STRATEGIC STUDIES QUARTERLY

Deterrence Stability in the Cyber Age

Edward Geist

Abstract

Technical and operational realities make it prohibitively difficult to adapt a Cold War paradigm of "deterrence stability" to the new domain of cyber warfare. Information quality problems are likely to forestall the development of a cyber equivalent of the strategic exchange models that assessed deterrence stability during the Cold War. Since cyberspace is not firmly connected to geographic space the way other domains are, modeling is extremely difficult, muddling the neat conceptual distinctions between "counterforce" (military) and "countervalue" (civilian) targets. These obstacles seriously complicate US planning for a credible cyber "assured response" and present substantial challenges to potential adversaries contemplating cyber attacks against US interests. To create a maximally effective deterrent against cyber threats, the United States should seek to maximize the challenges for possible opponents by creating a cyber "strategy of technology," emphasizing resilience, denial, and offensive capabilities.

✳ ✳ ✳ ✳ ✳

On 19 March 2015, Adm Michael S. Rogers, head of US Cyber Command (USCYBERCOM), declared in testimony before the Senate Armed Services Committee that the United States needs to field offensive cyber capabilities. Complaining that the White House has not yet delegated authority to USCYBERCOM to deploy offensive tools, Rogers expressed his concern that "in the end, a purely defensive, reactive strategy will be both late to need and incredibly resource-intense," drawing the conclusion that "we need to think about: how do we increase our capacity on the offensive side to get to that point of deterrence?" The admiral's message found a ready audience among the committee members. Concurring that "I just think it's critical to develop an offen-

Edward Geist is a MacArthur Nuclear Security Fellow at the Center for International Security and Cooperation at Stanford University and former Stanton Nuclear Security Fellow at the RAND Corporation. He earned a PhD in history from the University of North Carolina and has published articles in the *Journal of Cold War Studies, Russian Review, Slavic Review*, and the *Bulletin of the History of Medicine*.

sive cyber-capability," Sen. Angus King (I-ME) went so far as to invoke Stanley Kubrick's classic 1964 film *Doctor Strangelove*. "If you build the doomsday machine, you've got to tell people you have it. Otherwise the purpose is thwarted."[1]

Should the "delicate balance of terror," as RAND strategist Albert Wohlstetter termed the Cold War nuclear standoff, be imported into the cyber domain? While the conceptual simplicity of "mutual assured destruction" seems intuitive, Wohlstetter's famous 1958 essay of that title offers a timeless warning to those who would presume that deterrence is either easy or straightforward. "Perhaps the first step in dispelling the nearly universal optimism about the stability of deterrence," he cautioned, "would be to recognize the difficulties in analyzing the uncertainties and interactions between our own wide range of choices and the moves open to the Soviets." Far from being a desirable end goal per se, in his view strategic deterrence was an unpalatable necessity. While deterrence constituted "a keystone of a defense policy," Wohlstetter implored, "it is only a part, not the whole," and he concluded that "we have talked too much of a strategic threat as a substitute for many things it cannot replace."[2]

In the wake of Wohlstetter's article, US defense analysts deployed a suite of increasingly sophisticated tools for gauging the delicate balance of terror. These models of how a nuclear exchange between the superpowers might play out in turn became a cornerstone of the field of deterrence stability. By fielding nuclear forces capable of mounting a devastating retaliation even in the aftermath of a well-planned preemptive strike, both the United States and the Soviet Union (USSR) would be deterred from risking nuclear war.

Can this Cold War paradigm of deterrence stability be adapted to the new domain of cyber warfare? However attractive this prospect might appear, technical and operational realities make it prohibitively difficult. In particular, information quality problems are likely to forestall the development of a cyber equivalent of the strategic exchange models that undergirded assessments of deterrence stability between the Cold War superpowers. The fact that cyberspace is not firmly connected to geographic space the way other domains are makes such models extremely difficult to construct and muddles neat conceptual distinctions between "counterforce" (military) and "countervalue" (civilian) targets. While these obstacles seriously complicate US planning for a credible cyber "assured response," they also present substantial challenges to potential adversaries contemplating cyber attacks against US interests. Without the ability to model the effects of cyber attacks, proper US policies and capabilities would likely dissuade rational actors from mounting assaults that might fail to have the intended effect while eliciting a devastating retaliatory response from the United States. Therefore, to create a maximally effective deterrent against cyber threats, the United States should seek to maximize these challenges for possible opponents.

Accomplishing this goal will require a comprehensive cyber "strategy of technology," emphasizing the goals of resilience (minimizing the probable damage from a successful attack), denial (minimizing the probability an attack will succeed), and offensive capabilities. Such an approach—robust enough to confront the most sophisticated state-level adversaries—would also be more effective than a deterrence strategy against nonstate actors that might not be dissuaded by rational strategic calculations. While ideally this framework would cover both US government entities and civilian property, its high upfront cost would likely limit initial federal investment to systems critical for executing US military operations and protecting essential civilian infrastructure. However, private industry should be encouraged to employ similar techniques to increase resilience of its own assets. In contrast to the Cold War, when the nature of strategic nuclear weapons made "deterrence by denial" an impossible dream, in cyberspace the United States can present potential adversaries with a highly obfuscated and constantly evolving attack surface, dissuading adversaries by undermining their faith in prospects of success.

Operations Research and the Cyber Domain

"Operations analysis" first emerged as a distinct field during the Second World War in response to new technologies that posed the same kind of unprecedented concerns for the military as emerging cyber capabilities do today. According to RAND analyst E. S. Quade, "the major impetus for this activity was provided by the introduction of new weapons systems based on, and requiring for their operation, technical know-how foreign to past military experience." Originally directed largely at tactical questions such as how to best employ or disrupt novel technologies such as radar, in the postwar years operations analysis evolved into "systems analysis" as researchers began to evaluate longer-term weapons development projects with much higher degrees of uncertainty. During the 1950s weapons system analysts, particularly at the RAND Corporation, began expanding their purview to investigate sweeping questions of strategy and national defense policy.[3]

Nuclear weapons presented merely the most novel hurdle to defense analysts during the early Cold War. They confronted a furiously evolving technological landscape in which entire new fields, such as digital computing, quickly transitioned from laboratory experiments to critical components of military hardware. The initial temptation to dismiss the technical competence of communist adversaries, furthermore, swiftly proved naïve. Confident predictions that the USSR would require at least a decade, if not more, to field its own nuclear weapon were dashed by the first Soviet atomic test in 1949. In 1953 the USSR tested a rudimentary deliverable thermonuclear weapon, arguably beating the United States on this front by several months. Aggressive Soviet pursuit of ballistic missile technology paid off spectacularly a few years later, when the USSR used the R-7 intercontinental ballistic missile (ICBM) to launch Sputnik in October 1957. While Soviet propagandists crowed that their artificial moon proved the regime was making good on the Bolshevik promise to "bring fairy tales to life," many Americans panicked in response to a widespread perception that the United States was losing its technical edge—and possibly the Cold War along with it.[4]

Defense analysts at RAND and elsewhere weaponized America's intellectual potential to counter the communist threat. They deployed—or, in many cases, conceived—the latest mathematical and technological innovations to make the problems of superpower conflict tractable. In addition to adapting tools originally conceived for economics and industrial management to questions of war and defense, systems analysts applied novel methods such as Monte Carlo simulations, linear programming, and primitive digital computers to "think about the unthinkable," as futurist Herman Kahn termed it.[5]

These intellectual currents coalesced into a new art known as "modeling" or "model-building," which in turn has served ever since as a foundation—often implicit—for much of strategic thought. Concepts such as *assured destruction* hinged on the assumption that one could model the course of a nuclear exchange accurately enough to predict that a sufficiently large retaliatory force would, in fact, survive a well-planned preemptive strike. Figures from across the strategic spectrum deployed models to justify their particular answer to the ever-controversial question of "how much is enough?" In time, an entire discipline of "deterrence stability" grew up around analyses of this type.

The concept of *deterrence stability* emerged out of the debate during the late 1950s and early 1960s about the merits of "mutual" or "minimal" deterrence. In contrast to the Eisenhower administration's declared policy of "Massive Retaliation," which held that the United States needed to maintain absolute strategic superiority over the USSR to make its deterrent threats credible, the proponents of mutual or minimal deterrence argued that a finite force would dissuade Soviet aggression so long as it was survivable. While eschewing demands for an arms buildup on the scale of the 1950s, the minimal deterrence framework did not provide a clear answer to just how large such a retaliatory force needed to be to deter the Kremlin effectively. In 1960 Daniel Ellsberg at RAND wrote an influential piece titled "The Crude Analysis of Strategic Choices" that offered an explicit formalization of Wohlstetter's concept of deterrence. By estimating the "payoffs" for US and Soviet "strike first" and "strike second" strategies, Ellsberg's model aimed to help elucidate which policy choices would discourage the USSR from attempting a first strike. "The precise effects of a change in military 'posture,' policy, or plans upon these [utility estimates] are, of course, hard to determine, uncertain, and subject to controversy," he noted, but "nevertheless, rough estimates are often made, and these are, in fact, the basis for most policy recommendations as to choices among military alternatives."[6]

Ellsberg's model provided the foundation for the analysis of strategic postures in terms of deterrence stability, and the tantalizing prospect of identifying what would be sufficient to deter the Kremlin soon found approval among policy makers. In 1971, Pres. Richard Nixon declared that "our policy remains . . . to maintain strategic sufficiency," which he defined as "the maintenance of forces adequate to prevent us and our allies from being coerced." Furthermore, "stability . . . also means numbers, characteristics, and deployments of our forces which the Soviet Union cannot reasonably interpret as being intended to threaten a disarming attack."[7] However, estimating just what it would take to achieve these goals proved to be fraught with difficulty, and in the 1970s and 1980s an immense amount of ink was spilled about how deterrence stability should be analyzed, modeled, and estimated. Despite widespread consensus about the overall assumptions of the deterrence stability framework, which could encompass a spectrum of strategic philosophies from minimal deterrence to war fighting, vociferous debate ensued about how to model the superpower nuclear balance and determine how many weapons would deter Soviet coercion without appearing threatening.[8]

Attempts to gauge the nuclear balance of terror between the superpowers employed a wide array of methodologies, but the assumed characteristics of nuclear weapons and delivery systems provided some common points of reference. In particular, nearly all of the models analyzed the problems of delivery system performance and target survivability in spatial terms. Furthermore, reconnaissance satellite photos and other intelligence data made it possible to estimate the number and probable characteristics of enemy bombers and missiles. While vociferous debates erupted between defense analysts in the United States over questions such as the exact yields of Soviet ICBM warheads and their accuracy, uncertainties for these values were well within an order of magnitude, and many of them made little impact on model outputs anyhow. From the metric of "equivalent megatonnage," which linearized the total destructiveness of superpower nuclear arsenals on the basis of the total area their warheads could theoretically expose to a certain blast overpressure, to the more sophisticated "counterforce potential" that incorporated accuracy to estimate an arsenal's total ability to hold hardened targets such as ICBM silos at risk, to full strategic exchange models that aimed to estimate how many weapons would be available to retaliate after a preemptive strike, analysts generally assumed that nuclear war could be reduced to measures of radii and area.

This commonality aside, models of strategic nuclear forces assumed a dazzling array of forms, but one in particular, the "sufficiency model," played an outsized role in public discussions of deterrence stability. As John A. Battilega and Judith K. Grange wrote in 1978, "strategic nuclear forces have given birth to a special class of models used to roughly assess the absolute and relative sufficiency of the U.S. strategic nuclear force posture, and, conversely, to assess the significance of foreign nuclear force postures." Typically falling "into the category of static or quasi-dynamic measures of effectiveness," the "primary use" of such models was "to provide a vehicle for the discussion of such concepts as strategic parity, deterrence, and stability." According to the authors, "the role of such models has evolved uniquely in connection with nuclear forces." Factors including "the definition of U.S. strategic deterrence objectives in ways which required relative comparisons with foreign adversaries, . . . the requirement to popularly debate, but in a semi-technical language, the major U.S. nuclear weapons programs, . . . [and] the requirement to think through major U.S. deterrence, strategy, and force-sizing options in a way which could be understood but which did not refer to historical experience with nuclear warfare" drove this evolution. Troublingly, this ubiquity sometimes led to these models being employed for purposes for which they were not necessarily suited: "these models are sometimes used as primary or secondary measures of effectiveness in force planning or force interaction," the authors noted, "but it should be remembered that the reason for this use stems from their historical evolution as sufficiency models."[9]

Useful as the concepts of deterrence stability and strategic sufficiency were in the policy debates of the late Cold War, by the 1990s their limitations became more and more apparent. Increasingly elaborate derivatives of Ellsberg's initial framework exacerbated a shortcoming Ellsberg admitted in 1961: the need to assign values to variables without any real-world justification for doing so.[10] Furthermore, deterrence stability and strategic sufficiency proved difficult to translate into the multipolar post–Cold War geopolitical landscape. In South Asia, the emergence of India and Pakistan as new nuclear powers offers a particularly pressing real-world countercase to elegant mathematical models of strategic stability. Unlike the Cold War superpowers, which both feared a preemptive nuclear strike, New Delhi and Islamabad both envision that nuclear use would grow out of an all-too-conceivable conventional confrontation along the countries' contested border. The additional presence of China, a long-established nuclear power, further complicates the regional strategic picture. The multiplicity of actors, along with the diversity of possible scenarios, makes it extremely challenging to model deterrence stability in this part of the world.[11] The limitations of such modeling approaches in the nuclear domain suggest that we should hesitate before importing them into emerging arenas, such as cyber warfare.

Modeling Cyber: Wrong but Useful

For better or for worse, we cannot construct sufficiency models to estimate deterrence stability in the cyber domain precisely because cyberspace differs so much from the conventional domains. Cyberspace is not measured in inches and miles, nor can the effectiveness of cyber weapons be reduced to a simple measure of destructive radius. Neither cyber weapons nor their potential targets have the sort of predictable evolution that nuclear weapons did during the Cold War. Qualitatively, new weapons such as ICBMs only appeared after years of warning and usually took at least a few years beyond that to become truly operational. Furthermore, although delivery systems became more accurate and hardened targets grew marginally more survivable, the effects of nuclear weapons remained constant, even if scientific understanding of them continued a fitful evolution. By contrast, a radical new cyber weapon with never-before-seen effects could appear overnight, or a timely patch or upgrade might render a well-designed cyber attack impotent. The disconnect between cyberspace and physical space also makes it difficult to distinguish between counterforce and countervalue targets or to restrict collateral damage. As the Stuxnet case dramatically demonstrated, it can be difficult to construct a powerful cyber weapon without running the risk it will affect systems other than its intended targets. In light of such uncertainties, it is very hard indeed to imagine a cyber equivalent to the Cold War models that estimated the superpowers' relative nuclear might.

This is not to say that comprehensive models of cyberwar are impossible to build. Such models can and should be created, but the qualitative characteristics of cyberspace and the uncertainties involved render them unable to provide the kind of confident predictions essential to make assessments of strategic stability. As the eminent British statistician George E. P. Box famously put it, "essentially, all models are wrong, but some are useful."[12] What are the challenges of modeling cyber conflict, and to what purposes can such models reasonably be put?

Unfortunately, models of cyber war require a vastly higher level of sophistication than Cold War nuclear strategic models to be useful. Most models of nuclear conflict, such as the Arsenal Exchange Model, estimate the effects of attack on the basis of intersecting probability distributions in a two- or three-dimensional space.[13] Using the circular coverage function, estimates of delivery vehicle accuracy and target hardness can readily produce a probability estimate that the target will be destroyed. This calculation could be carried out using a slide rule, and in the early years of the Cold War, it usually was. Models used for estimating strategic stability generally neglected the temporal element altogether. By contrast, the effects of cyber attacks can only be modeled through the use of dependency graphs. Computers and networks are targeted for cyber attack specifically because they are (or are perceived to be) connected to some type of resource or activity that the attacker hopes to interrupt, manipulate, or disrupt. Mathematically, such systems can be treated as directed graphs with edges representing the influence of different parts of the network upon each other. Since these influences can only travel forward in time, the system should be treated as a directed acyclic graph in which each node in the network is represented by a different node in the graph for each moment in the system's evolution. Furthermore, each of these nodes is likely to react differently depending on its internal state. Clearly, this is not the sort of problem one can readily solve with a slide rule![14]

Thankfully, there exists a variety of computational approaches that can be applied to create models of system response to cyber attacks. So long as the system is not too large, it should be possible to use object-oriented programming to simulate the dependency graphs explicitly. In fact, the first object-oriented programming language, Simula, was invented in the 1960s for simulation purposes. An object-oriented cyber-attack model could be as finely detailed as its builders cared to make it and as extensive as available computing resources would allow. This could facilitate the use of such models to investigate possible interactions between cyber, kinetic, and nuclear attacks. Despite these attractions, an object-oriented approach is liable to require tremendous amounts of analyst manpower to construct, and it is not the only possible way to model cyber war. Finite element analysis, for instance, might be adapted to model certain kinds of cyber attacks.[15]

In addition to their relative complexity, models of cyber war are likely to be extremely sensitive to the information used to construct them. The structure of the dependency graph and the reaction of its nodes to particular stimuli depending on their state are likely to result in huge qualitative differences in the output results. In contrast to a nuclear attack, where one would hardly expect a single nuclear burst to destroy dozens of discrete targets simultaneously, in the cyber domain a well-placed attack on a vulnerable node might cause the prompt failure of all its dependencies. However, both the dependencies of any particular node—as well as its vulnerabilities—may be extremely difficult to ascertain in advance. Without good-quality intelligence about both of these factors, models of cyber attack cannot have predictive value.

What purpose, then, can such models serve? The above qualities make cyber models potentially useful for operations planning for theater campaigns but of dubious utility for creating broader political-military policy strategies. In the operational realm, models of cyber attack could be useful for research purposes even when constructed on a purely notional basis. For instance, such models could be created specifically to explore the possible dynamics of multidomain operations combining cyber with nuclear or kinetic operations. By offering a concrete framework in which to investigate various hypotheses about how such interactions could play out, these simulations could provide invaluable insights—even if they could not predict the success of any particular operation. These lessons could then be applied to reduce the cyber vulnerabilities of the United States and its strategic partners. With the benefit of sufficient information about target systems, such models could also be employed for operational planning, although the considerable amount of effort needed to construct the model and the potentially limited shelf life of the reconnaissance data are apt to make this extremely challenging.

However, for strategic assessment, models of cyber attack are dubious at best and liable to be downright harmful. The analytic categories that made models useful for studying nuclear deterrence stability translate poorly into the cyber domain. If there is a cyber analog of assured destruction, policy makers can never count on it due to the immense uncertainties that would be attendant on the construction of a cyber strategic model. Furthermore, the data collection necessary to implement such a model would itself be fraught with peril, as it would require making a comprehensive assessment of all US cyber vulnerabilities. Should such an assessment, or even a fraction of it, fall into the hands of an adversary, the damage to US security would be astronomical.

The Implausibility and Undesirability of Cyber Assured Destruction

The intrinsic uncertainties of planning cyber offensives have not dissuaded some observers from insisting that in cyberspace, the timeworn maxim "the best defense is a good offense" applies more than ever. "Although the United States must demonstrate that it has in its toolkit the requisite items for use against hostile parties when necessary, there has not been a clear cut public demonstration of cyber dominance to date of which the US has definitively taken and actively sought ownership," complained Frank J. Cilluffo, Sharon L. Cardash, and George C. Salmoiraghi in a 2012 article. "Against this background, should the United States consider engaging in the digital equivalent of an aboveground nuclear test?" This drastic measure, the authors asserted, "is not to be dismissed out of hand, . . . [as] if conducted with care (commensurate to the enormity of the exercise) [it] may be instrumental to deterring hostile actors."[16]

The widespread inclination to conceptualize cybersecurity problems in a framework analogous to that developed to characterize the superpowers' nuclear stalemate is all the more unaccountable given that Cold War nuclear strategists hardly considered apocalyptic possibilities as something to be welcomed. US and Soviet scientists alike expended herculean efforts attempting to craft viable defenses against nuclear attack, only to stumble in face of insurmountable technical obstacles. Deterrence constituted an unpalatable necessity that American and Soviet leaders found themselves compelled to embrace.

Does cyber attack share the characteristics that made deterrence the least-negative option in the nuclear domain? Some official assessments have asserted as much. In 2012 the Defense Science Board (DSB) concluded "the cyber threat is serious, with potential consequences similar in some ways to the nuclear threat of the Cold War." Characterizing an "existential cyber attack" as "capable of causing sufficient wide-scale damage for the government potentially to lose control of the country," the DSB asserted that this might be accomplished by adversaries who "can invest large amounts of money (billions) and time (years) to actually create vulnerabilities in systems, including systems that are otherwise strongly protected." While thankfully such "capabilities are today limited to just a few countries such as the United States, China, and Russia," the DSB asserted that "since it will be impossible to fully defend our systems against [such] threats, deterrence must be an element of an overall risk reduction strategy."[17]

However, accounts no less authoritative have discounted the probability of existential cyber attacks. Director of National Intelligence James R. Clapper reported to the Senate Armed Services Committee on 26 February 2015 that while "cyber threats to US national and economic security are increasing in frequency, scale, sophistication, and severity of impact . . . the likelihood of a catastrophic attack from any particular actor is remote at this time." In Clapper's assessment, "Rather than a 'Cyber Armageddon' scenario that debilitates the entire US infrastructure, we envision something different." Instead of a digital apocalypse engineered by Russia or China, Clapper foresaw "an ongoing series of low-to-moderate level cyber attacks from a variety of sources over time, which will impose cumulative costs on US economic competitiveness and national security."[18] With no Armageddon in prospect, does it really make sense to seek a cyber assured-destruction capability?

In any case, the would-be instigators of an existential cyber attack would find themselves stymied by the modeling challenges thus outlined. A cyber assault capable of causing the government to lose control over part of the country would almost certainly require mounting sophisticated attacks against multiple systems simultaneously, quite possibly in coordination with kinetic actions against critical targets. However, given the difficulty of assembling reliable intelligence essential to plan such an attack, much less model its likely dynamics, how would an adversary have sufficient confidence of its chances of success? Thus, only an extremely desperate or foolhardy opponent would be likely to take such a course of action—precisely the kind of less-than-rational actor who might not be deterred in any case. It is therefore hardly surprising that Paul K. Davis, one of the United States' most-experienced model builders, opined in a RAND working paper that "deterrence by itself is a fragile basis for strategic thinking." In his view, "hoping for deterrence with today's reality would be like grasping for straws. Deterrent measures should definitely be part of strategy, but the focus should be elsewhere."[19]

A Cyber Strategy of Technology

If deterrence based upon assured destruction cannot serve as the centerpiece of US cyber strategy, what can? Fortunately, the same fundamental challenges that complicate our efforts to model the effectiveness of offensive cyber operations also bedevil our probable opponents. With foresight, the United States can craft a strategy that aims to forestall cyber attack by exacerbating these difficulties as much as possible for would-be attackers. Through a combination of increasing the resilience of US systems and undertaking measures intended to obstruct and confuse enemy intelligence-gathering efforts, the United States can dissuade both state and nonstate adversaries from attempting the most audacious cyber attacks by denying them confidence in their likely success.

In 1970 Stefan T. Possony and J. E. Pournelle expressed the concern that the USSR, unlike the United States, pursued a "technological strategy" that might deliver them victory in the superpower rivalry without firing a shot. Despite Soviet economic and technical inferiority, the ability to focus a greater share of its more limited resources on military research and development, as well as simply steal technologies from the West when convenient, might allow the Kremlin to field superior forces—particularly if the United States allowed itself to become complacent. Defining "technological warfare" as "the direct and purposeful application of the national technological base and of specific advances generated by that base to attain strategic and tactical objectives," Possony and Pournelle declared that "genuine Technological War aims at reducing the use of firepower in all forms to a minimum."[20] Emphasizing that "like all wars, the Technological War requires a deliberate strategy," they suggested that the aim of such a strategy ought to be "to make the enemy counter each move that you make, and to dance to your tune."[21]

The United States should adopt such a "strategy of technology" to address the cyber threats of the twenty-first century. This strategy should comprise three basic strands. The first of these, *resilience*, aims to protect critical US infrastructure by increasing its ability to withstand enemy action. The second, *dissuasion by denial*, aims to complicate planning for attacks on US cyber assets by increasing the difficulty of intelligence collection and analysis for potential adversaries. The third component consists of a comprehensive *offensive cyber capability*—not as a standalone deterrent; however, because if opponents take steps similar to those outlined above, this deterrent will have serious credibility problems. Instead, the offensive cyber capability will serve two purposes. First, the United States must possess a firm grasp of the "state of the art" in offensive cyber techniques so as to identify essential measures for the resilience and denial missions. Second, the offensive capability needs to complement US defense planning for conventional, space, and nuclear operations.

To improve the resilience of its own and civilian cyber systems, the Department of Defense (DOD) should partner with private industry in a long-term effort to reduce the vulnerabilities exploited by cyber attacks. While eliminating all vulnerabilities is an unattainable goal, US security would benefit from potential adversaries possessing a less plentiful choice of attack vectors. There is good reason to believe that the barriers to secure software and hardware are primarily institutional and cultural in origin rather than technical. Many older codebases were developed in an era when present-day security challenges were totally inconceivable, and traditional software engineering practices deemphasized security concerns in favor of controlling costs and meeting deadlines. While the DOD began funding research into methods to prove the correctness of programs starting in the 1960s, these made almost no impact on the way either the US defense sector or private industry developed their systems, in part because such research took many decades to bear fruit. Researchers initially hoped to develop techniques that could be applied to software written in existing programming languages, only to find that proving the correctness of even the most trivial program in a language such as FORTRAN was forbiddingly difficult. Provably correct programs, it turned out, would require a paradigmatically different approach to programming and hardware engineering. Academic researchers began developing such techniques in the 1970s, but these remained impractical for many decades as both theory and implementation slowly improved. Furthermore, there was little demand for secure systems and software until relatively recently. While the DOD sought them out for certain applications, the private sector saw little need to pay the extra expense for features that appeared totally superfluous. With a captive defense market and the ubiquitous cost-control problems, there was little incentive to either produce secure systems and software at an affordable price point or to develop the human capital and technology needed for doing so. Fortunately, there is good reason to believe that there is a way to surmount these obstacles.

Recognizing that present-day approaches will not be adequate to meet the future needs of the US military, in 2012 the Defense Advanced Research Projects Agency (DARPA) embarked on a program to pioneer the creation of secure combinations of hardware and software on the basis of formal methods such as theorem-proving. Dubbed "High-Assurance Cyber-Military Systems," this project aims "to create technology for the construction of high-assurance cyber-physical systems, where high assurance is defined to mean functionally correct and satisfying appropriate safety and security properties."[22] As a demonstration, the DARPA developed a remote-controlled drone quadcopter so secure that its "red team" of hackers could not discover any vulnerabilities in it even after the opportunity to study the complete source code for a period of six weeks.[23] This feat suggests that secure software and hardware are not just a pipe dream, but it requires a development process very alien to usual practices. Widespread adoption of such technology, even just for defense purposes, will require the establishment of a whole new culture of system development, including training large numbers of programmers and engineers in radically different ways of thinking. This transition would be difficult and expensive, but it may prove the only way to protect US assets from increasingly sophisticated cyber attacks.

Private industry is also devoting increasing attention to the possibility of employing formal methods to limit its cyber vulnerabilities. Facing increasingly steep liabilities from cyber attacks against commercial interests, in recent years the technology industry has invested ever-greater resources in qualitatively improved software engineering techniques that greatly reduce the incidence of such vulnerabilities. For instance, the Mozilla Foundation has been aggressively developing Rust, a systems programming language that aims to liberate coders from the manual memory management that so often introduces serious security holes into software.[24] Another promising approach is the use of functional programming languages such as Haskell, which aim to enable the creation of nontrivial programs with provably correct behavior by forcing software to be written in accordance with strict mathematical formalisms. While radically different from the imperative programming style familiar to most programmers, this type of functional programming has attracted growing attention from security researchers because it promises to enable the creation of software with a radically reduced number of security vulnerabilities.[25] Although the Department of Homeland Security has ongoing programs to encourage more secure software engineering practices, the DOD—thanks to its extensive purchasing power—can help accelerate the development and adoption of these technologies and the replacement of vulnerability-ridden legacy code.[26]

Although more challenging to alleviate, hardware vulnerabilities often result from similar legacy issues and engineering oversights. US civilian cyber infrastructure grew organically out of technologies that were originally engineered prior to the emergence of the kind of security threats that are all too common today. Decades later, this heritage provides adversaries with a wide array of hardware exploits to compromise US systems. A transition to fundamentally more secure technologies would be both lengthy and highly disruptive, possibly requiring a fundamental reconceptualization of the internet's technical underpinnings but might in the long term prove necessary to protect US interests. As a consequence, the DOD should subsidize efforts to develop qualitatively more secure network hardware for its own use and encourage similar efforts by the private sector with the goal of protecting civilian systems that support military operations, and ultimately, the United States as a whole.

To deny potential adversaries easy access to critical systems, the United States can shroud accurate knowledge about known or suspected vulnerabilities in a fog of disinformation and noise. Although not practical in all cases, there is little reason why systems of particular concern, such as military command and control and civilian power grids, could not create a vast number of decoys that ape their signature in cyberspace. If particularly well-engineered, these systems will appear similar enough to the real thing to fool would-be cyber attackers that they have penetrated their target—while feeding these adversaries carefully prepared disinformation intended to either deceive them about real vulnerabilities or to encourage them to commit mistakes revealing their identity and intentions.[27] Confronted by a large number of such decoys, hackers would be hard-pressed to discern reality from willful falsehood, greatly increasing the difficulty of conducting the technical reconnaissance that makes complex cyberattacks possible. The United States can make this strategy even more effective by undertaking technical measures increasing the rate at which the "real" attack surface changes. While attempting to obscure all systems in this fashion would be far too expensive and crowd out legitimate network traffic, the defense community might be able to forge a productive partnership with private industry to craft the requisite technology base, as that sector also has select assets to protect.

Finally, given the increasing use of information technology by potential adversaries, the United States should develop offensive cyber capabilities to complement military operations in other domains and to identify and ameliorate US vulnerabilities. In a future conflict, the ability to compromise enemy assets by exploiting cyber vulnerabilities could make victory less costly in terms of both blood and treasure. Furthermore, without a state-of-the-art cyber offensive capability comparable to that possessed by potential adversaries, red teaming against our own systems will be of unacceptably low quality. However, while the

United States should cultivate offensive cyber capabilities, it would be a mistake to develop these around the goal of deterrence, given that the qualitative nature of the cyber domain poses forbidding obstacles to escalation control. Without reliable models to assess the relative strength of different states' offensive cyber capabilities, or estimate the effects of cyber attacks, the concept of deterrence stability makes little sense in cyberspace. SSQ

Notes

[1] Ellen Nakashima, "Cyber Chief: Efforts to Deter Attacks against the U.S. Are Not Working," *Washington Post*, 19 March 2015, http://www.washingtonpost.com/world/national-security/head-of-cyber-command-us-may-need-to-boost-offensive-cyber-powers/2015/03/19/1ad79a34-ce4e-11e4-a2a7-9517a3a70506_story.html.

[2] Albert Wohlstetter, *The Delicate Balance of Terror* (Santa Monica, CA: RAND, 1958), http://www.rand.org/about/history/wohlstetter/P1472/P1472.html.

[3] E. S. Quade, "Introduction," in *Systems Analysis and Policy Planning: Applications in Defense* (New York: Elsevier, 1968), 2–3.

[4] For an account of the "missile gap" panic that followed Sputnik, see Peter J. Roman, *Eisenhower and the Missile Gap* (Ithaca, NY: Cornell University Press, 1995).

[5] On the early institutional history of RAND, see Bruce L. R. Smith, *The RAND Corporation: Case Study of a Nonprofit Advisory Corporation* (Cambridge, MA: Harvard University Press, 1966). On the study of nuclear war at RAND in the early Cold War, see Fred Kaplan, *The Wizards of Armageddon* (Stanford, CA: Stanford University Press, 1983).

[6] Daniel Ellsberg, *The Crude Analysis of Strategic Choices* (Santa Monica, CA: RAND, 1960). A condensed version of Ellsberg's RAND report appeared as "The Crude Analysis of Strategy Choices," *American Economic Review* 51, no. 2 (May 1961): 472–78.

[7] Richard M. Nixon, *Public Papers of the Presidents of the United States, Richard Nixon: Containing the Public Messages, Speeches, and Statements of the President 1971* (Washington, DC: Government Printing Office, 1972), 310.

[8] Although both advocates of arms limitation employed Ellsberg's framework to advance their arguments, attacks on its underlying assumptions emerged by the early 1970s. For instance, see Douglas E. Hunter, "Some Aspects of a Decision-Making Model in Nuclear Deterrence Theory," *Journal of Peace Research* 9, no. 3 (1972): 209–22.

[9] John A. Battilega and Judith K. Grange, eds., *The Military Applications of Modeling* (Wright-Patterson AFB, OH: Air Force Institute of Technology Press, 1981), 245–46.

[10] One important example of such an elaboration is Glenn Kent and David Thaler's first-strike stability model. Glenn A. Kent and David A. Thaler, *First-Strike Stability: A Methodology for Evaluating Strategic Forces* (Santa Monica, CA: RAND, 1989). For a critique of the Kent/Thaler model, see Stephen J. Cimbala and James Scouras, *A New Nuclear Century: Strategic Stability and Arms Control* (Westport, CT: Praeger, 2002), 1–23.

[11] For a recent assessment of this subject, see the essays in Michael Krepon and Julia Thompson, eds., *Deterrence Stability and Escalation Control in South Asia* (Washington, DC: Stimson Center, 2013).

[12] G. E. P. Box and N. R. Draper, *Empirical Model-Building and Response Surfaces* (New York: John Wiley and Sons, 1987), 424.

13. Battilega and Grange, *Military Applications of Modeling*, 283–89.

14. Ibid., 381–400. Models with these characteristics were developed during the Cold War to assess the survivability of US command, control, and communications (C3) systems. Historical examples included the Minimum Essential Emergency Communications Network (MEECN) and STRAT Command, used by the USAF to evaluate C3 links to the strategic bomber force.

15. Finite element analysis is widely employed in engineering for analyzing complex problems. It works by subdividing a complex system of partial differential equations (PDE) into smaller subdomains that can be approximated by a simpler subset of the PDEs. Certain applications of finite element analysis in science and engineering suggest that it may be efficacious for modeling certain types of cyber attacks. For instance, the use of the technique to model the spread of infectious disease could be analogous to the spread of malware across a network of heterogeneous systems. See for instance Joshua P. Keller, Luca Gerardo-Giorda, and Alessandro Veneziani, "Numerical Simulation of a Susceptible–Exposed–Infectious Space-Continuous Model for the Spread of Rabies in Raccoons across a Realistic Landscape," *Journal of Biological Dynamics* 7, Supplement 1 (2013): 31–46.

16. Frank J. Cilluffo, Sharon L. Cardash, and George C. Salmoiraghi, "A Blueprint for Cyber Deterrence: Building Stability through Strength," *Military and Strategic Affairs* 4, no. 3 (December 2012): 15–16.

17. Department of Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat* (Washington, DC: Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics, 2013), 2, 6.

18. James R. Clapper, director of national intelligence, "Statement for the Record Worldwide Threat Assessment of the US Intelligence Community Senate Armed Services Committee," 26 February 2015, http://www.dni.gov/files/documents/Unclassified_2015_ATA_SFR_-_SASC_FINAL.pdf.

19. Paul K. Davis, "Deterrence, Influence, Cyber Attack, and Cyberwar" (working paper WR-1049, RAND, June 2014), 1.

20. Stefan T. Possony and J. E. Pournelle, *The Strategy of Technology: Winning the Decisive War* (Cambridge, MA: Dunellen, 1970), 4, 8.

21. Ibid., 5, 15.

22. John Launchbury, "High-Assurance Cyber Military Systems (HACMS)," Defense Advanced Research Projects Agency (DARPA), no date, http://www.darpa.mil/program/high-assurance-cyber-military-systems.

23. Kathleen Fisher, "Using Formal Methods to Enable More Secure Vehicles: DARPA's HACMS Program" (presentation, Tufts University, 16 September 2014), http://wp.doc.ic.ac.uk/riapav/wp-content/uploads/sites/28/2014/05/HACMS-Fisher.pdf.

24. Mozilla Foundation, "The Rust Programming Language," no date, http://www.rust-lang.org/.

25. For instance, see David Terei, Simon Marlow, Simon Peyton Jones, and David Mazières, "Safe Haskell," *Proceedings of the 5th Symposium on Haskell*, September 2012, 137–48.

26. Department of Homeland Security (DHS), "Build Security In," no date, https://buildsecurityin.us-cert.gov/. DHS also operates the Continuous Diagnostics and Mitigation program, which aims to provide "federal departments and agencies with capabilities and tools that identify cybersecurity risks on an ongoing basis, prioritize these risks based upon potential impacts, and enable cybersecurity personnel to mitigate the most significant problems

first." See DHS, "Continuous Diagnostics and Mitigation," 14 September 2015, http://www.dhs.gov/cdm.

27. "Social engineering"—manipulating individuals to divulge sensitive information—is a critical part of many cyber attacks, but disinformation could mislead adversaries into compromising themselves.

Disclaimer

The views and opinions expressed or implied in SSQ are those of the authors and are not officially sanctioned by any agency or department of the US government. We encourage you to send comments to: strategicstudiesquarterly@us.af.mil.