DOI: 10.1111/j.1539-6924.2010.01455.x

Perspective

Reasons for Secrecy and Deception in Homeland-Security Resource Allocation

Jun Zhuang[1],* and Vicki M. Bier[2]

In this article, we explore reasons that a defender might prefer secrecy or deception about her defensive resource allocations, rather than disclosure, in a homeland-security context. Our observations not only summarize and synthesize the results of existing game-theoretic work, but also provide intuitions about promising future research directions.

KEY WORDS: Deception; game theory; private information; secrecy; truthful disclosure

1. INTRODUCTION

Unlike natural disasters, attackers are intelligent and adaptive.[1,2] Therefore, while investments in protection from natural disasters are usually disclosed to the public, anti-terrorism investments are not always disclosed. Understanding when and how such investment information should be disclosed is a challenging issue for governments and other organizations facing terrorist attacks or other intentional threats.

Traditionally, security-related information such as defensive resource allocations is often kept secret. Preventing potential attackers from obtaining such information may make defenders feel safer. Moreover, there is a long tradition of deception in the military arena, as well as in business and capital ventures. However, while it is generally considered legitimate to use secrecy or deception in military strategies, these might not always be viewed as viable or legitimate strategies for homeland security. For example, even if it is optimal to deploy false defenses in order to deter military attacks, it may be politically infeasible to rely on this strategy in the context of homeland security, since members of the public might dislike being deceived or misled, and may want to know how they are being protected, even if that makes the defenses less effective.

In this article, we explore reasons that a defender might prefer secrecy or deception, rather than truthful disclosure, in a homeland-security context. In particular, Section 2 reviews the literature on secrecy and deception. Section 3 focuses on the case when the defender has private information; that is, when the attacker is uncertain about some defender characteristics, and can use information disclosed by the defender to infer those characteristics. Section 4 considers the benefits of partial secrecy and partial defense when the success probability of an attack is a nonconvex function of the defensive investment. Section 5 lists a few other reasons for secrecy and deception. Finally, Section 6 concludes this article, and suggests some future research directions. Our observations summarize the results of past game-theoretic work, but also provide intuitions about promising future research directions.

[1] Department of Industrial and Systems Engineering, University at Buffalo, the State University of New York, NY, USA.

[2] Department of Industrial and Systems Engineering, University of Wisconsin–Madison, WI, USA.

*Address correspondence to Jun Zhuang, Department of Industrial and Systems Engineering, University at Buffalo, the State University of New York, Buffalo, NY 14260, USA; tel: 1-716-645-4707; fax: 1-716-645-3302; jzhuang@buffalo.edu.

0272-4332/10/0100-1737$22.00/1 © 2010 Society for Risk Analysis

2. LITERATURE REVIEW

In the theoretical literature, disclosure of player actions is often found to be preferable to secrecy. Levy notes that the benefits of truthful disclosure include "enhanced accountability, enhanced predictability, and the provision of expert information to the economy."[3] For example, in the classic game of Prisoner's Dilemma, disclosure increases coordination between players, and therefore increases the benefits to both parties over time in repeated games. However, homeland security is not a coordination game, and intuition suggests that secrecy or even deception may sometimes be desirable.

Recent game-theoretic research has also indicated that publicizing defensive information instead of keeping it secret may help to deter attacks.[4,5] Similarly, Edmonds has argued that classifying too much information could hurt national security by impeding information sharing between first responders or others responsible for security, thereby decreasing the effectiveness of defenses.[6] Thus, it appears that at least under some circumstances, there can be merits to releasing defensive information.

Secrecy and deception have, of course, been investigated in military analyses,[7] psychology,[8] and computer science,[9,10] as well as economics and political science.[11,12] However, few of these studies have focused specifically on disclosure of resource allocations.

The Merriam-Webster online dictionary defines "secrecy" as "the habit or practice of keeping secrets or maintaining privacy or concealment."[13] Bok describes secrecy as "intentional concealment."[14] These definitions are straightforward and widely accepted (with minor modifications) in most fields, including the military, computer security, economics, political science, and homeland security. Tefft surveys secrecy from a variety of perspectives, including the comparison with privacy, the politics of secrecy, secrecy in business, and bureaucratic secrecy.[15]

Secrecy has been sometimes modeled as simultaneous play in game theory,[2] since in a simultaneous game, each player moves without knowing the moves chosen by the other players. Note that this does not actually require both players to make their decisions at the same time; the players can be viewed as being engaged in a simultaneous game as long as neither party knows the other's decision at the time it makes its own decision.

By contrast, although the dictionary definition of "deception" ("the act of deceiving")[16] is still straightforward, different fields have focused on different aspects of deception. In particular, the Joint Chiefs of Staff define "military deception" as "those actions executed to deliberately mislead adversary decision makers as to friendly military capabilities, intentions, and operations, thereby causing the adversary to take specific actions (or inactions) that will contribute to the accomplishment of the friendly mission."[7] This definition covers deception about both the defender's private information (capabilities and intentions), and the defender's planned future actions (operations).

Many other kinds of deception have been discussed in the literature. Some researchers model deception as sending noisy or imperfect signals to mislead one's opponents. For instance, Hendricks and McAfee and Oliveros use the Normandy invasion as an example to argue that the first mover (the Allies) optimally allocated resources to targets that they did not intend to attack in order to mislead the Germans about their true landing place by sending noisy signals.[17,18] Similarly, Brown *et al.* define deception in a zero-sum attacker-defender game as occurring when the defender discloses partial defenses in an attempt to route attacks to heavily defended locations.[19]

In this article, we define deception as disclosure of a different level of security investment than a defender has actually made. This is perhaps most similar to the definition given by Board in the context of the Trojan War; in particular, he defines deception as a situation in which "one player tricks another into believing that she has done something other than what she actually did."[20]

It is difficult to find specific examples of government secrecy and deception in homeland-security resource allocation, in part because these phenomena are often classified. In general, the use of government secrecy and/or deception generates criticism as being contrary to democracy,[21] and interfering with accountability.[22,23] There are also ethical concerns regarding government uses of secrecy and deception. Nevertheless, government secrecy and deception with regard to resource allocation have occurred in the context of homeland security; one example is the use of fake security cameras and guards to deter possible attacks. In the next section, we provide a detailed discussion of one possible motivation for defender secrecy or deception—namely, to exploit the attacker's uncertainty about the defender private information.

3. TAKING ADVANTAGE OF DEFENDER PRIVATE INFORMATION

Attacker uncertainty about defender private information can create opportunities for either defender secrecy or deception. In particular, defender private information can include knowledge about the existence of a particular target, about defense costs, about the effectiveness of defensive measures, and/or about asset valuations and vulnerabilities.

Many targets of interest (e.g., the Sears Tower, the Pentagon, and the Golden Gate Bridge) are well known to potential attackers, and attackers may already have reasonable estimates of their valuations. For targets like these, where the asset valuations are well known, Shapiro and Siegel show that it would be better to release defensive information rather than keeping it secret; in fact, it seems plausible that defenders may even have incentives to deceive by overstating their defenses in such cases in order to deter attacks.[24] By contrast, if the terrorists do not even know about the existence of a target, Shapiro and Siegel show that defenders should not release target-specific information. For example, a government would presumably have no reason to disclose the defenses of a top-secret nuclear facility if the attacker did not even know that the target existed. (Note that this falls outside the range of phenomena usually studied using game theory, which typically assumes that the attacker knows about the existence of the targets, even if the defender keeps her defensive investments secret.)

Generalizing the above results by Shapiro and Siegel, we hypothesize that it may be desirable to disguise valuable targets whose true value is not well known as being of low value, by understating their defenses. To illustrate, consider a theft game in a neighborhood consisting of both rich and poor households. Naturally, thieves are more interested in rich households. Thieves are in general uncertain about the type of each household, but may attempt to infer that type from observable defenses such as burglar alarms. In a wealthy neighborhood, where the thieves already know that most households are rich, there is relatively little information to be gained by thieves observing that a particular household has a burglar alarm or other defenses. Therefore, the deterrent value of disclosing such defenses (to convince thieves that they are unlikely to succeed) could be expected to outweigh the risk of revealing that the defended household is wealthy. By contrast, in a poor neighborhood, where most households have nothing valuable to protect, an elaborate system of defenses might make a household that has expensive valuables to protect dangerously conspicuous. In such neighborhoods, those few households that do have assets of significant value might prefer "secret" or relatively unobservable defenses (such as bedroom safes) overeasily observable defenses (such as barred windows or fierce guard dogs).

Zhuang developed a rigorous game-theoretic model of the above scenario, but, perhaps surprisingly, it did not yield secrecy at optimality.[25] Perhaps this is because Zhuang considered only pure strategies (i.e., without randomized choice of targets or randomized allocation of defenses). We speculate that we may be able to find secrecy and deception in this scenario if we allow randomized strategies; that is, if the thief attacks only with some probability (rather than for certain). In this case, rich households may want to deploy secret defenses in order to protect themselves if a theft does occur without attracting higher theft probabilities by revealing that they are high-value targets. We also expect that we may find secrecy in this type of model if we allow the thief to be nonstrategic (e.g., if we depart from the standard game-theoretic assumption that the thief is able to make perfectly accurate inferences from any observed defense; see Hao *et al.*, for a game in which one player may be nonstrategic with nonzero probability).[26]

Brown *et al.* find a benefit of secrecy in a zero-sum attacker-defender game in the context of ballistic missile deployment, perhaps because their attacker is not fully strategic in the above sense. In particular, in Brown *et al.*, the attacker is assumed to believe that locations are totally undefended if the defender chooses to keep the defenses of those locations secret;[19] Brown *et al.* treat secrecy similarly in the case of interdicting nuclear projects.[27] This avoids the unrealistic game-theoretic assumption of perfect attacker inference, but may be overly optimistic, since realistically, attackers may have prior beliefs about the likely levels of target defenses even when no defense has been observed.

Powell studies an attacker-defender, multiple-target game where the defender has private information about asset vulnerabilities (instead of valuations).[28] He argues that investing more in defense could be a signal to the attacker that the heavily defended targets are more vulnerable, and therefore may increase their probability of being attacked. He finds an equilibrium in which the defender allocates her defensive investments without regard to the vulnerability of the various targets, so that the attacker cannot infer their vulnerability. In particular, according to Powell, this type of equilibrium will hold when the more vulnerable sites are also more costly to protect. In this case, the defender will be better off keeping their vulnerability secret, rather than tipping off the terrorists to their vulnerability by spending a lot on defending them.

Hausken and Levitin study an attacker-defender game where the defender builds both genuine and false targets (in a series system), and the attacker chooses the targets to attack in such a way as to maximize the system vulnerability (given the attacker's lack of knowledge about which targets are false).[29] False targets can clearly be viewed as a type of defender deception. Since Hausken and Levitin assume that the attacker cannot distinguish between genuine and false targets, such deception gives the defender some advantages.

For another example of private information, Zhuang and Bier provide numerical examples for a game in which the defender's signaling cost is private information; that is, a defender may have either high or low costs of deception and secrecy, and the attacker is not sure whether those costs are high or low.[30] They find that the defender with lower deception and secrecy costs would in some cases prefer to overstate her defenses in order to mimic the defender with higher deception and secrecy costs, and therefore achieve attack deterrence at low cost by free riding on the other defender. In other cases, the defender with lower secrecy and deception costs was found to prefer secret defenses over disclosed defenses in order to prevent the defender with higher signaling costs from free riding (and thereby attracting attacks against both types of defenders). It is interesting to note that equilibriums involving secrecy and deception in Zhuang and Bier occur only at intermediate ranges of parameter values. In particular, when the asset valuation is high, then the defender will defend, and truthfully disclose that defense; similarly, when the asset valuation is low, the defender will not defend, and will truthfully disclose the lack of defense (since the attacker would be able to infer the lack of defense anyway from the low asset value). Only for intermediate asset valuations (where defense is almost or just barely justified on its own) does secrecy or deception play a role at equilibrium.

4. INCREASING THE COST EFFECTIVENESS OF DEFENSE

Intuitively, it seems plausible that revealing the defender's resource allocations could make those defenses less effective, and therefore attract more attacks and/or increase the success probability of attacks. In particular, disclosure of which targets are protected, and which specific defensive technologies are used, could in principle make it easier for the attacker to evade and/or overcome those defenses. For example, if every U.S. post office had installed anthrax sterilization equipment, and announced that information to the public (and therefore to terrorists), potential attackers might just deliver anthrax by private couriers such as Federal Express or United Parcel Services. Thus, the defender could have spent many millions of dollars, but made the attacker's job only negligibly more costly (by making the attacker pay the slightly higher shipping fees charged by private couriers). By contrast, secret sterilization equipment could presumably have been effective against future anthrax attacks through the mail. Similarly, when the U.S. Transportation Security Administration announced that liquids could no longer be taken through security at airports, one could imagine that prospective attackers might just have switched to solid explosives.

Again, Zhuang developed a game-theoretic model to investigate the above hypotheses, but was not able to find secrecy or deception at equilibrium in fully endogenous models[25] (although we expect that we might find such behavior in models allowing realistic nonstrategic behavior, such as that by Hao *et al.*[26]). However, secrecy was found to increase the effectiveness of defensive investments in cases where the success probability of an attack is a nonconvex function of the level of defensive investment. To understand why, suppose that we have a finite defensive budget, and there are two identical targets to defend. Assume further that the success probability of an attack is a nonconvex function of the defensive investment, so that small investments do little to reduce the success probability, and only a large investment is highly effective. In this case, distributing the budget evenly between the two targets might only leave both targets vulnerable; therefore, the defender might instead prefer to defend only one target (spending the entire budget on one target). In this case, the defender would clearly prefer secrecy to disclosure, since disclosure of the defended location would just attract attacks against the undefended target. As before, Dighe *et al*. find such secrecy only for intermediate asset values—too low to justify high levels of defense in both targets, but high enough for secret defense of a single target to be credible to attackers.(31)

Note that such nonconvex functions can easily arise in practice; for example, in cases where there is a minimum investment needed for defenses to be highly effective (say, because a more effective technology is costly, or because defenses are discrete rather than continuous). One practical example of a discrete defensive measure is the use of Lojack, which can be used to track down stolen cars. Ayres and Levitt describe the results of an empirical study of the benefits of Lojack.(32) The results of Ayres and Levitt suggest that the availability of information regarding the approximate market penetration of Lojack may help to reduce car thefts by deterring potential attackers, as long as the installation of Lojack on any given car is kept secret. By contrast, disclosing which cars are protected by Lojack would just put unprotected cars at greater risk.

Another illustrative application of this result is the case of onboard air marshals.(31) If the information about which planes have air marshals is kept secret, but the number of air marshals is known (at least approximately), then having air marshals on only some planes could result in attack deterrence. The importance of secrecy was highlighted by the Federal Law Enforcement Officers Association, when the cover of air marshals was allegedly in danger due to a rigid dress code.(33) This could present a significant risk if it increased the chance that terrorists could identify the air marshals, overpower them, and hijack a plane.(34)

Finally, preventing terrorists and weapons of mass destruction from entering the United States is clearly a top priority, but 100% inspection of container freight has long been controversial. If 100% inspection turns out to be infeasible or undesirable due to cost considerations, Bier and Haphuriwat suggest that secrecy could potentially be used to achieve effective attacker deterrence with only partial inspection.(35)

5. OTHER REASONS FOR SECRECY AND DECEPTION

Disclosure of defenses could make highly defended targets inherently more attractive and prestigious to attackers, in which case secrecy could be preferable to disclosure of defenses. For example, in computer security, "script kiddies" may be willing to attack lightly defended targets, but more serious hackers would presumably gain prestige only by attacking well-defended targets. Similarly, in the Iraq War, the United States has largely been unable to announce "good news" like the rebuilding of schools or hospitals, since doing so has often invited immediate attacks; in other words, disclosure of good news apparently makes the newly rebuilt targets more attractive to attackers. While disclosure of rebuilding efforts is different from disclosure of defensive investments, it seems plausible that the two types of disclosure could have similar effects (i.e., leading to increased attacker target valuations) in some cases.

Moreover, in most work to date, the cost of implementing truthful disclosure is assumed to be lower than the costs of secrecy and deception (since defenders usually must spend some extra effort to keep their actions secret, or to deceive their opponents).(30,36) Under this assumption, secrecy and deception will be preferred only because of their endogenous effects on attacker behavior (if at all). However, in some situations, the costs of secrecy or deception could actually be lower than the cost of truthful disclosure; for example, if truthful disclosure is politically costly. In such cases, secrecy and deception would naturally be preferred.

6. CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

In this article, we have explored some reasons that a defender might prefer secrecy or deception about her defensive resource allocations, rather than truthful disclosure, in a homeland-security context. When the defender has private information (i.e., when the attacker does not know some defender characteristics, such as asset vulnerabilities), secrecy and/or deception can be strictly preferred by some types of defenders in order to mimic defender types that are of less interest to attackers (e.g., defender types that are less vulnerable), or to distinguish themselves from defender types that are of greater interest to attackers. We also found that secrecy could increase the effectiveness of defensive investments in cases where the success probability of an attack is a nonconvex function of the level of defensive investment in order to avoid the effects of diminishing marginal returns and achieve greater overall security. In general, secrecy and deception seem to be desirable mainly when defense is only marginally justified—not so cost effective as to be clearly justified, and not so ineffective as to be clearly not worthwhile.

However, equilibriums involving secrecy and deception were surprisingly rare and difficult to obtain (e.g., occurring only for narrow ranges of parameter values, compared to the frequency with which secrecy and deception are observed in practice). We suspect this may be because of some of the more unrealistic assumptions of game theory (e.g., common knowledge, full rationality). Therefore, it would be worthwhile to develop models of optimal strategies for rational defenders when facing nonstrategic (but behaviorally realistic) players.

ACKNOWLEDGMENTS

This research was supported by the U.S. Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events under grant number 2007-ST-061-000001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the U.S. Department of Homeland Security. We also thank Area Editor Yacov Y. Haimes (University of Virginia) and three anonymous referees for their helpful comments.

REFERENCES

1. Bier VM. Game-theoretic and reliability methods in counterterrorism and security. Pp. 17–28 in Wilson A, Limnios N, Keller-McNulty S, Armijo Y (eds). Mathematical and Statistical Methods in Reliability, Series on Quality, Reliability and Engineering Statistics. Singapore: World Scientific, 2005.
2. Zhuang J, Bier VM. Balancing terrorism and natural disasters—Defensive strategy with endogenous attacker effort. Operations Research, 2007; 55(5):976–991.
3. Levy G. Decision making procedures for committees of careerist experts. American Economic Review, 2007; 97(2):306–310.
4. Sandler T, Arce DG. Terrorism and game theory. Simulation & Gaming, 2003; 34:319–337.
5. Bier VM, Oliveros S, Samuelson L. Choosing what to protect. Journal of Public Economic Theory, 2007; 9(4):563–587.
6. Edmonds S. Porter Goss' Op-Ed: "Ignoturn per Ignotius"!, 2006. Available at: http://www.truthout.org/article/sibel-edmonds-porter-gosss-op-ed-ignoturn-ignotius, Accessed March 2010.
7. Joint Chiefs of Staff. Joint Publication 3–13.4. Joint Doctrine for Military Deception, 1996. Available at: http://www.c4i.org/jp3_13_4.pdf, Accessed March 2010.
8. DePaulo BM, Wetzel C, Sternglanz RW, Wilson MJW. Verbal and nonverbal dynamics of privacy, secrecy, and deceit. Journal of Social Issues, 2003; 59(2):391–410.
9. Swire PP. ArXiv Computer Science E-Prints. What Should be Hidden and Open in Computer Security: Lessons from Deception, the Art of War, Law, and Economic Theory, 2001. Available at: http://arxiv.org/ftp/cs/papers/0109/0109089.pdf, Accessed March 2010.
10. Swire PP. A model for when disclosure helps security: What is different about computer and network security? Journal on Telecommunications and High Technology Law, 2004; 2:1–38.
11. Brams SJ. Deception in 2×2 games. Journal of Peace Science, 1977; 2:171–203.
12. Brams SJ, Zagare FC. Deception in simple voting games. Social Science Research, 1977; 6:257–272.
13. Merriam-Webster, Inc. Merriam-Webster Online Dictionary, 2009. Available at: http://www.m-w.com/dictionary/secrecy, Accessed March 2010.
14. Bok S. Secrets: On the Ethics of Concealment and Revelation. New York: Pantheon Books, 1982.
15. Tefft SK (ed). Secrecy: A Cross-Cultural Perspective. New York: Human Sciences Press, 1997.
16. Merriam-Webster, Inc. Merriam-Webster Online Dictionary, 2009. Available at: http://www.m-w.com/dictionary/deception, Accessed March 2010.
17. Hendricks K, McAfee P. Feints. Journal of Economics & Management Strategy, 2006; 15(2):431–456.
18. Oliveros S. Equilibrium bluffs: A model of rational feints. University of California-Berkeley, Haas School of Business, Working paper, 2010.
19. Brown G, Carlyle M, Diehl D, Kline J, Wood K. A two-sided optimization for theater ballistic missile defense. Operations Research, 2005; 53(5):263–275.
20. Board OJ. The deception of the Greeks: Generalizing the information structure of extensive form games. Greek Economic Review, 2002; 22:1–16.
21. Rourke FE. Secrecy and Publicity: Dilemmas of Democracy. Baltimore, MD: Johns Hopkins Press, 1961.
22. Schneier B. Secrets and Lies: Digital Security in a Networked World. Hoboken, NJ: John Wiley & Sons, 2000.
23. Maskin E, Tirole J. The politician and the judge: Accountability in government. American Economic Review, 2004; 94(4):1034–1054.
24. Shapiro JN, Siegel DA. Is this paper dangerous? Balancing secrecy and openness in counterterrorism. Security Studies, 2010; 19(1):66–98.
25. Zhuang J. Modeling Secrecy and Deception in Homeland Security Resource Allocation, Ph.D. dissertation, Department of Industrial and Systems Engineering, University of Wisconsin–Madison, 2008.
26. Hao M, Jin S, Zhuang J. Robustness of optimal defensive resource allocations in the face of less than fully rational attackers. Pp. 886–891 in Proceedings of the 2009 Industrial Engineering Research Conference, 2009.
27. Brown GG, Carlyle WM, Harney RC, Skroch EM, Wood RK. Interdicting a nuclear-weapons project. Operations Research, 2009; 57(4):866–877.
28. Powell R. Allocating defensive resources with private information about vulnerability. American Political Science Review, 2007; 101(4):799–809.
29. Hausken K, Levitin G. Protection vs. false targets in series systems. Reliability Engineering & System Safety, 2009; 94(5):973–981.
30. Zhuang J, Bier VM. Secrecy and deception at equilibrium, with applications to anti-terrorism resource allocation. Defence and Peace Economics, 2010.
31. Dighe NS, Zhuang J, Bier VM. Secrecy in defensive allocations as a strategy for achieving more cost-effective attacker deterrence. International Journal of Performability Engineering, 2009; 5(1):31–43.

32. Ayres I, Levitt S. Measuring the positive externalities from unobservable victim precaution: An empirical analysis of Lojack. Quarterly Journal of Economics, 1998; 113(1):43–77.
33. Hudson A. Washington Times. Air Marshals—Secrecy Ruined by Dress Code, July 9, 2004. Available at: http://www.washingtontimes.com/news/2004/jul/09/20040709–121013-3063r/, Accessed March 2010.
34. Frank T. Cover blown, but air marshal still flies. USA Today, June 1, 2006.
35. Bier VM, Haphuriwat N. Analytical method to identify the number of containers to inspect at U.S. ports to deter terrorist attacks. DOI: 10.1007/510479-009-0665-6. Available at: http://www.springerlink@com/content/817ju81220214027/Year=2009.
36. Zhuang J, Bier VM, Alagoz O. Modeling secrecy and deception in a multiple-period attacker-defender signaling game. European Journal of Operational Research, 2010; 203(2):409–418.