Risk Analysis, Vol. 24, No. 4, 2004

Thinking about How Many Guards Will Do the Job

Matthew Bunn*

1. INTRODUCTION

Scott Sagan has provided a fascinating theoretical treatment of the problem of how many guards to assign to protecting nuclear facilities (Sagan, this issue). Sagan's analysis clearly demonstrates that it is wrong to think about the problem as though the guards were redundant technical components of a system who do not affect each other's performance. Although Sagan does not draw out the broader implications here, the issues he raises arise in a far wider array of situations—in virtually any system where there are decisions to be made as to how many people to assign to security-critical or safety-critical tasks.

Indeed, there are several additional reasons, not mentioned by Sagan, to believe that guards' per-unit effectiveness may be reduced as more guards are added, particularly if the addition is done in a hurry:

- If the increase in number of guards on duty is accomplished simply by making the existing guard force work longer hours, the guards may be overtired and the imposition may reduce their morale.
- Pressure to hire more guards quickly may lead to reduced intensity of background checks, lowered standards for accepting individuals into the guard force (particularly as large numbers of facilities compete with each other for potential recruits in order to meet new standards), and decreased training time—all of which could reduce average individual effectiveness (and, in the case of background checks, increase the danger of ending up with an insider attacker as a member of the guard force).
- Facilities concerned about the impact on profits resulting from increases in the guard force may have more incentive to hold down salaries and benefits, which could reduce guards' crucial belief that they are valued members of a team working toward a common objective.

There is some evidence, based on interviews with members of guard forces and other sources of information, that all three of these effects have in fact been occurring at some U.S. commercial nuclear power plants since the Nuclear Regulatory Commission (NRC) ordered increased security measures following the 9/11 attacks (Project on Government Oversight, 2002).

As Sagan acknowledges, however, his arguments (and, by extension, the additional ones just offered) only prove that the problem of guard force size needs to be thought through carefully; they do not demonstrate that there would not be a substantial security benefit from adding more guards at nuclear plants and other dangerous industrial or military facilities in the United States and elsewhere. In short, Sagan provides a warning, but does not actually recommend that guard forces should not be increased. To determine when and where guard forces should be increased, three key issues in particular would need to be addressed:

- Whether there are factors that would point toward *increased*, rather than decreased, per-unit effectiveness as the number of guards is increased;
- Whether the factors Sagan has identified, or other factors reducing per-unit effectiveness with increasing guard force size, are in fact strong enough, at the particular points on the curve of guard force size vs. effectiveness that now exist, to greatly reduce the value of adding additional guards (or even to make the change

* Managing the Atom Project, Belfer Center for Science and International Affairs, Kennedy School of Government, Harvard University, Cambridge, MA 02138, USA; matthew_bunn@harvard.edu.

0272-4332/04/0100-0949$22.00/1

from adding more guards negative rather than positive); and
- Whether a percentage "effectiveness" in defending against an unspecified (and presumably unchanging) threat—which by definition cannot be higher than 100%—is the right way to think about the problem.

Fortunately, in making practical policy decisions regarding the number of guards to assign to different types of critical facilities, the world does not have to rely solely on reasoning from first principles: it is possible to collect limited, but nonetheless helpful, real-world data on the performance of guard forces as a function of their size, which can contribute to identifying the most effective approaches to securing such facilities. More broadly, of course, as Sagan's article makes clear, a broad range of technical and organization factors in addition to guard force size affect the capability of security systems to provide protection against particular levels of threat.

2. FACTORS THAT MAY LEAD TO INCREASED PER-UNIT GUARD EFFECTIVENESS

While Sagan identifies a number of factors that might lead to reduced per-unit guard effectiveness as the size of guard forces increased, there are also factors that might point in the opposite direction.

For example, currently the guards at some U.S. nuclear power plants feel that if terrorists ever attacked, they would be hopelessly outgunned, and there would be no point in sacrificing their lives in a hopeless cause—therefore, they say they would probably run, and they expect that their colleagues would as well (Project on Government Oversight, 2002). If the guard forces were increased (in numbers and capabilities) to the point that they felt they would definitely be able to fight off a plausible terrorist attack successfully, one would expect they would be much more likely to stand and fight. The military literature is replete with cases of forces that saw themselves as completely outnumbered and outgunned breaking and running—in Sagan's terms, having their per-unit effectiveness collapse to zero. (The more unusual case, in which a seriously outgunned unit decides to stand and fight, is the subject of countless books and movies.) In this case, it is the *perception* held by the guards of the likely balance between the defense force's capabilities and the capabilities of terrorists who might attack that is critical to effectiveness.

Similarly, while Sagan points out correctly that having more guards increases the chance of having someone cooperating with the terrorists among the guards, having more guards may also make it possible to deal with such an insider threat more effectively. With a small group, for example, there might not be enough people to allow the use of a "two-man-rule" on critical guard functions, whereas with a larger group, there would be enough people to ensure that where a single insider could cause a catastrophe, two (or more) people were always present. (Such a two-man rule is routinely enforced in handling nuclear weapons; indeed, Russia's nuclear weapons guard forces employ a three-man rule.) Along the same lines, while one insider might represent a catastrophic "common-mode" failure for a small guard force, as Sagan describes, it probably would not for a larger guard force, which had more dominant superiority over the attacking terrorists. If the guards are five people fighting three attacking outsiders, and one of the five defenders turns out to be on the attackers' side, the guard force is likely to be in serious trouble, but if the guard force has 20 people fighting the same number of outsiders, a single turncoat among the defenders should not be able to make much difference in the outcome.

3. QUESTIONS ON THE STRENGTH OF THE FACTORS

Sagan uses a number of figures to demonstrate what the effect of the negative factors he describes would be *if* certain strengths of those negative factors are assumed. But there is no data publicly available on what the strengths of these factors really are, at various guard force sizes. The reality is that we do not know what the shape of the curve of "guard force effectiveness as a function of size" is—and that shape may well be different for different types of facilities, different national and organizational cultures, and so on.

Fig. 1 shows the impact of changes in Sagan's assumptions. Sagan's "shirking" model is a subtractive one, in which adding each additional guard subtracts 15% from the effectiveness of each member of the guard force. Total guard force effectiveness collapses to zero at six guards, the point at which individual effectiveness also reaches zero (having started at 75%). Simply modifying this to a multiplicative model with the same percentage—that is, adding each additional guard multiplies the individual effectiveness of each guard force member by one

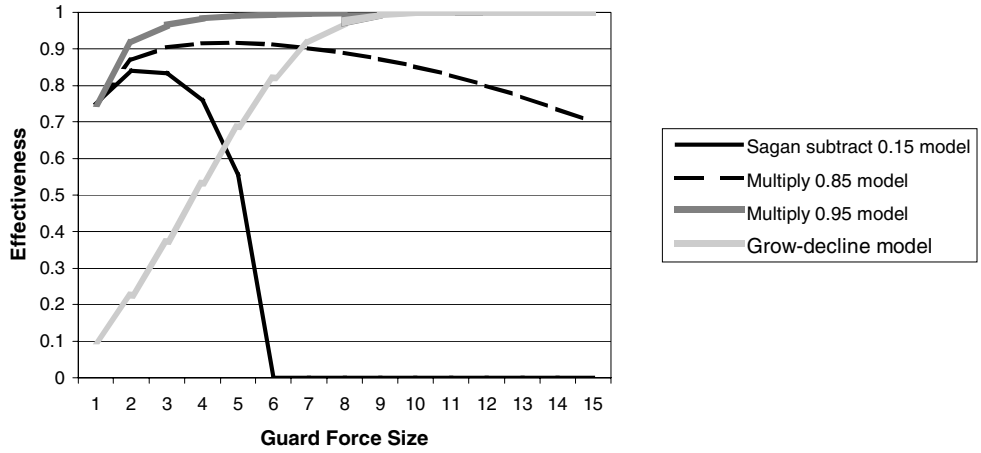

Fig. 1. Effect of varying assumptions on unit effectiveness as a function of guard force size.

minus 15 percent—changes the picture dramatically. As shown in the dotted line in Fig. 1, effectiveness now rises far above its peak in Sagan's model, and while it declines eventually, it never collapses. Such a multiplicative model is more plausible than a subtractive one, as the existence of large numbers of facilities around the world secured with large numbers of guards suggests that those with experience securing such facilities do not believe that individual effectiveness collapses to zero at just a handful of guards.

If the multiplier is 0.95—so that each additional guard reduces his or her colleagues' effectiveness by only 5%—then effectiveness rises rapidly to essentially 100% and remains there until the guard force size is well off this chart. This is shown in the dark gray line in Fig. 1. Finally, one could imagine models that were more dramatically different from Sagan's. For example, one might assume that the effectiveness of a single guard in protecting a large nuclear facility would be very low, and this per-unit effectiveness would improve as more guards were added (for the reasons alluded to above), until eventually it reached a peak and began to decline, as the positive per-unit benefits of adding more guards began to be overwhelmed by the negative impacts Sagan describes. The light gray line shows a version of such a model in which initial per-unit effectiveness is 10%, which is multiplied by 1.2 for each member of the guard force until the force reaches 10 members, at which point it begins to decline, being multiplied by 0.95 with the addition of each further guard force member. As can be seen, this model suggests a dramatic security improvement from adding more guards, going quickly from only 10% to almost 100% at nine members (a point at which individual effectiveness is still only 43%).

Obviously, these curves have very different shapes, with very different implications. Key questions to ask include: In the real world, is the shape of this graph such that there is a peak, after which effectiveness declines (as in the first two lines of Fig. 1), or a level at which the curve flattens out and additional guards offer little benefit (as in the second two lines)? And if so, how big are current guard force sizes compared to the size at which further increases would offer little benefit (or might even reduce effectiveness)? In other words, is there an elbow in the curve, and are we there yet? From first principles alone, we have no basis for judging the answers.

4. CONCEPTUAL PROBLEMS WITH "EFFECTIVENESS" AGAINST AN UNSPECIFIED THREAT

The conceptual structure of Sagan's analysis is based around the notion of individual "effectiveness" combining to overall unit "effectiveness" in dealing with an unspecified threat. While this approach offers useful insights, it may not be the best way to conceptualize the problem. Guard force "effectiveness" depends on the threat the force is supposed to cope with: a guard force that is highly effective against three well-armed and well-trained attackers may be grossly insufficient against 10. Hence, rather than being conceptualized as increasing to an "effectiveness" close to 100% against some fixed threat, security could be conceptualized as increasing continuously—a larger guard force might not provide much additional

protection against a small threat, but would offer protection against larger threats that would not be possible for a smaller force.

The way this is usually addressed in nuclear security systems is to define a specific magnitude of threat that security systems (including, as one component, the guard force) are required to be able to defeat, and then require that facilities have systems designed to have high effectiveness in defeating that "design basis threat" (DBT).[1] The DBT is defined both in numbers and in capabilities of likely attackers. The DBT, a facility's defenses must meet depends on the dangers the facility poses: U.S. nuclear weapons facilities, for example, have a larger DBT (and hence larger guard forces) than do commercial nuclear power plants, and some lower-risk nuclear facilities have no regulatory DBT at all. Setting the size and capabilities of the DBT is inevitably a political process, balancing how much facility operators will be required to pay for security against how much residual risk society will be required to accept of an attack the security system will be unable to defeat. After a highly contentious process, for example, in April 2003 the NRC finally increased the DBT for U.S. nuclear power plants, following the 9/11 attacks (Nuclear Regulatory Commission, 2003); while the new DBT is classified, it is reported to be far smaller than the four teams of 4–5 well-trained and suicidal attackers that struck on September 11, 2001. As a general rule, the larger the DBT, the larger the guard force required to be able to be able to defeat it reliably.

In the U.S. system, the adequacy of the combination of security hardware and guard forces put in place for a particular facility is assessed by two principal means. The first is computer modeling, which is used for both designing new security systems and evaluating existing ones. For example, a commonly used software package called ASSESS (the Analytical System and Software for Evaluating Safeguards and Security) includes a subcomponent (called BATTLE) that specifically models battles between terrorist attacking forces and guard forces to assess the probability that the guard force will succeed in fighting off an attacking force of a particular size (Sandia National Laboratory, personal communications, 2002). This subcomponent is based on actual data from past small-unit military engagements; it is not explicitly designed to model factors that may reduce or increase per-unit guard effectiveness (though if such factors affected the performance of the small units in the data set, that would be included in the model). Hence the predictable result is that in ASSESS simulations, adding more guards increases guard force effectiveness—and, as one would expect, if the postulated attacking force is larger, the guard force must also be larger to have a good chance of winning the engagement.

The second method is exercises and performance tests, in which mock attackers actually attempt to break into the system; such tests could be used to provide at least limited data on the real-world impact of factors such as those described in Sagan's article on guard force effectiveness as a function of size, as described in more detail below.

5. SECURITY: MORE THAN JUST NUMBERS

As Sagan's article makes clear, the effectiveness of a guard force (and of the overall security system of which it is a part) involves far more than the number of guards on duty at any given time. The system must include effective measures to detect, delay, and defeat potential attackers; after all, if attackers are never detected, or are not delayed and so are able to finish their attack before the defenders can respond, the best guard force in the world will not solve the problem. More broadly, however, security is a problem of organizational effectiveness—both among the guard force and among all the other personnel at the facility whose roles may be critical to achieving high security. High security, like high safety, is something that results from both technology and organizational culture; to engineer safety or security one therefore has to think through the human-behavior and organizational-response factors that affect safety and security.[2]

[1] Designing, building, and operating a system that can reliably beat a specified DBT tends to be more difficult to do than it sounds, because rather than protecting against random natural hazards and human errors (as one does to ensure the safety of a system), one has to protect against intelligent human beings, who may think of an attack strategy that did not occur to the designers of the defensive system—and one has to fight against complacency among those guarding against events that virtually never occur. Thus, a considerable amount of creativity is required in designing effective security systems, and mechanisms such as "red teams" pretending to be attackers are often used to help identify potential weaknesses. For a useful introduction to these issues, see Garcia (2001).

[2] As far as the author is aware, no in-depth analysis of the organizational factors that contribute to high levels of security has yet been published. An outstanding summary of the literature on how organizational factors contribute to or undermine *safety* can be found in Reason (1997).

6. COLLECTING REAL-WORLD DATA

Fortunately, the world does not have to rely solely on theoretical arguments that seem to point in many different directions. While it is impossible to collect much "real" data on the ability of guard forces of various sizes to fight off terrorists (since such battles do not happen very often), it is very possible to collect data from tests and exercises.

Both U.S. nuclear power plants regulated by the NRC and Department of Energy (DOE) facilities (whose security is regulated by an internal DOE body) are required to undergo occasional "performance tests" to assess the ability of their security systems and guard forces to defeat an attack corresponding to their DBT. To avoid people actually being shot during the tests, both the attackers and the defenders are equipped with nonlethal weapons similar to (but more sophisticated and realistic than) those used for "laser tag." Hence there is always some lack of realism in the tests, since the defenders are well aware of when the test is occurring, and the fact that it is only a test. Nevertheless, these tests have proved to have enormous value, including in identifying weaknesses, which are then addressed (Bukharin, 2000; Bukharin *et al*., 2000). Informed by a theoretical understanding of organizational factors to look for—and of which factors (such as social shirking) are likely to be minimized by the artificial circumstances of the tests—it would be very possible to collect data from the performance of actual guard forces in such tests to help resolve the questions Sagan raises.

7. RECOMMENDATIONS

From these considerations, several recommendations can be drawn:

- Assessments of how many guards are required at particular dangerous facilities should be informed by an understanding of the ways in which force size may affect organizational effectiveness (and of the other major factors that relate to organizational effectiveness).
- Additional data should be drawn from ongoing performance tests to assess the effect of force size (and other key organizational factors) on guard force effectiveness, and this data should be factored into computer models and other tools used to help determine appropriate guard force sizes for particular facilities. Exercise design for this purpose should be done by teams independent of the actual guard forces, who are unlikely to be willing to seriously contemplate that factors such as social shirking could significantly undermine their performance.
- Nuclear facilities and other especially dangerous facilities should be required to be able to defend against DBTs comparable to the capabilities terrorists have already demonstrated in recent attacks, using security systems and guard forces designed and sized with tools incorporating the best theoretical understandings and data available at the time.

REFERENCES

Bukharin, O. (2000). Physical protection performance testing: Assessing U.S. NRC experience. *Journal of Nuclear Materials Management*, Summer, 21–27.

Bukharin, O., Bunn, M., & Luongo, K. N. (2000). *Renewing the Partnership: Recommendations for Accelerated Action to Secure Nuclear Material in the Former Soviet Union*. Washington, DC: Russian-American Nuclear Security Advisory Council. Available at http://bcsia.ksg.harvard.edu/BCSIA_content/documents/mpca2000.pdf, May 28, 2003.

Garcia, M. L. (2001). *The Design and Evaluation of Physical Protection Systems*. Woburn, MA: Butterworth-Heinemann.

Nuclear Regulatory Commission. (2003). *NRC Approves Changes to the Design Basis Threat and Issues Orders for Nuclear Power Plants to Further Enhance Security*. Washington, DC: NRC. Available at http://www.nrc.gov/reading-rm/doc-collections/news/2003/03-053.html, May 28, 2003.

Project on Government Oversight. (2002). *Nuclear Power Plant Security: Voices from Inside the Fences*. Washington, DC: Project on Government Oversight. Available at http://www.pogo.org/p/environment/eo-020901-nukepower.html, May 28, 2003.

Reason, J. (1997). *Managing the Risks of Organizational Accidents*. Aldershot, UK: Ashgate.

Sagan, S. D. (this issue). The problem of redundancy problem: Why more nuclear security forces may produce less nuclear security. *Risk Analysis*.