Chapter 2
Laying an Intellectual Foundation for Cyberdeterrence: Some Initial Steps

Herbert Lin

Abstract This paper considers the basic question of how to effectively prevent, discourage, and inhibit hostile activity against important U.S. information systems and networks. It contains four main sections (Sections 2.1–2.3 of this paper are essentially a reproduction of The NRC letter report for the committee on deterring cyberattacks: informing strategies and developing options for U.S. policy, available at http://www.nap.edu/openbook.php?record_id=12886&page=2, 2010. Section 2.4 is based on material contained in National Research Council, in Proceedings of a workshop on deterring cyberattacks: informing strategies and developing options for U.S. policy, 2010). Section 2.1 describes a broad context for cybersecurity, establishing its importance and characterizing the threat. Section 2.2 sketches a range of possible approaches for how the nation might respond to cybersecurity threats, emphasizing how little is known about how such approaches might be effective in an operational role. Section 2.3 describes a research agenda intended to develop more knowledge and insight into these various approaches. Section 2.4 provides a summary of 15 papers by individual authors that address various aspects of the research agenda.

2.1 The Broad Context for Cybersecurity[1]

An important policy goal of the United States is to prevent, discourage, and inhibit hostile activity against the important information technology systems of the United States. This paper considers the threat of cyberattack, which refer to the deliberate use of cyber operations—perhaps over an extended period of time—to alter, disrupt,

[1]The discussion in this section is based on Chap. 1 of National Research Council (2009) and Chap. 2 of National Research Council (2007).

H. Lin (✉)
National Research Council, Washington, USA
e-mail: hlin@nas.edu

J. Krüger et al. (eds.), *The Secure Information Society*, DOI 10.1007/978-1-4471-4763-3_2,
© Springer-Verlag London 2013

deceive, degrade, usurp, or destroy adversary computer systems or networks or the information and/or programs resident in or transiting these systems or networks.[2] Cyberattack is not the same as cyber exploitation, which is an intelligence-gathering activity rather than a destructive activity and refers to the use of cyber operations—perhaps over an extended period of time—to support the goals and missions of the party conducting the exploitation, usually for the purpose of obtaining information resident on or transiting through an adversary's computer systems or networks.

Cyberattack and cyber exploitation are technically very similar, in that both require a vulnerability, access to that vulnerability, and a payload to be executed. They are technically different only in the nature of the payload to be executed. These technical similarities often mean that a targeted party may not be able to distinguish easily between a cyber exploitation and a cyberattack.

Because of the ambiguity of cyberattack and cyber exploitation from the standpoint of the targeted party, the term "cyberintrusion" will be used to refer to a hostile cyber activity where the nature of the activity is not known (that is, an activity that could be either a cyberattack or a cyber exploitation).

The range of possibilities for cyberintrusion is quite broad. A cyberattack might result in the destruction of relatively unimportant data or the loss of availability of a secondary computer system for a short period of time—or it might alter top-secret military plans or degrade the operation of a system critical to the nation, such as an air traffic control system, a power grid, or a military command and control system. Cyber exploitations might target the personal information of individual consumers or critical trade secrets of a business, military war plans, or design specifications for new weapons. Although all such intrusions are worrisome, some of these are of greater significance to the national well-being than others.

Intrusions are conducted by a range of parties, including disgruntled or curious individuals intent on vandalizing computer systems, criminals (sometimes criminal organizations) intent on stealing money, terrorist groups intent on sowing fear or seeking attention to their causes, and nation-states for a variety of national purposes. Nation-states can tolerate, sponsor, or support terrorist groups, criminals, or even individuals as they conduct their intrusions. A state might tolerate individual hackers who wish to vandalize an adversary's computer systems, perhaps for the purpose of sowing chaos. Or it might sponsor or hire criminal organizations with special cyber expertise to carry out missions that it did not have the expertise or the capability to undertake. Or it might provide support to terrorist groups by looking the other way as those groups use the infrastructure of the state to conduct Internet-based operations. In times of crisis or conflict, a state might harbor (or encourage, or control, or fail to discourage) "patriotic hackers" or "cyber patriots" who conduct hostile cyberintrusions against a putative adversary. Note that many such actions would also be plausibly deniable by the government of the host state.

[2]This report does not consider the use of electromagnetic pulse (EMP) attacks. For a comprehensive description of the threat from EMP attacks, see *Report of the Commission to Assess the Threat to the United States from Electromagnetic Pulse* (*EMP*) *Attack*, available at http://www.globalsecurity.org/wmd/library/congress/2004_r/04-07-22emp.pdf.

The threats that adversaries pose can be characterized along two dimensions—the sophistication of the intrusion and the damage it causes. Though these two are often related, they are not the same. Sophistication is needed to penetrate good cyberdefenses, and the damage an intrusion can cause depends on what the adversary does after it has penetrated those defenses. As a general rule, a greater availability of resources to the adversary (e.g., more money, time, talent) will tend to increase the sophistication of the intrusion that can be launched against any given target and thus the likelihood that the adversary will be able to penetrate the target's defenses.

Two important consequences follow from this discussion. First, because nation-state adversaries can bring to bear enormous resources to conduct an intrusion, the nation-state threat (perhaps conducted through intermediaries) is the most difficult to defend against. Second, stronger defenses reduce the likelihood but cannot eliminate the possibility that even less sophisticated adversaries can cause significant damage.

2.2 A Range of Possibilities

The discussion below focuses primarily on cyberattacks as the primary policy concern of the United States, and addresses cyber exploitation as necessary.

2.2.1 The Limitations of Passive Defense and Some Additional Options

The central policy question is how to achieve a reduction in the frequency, intensity, and severity of cyberattacks on U.S. computer systems and networks currently being experienced and how to prevent the far more serious attacks that are in principle possible. To promote and enhance the cybersecurity of important U.S. computer systems and networks (and the information contained in or passing through these systems and networks), much attention has been devoted to passive defense—measures taken unilaterally to increase the resistance of an information technology system or network to attack. These measures include hardening systems against attack, facilitating recovery in the event of a successful attack, making security more usable and ubiquitous, and educating users to behave properly in a threat environment (National Research Council 2007).

Passive defenses for cybersecurity are deployed to increase the difficulty of conducting the attack and reduce the likelihood that a successful attack will have significant negative consequences. But experience and recent history have shown that they do not by themselves provide an adequate degree of cybersecurity for important information systems and networks.

A number of factors explain the limitations of passive defense. As noted in previous NRC reports (National Research Council 2002, 2007), today's decision-making

calculus regarding cybersecurity excessively focuses vendor and end-user attention on the short-term costs of improving their individual cybersecurity postures to the detriment of the national cybersecurity posture as a whole. As a result, much of the critical infrastructure on which the nation depends is inadequately protected against cyberintrusion.

A second important factor is that passive defensive measures must succeed every time an adversary conducts a hostile action, whereas the adversary's action need succeed only once. Put differently, attacks can be infinitely varied, whereas defenses are only as strong as their weakest link. This fact places a heavy and asymmetric burden on a defensive posture that employs only passive defense.

Because passive defenses do not eliminate the possibility that an attack might succeed, it is natural for policy makers to seek other mechanisms to deal with threats that passive defenses fail to address adequately. Policy makers understandably aspire to a goal of preventing cyberattacks (and cyber exploitations as well), but most importantly to a goal of preventing *serious* cyberattacks—cyberattacks that have a disabling or a crippling effect on critical societal functions on a national scale (e.g., military mission readiness, air traffic control, financial services, provision of electric power). In this context, "deterrence" refers to a tool or a method used to help achieve this goal. The term "deterrence" itself has a variety of connotations, but broadly speaking, deterrence is a tool for dissuading an adversary from taking hostile actions.

Adversaries that might conduct cyberintrusions against the United States span a broad range and may well have different objectives. Possible adversaries include nation-states that would use cyberattacks to collect intelligence, steal technology, or "prepare the battlefield" for use of cyberattacks either by themselves or as part of a broader effort (perhaps involving the use or threat of use of conventional force) to coerce the United States; sophisticated elements within a state that might not be under the full control of the central government (e.g., Iranian Revolutionary Guards); criminal organizations seeking illicit monies; terrorist groups operating without state knowledge; and so on.

In principle, policy makers have a number of approaches at their disposal to further the broad goal of preventing serious cyberattacks on the United States. In contrast to passive defense, all of these approaches depend on the ability to attribute hostile actions to specific responsible parties (although the precise definition of "responsible party" depends to a certain extent on context).

The first approach, and one of the most common, is the use of law enforcement authorities to investigate cyberattacks, and then identify and prosecute the human perpetrators who carry out these attacks. Traditionally, law enforcement actions serve two purposes. First, when successful, they remove such perpetrators from conducting further hostile action, at least for a period of time. Second, the punishment imposed on perpetrators is intended to dissuade other possible perpetrators from conducting similar actions. However, neither of these purposes can be served if the cyberattacks in question cannot be attributed to specific perpetrators.

In a cyber context, law enforcement investigations and prosecutions have had some success, but the time scale on which such activities yield results is typically

on the order of months, during which time cyberattacks often continue to plague the victim. As a result, most victims have no way to stop an attack that is causing ongoing damage or loss of information. In addition, the likelihood that any given attack will be successfully investigated and prosecuted is low, thus reducing any potential deterrent effect. Notwithstanding the potential importance of law enforcement activities for the efficacy of possible deterrence strategies, law enforcement activities are beyond the scope of this report and will not be addressed further herein.

A second approach relies on deterrence as it is classically understood. The classical model of deterrence (discussed further in Sect. 2.2.2) seeks to prevent hostile actions through the threat of retaliation or responsive action that imposes unacceptable costs on a potential adversary or denies an adversary the benefits that may result from taking those hostile actions. Deterrence thus includes active defense, in which actions can be taken to neutralize an incoming cyberattack.

A third approach takes note of the fact that the material threat of retaliation underlying deterrence is not the only method of inhibiting undesirable behavior. Behavioral restraint (discussed further in Sect. 2.2.3) is more often the result of formal law and informal social norms, and the burden of enforcement depends a great deal on the robustness of such rules and the pressures to conform to those rules that can be brought to bear through the social environment that the various actors inhabit.

These approaches—and indeed an approach based on passive defense—are by no means mutually exclusive. For example, some combination of strengthened passive defenses, deterrence, law enforcement, and negotiated behavioral restraint may be able to reduce the likelihood that highly destructive cyberattacks would be attempted and to minimize the consequences if cyberattacks do occur. But how well any of these approaches can or will work to prevent cyberattacks (or cyberintrusions more broadly) is open to question, and indeed is a topic in need of serious research.

2.2.2 Classical Deterrence[3]

Many analysts have been drawn to the notion of deterring hostile activity against important IT systems and networks, rather than just defending against such activity. Deterrence seems like an inevitable choice in an offense-dominant world—that is, a world in which offensive technologies and tactics are generally capable of thwarting defensive efforts. As noted in Sect. 2.2.1, a major difficulty of defending against hostile actions in cyberspace arises from the asymmetry of offense versus defense.

Deterrence was and is a central construct in contemplating the use of nuclear weapons and in nuclear strategy. Because effective defenses against nuclear weapons are difficult to construct, using the threat of retaliation to persuade an adversary to refrain from using nuclear weapons is regarded by many as the most plausible and effective alternative to ineffective or useless defenses. Indeed, deterrence

[3]The discussion in Sect. 2.2.2 is based on Chap. 9 of National Research Council (2009).

of nuclear threats in the Cold War establishes the paradigm in which the conditions for successful deterrence are largely met.

Although the threat of retaliation is not the only possible mechanism for practicing deterrence, such a threat is in practice the principal and most problematic method implied by use of the term.[4] Extending traditional deterrence principles to cyberattack (that is, cyberdeterrence) would suggest an approach that seeks to persuade adversaries to refrain from launching cyberattacks against U.S. interests, recognizing that cyberdeterrence would be only one of a suite of elements of U.S. national security policy.

But it is an entirely open question whether cyberdeterrence is a viable strategy. Although nuclear weapons and cyber weapons share one key characteristic (the superiority of offense over defense), they differ in many other key characteristics, and the section below discusses cyberdeterrence and when appropriate contrasts cyberdeterrence to Cold War nuclear deterrence. What the discussion below will suggest is that nuclear deterrence and cyberdeterrence do raise many of the same questions, but indeed that the answers to these questions are quite different in the cyber context than in the nuclear context.

The U.S. Strategic Command formulates deterrence as follows (U.S. Department of Defense 2006):

> Deterrence [seeks to] *convince adversaries* not to take *actions that threaten U.S. vital interests* by means of decisive influence over their decision-making. Decisive influence is achieved by *credibly threatening* to *deny benefits* and/or *impose costs*, while *encouraging restraint* by convincing the actor that restraint will result in an *acceptable outcome*.

For purposes of this report, the above formulation will be used to organize the remainder of this section, by discussing at greater length the words in italics above. Nevertheless, there are other plausible formulations of the concept of deterrence, and these formulations might differ in tone and nuance from that provided above.

2.2.2.1 Convince

At its root, convincing an adversary is a psychological process. Classical deterrence theory assumes that actors make rational assessments of costs and benefits and refrain from taking actions where costs outweigh benefits. But it assumes unitary actors (i.e., a unitary decision maker whose cost-benefit calculus is determinative for all of the forces under his control), and also that the costs and benefits of each actor are clear, well-defined, and indeed known to all other actors involved, and further that these costs and benefits are sufficiently stable over time to formulate and implement a deterrence strategy. Classical deterrence theory bears many similarities

[4] Analysts also invoke the concept of deterrence by denial, which is based on the prospect of deterring an adversary through the prospect of failure to achieve its goals—facing failure, the adversary chooses to refrain from acting. But denial is—by definition—difficult to practice in an offense-dominant world.

to neoclassical economics, especially in its assumptions about the availability of near-perfect information (perfect in the economic sense) about all actors.

Perhaps more importantly, real decisions often take place during periods of crisis, in the midst of uncertainty, doubt, and fear that often lead to unduly pessimistic assessments. Even a cyberattack conducted in peacetime is more likely to be carried out under circumstances of high uncertainty about the effectiveness of technology on both sides, the motivations of an adversary, and the effects of an attack.

In addition, cyber conflict is relatively new, and there is not much known about how cyber conflict would or could evolve in any given situation. History shows that when human beings with little hard information are placed into unfamiliar situations in a general environment of tension, they often substitute supposition for knowledge. In the words of a former senior administration official responsible for protecting U.S. critical infrastructure, "I have seen too many situations where government officials claimed a high degree of confidence as to the source, intent, and scope of a [cyber]attack, and it turned out they were wrong on every aspect of it. That is, they were often wrong, but never in doubt" (National Research Council 2009, p. 142).

As an example, cyber operations that would be regarded as unfriendly during normal times may be regarded as overtly hostile during periods of crisis or heightened tension. Cyber operations X, Y, and Z undertaken by party A (with a history of neutrality) may be regarded entirely differently if undertaken by party B (with a history of acting against U.S. interests). Put differently, reputations and past behavior matter—how we regard or attribute certain actions that happen today will depend on what has happened in the past.

This point has particular relevance as U.S. interest in obtaining offensive capabilities in cyberspace becomes more apparent. The United States is widely regarded as the world leader in information technology, and such leadership can easily be seen by the outside world as enabling the United States to conceal the origin of any offensive cyber operation that it might have conducted. That is, many nations will find it plausible that the United States is involved in any such operation against it, and even if no U.S.-specific "fingerprints" can be found, such a fact can easily be attributed to putative U.S. technological superiority in conducting such operations.

Lastly, a potential adversary will not be convinced to refrain from hostile action if it is not aware of measures the United States may take to retaliate. Thus, some minimum of information about deterrence policy must be known and openly declared. This point is further addressed in Sect. 2.2.2.4.

2.2.2.2 Adversaries

In the Cold War paradigm of nuclear deterrence, the world is state-centric and bipolar. It was reasonable to presume that only nation-states could afford to assemble the substantial infrastructure needed to produce the required fissile material and develop nuclear weapons and their delivery vehicles. That infrastructure was sufficiently visible that an intelligence effort directed at potential adversaries could keep track of

the nuclear threat that possible adversaries posed to the United States. Today's concerns about terrorist use of nuclear weapons arise less from a fear that terrorists will develop and build their own nuclear weapons and more from a fear that they will be able to obtain nuclear weapons from a state that already has them.

These characteristics do not apply to the development of weapons for cyberattack. Many kinds of cyberattack can be launched with infrastructure, technology, and background knowledge easily and widely available to nonstate parties and small nations. Although national capabilities may be required for certain kinds of cyberattack (such as those that involve extensive hardware modification or highly detailed intelligence regarding truly closed and isolated system and networks), substantial damage can be inflicted by cyberattacks based on ubiquitous technology.

A similar analysis holds for identifying the actor responsible for an attack. In the nuclear case, an attack on the United States would have been presumed to be Soviet in origin because the world was bipolar. In addition, surveillance of potential launch areas provided high-confidence information regarding the fact of a launch, and also its geographical origin—a missile launch from the land mass of any given nation could be safely attributed to a decision by that nation's government to order that launch.

Sea-based or submarine-based launches are potentially problematic in this regard, although in a bipolar world, the Soviet Union would have been deemed responsible. In a world with three potential nuclear adversaries (the United States, Soviet Union, and China), intensive intelligence efforts have been able to maintain to a considerable extent the capability for attributing a nuclear attack to a national power, through measures such as tracking adversary ballistic missile submarines at sea. Identification of the distinctive radiological signatures of potential adversaries' nuclear weapons is also believed to have taken place.

The nuclear deterrence paradigm also presumes unitary actors, nominally governments of nation-states—that is, it presumes that the nuclear forces of a nation are under the control of the relevant government, and that they would be used only in accordance with the decisions of national leaders.

These considerations do not hold for cyberattack, and for many kinds of cyberattack the United States would almost certainly not be able to ascertain the source of such an attack, even if it were a national act, let alone hold a specific nation responsible. For example, the United States is constantly under cyberattack today, and it is widely believed (though without conclusive proof) that most of these cyberattacks are not the result of national decisions by an adversary state, though press reports have claimed that some are.

In general, prompt technical attribution of an attack or exploitation—that is, identification of the responsible party (individual? subnational group? nation-state?) based only on technical indicators associated with the event in question—is quite problematic, and any party accused of launching a given cyberintrusion could deny it with considerable plausibility. Forensic investigation might yield the identity of the responsible party, but the time scale for such investigation is often on the order of weeks or months. (Although it is often quite straightforward to trace an intrusion

to the proximate node, in general, this will not be the origination point of the intrusion. Tracing an intrusion to its actual origination point past intermediate nodes is what is most difficult.)

Three factors mitigate to some (unknowable) degree this bleak picture regarding attribution. First, for reasons of its own, a cyberattacker may choose to reveal to its target its responsibility for a cyberattack. For example, it may conduct a cyberattack of limited scope to demonstrate its capability for doing so, acknowledge its responsibility, and then threaten to conduct a much larger one if certain demands are not met.[5]

Second, over time a series of cyberintrusions might be observed to share important technical features that constitute a "signature" of sorts. Thus, the target of a cyberattack may be able to say that it was victimized by a cyberattack of type X on 16 successive occasions over the last 3 months. An inference that the same party was responsible for that series of attack might under some circumstances have some plausibility.

Third, the target of a cyberattack may have nontechnical information that points to a perpetrator, such as information from a well-placed spy in an adversary's command structure or high-quality signals intelligence. If such a party reports that the adversary's forces have just launched a cyberattack against the United States, or if a generally reliable communications intercept points to such responsibility, such information might be used to make a plausible inference about the state responsible for that attack. Political leaders in particular will not rely only on technical indicators to determine the state responsible for an attack—rather, they will use all sources of information available to make the best possible determination.

Nevertheless, it is fair to say that absent unusually good intelligence information, high confidence in the attribution of a cyberattack to a nation-state is almost certain to be unattainable during and immediately after that attack, and may not be achievable for a long time afterward. Thus, any retaliatory response to a cyberattack using either cyber or kinetic weaponry may carry a significant risk of being directed improperly, perhaps with grave unintended consequences.

2.2.2.3 Actions that Threaten U.S. Vital Interests

What actions is the United States trying to deter, and would the United States know that an action has occurred that threatens its vital interests?

A nuclear explosion on U.S. territory is an unambiguously large and significant event, and there is little difficulty in identifying the fact of such an explosion. The United States maintains a global network of satellites that are capable of detecting

[5]Of course, a forensic investigation might *still* be necessary to rule out the possibility that the putative attacker was only claiming responsibility for the attack when in fact it had no real ability to conduct the attack on its own. To mitigate the possibility that it might not be believed, the party claiming responsibility could leave a "calling card" in the wake of an attack whose contents only it could know.

and locating nuclear explosions in the air and on the ground, and a network of seismic sensors that provide additional information to localize nuclear explosions. Most importantly, a nuclear explosion would occur against the very quiet background of zero nuclear explosions happening over time.

But U.S. computer and communications systems and networks are under constant cyberintrusion from many different parties, and against this background noise, the United States would have to notice that critical systems and networks were being attacked and damaged. A cyberattack on the United States launched by an adversary might target multiple sites—but correlating information on attacks at different sites against a very noisy background to determine a common cause is today technically challenging. Target sets may be amorphous and complex, especially when massively complex and globally scaled supply chains are involved. And the nature of a questionable event (an intrusion) is often in doubt—is it an attack or an exploitation? If an attack, does a destructive cyberattack take place when the responsible software agent is *implanted* in a critical U.S. system, or when it is *activated*? Even knowing the effect or impact of an attack or exploitation is difficult, as the consequences of some intrusions will play out only over an extended period of time. (For example, an attack may be designed to have no immediate impact and only later to show destructive consequences.)

Another profound difference between the nuclear and cyber domains is that nuclear weapons are not thought to target individual private sector entities—it would be highly unusual for a major corporation, for example, to be the specific target of a nuclear weapon. By contrast, major corporations are subject to cyberattacks and cyber exploitations on a daily basis. This difference raises the question of whether deterrence of such intrusions on individual private sector entities (especially those that are regarded as a part of U.S. critical infrastructure) is an appropriate goal of U.S. policy—as suggested by recent allegations of Chinese cyberintrusions against human rights activists using Google's gmail.com service and against multiple private sector companies in the United States seeking important intellectual property of these companies (Cha and Nakashima 2010). The question is important, because targeted private entities might seek to defend themselves by retaliating against attackers or cyber spies, notwithstanding criminal prohibitions, with consequences damaging to U.S. national interests.

The question is important for a number of reasons. First, U.S. military forces have not been used in recent years to support the interests of specific private sector entities, at least not as a matter of declared public policy. Thus, an explicit threat to respond with force, whether cyber or otherwise, to a cyberattack on an individual private sector entity would constitute a major change in U.S. policy. Second, targeted private entities might seek to defend themselves by retaliating against attackers or cyber spies, even though such actions are currently illegal under U.S. law, and such retaliation by these entities might well have consequences damaging to U.S. national interests.

2.2.2.4 Credible Threat

A credible threat is one that an adversary believes can and will be executed with a sufficiently high probability to dissuade the adversary from taking action. (The definition of "sufficiently high" is subject to much debate and almost certainly depends on the specific case or issue in question. In some cases, even a low absolute probability of executing the deterrent threat is sufficient to dissuade.) In the nuclear domain, the United States developed strategic forces with the avowed goal of making them survivable regardless of what an adversary might do. Survivability means that these forces will be able to execute the retaliatory threat for which they are responsible under any possible set of circumstances. In addition, the United States conducts many highly visible military training exercises involving both its conventional and nuclear forces, at least in part to demonstrate its capabilities to potential adversaries.

On the other hand, U.S. capabilities for offensive cyber operations are highly classified, at least in part because discussing these capabilities in the open may point the way for adversaries to counter them. That is, at least some capabilities for conducting offensive cyber operations depend on a vulnerability that an adversary would be able to fix, if only he knew about it. To the extent that U.S. capabilities for cyber operations are intended to be part of its overall deterrent posture, how should the United States demonstrate those capabilities? Or is such demonstration even necessary given widespread belief in U.S. capabilities?

A credible deterrent threat need not be limited to a response in kind—the United States has a wide variety of options for responding to any given cyberattack, depending on its scope and character; these options include a mix of changes in defense postures, law enforcement actions, diplomacy, economic actions, cyberattacks, and kinetic attacks.[6]

Another dimension of making a threat credible is to communicate the threat to potential adversaries. A nation's declaratory policy underpins such communication and addresses, in very general terms, why a nation acquires certain kinds of weapons and how those weapons might be used. For example, a declaratory policy of the United States regarding nuclear weapons was stated in the National Military Strategy of 2004 (Joint Chiefs of Staff 2004):

> Nuclear capabilities [of the United States] continue to play an important role in deterrence by providing military options to deter a range of threats, including the use of WMD/E and large-scale conventional forces. Additionally, the extension of a credible nuclear deterrent to allies has been an important nonproliferation tool that has removed incentives for allies to develop and deploy nuclear forces.

[6]Chapter 1 of National Research Council (2009). As illustrations, a change in defensive posture might include dropping low-priority services, installing security patches known to cause inconvenient but manageable operational problems, restricting access more tightly, and so on. Law enforcement actions might call for investigation and prosecution of perpetrators. Diplomacy might call for demarches delivered to a perpetrator's government or severing diplomatic relations. Economic actions might involve sanctions.

For the use of cyber weapons, the United States has no declaratory policy, although the DOD Information Operations Roadmap of 2003 stated that "the USG should have a declaratory policy on the use of cyberspace for offensive cyber operations."[7]

Lastly, a "credible threat" may be based on the phenomenon of blowback, which refers to a bad consequence affecting the instigator of a particular action. In the cyberattack context, blowback may entail direct damage caused to one's own computers and networks as the result of a cyberattack that one has launched. For example, if Nation X launched a cyberattack against an adversary using a rapidly multiplying but uncustomized and indiscriminately targeted worm over the Internet, the worm might return to adversely affect Nation X's computers and networks. Blowback might also refer to indirect damage—a large-scale cyberattack by Nation X against one of its major trading partners (call it Nation Y) that affected Nation Y's economic infrastructure might have effects that could harm Nation X's economy as well. If concerns over such effects are sufficiently great, Nation X may be deterred (more precisely, self-deterred) from conducting such attacks against Nation Y (or any other major trading partner). Blowback may sometimes refer to counterproductive political consequences of an attack—for example, a cyberattack launched by a given government or political group may generate a populist backlash against that government or group if attribution of the attack can be made to the party responsible.

For blowback to be the basis of a credible threat, the dependencies that give rise to blowback should be apparent (or at least plausible) to a potential attacker. (As a possible example, it may be that given massive Chinese investment in U.S. securities, the Chinese have a large stake in the stability of U.S. financial markets, and thus might choose to refrain from an attack that might do significant harm to those markets.)

2.2.2.5 Denying Benefits

The ability to deny an adversary the benefits of an attack has two salutary results. First, an attack, if it occurs, will be futile and not confer on the adversary any particular advantage. Second, if the adversary believes (in advance) that he will not gain the hoped-for benefits, he will be much less likely to conduct the attack in the first place.

In the nuclear domain, ballistic missile defenses are believed to increase the uncertainty of an attack's success. For this reason, they need not be perfect—only good enough to significantly complicate an adversary's planning to the point at which it becomes impossible to carry out an attack with a high probability of success.

In the cyber domain, a number of approaches can be used to deny an adversary the benefits of an attack. Passive defenses can be strengthened in a number of ways, such as reducing the number of vulnerabilities present in vital systems, reducing

[7] Available at http://www.gwu.edu/~nsarchiv/NSAEBB/NSAEBB177/info_ops_roadmap.pdf.

the number of ways to access these systems, configuring these systems to minimize their exposed security vulnerabilities, dropping traffic selectively, and so on. Properties such as rapid recoverability or reconstitution from a successful attack can be emphasized.

Active defense may also be an option. Active defense against an incoming cyberattack calls for an operation, usually a cyber operation, that can be used to neutralize that incoming attack. A responsive operation (often described within the U.S. military as a "computer network defense response action") must be conducted while the adversary's cyberattack is in progress, so that there is an access path back to the facilities being used to mount the attack. In practice, active defense is possible only for certain kinds of cyberattack (e.g., denial-of-service attacks) and even then only when the necessary intelligence information on the appropriate targets to hit is available to support a responsive operation.

On the other hand, whether improvements in denying benefits are sufficient to deter a cyber adversary is open to question. Experience to date suggests that strengthening a system's passive defense posture may discourage the casual attacker, but will only suffice to delay a determined one. That is, the only costs to the attacker result from the loss of time and thus an increased uncertainty about its ability to conduct a successful attack on a precise timetable. Such uncertainty arguably contributes to deterrence if (and only if) the action being deterred is a necessary prelude to some other kind of attack that must also be planned and executed along a particular timetable.

2.2.2.6 Imposing Costs

Costs that may be imposed on an adversary typically involve the loss of assets or functionality valued by the adversary.

In the nuclear case, the ability to attribute an attack to a national actor, coupled with a knowledge of which specific states are nuclear-capable, enables the United States to identify target sets within each potential nuclear adversary, the destruction of which the United States believes would be particularly costly to those adversaries.

In the context of cyberattack, an attacker determined to avoid U.S. retaliation may well leave a false trail for U.S. forensic investigators to follow; such a trail would either peter out inconclusively or even worse, point to another nation that might well see any U.S. action taken against it as an act of war. (Catalytic conflict, in which a third party instigates mutual hostilities between two nations, is probably much easier in cyberspace than in any other domain of potential conflict.)

That said, the ability to attribute political responsibility for a given cyberattack is the central threshold question.

If responsibility cannot be attributed, the only hope of imposing any costs at all lies in identifying an access path to the platforms involved in launching the cyberattack on U.S. interests. For example, if it is possible to identify an access path to the attacking platforms in the midst of an ongoing cyberattack, knowledge of

the national (or subnational) actor's identity may not be necessary from a technical perspective to neutralize those platforms. (An analogy would be an unidentified airplane dropping bombs on a U.S. base—such an airplane could be shot down without knowing anything about the airplane or its pilot other than the fact that it was dropping bombs on a U.S. base.) Under these circumstances, a strike-back has some chance of neutralizing an incoming cyberattack even if the identity of the adversary is not known. By developing capabilities to deny the adversary a successful cyberattack through neutralization, the United States might be able to deter adversaries from launching at least certain kinds of cyberattack against the United States. Yet neutralization is likely to be difficult—destroying or degrading the source of a cyberattack while the attack is in progress may simply lead the adversary to launch the attack from a different source. It is also extremely likely that the attacking platforms will belong to innocent parties.

The attacking platforms may also be quite inexpensive—personal computers can be acquired for a few hundred dollars, and any software used to conduct an attack is virtually free to reproduce. Thus, the attacking platforms may not be assets that are particularly valuable to the attacker. Intermediate nodes that participate in an attack, such as the subverted computers of innocent parties used in a botnet, cost nothing from a capital standpoint, although they do represent some non-zero cost to the attacker of electronically capturing and subverting them.

The location(s) of the attacking platforms may be valuable to the attacker—more precisely, keeping such locations secret may be important to the attacker. But an adversary that chooses to conduct a cyberattack using platforms located in a particular location has also probably made the choice that he is willing to lose that secret location.

If responsibility can be attributed to a known actor, the range of possibilities for response becomes much larger. For example, if a nation-state can be identified as being responsible, anything of value to that state can be attacked, using any available means.[8] Indeed, options for responding to cyberattacks span a broad range and include a mix of changes in defensive postures, law enforcement actions, diplomacy, economic actions, and kinetic attacks, as well as cyberattacks.[9] Further, if

[8] One particular option deserves mention along these lines. As noted earlier, the U.S. Joint Chiefs of Staff wrote in 2004 that "Nuclear capabilities… [provide] military options to deter a range of threats, including the use of WMD/E and large-scale conventional forces." The same document defines WMD/E as follows: "The term WMD/E relates to a broad range of adversary capabilities that pose potentially devastating impacts. WMD/E includes chemical, biological, radiological, nuclear, and enhanced high explosive weapons as well as other, more asymmetrical 'weapons.' They may rely more on disruptive impact than destructive kinetic effects. For example, cyberattacks on U.S. commercial information systems or attacks against transportation networks may have a greater economic or psychological effect than a relatively small release of a lethal agent." Although the use of nuclear weapons against a known adversary could indeed impose very substantial costs, the threat to use nuclear weapons in response to any kind of cyberattack on the United States would not be credible to all adversaries.

[9] Some of these potential responses are less escalatory (e.g., changes in defensive postures); others, more so (e.g., retaliatory cyberattacks or kinetic attacks). Implementing less escalatory responses

individual/personal responsibility can be ascertained (or narrowed to a sufficiently small group of individuals), severe penalties could also be imposed, ranging from law enforcement prosecutions to permissible kinetic responses.

A variety of considerations might apply to choosing the appropriate retaliatory mode. For example, a "tit-for-tat" retaliatory response against an adversary might call for a cyberattack of comparable scale against a comparable target. However, a threat to do so might not be credible if the United States has a great deal to lose from such an action, thus throwing doubt on the viability of an "in-kind" deterrence strategy. On the other hand, a near-peer competitor might well be deterred from launching a large-scale cyberattack by the knowledge that it too would have much to lose if the United States launched an in-kind counterattack.

It may even be the case that when the responsible party is known, a responsive cyberattack is among the least useful tools for responding. Because a cyber adversary knows the time of his cyberattack, he can take action to mitigate the costs that the United States will attempt to impose following his attack. For example, the adversary can take steps in advance to invalidate the intelligence information on cyber targets that the defender has already collected on him, thus strengthening its defensive posture. Such an action could force the United States into either a nonselective retaliation or a retaliation delayed until new intelligence information can be collected. In the first case, the United States may not be willing to risk the large-scale escalation that might accompany a non-selective retaliatory cyberattack, and in the second case, the adversary may have already achieved its objectives by the time a new retaliatory strike can be planned.

Whether the *prompt* imposition of costs is necessary for deterrence is another unknown. U.S. nuclear forces and their command and control are structured to support prompt responses (in part because of a "use-it-or-lose-it" concern not necessarily present in a cyber context), and such a structure is believed to be an important element of deterring nuclear attack against the United States.

By contrast, the relationship between the pace at which responses are made and the deterrent effect of such responses in a cyber context is not well understood. Although a prompt response to an incoming cyberattack may have a number of possible benefits (e.g., a demonstration of resolve, an earlier termination of the damage resulting from an attack), such a response also raises the risk that a response may be misdirected or even undertaken mistakenly. There may be more to gain by seeking more information and being more confident about the necessary attributions.

2.2.2.7 Encouraging Restraint

Under the Cold War paradigm of nuclear deterrence, the technical prerequisite to encourage restraint on an adversary's part was the ability to execute a devastating

would seem to require lower levels of authority than would more escalatory responses, and thus would be more easily undertaken.

response no matter what the adversary did first. In particular, the existence of a powerful ballistic missile submarine force was regarded as the element of force structure that precluded a successful counterforce first strike by an adversary. More abstractly, it was the existence of a secure second-strike capability that was the foundation of encouraging restraint on the adversary's part.

In the cyber environment, there appears to be no realistic possibility of a targeted counterforce attack that will eliminate a nation's ability to execute offensive operations in cyberspace. Cyberattack forces are too easily dispersed (indeed, can operate covertly in other nations) and can launch attacks from myriad venues. (A broad and indiscriminate attack on the Internet infrastructure—analogous to a countervalue strike—might make it hard to mount a response in kind, at least until Internet services were restored.)

But it is still an open question if a secure second-strike cyberattack capability is an enabling condition for encouraging restraint on an adversary's part. That is, does the existence of a secure U.S. cyberattack capability contribute materially to encouraging an adversary to refrain from conducting offensive operations against the United States in cyberspace? Or could other U.S. capabilities for responding compensate for any shortfall in U.S. cyberattack capabilities? A related question is whether U.S. cyberattack capabilities contribute to deterring hostile adversary actions outside cyberspace. In this context, pre-emption to eliminate an adversary's cyberattack capabilities does not seem likely or plausible, although U.S. cyberattack capabilities could be used to disrupt an adversary's impending kinetic attack.

Restraint is also a concept that is relevant to escalation after conflict has begun. That is, after conflict has broken out (whether in cyberspace or kinetically), policy makers will seek to deter an adversary from escalating the conflict to greater levels of violence. In general, deterring escalation requires that the adversary believe that escalation will result in a worse outcome than maintaining the status quo, which implicitly requires that the United States have reserve capabilities (whether cyber or kinetic) that can produce such an outcome.

2.2.2.8 Acceptable Outcome

Whatever else it may be, an acceptable outcome surely involves a cessation of hostilities. A cessation of hostilities necessarily involves the transmission of orders from the cognizant political authority to its "shooters" to refrain from undertaking further offensive actions. A reciprocal or mutual cessation of hostilities involves both sides taking such action, and one party's cessation is generally conditional on the other side's cessation. Each party must therefore be convinced that the other side has ceased or will cease hostilities.

When conventional or nuclear conflict is involved, a cessation of hostilities is reasonably easy to recognize—no more missiles fly, no more nuclear weapons explode, and so on. But when cyber conflict is involved, recognizing a cessation of hostilities is quite problematic.

For example, given that there exists a background level of ongoing cyberattacks affecting the United States, how would the United States recognize that an adversary had ceased its cyberattacks? What evidence would be acceptable as proof positive that an adversary was complying with a cyber cease-fire?

Cessation of hostilities may also call for the removal of destructive elements emplaced in an adversary's information technology infrastructure. For example, if the United States had implanted Trojan horse software agents useful for cyberattack in an adversary's infrastructure, it might be obliged to remove them or render them harmless under the terms of a cease-fire. This could entail either some direct communications between the United States and these agents (which could be monitored and thus could reveal sensitive operational secrets of the United States) or keeping track of where such agents were implanted. Autonomous attack agents that require no further command direction after deployment and replicate themselves as they spread through adversary networks are particularly problematic in this regard.

Finally, both sides may have actors under their nominal jurisdiction that do not necessarily respond to national decisions to cease and desist. For example, in the aftermath of the August 2001 incident in which a Chinese fighter airplane was destroyed and a U.S. reconnaissance airplane forced to land on Chinese territory, private individuals on each side (so-called "patriotic hackers") began to conduct cyberattacks against various web sites of the other. In ordinary kinetic hostilities, private individuals do not generally have the physical wherewithal to participate directly in combat operations. But where cyberattack is concerned, they often do, and "combat operations" takes on an expanded meaning of "operations that damage or destroy adversary information technology or information."

2.2.2.9 Observations About Cyberdeterrence

An analysis of cyberdeterrence as traditionally conceived requires a knowledge of the specific adversary being deterred, the undesirable action to be deterred, the specific threat that constitutes the basis for deterrence, and the target(s) against which the threat is to be exercised.[10] These factors are not independent—for example, the nature of the relevant specific threat and target set for effective deterrence of a nation-state may well be different than that for a terrorist group, because what is both valuable and vulnerable to the former adversary (e.g., targets of economic significance) may not be to the latter (which does not have targets of economic significance and may not care if such targets are destroyed in its host nation). In short, a generalized cyberdeterrence strategy that does not account for individual adversaries and hostile actions is less likely to succeed than one that is appropriately tailored. Of course, the price for tailored deterrence is high—a great deal of knowledge and intelligence about specific adversaries is necessary to execute such a strategy.

[10] See Box 9.1 (National Research Council 2009).

Where cyberattacks launched by nation-states are at issue, cyberdeterrence should not be conceptualized as being necessarily separate from other spheres of potential conflict. Although it is possible that conflict between nations might occur entirely within cyberspace, there is no reason to presume that a sufficiently serious cyberattack would not have consequences in physical space. One reason, of course, is that computer systems and the physical world often do interact—computer systems control physical artifacts and accept data from the physical world. Adversary cyberattacks may also be accompanied by other hostile behavior, such as kinetic attacks or adverse economic actions.

The threats that are at the center of deterrence need not be limited to in-kind responses. Options for responding to cyberattacks on the United States span a broad range and include a mix of changes in defensive postures, law enforcement actions, diplomacy, cyberattacks, and kinetic attacks, and there is no reason that a retaliatory cyberattack would necessarily be favored over a retaliatory kinetic attack.

There is also a broad range of conflict scenarios to which cyberdeterrence may be applicable. For example, analysts often refer to strategic or tactical conflict between adversaries. A large-scale use of cyberattack against the critical infrastructure of a nation (e.g., against its electric grid, against its financial systems) might well be regarded as strategic in nature, whereas a cyberattack against an air defense radar system would almost certainly be regarded as tactical. Such different scenarios, or scenarios located at any point along this continuum of potentially deterrable cyberattacks, may well pose different challenges for how and to what extent deterrence is relevant to them. (For example, there may well be differences in the nature of the relevant deterrent threat or the likelihood that the deterrent threat would be carried out.)

The feasibility of cyberdeterrence and of international regimes to constrain cyberattacks on the United States is profoundly affected by the fact that the technology for cyberattacks is broadly and inexpensively available to everyone, nation-states and subnational entities down to the level of single individuals. Such broad availability means that the assumption of unitary actors is not necessarily valid.

Furthermore and as mentioned in Sect. 2.2.2.4, an environment in which certain critical infrastructures are highly interconnected across national boundaries leaves open a possibility (of unknown magnitude) that a cyberattack conducted in one nation may have global effects, including effects on the instigating nation. Perhaps the most prominent example is the existence of myriad cross-border links between financial institutions, and the consequent possibility that the U.S. financial sector (for example) might be harmed from an attack against another country's financial system.

Lastly, the private sector has a direct stake in U.S. cyberattack policy—uniquely more so than for policy regarding most other kinds of military action because of the extent of private sector ownership and operation of many of the national critical infrastructure systems that must be protected. In addition, to the extent that policy needs require certain cyberattacks to be carried out, private sector cooperation may well be required. (At the very least, accidental or inadvertent interference with a U.S. government cyberattack will have to be avoided.) And as noted in Sect. 2.2.2.3,

questions arise about whether deterrence of cyberattacks against individual private sector entities is properly a component of U.S. policy. An answer in the affirmative will raise the question of whether granting private sector entities the right to engage in active defense as a response to cyberattacks directed at them would enhance or detract from cyberdeterrence.

2.2.3 International Regimes that Limit or Require Certain Behaviors

The preceding discussion suggests that at the very least, classical deterrence theory (as construed for deterring nuclear attacks on the United States) is quite problematic when applied to cyberattacks on the United States because many of the conditions necessary for nuclear deterrence are absent from the cyber domain.

Whether a deterrence framework can be developed for the cyber domain is open to question, and indeed is one primary subject of the papers to be commissioned for this project. But whatever the useful scope for deterrence, there may also be a complementary and helpful role for international legal regimes and codes of behavior designed to reduce the likelihood of highly destructive cyberattacks and to minimize the realized consequences if cyberattacks do occur. That is, participation in international agreements may be an important aspect of U.S. policy.

In the past, nations have pursued a variety of agreements intended to reduce the likelihood of conflict and to minimize the realized consequences if conflict does occur (and also to reduce the financial costs associated with arms competitions) under the broad rubric of arms control. To achieve these objectives, arms control regimes often seek to limit capabilities of the signatories or to constrain the use of such capabilities. Thus, in the nuclear domain, agreements have (for example) been reached to limit the number and type of nuclear weapons and nuclear weapons platforms of the signatories—a limitation on capability that putatively reduces the destructiveness of conflict by limiting the capabilities on each side.

Agreements have also been reached for purposes of constraining the use of such capabilities—for example, the United States and Russia are parties to an agreement to provide advance notice to each other of a ballistic missile launch. Other proposed restrictions on use have been more controversial—for example, nations have sometimes sought agreement on "no first use of nuclear weapons." Agreements constraining the use of such capabilities are intended to reduce the possibility of misunderstandings that might lead to conflict and thus reduce the likelihood of conflict.

Lastly, international legal regimes and codes of behavior can make certain kinds of weapons unacceptable from a normative standpoint. For example, most nations today would eschew the overt use of biological weapons, and thus the likelihood of such use by any of these nations is lower than it would be in the absence of such a behavioral norm.

In the present case (that is, in thinking about ways to prevent cyberattacks of various kinds), one of the most powerful rationales for considering international agreements in the cyber domain is that all aspects of U.S. society, both civilian and

military, are increasingly dependent on information technology, and to the extent that such dependencies are greater for the United States than for other nations, restrictions on cyberattack asymmetrically benefit the United States. Proponents of such agreements also argue that aggressive pursuit of cyberattack capabilities will legitimize cyberattack as a military weapon and encourage other nations to develop such capabilities for use against the United States and its interests, much to its detriment.

Objections to such regimes usually focus on the difficulty (near-impossibility) of verifying and enforcing such an agreement. But the United States is a party to a number of difficult-to-enforce and hard-to-verify regimes that regulate conflict and prescribe rules of behavior—notably the Biological Weapons Convention (BWC). In recent years, the BWC has been criticized for lacking adequate verification provisions, and yet few policy makers suggest that the convention does not further U.S. interests.

In the cyber domain, meaningful agreements to limit acquisition of cyberattack capability are unlikely to be possible. Perhaps the most important impediment to such agreements is the verification issue—technology development for cyberattack and the testing of such technology would have few signatures that could be observed, even with the most intrusive inspection regimes imaginable.

Agreements to constrain cyberattack capabilities are also problematic, in the sense that little can be done to verify that a party to such an agreement will in fact restrict its use when it decides it needs to conduct a cyberattack. On the other hand, such agreements have a number of benefits.

- They help to create international norms regarding the acceptability of such behavior (and major nation-states tend to avoid engaging in broadly stigmatized behavior).
- They help to inhibit training that calls for such use (though secrecy will shield clandestine training).
- The violation of such agreements may be detectable. Specifically, cyberattacks that produce small-scale effects may be difficult to detect, but massively destructive attacks would be evident from their consequences, especially with appropriate rules to assist forensic assessment. If a violation is detected, the violator is subject to the consequences that follow from such detection.

Lastly, even though the development of regimes constraining use would address only cyberattacks associated with nation-states, they could have significant benefit, as nation-states do have advantages in pursuing cyberattack that most nonstate-supported actors do not have. Although such regimes would not obviate the need for passive defenses, they could be useful in tamping down risks of escalation and might help to reduce international tensions in some circumstances.

As illustrations of regimes constraining use, nations might agree to confidence-building measures that committed them to providing mutual transparency regarding their activities in cyberspace, to cooperate on matters related to securing cyberspace (e.g., in investigating the source of an attack), to notify each other regarding certain activities that might be viewed as hostile or escalatory, or to communicate directly

with each other during times of tension or crisis. Agreements to eschew certain kinds of cyberattack under certain circumstances could have value in reducing the likelihood of kinetic conflict in those cases in which such cyberattacks are a necessary prelude to a kinetic attack.

Limitations on cyber targeting (e.g., no cyberattacks on civilian targets; requirements that military computers be explicitly identified; no first use of cyberattack on a large scale; or no attacks on certain classes of targets, such as national power grids, financial markets or institutions, or air traffic control systems) could prevent or reduce the destructiveness of an attack, assuming that collateral and/or cascading damage could be limited. Agreements (or unilateral declarations) to abide by such agreements might be helpful in establishing appropriate rules of conduct (norms of behavior) and a social structure to enforce those rules.

On the other hand, U.S. policy makers and analysts have not seriously explored the utility and feasibility of international regimes that deny the legitimacy of cyberattacks on critical infrastructure assets, such as power grids, financial markets, and air traffic control systems.[11] How useful would such a regime be, especially applied in concert with a significantly improved cyberdefensive posture for these assets? How would difficulties of verification and enforcement affect relative national military postures and the credibility of the regime? What meaningful capabilities would the United States be giving up if it were to agree to such a regime? These and other related questions find few answers in the literature. The feasibility of these or other regimes to limit use of cyberattack is unclear, especially in light of the difficulties of working out the details of how the regime would actually operate. It is for this reason that research is needed to explore their feasibility.

Agreements in a cyber context might also usefully address important collateral issues, such as criminal sanctions or compensation for damages sustained under various circumstances. They might also require signatories to pass national laws that criminalize certain kinds of cyber behavior undertaken by individuals and to cooperate with other nations in prosecuting such behavior, much as the Convention on Cyber Crime has done.[12]

There are a number of major complications associated with arms control regimes for cyberattack. These include:

- The functional similarity between cyber exploitation and cyberattack. That is, from the target's perspective, it may be difficult or impossible to distinguish between a cyber operation intended for attack and one intended for exploitation. Restrictions on cyberattack will almost certainly restrict cyber exploitation to a large degree, and nations—including the United States—may well be loath to surrender even in principle any such capability for gaining intelligence.

[11] Indeed, the United States has until recently avoided discussions on military uses of cyberspace. In December 2009, it was publicly reported that the United States had begun to engage with Russian officials and with UN officials (see Markoff and Kramer 2009), although the emphasis of the United States in these talks was apparently directed toward combating Internet crime and as a collateral effect strengthening defenses against any militarily-oriented cyberattacks.

[12] See http://conventions.coe.int/Treaty/EN/Treaties/html/185.htm.

- The lack of state monopoly over cyber weapons. For kinetic weaponry, the destructiveness and potency of any given weapon has some significant correlation with the extent to which it is only available to nation-states—almost everyone has access to rifles, whereas jet fighters and submarines are mostly restricted to nations. For cyber weapons, this correlation is far less strong, and private parties can and do wield some cyber weapons that can be as destructive and powerful as some of those wielded by nation-states. Although as a rule nation-states do have major operational advantages in conducting cyberattacks (e.g., intelligence agencies that can support cyberattack), nonstate actors are certainly capable of acquiring cyber weaponry that can cause enormous damage.
- "Positive inspection" arrangements to increase the confidence that each side is abiding by an agreement not to engage in proscribed activities could be easily thwarted or circumvented. One primary reason is that the footprint of personnel and equipment needed to conduct cyber operations is small, and thus could be located virtually anywhere in a nation (or even in another nation).
- In contrast to nuclear weapons, the private sector has essentially unlimited access to most of the technology that underlies cyberattack weapons, and the scope for destructive use varies over a much wider range. Thus, an extraordinary degree of intrusiveness would be required to impose controls on the private acquisition and use of cyber weapons. It would be impractical and unacceptable, not to mention futile, to subject every personal computer and all forms of electronic communication to inspection to ensure that cyber weapons are not present on computers or concealed within e-mails. On the other hand, special rules might help to regulate access to the operations of critical social infrastructure in order to improve the attribution of parties that come into contact with them.
- The inherent anonymity of cyberattacks, mentioned above, greatly complicates the attribution of responsibility for an attack, and thus it is difficult to hold violators of any agreement accountable. Any alleged violation could simply be met with a strongly worded denial, and unambiguous evidence supporting the allegation would be hard to provide. Moreover, behavioral norms are generally much harder to instill and enforce in an environment in which actors can act anonymously.

  Suggestions are often made to create a parallel Internet (call it an SAI, for strongly authenticated Internet) that would provide much stronger authentication of users than is required on today's Internet and would in other ways provide a much more secure environment.[13] If important facilities, such as power grids and financial institutions, migrated to an SAI, accountability for misbehavior would be much greater (because of the lack of anonymity) and the greater security of the

[13]For example, the White House Cyberspace Policy Review of May 2009 called for the nation to "implement, for high-value activities (e.g., the Smart Grid), an opt-in array of interoperable identity management systems to build trust for online transactions" (White House 2009). More recently, a trade press article reported on the intent of the Defense Information Systems Agency of the U.S. Department of Defense to establish an enclave for its unclassified networks that is isolated from public Internet access (Corrin 2010).

environment would mean that only very sophisticated parties could mount attacks on it or within it.

Although the availability of an SAI would certainly improve the security environment over that of today, it is not a panacea. Perhaps most importantly, SAI users would immediately become high-priority targets to be compromised by nontechnical cyberattacks. A compromised SAI user would then become an ideal platform from which to launch IT-based cyberattacks within the SAI—and in particular, would become an ideal jumping-off point for slowly and quietly assembling an array of computing resources that can be used for attack—all of which would be on the SAI. In addition, experience with large networks indicates that maintaining an actual air-gap isolation between an SAI and the standard Internet or dial-up or wireless connections would be all but impossible—not for technical reasons but because of a human tendency to make such connections for the sake of convenience.

- Subnational groups can take action independently of governments. Subnational groups may be particularly difficult to identify, and are likely to have few if any assets that can be targeted. Some groups (such as organized hacker groups) regard counterattacks as a challenge to be welcomed rather than a threat to be feared. Finally, a subnational group composed of terrorists or insurgents might seek to provoke retaliation in order to galvanize public support for it or to provoke anti-American sentiments in its supporting public.

  This last point is particularly relevant to any international agreements or regime that the United States might deem helpful in reducing cyberattacks against it—any legal agreement or regime must be respected by all parties, including the United States. If the United States wishes other nations to eschew certain actions or to abide by certain behavioral requirements or to grant it certain rights under certain circumstances, it too must be willing to do the same with respect to other nations.

  As an example, some analysts have suggested that it is an appropriate strategy for the United States to seek the right to retaliate against a nation for offensive acts emanating from within its borders, even if that nation's government denies responsibility for those attacks and asserts that those responsible are nonstate actors. Doing so, they argue, would give states an incentive to crack down on harmful private offensive actors in its borders. On the other hand, it is not clear that it is in the U.S. interest for the United States to be subject to such a regime, given that parties within the United States are themselves responsible for conducting many cyberattacks against the rest of the world. Any solution proposed for other nations must (most probably) be tolerable to the United States as well, but accepting such consequences may be politically, or economically, or legally infeasible.

It should also be noted that the traditional arms control agreements are not the only form of agreement that might be helpful (National Research Council 2009, Chap. 10). For example, nations have sometimes agreed on the need to protect some area of international activity such as airline transport, telecommunications, maritime activities, and so on, and have also agreed on standards for such protection.

They may declare certain purposes collectively with regard to a given area of activity on which they agree, often in the form of a multilateral treaty, and then establish consensus-based multilateral institutions (generally referred to as "specialized agencies" composed of experts rather than politicians) to which to delegate (subject to continuous review) the task of implementing those agreed purposes.

It has sometimes been easier to obtain agreement among the nations involved on standards and methods concerning the civilian (commercial) aspects of a given activity than to obtain agreement on the military (governmental) aspects of the same activity (Sofaer and Goodman 2000). For example, civil aviation is regulated internationally through agencies that have promulgated numerous agreements and regulations, all by consensus. Over the years, some precedents, and some forms of regulation, have been established, again largely by consensus, that have enhanced the protection of civilian aviation and reduced the uncertainties regarding governmental (military) aviation. A similar pattern of international regulation has resulted in increased maritime safety.

In both areas, states have agreed to criminalize terrorist attacks, and to prosecute or extradite violators. These commitments have not uniformly been kept, but security has been enhanced in these areas of international commerce because of the virtually universal support given to protecting these activities from identified threats. It is an open question whether such an approach might enhance cybersecurity internationally, whether or not it excludes any direct application or restriction on the national security activities of signatories.

2.2.4 Domestic Regimes to Promote Cybersecurity

Law enforcement regimes to prosecute cyber criminals are not the only ones possible to help promote cybersecurity. As noted in *Toward a Safer and More Secure Cyberspace*, the nation's cybersecurity posture would be significantly enhanced if all owners and operators of computer systems and networks took actions that are already known to improve cybersecurity. That is, the nation needs to do things that the nation already knows how to do.

What that report identified as a critical problem in cybersecurity was a failure of action. That report attributed the lack of adequate action to two factors—the fact that decision makers discount future possibilities of disaster so much that they do not see the need for present-day action (that is, they weigh the immediate costs of putting into place adequate cybersecurity measures, both technical and procedural, against the potential future benefits (actually, avoided costs) of preventing cyber disaster in the future—and systematically discount the latter as uncertain and vague) and the additional fact that the costs of inaction are not borne by the relevant decision makers (that is, the nation as a whole bears the cost of inaction, whereas the cost of action is borne by the owners and operators of critical infrastructure, which are largely private-sector companies).

Accordingly, that report called for changes in the decision-making calculus that at present excessively focuses vendor and end-user attention on the short-term costs

of improving their cybersecurity postures. The report did not specify the nature of the necessary changes, but rather noted the need for more research in this area to assess the pros and cons of any given change.

The present report reiterates the importance of changing the decision-making calculus described above, but suggests that developing the necessary domestic regime (including possibly law, regulation, education, culture, and norms) to support a new calculus will demand considerable research.

2.3 A Possible Research Agenda

Although the preceding section seeks to describe some of the essential elements of cyberdeterrence, it is sobering to realize the enormity of intellectually unexplored territory associated with such a basic concept. Thus, considerable work needs to be done to explore the relevance and applicability of deterrence and prevention/inhibition to cyber conflict. At the highest level of abstraction, the central issue of interest is to identify what combinations of posture, policies, and agreements might help to prevent various actors (including state actors, nonstate actors, and organized criminals) from conducting cyberattacks that have a disabling or a crippling effect on critical societal functions on a national scale (e.g., military mission readiness, air traffic control, financial services, provision of electric power).

The broad themes described below (lettered A-H) are intended to constitute a broad forward-looking research agenda on cyberdeterrence. Within each theme are a number of elaborating questions that are illustrative of those that would benefit from greater exploration and analysis. Thoughtful research and analysis in these areas would contribute significantly to understanding the nature of cyberdeterrence.

A. Theoretical Models for Cyberdeterrence

1. Is there a model that might appropriately describe the strategies of state actors acting in an adversarial manner in cyberspace? Is there an equilibrium state that does not result in cyber conflict?
2. How will any such deterrence strategy be affected by mercenary cyber armies for hire and/or patriotic hackers?
3. How does massive reciprocal uncertainty about the offensive cyberattack capabilities of the different actors affect the prospect of effective deterrence?
4. How might adversaries react technologically and doctrinally to actual and anticipated U.S. policy decisions intended to strengthen cyberdeterrence?
5. What are the strengths and limitations of applying traditional deterrence theory to cyber conflict?
6. What lessons and strategic concepts from nuclear deterrence are applicable and relevant to cyberdeterrence?
7. How could mechanisms such as mutual dependencies (e.g., attacks that cause actual harm to the attacker as well as to the attacked) and counterproductivity (e.g., attacks that have negative political consequences against the attacker) be

used to strengthen deterrence? How might a comprehensive deterrence strategy balance the use of these mechanisms with the use of traditional mechanisms such as retaliation and passive defense?

B. Cyberdeterrence and Declaratory Policy

8. What should be the content of a declaratory policy regarding cyberintrusions (that is, cyberattacks and cyberintrusions) conducted against the United States? Regarding cyberintrusions conducted by the United States? What are the advantages and disadvantages of having an explicit declaratory policy? What purposes would a declaratory policy serve?
9. What longer-term ramifications accompany the status quo of strategic ambiguity and lack of declaratory policy?
10. What is the appropriate balance between publicizing U.S. efforts to develop cyber capabilities in order to discourage/deter attackers and keeping them secret in order to make it harder for others to foil them?
11. What is the minimum amount and type of knowledge that must be made publicly available regarding U.S. government cyberattack capabilities for any deterrence policy to be effective?
12. To the extent that a declaratory policy states what the United States will not do, what offensive operational capabilities should the United States be willing to give up in order to secure international cooperation? How and to what extent, if at all, does the answer vary by potential target (e.g., large nation-state, small nation-state, subnational group, and so on)?
13. What declaratory policy might help manage perceptions and effectively deter cyberattack?

C. Operational Considerations in Cyberdeterrence

14. On what basis can a government determine whether a given unfriendly cyber action is an attack or an exploitation? What is the significance of mistaking an attack for an exploitation or vice versa?
15. How can uncertainty and limited information about an attacker's identity (i.e., attribution), and about the scope and nature of the attack, be managed to permit policy makers to act appropriately in the event of a national crisis? How can overconfidence or excessive needs for certainty be avoided during a cyber crisis?
16. How and to what extent, if at all, should clear declaratory thresholds be established to delineate the seriousness of a cyberattack? What are the advantages and disadvantages of such clear thresholds?
17. What are the tradeoffs in the efficacy of deterrence if the victim of an attack takes significant time to measure the damage, consult, review options, and most importantly to increase the confidence that attribution of the responsible party is performed correctly?
18. How might international interdependencies affect the willingness of nations to conduct certain kinds of cyberattack on other nations? How can blowback be

exploited as an explicit and deliberate component of a cyberdeterrence strategy? How can the relevant feedback loops be made obvious to a potential attacker?

19. What considerations determine the appropriate mode(s) of response (cyber, political, economic, traditional military) to any given cyberattack that calls for a response?
20. How should an ostensibly neutral nation be treated if cyberattacks emanate from its territory and that nation is unable or unwilling to stop those attacks?
21. Numerous cyberattacks on the United States and its allies have already occurred, most at a relatively low level of significance. To what extent has the lack of a public offensive response undermined the credibility of any future U.S. deterrence policy regarding cyberattack? How might credibility be enhanced?
22. How and to what extent, if at all, must the United States be willing to make public its evidence regarding the identity of a cyberattacker if it chooses to respond aggressively?
23. What is the appropriate level of government to make decisions regarding the execution of any particular declaratory or operational policy regarding cyberdeterrence? How, if at all, should this level change depending on the nature of the decision involved?
24. How might cyber operations and capabilities contribute to national military operations at the strategic and tactical levels, particularly in conjunction with other capabilities (e.g., cyberattacks aimed at disabling an opponent's defensive systems might be part of a larger operation), and how might offensive cyber capabilities contribute to the deterrence of conflict more generally?
25. How should operational policy regarding cyberattack be structured to ensure compliance with the laws of armed conflict?
26. How might possible international interdependencies be highlighted and made apparent to potential nation-state attackers?
27. What can be learned from case studies of the operational history of previous cyberintrusions? What are the lessons learned for future conflicts and crises?
28. Technical limitations on attribution are often thought to be the central impediment in holding hostile cyber actors accountable for their actions. How and to what extent would a technology infrastructure designed to support high-confidence attribution contribute to the deterrence of cyberattack and cyber exploitation, make the success of such operations less likely, lower the severity of the impact of an attack or exploitation, and ease reconstitution and recover after an attack? What are the technical and nontechnical barriers to attributing cyberintrusions? How might these barriers be overcome or addressed in the future?

D. Regimes of Reciprocal/Consensual Limitations

29. What regimes of mutual self-restraint might help to establish cyberdeterrence (where regimes are understood to include bilateral or multilateral hard-law treaties, soft-law mechanisms [agreements short of treaty status that do not require ratification], and international organizations such as the International Telecommunication Union, the United Nations, the Internet Engineering Task Force, the Internet Corporation for Assigned Names and Numbers, and so on)?

Given the difficulty of ascertaining the intent of a given cyber action (e.g., attack or exploitation) and the scope and extent of any given actor's cyber capabilities, what is the role of verification in any such regime? What sort of verification measures are possible where agreements regarding cyberattack are concerned?

30. What sort of international norms of behavior might be established among like-minded nations collectively that can help establish cyberdeterrence? What sort of self-restraint might the United States have to commit to in order to elicit self-restraint from others? What might be the impact of such self-restraint on U.S. strategies for cyber conflict? How can a "cyberattack taboo" be developed (perhaps analogous to taboos against the use of biological or nuclear weapons)?
31. How and to what extent, if any, can the potency of passive defense be meaningfully enhanced by establishing supportive agreements and operating norms?
32. How might confidence-building and stability measures (analogous to hotline communications in possible nuclear conflict) contribute to lowering the probability of crises leading to actual conflict?
33. How might agreements regarding nonmilitary dimensions of cyberintrusion support national security goals?
34. How and to what extent, if at all, should the United States be willing to declare some aspects of cyberintrusion off limits to itself? What are the tradeoffs involved in foreswearing offensive operations, either unilaterally or as part of a multilateral (or bilateral) regime?
35. What is an act of war in cyberspace? Under what circumstances can or should a cyberattack be regarded as an act of war.[14] How and to what extent do unique aspects of the cyber realm, such as reversibility of damage done during an attack and the difficulty of attribution, affect this understanding?
36. How and to what extent, if any, does the Convention on Cyber Crime (http://conventions.coe.int/Treaty/EN/Treaties/html/185.htm) provide a model or a foundation for reaching further international agreements that would help to establish cyberdeterrence?
37. How might international and national law best address the issue of patriotic hackers or cyber patriots (or even private sector entities that would like to respond to cyberattacks with cyber exploitations and/or cyberattacks of their own), recognizing that the actions of such parties may greatly complicate the efforts of governments to manage cyber conflict?

E. Cyberdeterrence in a Larger Context

38. How and to what extent, if at all, is an effective international legal regime for dealing with cyber crime a necessary component of a cyberdeterrence strategy?
39. How and to what extent, if at all, is deterrence applicable to cyberattacks on private companies (especially those that manage U.S. critical infrastructure)?

[14]The term "act of war" is a colloquial term that does not have a precise international legal definition. The relevant terms from the UN Charter are "use of force," "threat of force," and "armed attack," although it must be recognized that there are no internationally agreed-upon formal definitions for these terms either.

40. How should a U.S. cyberdeterrence strategy relate to broader U.S. national security interests and strategy?

F. The Dynamics of Action/Reaction

41. What is the likely impact of U.S. actions and policy regarding the acquisition and use of its own cyberattack capabilities on the courses of action of potential adversaries?
42. How and to what extent, if at all, do efforts to mobilize the United States to adopt a stronger cyberdefensive posture prompt potential adversaries to believe that cyberattack against the United States is a viable and effective means of causing damage?

G. Escalation Dynamics

43. How might conflict in cyberspace escalate from an initial attack? Once cyber conflict has broken out, how can further escalation be deterred?
44. What is the relationship between the onset of cyber conflict and the onset of kinetic conflict? How and under what circumstances might cyberdeterrence contribute, if at all, to the deterrence of kinetic conflict?
45. What safeguards can be constructed against catalytic cyberattack? Can the United States help others with such safeguards?

H. Collateral Issues

46. How and to what extent do economics and law (and regulation) affect efforts to enhance cybersecurity in the private sector? What are the pros and cons of possible solution elements that may involve (among other things) regulation, liability, and standards-setting that could help to change the existing calculus regarding investment strategies and approaches to improve cybersecurity? Analogies from other "protection of the commons" problem domains (e.g., environmental protection) may be helpful.
47. What are the civil liberties implications (e.g., for privacy and free expression) of policy and technical changes aimed at preventing cyberattacks, such as systems of stronger identity management for critical infrastructure? What are the tradeoffs from a U.S. perspective? How would other countries see these tradeoffs?
48. How can the development and execution of a cyberdeterrence policy be coordinated across every element of the executive branch and with Congress? How should the U.S. government be organized to respond to cyber threats? What organizational or procedural changes should be considered, if any? What roles should the new DOD Cyber Command play? How will the DOD and the intelligence community work together in accordance with existing authorities? What new authorities would be needed for effective cooperation?
49. How and to what extent, if any, do private entities (e.g., organized crime, terrorist groups) with significant cyberintrusion capabilities affect any government policy regarding cyberdeterrence? Private entities acting outside government control and private entities acting with at least tacit government approval or support should both be considered.

50. How and to what extent are current legal authorities to conduct cyber operations (attack and exploitation) confused and uncertain? What standards should govern whether or not a given cyber operation takes place? How does today's uncertainty about authority affect the nation's ability to execute any given policy on cyberdeterrence?
51. Cyberattack can be used as a tool for offensive and defensive purposes. How should cyberattacks intended for defensive purposes (e.g., conducted as part of an active defense to neutralize an incoming attack) differ from those intended for offensive purposes (e.g., a strategic cyberattack against the critical infrastructure of an adversary)? What guidelines should structure the former as opposed to the latter?

Research contributions in these areas will have greater value if they can provide concrete analyses of the offensive actors (states, criminal organizations, patriotic hackers, terrorists, and so on), motivations (national security, financial, terrorism), actor capacities and resources, and which targets require protection beyond that afforded by passive defenses and law enforcement (e.g., military and intelligence assets, critical infrastructure, and so on).

2.4 Deterring Cyberattacks: Informing Strategies and Developing Options for U.S. Policy

On June 10–11, 2010, the National Research Council held a workshop entitled "Deterring Cyberattacks: Informing Strategies and Developing Options for U.S. Policy." During this workshop, a number of papers, related to this topic and commissioned by the National Research Council, were presented. These papers were revised and then printed in the published proceedings of this workshop. In addition, the NRC sponsored a prize competition for papers that addressed one or more of the questions raised in Sect. 2.3 above. Two of these papers were singled out for recognition and were included in the published proceedings.

This section contains summaries of the papers published in the proceedings. These summaries were contributed by the authors of those papers. The groupings below reflect the way in which papers were groups in the published proceedings.

2.4.1 Group 1—Attribution and Economics

Introducing the Economics of Cybersecurity: Principles and Policy Options by Tyler Moore

1. Many of the problems plaguing cybersecurity are economic, and modest interventions that align stakeholder incentives and correct market failures can improve our nation's cybersecurity posture substantially.

2. Government should engage Internet service providers (ISPs) in the malware-remediation process by offering exemption from liability for the harm caused by their customers' infected machines in exchange for assisting with the cleanup. The costs of cleanup should be split between ISPs, software firms, and the government, and infection reports should be published on data.gov to encourage better measurement and accounting of these harms.
3. Better data on security incidents are needed to motivate optimal private sector cybersecurity investment. To that end, aggregated reports of online-banking incidents and losses from banks should be collected and aggregated statistics published on data.gov.

Untangling Attribution by David Clark and Susan Landau

1. The occasions when attribution at the level of an individual person is useful are few.
2. Attribution of multistage attacks requires tracing a chain of attribution across several machines.
3. Multistage attacks pose a prime problem for the research community. They should be of central attention to network researchers rather than, for example, the problem of designing highly robust top-down identity schemes.
4. Internet protocol (IP) addresses are more useful than is sometimes thought as a basis of various kinds of attribution.

A Survey of Challenges in Attribution by W. Earl Boebert

1. The Internet has intrinsic features and extrinsic services that support anonymity and inhibit forensic attribution of cyberattacks, and this situation is expected to worsen.
2. Even if perfect forensic attribution were achieved, it would not have a substantial deterrent effect in most cases in which serious disruptive cyberattacks are contemplated by parties hostile to the United States.
3. Alternatives to forensic attribution include counterattack ("hack-back") and sustained, aggressive covert intelligence-gathering and subversion of potential attackers. Such methods promise a greater deterrent effect than forensic attribution. The obstacles to them are primarily nontechnical.

2.4.2 Group 2—Strategy, Policy, and Doctrine

Applicability of Traditional Detserrence Concepts and Theory to the Cyber Realm by Patrick M. Morgan

1. We are fortunate in still being in the early stages of devising responses to the cyber attack threat, and it is hoped that this means that we will avoid the mistakes of

our frantic early—Cold War responses to the Soviet bloc threat—including excessive development of our nuclear-weapons arsenal, the adoption of an unsupportable basic nuclear strategy, and excessive readiness to use nuclear weapons early in any conflict.
2. We must be particularly concerned about the possibility of a strategic surprise first-strike cyber attack in the long run. It is unclear whether such capabilities in cyberspace will ever be developed (they might be), and such an attack would be extremely difficult to detect in advance, so sensitivity to the possibility of one would lead to all sorts of high-alert, potentially overreactive postures on the part of the United States (and a possible opponent)—the worst situation for keeping deterrence stable, as was noted in the study of deterrence during the Cold War.
3. International cooperation to deal with the cyber attack threat is needed because cyberspace is a transnational resource, intended not to be threatening but to be helpful and liberating. It creates a high level of interdependence, so threats that emerge from it must be approached in a multilateral, cooperative fashion. It involves a greater degree of interdependence than that experienced by the antagonists during the Cold War, which led them to develop many elaborate arms-control measures. It is like the interdependence and international cooperation that are now being used or pursued to deal with international terrorism, global warming, and international epidemics and thus is well within our capacities.

Categorizing and Understanding Offensive Cyber Capabilities and Their Use by Gregory Rattray and Jason Healey

1. Offensive cyberoperations can be characterized in many ways. For example, they may be overt, covert, or somewhere in between; or the attacker and the defender or neither can be national military or can be a group with many different kinds of relationships.
2. Many (perhaps even most) of the forms of offensive operations have yet to be seen, and the future of conflict in cyberspace is likely to be very different from the past.
3. The battles of the cyber future may not be "cyber Pearl Harbors" or "digital 9/11s" but may be more analogous to a force-on-force Battle of Britain, a massive support to kinetic operations like the Battle of St. Mihiel, or a long, hard slog over years like the war in Vietnam.

A Framework for Thinking About Cyber Conflict and Cyber Deterrence with Possible Declatory Policies for These Domains by Stephen J. Lukasik

1. A set of long-range security goals suggest 11 unilateral U.S. declarations to initiate processes for the protection of the cybercommons. The declarations are based on the accepted structure of sovereign states as the mechanism to propagate the objectives through eventual international agreements.
2. The declarations assign to each sovereign jurisdiction the responsibility for eliminating the distribution of malware and the capturing of computers for use as botnets in it and the responsibility for attaching a state label to each packet leaving it.

3. The declarations attach to each state that allows harmful packets to leave it potential complicity for any harm suffered by a recipient of the packets. It calls for adjudication of disputes arising from such allegations with appropriate international mechanisms recognized by the parties to such disputes.
4. The declarations imply that attack attribution need go only deep enough to identify the sovereign entities that allowed harmful packets to leave them. But holding "innocent" transit states complicit requires all states to inspect packets coming into them for potential harm and, by implication, to reject them.

Pulling Punches in Cyberspace by Martin Libicki

1. The laws of war do not map very well into cyberspace, because of the potentially large differences between what operations were intended to do, what they actually do, and what they have been perceived to do.
2. Several of the factors that should persuade a state to pull its punches in cyberspace, such as the difficulty of reconciling operations with a state's narrative, the fear of escalation, and the occasional need to take back an action, apply in the physical world but are strongly influenced by the many ambiguities of cyberspace operations.
3. A sub-rosa response to an attack of uncertain effect and attribution has much to recommend it, but it means abjuring attacks on many types of targets. Reliance on sub-rosa responses can promote a lack of accountability among operators.

2.4.3 Group 3—Law and Regulation

Cyber Operations in International Law: The Use of Force, Collective Security, Self-Defense and Armed Conflicts by Michael N. Schmitt

1. The law governing when a cyber operation is a violation of the prohibition of the use of force in the UN Charter and customary international law is unclear. Thus, policy looms large, especially as one may not be able to predict accurately whether other States will deem a given action a violation.
2. The law governing when a State may respond kinetically in self-defense pursuant to Article 51 of the UN Charter and customary international law is relatively clear: the attack must cause (or be intended to cause) death, injury, or damage to property before such a response is lawful. However, States are unlikely to accept that limit in the face of a cyber operation that does not have such consequences when directed against critical assets. Thus, the law should be expected to evolve as State expectations and attitudes crystallize.
3. The law of armed conflict is generally adequate to handle a cyber operation mounted during hostilities. The major point of contention is whether an attack directed against the civilian population or civilian objects is unlawful if it does not injure or kill civilians or damage civilian property. In the view of the author, such operations are lawful.

Cyber Security and International Cooperation by Abe Sofaer

1. Cyber insecurity is an important and expensive problem that is inherently transnational, adversely affects all users worldwide, and is caused by many major players, including parties inside the United States.
2. No state (or group of like-minded states) will be able to deal effectively with all the major aspects of cyber insecurity through defensive and offensive measures.
3. International cooperation is likely to contribute to enhancing cybersecurity in some but not all aspects of current concern through agreements that avoid attempts to regulate inappropriate areas of concern (espionage and aspects of warfare), that seek objectives and use methods consistent with U.S. political and privacy values, and that maintain current, private, professional standard-setting activity rather than transferring such functions to government officials, national or international.

The Council of Europe Convention on Cybercrime by Michael A. Vatis

1. The Council of Europe's Cybercrime Convention has been an effective tool for fostering international cooperation on investigations involving computers and digital evidence. Because of the Convention, more countries have passed substantive laws addressing cybercrime and improved their cyber investigation capabilities, and parties to the Convention assist each other more rapidly and frequently.
2. The principal shortcomings of the Convention are its narrow membership (mostly European countries and the United States; Russia and China are not parties) and the lack of an enforcement mechanism if a country refuses to lend assistance when requested.
3. The Convention therefore could be made more effective by increasing its membership and by imposing costs of some sort on states that refuse cooperation without a legitimate, credible reason. While getting parties to agree to impose any kind of sanctions on uncooperative states seems unrealistic, public exposure of a state's lack of cooperation might have some salutary effect. Moreover, the U.S. could announce that, in the case of highly damaging attacks, it reserves the right to engage in unilateral self-help (such as cross-border searches of computers, or perhaps even counter-attacks on computers responsible for the attacks on computers in the U.S.) when the country from which the attacks appear to be emanating refuses to cooperate and provides no legitimate, credible reason.

2.4.4 Group 4—Psychology

Decision Making Under Uncertainty by Rose McDermott

1. Psychological factors are a critical part of understanding the perception of threat, and the kinds of systematic biases that can influence decision makers when they contemplate how to respond.

2. Overconfidence presents a pervasive and endemic problem for decision makers with regard to attribution in particular.
3. The anonymous nature of cyberspace and the speed with which processes of social contagion can spread information like a virus highlights the fact that deterrence no longer offers a viable strategic response for the uncertainty which characterizes this domain; rather, analogies drawn from the spread of infectious disease provides a more helpful model in thinking about designing more effective response strategies.

2.4.5 Group 5—Organization of Government

The Organization of the United States Government and Private Sector for Achieving Cyber Deterrence by Paul Rosenzweig

1. The potential U.S. government responses to a cyber incident span the whole of government and are not limited to cyber responses.
2. Private sector cybersecurity suffers from the "tragedy of the commons," so some form of collective response is essential.
3. Global supply chain security is weak, and a substantial threat from hardware intrusions has yet to be systematically addressed.
4. Policy makers should consider formalizing public–private cybersecurity cooperation through a publicly chartered nonprofit government corporation akin to the American Red Cross.

2.4.6 Group 6—Privacy and Civil Liberties

Civil Liberties and Privacy Implications of Policies to Prevent Cyberattacks by Robert Gellman

1. The civil liberties and privacy implications of potential policies and processes to prevent cyber attacks raise a host of unbounded, complex, difficult, and contested legal and constitutional issues.
2. Cyber-attack prevention activities will at times make use of the surveillance authority given to the federal government, and the law of surveillance is famously complex. One particularly important element is the absence of a constitutionally recognized expectation of privacy in a person's records held by a third party. The growing importance of third-party storage on the Internet and the technological obsolescence of many privacy statutes increases the tension between communication privacy and cyber attack prevention activities based on surveillance.
3. Anonymity on the Internet is prized by many Internet users for various reasons. A general constitutional right to anonymity has not been clearly defined, and conflicts are likely to arise between cyber attack prevention activities that attempt to

identify users and the interests of those who seek anonymity for whistleblowing, political, or other purposes.
4. The Privacy Act of 1974, the main information-privacy law applicable to the federal government, implements principles of fair information practice. The act, which applies to intelligence and law enforcement agencies, strikes a balance between competing objectives by allowing a partial exemption for the agencies. Similar exemptions would probably be available for cyber attack prevention activities.
5. Licensing of computer users, computers, or computer software is a possible response to cyber attack prevention needs. The United States has experience in licensing people and equipment in a way that generally balances due process interests of individuals with the government's need to function. However, a governmentally established identification or authorization prerequisite to general Internet access would be controversial. The authority of the federal government under the Commerce Clause (in Article I, Sect. 8 of the U.S. Constitution) is likely to clash with First Amendment interests, with much depending on the specific details of any regulatory scheme.

2.4.7 Group 7—Contributed Papers

Targeting Third Party Collaboration by Geoff Cohen
Note: Cohen's paper was awarded First Prize in the NRC Prize Competition for Cyberdeterrence Research and Scholarship "for original first steps in addressing the problem of third-party contributors to cyberinsecurity."

1. Existing cybercrime against U.S. private and public interests is a more pressing threat than future cyberwar.
2. Successful cyberattacks can only occur with the (possibly unwitting) collaboration of many US-based third-party infrastructure providers, such as ISPs, network operators, certification authorities, hosting providers, name registrars, and private individuals.
3. Law and policy need to be adjusted to encourage or enforce more aggressive monitoring, notification, and resolution of computer security issues, across all third party participants.

Thinking Through Active Defense in Cyberspace by Jay P. Kesan and Carol M. Hayes
Note: Kesan and Hayes' paper was awarded Honorable Mention in the NRC Prize Competition for Cyberdeterrence Research and Scholarship "for raising important issues regarding active defense in cyberspace."

1. Is active defense technologically feasible? Active defense technology exists and has been steadily improving in accuracy, but it may need further improvements before an active defense system can be implemented.

2. When would active defense be appropriate? Given various legal and practical considerations, active defense is probably most suitable as a response to denial-of-service attacks.
3. Who should be in control of active defense? For the purpose of consistency in implementation and to avoid escalation problems, the government should oversee active defense rather than having each firm responsible for making decisions about cyber counterstrikes case by case. Legal concerns and alternatives should also be considered, as should a potential process for an active defense program.
4. How can innocent third parties be protected? Liability rules should be in place to protect oblivious intermediaries whose systems are inadvertently harmed by cyber counterstrikes aimed at an attacker who had compromised the intermediary systems.

References

Cha, A. E., & Nakashima, E. (2010). Google China cyberattack part of vast espionage campaign, experts say. *Washington Post*. 14 January 2010.

Corrin, A. (2010). DISA to establish safe haven outside the Internet. *DefenseSystems.com*. 12 February 2010. Available at http://defensesystems.com/articles/2010/02/12/disa-dmz.aspx?s=ds_170210.

Joint Chiefs of Staff (2004). *The national military strategy of the United States of America*. Available at http://www.strategicstudiesinstitute.army.mil/pdffiles/nms2004.pdf.

Markoff, J., & Kramer, A. E. (2009). U.S. and Russia open arms talks on web security. *New York Times*. 13 December 2009. Available at http://www.nytimes.com/2009/12/13/science/13cyber.html.

National Research Council (2002). In *Cybersecurity today and tomorrow: pay now or pay later*. Washington: The National Academies Press.

National Research Council (2007). In *Toward a safer and more secure cyberspace*. Washington: National Academies Press.

National Research Council (2009). In *Technology, policy, law, and ethics regarding U.S. acquisition and use of cyberattack capabilities*. Washington: National Academies Press.

National Research Council (2010). In *Proceedings of a workshop on deterring cyberattacks: informing strategies and developing options for U.S. policy*. Washington: National Academies Press (material from the NRC is reprinted with permission, courtesy of the National Academies Press, Washington).

Sofaer, A. D., & Goodman, S. E. (2000). *A proposal for an international convention on cyber crime and terrorism*, Center for International Security and Cooperation, Stanford University. August 2000.

The NRC Letter Report for the Committee on Deterring Cyberattacks: Informing Strategies and Developing Options for U.S. Policy (2010). Available at http://www.nap.edu/openbook.php?record_id=12886&page=2.

U.S. Department of Defense (2006). Deterrence operations: joint operating concept, Version 2.0, December 2006. Available at http://www.dtic.mil/futurejointwarfare/concepts/do_joc_v20.doc.

White House (2009). Cyberspace Policy Review. Available at http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf.