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

The Problem of Redundancy Problem: Why More Nuclear Security Forces May Produce Less Nuclear Security†

Scott D. Sagan*

1. INTRODUCTION

After the September 11, 2001 terrorist attacks on the World Trade Center and the Pentagon, many scholars, journalists, and public officials expressed fear about the security of nuclear facilities in the United States. Terrorists could attack military bases, weapons in transit, or nuclear weapons production and dismantlement plants in order to steal a weapon or its components. Terrorists might attack nuclear power reactors, nuclear materials storage sites, nuclear waste transportation vehicles, or nuclear research facilities, with two basic motives in mind: to cause a conventional explosion, spreading radioactive materials in the area; or to seize the nuclear materials, which could be used for building either a dirty bomb (a radiological weapon) or, conceivably, a primitive nuclear bomb. These fears were highlighted in President George Bush's January 2002 State of the Union address, in which he reported that diagrams of American nuclear plants were discovered in al-Qaeda hideouts in Afghanistan.(1) Senior U.S. intelligence officials also revealed that Osama bin Laden had sent operatives to try to purchase stolen nuclear materials and that there was "pretty convincing evidence" that al-Qaeda operatives had been "casing" nuclear power plants in the United States prior to the September 11 attacks.(2) In January 2002, U.S. intelligence agencies issued a warning, based on an interrogation of a captured terrorist, of a possible attack on a nuclear power plant or Department of Energy (DOE) nuclear weapons facility.(3) Then, in June 2002, the Justice Department announced that it had arrested an American citizen who had joined al-Qaeda in Pakistan and was sent back to the United States to develop and execute a plan to seize nuclear materials and use them in a radiological bomb attack.(4)

The reaction to this new terrorist threat has been strong and predictable. Emergency efforts began immediately to deploy more security forces to protect U.S. nuclear facilities after the September 11 attacks. In October 2001, the Nuclear Regulatory Commission (NRC), according to one of its members, advised U.S. state governors to "request additional security patrols or posts, using local law enforcement, state police or National Guard if needed, in addition to all their own people" to provide extra protection for nuclear power plants.(5) In March 2002, the Secretary of Energy requested over $138 million in supplementary funds to hire new guard forces and provide better physical security for DOE nuclear facilities.(6) Concerns about terrorism were further increased when it was revealed that in DOE training exercises, mock terrorists had successfully stolen or seized plutonium and other sensitive nuclear materials from facilities at Los Alamos and at Rocky Flats.(7) The President's Foreign Intelligence Advisory Board raised similar alarm bells when it reported that it had taken DOE officials 35 months "to write a work order to replace a lock at a weapons lab facility containing sensitive nuclear information" and 45 months "to correct a broken doorknob that was sticking in an open position and allowing access to sensitive sites."(8)

Experts testifying to Congress have strongly advocated adding more security guards and patrols at nuclear facilities to prevent nuclear terrorism. For example, Paul Leventhal, the President of the Nuclear Control Institute, recommended to Congress that the regulatory minimum of five guards per site be increased to a number capable of defeating a terrorist attack in excess of the 19 terrorists involved in the

† Winner of Columbia University's Institute for War and Peace Studies 2003 best paper in Political Violence prize.

* Center for International Security and Cooperation, Encina Hall, Stanford University, Stanford, CA 94301-6165; ssagan@stanford.edu.

0272-4332/04/0100-0935$22.00/1

September attacks, and other security experts recommended to Congress that security forces as large as 30–40 guards be stationed at each nuclear plant or weapons facility in the United States.[9] The increased threat of nuclear terrorism, it is argued, must be met with a countervailing increase in nuclear security personnel.

There are understandable incentives for organizational leaders to want to devote more resources and more personnel to address dangerous problems when they are seen to develop. From a political perspective, action must be taken after a major disaster, at a minimum, to let insiders and outsiders see that top officials are doing something to prevent a reoccurrence. If the causes of the problem are uncertain, however, the appropriate reaction is unclear. This article analyzes how we should think about nuclear security and the emerging terrorist threats. It presents a warning about the most simple, and most tempting, solution to our new nuclear terrorism problem: to add more security forces to protect power plants, weapons facilities, and nuclear storage sites. The article uncovers the dark side of redundancy by focusing on how efforts to improve nuclear security can inadvertently backfire, increasing the risks they are designed to reduce.

2. REDUNDANCY AND RELIABILITY

The use of redundancy in its many forms is a common strategy used to make more reliable systems out of inherently imperfect parts. Redundancy theory in engineering demonstrates how even unreliable components, *if* independent and connected in a parallel manner, can lead to rapid increases in overall system reliability. A large number of social scientists and security analysts have, therefore, called for the widespread use of redundancy as one of the necessary requirements of "high reliability organizations."[10] Todd La Porte, for example, argues that high reliability organizations "are characterized especially by flexibility and redundancy in pursuit of safety and performance."[11] Jonathan Bendor notes that "the most basic argument for redundancy rests on the practicality of increasing a system's reliability without increasing the reliability of its constituent elements" and provides a compelling analogy:

> Suppose an automobile had dual breaking (sic) circuits: each circuit can stop the car, and the circuits operate independently so that if one malfunctions it does not impair the other. If the probability of either one failing is 1/10, the probability of both failing simultaneously is $(1/10)^2$, or 1/100. Add a third independent circuit and the probability of the catastrophic failure of no brakes at all drops to $(1/10)^3$, or 1/1,000.[12]

In short, according to "high reliability theorists," organizations that seek to transcend Murphy's Law must adopt the following motto: if it can go wrong, it will go wrong, and therefore every "it" must have a redundant backup so that the whole system does not fail.

The beauty of redundancy is illustrated in Fig. 1. It should, therefore, come as no surprise that the use of redundancy is so common in organizations that manage hazardous technologies such as U.S. Navy aircraft carriers,[13] nuclear weapons command and control,[14] critical computer software,[15] and the air

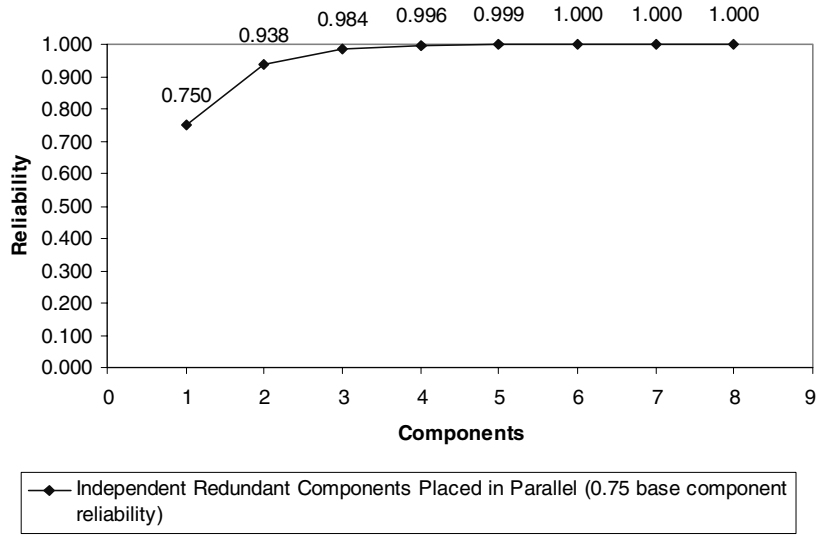

Fig. 1. The benefits of redundancy for reliability.

traffic control network.[16] As long as components are imperfect, independent and parallel redundancies can reduce the risks of system failure to extraordinarily low levels. Indeed, in Fig. 1, system reliability was so high that Excel rounded it up to 1.000 reliability.

In contrast, scholars in the "normal accidents theory" school have argued that organizations that exhibit both high degrees of interactive complexity and tightly coupled operations will suffer serious accidents despite their efforts to maintain high reliability and safety.[17] Complexity leads to hidden failure modes that no one predicts and that no one identifies quickly when they produce incidents; tight coupling means that when one thing goes wrong, others do quickly since there is little slack in the system to continue safe production. These "normal accident" theorists have further argued that since adding redundancy can increase the complexity of a system, efforts to increase safety and security through the use of redundant safety devices may actually backfire, inadvertently making systems fail more often.[18]

This article further develops this theory, both by providing new arguments about potential counterproductive effects of redundancy and by presenting new empirical examples of the problem of redundancy problem. Three serious problems are analyzed: (1) the catastrophic common-mode error problem; (2) the social shirking problem; and (3) the overcompensation problem. Each will be examined in turn, with presentations of the logic of the argument, as well as empirical examples. I will then discuss the implications of each of these problems with redundancy for the crucial current policy question: Will more nuclear security forces provide improved nuclear security against terrorist threats?

3. REDUNDANCY BACKFIRES THROUGH COMMON-MODE ERRORS

The first problem with redundancy is that adding extra components can inadvertently create a catastrophic common-mode error (a fault that causes all the components to fail). In complex systems, independence in theory (or in design) is not necessarily independence in fact. As long as there is some possibility of unplanned interactions between the components leading to common-mode errors, however, there will be inherent limits to the effectiveness of redundancy as a solution to reliability problems. The counterproductive effects of redundancy when extra components present even a small chance of producing a catastrophic common-mode error can be dramatic, as illustrated in Fig. 2 below.

This danger is perhaps most easily understood through a simple example from the commercial aircraft industry. Aircraft manufacturers have to determine how many engines to use on jumbo jets. Cost is clearly an important factor entering their calculations. Yet so is safety, since each additional engine on an aircraft both increases the likelihood that the redundant engine will keep the plane in the air if all others fail in flight and increases the probability that a single engine will cause an accident, by blowing up or starting a fire that destroys all the other engines and the aircraft itself. In Fig. 2, I assume that 40% of the time that each engine fails, it does so in a way (such as starting a catastrophic fire) that causes all the other

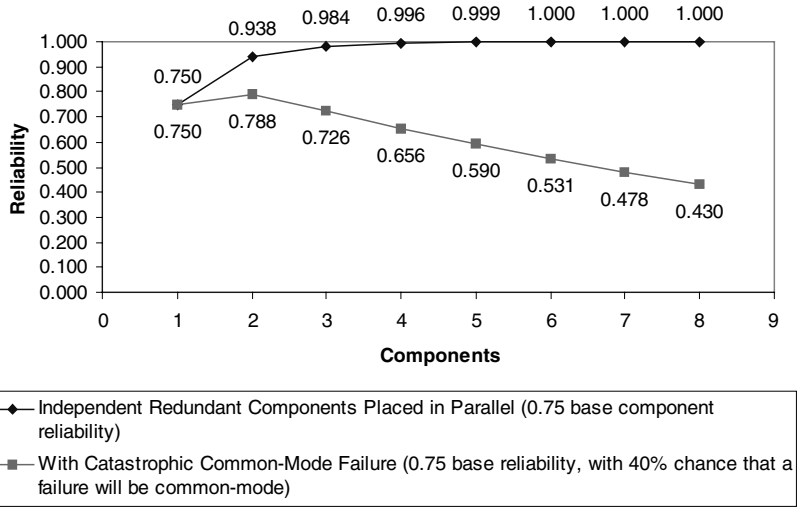

Fig. 2. Catastrophic common-mode errors.

engines to fail as well. Aircraft manufacturers make similar calculations in order to estimate how many engines would maximize safety. Boeing, for example, used such an analysis to determine that, given the reliability of modern jet engines, putting two engines on the Boeing 777, rather than three or more engines as exist on many other long-range aircraft, would result in lower risks of serious accidents.[19]

In more complex systems or organizations, however, it is often difficult to know when to stop adding redundant safety devices because of the inherent problem of predicting the probabilities of exceedingly rare events. It should, therefore, not be surprising that many serious accidents with hazardous technologies, even in organizations that appear to be highly reliable in general, are caused by redundant safety devices designed to reduce such risks. The "safety device accident" phenomenon has perhaps been most often witnessed in nuclear power plant accidents. The October 1966 near-meltdown accident at the Fermi reactor near Monroe, MI, for example, was caused by an emergency safety device, a piece of zirconium plate, that had been placed inside the reactor to reduce risk that materials from the core would burn through the containment walls in an accident. The zirconium plate broke off, however, and blocked a pipe, stopping the flow of coolants into the reactor core. To make matters worse, this safety device had been installed at the last stage in the construction of the reactor and was therefore *not* on the final "as built" set of blueprints. The power plant operators during the accident, therefore, could not figure out what was blocking the flow of coolants.[20]

3.1. The Insider Threat as a Common-Mode Failure

This kind of problem exists whenever you add security forces to protect a critical site. Who should guard the guardians? If there is any danger of an "insider threat"—that is, a new guard being the terrorist one is trying to protect against—then at some point adding redundancy can backfire. Unfortunately, organizations that pride themselves on high degrees of personnel loyalty can be biased against accurately assessing and even discussing the risk of insider threats and unauthorized acts.

A dramatic case in point is the failure to remove Sikh bodyguards from Indira Gandhi's personal security unit after she had instigated a violent political crackdown on Sikh separatists in 1984. Increased security personnel were deployed at the Prime Minister's residence after a series of death threats were made against the prime minister and her family. According to H. D. Pillai, the officer in charge of Gandhi's personal security, "the thrust of the reorganized security . . . was to prevent an attack from the outside": "What we did not perceive was that an attempt would be made inside the Prime Minister's house."[21] When it was suggested by other officials that Sikh bodyguards should be placed only on the outside perimeter of the Prime Minister's compound, Mrs. Gandhi felt that this could not be done without damaging her political reputation: "how can I claim to be secular if people from one community have been removed from within my own house," she claimed.[21] Two Sikh guards—one long-standing bodyguard for Gandhi, and the other a new guard added given the emergency—conspired together and assassinated Mrs. Gandhi on October 31, 1984.

How many guards should we have at nuclear facilities? If the problem is only one of maximizing the probability that at least one guard will identify and disrupt a terrorist attack or an attempt to steal nuclear materials, then the obvious solution—an increased number of security personnel—will be appropriate. But calculations must also include an accurate assessment of insider threats. Are the leaders of the U.S. nuclear weapons facilities and nuclear power plants susceptible to Mrs. Gandhi's blindspot, an inability to see and deal with the insider threat problem?

There is both good and bad news here. On the positive side, the NRC "design basis threat" for protecting nuclear facilities prior to the September 11 attacks did include a consideration of potential insider threats. The minimum requirement for five guards to be deployed at each nuclear research or power facility was based on an assumption of three terrorists "along with a single insider capable of participating in a violent attack." There were thus four terrorists in the model threat and the NRC, therefore, set the regulation at five security guards to make sure that the guard units had a one-man (or woman) measure of superiority.[22]

On the negative side, however, this NRC design basis threat ignored the possibility that the added guards themselves might be the insider threat. Leaders of the nuclear power industry and regulators insist that the insider threat problem is not a serious one since security guards and others with access to critical areas in nuclear facilities are, it is claimed, thoroughly vetted through intense background checks, random drug and alcohol tests, and security management

programs, like the Continuous Behavior Observation Program, which ensures that supervisors and colleagues will report on any suspicious behavior.[23] Yet there are several reasons to question whether such programs have eliminated the insider threat problem. First, the criteria used to assess suspicious behavior are suspicious. For example, security personnel of at least one nuclear weapons facility were known to have ties with members of anti-government right-wing militia groups. After the head of operations of the Rocky Flats nuclear security force, a private security contractor, reported this to the DOE, however, the Secretary of Energy wrote: "It is not illegal for anyone to belong to a militia organization. Membership by itself in one of these organizations does not constitute a basis for denying a security clearance or employment in a security position."[24] Second, the DOE's Office of Nuclear Safety has reported in the past on several incidents of insider sabotage of nuclear safety equipment in the U.S. nuclear weapons complex, the perpetrators of which were never determined: these sabotage attempts include the cutting of electrical wires at the Idaho advanced test nuclear reactor and the loosening of hydrogen feed lines at a plutonium facility at the Los Alamos National Laboratory.[25] Third, the process of background checks and rules for unaccompanied access to critical areas in nuclear power facilities are by no means fool proof. The NRC acknowledges, for example, that full information is often lacking on the criminal or medical records of foreign-born applicants for unaccompanied access to nuclear sites, since inadequate records in the country of birth often exist and less than complete foreign government cooperation with the FBI is common.[26] Moreover, even after the September 11 attacks, the NRC continued to permit individual nuclear plant operators to grant temporary access for individuals to critical sites before the full FBI background screening was completed, and thus, as late as March 2002, an individual who was later discovered to have lied about his past criminal record was granted unaccompanied access to critical nuclear plant areas.[27]

In short, it cannot be safely assumed that nuclear security guard forces and other individuals inside nuclear facilities are immune to penetration by domestic or foreign terrorist organizations. Improving efforts to guard against insider threats through better screening of all personnel in nuclear facilities will be difficult and presents complex civil liberty questions. Yet simply ignoring that an insider threat problem exists will not make it go away.

4. REDUNDANCY BACKFIRES THROUGH SOCIAL SHIRKING

The second way in which redundancy can backfire is when diffusion of responsibility leads to "social shirking." This common phenomenon—in which individuals or groups reduce their reliability in the belief that others will take up the slack—is rarely examined in the technical literature on safety and reliability because of a "translation problem" that exists when transferring redundancy theory from purely mechanical systems to complex organizations. In mechanical engineering, the redundant units are usually inanimate objects, unaware of each other's existence. In organizations, however, we are usually analyzing redundant individuals, groups, or agencies, backup systems that are aware of one another. Such awareness clearly can influence each unit's reliability. Organizational theorists' reliance on engineering analogies can be highly misleading in this regard. Bendor's comparison of redundant actors to "dual breaking (sic) circuits" in automobiles, for example, focuses attention away from social interaction, since brakes in cars are not aware of one another. In other cases, social interaction is simply assumed to increase the reliability of each component: increasing competition is meant to force each unit to work harder. Under many circumstances, however, the opposite can be true: awareness of other redundant units can decrease system reliability if it leads an individual or subunit to shirk off unpleasant duties because it is assumed that someone else will take care of the problem.

Fig. 3 is a simple illustration of how adding redundancy can reduce system reliability if diffusion of responsibility inadvertently decreases component reliability. Imagine that the probability that any individual witness to a violent crime will call the police is 75% if he or she is the sole witness, but that every time another witness is added to the scene, the likelihood that each person will report the crime decreases by 15%. If each witness believes that his or her phone call is less necessary because of the presence of other witnesses, system reliability could decrease dramatically.

The existence of this kind of social shirking among individuals is well documented in the large social-psychology literature on "unresponsive bystanders." This research, sparked by the failure of 38 witnesses to report the 1964 murder of Kitty Genovese in Queens, has shown that individuals are *less* likely to report others' criminal behavior or to intervene in medical emergencies if they know that there are other witnesses

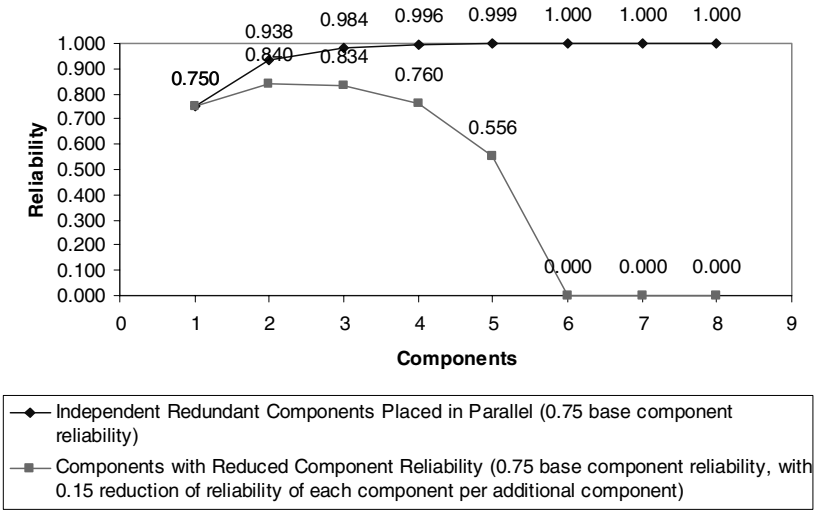

Fig. 3. Social shirking.

to the activity.[28] To give one dramatic example, in one so-called "Lady in Distress" experiment, 70% of the subjects who were on their own tried to get help when hearing a woman being assaulted in another room, while only 40% of subjects who knew that there was someone else also witnessing the attack responded at all.[29]

4.1. Diffusion of Responsibility in Military and Security Guard Units

Members and leaders of elite organizations that pride themselves on duty and responsibility would not easily acknowledge that one individual's awareness of others also doing a job could lead the first individual to become less reliable himself. Yet this social shirking phenomenon clearly occurs, even in elite military units. For example, two cases of redundancy backfiring by producing reduced component reliability were root causes of the April 1994 accidental shoot-down of two U.S. Army Black Hawk helicopters by two U.S. Air Force F-15 pilots in Iraq, as demonstrated in Scott Snook's book, *Friendly Fire*.[30]

The first case concerns the behavior of the crew in the Airborne Warning and Command Systems (AWACS) aircraft flying above the combat zone. There were 19 crew members aboard, each of whom knew that two U.S. helicopters were flying through the "no-fly zone" carrying 26 American and NATO peacekeepers. Yet, when the F-15 lead pilot below them announced on radio circuits that he spotted two helicopters, not one of the AWACS crew members alerted the F-15 pilots about the existence of the U.S. Black Hawks in the area. Why? As suggested in the following interviews from the investigation, diffusion of responsibility meant that everyone, and therefore no one, was responsible for doing the job:

> Investigator: Who's responsible on the AWACS aircraft for going through the procedures that the General [Major General Andrus] just described in trying to, in layman's terms, identify the hits there?
>
> Military Crew Commander: *Everybody is*.
>
> Investigator: Who has primary responsibility?
>
> Military Crew Commander: I would have everybody looking at it.
>
> Investigator: In the tactical area of operation on board the AWACS, who has command, control, and execution responsibilities for ATO tasked missions?
>
> Military Crew Commander: That's a very general question. The answer would be *everybody on position on the AWACS crew*.
>
> Investigator: Who is responsible for tracking helicopters that are tasked, according to the ATO, Air Tasking Order?
>
> Senior Director: No one is responsible . . . (emphasis added throughout) (pp. 119, 120, 126).[30]

The second example from this "friendly fire" incident is the fact that the two F-15 pilots were supposed to confirm that the unexpected helicopters they discovered in a "no-fly zone" in Iraq were indeed Iraqi "Hind" helicopters and not U.S. or allied helicopters. After the lead pilot called out that the helicopters were Hinds, he therefore sought confirming evidence from his wingman. This second pilot could not confirm the identity, but called out "Tally Two" to indicate that he saw two helicopters. According to the military investigation, "the F-15 flight lead understood his wingman's transmission to mean that he confirmed the identification" (vol. 1, p. 22).[31] He, therefore, issued an order to shoot down the helicopters. The F-15

wingman did not stop the engagement. His later testimony in court reveals the following:

> 120Q: Perhaps we should ask first, did you positively identify the helicopters?
>
> 120A: I never came out and said that they—positively ID'd as Hinds. I came in on that ID pass—I saw the high engines, the sloping wings, the camouflaged body, no fin flashes or markings, I pulled off left, I called "Tally Two." I did not identify them as hostile—I did not identify them as friendly. I expected to see Hinds based on the call my flight leader had made. I didn't see anything that disputed that. I've played that particular sequence over in my mind a couple of hundred times. I don't believe I ever came off and called "Tally Hind." I called "Tally Two" at that point and the ID was based on what my flight leader called. (vol. 12, pp. 27, 28)[31]

Had either F-15 pilot been on his own, he would have likely made one or more additional passes to confirm that these were indeed enemy helicopters. Instead, both relied on the other to confirm their identity, when in fact neither did.

There is no sure-fire way to combat this problem. On the one hand, the traditional method of ensuring that one lead individual has primary responsibility for the operation certainly increases the chances that one person is paying attention and reacts when problems are identified. On the other hand, centralization of responsibility can decrease the chances that problems will in fact be identified. One of the key benefits of redundancy can be lost, since the secondary individuals may defer to the lead individual and have their reliability reduced.

How many guards should we have at nuclear facilities? Congressional testimony and government studies on this issue display little awareness of the existence of the social shirking or diffusion of responsibility phenomenon. Nuclear security guards, like pilots and air crew members, do not like to acknowledge that their reliability is influenced by the knowledge that others are also doing the same. But there is no reason to suspect that nuclear security personnel are immune to this all too human problem. Simply adding more individuals to nuclear security forces could, therefore, actually reduce the effectiveness of the whole unit in identifying or defending against a terrorist attack.

5. REDUNDANCY BACKFIRES THROUGH OVERCOMPENSATION

The third basic way in which redundancy can be counterproductive is when the addition of extra components encourages individuals or organizations to increase production in dangerous ways. In most settings, individuals and organizations face both production pressures and pressure to be safe and secure. If improvements in safety and security, however, lead individuals to engage in inherently risky behavior—driving faster, flying higher, producing more nuclear energy, etc.—then expected increases in system reliability could be reduced or even eliminated.

The offsetting or compensation behavior phenomenon has been widely studied. Research demonstrates, for example, that laws requiring "baby-proof" safety caps on aspirin bottles have led to an increase in child poisoning because parents leave the bottles outside the medicine cabinet.[32] Similar studies have suggested that the increased use of ski helmets has not led to decreases in head injuries in accidents on the slopes because many skiers with helmets just go faster down more treacherous terrain.[33] The literature on the effects of safety devices (such as airbags and seat belts) on automobile accident rates has also demonstrated that many, though not all, drivers are more reckless when driving in "safer" cars.[34] Fig. 4–6 illustrate three possibilities concerning offsetting behavior and can best be understood by thinking about the behavior of skiers driving up to Lake Tahoe ski resorts from San Francisco in the winter. Imagine that the probability of a serious automobile crash is 0.25 per lifetime for drivers who drive at the speed limit of 50 miles per hour when heading for the slopes every weekend, but increases when driving speed is increased. Fig. 4 demonstrates how the benefits of redundancy can lead to "rational" offsetting behavior: if the probability of accidents increases by only 20% for every 25 miles an hour increase in speed, the "rational driver" can drive 25% faster and be still be safer each time a new safety device is added to his or her car. Fig. 5 shows, however, that with only a small change in the increase in driving speed (from 25% to 35%), the same "rational" driver will now increase the likelihood of suffering a fatal car crash with each new safety device. Fig. 6 illustrates an even more pernicious "double effect" phenomenon in which the reliability of *each* component goes down (since seat belts, brakes, and air bags are less effective at higher speeds) and the inherent risk of an accident goes up with increase in speed.

5.1. Overcompensation in Organizations

In theory, analysts should be able to calculate the interactive effects of safety devices and increase in speed. Overcompensation can occur in organizational

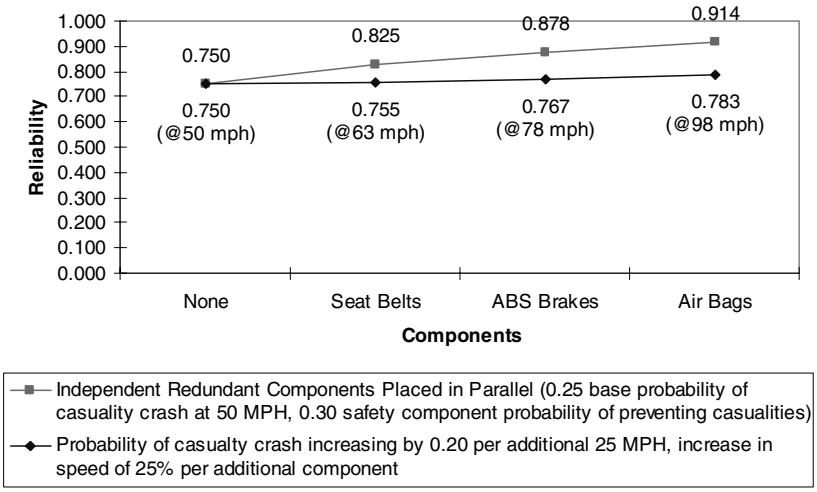

Fig. 4. Rational compensation.

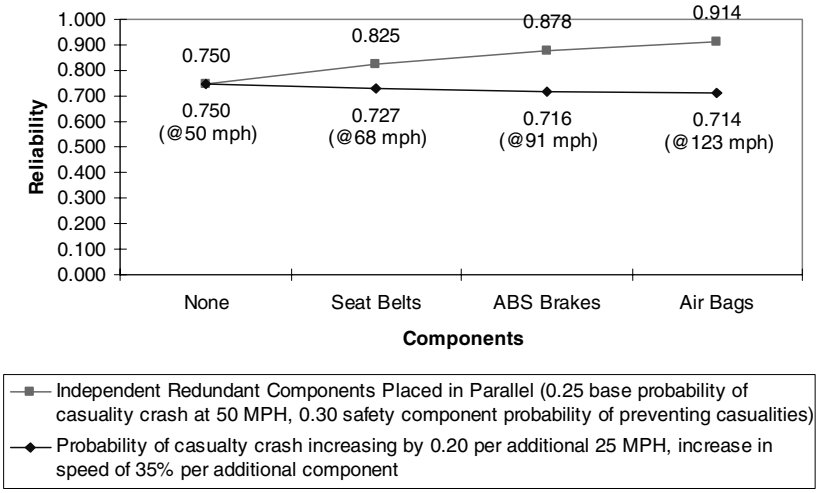

Fig. 5. Overcompensation.

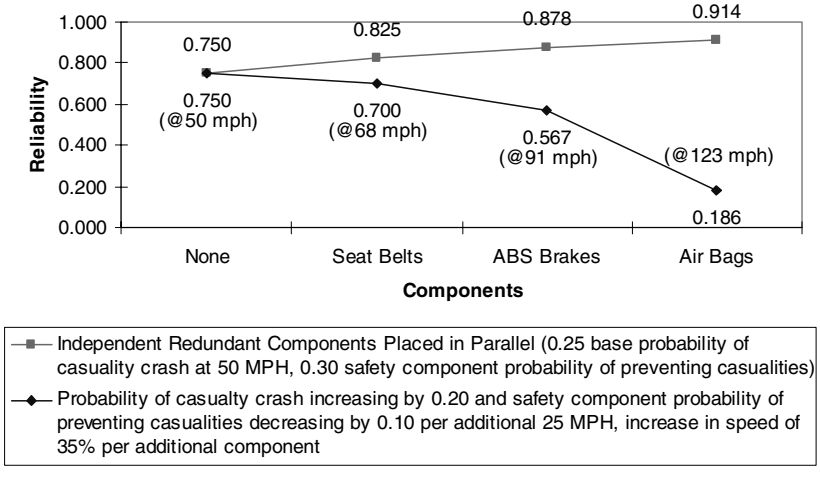

Fig. 6. Overcompensation and reduced component reliability.

settings, however, because operators and leaders cannot adequately judge where they are on the curves outlined above. They may not want to go below the knee of the curve, but they may believe that they have made safe decisions, when in fact they have already altered their operating conditions in ways that make their behavior, in the shadow of redundancy, less safe.

A dramatic case in point is the January 1986 space shuttle *Challenger* explosion. A strong consensus about the basic technical cause of the accident emerged soon afterward with the publication of the Rogers Commission report: the unprecedented cold temperature at the Kennedy Space Center at the time of launch caused the failure of two critical O-rings on a joint in the shuttle's solid rocket booster, producing a plume of hot propellant gases that penetrated the shuttle's external fuel tank and ignited its mixture of liquid hydrogen and oxygen. In contrast to the technical consensus, a full understanding of why NASA officials and Morton Thiokol engineers decided to launch the shuttle that day, despite the dangerously cold weather, has been elusive. The *Challenger* launch decision can be understood as a set of individuals overcompensating for improvements in space shuttle safety that had been produced through the use of redundant O-rings. This overcompensation interpretation differs significantly from both the traditional arguments that "production pressures" forced officials to break safety rules and consciously accept an increased risk of an accident to permit the launch to take place and Diane Vaughan's more recent argument, which focuses instead on how complex rules and engineering culture in NASA created "the normalization of deviance" in which risky operations were accepted unless it could be proven that they were extremely unsafe.

The production pressures explanation—that high-ranking officials deliberately stretched the shuttle flight safety rules because of political pressure to have a successful launch that month—was an underlying theme of the Rogers Commission report and is still a widely held view today.[35] The problem with the simple production pressure explanation is that Thiokol engineers and NASA officials were perfectly aware that the resilience of an O-ring could be reduced by cold temperature and that the potential effects of the cold weather on shuttle safety were raised and analyzed, following the existing NASA safety rules, on the night of the *Challenger* launch decision. Vaughan's argument focuses on a deeper organizational pathology: "the normalization of deviance." Engineers and high-ranking officials had developed elaborate procedures for determining "acceptable risk" in all aspects of shuttle operations. These organizational procedures included detailed decision-making rules among launch officials and the development of specific criteria by which to judge what kinds of technical evidence could be used as an input to the decision. The Thiokol engineers who warned of the O-ring failure on the night before the launch lacked proper engineering data to support their views and, upon consideration of the existing evidence, key managers, therefore, unanimously voted to go ahead with the launch. Production pressures were not the culprits, Vaughan insists. Well-meaning individuals were seeking to keep the risks of an accident to a minimum, and were just following the rules (p. 386).[36] The problem with Vaughan's argument, however, is that she does not adequately explain why the engineers and mangers followed the rules that night. Why did they not demand more time to gather data, or protest the vote in favor of a launch, or more vigorously call for a postponement until that afternoon when the weather was expected to improve?

The answer is that the *Challenger* accident appears to be a tragic example of overcompensation. There were two O-rings present in the critical rocket booster joint: the primary O-ring and the secondary O-ring were listed as redundant safety components because they were designed so that the secondary O-ring would seal even if the first leaked because of "burn through" by hot gasses during a shuttle launch. One of the Marshall space center officials summarized the resulting belief: "We had faith in the tests. The data said that the primary would always push into the joint and seal . . . . And if we didn't have a primary seal in the worst case scenario, we had faith in the secondary" (p. 105).[36] This assumption was critical on the night of January 27, 1986 for all four senior Thiokol managers reversed their initial support for postponing the launch when a Marshall Space Center official reminded them of the backup secondary O-ring. "We were spending all of our time figuring out the probability of the primary seating," one of the Thiokol managers later noted:

> [t]he engineers, Boisjoly and Thompson, had expressed some question about how long it would take that [primary] O-ring to move, [had] accepted that as a possibility, not a probability, but it was possible. So, if their concern was a valid concern, what would happen? And the answer was, the secondary O-ring would seat (p. 320).[36]

In short, the *Challenger* decision makers failed to consider the possibility that the cold temperature would reduce the resilience of both O-rings in the booster joint since that low probability event had not been wit-

nessed in the numerous tests that had been conducted. That is, however, exactly what happened on the night of unprecedented cold temperatures. Like many automobile drivers, these decision makers falsely believed that redundant safety devices allowed them to operate in more dangerous conditions without increasing the risk of a catastrophe.

How many guards should we have at U.S. nuclear facilities? An awareness of the overcompensation should remind nuclear security analysts to beware of overconfidence inadvertently producing riskier behavior. The analogy to "driving faster" would be to keep more vulnerable nuclear storage sites open or even create new nuclear facilities, in the belief that increased guard units had made all of them more secure. The analogy to the "double effect" in the automobile safety model is the likelihood, for example, that the same number of nuclear guards will be less effective in protecting more than one nuclear materials storage site at a given nuclear facility.

The overcompensation problem thus suggests one final warning that needs to be kept in mind during the ongoing debate on security against nuclear terrorism. Predicted increases in nuclear security forces should not be used as a justification of maintaining inherently insecure facilities or increasing the numbers of nuclear power plants, storage sites, or weapons facilities. Unfortunately, there are signs that this may happen. For example, in 1999, the President's Foreign Intelligence Advisory Board recommended that DOE significantly consolidate its large stockpile of sensitive nuclear materials, but in April 2002, the White House refused to provide the $41 million special funding request by the DOE to make this happen.[8,39] In May 2002, the Tennessee Valley Authority voted to restart a long mothballed nuclear reactor at Browns Ferry, Alabama, despite the heightened concerns about nuclear security after the September 11 attacks, and the later December 2001 terrorist warnings.[38] In December 2002, the NRC ruled that a set of companies seeking to build a new factory to turn weapons' plutonium into reactor fuel did not need to include an assessment of the risk of terrorist attacks in the license application because improvements to security at all U.S. nuclear facilities had already been instituted after the World Trade Center and Pentagon attacks.[39] If the predicted improvements in nuclear security after September 11 lead officials to expand the number of sites that need to be protected, however, it is not clear that the net effect of such "improvements" will actually be enhanced nuclear security.

6. CONCLUSION: THINK TWICE ABOUT REDUNDANCY

The central theoretical insight presented in this article is that organizational efforts to increase reliability and security through redundancy can backfire in numerous and complex ways. The implication of the argument, however, is *not* that redundancy never works in efforts to improve reliability and security. Moreover, the central policy lesson is *not* that the U.S. government should reject all proposals to place more security forces at nuclear facilities, given the heightened terrorist threat after the September 11, 2001 attacks. Instead, the lesson is that we need to be smarter in the way we think about redundancy.

This article presented a simple set of warnings. First, and most obviously, this is a warning against the knee-jerk reaction to throw more resources and add more people to address a problem when a crisis makes emergency efforts appear necessary. Unfortunately, organizations too often can have incentives to do exactly that. The DOE requested $138 million in emergency funds to improve the security of weapons, weapons materials, and radioactive waste soon after the September 11 attacks, but 93% of the DOE request was rejected by the White House Office of Management and Budget on grounds that the DOE had not done enough research on how the money would be used effectively and, especially, had not developed a new design basis threat against which to plan and measure new security efforts.[37] This has produced a heated debate in Congress on how best to improve nuclear security and, unfortunately, some congressmen have insisted that adding more nuclear security guards is the simple solution to our security problem.[40] Hopefully, this article can remind us all that this is not the case. Engineering and social science perspectives can be joined together to improve the quality of what should be a continuing debate about how best to protect U.S. nuclear facilities from terrorists.

The second warning, however, is to remember that low probability events happen all the time. None of the incidents described in this article were considered likely before they occurred and that fact should focus attention on the danger of misestimating the risks of redundancy. Organizations can too easily wash out estimates of low-probability events by transforming them into assumptions of impossibility. An important function of organizations is to create beliefs and rules about judging which future scenarios are probable and which are not. To focus the attention of its

members on "real" problems, organizations must decide which events are likely and which are improbable, and which are worth worrying about. Once events are considered sufficiently improbable, they can be left off the agendas, excluded from planning contingencies, and tucked away off individuals' mental radar screens. Like the Excel program that rounded off the reliability estimates in Fig. 1 above, organizational assumptions "round off" low probability events. Once these estimates become excluded from common view, they are transformed and the event is no longer considered improbable, but instead impossible.

The first step toward wisdom in this area is a simple recognition of the three problems with redundancy and their implications. If there are inherent limits to the reliability that redundancy can offer, we must not delude ourselves into believing that complex and tightly coupled organizations can be made *perfectly* reliable if only they try hard enough and build sufficient back-up security forces and safety devices. A deeper awareness of the three pathways by which redundancy can lead to increased risks—catastrophic common-mode errors, social shirking, and overcompensation—will not, of course, necessarily lead to more accurate organizational estimates of how much redundancy is enough and how much is too much. But it should help analysts and officials alike question the common, but false, intuition that actions taken to improve security will always have a positive effect.

ACKNOWLEDGMENTS

The author would like to thank the following individuals for their helpful comments on earlier drafts: Lee Clarke, Alexander Montgomery, Charles Perrow, Scott Snook, and Diane Vaughan.

REFERENCES

1. Bush, G. W. (2002). *State of the Union Address*. Available at http://www.whitehouse.gov/news/releases/2002/01/20020129-11.html, January 29, 2002.
2. Borenstein, S. (2002). Concern heightened over nuclear attacks. *San Jose Mercury News*, A3, January 31, 2002; Shanker, T. (2002). U.S. analysts find no sign Bin Laden had nuclear arms. *New York Times*, February 26, A1.
3. Gertz, B. (2002). Nuclear plants targeted. *Washington Times*, January 31, A1.
4. Risen, J., & Shenon, P. (2002). Traces of terror: The investigation; U.S. says it halted Qaeda plot to use radioactive bomb. *New York Times*, June 11, A1.
5. Hebert, H. J. (2001). Seven states using national guard to help secure reactors. *Associated Press*, November 1.
6. Abrahamson, S. (2002). *Letter to OMB Director.* Available at http://www.house.gov/markey/iss_terrorism_ltr020314.pdf, March 14, 2002.
7. The Project on Government Oversight. (2001). *U.S. Nuclear Weapons Complex: Security at Risk.* Available at http://www.pogo.org/nuclear/security/2001report/reporttext.htm.
8. Special Investigative Panel. (1999). *Science at its Best; Security at its Worst, the President's Foreign Intelligence Advisory Board* (p. 18).
9. Leventhal, P. (2001). *A Review of Security Issues at Nuclear Power Plants*. Nuclear Control Institute and Committee to Bridge the Gap before the House Committee on Energy and Commerce Subcommittee on Oversight and Investigations. Available at http://energycommerce.house.gov/107/hearings/12052001Hearing435/Leventhal746print.htm, December 5, 2001.
10. Landau, M. (1969). Redundancy, rationality, and the problem of duplication and overlap. *Public Administration Review*, *4*, 346–358; Roberts, K. H. (1990). Some characteristics of one type of high reliability organization. *Organization Science*, *2*, 160–176; La Porte, T. R., & Consolini, P. M. (1991). Working in practice but not in theory: Theoretical challenges of high reliability organizations. *Journal of Public Administration Research and Theory*, *1,* 19–47; Heimann, C. F. L. (1995). Different paths to success: A theory of organizational decision making and administrative reliability. *Journal of Public Administration Research and Theory*, *1*, 45–71.
11. La Porte, T. R. (1996). High reliability organizations: Unlikely, demanding, and at risk. *Journal of Contingencies and Crisis Management*, *4*, 63.
12. Bendor, J. B. (1987). *Parallel Systems* (pp. 26–27). Berkeley: University of California Press.
13. Rochlin, G. I., La Porte, T. R., & Roberts, K. H. (1987). The self-designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review*, *40*, 76–90.
14. Sagan, S. D. (1993). *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton, NJ: Princeton University Press.
15. Levenson, N. (1995). *Safeware*. Reading, MA: Addison-Wesley.
16. La Porte, T. R. (1988). The United States air traffic system: Increasing reliability in the midst of rapid growth. In R. Mayntz & T. P. Hughes (Eds.), *The Development of Large Technical Systems* (pp. 215–244). Boulder, CO: Westview Press.
17. Perrow, C. (1984). *Normal Accidents: Living with High Risk Technologies*. New York: Basic Books; Sagan, S. D. (1993). *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton, NJ: Princeton University Press; Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA.* Chicago: University of Chicago Press.
18. Reason, J. (1997). *Managing the Risks of Organizational Accidents.* United Kingdom: Ashgate Publishing; Sagan, S. D. (1994). Toward a political theory of organizational reliability. *Journal of Contingencies and Crisis Management, 4*, 228–240; Perrow, C. (1999). Organizing to reduce the vulnerabilities of complexity. *Journal of Contingencies and Crisis Management*, *3*, 150–155.
19. Taylor, R. W. (1990). *Twin-Engine Transports: A Look at the Future*. Seattle, WA: Boeing Corporation Report; Del Valle, C., & Schroeder, M. (1996). Did the FAA go easy on Boeing? *Business Week*, January 29, 56–60.
20. Fuller, J. G. (1975). *We Almost Lost Detroit.* New York: Reader's Digest Press; Mosey, D. (1990). Reactor accidents: Nuclear safety and the role of institutional failure. *Nuclear Engineering International Special Publications*, p. 48.
21. Sarin, R. (1990). *The Assassination of Indira Gandhi* (p. 19). New Delhi: Penguin.

22. Hirsch, D. (2002). The NRC: What, me worry? *Bulletin of Atomic Scientists*, *58*(1), 38–44.
23. Jelter, J. (2002). Tight security rings U.S. nuclear power plants. *Reuters.* Available at http://www.planetark.org/avantgo/dailynewsstory.cfm?newsid=13996, January 11, 2002.
24. Carrier, J. (1997). Flats security lax, ex-officials warn. *Denver Post*, May 20, A1; Pena, F. (1998). *Letter to Congressman Edward Markey*. Available at http://www.house.gov/markey/iss_terrorism_ltr980421.pdf.
25. Office of Nuclear Safety. (1993). *New Directions in Nuclear Safety and Management* (pp. 34–35). Department of Energy.
26. Markey, E. (2002). *Security Gap: A Hard Look at the Soft Spots in Our Civilian Nuclear Reactor Security,* Enclosure 1 (p. 9). Available at http://www.house.gov/markey/iss_nuclear_ltr020325a.pdf, March 25, 2002.
27. Montgomery, B. (2002). Oconee nuke plant had worker with criminal record. *Greenville News*, May 10, A1.
28. Darley, J. M., & Latane, B. (1968). Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology*, *4*, 377–383; Latane, B., & Nida, S. (1981). Ten years of research on group size and helping. *Psychological Bulletin*, *2*, 308–324.
29. Latane, B., & Darley, J. M. (1970). *The Unresponsive Bystander: Why Doesn't He Help?* New York: Meredith Corporation.
30. Snook, S. (2000). *Friendly Fire*. Princeton, NJ: Princeton University Press, 2000.
31. Andrus, J. G. (1994). *AFR 110-14 Aircraft Accident Investigation Board Report, US Army UH-60 Black Hawk Helicopters 87-26000 and 88-26060*, ("Andrus Report"). U.S. Air Force.
32. Viscusi, W. D. (1984). The lulling effect: The impact of child-resistant packaging on aspirin and analgesic ingestions. *American Economic Review*, *2*, 324–327; Viscusi, W. D. (1985). Consumer behavior and the safety effects of product safety regulation. *Journal of Law and Economics*, *3*, 527–553.
33. Lichtenstein, G., & Isham, J. (2003). Helmets do not make the ski slopes safer. *New York Times*, January 19, A7.
34. Peltzman, S. (1975). The effects of automobile safety regulation. *Journal of Political Economy*, *4*, 677–725; Wilde, G. J. S. (1982). The theory of risk homeostasis: Implications for safety and health. *Risk Analysis*, *4*, 209–225; Crandall, R. W., & Graham, J. D. (1984). Automobile safety regulation and offsetting behavior: Some new empirical estimates. *American Economics Review*, *2*, 328–331; Bloomquist, G. (1988). *The Regulation of Motor Vehicles and Traffic Safety.* Boston: Kluwer Academic Publishers; Peterson, S., Hoffer, G., & Millner, E. (1995). Are drivers of air-bag-equipped cars more aggressive? A test of the offsetting behavior hypothesis. *Journal of Law and Economics*, *38*, 251–264.
35. Heimann, C. F. L. (1997). *Acceptable Risks: Politics, Policy, and Risky Technologies*. Ann Arbor, MI: University of Michigan Press.
36. Vaughan, D. (1996). 386.
37. Wald, M. L. (2002). White House cut 93% of funds sought to guard atomic arms. *New York Times*, April 23, A8.
38. Firestone, D. (2002). Utility board votes to restart nuclear reactor in Alabama that has been idle since 1985. *New York Times*, May 17, A10.
39. Wald, M. L. (2003). N.R.C. excludes terrorism as licensing consideration. *New York Times*, January 7, A11.
40. Markey, E. (2002). *Nuclear Weapons Lab Security Guard Force Cut By 40% Markey Finds*. Press Release, August 20, U.S. Congress, Office of Edward Markey.