Improvised Explosive Devices: Basic Research Opportunities.* Committee on Defeating Improvised Explosive Devices: Basic Research to Interrupt the IED Delivery Chain, National Research Council. Washington, DC: National Academies Press.

Rosendorff, B. Peter and Todd Sandler. 2004. Too Much of a Good Thing? The Proactive Response Dilemma. *Journal of Conflict Resolution* 48 (5): 657-71.

Sageman, Marc. 2004. *Understanding Terror Networks.* Philadelphia: University of Pennsylvania Press.

Sageman, Marc. 2008. *Leaderless Jihad: Terror Networks in the Twenty-First Century.* Philadelphia: University of Pennsylvania Press.

Sandler, Todd and Kevin Siqueira. 2009. Games and terrorism: recent developments. *Simulation and Gaming*, 40 (2): 164-192.

Silber, Mitchell D. and Arvin Bhatt. 2007. *Radicalization in the West: The Homegrown Threat.* New York: New York City Police Department.

Wein, Lawrence M., David L. Craft and Edward H. Kaplan. 2003. Emergency response to an anthrax attack. *Proceedings of the National Academy of Sciences of the United States of America* 100 (7): 4346-4351.

Wright, Lawrence. 2004. *The Looming Tower.* New York: Knopf.

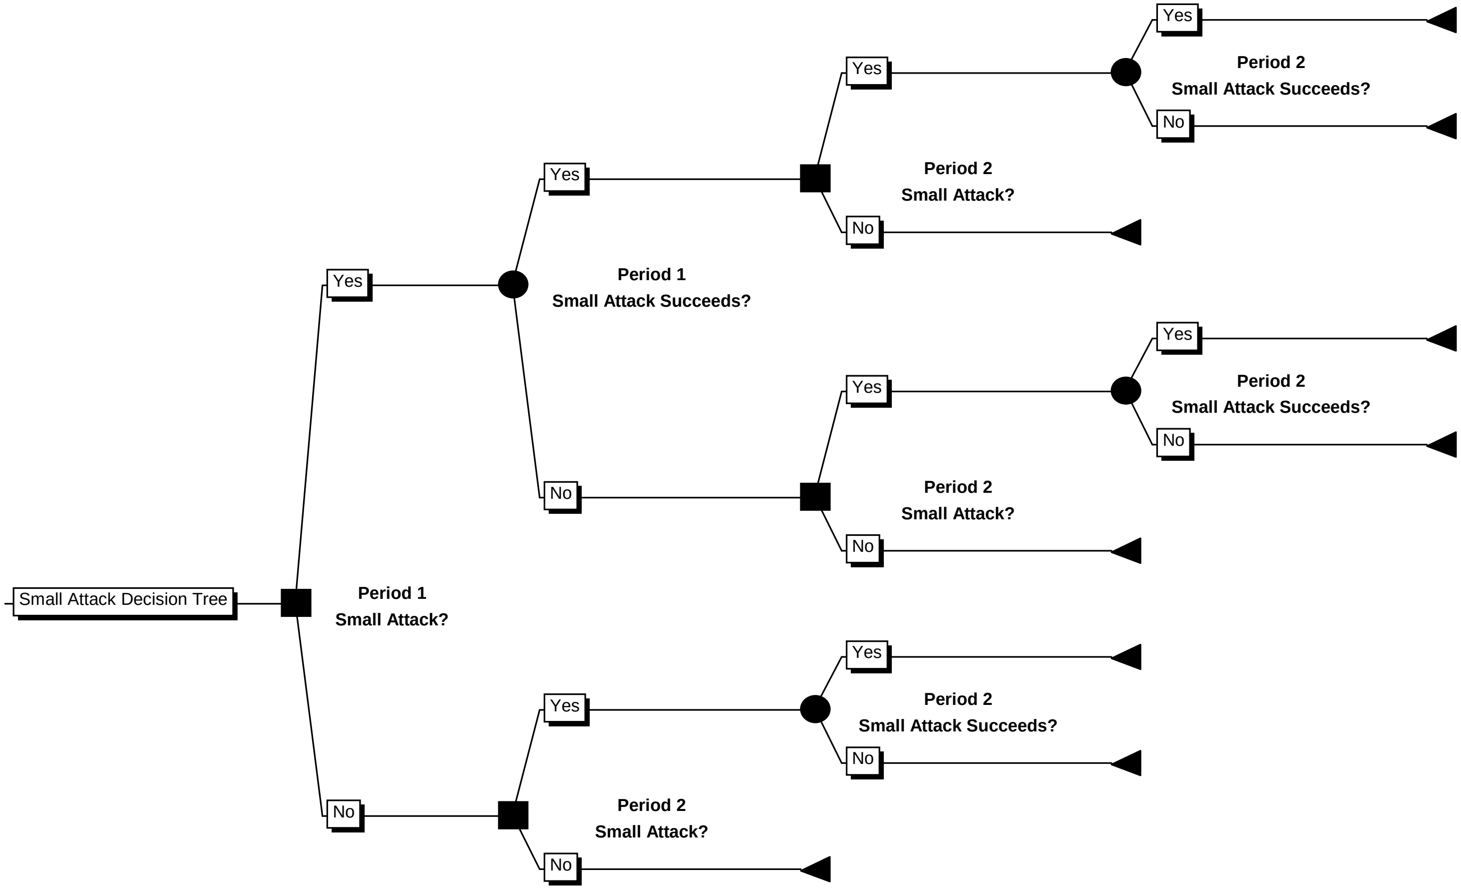

Figure 1: Decision Tree For Small Attacks Model

Figure 2: Optimal Decisions and Value Function for Small Attacks Model

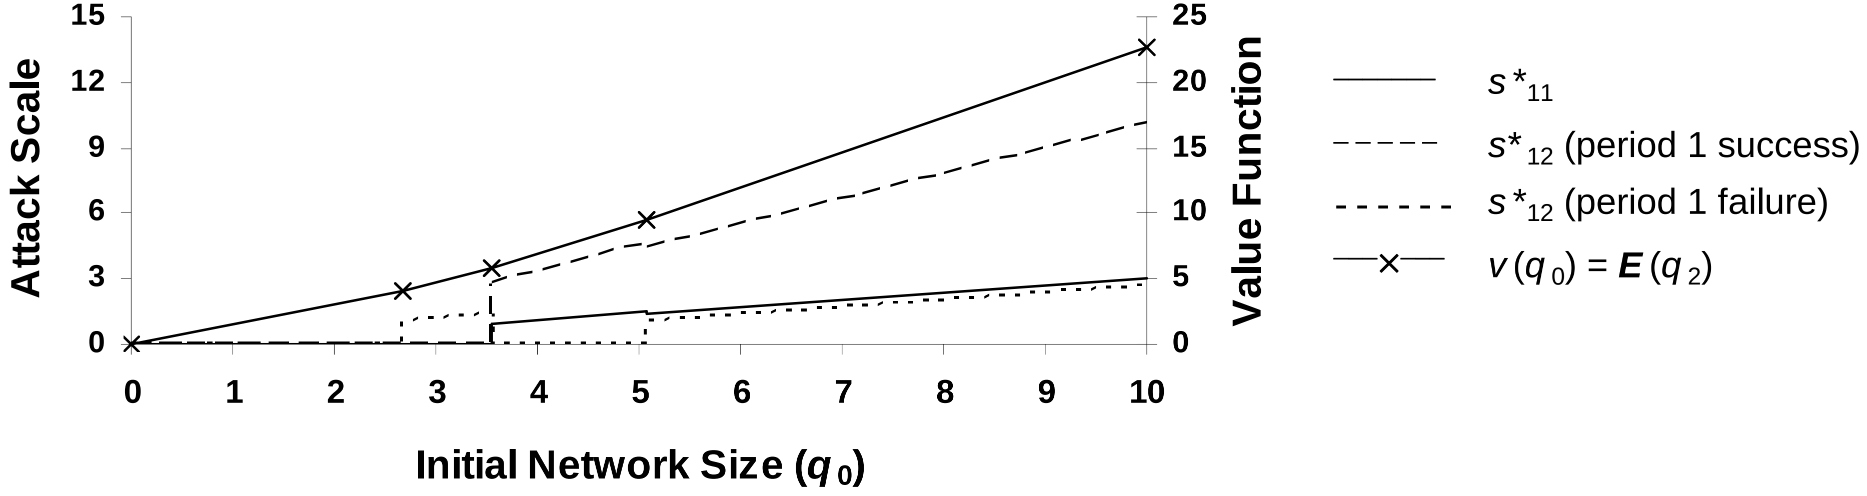

Figure 3: Optimal Decisions and Value Function for Complete Model

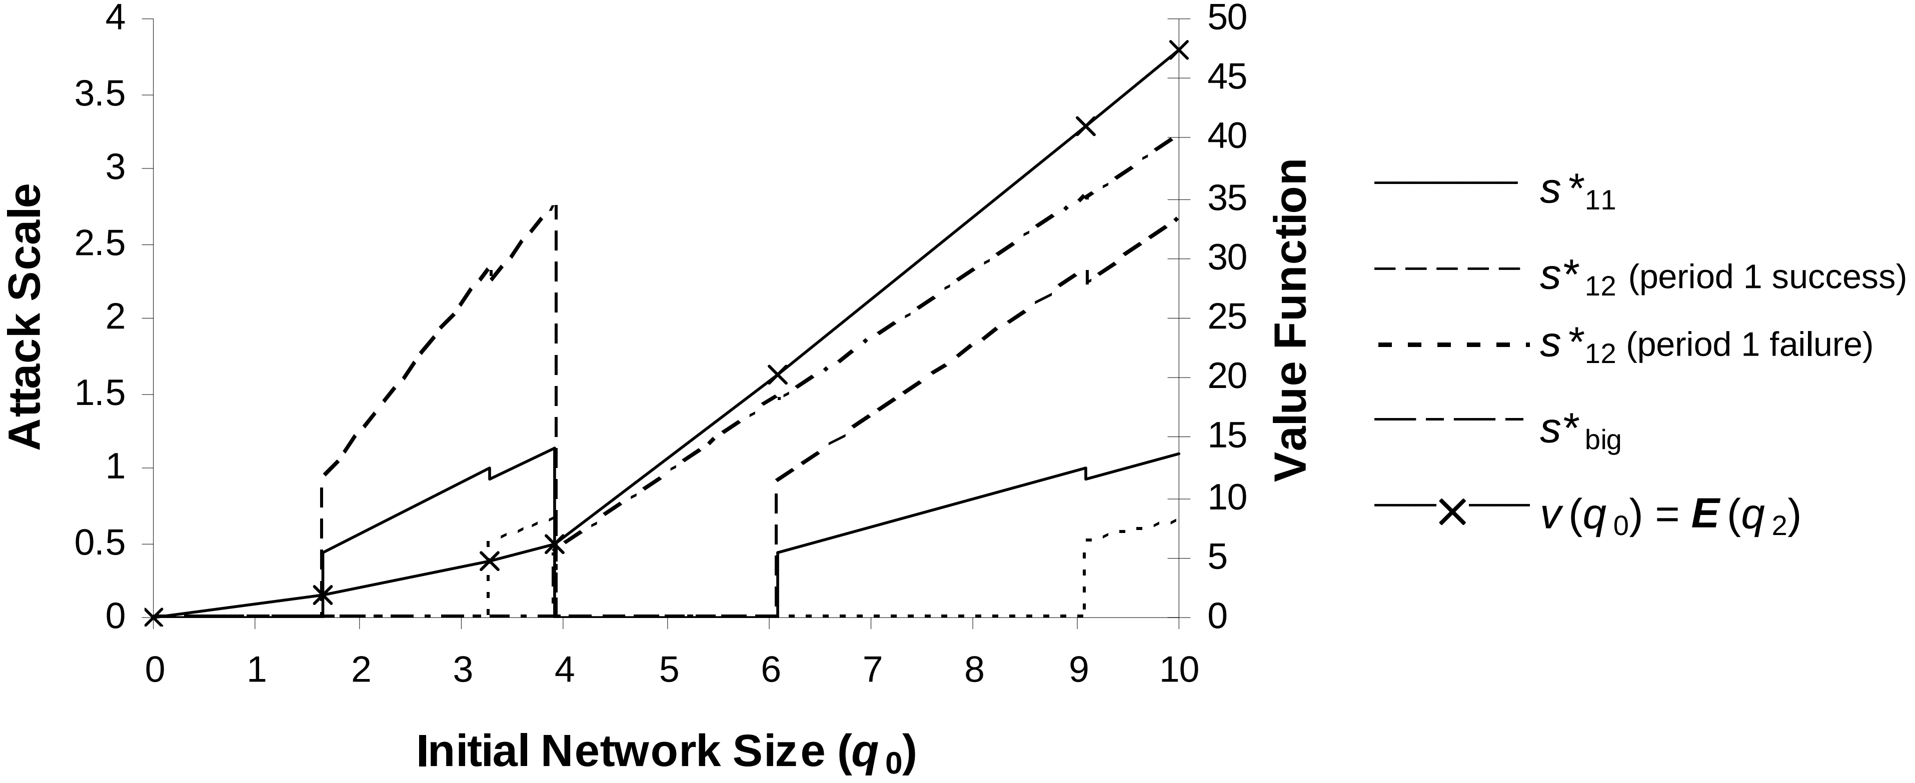

Figure 4: Final Terror Network Size E[q_2]
Controlling For Terrorist Recruitment Rate *r*

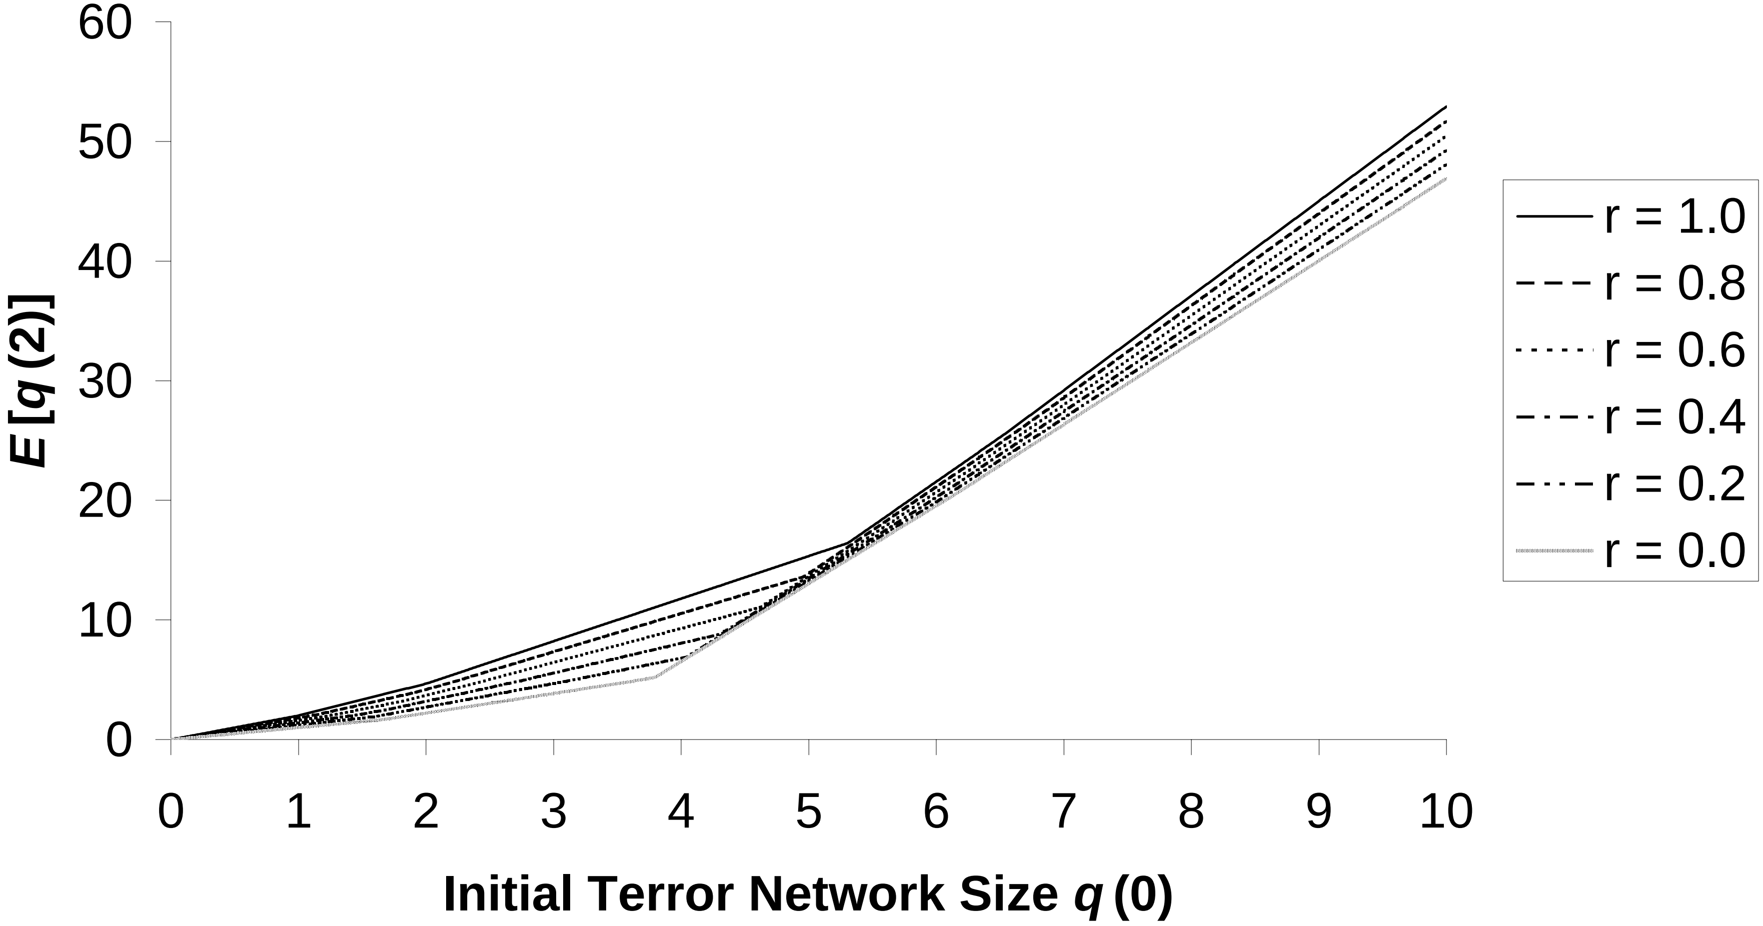

Figure 5: Final Terror Network Size $E[q_2]$ Controlling For Failed Detection Rate θ And Terrorist Recruitment Rate *r*

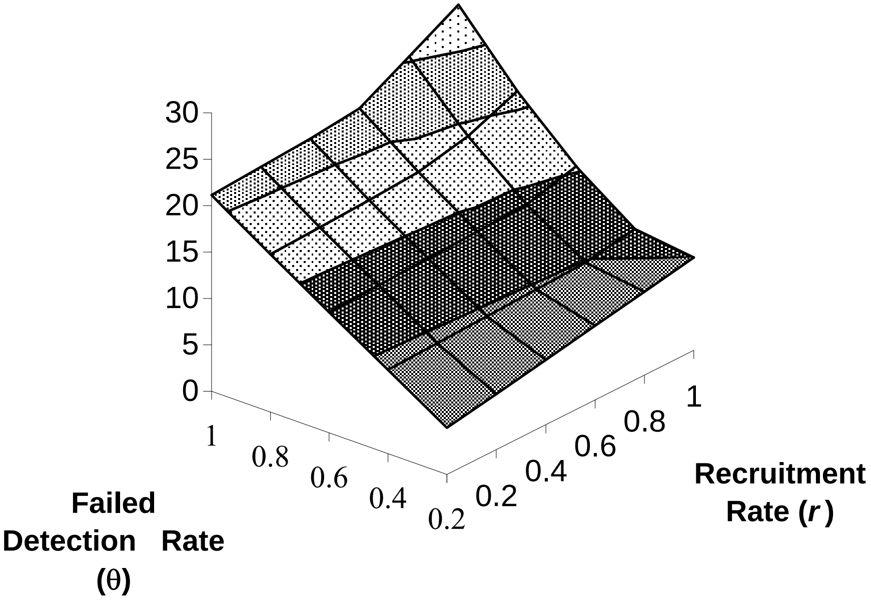
Failed Detection Rate (θ)
Recruitment Rate (r)