

# Stanton Nuclear Security Fellows Seminar

---

## PANEL 3: Theoretical Perspectives

### 1. Rebecca Ward, BCSIA

#### *A Game Theoretic Approach to Safeguards Selection and Optimization*

The risk of nuclear proliferation has elevated in concert with increased global interest in nuclear power and a divergence in commercial fuel cycle technologies. Together these factors have placed heavy demands on the International Atomic Energy Agency (IAEA), the organization tasked with verification of peaceful nuclear activities.

While the IAEA's workload continues to grow, its available resources available remain relatively stagnant. Even against an increased threat backdrop, the IAEA's verification budget rose from approximately \$145 million in 2007 to \$160 million in 2010, a rate only marginally higher than inflation.<sup>1</sup> This budget constraints have spurred efforts by the IAEA to increase efficiency by moving away from the traditional, prescriptive safeguarding approach, where the application of safeguards varies little from state to state regardless of perceived threat or size of the nuclear program. Traditional safeguarding strategies also rely heavily on on-site inspections, which are costly and can be inefficient.

In an effort to improve safeguarding efficiency, researchers have developed many analytical techniques to perform proliferation pathway analysis, primarily based on probabilistic risk assessment techniques. While probabilistic techniques are valuable for describing fundamentally random events, like natural disasters, they may be ill equipped addressing adversarial problems, like proliferation. A 2010 National Research Council report questions the use of probabilistic techniques for adversarial risk analysis, noting that data in this area is too scarce to characterize adequately the threat or consequences of an attack. Further, the study suggests that probabilistic techniques may not fully capture the behavior of intentional actors, like a malevolent state or terrorist.<sup>2</sup> Instead the report suggests that a game theoretic approach to intelligent risk analysis may be more appropriate.

To that end, my work develops a game theoretic model of the interaction between a state-facilitated insider and an inspectorate tasked with detecting illegal behavior. The threat scenario assumes a state possessing multiple fuel cycle facilities—a gas centrifuge enrichment plant (GCEP) and an aqueous reprocessing facility. The adversary is either the state or some faction of insiders in collusion with the state who therefore enjoy access to state resources. In opposition to the adversary is an inspectorate, on which a budget constraint is imposed. The inspectorate's goal is to detect any misuse or diversion at

---

<sup>1</sup> “IAEA Regular Budget for 2010”, IAEA, accessed November 1, 2012, <http://www.iaea.org/About/budget.html>.

<sup>2</sup> Committee to Review the Department of Homeland Security's Approach to Risk Analysis and Nation Research Council, “Review of the Department of Homeland Security's Approach to Risk Analysis” (Washington D.C.: The National Academies Press, 2010).

the facilities, and the adversary seeks to perpetrate such a misuse or diversion while evading detection. At present the model focuses implicitly on international facilities, but is being designed such that modular changes in input data could allow for the generation of meaningful results regarding insider theft at domestic nuclear facilities.

The substantive fruit of my work is the development of a computational tool that suggests optimal safeguarding strategies for a given inspector budget. The tool itself is a computer program written in the C programming language and executable on any personal computer. Hardcoded into the program is information about a suite of safeguarding options at GCEP and reprocessing facilities, based on current IAEA practices and safeguards technology currently under development. Each of these safeguarding options has associated tunable parameters that can be modified to change the inspectors' probability of detecting a diversion. For example, mass balance verification is a safeguarding technique routinely applied at GCEP facilities by the IAEA. This tool contains two parameters associated with mass balance verification—frequency and size of inspection team—that can be adjusted to increase detection probability. The trade-off, however, is that increasing inspection frequency or size of inspection team costs more, and the defender has a limited budget. The budget is selected by the user of the tool from a pre-determined set of options. Similarly, information about diversion and misuse strategies at the facilities is encoded into the tool. These options are based on diversion strategies commonly presented in the literature and proliferation strategies that IAEA techniques are explicitly designed to counter. For a given inspector budget, the model outputs the optimal inspector safeguarding strategy, the optimal adversary diversion/misuse strategy, and the overall probability that the diversion/misuse will be detected.

As mentioned previously, at the heart of the model is a game theoretic solver. A two-person, zero-sum, simultaneous-play game was formulated to describe the adversarial scenario outlined above. The simultaneous-play nature of the game implies that both the inspector and the adversary know *a priori* the options available to the other player, but they do not know what strategy that player will select. Simultaneous play was assumed because it faithfully describes many critical safeguarding techniques, such as random inspections. The equilibrium strategies for both players are found by solving the game, which is done using the fictitious play algorithm. Fictitious play is a myopic, iterative algorithm that converges to equilibrium in the infinite limit for two-person, zero-sum games.<sup>3</sup>

While the game theoretic solver performs the optimization, a simulation model of the enrichment and reprocessing facilities is used to populate the game's payoff matrix. The game solver selects a defender-attacker strategy pair and passes it to the simulation model. Using the parameters defined by the strategies, the simulation model computes the scenario detection probability and weights this probability by the attractiveness and quantity of material obtained. This value is returned to the game as the payoff for that strategy pair. In this way the solver iteratively calls upon the simulation model to calculate payoffs for different strategy pairs and uses these payoffs to compute the equilibrium payoff (and strategies). Inputs to the simulation model are derived from three sources: open-source literature,

---

<sup>3</sup> Julia Robinson, "An Iterative Method of Solving a Game," *The Annals of Mathematics*, Second Series, Vol. 54, No. 2 (1951): 296-301.

where available (e.g. vulnerability assessment for seals); first principles, when applicable (e.g. gamma detectors); and analogous measures or techniques in other fields that are readily available in the literature (e.g. video surveillance at train stations). While a substantial effort was made to find defensible, albeit not “correct” values for detection probability inputs, no such effort was made to calculate costs. Costs used in the model are invented and intended to have value only relative to one another. For example two inspectors cost twice as much as one, but no effort was made to quantify the absolute cost of two inspectors. The inspectorate’s budget is imposed by the game, and is done so by eliminating strategies whose cost exceeds the inspector’s budget.

The imposition of inspector budget is a critical element of this work for two reasons: it produces a higher fidelity model of the real situation faced by the IAEA and other resource-constrained inspectorates, providing for more realistic strategy recommendations; and it creates a platform to compute the cost efficiency of safeguarding measures. Cost efficiency is effectively the amount incremental detection probability provided by a technique or technology per dollar spent. One can imagine this metric being of particular utility when considering investment decisions in proposed safeguarding technology. This computational tool can be used to calculate the break-even cost for a novel safeguarding technology, or the cost above which the inspector does not select the technology, because the detection probability it contributes is not worth the money costs. Such an analysis could be performed before making safeguards R&D investment decisions, thus providing a rational basis for decision-making and reducing risk.

In addition to providing valuable information for safeguards R&D investments, this model generates output that can be of use to the IAEA and other cost-constrained inspectorates. The IAEA has very little latitude in changing its safeguarding strategies because in many cases it is legally constrained. These legal constraints were intentionally left out of this model, so as to explore the entire strategy space without restriction. Thus while strategies recommended by this model may not be of literal use to the IAEA, the values implied by these strategy recommendations may. Selecting the optimal safeguarding strategy requires many forced trade-offs (e.g. given a certain inspection budget, it is better to inspect more frequently, or with a larger team less frequently), and the results of these trade-offs hold useful information about the relative value of different safeguards parameters. Further, this tool outputs optimal safeguarding strategies, given certain adversary options. By altering adversary specifications and observing the resultant changes in the optimal inspector strategy, much can be learned about the relationship between adversary capability and intent and ideal safeguarding strategy.

The aim of this research is to model faithfully the adversarial dynamic between a cost-constrained inspectorate and a proliferant state and recommend safeguarding strategies accordingly, so as to improve inspector efficiency and reduce the global risk of proliferation.

## 2. Andrew Coe, CFR

Governments face difficult, important decisions about how to prevent or respond to the proliferation of nuclear weapons. Unfortunately, there is no consensus on how the United States should deal with proliferating states such as Iran and North Korea, or on what it should seek in negotiations over revisions to the non-proliferation regime. The confusion stems in part from the absence of a systematic understanding of why the interactions between the US and proliferating states sometimes end in a negotiated deal and other times end in costly containment or even war. And there is even less understanding of what kinds of revisions to the non-proliferation regime might be not only desirable but also *agreeable* for most states.

My aim is to redress these deficiencies by developing rigorous theories of these interactions and gathering the data necessary to test them. The goal is to build a coherent, validated framework for understanding these interactions that should enable policy-makers to make better-informed decisions. Below, I've described three associated papers that I will work on while a Stanton Fellow. The first investigates why it is that bargaining over weapons programs sometimes breaks down, using this to understand the roots of the conflict between Iraq and the US after the first Gulf War. The second, to be done with my collaborator Jane Vaynman, explores the evolution of the non-proliferation regime since the end of the Cold War. These first two papers involve highly technical work and are thus intended for academic audiences. However, the third paper uses the first two, and earlier academic work, as rigorous foundations for analysis that is directly relevant to US policy, and is intended for policy audiences. The academic work allows us to make well-founded projections of how the non-proliferation regime might evolve in the future, and to explore how the US might go about shaping this evolution.

### **The Limits of Non-Proliferation and the Origins of the Iraq War**

Why do negotiations between the US and a state suspected of pursuing weapons of mass destruction sometimes fail, and lead to costly containment or even war? In principle, both sides would be better off if they could avoid these costs by agreeing on a suitable deal: non-proliferation in exchange for certain concessions from the US, such as a security guarantee. Of course, if it were hard for the US to be sure that a proliferant was not covertly pursuing WMD, then this would undermine the viability of a deal, but again in principle, it is always possible for the US to be sure. For example, even in the case of Iraq, the US could have blanketed the whole country in inspectors, ensuring that any suspicious site could be continuously watched. And yet this didn't happen. Why?

This paper will analyze two potential obstacles to such deals: the (potential) trade-off between the proliferant's regime security and the US's confidence that it is not cheating on inspections; and the costs to the US of the continuation of the proliferant's regime. Both of these factors have been noticed before, but their relationship to the overall viability of a deal is not understood. My approach is to develop a game-theoretic model of the negotiations between the US and a potential proliferant. In it, the two sides bargain over the issues in dispute between them, while the proliferant may covertly try to develop weapons of mass destruction, and the US may take steps to render its efforts more expensive or less likely to succeed. The two sides can also agree to an inspections deal, which the proliferant might then try to evade. Using this model, I can determine the conditions under which the two

obstacles will make it impossible to enforce a deal, or render the US unwilling to offer a viable one, and when the latter will lead to containment versus war.

This analysis should enable me to offer a coherent explanation for how it is that the US and Iraq wound up locked in a costly confrontation throughout the 1990s, and why it is that, when containment came to an end, the US chose war rather than toleration. Remarkably, despite all the books and articles written on the origins of the 2003 Iraq War, there is still no explanation of the war that is consistent with all (or even most of) the facts and that makes clear why the situation evolved the way it did. Comparing the fit between these facts and the predictions of this model provides a good test of its validity, and is also interesting in its own right as a means of understanding the long road to the war. It is particularly important to understand whether the Iraq War is best thought of as some kind of aberration—owing to the unusual characteristics of the Bush administration or the reaction to 9/11—or instead something that might occur again.

The analysis should also allow an assessment of some fundamental difficulties in the non-proliferation enterprise, at least as it is currently institutionalized. Much of the discussion of how to strengthen non-proliferation has revolved around where to draw the line between permitted and prohibited nuclear activities, and how intrusive IAEA inspections should be. What if instead there was a different way to structure the relationship between the IAEA and concerned powers such as the US, so as to lessen the tradeoff between a proliferant's security and the intrusiveness of the inspections it must tolerate? Thus, the analysis might yield new ideas for reforms of the non-proliferation regime in order to overcome these difficulties.

### **Unipolarity and the Difficulties of the Nuclear Non-Proliferation Regime**

*In this paper, Jane Vaynman and I will investigate the evolution of the non-proliferation regime since the end of the Cold War. How do we understand its successes and failures with respect to individual states' pursuit of nuclear weapons, and others' support thereof? Why have revisions to the regime that would strengthen it been resisted? What lies behind the difficulty of achieving international consensus on how to deal with states like Iran and North Korea?*

This is a natural extension of some earlier work Jane and I did on the origins and stability of the regime during the Cold War. Our theory is that the Cold War successes of the NPT were built on two underlying factors: the willingness of many states to eschew nuclear weapons so long as others did, and the collusion between the superpowers to pressure any other states into doing so as well. Our intent is to assess how the relevant inputs identified by the theory changed after the Cold War ended, and then derive predictions for what should have happened with respect to behavior under the regime. We can then compare these predictions to the historical record of what actually happened. If they match well, then the theory is supported and we will have a rigorous basis for projecting the future evolution of the regime, and for assessing the pros and cons of alternative policies for revising it (see the description of the next paper, below).

Since we have already developed the theory, the bulk of the work lies in figuring out how to measure the relevant inputs, and in assembling a comprehensive dataset of states' behaviors under the regime

since the end of the Cold War. According to our theory, the inputs that must be measured are: how the dimensions of policy that are most salient internationally have changed (e.g., the shift from capitalism vs. communism to democracy vs. dictatorship); how the preferences of individual states over these policies have changed (e.g., the liberalization of many former members of the Soviet bloc); how the value to potential clients of security guarantees from potential patrons has changed (e.g., Russia can no longer provide credible guarantees to countries not in its neighborhood); and finally how the efficacy of the regime's machinery has changed (e.g., how good the IAEA and Western intelligence services are at detecting "cheating"). It is easy to make informal judgments about how these inputs have changed, but our challenge is to develop objective measurements of each.

Once we have these measurements in hand, we will derive and test predictions from the theory about several aspects of states' behavior. Which states do we expect to pursue nuclear weapons, and which will eschew them? Which states will begin or end an alliance with a patron (e.g., the US), and which will seek to renegotiate the terms of their alliance? Finally, which states will support certain revisions to the regime, and which will resist them? While some of this data has already been collected, we will need to construct an original dataset of states' expressed preferences over various mooted revisions to the regime.

### **The Past, Present, and Future of the Nuclear Non-Proliferation Regime**

Please see Jane Vaynman's memo for a description of our policy paper.

### 3. Jane Vaynman, CFR

Adversarial states compete for resources, power and influence, but they also often cooperate to restrain that competition using a variety of institutional arrangements. Since 1945, that competition has included a new threat, nuclear weapons. Not long after, states responded with a new institution, the Nuclear Nonproliferation Treaty (NPT). The US and USSR/Russia have also signed a series of bilateral treaties intended to limit their nuclear arsenals and lower the risk of use. In my research, I seek to understand how risks of nuclear war and nuclear proliferation have been addressed through institutions. What do we know about the forms and functions of existing security institutions that can help us craft future cooperation with nuclear-armed adversaries or improve on structures we have today to limit proliferation?

To address these questions, I am working on two related research projects, both of which include products for academic and policy audiences. The first analyzes nuclear arms control in the broader context of adversarial cooperation, including treaties to limit conventional armaments, stop WMD proliferation, or prevent access to strategic territories. Here, I focus on how states design agreements which share information with their adversaries. The second project focuses on one key institution, the NPT regime. In this collaborative project, Andrew Coe and I explore the stability of the NPT and its uncertain future as an effective institutional response to proliferation. In his memo, Andrew describes the academic paper we are working on, while I discuss our policy-oriented paper below.

#### *Enemies in Agreement: Information Sharing in Adversarial Security Cooperation*

Since the 1960's, countries have used formal agreements to limit nuclear weapons. Yet the experience with arms control is much broader than a nuclear lens suggests. It reveals a diverse range of options in how cooperation agreements (nuclear or otherwise) have been designed, particularly when it comes to the delicate task of sharing information with adversaries. Many scholars and practitioners will agree that agreements with monitoring and verification measures are more likely to be effective, as cheating can be deterred or readily detected. However, some treaties fail to include these provisions, and we actually observe that in creating security agreements states include a variety of features, from extensive monitoring to barely self-reporting. Why are states sometimes able to create intrusive and highly enforceable treaties, but at other times appear to settle for less effective mechanisms? The goal of this project is to understand not only why adversarial states sometimes choose to cooperate, but also why cooperation takes the form that it does, particularly with regard to how states share information.

I argue that states design different kinds of treaties, including ones that seem weak or unenforceable, because of different beliefs they have about each party's incentives to cheat on an agreement. How certain states are of their beliefs is the core story. I develop a theory analyzing the tradeoffs between benefits of cooperation, costs of enforcing it, advantages of being the state which cheats, and risks of being cheated on. I show that when beliefs are highly certain (that the adversary definitely has low incentives to cheat or high ones) we should expect to see treaties with little enforcement or no treaties at all. But, when states are uncertain about the adversary's incentives (incentives might be high or low), under some conditions they prefer to essentially hedge and pick mid-range options for information sharing provisions, forming semi-enforceable treaties. In several particularly interesting cases, when

uncertainty about likely cheaters increases, treaties with high monitoring are signed. In other words, uncertainty actually allows for some kinds of agreements to happen where they might not have otherwise.

Uncertainty about the opponent's incentives to cheat on an agreement can be driven by a number of factors, including changing power, new technology, etc. In my project, I focus on one key source of this uncertainty: domestic political volatility. New leadership in an opponent state or powerful opposition parties can create uncertainty about foreign policy directions. For example in the case of the Intermediate Nuclear Forces Treaty, the shifts created by Gorbachev's leadership contributed to greater uncertainty among US policy makers about Soviet intentions. At the same time, documents reveal that Gorbachev's government had concerns about the role of other US domestic actors who opposed the Reagan security agenda. I argue that uncertainties created by these domestic factors influence both the ability of states to agree to arms control and affect the form of the agreement that they conclude.

To test these ideas, I pursue empirical analysis of nuclear and other security agreements on both quantitative and qualitative dimensions. First, I develop an original dataset of all adversarial security cooperation agreements. To date, no data have been systematically collected for such agreements. The dataset, which includes about 250 agreements from 1800-2010, allows me both to understand the full range of variation in agreement design, and assess how agreements on nuclear weapons related issues compare to other treaties. Having coded information provisions in each treaty, I will next be using statistical techniques to evaluate my theoretical claims on the relationship between domestic politics, uncertainty, and agreement design.

Second, I am investigating several key cases in nuclear cooperation, using declassified primary sources, secondary analysis, and interviews with decision makers. As previewed above, the case of the INF Treaty provides the strongest illustration of the dynamic laid out in the theory - as uncertainty increases between two states; they are able to create a highly monitored agreement where one was previously not attainable. A second case study extends the analysis to the multilateral context by comparing the Nuclear Non-Proliferation Treaty and the Biological Weapons Convention, which both seek to limit the spread unconventional weapons but vary considerably in their institutional structures – the BWC has no monitoring mechanism while the NPT features extensive inspections. Finally, I test my theory outside the US-Russia context by examining nuclear confidence building agreements between India and Pakistan.

My project has implications for further research in security studies, institutions scholarship and policy making. For an academic audience, I provide theory development which builds on recent work in the treaty design literature, as well as testing using original data. For policy audiences, my project brings an institutions lens to what has traditionally been a security topic, suggesting new ways of looking at questions in the nuclear sphere. If decision makers were better informed about the sources of and constraints on agreement design choices, they might be able to create better agreements. For example, as the US and Russia move towards the next round of arms control negotiations, perhaps on tactical nuclear weapons, how will political opposition within Russia affect the kind of agreement that Russia is willing to – or able to – sign with the US? Will the uncertainty created by domestic unrest allow for more



security information to be exchange between old rivals? As another example, one clear extension of my project is to consider how adversarial states design cooperation when the source of the threat is not one another but a non-state actor, including agreements on terrorism, smuggling of dual-use technology, or even cyber-attacks. My research will give better insight on the factors which drive state choices over information sharing, and will allow me to develop specific recommendations for policy makers on how to best accomplish cooperation goals in both traditional nuclear security and emerging areas.

### *Alternative Futures of the Nuclear Nonproliferation Regime*

The future of the NPT is murky – new nuclear aspirants challenge its effectiveness, growing civilian programs challenge its rules, and attempts at reform have stalled from one review conference to the next. Scholars and policy makers have suggested ways to reform the treaty, or even revise its basic bargain, but there is little analysis on how these options might play out longer term, as well as a tendency to overlook how different types of states play roles in regime. What forms might this institution take in the future? How will the evolution of the NPT, and the way the United States supports its enforcement, affect the likelihood of proliferation?

In this paper, Andrew Coe and I apply our theory of NPT stability to an assessment of the treaty's future – or more accurately, to several alternative futures which policy makers might see emerging in the next decade. The theory, developed and tested in separate academic papers, allows us to provide a policy analysis of key elements in the regime, including participation in the treaty, major powers and clients, and trade-offs which motivate nuclear ambitions in some states. Our goal in the broader project is to build a new understanding of how states cooperate on nuclear weapons restraint, and this paper takes a first step in applying that perspective to policy questions. The audience for this project is both policy makers and social scientists working on policy-relevant research, and we will seek publication in a venue which is accessible by both groups.

We lay out three alternative futures: (1) the new duopoly; (2) great power pullback; and (3) US the enforcer. First, we imagine a scenario where a rising China becomes increasingly interested in dominating the policies of its neighbors. To maintain this dominance, China will prefer non-nuclear client states and become an increasingly willing supporter of NPT enforcement. The superpower collusion of the Cold War (discussed in our first academic paper) may be reborn, and while the rise of China might bring numerous other security concerns for the US, with respect to nuclear proliferation, stability and enforcement would actually be improved.

Second, we propose a scenario where the NPT regime is maintained less by its institutional rules and more by a shift in state interests in nuclear weapons. In this future, the US is willing to give smaller states greater autonomy as regional powers, pulls away from direct support of democratization efforts, and may even decrease US military presence abroad. With fears of regime change and US unilateral action diminished, states will still use civilian nuclear programs to hedge, but few will see a clear security need to develop nuclear weapons capabilities.

In the third scenario, the US is not willing to pull back from other global policy interests while seeking nonproliferation goals. States continue to have strong incentives for a breakout nuclear capability, and

the US often unilaterally enforces the NPT. The US uses greater coercive resources to police nuclear aspirants, and we would expect strong efforts by the US and its allies to add tools to the regime, such as pressuring states to accept the Additional Protocol. We suggest that this maybe the most dangerous scenario, as it is not clear whether US efforts will be able to match the pace of growing incentives among weaker states to turn clandestine programs or even hedging positions into nuclear weapons capabilities as quickly as possible.

In the course of the analysis, we identify key policy decisions being made today or in the immediate future which may contribute to these scenarios unfolding. Our goal is not to argue that certain decisions are more advisable, but rather provide policy makers with a way to understand the connections between today's NPT regime challenges, the options they face, and the implications of those policies for future nuclear proliferation dynamics.

## 4. Jeff Kaplow, RAND

### *The Credibility of the Nuclear Nonproliferation Regime*

Many in the nonproliferation community have attempted to sound the alarm about the decline of the Nuclear Nonproliferation Treaty (NPT), often framing their arguments in terms of the credibility of the regime. These analysts argue that efforts by some states to circumvent the treaty, ongoing non-compliance within and outside of the treaty, and the failure to punish such non-compliance, make others less likely to comply with the regime in the future. Commonly cited harbingers of the regime's decline include the USIndia nuclear cooperation deal; the development of weapons outside the treaty by Israel, India, and Pakistan; the development of weapons inside the treaty by North Korea, Iran, and Iraq; the exit of North Korea from the treaty; and the failure of the international community to punish these transgressions.

Few of these analysts, however, explicitly ask the underlying question: does the track record of an international security treaty affect future compliance? The answer is important for decision-makers who must weigh short-term policy goals against their long-term and uncertain impact on the credibility of the nonproliferation regime. For example, should the United States make deals with non-NPT states that will help to safeguard nuclear materials, even at the cost of lending legitimacy to those outside the regime? Should Iran's non-compliance be treated as a unique case—subject to as much negotiating flexibility as possible—or should the international community worry that concessions will set a dangerous precedent for future proliferators?

If the past performance of the regime does help to determine its future effectiveness, it is not clear that all blemishes on the regime's record will have the same impact on future compliance. There are at least three varieties of non-compliance with the regime. Abstention is the failure of important states, such as India, Pakistan, and Israel, to join the regime in the first place. The presence of these nuclear states outside of the regime, and the fact that they have largely avoided sustained punishment from the international community, casts a shadow over those that have chosen to join and comply with the regime. Violation is the pursuit of nuclear weapons while a member of the NPT, which sets a dangerous precedent for other member states. Finally, exit is the decision to leave the NPT and pursue nuclear work free from legal constraints. Each form of non-compliance raises unique issues of precedent and the violation of international law.

### *Research Design*

My proposed research design has three parts: detailed case studies of the ways in which states have responded to questions of regime compliance in the past, a quantitative analysis of the role that regime performance plays in the decisions of states to engage in nuclear weapons efforts, and a survey of policy elites to determine what forms of noncompliance are most likely to affect the credibility of the regime. Each contributes to a test of the hypothesis that the perceived effectiveness of the nonproliferation regime today affects state compliance with the regime in the future.

For the case studies, I will select a small number of events that seem to cast doubt on the efficacy of the nonproliferation regime or that appear to bolster the regime. For each episode, I will examine in detail

the reaction of states at the time and for several years afterward, with a focus on how the event colored state views on the future effectiveness of the treaty and whether it prompted states to question their own treaty compliance. Candidates for case studies include the exposure of nuclear weapons efforts by states that were party to the NPT (such as South Korea, Taiwan, Iraq, or Iran), the failure to include major nuclear states within the NPT (such as the Pakistani and Indian nuclear tests), the decision by North Korea to exit the treaty, and the decision by South Africa to give up its nuclear weapons program and join the NPT.

I also will conduct statistical tests to determine whether episodes of non-compliance with the regime are associated with an increased probability that other states will themselves pursue nuclear weapons. This will require a new data collection effort to identify when each incidence of suspected non-compliance first became known among other NPT member states. I plan to gather these data from government archives, searches of news sources, and the rich secondary literature on specific nuclear weapons efforts.

Finally, I plan to incorporate survey data that provides new insight into the way policy elites connect treaty violations to the treaty's future effectiveness. A project at the Laboratory on International Law and Regulation (ILAR) at the University of California, San Diego, has fielded a web-based survey that asks those with a policy background to assess the effectiveness of a hypothetical nonproliferation treaty under several different scenarios. While the availability of this data source will depend on survey response rates and other factors, these data could provide useful context to ongoing policy arguments about how seriously to take threats to the nonproliferation regime's credibility.

### *Contribution to the Policy Process*

Strengthening the nuclear nonproliferation regime is an important US policy goal. Policymakers may have to make trade-offs, however, between defending the credibility of the regime and achieving other foreign policy goals. There are many examples of more immediate policy goals trumping America's long-term interest in maintaining the credibility of the nonproliferation regime: ending nuclear-related sanctions on India and Pakistan, signing a nuclear cooperation agreement with India, declining to press Israel to join a nuclear-weapons free zone in the Middle East, and failing to strongly punish North Korea's violation of and exit from the NPT.

My research will contribute to efforts to understand the importance of the credibility of the nonproliferation regime by examining the fragility of the regime: is it true that violations of the regime today will make the regime less effective in the future? This work will also offer some insight into what forms of non-compliance might put the regime most at risk. Together, the answers to these questions may allow policymakers to better calculate the costs and benefits of nonproliferation decisions that play an important role in maintaining international peace and security.