

Underfunding in Terrorist Organizations

JACOB N. SHAPIRO

Stanford University

DAVID A. SIEGEL

Florida State University

A review of international terrorist activity reveals a pattern of financially strapped operatives working for organizations that seem to have plenty of money. To explain this observation, and to examine when restricting terrorists' funds will reduce their lethality, we model a hierarchical terror organization in which leaders delegate financial and logistical tasks to middlemen, but cannot perfectly monitor them for security reasons. These middlemen do not always share their leaders' interests: the temptation exists to skim funds from financial transactions. When middlemen are sufficiently greedy and organizations suffer from sufficiently strong budget constraints, leaders will not fund attacks because the costs of skimming are too great. Using general functional forms, we find important nonlinearities in terrorists' responses to government counterterrorism. Restricting terrorists' funds may be ineffective until a critical threshold is reached, at which point cooperation within terrorist organizations begins to break down and further government actions have a disproportionately large impact.

Shortly after noon on February 26, 1993, approximately 1,500 pounds of homemade explosives were detonated on the second level of the World Trade Center parking structure. Six people were killed and more than a thousand were injured. It could have been much worse; the F-350 Econoline van rented by the terrorists had a cargo capacity of over 4,000 pounds. Ramzi Yousef, the leader of the attack, later testified he did not have enough money to purchase sufficient materials to build a larger bomb. He also claimed the attack was rushed because his cell ran out of money (Levitt 2002). Yousef's cell in New York was underfunded at a time when Al Qaeda businesses in Sudan were generating enough income to support thousands of followers, run military training camps, and provide large kickbacks to the Sudanese government (Gunaratna 2002:38–42).

Similar underfunding problems have plagued other international terrorist organizations. Idris, the chief logistician for the 2002 Bali bombing, was arrested 3 weeks after the bombing. The police tracked him down through his involvement in a bank robbery intended to raise funds for subsequent operations. A detailed study of the money sent to Idris to fund the Bali bombing reveals that he received much

Authors' note: The authors thank Jim Breckenridge, Ethan Bueno de Mesquita, James Fearon, Jeanne Giraldo, Harold Trinkunas, the participants in the CISAC Social Science Seminar, and several anonymous reviewers for tremendously helpful comments and criticisms. All errors are our own.

more than he disbursed for that attack, making his subsequent involvement in risky criminal activities puzzling (Abuza 2003b:54–55). Indeed, there is a recurring pattern of financially strapped operatives working for terrorist organizations that seem to have plenty of money. We seek to explain this pattern.

The puzzle for traditional perspectives on terrorist financial and logistical systems is that groups, which are purportedly organized to carry out attacks often provide inadequate funds to their operatives. Standard accounts stress the efficiency with which terrorist financial networks distribute funds while operating through a variety of covert channels. They describe terrorist organizations as able to shift agilely between multiple avenues to raise and use funds with no mention of transaction costs or organizational infrastructure requirements. We are told that “al-Qaeda is notably and deliberately decentralized, compartmentalized, flexible, and diverse in its methods and targets . . . Al-Qaeda’s financial network is characterized by layers and redundancies. It raises money from a variety of sources and moves money in a variety of manners” (Weschler and Wolosky 2002:6). Reports from the multinational Financial Action Task Force on Money Laundering (2004:6), the Asia/Pacific Group on Money Laundering (2004:32), and others provide a similar narrative (Brisard 2002:6; Singapore Ministry of Home Affairs 2003:6).

From this perspective, it is hard to imagine why a group would underfund operational elements. As an alternative, suppose that the members of a terrorist support network, middlemen, were not uniformly driven by mission accomplishment, but that some were driven by monetary rewards. Because the network is covert, informational asymmetries abound, exacerbating the principal–agent dilemmas found in any hierarchical organization, creating numerous opportunities for middlemen to appropriate resources for personal use. In this scenario a greedy middleman need only pass on as much money to operatives as is required to achieve an acceptable number of successful attacks.¹ This rough sketch, which our model formalizes, can explain underfunding in hierarchical terrorist groups.

A game-theoretic approach is particularly attractive for dealing with this puzzle because it places the organizational dilemmas faced by terrorist groups into the starkest possible contrast (Victoroff 2005). Developing formal models of terrorism can also help explain otherwise puzzling patterns. The best example of such work is Bueno de Mesquita (2005b), which reconciles the seemingly contradictory findings that while terrorist operatives are neither poor nor lacking in education, poverty, and lack of economic opportunity are positively correlated with terrorism. In the case of our model, a formal presentation is required because the interactions are simply too complicated to understand decisively in an informal analysis.

Previous formalized rational-choice analysis of terrorism has focused on two areas: (1) the interaction between governments and terrorist groups; and (2) groups’ internal dynamics. In the first area, much of the early work focused on the signaling dynamics of terrorism. Lapan and Sandler (1993) present a model in which the optimal government strategy depends on the resources of the terrorist group. The terrorist group can use the scale of attacks to send a signal about these resources and has incentives to misrepresent its resources to gain concessions from the government. Overgaard (1994) presents a more subtle signaling model in which terrorist resources are renewable between periods and the terrorists have positive alternative uses for resources, alternatives such as providing social services. Overgaard finds that if concessions are ruled out for exogenous political reasons, then only a pooling equilibrium exists and the government learns nothing from the scale of attacks. More recently, Sandler (2003) explains the under-provision of international counter-terrorism as a collective action problem. Frey and Luechinger (2004)

¹ “Acceptable” here means the probability that the middleman is not fired or killed is sufficiently small relative to the rewards he obtains.

use a simple supply and demand model to examine the relative merits of deterrence and decentralization of critical infrastructure as counter-terror policies.

In one of the first formal rational-choice analyses of terrorist groups' internal dynamics, Chai (1993) examines why people participate in covert anti-government organizations even though their chances of getting particularistic benefits out of doing so are quite slight given the low likelihood of overthrowing the government. His analysis helps explain the common requirement among European terrorist groups that new members commit a violent act as part of the recruiting process. Siqueira (2005) models the interaction between militant and political factions within a dissident movement. His work highlights how the external manifestations of competition between factions depend on the nature of the externalities each factions' actions create for the other faction. Berman and Laitin (2005) examine why radical religious groups are particularly well suited to using suicide attacks. Recent analysis of how terrorist conflicts end formalizes Crenshaw's (1991) arguments about how the internal dynamics of terrorist organizations can cause difficulties for the peace process. Bueno de Mesquita (2005c) models the problems created by differing levels of ideology within a terrorist organization trying to make peace. His analysis outlines the strategic considerations underlying the tendency of radical factions to break off from groups engaged in peace negotiations.

While the level of sophistication in organizational analysis of terrorism has increased, no one has yet explored the challenges that heterogeneity poses for terrorist support systems and how those challenges affect groups' abilities to kill people.² Our work begins to address this by focusing on the relationship between terrorist leaders and the middlemen they must use for security reasons.

Beyond contributing to a growing body of literature on the organizational dynamics of terrorist groups, our analysis has implications for important policy debates. It speaks to the treatment of possible informers, to the interpretation of the apparent success or failure of counter-terror policies, and to more fundamental questions over the role of counter-terrorist financing in the larger counter-terrorism efforts. A great deal of energy has been spent over the last 5 years to clamp down on terrorist financing. UN resolutions have been passed, international standards set, technical assistance programs started, and financial laws changed in over 100 countries (Francis 2004; Weschler and Wolosky 2004:9–11). Legal changes undertaken as part of this campaign have created serious threats to treasured civil liberties in Western democracies and appear to have been highly counter-productive from an intelligence standpoint (Donohue 2006). Asset seizures, and the threat thereof, have crippled Islamic charities that provide invaluable social services in areas of failed governance. Efforts to limit terrorists' use of alternative remittance systems have had disastrous social welfare implications in some of the world's poorest areas (Medani 2002).

Yet all this activity may have little positive effect. Terrorist attacks are cheap compared with the budgets of terrorist organizations, at least when fixed costs are not taken into account.³ Of the 23 international attacks attributed to Al Qaeda since 1995, only the World Trade Center attacks appear to have cost the group even 1%

² A more general problem which is just now being addressed in the academic literature is that most studies of terrorism fail to consider the great heterogeneity of terrorists. This problem is explored in great detail in Victoroff (2005).

³ Most estimates are that terrorist groups spend less than 10% of their income on actually conducting attacks. For example, estimates of the centrally managed portion of Al Qaeda's budget range from \$16 million to under \$50 million a year. See Second Report of the [UN] Monitoring Group (2002:12, 27). The 1993 World Trade Center attack cost less than \$20,000, according to its mastermind Ramzi Yousef, and the Bali bombing cost at most \$35,000 (Lee 2002:4) (see also Abuza 2003b). Given these numbers, we concur with the CIA estimate that Al Qaeda's spending on actual operations was quite small. (National Commission on the Terrorist Attacks on the United States 2004b:11).

of its estimated annual budget.⁴ Given how hard it is to stop terrorist funds, and how cheap attacks appear to be, one must wonder whether all this effort is for naught. Our analysis informs this debate by identifying how counter-terrorist financing efforts can limit attacks without necessarily making attacks unaffordable. Specifically, seizing funds, and thereby tightening terrorists' budget constraints, can increase friction within organizations to the point where leaders choose to spend on activities other than attacks. However, before this constraint is reached, seizing funds is unlikely to reduce attacks.

The remainder of the paper proceeds as follows. The next section provides empirical and theoretical motivation for our model of a hierarchical terrorist group. The third section formally presents the model. The fourth section develops the equilibrium strategies and derives comparative statics. Readers not interested in the mathematical development can omit this section without significant loss of comprehension. The penultimate section discusses the results, providing greater intuition about the model through a computational illustration of our predicted outcomes. Using general functional forms we find important nonlinearities not revealed by the comparative statics. The last section concludes.

Motivation

Because of the covert nature of their work, terrorist networks must operate with fewer checks and balances than most financial organizations require. Indeed, the cellular structure of terrorist networks so often cited in the literature necessarily implies that leaders will be poorly informed about the actions of their subordinates. If we assume that all the members of the network are uniformly committed to the cause and that they all agree on how best to advance the group's political goals, then there is no inconsistency here. However, suppose leaders, middlemen, and operational cadres have divergent preferences over spending. Then the informational asymmetries created by the secretive nature of terrorist networks leads to myriad opportunities for spending money differently than leaders would like.

Of course, the real-world division of labor is not always so stark. The level of specialization can vary over time and between groups. Al Qaeda and affiliates used to have quite defined organizational roles with a strong distinction between support and operational roles.⁵ However, since losing their refuge in Afghanistan, Al Qaeda and affiliates may have shifted to a less hierarchical system. In Madrid and Casablanca the same members appear to have engaged in logistical tasks and conducted operations.⁶ Moreover, the level of specialization can be a strategic choice. Resource-poor groups must be efficient to survive, while wealthy organizations may not be concerned with inefficiencies so long as they can meet their political goals.

These subtleties aside, using an agency theory framework can help explain otherwise puzzling behaviors. Consider the following informal example where the relationship between terrorist leaders and their financial system is described in terms of a principal-agent relationship. The principals, terrorist leaders, must delegate certain tasks—raising funds and distributing them to operational elements—to their agents, the financial network. Delegation entails a risk that if the agents' preferences differ from those of the principal, the agents will not carry out their tasks exactly as the principal would like; they may "shirk" by retaining some portion of the funds intended to support an attack.

⁴ On attacks, see the RAND/St. Andrews data provided by the Memorial Institute for the Prevention of Terrorism. Available at <http://www.tkb.org>. (Accessed February 13, 2006). On costs, a brief summary is provided by Prober (2005).

⁵ For example, the planning and bomb making for the African Embassy bombings were conducted by individuals who left the country shortly before the actual attacks.

⁶ M. Sageman, private communication, January 8, 2005.

Because monitoring the agents entails a security cost, the principals can only observe the outcome of the agents' actions: whether an attack succeeds or fails. This outcome is probabilistic, and thus provides only a noisy signal as to whether the agent passed on all the funds or not. Thus there is some space for agents to skim funds, thereby reducing the principal's utility. This strategy is feasible because the principal can neither perfectly monitor nor punish the agent with certainty. Clearly an agency theory approach can help explain the observed pattern of underfunding, but is there any evidence that these kinds of problems actually occur?

Many would cite Al Qaeda as the hard case for testing the idea that terrorist organizations suffer from agency problems and other organizational pathologies. Al Qaeda reportedly had a significant vetting process for membership and is generally discussed as a remarkably cohesive group (Testimony of FBI Agent 2001; Jenkins 2002:5). However, given that Al Qaeda is a relatively small group, operating small cells, over vast distances, in areas of the world with poor communications infrastructure, it is in some sense an easy case for agency problems. Such problems should exist here if the middlemen have preferences that diverge from those of the leadership. In this light, consider the following e-mail written by Ayman Al-Zawahiri, Al Qaeda's second-in-command, to a Yemeni cell on February 11, 1999:

... With all due respect, this is not an accounting. It's a summary accounting. For example, you didn't write any dates, and many of the items are vague. The analysis of the summary shows the following:

1 You received a total of \$22,301. Of course, you didn't mention the period over which this sum was received. Our activities only benefited from a negligible portion of the money. This means that you received and distributed the money as you pleased ...

2 Salaries amounted to \$10,085, 45 percent of the money. I had told you in my fax ... that we've been receiving only half salaries for 5 months. What is your reaction or response to this?

3 Loans amounted to \$2,190. Why did you give out loans? Didn't I give clear orders to Muhammad Saleh to ... refer any loan requests to me? We have already had long discussions on this topic ...

4 Why have guesthouse expenses amounted to \$1,573 when only Yunis is there, and he can be accommodated without the need for a guesthouse?—(Cullison 2004).

Al Qaeda, in fact, faces recurring difficulties due to the divergent motivations of its membership. Two examples from the East Africa Embassy bombings are instructive. Jamal Ahmed Al-Fadl, who testified in the African Embassy bombing trial, had stolen money from Al Qaeda, got caught, went on the run, and approached the U.S. government in an attempt to save himself and his family. L'Hussein Kherchtou, a member of the Nairobi team, testified for the government because he was so appalled at the un-Islamic embezzlement practiced by senior members of his team. Thus even in the hard case, we see divergent motivations creating agency problems.⁷

Agency problems over spending occur in smaller, more localized organizations as well.⁸ During the Christian-Muslim violence in Poso in late 2000, a relatively senior Jemaah Islamiyah (JI) member arranged to raise funds from oil company workers to be channeled through one local militia, KOMPAK-Solo, to JI and another local militia, Mujahidin KOMPAK. The workers were so concerned about the probity of these transfers that they appointed an auditor to oversee the funds (International Crisis Group 2004:9–10). This auditor's involvement exacerbated tensions between

⁷ For a sample of captured Al Qaeda documents highlighting agency problems, see Felter et al. (2006).

⁸ Agency problems over strategy and tactics have bedeviled terrorist organizations since the 1890s (Felter et al. 2006:11–21).

Mujahidin KOMPAK and JI, and the relationship deteriorated to the point that formal authorization by senior leadership was required for members of the groups to share weapons, something they had done freely during the early months in Poso. In other words, JI and Mujahidin KOMPAK suffered from the type of agency problems our model highlights.

Game

The interaction we describe is that between a terrorist boss, who controls the flow of money and desires some impact from this expenditure, and the middlemen he hires to perform the logistical support for those who will actually carry out the operations. Upon first glance, this is simply a standard principal-agent dynamic, in which the boss principal hires the middleman agent to do a job for him, though a particularly grisly one. However, this ignores the essential covert nature of the terrorist organization, in which monitoring is all but completely limited to whether or not an attack has succeeded. Moreover, contingent payment—the means by which the principal typically obtains his desired behavior—is either infeasible because middlemen do not have the capital to finance attacks, or is not practiced.

The general pattern of funds transfer seems to be that leaders provide funding in either an initial block grant, or on a need-to-have basis. The block grant method appears to have been used for the three Thai members of JI who were carrying \$50,000 to support operations when they were arrested in Phnom Penh in May 2003 (Abuza 2003a:194). Both the Bali bombings and the September 11 attacks followed the need-to-have model.⁹ Although there is some evidence that Hambali, the key middleman for the 2003 J. W. Marriott bombing in Jakarta, received a substantial bonus following the attack,¹⁰ we have been unable to find a single case in which a middleman or facilitator financed an attack and received payment from the group leader after the attack had succeeded, the contingent payment model.

The difficulty in monitoring agents, along with the ability of agents to go to the government, makes the application of a punishment difficult—at least in the international context—leaving little the principal can do to condition his agent. In a covert system, the agent holds an inherent threat over the organization. If he is too dissatisfied with his punishment, he can go to the authorities, as Jamal Ahmed Al-Fadl did. Because agents have exactly this option, punitive strategies should only exist where the organization can wield a credible threat of violence over the agent. It is possible for some localized groups, such as the IRA, to use murder to condition their agents. Yet even in such cases, punishment is problematic. One study of Palestinian groups operating in Lebanon in between 1976 and 1983 revealed that groups that punished disloyalty had to pay higher wages to their agents (Adams 1986:86). It seems quite likely that financial agents operating in foreign countries, such as the Yemeni recipient of Al-Zawahiri's e-mail, will be less susceptible to punishment strategies. That agent responded to being called out by quitting the network, illustrating the difficulties transnational groups face in using punishment strategies (Cullison 2004).

What the boss can do is refuse to use that middleman again for subsequent attacks, denying him future gains from participation.¹¹ As we will show below, if middlemen are patient enough, then this incentive proves sufficient to motivate even agents solely interested in pecuniary gain to act at least partially in the boss' interests.

⁹ For the Bali bombing figures, see Abuza (2003b). For the September 11 attacks, see National Commission on the Terrorist Attacks Upon the United States (2004a: section 5.4). See also Roth (2004: appendix A).

¹⁰ P. Williams, private communication, April 11, 2005.

¹¹ The Damascus-based leaders of Palestinian Islamic Jihad (PIJ) reportedly use exactly this method to condition their middlemen according to the sometimes-reliable *DEBKA-Net-Weekly* (Bounty for Murder 2004).

Actors

Our model starts from the position that a boss wants to maximize political impact. We can think of the Al Qaeda leaders whose expressed goal was to compel the United States to withhold its support from apostate Arab regimes such as the Saudi royal family (Al-Zawahiri 2001). In Zawahiri's writings, the purpose of terrorist attacks is twofold. First, to rally support around the movement by setting an example for others to follow. Second, to impose costs on the United States for policies that contradict the groups' goals. Both fall under the larger goal of achieving political impact.

A leader gets political impact by providing a basket of goods that includes attacks, as well as more prosaic goods such as social welfare services, ideological communications, training camps, and the like. Evidence from the occupied territories suggest that both Hamas and the PLO vary their outputs of each type of legitimacy-generating activity according to what garners the most support at a given moment in time (Rees 2001; Gunning 2004:242–243). Our model captures this substitution dynamic in a rough way. In our model, the boss simply gets utility from attacks and some disutility from spending money on attacks that could otherwise be spent on all the other goods that terrorist groups can produce. By varying the intensity of this disutility, we can account for relative preferences between attacks and other goods.

Now consider the preferences of the middlemen. We assume individuals join terrorist organizations when the utility of doing so is at least as good as that provided by their next best option. Utility is composed of two components. First, individuals get utility from the impact of their actions in furthering the group's goals, from doing what they believe is right. We simplify this to the probability that an attack succeeds. Second, individuals get utility out of monetary compensation: the money they take in that does not go to an attack. Each individual weighs these two components such that the sum of the weights is one. At the extremes are individuals who are purely motivated by impact, suicide bombers perhaps, and those motivated purely by money.

There is good empirical evidence that the preferences of middlemen are not always aligned with those of leaders and operational elements. Mid-level managers of organizations such as Harakat ul-Mujahedin (HUM), a Pakistani militant group focused on Kashmir, often live luxurious lives far beyond what their followers can afford (Stern 2003:213–216). Captured PLO documents show that those who plan attacks are paid eight times as much as is given to the families of those who die carrying out the attacks (Israeli Defense Forces 2002). People running criminal fund-raising operations in the United States for Hezbollah drive luxury cars and live in upper-middle class neighborhoods (Farah 2004:164).

This preference divergence should be expected as terrorist groups suffer from two adverse selection problems with respect to their financial and logistical operatives. The first is that those likely to survive long in terrorist networks tend to be less ideologically committed as they are less likely to volunteer for the most dangerous missions.¹² The second is that because participation as a financier or logistician is less risky than participating as a local leader or operator, middlemen in terrorist organizations will tend to be less committed. We examine each in turn.

One striking pattern that emerges from a close examination of terrorist organizations is that financial network members face dramatically lower risks than local leaders or tactical operatives. Beyond not being asked to participate in risky or inherently fatal ventures, they are less likely to be targeted by government forces. When targeted, they are less likely to be killed. And when arrested, they face more lenient treatment.

¹² We thank Ethan Bueno de Mesquita for pointing this out.

We assessed the risks of participating at different levels using Sageman's (2004) sample of 366 participants in the global Salafi jihad: Al Qaeda, affiliated organizations, and some operating outside of traditional organizations. Using open-source material we collected data on individuals' operational roles, when they left the jihad, and how they left.¹³ According to these data, between 1997 and 2003, financiers were rarely killed and their chances of being arrested rate was 10–20% lower than that of tactical operators, with 2002 being the only exception.

Even when government succeeds in capturing logisticians and other support network members, these individuals face dramatically lower consequences than operators. Only one of the 32 financiers and logisticians removed from the global Salafi jihad between January 2001 and December 2003 was killed. A particularly telling example is the Jemmaah Islamiyah (JI) cell, which was broken up in Singapore in late 2001. The cell provided fund-raising services to JI and was engaged in making logistical arrangements for an Al Qaeda attack in Singapore. Of the 30 + people arrested, the 13 engaged in direct logistical support each received 2 years in prison. Those engaged in fund-raising activities were released but not permitted to leave the country (Ressa 2003:158–160).

With these facts in mind, consider a hierarchical organization where individuals come up through the ranks, starting out in subordinate roles and moving into management roles as local leaders, financial facilitators, or logisticians.¹⁴ Throughout their careers, these individuals will have opportunities to volunteer for risky missions.¹⁵ Those most likely to do so will be those who place the highest weight on impact. Thus, the longer individuals remain in the organization, and the further they move up the management structure, the more likely they are to place a heavy weight on monetary rewards.¹⁶

Even without these adverse selection processes, there are reasons to expect divergence. The lenient treatment observed for support network members means that the threshold level of risk acceptance and commitment required for participation in support activities is much lower than for participation in tactical roles. We make the reasonable assumptions that there is a distribution of weights in the population of potential members and that the population of true believers is limited.¹⁷ Thus, given set wages for different activities, some individuals might participate in support activities while balking at other roles. Seeking to maximize operational capability, a rational organization would concentrate such individuals in support roles, freeing up the true believers for riskier operational duties. These personnel decisions would then lead to consistent differences in ideology between leaders and middlemen.¹⁸

¹³ Sageman (2004) uses data on 172 of these individuals. We conducted independent coding of operational roles for those participating between 1997 and 2003.

¹⁴ This progression need not happen within one organization. For example, much of the leadership in JI have been waging jihad together, at varying levels of intensity, since the mid-1980s (International Crisis Group 2003b: 7–9).

¹⁵ Evidence from trial transcripts and other sources suggest that volunteerism is a primary method of selection with leaders choosing from among volunteers. The selection process for the 9/11 attacks is discussed in National Commission on the Terrorist Attacks Upon the United States (2004a).

¹⁶ Weinstein (2005) discusses a slightly different adverse selection problem in his work on rebel groups in Sierra Leone. He posits that a wage that brings high-quality recruits will also bring in low-quality individuals. As leaders are unable to observe the recruits type, they face an adverse selection problem.

¹⁷ Considering the relatively low frequency of attacks by Islamist terrorist organizations given their substantial worldwide membership, we believe this assumption is reasonable. As the Washington DC sniper attacks proved, sowing terror through small-scale attacks is eminently feasible and can be done with very small expenditures. Given the ready availability of fire-arms in the United States, and the thousands of prospective terrorists who have received small-arms training since Osama Bin Laden issued his first fatwa calling for worldwide attacks against the United States in 1998, the lack of frequent, small-scale attacks is puzzling. One plausible explanation is that the population of true believers is actually fairly limited.

¹⁸ Shapiro (2007) examines these selection and recruiting dynamics in more detail.

There is good empirical evidence that terrorist middlemen often have preferences that are not aligned with their leadership. The adverse selection process and differential levels of risks faced by those filling different roles provide a mechanism for understanding why. Our model incorporates this insight explicitly into the middlemen's preferences and analyzes how this preference divergence affects groups' abilities to kill.

Model

To bring this dynamic into sharp focus, we begin with a fairly simple model that nevertheless captures the mechanisms at work here. There are only two actors: a terrorist boss whom we refer to as "B," and a single middleman, called "M." We assume that each period of the infinitely repeated game consists of the middleman's disbursing some fraction of the money he is given in order to fund an attack. The likelihood of a successful attack is increasing in this fraction; thus, the boss would prefer that all the money he hands to M be spent. On the other hand, if M has some innate preference for money—that is, he cares about this in addition to the success of the attack—then he has the incentive to skim some of the funds for himself.

Game Form

Each period consists of the following sequence of events: B gives w_0 to M, M spends $w_0 - x$ on the operation, keeping x for himself, and the attack succeeds with probability $p(w_0 - x; \alpha, \beta)$ with $p(\cdot)$ increasing, eventually concave, and twice differentiable in $w_0 - x$. Here α and β define the operational environment. α provides the baseline funding level necessary to achieve a certain success rate, and so represents the difficulty of successfully completing an attack. β determines the rate at which changes in funding are translated into success probabilities, and so represents the sensitivity of the success rate to funding.¹⁹ B observes only the outcome of the attack—whether or not it succeeds—and decides based on this whether or not to utilize M in the next period. We assume that the pool of potential middlemen is sufficiently large that B need not use a discarded M again, and that there is no cost (such as being ratted out by your former middleman) to changing M. For now, we also assume that all M are identical. This assumption is strong and will be relaxed in future work, but for now it allows us to focus on the essential aspects of the interaction before branching out into behavior across populations. The stage game just described is repeated infinitely, though in practice no particular M is retained for much of this time.²⁰

Actors

We begin with the boss, B, who is assumed to care primarily about the success probability of the attack, p ,²¹ but also about the amount spent on the attack. We assume a separable expected utility $E[U] = p - H(w_0)$ for each period. $H(0) \geq 0$, with $H' > 0$, $H'' \geq 0$, determines B's disutility from the loss of funds that go into the attack; w_0 is the amount given to M at the beginning of each period to use to fund the attack that period. The larger H' , the more money the boss would like to allocate to other goods such as social service provision. Different groups with

¹⁹ One can think of α and β as parameters of a location-scale family of cumulative density functions that defines the relationship between spending and success probability.

²⁰ This model is similar to that presented in Ferejohn (1986) insofar as B observes a noisy signal about M's performance, and hence his type, and has to choose retrospectively whether to hire B again in the next period. Here, however, in order to produce a signal—and achieve an appreciable chance of success—B must fund M before the signal occurs, resulting in a substantial change in both B's strategy space and M's decision calculus.

²¹ In this section and the following one we suppress the parameters of p for notational simplicity.

different goals would have different dependences of H on w_0 . For example, Hamas is estimated to spend somewhere in the vicinity of 90% of its annual budget on providing social services, so less than 10% is allocated to actual attacks (International Crisis Group 2003a:13). Suppose Hamas provided enough funds so that it almost always succeeded. In that case, this revealed preference would imply that w_0 spent on attacks would yield no more utility than $9 \cdot w_0$ spent elsewhere.

The middleman, M , also cares about both money and the likelihood of a successful attack, and is assumed to maximize an expected utility $E[V(x,p)] = \gamma v(x) + (1 - \gamma)p$ per period. $x \leq w_0$ is the amount M skims off the top, so that $w_0 - x$ is the actual amount used to fund each attack. For simplicity, we have again used a separable form for this utility, with $v' > 0$, $v'' \leq 0$. γ parameterizes the degree of pecuniary interest on the part of M . When $\gamma = 1$, the agent cares nothing for the cause; when $\gamma = 0$ the cause is everything. Notice that this framework makes no inherent claims about the type of middlemen a group will have. Thus, it can be used to analyze a wide variety of different organizations.

In this model we assume that B knows γ . The future will relax this assumption, but we believe it is justified here for two reasons. First, this model addresses the case of a homogeneous pool of M where we assume B has had time to develop accurate beliefs about the pool. Second, in many groups, the members work together for long periods of time. For example, most identified members and supporters of JI are also members of a long-standing informal network interested in establishing Islamic law in Indonesia (International Crisis Group 2002). Through their repeated interactions, we expect that the leaders would be able to develop fairly accurate beliefs about the commitment of their prospective middlemen.

Taken together, these parameters describe a particular period's payoff. We capture how utilities compare across time through the actors' discount factors. These parameters usually take one of two interpretations: either individuals value future payoffs less than present ones, or the game itself has some probability of ending each period, denying the accrual of future utility. We are agnostic between the two interpretations. In this model, the second interpretation can be understood as follows: δ is the probability an agent is not captured or killed. This probability is generated by something exogenous to this model: the skill level of the government. In a sense, values for this interpretation are already built into our model; B is guaranteed to be around in each period, while M faces a chance of being let go based on the outcome of his task. As such, we believe the first interpretation is the more plausible one for this model, even if it is empirically harder to identify.²²

Defining δ_B and δ_M as the discount factors, we obtain the two utilities: $E[U] = \sum_{t=0}^{\infty} \delta_B^t (p_t - H(w_{0,t}))$ and $E[V] = \sum_{t=0}^{\infty} 1_{\text{hired}} \delta_M^t (\gamma v(x_t) + (1 - \gamma)p_t)$, where 1_{hired} is an indicator function that equals 1 if M has been hired that period, and 0 if not.

Results

As we assume that B may costlessly choose a new, identical M whenever he chooses not to retain an old one, every period looks identical to him.²³ Because B therefore faces the same payoff-related incentives in every period, we make the reasonable simplifying assumption that B utilizes a stationary strategy, one that is constant in time.²⁴ As the strategy space for B consists of the amount of money paid to M , w_0 ,

²² It is easier to observe survival rates than it is to estimate discount factors from observed behavior.

²³ The plausibility of the assumption of identical M varies depending on the competitive environment for a group. Where government enforcement is lax, moving funds to operational cells requires no particular skill, and so most any agent can fulfill the task. When a group is under significant government pressure, moving funds can require substantial machinations, and thus the population of adequately skilled middlemen who care sufficiently about the cause may be quite small.

²⁴ We view this as reasonable because the only restriction it imposes is that B cannot condition on nonpayoff-related histories.

and rules for when to retain M in the face of both a successful attack and a failed one, which we will call q_S and q_F , respectively, this implies that each will be constant over time.

The retention rules, q_S and q_F , arise from the circumstances of the interaction. The lack of information for B beyond an attack’s success or failure implies the form of these retention rules can only depend on the history of such information for each agent. Note that this allows for steadily increasing (or decreasing) likelihoods of retention as successes (or failures) mount. As we are considering stationary strategies, however, we will assume the most basic of such rules: retain with probability q_F upon each failure, and retain with probability q_S upon each success. In this case, as M responds to B’s actions within a completely static environment, we will assume that M’s equilibrium strategy—how much to skim, x —will also be time independent. Our equilibrium concept is thus subgame perfect Nash equilibrium in stationary strategies.

We further assume that M cannot borrow money to add to the success of an attack. Of course, M also cannot skim more than is given to him. Under these assumptions we are presented with a constrained optimization problem in which B maximizes his utility, conditional on M’s maximizing his utility, subject to the constraint that $x \in [0, w_0]$. The full optimization breaks the problem up into four parts: the case where all the money is skimmed, the “honest” case where no money is skimmed, the interior or “skimming” solution, which generally obtains for greedy middlemen and rich bosses focused on attacks, and the “transition” region between “honest” and “skimming” regimes, in which no skimming occurs, but the boss nevertheless must consider his agent’s divergent interests.

In the next subsection we derive solutions for the equilibrium actions of both boss and middleman and briefly discuss how changing the assumptions of the model would alter this equilibrium. In the one following we examine how the equilibrium outcome changes with the model’s parameters, introducing general functional forms for p , H , and v to aid our intuition. Finally, “Discussion” adds a budget constraint for the boss in order to explore the model’s meaning in real-world terms and highlight some important nonlinearities not captured fully in the comparative statics provided in this section.²⁵

Equilibrium Strategies

Ignoring constraints for the moment, B solves the maximization problem:²⁶

$$\max_{w_0, q_S, q_F} p - H(w_0), \tag{1}$$

subject to:

$$\max_x \gamma v(x) + (1 - \gamma)p + \delta_M(pq_S + (1 - p)q_F)C(x^*, w_0, q_S, q_F; \gamma, \delta_M), \tag{2}$$

where $C(x^*, w_0, q_S, q_F; \gamma, \delta_M)$ is M’s continuation value.²⁷ Note that C depends on the future actions of M in equilibrium. We solve the game using backward induction,

²⁵ The budget constraint is not included in this section as here we only consider the interior solution of the constrained maximization problem.

²⁶ B’s optimization problem is set up as a static one because (1) he chooses an identical M whenever one is fired; and (2) we are considering stationary strategies. These imply that every period is identical to B, and so there is no incentive for B to alter his strategy after successes or failures. B’s discount factor thus does not impact the analysis.

²⁷ We can separate M’s utility given at the end of the previous section into terms corresponding to M’s present utility and M’s expected future utility, given that M will be playing his best response (the equilibrium value s^*), as we have assumed stationarity. Comparing this separated utility with Equation (2) implies that $C(x^*, w_0, q_S, q_F; \gamma, \delta_M)$ equals $(\gamma v(x^*) + (1 - \gamma)p(w_0 - x^*)) \sum_{i=0}^{\infty} \delta_M^i (p(w_0 - x^*)q_S + (1 - p(w_0 - x^*))q_F)^i$. As the summand is less than one, this yields (3) when summed.

beginning with the middleman’s problem, which is to choose the x^* that maximizes:

$$C = \frac{\gamma v(x) + (1 - \gamma)p}{1 - \delta_M(pq_S + (1 - p)q_F)}. \tag{3}$$

The form of C is suggestive: increasing q_S and decreasing q_F would both seem to increase C ’s dependence on p . As dependence on p correctly aligns the incentives of B and M, it is in B’s best interests to condition M’s utility on p as much as is possible. Thus we would expect that, in equilibrium, B would always rehire upon a success, and always fire upon a failure. We show this formally with the following two lemmas:²⁸

Lemma 1: x^* is decreasing in q_S , and the optimal $q_S^* = 1$.

Lemma 2: x^* is increasing in q_F , and the optimal $q_F^* = 0$.

We can therefore eliminate the retention rules from further analysis.²⁹ This simplifies (3), as well as the first-order condition that arises from it:³⁰

$$\frac{v'(x^*)}{p'} - \frac{\delta_M p v'(x^*)}{p'} - \delta_M v(x^*) = \frac{1 - \gamma}{\gamma}. \tag{4}$$

Equation (4) implicitly defines x^* . Recall that p is also a function of x^* and w_0 ; we have left off its dependence for readability.

Noting that $\frac{dp}{dw_0} = \frac{\partial p}{\partial w_0} + \frac{\partial p}{\partial x^*} \frac{dx^*}{dw_0}$, we can use the equilibrium value in (4) with (1) to find the second first-order condition:

$$p'(1 - \frac{dx^*}{dw_0}) = H'(w_0). \tag{5}$$

With (4) and (5) in hand, we can now find the full equilibrium, which consists of the pair (x^*, w_0^*) , along with the p resulting from these. There are four regimes in which an equilibrium might lie: (i) the “breakdown” region where M prefers to appropriate all the funds he is given; (ii) the interior “skimming” region where skimming is sustainable in equilibrium; (iii) the “transition” region between (ii) and (iv) in which M takes nothing in equilibrium, but would were he to be given more funds; and finally (iv) the “honest” region where M takes nothing regardless of how much he is given. Cutoffs between the regimes are in general complex functions of the parameters. As we are interested primarily in γ , however, we will only determine cutoff values in this parameter, but one should note that this is only a one-dimensional picture of the cutoffs in the larger parameter space.

First consider the case where the constraint $x = w_0$ binds, so M prefers to take the money and run. In this regime B would obtain the same minimal success probability— $p(0; \alpha, \beta)$ —as he would by not funding M at all, and so in equilibrium no money changes hands. While the equilibrium in this case is therefore trivial, *when* this breakdown of cooperation occurs is of vital importance. We find that this occurs

²⁸ Proofs for these and all other results in this section can be found in the Mathematical Appendix located on the Web at http://mailer.fsu.edu/~dsiegel/Research/Shapiro_Siegel_Underfunding_Math_Appendix.pdf.

²⁹ It is important to keep in mind that this simplification arises due to the homogeneity of the pool of middlemen and the lack of switching costs. Future work will relax these assumptions, resulting in more complex retention rules.

³⁰ In an infinitely repeated game, the continuation value is identical to the value of the game at its outset. Given our assumption of stationarity we may thus maximize (3) directly. A slightly more intuitive approach, but one that entails more notation, would be to calculate a first-order condition from M’s full utility, separated into present and future components, as discussed in a previous footnote. This yields $\gamma v'(x^*) - (1 - \gamma)p'(w_0 - x^*) - \delta_M(p'(w_0 - x^*)q_S - p'(w_0 - x^*)q_F) \times (\gamma v(x^*) + (1 - \gamma)p(w_0 - x^*))[(1 - \delta_M(p(w_0 - x^*)q_S + (1 - p(w_0 - x^*)q_F))]^{-1} = 0$. The total expected future payoff, contained in M’s continuation value, is thus explicitly compared with the marginal present benefit. Utilizing Lemmas 1 and 2, multiplying through by $\frac{1 - \delta_M p}{\gamma}$, and simplifying yields the same (4) as we obtained by maximizing (3), as it must.

in situations where the money granted to each M is limited, usually due to a budget constraint, and so we put off further discussion of this case until the next section.

Next, consider the regime in which neither constraint on M binds, so there is some potential for skimming. In order to analyze (5) in this case, we must first implicitly differentiate the equilibrium value x^* . The required derivative is

$$\frac{dx^*}{dw_0} = \frac{v'(x^*)[\delta_M(p')^2 + (1 - \delta_M p)p'']}{(1 - \delta_M p)[v''(x^*)p' + v'(x^*)p'']} \tag{6}$$

Combining (5) and (6) yields the condition:

$$H'(w_0^*) = \frac{(p')^2[v''(x^*)(1 - \delta_M p) - \delta_M v'(x^*)p']}{(1 - \delta_M p)[v''(x^*)p' + v'(x^*)p'']} \tag{7}$$

Equations (4) and (7) implicitly define (x^*, w_0^*) in the “skimming” region.

Now consider the “transition” region. Here x^* is zero, but B cannot simply provide his optimal level of funding, and so achieve his optimal chance of success, as he can in the “honest” regime described below. Here M will skim if B pays some w greater than the equilibrium value of w_0 , the point at which M’s utility from starting to skim balances exactly against the lowered probability of success this skimming would cause. More funding would allow M to skim while also increasing p ; mathematically this means that $\frac{dx^*}{dw_0} > 0$ in this regime. Thus γ_1 , the cutoff between the “skimming” and the “transition” regimes, is the highest value of γ such that $x^* = 0$, and is obtained by solving (4) with $x = 0$ for γ . This is

$$\gamma_1 = \left[1 + v'(0) \frac{[1 - \delta_M p_1]}{p'_1} - v(0)\delta_M \right]^{-1} \tag{8}$$

where p_1 is the equilibrium value of the interior solution obtained by decreasing the value of γ until M just stops skimming. The probability of success, $p_1(w_0)$, in the “transition” region is given implicitly by (4) with $x = 0$. It is worth noting again that the existence of the conflict of interest depresses funding levels even though no actual skimming takes place. We will see this graphically in “Discussion.”

Finally, consider the case where the constraint $x = 0$ binds M regardless of how much money he is given: M is sufficiently motivated by the cause so as to give up graft entirely. In this case x^* is unchanging regardless of how much funding M receives—it is simply 0 always—and so (5) becomes

$$p'_0 = H'(w_0) \tag{9}$$

This specifies p_0 , the equilibrium value of the success probability in the honest regime.

As B maximizes his own utility without regard to M’s desires here, p_0 is the ideal success rate for B, taking into account the other potential uses of his money. Not surprisingly, we find that this best possible situation for B occurs in general when middlemen are less desirous of money, at least in relation to the mission’s success. Define γ_0 as the cutoff between “transition” and “honest” regimes. This is the point at which both $x^* = 0$ and $\frac{dx^*}{dw_0} = 0$, and can be found by substituting p_0 for p_1 in (8).

Together, (4)–(9) provide equations that specify both x^* and w_0^* analytically in all four regimes, and give the cutoffs in γ between the last three as well. In the next subsection we will build an intuition for the model’s behavior by signing comparative statics of interest under the assumption of specific, empirically reasonable functional forms. Before doing so, we briefly consider the effect certain variations of the model would have on this equilibrium.

There are two slight alterations to the present model that can be productively examined without going too far afield: a change to the way in which B pays

M, and a change to the signal that B receives. We examine the logic of each in turn.³¹

Presently the boss is constrained to fund middlemen through what may be thought of as a block grant: both “salaries” and operational resources are taken from the same sum. While we justified this assumption in “Motivation,” it is not unreasonable to ask what would change if a payment, possibly contingent on the success of the attack, were offered in addition to the normal level of funding.

For the noncontingent payment, the answer is very little. When a change does occur, it is uniformly negative for the boss, implying that such a payment would never be used in equilibrium. This occurs because a noncontingent payment is indistinguishable from graft for M: increasing such a payment decreases graft equally much. Thus, this method cannot improve the outcome for B, but can hurt him because an M who would not want to engage in graft would also prefer the noncontingent payment be spent on operations. As B would rather have all the funding be usable for an attack—unless a different funding mechanism could spur less graft—there is no reason for B ever to use this mechanism. Essentially, the model mirrors what we see in reality: graft is often an implicit salary.³² Notice, though, that it is a salary only imperfectly defined; the middleman’s attributes dictate what he gets, as opposed to the boss’s choice.

In contrast, the existence of a contingent payment alters the outcome substantially. By providing a payment that depends on the success of the attack, the boss provides strong incentives to the middleman to minimize graft, as now success yields money in addition to its other benefits. Thus, while such a payment is not used in the regime where no graft takes place, it is used substantially when graft would otherwise be seen, with just enough of a payment so as to make M skim basically nothing off of the top of the block grant. Because receiving such a payment depends on p , which is decreasing in x , the net effect of contingent payments is to further align the motivations of the middlemen with those of B. Allowing such payments lets the boss achieve roughly the same outcome at less cost, lessening the budget constraint at which cooperation breaks down. The fact that we do not observe such payment is testament either to the fact that contingent payments entail extra communications, and thus a security cost, or, less positively, to our inability to discern terrorist organizations’ behavior.

In a different vein, we ask: What if the boss can roughly determine the level of graft in M, perhaps by noting recent changes in his middlemen’s lifestyles? Though the mathematics quickly get complex, we can discern the broad strokes of this change’s effect from the form of M’s utility, equation (3). Adding a signal of the level of M’s skimming impacts only B’s decision as to whether or not to fire a particular M in equilibrium,³³ effectively making δ a function of x . As this $\delta(x)$ will be decreasing in x —implying that agents who skim more have more of a chance of being discovered and fired—such a signal of skimming will provide an additional incentive for agents not to skim, with roughly the same impact that an equivalent decrease in δ would have. By not including such a signal, we are thus treating the worst-case scenario for the boss.

³¹ Because they add little to the overall analysis, we do not include the analytical and computational results that underlie this logic. They are available from the authors upon request.

³² Graft as an implicit salary could explain high levels of corruption within Palestinian militant organizations before the Oslo Accords.

³³ Lowering the level of w_0 is less efficient, as it results in a lower level of p for B than would otherwise be obtainable by increasing the likelihood of a firing.

Comparative Statics

As we saw in the previous subsection, the way in which M responds to changes in funding levels is of central importance to the decision making of both boss and middleman. The equation specifying this in the interior regime (when x is constrained it clearly does not vary) is given in (6), which tells us whether middlemen increase or decrease their skimming when the level of funding is increased. As the next result makes clear, the answer depends upon several factors.

Proposition 1: *Assume that p is increasing and concave over the region in which an interior solution obtains and that $v(x)$ is increasing. Then there exists an $\varepsilon > 0$ such that for all $\delta_M \leq \varepsilon, x^*$ is increasing in w_0 .*

Proposition 1 tells us that, when agents do not care much about the payoffs from future periods (a setting we can mimic either by setting $\delta_M = 0$ or $C = 0$), increasing the level of funding increases the amount skimmed. This is intuitive: a greedy and shortsighted agent has little motivation to spend money on an attack that provides a low level of utility, even if it should succeed. Unfortunately, we cannot say more than this; (6) can be both positive and negative in general, depending on γ, δ_M , and the functional form of p . This indeterminacy propagates through due to the centrality of M 's skimming decision, and we find that, without further assumptions, we cannot sign relevant comparative statics.³⁴

In order to circumvent this indeterminacy, we must make additional assumptions on our functions v, H , and p . In what follows, we assume what we believe are reasonable functional forms for each of the three, and derive a full range of comparative statics using these. Before doing so, we briefly introduce our functional forms, accompanied by empirical justifications of our choices whenever possible.³⁵

For p , we assume that an attack succeeds with probability

$$p(w_0 - x; \alpha, \beta) = \frac{e^{\beta(w_0 - x - \alpha)}}{1 + e^{\beta(w_0 - x - \alpha)}}, \tag{10}$$

where α and β define the difficulty of successfully completing the attack and the sensitivity of the success rate to funding, respectively. This form for the probability function has three distinct advantages. First, it matches our intuition: for low levels of spending, success rates increase slowly; at moderate levels of spending, each additional unit of spending greatly enhances the probability of success; once spending levels reach a certain level, there are diminishing marginal returns in terms of success. Second, it asymptotes nicely at 0 and 1, so we do not need to truncate it artificially. Third, using this form the parameters of our model can be estimated using a standard logistic regression, although doing so requires data on how much is being spent on attacks and on the rate of both successful and failed attacks.³⁶

³⁴ In fact, even when $\delta_M = 0$, the sign of many comparative statics is in general indeterminate.

³⁵ It should be noted that, though we cannot sign derivatives for the fully general case, our results do not rest entirely on the exact functional forms of p, H , or v . Holding the other two functions constant, as long as p is always increasing and concave over the region for which a "skimming" equilibrium exists, we can reproduce nearly all the results below. Most of the comparative statics continue to hold as well as long as H is a convex function. The only real difference is that the boss is less likely to increase funding to greedy agents, so that p is no longer constant with γ in equilibrium. Letting v be nonlinear, but concave, results in greater complexity, and here we can only say that the middleman still reacts to changes in the parameters in the same way, holding funding constant, only he now values the future more due to his risk aversion. Because we believe the simplifications we use are reasonable and do not feel that the significantly increased algebraic complexity entailed in relaxing them yields additional insight, we do not report these more general results here. They are available upon request.

³⁶ Several datasets provide information on terrorist successes and failures in the Palestinian–Israeli conflict. Unfortunately no time-series data exist on the costs of terrorist attacks. For data on failed and foiled attacks from October 2000 to August 2004, see Merari (2004). For data that includes successful attacks, Israeli interdictions, and targeted killings by Israel, see Kaplan et al. (2005).

For v , we assume M's overall level of wealth is large enough that his decisions regarding any one attack are made within the linear section of his utility curve.³⁷ Thus we define $v(x) = bx$.

Finally, for H we again make the empirically reasonable assumption that the range of money B pays out to M in equilibrium is a relatively minor fraction of B's overall wealth.³⁸ This allows us to treat B's utility function as linear in this range, so that $H(w_0) = cw_0$. Using these three functions with (4)–(9) yields the following lemma.

Lemma 3: *In the interior, "skimming" region, and thus in the "transition" and the "honest" regions as well, $p > \frac{1}{2}$. Further, $\frac{dx^*}{dw_0} > 0$ whenever $\delta_M < \frac{2p-1}{p^2}$, and less than or equal to zero otherwise.*

Lemma 3 states that, when there is no budget constraint, B will not fund M unless the probability of success is at least one-half. Further, there are numerous circumstances in which M's skimming becomes *less* of an issue the *more* money he receives from B. The reason for this perhaps counterintuitive result can be seen in M's utility function, given in (3). M discounts the future according to both δ_M and the probability that he will be rehired, which here is just p .³⁹ The form of the utility is therefore very sensitive to changes in p , producing increasing returns for M's decreasing the amount he skims in certain regions. Owing to the functional forms of p and M's utility, these increasing returns occur when p is close to $\frac{1}{2}$, and when there are high discount factors. In this region, therefore, a small increase in funding for M can lead to a substantial increase in the likelihood that an attack would succeed, a result we discuss further in the next section. There we also discuss the impact of a budget constraint in more depth. Such a constraint actually makes B's decision making easier in the region in which it binds, in that B's decision devolves to whether or not to fund with the entire budget or with nothing at all.

Proposition 2: *Increasing the disutility B obtains for funding attacks— c —decreases the level of funding B offers M and decreases the probability of success of the attack in both "skimming" and "honest" regions, while having no impact on either in the "transition" region. Increasing c does alter the range in γ over which the transition region obtains, though, increasing both γ_0 and γ_1 . (Recall that γ_0 is the cutoff between "transition" and "honest" regimes, and γ_1 is the cutoff between "skimming" and "transition" regimes).*

The statement that increasing B's disutility for funding decreases funding in equilibrium is hardly surprising. Lemma 3 makes the rest of the statement less obvious, however, as the amount of skimming can be either increasing or decreasing in funding depending on δ_M . Proposition 2 cuts through the indeterminacy and states assuredly that the more B values actions other than attacks, the less likely attacks are to succeed.

The results involving the transition region are more complex, but may be understood by recalling that this region occurs when B prefers not to increase his funding, because then M would start skimming. Thus the equilibrium values of p

³⁷ Linearity is a common assumption, is extremely tractable, and no other functional form, such as $v(x) = \sqrt{x}$, is obviously better. Moreover, we rarely see individuals getting rich by supporting terrorism, or using money from supporting terrorism as their sole source of income. Where they do, it is on the fundraising end through skimming money from fundraising ventures. For examples of this, see Adams (1986:85, 103).

³⁸ Recall we showed in the introduction that attacks appear to consume a small portion of groups' overall resources.

³⁹ Note that p is only equivalent to the probability of being rehired by virtue of Lemmas 1 and 2. Unlike the rehiring probabilities q_S and q_F , p depends explicitly on both x and w_0 in equilibrium. M therefore will not react to a high probability of being rehired by skimming more, as this would cause his chance of being rehired to decrease, resulting in a less favorable outcome for M.

and w_0 are independent of c in this region. Despite this, c does have an effect; increasing c means that B's optimal funding level is lowered, increasing γ_0 as he is able to achieve this lower level with more greedy agents. Further, as he wants less to subsidize skimming, γ_1 increases as well.

Proposition 3: *Increasing M's value for money or M's greed— b or γ —in the “skimming” region increases both the funding B offers M and the amount M skims, while leaving the probability of the attack's success constant. In the “transition” region, increasing both decreases both funding and success probability, while in the “honest” region, altering either parameter has no effect on funding or the probability of success. In addition to these within-region variations, raising b lowers both cutoffs γ_0 and γ_1 so that less greedy agents are required to ensure better scenarios for B.*

Again the statement of the proposition contains more and less surprising results. While it is intuitive that increasing M's greed or value for money should yield additional skimming, it is by no means obvious that B should simply replace the extra money lost to skimming, so as to keep the likelihood of success the same.⁴⁰ In the absence of a budget constraint, utilizing greedier agents does not alter the outcome of the attempted attack in the “skimming” region. What it *does* alter, as we shall see in the next section, is what level of budget constraint binds, and thus the point at which cooperation within the organization breaks down.

Proposition 4: *Increasing the degree to which M values the future, his discount rate δ_M , has an indeterminate effect on both skimming and funding levels in the “skimming” region, but increases the probability of success of an attack. It raises both funding and the probability of success in the “transition” region, and has no effect on either in the “honest” region, though it does raise γ_0 , allowing more greedy agents to be used in the “honest” region.⁴¹*

The feedback mechanisms discussed after Lemma 3 make signing comparative statics involving δ_M difficult: increasing M's patience leads to lower levels of skimming, which, depending on other parameters, can lead to decreased funding, which can then lead to increased skimming, and so on. Despite this, the effect on the mission's chance of success is clear: a more patient M implies a greater likelihood of success.

Proposition 5: *In all regions in which funding is provided, creating a more difficult environment for terrorist activity through increasing α increases levels of funding by the same amount, leaving skimming and the probability of success unchanged. Increasing the sensitivity of the success probability to the net level of funding through increasing β decreases skimming and has an indeterminate, though usually negative, effect on funding in the “skimming” region. Increasing β increases the success probability in equilibrium in all regions in which funding is provided and also raises both γ_0 , allowing more greedy agents to be used in the “honest” region, and γ_1 , so that more agent types fall into the “transition” region.*

The difference between the equilibrium dependence on each type of environmental parameter clarifies our choice of functional form for p . The first, α , has no effect on the likelihood of success until the budget constraint binds,⁴² but increasing α does directly decrease the slack in B's budget. Thus we would expect a nonlinear impact of increasing α : no effect until the budget constraint is reached, and a big effect afterward as cooperation breaks down. This is what we see in the next section.

⁴⁰ This is due to B's assumed linear preference for money, as discussed briefly in footnote 35.

⁴¹ Increasing δ_M also increases γ_1 for all $\delta_M > \frac{1}{2}$.

⁴² This again is due to the linearity of B's disutility for money.

In contrast, β has a more gradual impact, in that increasing it makes the cost of underfunding greater for M and provides better information to B about the likelihood of M's failing to fund the attack sufficiently.

Discussion

Thus far we have shown that an equilibrium of the game exists and illustrated analytically how it changes when the parameters shift. Here we flesh these results out by exploring computationally the effect of a budget constraint on the boss–middleman interaction⁴³ utilizing the specific functional forms discussed in “Comparative statics.” This illustration reveals important nonlinearities in terrorists' responses to government action under very reasonable assumptions about functional forms and parameterizations.

We begin by choosing values for the operational environment, α and β . As estimating these parameters for a particular conflict is impossible without data on the costs of attacks, we arbitrarily set $\alpha = 600$ and $\beta = 0.005$. With $\beta = 0.005$ groups have a 10% success rate when funding per attack is at \$300.

Given extant data on Al Qaeda and affiliated groups such as JI, we were able to use realistic values for δ_M . Here we consider the interpretation that δ represents the probability that the game will not be ended by an exogenous shock such as government action. We found that the average survival rate for middlemen between 1997 and 2003 was approximately 83%.

We set c , the boss' marginal disutility for spending on attacks, so that B's utility is approximately zero for the case of p near 1 with $\gamma = 0$, so that zero utility corresponds to receiving a particular payoff from the least greedy agent. This yields an estimate of $c = 1/(4\alpha)$, which is low enough that the boss would want to fund an attack up to his ideal level if skimming by the middlemen were not an issue. We find this reasonable, as bosses with higher values of c than this would be uniformly less willing to fund attacks in all cases, and thus be less of the type we are studying here. Nevertheless, we do discuss what happens if c increases below.

We use a similar strategy to calculate b . As we have a different parameter, γ , to describe M's relative utility for increasing success probabilities and skimming money, all b must do is put dollar amounts on the same footing as probabilities, which are bounded by 0 and 1. Accordingly, we calculate b by equating M's utility over a range of parameters, setting a perfectly greedy middleman's equilibrium utility at $p = \frac{1}{2}$, equal to a perfectly committed M's utility, also at $p = \frac{1}{2}$. This yields an estimate of $b = \frac{\beta\delta}{8(1-\delta/2)}$. While not perfect, we know from the previous section that positive changes to this parameter have consistent positive effects on the equilibrium levels of funding and skimming for $p > \frac{1}{2}$, and so our exact choice is not important.

To summarize, we use the following parameters to explore the model:⁴⁴

- $\alpha = 600$.
- $\beta = 0.005$.
- $\delta_M = 0.83$.
- $c = 0.00042$.
- $b = 0.00089$.

⁴³ The budget constraint places a limit on how much the boss can allocate to funding the middleman. The computational procedure for finding the equilibrium does not consider funding levels above the budget constraint.

⁴⁴ We did not choose set values for γ , as that is our key causal variable, and in all the illustrations we either explore its full range or estimate results for $\gamma = \{0.25, 0.5, 0.75, 0.9, 0.1\}$. We do not consider the perfectly committed middleman, as the case is trivial in this model. In future models where the boss wants to achieve a specific level of impact, rather than maximizing his success probability, we will consider the case where a boss gets more impact than he wants by hiring a perfectly committed agent, an agent for whom $\gamma = 0$.

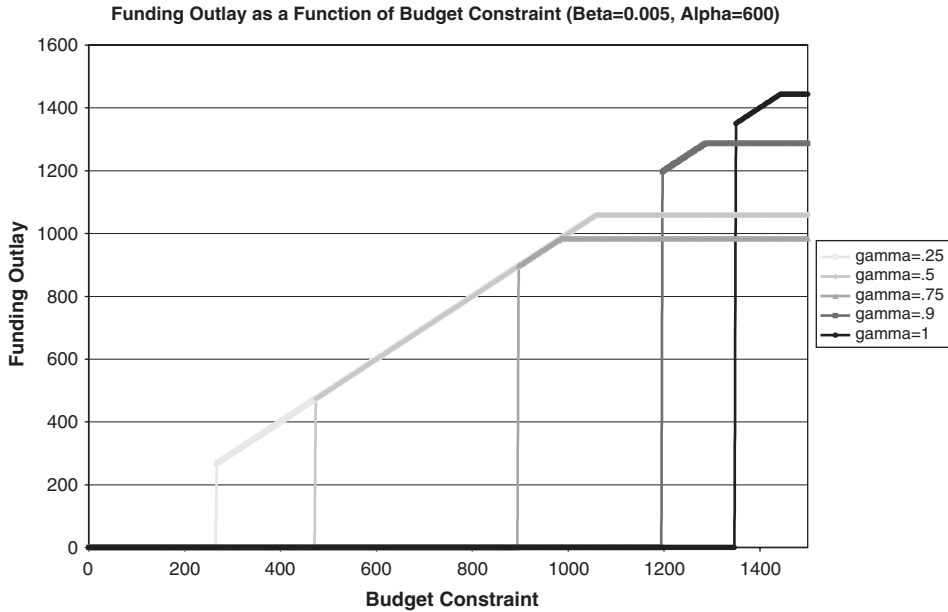


FIG. 1. Funding Outlay as a Function of Budget Constraint ($\beta = 0.005, \alpha = 600$)

Using these estimates, we solve numerically for the constrained maximization problem described in the previous section. This amounts to maximizing M’s utility over the range $x \in [0, w_c]$, where w_c is the budget constraint in effect, and then maximizing B’s utility over the range $w_0 \in [0, w_c]$, for a range of γ .⁴⁵ Allowing w_c to vary from zero to a point high enough such that it would not bind anyone’s actions—in this case, to \$1,500—provides us with the ability to explore the entire range of potential interactions, from the point at which M would rather just take the money and run, to the point where B has achieved a sufficient success rate and chooses to spend the rest elsewhere. The results of this are summarized in Figures 1–4.

Figure 1 shows the amount B spends as a function of his budget constraint and how venal his middleman is. Figure 2 shows the amount M skims, and Figure 3 shows the probability that the attack will succeed, both as a function of w_c and of γ .

There are two major things to note here. First, the equilibrium level of funding in Figure 1—that which occurs at a budget constraint of \$1,500, and sometimes considerably less—does not linearly increase with venality. While that is what we would intuitively expect, as greedier middlemen might require more money to overcome their increased skimming and ensure a reasonable level of success, Figure 2 shows why it is not the case. Below some level of venality, it no longer pays for M to skim at all, and the boss can eke out an additional 10% in the success rate by upping the amount of money he pays to M relative to what he would be willing to do if skimming were an option for M. The existence of this cutoff in γ —derived analytically in the previous section—leads to a clear policy implication: state action that serves to remove the most venal of the available middlemen can lead to a sizable *increase* in the likelihood of terrorist success if those middlemen are replaced by more committed types. Increasing the risks to middlemen is likely to have a similar impact by pushing less committed individuals to give up the cause. This

⁴⁵ Essentially, w_c serves as an endpoint in the players’ strategy space in the computational search for an equilibrium.

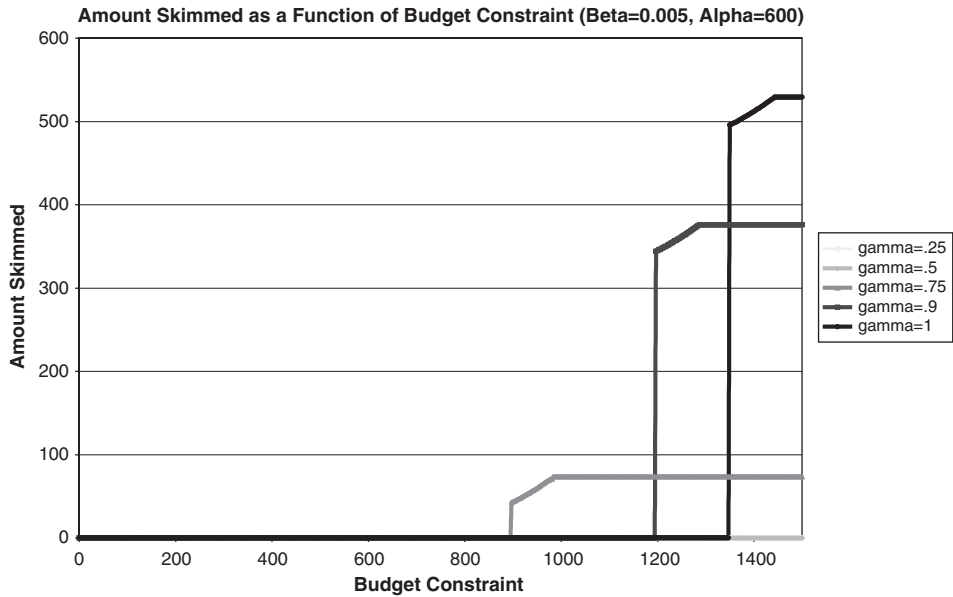


FIG. 2. Amount Skimmed as a Function of Budget Constraint ($\beta = 0.005$, $\alpha = 600$)

suggests that efforts to limit terrorists' finances should be focused on interdicting funds, rather than on capturing financiers.

Second, the behavior of each quantity of interest is decidedly nonlinear in the budget constraint. Figure 3 clearly illustrates this dynamic. For all levels of middleman greed, there is a constraint level below which the chance of an attack's succeeding becomes small due to the breakdown of cooperation between boss and middleman.⁴⁶ Essentially, there is a point at which B prefers to spend all the money on nonattack goods rather than on achieving the low success probability he can get given his budget and M's greed. However, once the constraint exceeds that level—which will be different for different values of the other parameters—the chance of success undergoes a discontinuous increase, up to 80% in the case of the more venal agents. Past this point, increasing the level of funds available to B causes a more gradual, nearly linear increase in all the variables, until an equilibrium is reached, whereupon B no longer desires to increase the likelihood of success. This leads to a second clear policy implication: interdicting terrorists' funds can provide dramatic decreases in the likelihood of an attack, even in cases where previous reductions in available funds seemed to have a more gradual impact.

Note that government actions, which make the operational environment more challenging, increasing α , should have effects similar to tightening the budget constraint. Thus we expect important nonlinearities in the returns to government counter-terrorism. For some budget constraints and levels of α , it may appear that counter-terror policies are having no effect. However, once government makes the environment hard enough so that the budget constraint starts to bite, there should be dramatic decreases in the level of terrorist activity. Government actors that do not recognize the existence of this nonlinearity may cease eventually effective actions too soon.

Figure 4 displays results of a different sort, mirroring the comparative statics described in the previous section. Here we examine how the equilibrium values of

⁴⁶ It is only non-zero due to the form of the success probability and the particular parameter estimates in use.

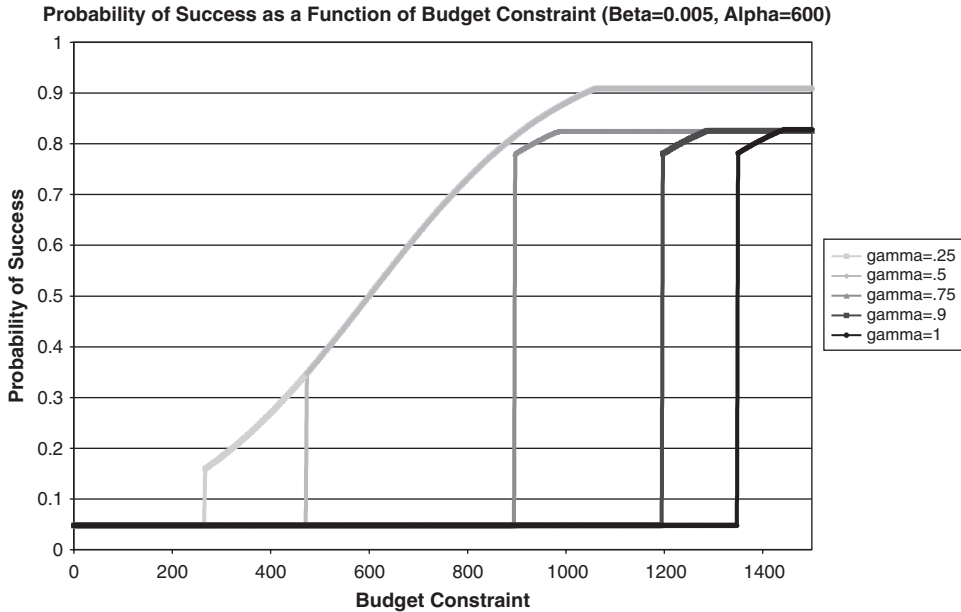


FIG. 3. Probability of Success as a Function of Budget Constraint ($\beta = 0.005, \alpha = 600$)

$p, x,$ and w_0 vary with γ when no budget constraint binds. We note several things. Past the cutoff point given earlier in equation (8), which for these parameters yields $\gamma_1 = 0.72$, the equilibrium values of both skimming and funding come from the interior “skimming” solution described in the previous section. Accordingly, they increase at the same rate with venality, and the probability of success therefore remains constant. Below this point, the situation is much different. M does not skim, and so the boss ups his financial outlay to achieve a higher success rate. This is the “transition” region discussed in the previous section, and here the success rate continues to climb with further decreases in γ until a second cutoff is reached, at $\gamma_0 = 0.65$. At this point B has achieved his optimal level of funding given his other uses for money—the level he would choose if he did not require M—and further decreases in γ have no effect on the equilibrium.⁴⁷

Because we are primarily interested in γ and the budget constraint, due in large part to their clear relation to policy, we do not produce similar figures for all the other parameters in the model. Thanks to the work of the previous section, however, we know what varying each of the remaining parameters would entail.

Increasing c decreases the probability of a successful attack, and shifts the cutoffs between regions to the right. It also decreases the area over which a non-zero equilibrium obtains, and a large enough c can cause cooperation to vanish entirely. Thus increasing the relative utility of nonviolent uses of money can have a disproportionate effect on reducing the rate of successful attacks. This finding thus provides an explicit organizational mechanism by which encouraging groups using terrorist tactics to enter the political arena—or to engage in social service provision—can reduce the level of violence.

Increasing b has almost the same effect as increasing venality, and so is of little additional interest. From the previous section, increasing δ_M increases the probability of success of an attack as long as M’s skimming or potential skimming has an

⁴⁷ Figure 4 also reinforces the point that removing the most venal middlemen may lead to sudden increases in terrorists’ success rates.




FIG. 4. Equilibrium Values with no Budget Constraint as a Function of γ ($\beta = 0.005$, $\alpha = 600$)

effect on B's funding levels. Further, increased patience mitigates increased greed, to some extent, in that increasing δ_M generally causes more greedy agents not to skim in equilibrium. We infer from the form of the utility for M, (3), that increasing δ_M results in a tighter budget constraint being necessary for cooperation to collapse.

Increasing α has no effect on the probability of success when there is no budget constraint, but it does cause an existing constraint to bind sooner, leading to the outcomes described above. Finally, increasing β increases the success probability in equilibrium, and reduces the types of M who choose to skim in equilibrium, shifting the skimming region to the right in Figure 4. Computational work indicates that increasing β also produces a wider range of budgets in which cooperation can occur. This is because an increase in β makes the transition from a high probability of success to a low probability more abrupt, effectively increasing the informativeness of an attack's success or failure which, much like the additional signal of skimming discussed in the previous section, is generally good for B.

Conclusion

We have presented a model of a hierarchical terrorist organization in which leaders must delegate financial and logistical tasks to middlemen for security reasons. However, these middlemen do not always share their leaders' interests. In particular, the temptation always exists to skim funds from any financial transaction. To counteract this problem, leaders can threaten to punish the middlemen. Because logisticians in terrorist organizations are often geographically separated from leaders, and because they can defect to the government if threatened, violence is rarely an effective threat. Therefore, leaders must rely on more prosaic strategies to solve this agency problem; we focused on leaders' ability to remove middlemen from the network, denying them the rewards of future participation.

Because our model is inherently hierarchical, its explanatory scope is limited to those groups, which operate through a defined chain of command. As such, the model's applicability to Al Qaeda and related groups may be diminishing. Before

the loss of their Afghan sanctuary, these groups operated through defined hierarchies with relatively clear distinctions between operational and supporting roles.⁴⁸ More recent attacks have followed a different, more bottom-up pattern. The London, Madrid, and Casablanca bombings were locally initiated and in all three cases some individuals undertook both logistical and operational roles. While this pattern is apparent in some groups, the majority of terrorist organizations worldwide remain hierarchically organized, and thus amenable to the principal-agent analysis we have conducted.

Our analysis yielded several important policy implications. The first is that the removal of the most venal middlemen, those who are the easiest to identify and will be most vulnerable to government incentives, actually makes the terrorist boss' problem easier and can lead to a jump in the number of successful attacks. This is in line with others' findings that as moderates leave an organization, it often becomes more violent (Bueno de Mesquita 2005a). The second is that reducing a group's available funds below a certain threshold can have a dramatic impact, even if previous reductions yielded only gradual effects. The third is that a similar threshold effect exists for counter-terror efforts that make it harder to conduct attacks. Such efforts may yield no result until they make the environment hard enough that it is not worth it for a group to attempt any attacks. The fourth is that efforts to restrict funding to terrorist organizations are unlikely to reduce attacks unless they can reach a threshold, which depends on the operational environment and the nature of the individuals in the group. This finding should give pause to policy makers who must balance spending on counter-terrorist financing efforts against other, less socially damaging, counter-terror efforts. Finally, our model strongly suggests that increasing the value of nonviolent uses of money, perhaps by encouraging entry into the political process, can have a disproportionate effect by making it impossible for leaders and middlemen to cooperate.

Overall, our main finding is that the level of greed among middlemen is central to the internal dynamics of terrorist groups. All of the boss' problems become harder when they have to rely on less ideologically committed individuals. While this result is by no means counter-intuitive, its implications deserve consideration. Rather than differentiate terrorist organizations solely on the basis of their strategic goals, or ideological style, our theories should take into account their internal make-up. Specifically, we should seek to understand how the motivations of members who do not volunteer for the riskiest activities impact groups' behavior. Doing so will lead to a more thorough, more effective understanding of terrorism.

References

- ABUZA, ZACHARY. (2003a) Funding Terrorism in Southeast Asia: The Financial Network of Al Qaeda and Jemaah Islamiyah. *Contemporary Southeast Asia* 25(2):169–199.
- ABUZA, ZACHARY. (2003b) Funding Terrorism in Southeast Asia: The Financial Network of Al Qaeda and Jemaah Islamiyah. NBR Analysis 14:5.
- ADAMS, JIM. (1986) *The Financing of Terror*. New York: Simon and Schuster.
- AL-ZAWAHIRI, AYMAN. (2001) Knights Under the Prophet's Banner, part 11. Al-Sharq al-Awsat (London) December 2, trans. Foreign Broadcast Information Service (FBIS). Available at http://www.fas.org/irp/world/para/ayman_bk.html. (Accessed March 1, 2005).
- ASIA/PACIFIC GROUP ON MONEY LAUNDERING. (2004) *APG Annual Typologies Report*. Sydney: Asia/Pacific Group on Money Laundering.
- BERMAN, ELI, AND DAVID LAITIN. (2005) Hard Targets: Theory and Evidence on Suicide Attacks. NBER Working Paper 11740.

⁴⁸ Felter (2006) examine captured Al Qaeda documents detailing the specific goals and responsibilities for those filling different roles within Al Qaeda.

- BOUNTY FOR MURDER OFFERED BY IRAN'S REVOLUTIONARY GUARDS, SYRIA, HIZBALLAH. (2004). *DEBKA-Net-Weekly*, February 26. Available at <http://www.debka-net-weekly.com>. (Accessed March 15, 2004).
- BRISARD, JEAN-CHARLES. (2002) Terrorism Financing: Roots and Trends of Saudi Terrorism Financing, Report Prepared for the President of the Security Council United Nations. New York: JCB Consulting.
- BUENO DE MESQUITA, ETHAN. (2005a) Conciliation, Commitment, and Counter-Terrorism. *International Organization* 59(1):145–176.
- BUENO DE MESQUITA, ETHAN. (2005b) The Quality of Terror. *American Journal of Political Science* 49(3):515–530.
- BUENO DE MESQUITA, ETHAN. (2005c) The Terrorist Endgame: A Model with Moral Hazard and Learning. *Journal of Conflict Resolution* 49(2):237–258.
- CHAI, SUN-KI. (1993) An Organizational Economics Theory of Antigovernment Violence. *Comparative Politics* 26(1):99–110.
- CRENSHAW, MARTHA. (1991) How Terrorism Declines. *Terrorism and Political Violence* 3(1):69–87.
- CULLISON, ALAN. (2004) Inside Al Qaeda's Hard Drive. *The Atlantic Monthly* 294:55–65.
- DONOHUE, LAURA K. (2006) Anti-Terrorist Finance in the United Kingdom and United States Part II: How Effective? *Michigan Journal of International Law* 27(4):303–435.
- FARAH, DOUGLAS. (2004) *Blood From Stones: The Secret Financial Network of Terror*. New York: Broadway Books.
- FELTER, JOSEPH, *et al.* (2006) *Harmony and Disharmony: Exploiting al-Qa'ida's Organizational Vulnerabilities*. Westpoint: Combating Terrorism Center.
- FEREJOHN, JOHN. (1986) Incumbent Performance and Electoral Control. *Public Choice* 50(1–3):5–25.
- FINANCIAL ACTION TASK FORCE. (2004) *Report on Money Laundering and Terrorist Financing Typologies: 2003–2004*. Paris: Financial Action Task Force.
- FRANCIS, DAVID R. (2004). The war on terror money. *Christian Science Monitor*, April 8.
- FREY, BRUNO S., AND SIMON LUECHINGER. (2004) Decentralization as a Disincentive for Terror. *European Journal of Political Economy* 20(2):509–515.
- GUNARATNA, ROHAN. (2002) *Inside Al Qaeda*. New York: Columbia University Press.
- GUNNING, JEROM. (2004) Peace With Hamas? *International Affairs* 80(2):233–255.
- INTERNATIONAL CRISIS GROUP. (2002) Al-Qaeda in Southeast Asia: The case of the 'Ngruki Network' in Indonesia. ICG Indonesia Briefing 1.
- INTERNATIONAL CRISIS GROUP. (2003a) Islamic Social Welfare Activism in The Occupied Palestinian Territories: A Legitimate Target? International Crisis Group Middle East Report 13.
- INTERNATIONAL CRISIS GROUP. (2003b) Jemaah Islamiyah in South East Asia: Damaged but Still Dangerous. ICG Asia Report 63.
- INTERNATIONAL CRISIS GROUP. (2004) Indonesian Backgrounder: Jihad in Central Sulawesi. ICG Asia Report 74.
- ISRAELI DEFENSE FORCES, MINISTRY OF PARLIAMENTARY AFFAIRS. (2002) The Involvement of Arafat, 42 PA Senior Officials and Apparatuses in Terrorism against Israel, Corruption and Crime. Available at <http://mfa.gov.il/mfa/go.asp?MFAH0lom0>. (Accessed March 14, 2005).
- JENKINS, BRIAN MICHAEL. (2002) *Countering Al Qaeda*. Santa Monica: RAND.
- 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* 28:225–235.
- LAPAN, HARVEY E., AND TODD SANDLER. (1993) Terrorism and Signaling. *European Journal of Political Economy* 9(3):383–397.
- LEE, RENSSLAER. (2002) Terrorist Financing: The U.S. and International Response. Congressional Research Service Report for Congress, RL31658.
- LEVITT, MATTHEW A. (2002) The Political Economy of Middle East Terrorism. *The Middle East Review of International Affairs (MERIA) Journal* 6(4):49–65.
- MEDANI, KHALID. (2002) Financing Terrorism or Survival? Informal Finance, State Collapse and the U.S. War on Terrorism. Middle East Report 32:2–9.
- MERARI, ARIEL. (2004) Suicide Terrorism in the Context of the Israeli-Palestinian Conflict, Paper presented at the National Institute of Justice Suicide Terrorism Conference, October 25–26, Washington, DC.
- NATIONAL COMMISSION ON TERRORIST ATTACKS UPON THE UNITED STATES. (2004a) *The 9/11 Commission Report*. New York: W.W. Norton and Company Inc.
- NATIONAL COMMISSION ON TERRORIST ATTACKS UPON THE UNITED STATES. (2004b) Staff Statement no. 15: Overview of the Enemy.

- OVERGAARD, PER BALTZER. (1994) The Scale of Terrorist Attacks as a Signal of Resources. *Journal of Conflict Resolution* 38(3):452–478.
- PROBER, JOSHUA. (2005) Accounting for Terror: Debunking the Paradigm of Inexpensive Terrorism. Washington Institute for Near East Policy, Policy Watch #1041.
- REES, MATT. (2001) How Hamas-Hezbollah Rivalry Is Terrorizing Israel. *Time*, April 23.
- RESSA, MARIA A. (2003) *Seeds of Terror*. New York: Free Press.
- ROTH, JOHN. (2004) *Monograph on Terrorist Financing: Staff Report to the Commission*. Washington, DC: National Commission on the Terrorist Attacks Upon the United States.
- SAGEMAN, MARC. (2004) *Understanding Terror Networks*. Philadelphia: University of Pennsylvania Press.
- SANDLER, TODD. (2003) Collective Action and Transnational Terrorism. *World Economy* 26:779–802.
- SECOND REPORT OF THE [UN] MONITORING GROUP. (2002) *Pursuant to Security Council Resolution 1390*. New York: United Nations.
- SHAPIRO, JACOB N. (2007) The Greedy Terrorist: A Rational Choice Perspective on Terrorist Organizations' Inefficiencies and Vulnerabilities. In *Terrorist Financing in Comparative Perspective*, edited by Harold Trinkunas and Jeanne K. Giraldo. Stanford: Stanford University Press.
- SINGAPORE MINISTRY OF HOME AFFAIRS. (2003) The Jemaah Islamiyah Arrests and the Threat of Terrorism.
- SIQUEIRA, KEVIN. (2005) Political and Militant Wings within Dissident Movements and Organizations. *Journal of Conflict Resolution* 49(2):218–236.
- STERN, JESSICA. (2003) *Terror in the Name of God: Why Religious Militants Kill*. New York: HarperCollins.
- TESTIMONY OF FBI AGENT. (2001) Testimony of FBI Agent John Anticev on Odeh, February, United States of America v. Usama bin Laden, et al. 5 (7) 98 Cr. 1023, 27, 1630–1638.
- VICTOROFF, JEFF. (2005) The Mind of the Terrorist: A Review and Critique of Psychological Approaches. *Journal of Conflict Resolution* 49(1):3–42.
- WEINSTEIN, JEREMY. (2005) Resources and the Information Problem in Rebel Recruitment. *Journal of Conflict Resolution* 49(4):598–624.
- WESCHLER, WILLIAM F., AND LEE S. WOLOSKY. (2002) *Terrorist Financing: Report of an Independent Task Force Sponsored by the Council on Foreign Relations*. New York: Council on Foreign Relations.
- WESCHLER, WILLIAM F., AND LEE S. WOLOSKY. (2004) *Update on the Global Campaign Against Terrorist Financing*. New York: Council on Foreign Relations.

