This article was downloaded by: *[Princeton University]*
On: *14 April 2010*
Access details: *Access Details: [subscription number 917274464]*
Publisher *Routledge*
Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK

Security Studies
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713636712

Is this Paper Dangerous? Balancing Secrecy and Openness in Counterterrorism

Jacob N. Shapiro [a];David A. Siegel [b]
[a] Princeton University, [b] Florida State University,

Online publication date: 23 February 2010

To cite this Article Shapiro, Jacob N. andSiegel, David A.(2010) 'Is this Paper Dangerous? Balancing Secrecy and Openness in Counterterrorism', Security Studies, 19: 1, 66 — 98
To link to this Article: DOI: 10.1080/09636410903546483
URL: http://dx.doi.org/10.1080/09636410903546483

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: http://www.informaworld.com/terms-and-conditions-of-access.pdf

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Security Studies, 19:66–98, 2010
Copyright © Taylor & Francis Group, LLC
ISSN: 0963-6412 print / 1556-1852 online
DOI: 10.1080/09636410903546483

Downloaded By: [Princeton University] At: 14:30 14 April 2010

Is this Paper Dangerous? Balancing Secrecy and Openness in Counterterrorism

JACOB N. SHAPIRO AND DAVID A. SIEGEL

We analyze a seemingly simple question: When should government share private information that may be useful to terrorists? Policy makers' answer to this question has typically been "it is dangerous to share information that can potentially help terrorists." Unfortunately, this incomplete response has motivated a detrimental increase in the amount of information government keeps private or labels "sensitive but unclassified." We identify two distinct types of private information that are potentially useful to terrorists and identify the range of conditions under which sharing each can enhance counterterrorism efforts. Our results highlight the complex trade-offs policy makers face in deciding how much openness is right in a world where protecting the people from terrorists has become a central duty of government.

Jacob N. Shapiro is Assistant Professor of Politics and International Affairs at Princeton University. His primary research interests are terrorism, insurgency, and security policy. Shapiro co-directs the Empirical Studies of Conflict Project. His research has been published in *International Security, International Studies Quarterly, Security Studies, Foreign Policy, Military Operations Research*, and a number of edited volumes.

David A. Siegel Assistant Professor of Political Science at the Florida State University. He has research interests in political violence, terrorism and collective action. His research has been published in the *American Journal of Political Science, International Studies Quarterly*, the *Journal of Theoretical Politics, Security Studies, The B.E. Journal of Theoretical Economics,* and *PS: Political Science & Politics*.

The authors thank Tom Christensen. James Fearon, Charles Perrow. Robert Powell, Mark Souva, Paul Stockton, Jim Wirtz, our excellent anonymous reviewers, and many students at the Naval Postgraduate School Center for Homeland Security and Defense (CHDS) for tremendously helpful comments and criticisms. We are especially grateful to the CHDS alumni network for responding so enthusiastically to our survey. All errors and omissions are our own.

"The liberties of a people never were, nor ever will be, secure, when the transactions of their rulers may be concealed from them."

—Patrick Henry, 1787

Does Patrick Henry's injunction—and similar admonitions by James Madison, Thomas Jefferson, and others—still ring true in an age when preventing terrorism has become a central duty of government? How open should government be when aggressive non-state actors seek to take advantage of information shared in the name of good governance or the public's right to know? Are certain kinds of information simply too dangerous to allow into the public realm? The consensus within much of the American government since 2001 has been that the threat of catastrophic terrorism demands less openness and increased attention to the potential costs of openly sharing government information.

In March 2002 then-Presidential Chief of Staff Andrew Card issued a memo to executive branch agencies instructing them to use Freedom of Information Act (FOIA) exemptions to withhold information whenever there was a legal basis to do so.[1] Card's memo reversed a long-standing bias toward openness, sending a strong signal that government should be more aggressive in keeping information private because of the threat of terrorism. This was hardly the first attempt by executive branch officials to reduce government openness. Policy makers have long struggled with the tension between the potential benefits of openness and the desire of government officials both to protect secrets and to shield the details of the policy-making process from public view.[2]

What Card's memo reflects is the fact that the potential for secrecy has become much greater since September 2001 as the scope of "national security" has expanded exponentially. In March 2003, Executive Order 13292 changed government classification standards to include "scientific, technological, or economic matters relating to the national security, which includes defense against transnational terrorism."[3] The addition of "transnational terrorism" was new, and the potential breadth of this redefinition becomes clear when considered alongside Homeland Security Presidential Directive Seven (HSPD-7). HSPD-7 states that "terrorists seek to destroy, incapacitate, or exploit critical infrastructure and key resources across the United States to

[1] Genevieve J. Knezo, *"Sensitive But Unclassified" Information and Other Controls: Policy and Options for Scientific and Technical Information* (Washington, DC: Congressional Research Service, 2006), 10.

[2] Daniel Patrick Moynihan and Larry Combest, *The Commission on Protecting and Reducing Government Secrecy* (Washington, DC: Government Printing Office, 1997).

[3] Knezo, *"Sensitive but Unclassified" Information and Other Controls*, 3.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

Downloaded By: [Princeton University] At: 14:30 14 April 2010

threaten national security . . . ."[4] Bounding the terms "critical infrastructure" and "key resources" is near impossible.

As we might expect given the vast discretion this definition of "national security" provides policy makers, there has been a dramatic decrease in the sharing of government information and of research funded by government.[5] Sensitive but Unclassified (SBU) information has become a new buzzword in American government, with over one hundred categories of SBU in use today. 67 percent of the 186 U.S. federal, state, local, and industry homeland security officials surveyed for this article report the use of SBU labels to control information has increased since 2000.[6] Even government officials find their ability to share information stymied by the proliferation of classification standards.[7] Meanwhile, directives suggest those charged with releasing information place renewed emphasis on not releasing information that might expose infrastructure, government, or the people to an increased risk of attack.[8] This movement away from information sharing generated resistance during the George W. Bush administration, but efforts to reestablish a norm of openness were generally stalled by ongoing security concerns.[9]

Despite its obvious normative, theoretical, and practical importance, the trade-off between government secrecy and openness has received scant attention in the political science literature. What scholarship there is has either

[4] "Homeland Security Presidential Directive Seven: Critical Infrastructure Identification, Prioritization, and Protection," Department of Homeland Security, 17 December 2003, available at http://www.dhs.gov/xabout/laws/gc_1214597989952.shtm

[5] Patricia McDermott and Emily Feldman, *Secrecy Report Card 2007: Indicators of Secrecy in the Federal Government*, OpenTheGovernment.org, available at http://www.openthegovernment.org/otg/SRC2007.pdf; and Rick Blum, *Secrecy Report Card 2005: Quantitative Indicators of Secrecy in the Federal Government*, OpenTheGovernment.org, available at http://www.openthegovernment.org/otg/SRC2005.pdf.

[6] Only five report it has decreased.

[7] The fixes embodied in the Information Sharing Environment (ISE) established under Executive Order 13388 focus on more effectively sharing information within government and pay short shrift to the value of sharing information with the public more generally. Thomas E. McNamara, "Information Sharing Environment Implementation Plan," Information Sharing Environment, accessed at http://www.ise.gov/docs/ise-impplan-200611.pdf, 24, 94, 8 March 2007; and Harold C. Relyea, *Security Classified and Controlled Information: History, Status, and Emerging Management Issues* (Washington, DC: Congressional Research Service. RL33494, 2006), 26–27.

[8] "FOIA Guide, 2004 Edition," Exemption One, accessed at http://www.usdoj.gov/oip/foi-act.htm, 7 April 2006.

[9] In the 109th Congress, for example, bills designed to strengthen the FOIA and to reestablish the principle that agencies should err on the side of openness died in committee in both the House and Senate. See Restore Open Government Act of 2005, HR 2331, 109th Cong., 1st sess., available at http://thomas.loc.gov/cgi-bin/bdquery/z?d109:h.r.02331; Open Government Act of 2005, S 394, 109th Cong., 1st sess., available at http://thomas.loc.gov/cgi-bin/bdquery/z?d109:s.00394. See also Relyea, *Security Classified and Controlled Information*; and John Cornyn, "Ensuring the Consent of the Governed: America's Commitment to Freedom of Information and Openness in Government," *LBJ Journal of Public Affairs* (Fall 2004). Three bills designed to encourage openness passed the House in the 110th Congress in July 2008—Reducing Over-Classification Act of 2007, HR 4806; Improving Public Access to Documents Act, HR 6193; and Reducing Information Control Designations Act, HR 6576—but were not passed into law.

focused on explaining why governments seek secrecy or on detailing problems this preference creates for researchers. David Gibbs and Francis Rourke follow the first path, developing theories about why governments seek to control information.[10] Rourke provides a thorough dissection of the causes of ongoing tension between secrecy and openness in democracies, but offers little to help policy makers find the right balance.[11] On the second path, Morton Halperin discusses the mechanics of FOIA implementation while Cyril Black focuses on how secrecy in government-sponsored research creates problems in planning future research.[12] There is little other work on this trade-off in the political science literature.[13] Dennis Thompson's analysis comes closest to ours.[14] He focuses on what he calls the basic dilemma of accountability: Democracy requires openness, but some policies require secrecy to be effective. His analysis touches implicitly on many of the issues we raise here, but where he deals with them from a normative standpoint, we focus on a more practical question: When will sharing information be a net benefit to society?

Unfortunately, the policy-oriented literature on information sharing does not make up for the thinness of the academic literature. These discussions tend to be highly polemical, focusing solely on either the costs of openness[15] or on the harm that is sure to come from excessive secrecy.[16] Even the more careful examinations of the issue often miss out on important aspects of the problem. An excellent RAND Corporation report on the security implications of publicly available geospatial (GIS) data, for example, effectively dissects potential malevolent uses of such information—for example, facilitating the targeting of GIS-guided munitions—and identifies a number of clear yes/no questions decision makers can answer to assess the information's potential

Downloaded By: [Princeton University] At: 14:30 14 April 2010

[10] David N. Gibbs, "Secrecy and International Relations," *Journal of Peace Research* 32, no. 2 (1995): 213-2-8; and Francis E. Rourke, "Secrecy in American Bureaucracy," *Political Science Quarterly* 72 (1957): 540–64.

[11] Francis E. Rourke, *Secrecy and Publicity: Dilemmas of Democracy* (Baltimore: The Johns Hopkins Press, 1961).

[12] Morton H. Halperin, "Freedom of Information and National Security," *Journal of Peace Research* 20, no. 1 (1983): 1-4; and Cyril E. Black, "Accessibility of Government-Sponsored Research in International Studies," *International Studies Quarterly* 14 (1970): 320–24.

[13] Searching article abstracts and titles in the Social Sciences Citation Index for the words "government" and "secrecy" returns only seven articles in academic political science journals between 1973 and 2007, none of which analyze the practical trade-off between openness and secrecy. Index available at http://scientific.thomsonreuters.com/cgi-bin/jrnlst/jlresults.cgi?PC=J&SC=UU.

[14] Dennis F. Thompson, "Democratic Secrecy," *Political Science Quarterly* 114, no. 2 (Summer 1999): 181-93.

[15] Alexander J. Breeding, *Sensitive But Unclassified Information: A Threat to Physical Security* (Bethesda, MD: SANS Institute, 2003); Evan M. Slavitt and Gregory D. Cote, "National Security vs. Public Disclosure: The War on Terrorism's Implications Upon Federal Emergency Planning and Right to Know Laws," National Security White Papers, The Federalist Society (December 2003).

[16] Steven Aftergood, "The Age of Missing Information," *Slate*, 27 March 2005, accessed at http://www.slate.com/id/2114963/ on 12 October 2009; and Lisa Graves, Senior Counsel for Legislative Strategy, American Civil Liberties Union, Testimony on S 394, the "OPEN Government Act," before the Terrorism, Technology, and Homeland Security Subcommittee of the United States Senate Committee of the Judiciary, on 15 March 2005.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

value to attackers.[17] While the report argues decision makers should consider GIS information's social value, it does not provide similarly clear questions decision makers can use to understand the positive uses or the social costs of keeping the same information private. The report thus offers incomplete guidance to policy makers charged with balancing the threat of harmful use against the potential gains from sharing geospatial data. In like fashion, Jacques Gansler and William Lucyshyn's otherwise-excellent overview of how to balance openness and security overlooks both the positive externalities of information sharing and the ways in which making information public can directly enhance counterterrorism efforts.[18]

The most thorough analysis of the issue has come in a number of reports to Congress by the Government Accountability Office (GAO) and the Congressional Research Service (CRS).[19] However, these organizations are not charged with examining the hard trade-offs implied by their detailed analyses. For example, John Moteff and Gina Steven's CRS analysis of the disclosure of information about critical infrastructure reveals a complex trade-off: Private firms are more likely to disclose information about vulnerabilities to government if they are offered liability protection; however, said liability protection makes them less likely to take actions to reduce the very vulnerabilities so recently revealed to government.[20] There is no discussion of how to resolve this dilemma, or how to think about other such trade-offs, in the CRS report or in the other government reports we surveyed. This is an understandable omission given the restricted scope of such analyses, but it means there is still a substantial gap in our understanding of the conditions under which government should share private information that is potentially useful to terrorists. As one federal official stated, "The balance between need-to-know

[17] John C. Baker et al., *Mapping the Risks: Assessing the Homeland Security Implications of Publicly Available Geo-Spatial Information* (Washington, DC: RAND Corporation, 2004).

[18] Jacques S. Gansler and William Lucyshyn, *The Unintended Audience: Balancing Openness and Secrecy: Crafting an Information Policy for the 21st Century* (College Park, MD: Center for Public Policy and Private Enterprise School of Public Policy, University of Maryland, 2004).

[19] Government Accountability Office, *Managing Sensitive Information: DOE and DOD Could Improve Their Policies and Oversight* (Washington, DC, 2006); Government Accountability Office, *Transportation Security Administration: Clear Policies and Oversight Needed for Designation of Sensitive Security Information* (Washington, DC, 2005); Government Accountability Office, *Critical Infrastructure Protection: Improving Information Sharing with Infrastructure Sectors* (Washington, DC, 2004); Government Accountability Office, *Information Sharing: Practices That Can Benefit Critical Infrastructure Protection* (Washington, DC, 2001); John D. Moteff and Gina Marie Stevens, *Critical Infrastructure Information Disclosure and Homeland Security* (Washington, DC: Congressional Research Service, RL31547, 2003); John D. Moteff, *Risk Management and Critical Infrastructure Protection: Assessing, Integrating, and Managing Threats, Vulnerabilities and Consequences* (Washington, DC: Congressional Research Service. RL32561, 2005); and Knezo, "Sensitive but Unclassified" Information and Other Controls.

[20] Moteff and Stevens, *Critical Infrastructure Information Disclosure and Homeland Security.*

and protection of information is extremely difficult and has not received adequate attention from policy makers."[21]

Addressing this problem requires a proper delineation of the types of information in play. We can identify three distinct kinds of information government may want to keep secret due to the threat of terrorism:

1. Organizational information that helps government better predict terrorist operating patterns, but can also help terrorists identify their own operational vulnerabilities. A good example of organizational information is the collection of more than a million documents captured during operations in Afghanistan, Iraq, and elsewhere held in the United States Department of Defense's Harmony Database. In theory, terrorists could analyze these documents to identify problems in their own organizations. In practice, a substantial number of these documents were declassified and used to identify a set of vulnerabilities inherent in covert organizations.[22]
2. General information that contributes to scientific research, to good governance, or to corporate accountability, but can also be analyzed to identify unknown vulnerabilities. Data on fiber-optic communication networks clearly fit in this category. In 2003, George Mason University doctoral student Sean Gorman used public information to map the fiber-optic network of the United States, identifying critical choke points in the country's telecommunications infrastructure.[23] While the specific sites Gorman identified would presumably be of great interest to terrorists, his analysis implicitly identified cost-effective ways to remove potentially crippling interdependencies between firms, making society more resilient against all manner of disasters.
3. Target-specific information that can help society better protect potential targets, but that reveals known vulnerabilities. Revealing data on the vulnerabilities of certain kinds of infrastructure, for example, can be a net benefit when the target would be inadequately defended absent that revelation. A series of GAO reports about weaknesses in defensive measures

[21] This quote and subsequent quotes are from a survey of 186 U.S. federal, state, local, and industry homeland security officials. We identify sources by respondent number to maintain anonymity. The survey is described in the section "Secrecy and Openness in Practice." Respondent 725732258.

[22] Joseph H. Felter et al., *Harmony and Disharmony: Exploiting al-Qa'ida's Organizational Vulnerabilities* (West Point: Combating Terrorism Center, 2006); Jacob N. Shapiro and Clinton Watts, ed., Al-Qa'ida's (Mis)Adventures in the Horn of Africa (West Point: Combating Terrorism Center, 2007); Brian Fishman et al., Sinjar Two: al-Qa'ida in Iraq's Foreign Recruiting, Finances, and Future (West Point: Combating Terrorism Center, 2008); and Brian Fishman, Dysfunction and Decline: Lessons Learned from Inside al-Qa'ida in Iraq (West Point: Combating Terrorism Center, 2009).

[23] Sean Gorman, "Networks, Complexity, and Security: The Role of Public Policy in Critical Infrastructure Protection" (PhD diss., George Mason University, 2004).

Downloaded By: [Princeton University] At: 14:30 14 April 2010

at commercial nuclear power plants, for example, played a key role in overcoming industry resistance to stricter security standards.[24]

The trade-offs involved in releasing organizational information are clear without any formal analysis, so we do not discuss them at length. Government officials should release organizational information whenever society is more effective than terrorists at utilizing it. This will generally be the case as there is a massive cognitive apparatus seeking to identify terrorist vulnerabilities, while the terrorists of greatest concern today rely on a relatively small coterie of public intellectuals publishing occasional articles on the internet.[25] As Bendor shows, multiple analysts working in parallel are strictly better than individual experts at a wide range of tasks, so long as they are sufficiently independent.[26] A similar logic stands behind arguments about the superior security of open source software, whose advocates claim that having multiple independent users evaluating security flaws—and other bugs—leads to vulnerabilities being identified and fixed faster than in traditional software, often before malign hackers can identify them.[27] Making organizational information public allows think tanks and academics to operate as just such independent elements in the government's cognitive apparatus.

To understand the specific conditions under which sharing the other types of information can make society better off, we explore the logic behind a series of strategic interactions.[28] This exercise is not intended to produce counterintuitive results. Instead, it forces us to consider explicitly all the trade-offs inherent in the problem. By identifying the conditions under which an idealized government seeking to maximize social welfare shares

Downloaded By: [Princeton University] At: 14:30 14 April 2010

[24] Government Accountability Office, *Nuclear Regulatory Commission Oversight at Commercial Nuclear Power Plants Needs to Be Strengthened* (Washington, DC, 2003); Government Accountability Office, *Preliminary Observations on Efforts to Improve Security at Nuclear Power Plants* (Washington, DC, 2004); Government Accountability Office, *Nuclear Power Plants Efforts Made to Upgrade Security, but the Nuclear Regulatory Commission's Design Basis Threat Process Should Be Improved* (Washington, DC, 2006); and Mark Holt and Anthony Andrews. Nuclear Power Plant Security and Vulnerabilities (Washington, DC: Congressional Research Service, RL34331, 2008).

[25] For a citation-mapping of jihadi public intellectuals, see William McCants and Jarret Brachman, *Militant Ideology Atlas* (West Point, NY: Combatting Terrorism Center, 2006). The most well-known of these public intellectuals is Abu Musab al-Suri. Brynjar Lia and Thomas Hegghammer "Jihadi Strategic Studies: The Alleged Al Qaida Policy Study Preceding the Madrid Bombings," *Studies in Conflict and Terrorism* 27, no. 5 (September/October 2004): 355–75, analyze his 2003 policy paper which some believe inspired the Madrid bombing.

[26] Jon Bendor, *Parallel Systems: Redundancy in Government* (Berkeley and Los Angeles: University of California Press, 1985).

[27] Analysts of encryption software have long argued that making the source code publicly available leads to more secure software because bugs are remedied which original developers may have missed. Scott A. Hissam, Daniel Plakosh, and Charles B. Weinstock, "Trust and Vulnerability in Open Source Software," *IEE Proceedings–Software* 149, no. 1 (2002): 47-51; Steve Weber, *The Success of Open Source*. (Cambridge, Harvard University Press, 2004); and David A. Wheeler, "Why Open Source Software/Free Software (OSS/FS, FLOSS, or FOSS)? Look at the Numbers!" accessed at http://www.dwheeler.com/oss_fs_why.html (2007), 31 January 2008.

[28] In an online appendix we provide a formal game-theoretic analysis of these interactions—available at http://myweb.fsu.edu/dsiegel/Research.

information in equilibrium, we implicitly identify the conditions under which releasing information of each type is likely to benefit society.

We begin by examining the strategic decision to share general information by studying the situation in which both government and terrorists are simultaneously searching for vulnerabilities—government to defend them and terrorists to attack them.[29] The degree to which potentially relevant general information is released alters the rates at which each actor discovers vulnerabilities, as when Gorman used publicly available information to identify previously unknown weaknesses in the U.S. communications infrastructure.[30] We then focus on target-specific information by building on Robert Powell to study the situation where the aforementioned search process has already occurred and government must release target-specific information about and allocate defensive resources—such as police officers or other security personnel—among vulnerabilities that it has discovered.[31]

As a general point, we find that when government is better able to make use of the resources of the larger society than the terrorists are able to make use of any information released, releasing information is often the preferred strategy. In each interaction there is a range of scenarios within which government should release information, even when we only consider the benefits to counterterrorism.[32] Releasing information becomes even more beneficial when we also consider the positive externalities to openness. Our results suggest a more nuanced approach than that typified by recent government actions, one in which government officials seek explicitly to balance the potential benefits to openness against the costs of information sharing. If openness should be favored from the narrow perspective of maximizing counterterrorism, then arguments about the public's right to know, or about core democratic principles, should carry even greater weight in policy discussions.

This paper proceeds as follows. The first section, "Types of Information and the Traditional View," examines the different types of information in detail and illustrates why the typical government approach to the problem is incorrect. The second section, "Strategic Information Release," analyzes the strategic interactions involved in releasing each type of information. A formal presentation of the analysis, with proofs, can be found in an online appendix.[33] The third section, "Secrecy and Openness in Practice," examines a survey of 186 U.S. homeland security officials to understand better how they actually make decisions about information sharing. The fourth section,

Downloaded By: [Princeton University] At: 14:30 14 April 2010

[29] We use target and vulnerability interchangeably in this paper. Strictly speaking, targets are physical entities while vulnerabilities are gaps in defense that permit a successful attack on such entities. We use the words interchangeably, as an entity not vulnerable to attack is not really a target.

[30] Sean Gorman, *Networks, Complexity, and Security.*

[31] Robert Powell, "Defending Against Terrorist Attacks with Limited Resources," *American Political Science Review* 101, no. 3 (August 2007): 527–41.

[32] Formally, there is a range of parameter values over which releasing information is an equilibrium action.

[33] The appendix is available at http://myweb.fsu.edu/dsiegel/Research.

"Policy Implications," concludes with recommendations for officials involved in setting information-sharing policy.

TYPES OF INFORMATION AND THE TRADITIONAL VIEW

The traditional approach to sharing private information in homeland security and counterterrorism has been to imagine the problem as one in which government knows the true value of targets, but the terrorist does not. The terrorist has to choose where to attack to maximize its expected benefit from the attack. Typically, optimal strategic behavior in this setting involves both sides' randomizing over targets. Terrorists must randomize so that government cannot guess which target they will attack, and if defensive resources are limited, governments must randomize so that terrorists cannot guess which potential targets are less well defended. When terrorists do not know the targets' true value, they will randomize incorrectly. Since suboptimal randomization by terrorists helps the government, it is generally viewed as unfavorable to release information, absent some call to the public's right to know.

Looking at the problem this way elides a number of important features. First, it deals only with information that pertains to specific targets, leaving out the question of when government should reveal more general information that might help identify unknown vulnerabilities before terrorists can discover them. Second, it cannot account for the incomplete nature of both sides' knowledge. In reality, neither terrorists nor government know all potential vulnerabilities and all possible targets. Third, it does not take into account the fact that some information can pay dividends across multiple targets, either to attacker or to defender. Finally, it misses the possibility of economies of scale to defense or of increases in the efficiency of resource use with the release of information.

These debates are not unknown in government. As one state homeland security official observed, "At the operations level there is always controversy between those who are sure that 'we will give terrorists ideas', as though the terrorists are incapable of thinking of them on their own, and those who think that if we have some ideas of our vulnerabilities that terrorists probably do also, therefore we should use knowledge of those vulnerabilities to reduce them."[34] That these debates are so often carried out in binary terms is evidence that developing a richer understanding of the problem is valuable. In order to do so the following subsections define the three key types of private information government can share, discuss how each can help terrorists attack society and outline how sharing them may enhance defensive efforts.

[34] Respondent 724229888.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

Downloaded By: [Princeton University] At: 14:30 14 April 2010

Organizational Information

Organizational information is that which helps government better predict terrorist operating patterns, but also can help terrorists identify their own operational vulnerabilities. Examples include the kinds of sensitive materials traditionally classified because sharing them might reveal what government knows about terrorists, or might compromise intelligence sources and methods, thereby reducing the future flow of intelligence. The negative ramifications of sharing such information are well understood with respect to traditional state-level adversaries. What is poorly understood is that releasing this kind of information has different implications for fighting non-state actors.

Successful counterterrorism must take into account the great variability of terrorist tactics; it is not enough to prepare the best response to the other side's average operating patterns. Against a terrorist threat, government must get it right for each cell. When the other side is as poorly understood as are terrorists, information sharing can help in two ways. First, local law enforcement officials who lack clearances and the systems for processing classified information are the front line in counterterrorism. The more they understand about the threat, the better they can do their jobs. While selectively sharing information with law enforcement officials is attractive in principal, continuing dissatisfaction with information sharing between levels of the U.S. government suggests parceling out information on a need-to-know basis is inherently problematic.[35] Second, sharing organizational information allows outside researchers—and government officials without security clearances—to contribute to a better understanding of terrorist organizations. The resulting benefits may well outweigh the costs of letting the other side know what government knows.

The best example of organizational information would be a dataset of instances in which government successfully broke up a terrorist cell or attack plan, leading to a failed attack. Releasing this particular type of organizational information is problematic in that terrorist groups only observe their own failures, not those of like-minded groups. Thus, making failure data public could greatly increase groups' understanding of the causes of failure, and hence enhance their operational capabilities. However, such a dataset would be invaluable to researchers. While analysts today can study patterns of terrorist success and government failure, there is no way to study government success and terrorist failure quantitatively. This lacuna means it is impossible to distinguish between the case when a low level of terrorist activity is driven by successful government counterterrorism, and when it

[35] In February 2009, more than seven years into efforts to enhance information flow, local, state, and federal officials we surveyed continued to complain about inadequate sharing within the law enforcement and first-responder communities. Respondents 724215968, 724215641, 724214470, and 724212480.

is driven by other factors such as political decisions by terrorist leaders or by groups' internal organizational dynamics. Without information on total terrorist attempts, there is simply no way to study rigorously the strategic interactions between governments and terrorists.

Similar issues arise with detailed case studies of terrorist financial practices. Sanitized case studies of terrorist financial transactions are provided in annual typology reports published by the Financial Action Task Force on Money Laundering (FATF) and in documents published by the Egmont Group, a counter-money laundering organization. Unfortunately, because these case studies are sanitized they cannot be linked to specific groups at specific times. This limits what terrorist financial agents can learn from them, but also means analysts are unable to determine how groups' political and operational environments influence their financial structure. Such information would be hugely useful for those charged with evaluating efforts to counter terrorist financing, efforts that have imposed huge costs on people in the poorest parts of the world.[36] Unfortunately, no such analysis is possible when publishing organizations feel offering specificity would tell terrorists too much about their own vulnerabilities.[37]

General Information

General information is that which can be analyzed to identify vulnerabilities unknown to government. Here we can think of data detailing the average loads on various electrical transmission lines, data that can be used to identify lines whose removal would lead to rolling blackouts.[38] Alternately, we can think of the kinds of publicly available information about milk

Downloaded By: [Princeton University] At: 14:30 14 April 2010

[36] Khalid M. Medani, "Financing Terrorism or Survival? Informal Finance and State Collapse in Somalia, and the U.S. War on Terrorism," *Middle East Report* 32, no. 2 (Summer 2002): 2–9.

[37] Releasing information can help correct misunderstandings that hinder counterterrorism efforts. For example, the frequent refrain that terrorists are irrational and motivated by religious fervor hinders government response. Law enforcement officials believing this line are less likely to focus attention on the secular middle-class individuals who make up the majority of terrorist operatives in some groups. See Alan B. Kreuger and Jitka Malečkova, "Education, Poverty and Terrorism: Is There a Causal Connection?" *The Journal of Economic Perspectives* 17, no. 4 (Autumn 2003): 119–44; Claude Berrebi, "Evidence About the Link Between Education, Poverty and Terrorism Among Palestinians," *Peace Economics, Peace Science and Public Policy* 13, no. 1 (2007), available at http://www.bepress.com/peps/vol13/iss1/2; Marc Sageman, *Understanding Terror Networks* (Philadelphia: University of Pennsylvania Press, 2004); Christine Fair and Bryan Shepherd, "Research Note: Who Supports Terrorism? Insights from Fourteen Muslim Countries," *Case Studies in Conflict and Terrorism* 29, no. 2 (January/February 2006); and C. Christine Fair, "Militant Recruitment in Pakistan: A New Look at the Militancy-Madrasah Connection," *Asia Policy* 1, no. 4 (Summer 2007).

[38] In August 2003 the failure of one power line south of Cleveland, Ohio, led to the cascading failure of the regional power infrastructure, affecting fifty million people and causing four to ten billion dollars in economic losses in the United States alone. Eric J. Lerner, "What's Wrong with the Electric Grid," American Institute of Physics, accessed at http://www.aip.org/tip/INPHFA/vol-9/iss-5/p8.html, 2003, 7 April 2006; and *Electric Power Annual*, 2003," Electric Information Agency, accessed at http://tonto.eia.doe.gov/FTPROOT/electricity/034803.pdf, 7 April 2006.

pasteurization that fed Lawrence Wein and Yifan Liu's analysis of vulnerabilities in the dairy supply chain.[39] Their analysis identified critical, relatively easily remedied vulnerabilities to bioterrorism and was presented to government and industry officials months before publication.[40] Despite this, officials at the Department of Health and Human Services (HHS) unsuccessfully attempted to block publication of the article on the grounds that it would provide too much information to potential attackers. This objection only makes sense if government knew the dairy industry had not addressed the vulnerabilities Wein and Liu identified. HHS' actions thus lend credence to the notion that sharing information can force industry to internalize more fully the costs of attacks on privately owned infrastructure.

An example of how sharing general information can—and did—help protective efforts is Wein and Manas Baveja's analysis of the U.S. Visitor and Immigrant Status Indicator Technology (US-VISIT), a fingerprint identification program designed to prevent known terrorists from entering the United States through legal channels.[41] The authors analyzed the system using performance information on fingerprint readers from the website of the National Institute of Standards and Technology (NIST), finding that terrorists could take advantage of the original system by reducing their fingerprint image quality. Requiring subjects to present more fingers if their prints were poor, however, would improve the system's performance even when the terrorists knew exactly how it worked. US-VISIT is adopting a version of this system. Importantly, this analysis would not have been possible if government had kept a tight hold on information about the characteristics of the biometric identification systems, information that could have helped terrorists before the fix was identified.

Existing discussions of when general information should be shared tend to focus on how many cognitive steps terrorists must make to take advantage of it. The argument is that the fewer knowledgeable personnel needed for terrorists to exploit information, the more dangerous it is to release.[42] Implicit in this argument is a belief that if disclosure is dangerous, the information should not be shared. But this is only one half of the equation. Much scientific information—especially biological information—cannot be readily divided

Downloaded By: [Princeton University] At: 14:30 14 April 2010

[39] Lawrence M. Wein and Yifan Liu, "Analyzing a Bioterror Attack on the Food Supply: The Case of Botulinum Toxin in Milk," *Proceedings of the National Academy of Sciences* 102, no. 28 (July 2005): 9,985–89.

[40] Rick Weiss, "Report Warns of Threat to Milk Supply: Release of Study Citing Vulnerability to Bioterrorism Attack Was Opposed by U.S. Officials," *Washington Post*, 29 June 2005, A08.

[41] Lawrence M. Wein and Manas Baveja, "Using Fingerprint Image Quality to Improve the Identification Performance of the U.S. Visitor and Immigrant Status Indicator Technology Program," *Proceedings of the National Academy of Sciences* 102, no. 21 (May 2005): 7,772–75.

[42] Gansler and Lucyshyn, *The Unintended Audience.*

Downloaded By: [Princeton University] At: 14:30 14 April 2010

up into safe and dangerous components.[43] That which is potentially useful is often potentially dangerous.[44]

The proper question is whether sharing general information contributes more to the expected losses from terrorist attacks than to homeland security, scientific research, good governance, and corporate accountability. As we will see, the answer to this question depends, in part, on both sides' analytical capabilities. If those on the side of government can use the information to identify and address vulnerabilities faster and more effectively than terrorists, then releasing information of this type is very likely to be to society's advantage.

Target-Specific Information

Target-specific information is that which reveals vulnerabilities known to government, ones of which terrorists may not be aware. Here we can think of information that gives the location of concentrations of hazardous chemicals or that identifies critical infrastructure nodes. Officials usually consider keeping this type of information secret an unalloyed good. However, this type of information can help government and private industry better protect potential targets, develop redundant systems, and recover from attack.[45]

Sharing information about known vulnerabilities can also lead to better allocations of homeland security resources by spurring industry to self-interested action to protect critical infrastructure. Spending money on rare events reduces profits. Thus, when the full social costs of such events exceed the private costs, industry officials fulfilling their fiduciary duty to maximize profits will spend less on protection than is socially optimal. Suppose, for example, that terrorists attack a chemical plant, releasing toxic gases. The owner can expect to pay for repairs to the plant, but is unlikely to be held liable for the full costs of the event, which could include long-term medical care, disability benefits, clean-up costs, and losses from reduced economic

[43] Stanley Falkow et al., *Seeking Security: Pathogens, Open Access, and Genome Databases* (New York: National Academies Press, 2004).

[44] For information on how genetic information with tremendous therapeutic potential also contains the necessary information for crafting racially targeted biological weapons, see Tonya Putnam, "Racial Weapons: An Essay on the Law, Ethics, and Politics of Biological Warfare in the Age of the Genome," paper presented at the annual meeting of the International Studies Association (San Diego, CA, March 2006).

[45] Whether it does so depends on the nature of the vulnerability identified. Making information on the chemicals used in industrial facilities readily available can clearly help first responders. Making public information about which electrical lines to attack to cripple the power grid may seem less helpful. Such information might not help local law enforcement much because there are many points of attack for any particular line, but could help terrorists a great deal. However, what sharing such information might do is spur power companies to self-interested action to protect the grid, perhaps by building redundant capacity. Government Accountability Office, *Information Sharing: Practices That Can Benefit Critical Infrastructure Protection.*

Downloaded By: [Princeton University] At: 14:30 14 April 2010

activity in the surrounding area. As one state homeland security official put it, "It is difficult to reduce a vulnerability if you cannot identify it for fear that it will come to the attention of terrorists."[46]

The tension between public and private costs is exacerbated under three conditions: (1) when industry has incomplete information about threats and vulnerabilities;[47] (2) when the public lacks the information necessary to force industry to internalize the social risks of their activities; and (3) when government regulators cannot make public their assessments, which would create political pressure for industry to cover the gap.[48] Sharing information eases all three.

Moreover, information about known vulnerabilities is critically important for making high-level choices about who should take responsibility for making resource-allocation decisions for homeland security. For example, Rudy Darken and Ted Lewis note that specific pieces of certain states' critical infrastructure reside outside of those states (Alaska's main telephone exchange is in Seattle) or are completely dependent on supplies coming over out-of-state infrastructure (the largest power plant in Missouri relies on coal coming by rail from Wyoming).[49] This pattern means that the current policy of pushing resource allocation decisions down to the state level will not adequately protect key pieces of critical infrastructure. Their analysis would not have been possible without knowledge about existing vulnerabilities, without target-specific information.

The likely effects of sharing target-specific information depend on who is responsible for protecting a particular target. We can make a distinction between three types of targets: (1) privately controlled targets, (2) targets controlled by state or local governments, and (3) targets controlled by the federal government. Of course, each type of target may be critically dependent on resources that must be protected by another party.[50] Where a target in one jurisdiction is critically dependent on resources that must be protected by another party, openness increases the likelihood that appropriate resources will be allocated to defense. Seattle officials, for example, did not understand the importance of their telephone exchange until the interdependency was brought to their attention by outside analysts.[51]

For privately controlled targets, target-specific information can lead to better protective measures in two ways. First, it may help identify critical

[46] Respondent 724229888.

[47] Stephen E. Flynn and Daniel B. Prieto, *Neglected Defense: Mobilizing the Private Sector to Support Homeland Security* (New York: Council on Foreign Relations Press, 2006), 15–16.

[48] Congressional Budget Office (CBO), *Homeland Security and the Private Sector* (Washington, DC: The Congress of the United States, 2004), 2-4.

[49] Rudy Darken and Ted G. Lewis, "Potholes and Detours in the Road to Critical Infrastructure Protection Policy," *Homeland Security Affairs* 1, no. 2 (Fall 2005).

[50] In addition to the previous examples, we can think of New York City's tragic dependence on airline passenger screeners in Boston.

[51] Ted G. Lewis, interview by Jacob Shapiro, 5 January 2007.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

interdependencies. Second, it can create public pressure for actions to mitigate the costs of an attack. Many critical substations in the nation's electrical power grid rely on customized transformers and other components.[52] Electrical companies have not historically kept spares on hand for these components because they rarely fail during normal operations. The problem is that many substations are in isolated areas, are thus vulnerable to attack, and replacing these large transformers can take several months.[53] Although this vulnerability has been inherent in the system for some time, it was only after the publication of many articles highlighting it that the power industry—and the Department of Homeland Security—began getting serious about stockpiling spare components.[54]

At first glance, it seems like the need for information sharing about specific targets should be less for state-controlled targets. Many states have been extremely active about identifying both targets and dependencies that need to be taken into account when planning protective measures.[55] However, this viewpoint misses out on the fact that sharing target-specific information can be extremely useful for guiding protective actions by states. Much has been written since 2003 on the importance of allocating protective resources on the basis of risk.[56] The most widely used means of doing so by organizations with a significant financial stake in the matter—reinsurance companies—is to use an algorithm that relies critically on expert opinion about target selection by terrorists, the likely impact of different attack modes, and the likelihood of attack.[57] Notice that target-specific information is critical for assessing the

[52] Leonard Anderson, "Power Grid Aims for Backup Supplies," *Reuters*, 2 June 2005.

[53] CBO, *Homeland Security and the Private Sector.*

[54] Public knowledge or heightened attention sometimes leads to temporary protective measures that last until private interests reassert themselves once the issue becomes less politically salient. Following a January 2004 train crash in South Carolina that released toxic gases killing nine people, Washington D.C. enacted a ban on the transportation of deadly chemicals along rail corridors passing close to the national capital area. See Philip Auerswald et al., "The Challenge of Protecting Critical Infrastructure," *Issues in Science and Technology* 22, no. 1 (Fall 2005), accessed at http://www.issues.org/22.1/auerswald.html, 6 April 2006. After several months of legal fights, the ban was overturned in court, although the city government continues to support a ban. See Carol D. Leonnig, "Judge Upholds Hazmat Rail Ban," *Washington Post*, 19 April 2005, B02; and Carol D. Leonnig, "Judge Demands To View Rail Plan D.C. Hazmat Cargo Ban At Issue in CSX Lawsuit," *Washington Post*, 22 September 2005, B01. On the electrical grid, see Stephen E. Flynn, *America–Still Unprepared, Still in Danger*, report of an Independent Task Force Sponsored by the Council on Foreign Relations (2002); Antonio Regalado and Gary Fields, "Blackout a Reminder of Grid's Vulnerability to Terror," *Wall Street Journal*, 15 August 2003, A4; R. James Woolsey and Rachel Belton, "We Must Face a Connected World's 'Butterfly Effect'," *Los Angeles Times*, 5 May 2004; Amy Abel, *Government Activities to Protect the Electrical Grid,* (Washington, DC: Congressional Research Service, RS21958, 2004); and Stephen E. Flynn, private communication with Jacob Shapiro, 20 October 2005.

[55] See for example the detailed process for identifying critical infrastructure dependencies outlined in Washington State, *Homeland Security Region 6 Critical Infrastructure Protection Plan*, 2005, accessed at http://www.metrokc.gov/prepare/docs/Region6CIP, 10 April 2006.

[56] For a summary, see John D. Moteff, *Risk Management and Critical Infrastructure Protection.*

[57] Henry H. Willis et al., *Estimating Terrorism Risk* (Santa Monica, CA: RAND Corporation, 2005); and Gordon Woo, private communication with Jacob Shapiro, 20 January 2006.

latter two. Sharing target-specific information can thus enhance states' ability to make judgments about the appropriate protective allocations, thereby enhancing the efficiency with which they use scarce resources.

For federally controlled targets, sharing target-specific information is less important. The federal government has many methods to share sensitive information internally, and senior decision makers can often generate the protective actions they desire. As such, sharing target-specific information is less likely to yield additional benefits. Where substantial benefits can accrue on federally controlled targets is in sharing general information that can be analyzed to identify vulnerabilities and that also can be used to enhance government response.

Our discussion so far highlights a series of trade-offs. None are simple to quantify and some are inherently unknowable. How would one assess the future value of a particular piece of scientific information when it is openly shared versus when it is closely held? What the above discussion illustrates is the necessity of taking multiple factors into account. The next section develops a more subtle understanding of when the net benefit of information sharing will be positive.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

STRATEGIC INFORMATION RELEASE

One strength of formal analysis is its ability to make explicit the trade-offs that decision makers should consider in a particular strategic interaction. In this paper, that interaction is straightforward: In addition to allocating resources to defend potential targets, government can release general or target-specific information. Doing so can improve the terrorists' strategic circumstances by illuminating previously unknown attack modes and by increasing the likelihood that an attack will succeed. Conventional wisdom since 9/11 has largely focused on these negatives, and thus access to information has been increasingly restricted. Releasing information, however, may also improve government's strategic situation. These positive effects of information-sharing have been undervalued in policy making and are rarely mentioned in analyses of the increasing use of the SBU label.

This section works through the logic of a series of strategic interactions to highlight important trade-offs between secrecy and openness. Our results provide a fundamentally prescriptive account of what government should do, rather than a descriptive or positive account of what it does do. By describing the range of settings within which an idealized government seeking to maximize social welfare should share information, we implicitly identify the conditions under which real government officials should do the same. As our models' collective focus is to illuminate these conditions, we concentrate

Downloaded By: [Princeton University] At: 14:30 14 April 2010

here on the models' implications, leaving their full mathematical treatment to the online appendix.[58]

Consider the strategic interaction between two players: government (G) and terrorists (T). T seeks to discover and attack the target that would have the greatest expected impact, taking into account its likelihood of success. G seeks to discover unknown vulnerabilities and address them before T can find and exploit them. G must balance the potential defensive gains from releasing information against the risk of alerting T to vulnerabilities or increasing T's chances of success against known targets.

For simplicity, we will say that G releases information and uses it to alter the likelihood of a successful attack via subsequent analysis and coordination. Substantively, there are two separate sets of actors treated under the umbrella of G. The first consists of those within government that have prior access to the relevant information and must make decisions on defensive allocations and information release. The second consists of those within the rest of the government and the larger society who lack prior access to the relevant information, but may use it in ways that improve government's defensive efficiency. Treating the interaction between these sets of actors as a black box allows us to stay narrowly focused on the strategic dynamics of information release.

Given this focus, we elide two important issues. The first is the potential for agency problems between G and the society it protects. We assume throughout that G has society's best interests at heart and so seeks to maximize expected social welfare. This assumption clearly does not hold for all members of government. For example, releasing information about a target that is subsequently attacked might be expected to create political backlash, motivating officials not to authorize the release even if the expected value to society writ large is purely positive.[59] While a full consideration of such agency problems would take us far afield, we do cover the impact of negative (or positive) backlash.

The second is the fact that if the government can release information to a small subset of society, effectively increasing the number of analysts while minimizing the risk of malign users seeing the information, then the trade-off between secrecy and openness is obviously less stark. The Bureau of Economic Analysis (BEA), for example, routinely makes confidential microdata available to academic researchers through the Special Sworn Employee program.[60] Including this option for G would not change our core arguments; when we discuss the possibility that releasing information can help

[58] Appendix available at http://myweb.fsu.edu/dsiegel/Research.

[59] Releasing such information might also allow for more effective evacuation and management plans, reducing post-attack casualties and thereby counteracting backlash.

[60] The Census Bureau similarly makes confidential census data available to outside researchers through the Center for Economic Studies.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

make defensive resources more efficient, for example, we are considering improvements relative to the baseline of no public release. One could, in principle, break this decision down first to the relative gains of releasing to a select few, and then to those arising from more public release. Our analysis still describes relevant trade-offs—any new possessor of such information entails additional risk, even if it is small—and so remains relevant to policy even when the possibility of bringing some experts into the government fold exists.

SELECTIVE INFORMATION RELEASE

As a practical matter, whether selectively sharing information makes sense for a particular system depends on the characteristics of that system. The appropriate question to consider is whether the marginal gains of moving from selective sharing to full openness outweigh the increased risks from doing so. For selective information sharing to capture most of the gains from full openness, there must be a way to match appropriately skilled experts with the relevant information. This matching process can happen in two ways. In the bottom-up approach, outside analysts propose projects that can use confidential government data to make progress on important problems. This is the model followed by the Census Bureau and the BEA. Notice that this approach requires that outside researchers be able to identify what information the government holds. Because the BEA and Census post their survey forms and sampling methodology online, researchers are able to learn exactly what variables have been collected on which entities, and thus can plan research and make proposals without ever seeing the micro-data that must be kept confidential.[61] No similar ability for outside analysts to know what the government knows exists in the counterterrorism realm, meaning bottom-up matching is unlikely to work well.

In the top-down approach, government officials reach out to particular analysts. Think here of government officials hiring outside experts to study a subject that is poorly understood within their part of the government.[62] For this method to work two conditions must be obtained: (1) government officials must understand the system well enough to identify the relevant experts, and (2) they must know what information these experts need. Whether or not these conditions are obtained depends in part on the type of information being released. For general information, there is substantial evidence suggesting neither condition is likely to be obtained in the poorly understood

[61] Dennis Fixler, Chief Statistician, Bureau of Economic Analysis, phone interview with Jacob Shapiro, 1 February 2008.

[62] The first draft of the White House report on Hurricane Katrina, for example, was written by a contractor.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

domain of counterterrorism, which includes protecting wide-ranging "critical infrastructure" and "key resources."[63] On the first condition, there is a great deal of confusion over what properly constitutes critical infrastructure, meaning there must be a great deal of uncertainty over who the relevant experts are.[64] On the second condition, officials within the federal government still have not identified the information they must share to protect critical infrastructure, a limited subset of the targets that must be protected from terrorists.[65]

In contrast, the conditions for selective release are likely to be obtained for target-specific information. Having already identified the vulnerability, government can release information directly to the target and potentially to already-chosen subject matter experts. These experts can then request further details, reducing the need for government to discern what is needed beforehand.

Even in the case of target-specific information, however, there are limitations to what selective release can accomplish. First, the problem of divergent preference between government and target continues to hold, particularly if the target is in the private sector. Unless the owner of the vulnerable site possesses, or can be made to possess, preferences in line with the social optimum, it will have cause to under-provide counterterrorism and/or post-terror damage abatement. Regulation is one possible solution, but one that has had limited success due to both noncompliance and agency capture.[66]

The history of the design basis threat (DBT) for commercial nuclear reactors is instructive here. The DBT specifies the characteristics of the enemy that security forces at nuclear power plants must be able to defeat. After the September 11 attacks, the Nuclear Regulatory Commission (NRC) conducted a review of the DBT for commercial power plants, and in April 2003 it proposed new rules that "increased the number of attackers, refined and expanded the list of weapons and equipment that might be used in an attack, and increased the maximum size of a vehicle bomb that plants must defend

[63] Neither condition was obtained, for example, in the case of the US-VISIT program whose managers had never considered that the system could be gamed in the ways Wein and Baveja identified and so never passed information on the system to outside analysts with the necessary expertise.

[64] Debates over the utility of the National Asset Database reveal that uncertainties identified in John D. Moteff and Paul Parfomak, *Critical Infrastructure and Key Assets: Definition and Identification* (Washington, DC: Congressional Research Service, RL32631, 2004) about criticality have not been remedied as of 2006, as per John D. Moteff, *Critical Infrastructure: The National Asset Database* (Washington, DC: Congressional Research Service, RL33648, 2006).

[65] Defense Science Board, *Report of the Defense Science Board Task Force on Critical Homeland Infrastructure Protection* (Washington, DC: Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics, 2007).

[66] This argument is no different than that frequently applied to environmental degradation. Billions of dollars in cleanup and damages by Exxon after the *Valdez* spill may have led to increased safety measures in shipping, but these measures were not taken before the spill despite regulation, and twenty years later damage from the spill lingers.

against."[67] In developing the new DBT, the NRC excluded certain weapons that intelligence assessments considered to be a threat to nuclear reactors because ". . . industry considered these adversary characteristics prohibitively expensive to defend against."[68] The key takeaway from this episode is not that the current DBT is inadequate; an analysis of that would be well beyond the scope of this paper. Rather, the point simply is that there is often tension between what government officials want to spent on security and what industry is willing to pay. When such tension exists, public exposure of vulnerabilities can create pressure for greater protective investments than private firms would make on their own.[69]

Second, once government has matched experts to information, officials must be able to set up secure, effective information-sharing systems. Here U.S. government efforts have largely failed, at least with respect to protecting critical infrastructure.[70] On 27 August 2004, President George W. Bush signed Executive Order 13356, which directed agencies to place a high priority on exchanging information relevant to preventing terrorist attacks.[71] Just over one year later the president issued two further orders intended to improve the sharing of information to prevent terrorism.[72] These orders paid particular attention to the need to standardize the SBU Labels, which were impeding the flow of information within the federal government and between the federal government and state and local governments.[73] In August 2007 the Information Sharing Environment Program Manager was still struggling to develop a plan to streamline the use of SBU labels.[74] By late 2008 a set of new directives placed most non-statutory information protection categories under the larger Controlled Unclassified Information (CUI) category but allows agencies to establish ad hoc markings and control procedures on an individual basis

Downloaded By: [Princeton University] At: 14:30 14 April 2010

[67] Government Accountability Office, *Nuclear Power Plants Efforts Made to Upgrade Security, but the Nuclear Regulatory Commission's Design Basis Threat Process Should Be Improved*.

[68] Ibid., 6.

[69] See Government Accountability Office, *Preliminary Observations on Efforts to Improve Security at Nuclear Power Plants*; and Matthew L. Wald, "Group Says Test of Nuclear Plant's Security Was Too Easy," *New York Times*, 16 September 2003, A7. These two sources could be used to infer the number of attackers commercial nuclear plants are prepared to defend against but played a role in motivating stronger action by the NRC.

[70] Defense Science Board, *Report of the Defense Science Board Task Force on Critical Homeland Infrastructure Protection,* (Washington, DC: Office of the Under Secretary of Defense For Acquisition, Technology, and Logistics, 2007).

[71] Executive Order no. 13356, Federal Register 69, no. 169, 1 September 2004.

[72] Executive Order no. 13388, Federal Register 70, no. 207, 27 October 2005; George W. Bush, *National Strategy for Information Sharing*, National Security Council, The White House, October 2007.

[73] The friction introduced by SBU labels should hardly be surprising given that the majority of federal agencies using various SBU controls do so based on internally generated guidelines that are often ambiguous as to what kinds of information qualify for specific controls, who has authority to apply those controls, and the rules for passing that information to others within the government. National Security Archive, *PSEUDO-SECRETS: A Freedom of Information Audit of the* U.S. *Government's Policies on Sensitive Unclassified Information* (Washington, DC: The George Washington University, 2006).

[74] Patricia McDermott and Emily Feldman, *Secrecy Report Card 2007*.

with no larger review.[75] This slow progress highlights a simple fact: Selective information sharing is hard in systems involving many actors, meaning the protective gains to selective information sharing are likely to fall well short of those from full openness.[76]

One local homeland security official neatly summarized the challenges of selective information sharing: "There are obvious reasons to keep sensitive information out of the hands of terrorists. However I find that, as a public employee involved in preparedness planning, I need specific information on sensitive issues. That information is frequently not accessible. It hinders my planning but the lack of specific information would not discourage a terrorist."[77] Many officials surveyed for this article expressed similar reservations. One state homeland security told us, "I believe secrecy is over-emphasized too often and that it prevents opportunities for state and local law enforcers to interdict information and potential suspects that may be related to terrorism."[78] That even law enforcement officials report challenges with selective information sharing should give pause to those arguing that information can be strategically allocated to all who can make good use of it.

The Impact of General Information

Before government can decide what to protect and terrorists can decide what to attack, both sides need to identify targets. We thus begin our analysis by focusing tightly on the competitive search problem that arises when both government and terrorists are trying to identify vulnerabilities. Here general information helps both players to discern previously unknown targets or modes of attack. On government's side, we can think of the kind of research conducted by Wein and his collaborators, or by Darken and Lewis. On the terrorists' side we can think of groups using publicly available information to develop new attack methods.[79]

Consider a multi-period interaction in which G has a single decision to make: how much general information to release before any search occurs. Assume that there is a finite set of potential targets and/or potential attack modes against these targets, only a fraction of which are known to G and T at the beginning of the game. In each period both G and T search for

[75] Patricia McDermott and Amy Fuller, *Secrecy Report Card 2008* (Washington, DC: OpenTheGovernment.Org, 2008).

[76] The great dissatisfaction with the Homeland Security Information Network (HSIN) provides another cautionary note on the difficulty of selective information sharing as per Lara Jakes Jordan, "Homeland Security Information Network Criticized," *Washington Post*, 10 May 2005, A06.

[77] Respondent 724216542.

[78] Respondent 724339460.

[79] For example, Aum Shinrikyo chemist Masami Tsuchiya reportedly figured out how to synthesize VX nerve gas only after discovering the formula for a VX precursor in a chemistry magazine as discussed in Brian Jackson et al., *Aptitude for Destruction: Volume 2, Case Studies of Organizational Learning in Five Terrorist Groups* (Santa Monica, CA: RAND Corporation, 2005).

Downloaded By: [Princeton University] At: 14:30 14 April 2010

additional targets and attack modes, T to attack and G to defend. General information makes this search more fruitful for both parties, and the more information G releases the faster both G and T find targets.

Searching has one benefit for each side. For T, it widens the pool of known targets, increasing the likelihood of finding a target that will yield greater impact if attacked; we call targets that yield greater impact high-value targets. For G, searching reveals vulnerabilities that T might target. Once vulnerabilities are revealed, G can more effectively defend them, decreasing the chance that T can successfully attack them. In this setting, T faces two costs the longer the search goes on. First, there may be a direct cost to search: It requires the use of resources to keep operatives in the field. Second, there is an indirect cost: The more time T spends looking for better targets, the more time G has to reduce vulnerabilities. Since there is a finite number of targets, at some point it cannot be beneficial for T to continue to search—G will know and will have defended all vulnerabilities by then. Thus, both the direct and indirect costs of search for T are increasing in time. The benefit of search for T, however, is constant under the plausible assumption that T does not get significantly better at identifying high-value targets over time.[80]

This means that there will be some point at which the cost to T of continuing to search outweighs the benefit of search, and T will stop and attack the best target it has identified. In this setting, T can use two different decision rules to decide when to stop searching and attack.[81] In the first, T must decide when to stop searching ahead of time. This is analogous to the situation in which a cell is given instructions to look for targets for a set period and then attack the best target it has seen. In the second, T must decide to attack or to continue searching in each period. Here T's stopping rule is conditional on the proportion of each type of target T has seen.

Regardless of which rule T uses, it turns out that whether or not G should release general information depends centrally on the relative efficiency with which G and T identify targets. If G discovers (and fixes) vulnerabilities much faster than T finds them for a given level of information, then it can be prohibitively costly for T to search at all, due to the increased likelihood that G will have already discovered the vulnerability it would otherwise attack. Attacks in this regime should happen with little delay and will more rarely strike high-value targets. At the other extreme, if T is much more efficacious than G at search—an extreme that we, for reasons detailed above, view as highly unlikely—than any release of general information is bad. In this case, T may engage in search and will achieve a better outcome. In between these extremes, G generally benefits more from releasing information the

[80] The general failure of terrorist cells to identify truly innovative modes of attack—a few high-salience examples aside—supports this assumption.

[81] We formally derive the optimal stopping points and outcomes for T and G under both rules in Propositions 1 and 2 in the appendix available at http://myweb.fsu.edu/dsiegel/Research .

Downloaded By: [Princeton University] At: 14:30 14 April 2010

better it is at using this information to discern vulnerabilities relative to T. One good metric for understanding the extent of G's cognitive advantage is to focus on the capabilities of terrorist groups with which G is concerned. Terrorist groups' analytical capacities vary in sensible ways with their size and the level of government pressure they face. Small underground groups like the Weather Underground in the 1970s typically have negligible research capabilities.[82] Larger groups with safe havens where they operate free from government pressure often devote personnel and resources to basic research. Al Qaeda did so in the 1990s in Afghanistan as did the Provisional Irish Republican Army in rural areas of the Republic of Ireland during the 1970s and 1980s.[83] By taking such obvious indicators into account, officials can make reasoned judgments about their cognitive advantage.

Though the interaction we have described is relatively simple, it reveals a basic fact. When G and T are engaged in a competitive search for targets, releasing general information can reduce G's expected losses even when this information helps terrorists find targets. Focusing only on T's side of the equation, as is typical, misses this important fact.

The Impact of Target-Specific Information

We now turn to the situation where the search process has already been completed so that G and T know the targets, or some subset of them. Our discussion's starting point is Powell's sequential resource allocation game.[84] Here the strategic interaction takes place over two periods. First G decides whether to release target-specific information among all N potentially targeted sites.[85] T observes G's policies—that is, the resource allocations and target-specific information released—for the sites it knows about.[86] Then

[82] Larry Grathwohl in Larry Grathwohl and Frank Reagan, *Bringing Down America* (New Rochelle, NY: Arlington House, 1976), 143-44 describes a remarkably amateurish bomb design proposed by senior Weatherman William Ayers.

[83] Shane O'Doherty, *The Volunteer: A Former* IRA *Man's True Story* (Scarborough, ON: HarperCollins Publishers Canada, 1993) for example, provides a vivid description of the development of the Provisional Irish Republican Army's (PIRA) gelignite bombs.

[84] The game Robert Powell analyzes in "Defending Against Terrorist Attacks with Limited Resources," illustrates the dynamics of resource allocation, but does not consider the role of information. Robert Powell, "Allocating Defensive Resources with Private Information about Vulnerability," *American Political Science Review* 101 (2007): 799–810 considers an aspect of information we do not, the fact that resource allocation itself can signal the level of vulnerability of a known site, but does not consider the more explicit forms of information that we do. In his model, informative signaling (in the form of a separating equilibrium) may sometimes be achieved if the more vulnerable sites are easier to defend on the margins. Though developed independently, our models also speak to some of the issues analyzed in Jun Zhuang and Vicki M. Bier, "Balancing Terrorism and Natural Disasters—Defensive Strategy with Endogenous Attacker Effort," *Operations Research* 55, no. 5 (September-October 2007): 976–91. Their main focus is on resource allocation decisions when terrorists choose a level of effort, rather than making a binary decision about whether or not to attack.

[85] We can think of these as the sites that G found during the earlier search process.

[86] That is, for those sites that it found during its search process.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

Downloaded By: [Princeton University] At: 14:30 14 April 2010

T attacks the target that grants it the highest expected benefit, taking G's strategy into account. The probability of a successful attack by T depends on both the target-specific information released about that site and the resources allocated to that site. Here the benefit to T for having earlier discovered a previously unknown vulnerability becomes apparent. Because G can allocate neither resources nor target-specific information to a site about which it does not know, T will have an easier time successfully attacking these vulnerabilities. Given two targets of equal value, T will attack the easier one.[87]

Sharing target-specific information has three effects in this interaction. First and most importantly, it can alter the chance of a successful attack at a single site. Second, it can increase the chance T knows about a site. Third, releasing target-specific information may occasion externalities. On the positive side, releasing information can satisfy the normative principle that citizens should be aware of the dangers created by industrial facilities near their homes. This was one of the major motivations for the Emergency Planning and Community Right-to-Know Act (EPCRA).[88] The law requires facilities handling certain chemicals to make annual reports to the EPA and local officials regarding the nature and average amount of hazardous materials stored on their sites. Public access to this data has been significantly reduced due to terrorism concerns.[89] On the negative side, public officials who have identified a specific target may suffer an additional political cost if that location is subsequently attacked. Some of these externalities come into play only when a site is successfully attacked; others are realized regardless of the outcome of the attack.

The simplest possible game in this setting illustrates the tension between secrecy and openness. Suppose there is only one target, and that resources are static, so that G just has the decision-theoretic problem to release some level of information. If G not does not release information (that is, the level is zero), then the chance that T knows about the site (and so can attack it) depends solely on the outcome of the prior search game. If G does release information, then we assume it becomes more likely that T learns about the site and the more information G releases. This effect of information release is purely negative for G, as it increases the chance that the site is attacked, and it is this effect that is typically the focus of public debate. However, there is also a second effect because releasing target-specific information alters

[87] This does not, of course, preclude T from repeatedly attacking a well-known vulnerability, such as airliners, if it views this as a particularly high-value target or if G is unable to defend this vulnerability sufficiently well despite its being common knowledge as per Jeremy Shapiro, Managing Homeland Security: Develop a Threat Based Strategy (Washington, DC: Brookings Institution, 2008). The important comparison for T is the expected utility derived from an attack on each target, which depends on both the target's value and its probability of being attacked successfully.

[88] Emergency Planning & Community Right-to-Know Act, P.L. 99-499, 42 USC 11001–11050.

[89] Moteff and Stevens, *Critical Infrastructure Information Disclosure.*

the ability of G to defend the site. This, as we argued above, is typically a positive effect: G is usually able to make better use of information than T. Thus releasing information will often lower the probability of a successful attack, even as it raises the probability that T will discover the site. If the utility gain from the positive effect outweighs the loss from the negative effect, G should release target-specific information.

This simple example illustrates that as long as openness can help defend targets, it cannot be right that releasing information that increases the chances of an attack is always bad. Adding realistic complications to the model does not alter this insight. In general, the defensive gains from openness are more likely to dominate whenever releasing information: (1) produces a small degree of expected backlash; (2) is likely to produce a large utility gain to G due to the decreased likelihood of a successful attack; and (3) does not greatly increase the likelihood that T discovers the site, perhaps because the probability that T already knows of the site before G releases information is large.

When will each of these conditions be obtained? Limited backlash is likely to occur when public support for government is high and the population believes the chance of attack is also high. Polling can provide an assessment of both beliefs.

Whether G should expect big gains from information release depends on the source of the costs of an attack. When major costs arise from the direct impact of the attack, such as the impact of explosives on a plane in flight, then protective gains have to come from either motivating better defense or innovations in defense. Given weight limitations, efforts to secure aircraft against internal explosions are unlikely to pay significant dividends. In general, the prospects for openness to help when costs are direct seem relatively low unless there are major gaps between public and private interests.

When major costs are cascading disruptions, as in an attack on a major piece of infrastructure or the failure of the screening system at an airport, then there are greater opportunities for analysis to identify cheap ways to add redundancy or to reduce interdependencies.[90] As a general rule, decision makers seeking to identify the potential gains from information sharing should ask, "is this information relevant to a system/target which could be made more robust if we understood it better?" Scientific and engineering analysis can help determine this in general. If the answer is "yes," then selective information sharing to government scientists can provide the motivation for further openness to the public.

A simple proxy for whether or not terrorists know of a vulnerability beforehand is whether or not the vulnerability is observable, and particularly whether it is casually observable. A directed internet search can provide a

[90] Indeed, both Sean Gorman, *Networks, Complexity, and Security* and Wein and Baveja, "Using Fingerprint Image Quality" do just that.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

secondary proxy. Suppose, for example, that the chemical detection swabs used at airports routinely missed a certain class of explosives. It would likely take substantial research to determine this fact and thus government would do best by not releasing information about the standards for the detection systems. In contrast, the fact that large amounts of volatile chemicals are transported over rail lines running blocks from the U.S. capital is easily discernable, and thus relevant information release is less likely to provide new information to terrorists.

An example of this in practice is work by academic physicists on screening for nuclear weapons in shipping containers.[91] This threat is well-known, so that terrorists are likely to know about it and backlash for revealing information is likely to be comparatively small, since the possibility of such an attack is already expected. Further, utility gains are potentially large due to the current inability to screen more than a small percentage of incoming containers at reasonable cost and the gains that could accrue via technological improvements. Openness, in the form of publishing basic information on the radiological characteristics of different materials, uses the academic incentive structure to get experts who would not otherwise study the problem to do so at little cost to government.

These dynamics play out in a slightly more complicated fashion when we consider the case where there are multiple targets. The intuition here is best developed by considering the simple case where G has a binary decision: to release information or not for each target. G has some prior belief about the probability T knows about each target; if G releases information, then the probability T knows about that target goes up. At the same time, the ability of G to defend the target also goes up if it releases information. We have seen that, for any target considered in isolation, G faces a straightforward trade-off between defense and discovery. When considering multiple sites the same trade-off at each site still holds, but now decisions at one site can shift attacks to other sites, as they become comparatively more attractive to T. For example, releasing information might make T's favorite site unpalatable due to a decrease in the probability of a successful attack, leading T to attack what had been its second-favorite site instead. This introduces complexity, in that G might not want to release information at a site to prevent T from instead attacking a different site that G values more. In the appendix, we describe an algorithm that takes into account these interdependencies and yields thresholds for G's prior beliefs on the likelihood that T knows about each site ahead of time.[92] If the chance that T knows about each site before G releases information exceeds these thresholds, G

[91] See, for example, J. I. Katz, "Detection of Neutron Sources in Cargo Containers" *Science and Global Security* 14, nos. 2 and 3 (December 2006): 145–49, which suggests an alternative screening technique that remedies challenges previously identified by physicists working for the U.S. government.

[92] Appendix available at http://myweb.fsu.edu/dsiegel/Research.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

releases information on that site, otherwise it does not. As we would intuitively expect, these thresholds are lower when G gets significant defensive improvements by sharing information.

Information and Resources

If releasing information can be beneficial, the key question is when will it be beneficial in a world of limited resources. To get a handle on this question we can consider the situation where G can release information on and allocate defensive resources to each site. T's probability of successfully attacking any given site depends on both. To understand this interaction we need to make a few assumptions. First, we assume that allocating resources does not alter the likelihood that T knows about the target.[93] Instead, the probability of T's knowing about a site is dependent solely on information released by G, and any prior information T might have. We believe this is a reasonable assumption; certainly not all resource allocation is observable to T, particularly when the site itself is unknown. Furthermore, T can observe resources going into a target without knowing how to make use of this information. We also assume that we are dealing with pure negative backlash, so that the costs to G if a site is attacked are larger if information is released about the site. This makes it less likely that government would release information, so this assumption biases our results away from the conclusion that information is released in equilibrium.

Analyzing this more complex setting shows that the addition of resource allocation strengthens the case for the release of target-specific information. The core assumption driving this finding is that information that contributes to defenses at a given site can have second-order effects by altering the marginal benefits to resource allocation at all sites. Consider the concrete example of sharing genetic information about disease genomes. This information has obvious utility to terrorists but can enable drug companies to speed development of effective vaccines for new strains.[94] This not only enables better defense against biological attack due to the vaccines, but also frees some of the resources that formerly were spent stockpiling prophylactic medicine to be used elsewhere. The degree to which this kind of dynamic obtains depends on the degree to which the release of information enables a more efficient use of resources at one or more sites. We believe that this efficiency gain will typically be realized for the reasons detailed above. If nothing else, information release can increase the set of options available to government, increasing the likelihood that a better method of defense might be found. However, when information has little impact on success,

[93] Robert Powell, "Allocating Defensive Resources with Private Information about Vulnerability," examines the case where it does.

[94] Falkow et al., *Seeking Security.*

Downloaded By: [Princeton University] At: 14:30 14 April 2010

Downloaded By: [Princeton University] At: 14:30 14 April 2010

or occasions few efficiency gains, it is less likely that its release will be beneficial.

An important implication of this dynamic is that G's decision about the total amount of information released will be affected by the amount of resources that are available. The greater those resources, the greater weight G should give to the efficiency gains in defense resulting from information release. In other words, increasing resources tips the balance inherent in releasing information toward protection and away from backlash and discovery. Even when T's knowledge of targets is uncertain, if information makes resource use more efficient, there will be cases in which government should release target-specific information. From this perspective, the fact that the massive increase in resources devoted to protecting against terrorism from 2002-2008 was accompanied by a large increase in secrecy provides evidence that U.S. government decisions about openness were not being made optimally if the goal was to maximize defense against terrorism.

SECRECY AND OPENNESS IN PRACTICE

One natural concern that follows from the section "Strategic Information Release" is whether the increase in secrecy observed by so many actually demonstrates an unwarranted bias in favor of secrecy.[95] It could be that government officials are analyzing the problem as we suggest, but that the relevant variables have changed in ways implying greater security is, in fact, optimal. To see whether this is the case we surveyed 186 U.S. federal, state, local, and industry homeland security officials and asked two kinds of questions. First, we asked them a simple open-ended question: "How do you think about the tradeoff between secrecy and openness in Homeland Security?" Second, we asked a number of questions designed to measure how they value the key variables discussed in the section "Strategic Information Release."

We sent the survey to 510 members of the alumni network of the Naval Postgraduate School's Center for Homeland Defense and Security (CHDS). Our 38 percent response rate is comparable to that yielded by other expert e-mail surveys.[96] The CHDS alumni network represents a broad cross section of federal, state, and local security officials and so is an appropriate population for addressing the question of how officials balance secrecy and openness in practice. Overall, 36 percent of our respondents work for the federal government, 24 percent for state governments, 34 percent for local governments,

[95] We thank the editors for suggesting this possibility and pushing us to address it.

[96] For example, the IR Scholar Survey detailed in Richard Jordan et al., "One Discipline or Many? TRIP Survey of International Relations Faculty in Ten Countries," 2009, College of William and Mary, Williamsburg, VA, available at http://irtheoryandpractice.wm.edu/projects/trip/Final_Trip_Report_2009.pdf, received a 42 percent response rate from potential U.S. respondents.

and 6 percent for private firms. Importantly, this is a population of officials who have explicitly sought out advanced professional education, meaning they should be biased toward more careful analysis of this problem than the average homeland security official. We expect this self-selection process to bias our results in favor of officials analyzing the problem in the ways suggested in the section "Strategic Information Release."

Three results stand out from our analysis. First, responses to the open-ended question show there is no clear consensus among officials about how to approach the trade-off between secrecy and openness. Some respondents unreservedly favored secrecy. One local official argued, "We need it (secrecy) especially in government law enforcement and defense. Terrorists both homegrown and international watch all our government and private web sites."[97] Others took a clear position in favor of openness, typically because they doubted the ability of officials to access required information without it. An industry official took the slightly jaundiced view that "terrorists/criminals will normally find the information regardless, but resource-limited agencies responsible for protection and preparedness will not have such diligence."[98] A few discussed the problem in terms that should now be familiar to the reader. One federal official eloquently argued, "This is a crucial balancing act which must be considered in a case-by-case basis. Each instance of debating the opening of homeland security information must consider whether sources will be divulged and whether releasing the information is likely to lead to saving lives or apprehending offenders."[99] While state and local officials were somewhat more likely to argue for openness on the grounds that excessive secrecy creates problems for law enforcement and preparedness agencies, there were no gaping federal/state/local divides apparent among our respondents. Officials at all levels described widely varying approaches to this issue.

Second, officials do not generally believe changes in the level of secrecy have increased security. Sixty-seven percent of our respondents report secrecy has increased since 2000, and 30 percent report it has remained the same. At the same time, 60 percent of our respondents report that the changes in information control have had no effect or a negative effect on the safety of society from terrorism. This perception is not consistent with the hypothesis that changes have been driven by a well-reasoned effort to increase security.

Third, and most importantly, officials do not believe the environment has changed in ways that imply increased security if they are thinking about the problem in the ways we advise. Specifically, we argue that the balance should tip to openness if (1) there are more protective resources available;

[97] Respondent 723214546.

[98] Respondent 724627525.

[99] Respondent 724716301.

Downloaded By: [Princeton University] At: 14:30 14 April 2010

Downloaded By: [Princeton University] At: 14:30 14 April 2010

(2) terrorists know more about vulnerabilities ex ante; (3) the efficiency of protective resources increases with openness; and (4) there are positive externalities to openness, that is, the public's right to know.

Respondents to our survey overwhelmingly believe (1) and (2) to be the case relative to the year 2000. Fully 97 percent of respondents report an increase in the "resources our governments (federal, state, and local) devote to protecting society from terrorism." A slightly smaller number, 79 percent, said there has been an increase in "terrorist organizations' knowledge about potential targets in the United States." Both point to greater openness as the optimal policy from a counterterrorism perspective.

To assess (3) and (4) we asked respondents to describe their decision process regarding an unspecified vulnerability. A substantial minority unambiguously agreed that information release could render protective resources more efficient. When asked, "How much does the possibility that independent analysis might reduce that vulnerability factor into your decision?" fully 45.5 percent of respondents answered "a lot," while another 45.5 percent respondents "a little." Thus 91 percent of our respondents saw the potential for independent analysis to help. A similarly large proportion of respondents, 85 percent, said they would take the public's "right to know" into account. When asked "How much does the public's 'right to know' factor into your decision?" Thirty-four percent responded "a lot," 52 percent responded "a little," and only 15 percent responded "not at all."

Whether taking these factors into account implies greater openness depends on the specifics of the vulnerability in question, but it is hard to argue there has been a meaningful decrease since 2001 in either the attention independent analysts pay to remedying vulnerabilities to terrorism (quite the opposite in fact) or the public's inherent right to know what government is doing. These results are thus hard to square with a uniform increase in secrecy being optimal from the perspective of the section "Strategic Information Release."

There is one possible way in which our survey results could be consistent with officials analyzing the problem as we suggest. The section "Strategic Information Release" notes that concerns with political backlash can motivate increased security. Our survey does find that 40 percent of respondents give "a lot" of weight to the potential for political backlash in making information sharing decisions and that fully 53 percent believe backlash is "very likely" to occur if they "were to release information about a vulnerability that was subsequently targeted." Our survey results are thus consistent with either (1) homeland security officials not analyzing this decision in the manner the section "Strategic Information Release" argues they should, or (2) concerns over political backlash trumping security from terrorism. In the first case, this paper can serve as a valuable corrective. In the second case, this paper illustrates there are security gains being lost to political considerations. Both are important.

POLICY IMPLICATIONS

Policy makers today face a devilish task: balancing the general desirability of openness in a democratic polity against the duty to protect the people from terrorism. Tension between secrecy and openness is not new to modern states, but the optimal trade-off is much harder to identify in the current security environment. In 1980, American officials were not concerned that Soviet forces would attack water treatment systems or seek to sabotage chemical plants. As security policy has become increasingly focused on non-state threats, understanding what types of information create unacceptable security risks has become much harder.[100] Policy makers thus have greater discretion in restricting information than in any time in recent history and so need better guidance on how to navigate the trade-offs between openness and secrecy. Our analysis can help.

We begin with the key point: Under a wide range of conditions, open sharing of government's private information can enhance efforts to protect citizens. This is true even when the information deals directly with specific targets. When the positive externalities of information sharing are taken into account, the set of conditions under which open information sharing benefits society becomes wider still. Our analysis puts to rest the overly simple conception that revealing vulnerabilities to the other side is strictly a poor idea. Instead, we must substitute a more nuanced picture, one that focuses on the relative gains to both sides of information sharing. In many cases, there are benefits to openness that outweigh the costs of revealing targets or of helping the other side operate more effectively.

A number of results emerge from our analysis of the strategic interaction around openness. First, it pays to release information that speeds the rate of target discovery when three conditions are obtained: (1) government can defend known targets better than unknown targets, (2) terrorists prefer to attack unknown targets, and (3) analysts on the side of government have a sufficiently large cognitive advantage over the terrorists. Assumption one and its corollary, assumption two, are supported by the record of terrorist groups' seeking to find innovative modes of attack, searching for the path of least resistance.[101] As government has access to cognitive potential vastly exceeding that of terrorist groups (assumption three), our analysis suggests much more attention should be paid to how openness can help government

Downloaded By: [Princeton University] At: 14:30 14 April 2010

[100] Recognizing the increased complexity of information policy, the 110th Congress created a subcommittee explicitly charged with examining how America guards and shares information. See Alexander Bolton, "Waxman to stir debate with transparency subcommittee," *The Hill*, 4 January 2007.

[101] As discussed in the section "Strategic Information Release," this does not imply that terrorists will never rationally target well-defended targets, such as airliners, if they are also of particularly high value. Walter Enders and Todd Sandler, "What Do We Know about the Substitution Effect in Transnational Terrorism." In Andrew Silke and G. Ilardi ed., *Researching Terrorism: Trends, Achievements, Failures* (Ilford, UK: Frank Cass, 2004).

Downloaded By: [Princeton University] At: 14:30 14 April 2010

identify vulnerabilities. If this information could be shared in such a way so that terrorists cannot see it, so much the better. However, the evidence strongly suggests selective-sharing of general information frequently, if not always, slows the rate of target discovery. Avoiding this decrease requires government to identify the relevant issue-area experts, know what information they need, and be able to share it with them effectively. All are difficult to accomplish given the range of issues involved in protecting society against terrorism. Further, even when selective-sharing is possible, the controls it requires limit the independence of the analysts and hence the value of their analysis.

Obviously, some information should be kept secret, and our analysis of information sharing also analyzes the boundaries of when government should share target-specific information. When the probability the terrorists already know about a target is small and the protective gains from information sharing are expected to be small, government should never reveal private information. It is hard to see, for example, how releasing information about which parts of government buildings are most vulnerable to vehicle bombs could help. Conversely, when the probability the terrorists already know about a target is large and the protective gains to information sharing are expected to be large, government should always release information. This seems likely to be the case for a wide variety of infrastructure targets that are (1) readily observable, and (2) critical components of larger economic and social systems.

In all cases, the threat of political backlash decreases the likelihood that government will release information, even when it would be socially optimal ignoring political considerations. From this perspective, governments more concerned with social welfare than with political gain (or with fewer agency problems) should be more open, more likely to share target-specific information.

A similar set of relative considerations apply when balancing the potential defensive gains from information sharing against the fact that openness helps terrorists identify targets. When a target is likely to be very poorly defended unless information is released, despite potentially being known to terrorists, it is extremely likely that government is better off with open information sharing. These conditions are most likely to be met with privately owned targets in weakly regulated industries.[102] For example, facilities using certain chemicals are required to report potential consequences of accidents to the Environmental Protection Agency. Such information has obvious utility for terrorists and so is withheld from the public. Terrorists, however, can observe the hazmat labels prominently posted on trucks and rail cars

[102] Our findings thus match and deepen the intuition embedded in CBO, *Homeland Security and the Private Sector* and Flynn and Prieto, *Neglected Defense* that encouraging openness increases the chances that private companies will undertake the appropriate level of counterterrorism protection.

coming into chemical facilities. Making the reports public could help politicians persuade the chemical industry to take protective measures at the most dangerous sites, measures the industry has aggressively resisted. The current policy amounts to trusting terrorists to be less observant than officials tasked with identifying terrorist surveillance efforts. It is entirely possible that this is not the optimal trade-off.

One of the hardest challenges facing policy makers is how to think about balancing the secondary effects of releasing information, enhancing scientific research and the like, against the more obvious costs. Our analysis suggests that when releasing information makes the use of defensive resources more efficient, it is more likely to be in government’s interest to do so. Thus, ceteris paribus, government should be especially reticent about keeping information private in areas where basic research is likely to contribute to protection. We note sadly the opposite has been the case since 2001. The deeper point revealed by our analysis is that it is not enough to take into account the direct effects of sharing information. Doing so misses a key part of the problem. Government officials should consider how information sharing can improve the efficiency of protective spending.

Finally, we find that if releasing general information can also help government defend against specific targets, it becomes even more beneficial to release all types of information. This multiplier effect to releasing certain kinds of information that are potentially useful to terrorists has not been previously identified, but can be central to policy decisions about what should be done with government’s private information.

The key takeaway from this analysis is that the current discussion of information sharing should be substantially deepened. One federal official neatly summarized a more nuanced approach: “Secrecy does not necessarily increase security. Although it may deny information to our adversaries, it also denies information to those who need access; perhaps decreasing our ability to protect ourselves; perhaps decreasing the level of trust our citizens have in our government . . . . It’s a tricky balance that needs constant oversight . . . ”[103] Putting these words into practice means carefully considering (1) the characteristics of different types of information; (2) their impact conditional on who is responsible for different targets; (3) the secondary effects of releasing information; and (4) the externalities, both positive and negative, arising from sharing each type of information. By taking such an approach, government will better meet its mandate to maintain an open society while protecting citizens from the threat of terrorism.

[103] Respondent 724452379.

Downloaded By: [Princeton University] At: 14:30 14 April 2010