THE NEW ATLANTIS
A Journal of Technology & Society

Why Not Nuclear Disarmament?

Christopher A. Ford

In accepting the Nobel Peace Prize in Oslo on December 10, 2009, President Barack Obama described as "urgent" the effort "to prevent the spread of nuclear weapons, and to seek a world without them." In announcing the award two months earlier, the Norwegian Nobel Committee made special note of "Obama's vision of and work for a world without nuclear weapons," a vision that "has powerfully stimulated disarmament and arms control negotiations." But realizing this vision will require detailed consideration of several tough questions—questions that the movement for global nuclear disarmament has not yet seriously addressed. How the world will actually get there, and how it might be expected to *stay* there, must be thought through before we embark on a project of radical disarmament.

Some disarmament advocates believe that the moment has not yet arrived to worry about the nuts and bolts of how a disarmed world would actually work, such as the means by which the international community would be able to deter regime "breakout" by a country keen on using the disarmament of others to make itself the planet's only nuclear weapons holder. Indeed, the prospect of total nuclear disarmament is so transcendentally appealing in some quarters that such practical challenges and fundamental questions are only of distant interest, if any at all. Some disarmament advocacy groups are so far removed from reality that they argue that reliance upon "deterrence" in security policymaking should end altogether—that it is possible not just to end the use of nuclear weapons in deterrence, but that we need never rely upon anything to "deter" anything else ever again.

But back down here on Earth, serious advocates of disarmament are increasingly coming to recognize that unless they can provide compelling responses to some crucial challenges, their goal will only further recede from sight. In part, the challenges are technical; in part, they are political; in part, they point to fundamental ethical concerns. They amount to more than just a critique of the disarmament movement's ends and proposed means. Taken together, these questions compose a broad agenda for disarmament research for the months and years ahead.

***Christopher A. Ford**, a* New Atlantis *contributing editor, is a senior fellow at the Hudson Institute. He previously served as U.S. Special Representative for Nuclear Nonproliferation, and as Principal Deputy Assistant Secretary of State for arms control, nonproliferation, and disarmament verification and compliance policy.*

Verification Challenges

The first and most common question arising in disarmament debates is how compliance could be verified. There is no clearly-defined rule for judging the uncertainty of verification; levels of uncertainty that might be entirely unacceptable in one context may be tolerable in another. The uncertainty stems from a mix of factors that ultimately all require judgment calls, such as the impact of a treaty violation on the military balance regulated by that treaty, the ease and speed with which other parties can respond to a violation after it has been detected, and the degree to which parties face incentives to cheat in the first place.

Here is how U.S. arms negotiator Paul Nitze defined the concept of "effective verification" in the 1980s, during the difficult negotiations with the Soviets over what became the Intermediate-Range Nuclear Forces Treaty: "If the other side moves beyond the limits of the treaty in any militarily significant way, we would be able to detect such violation in time to respond effectively, and thereby deny the other side the benefit of the violation." Later, during the negotiations on the original Strategic Arms Reduction Treaty (START), Secretary of State James Baker added the qualification that effective verification also entails being able to detect "patterns of marginal violations that do not present immediate risk to U.S. security." This so-called Nitze-Baker standard is as close to an official litmus test for effective verification as has ever been offered.

An outright ban on nuclear weapons would of course offer the greatest challenge for the effectiveness of verification. Global disarmament presents a very different situation than that confronting negotiators in the traditional context of U.S.-Russian arms control. To illustrate this difference, consider first a traditional arms control treaty: the new post-START agreement that President Obama and Russian President Dmitry Medvedev signed in Prague in April 2010. That treaty restricts each side to no more than 1,550 "operationally deployed" nuclear warheads. Even with agreement upon some augmented transparency and verification mechanisms, it will be very hard for the United States to be *sure* that Russia is not holding something back. But it might not be absolutely essential to have such certainty. Would U.S. national security interests be gravely imperiled if Russia in fact deployed a total of 1,566 warheads—an overage of merely one percent of the permitted total? Perhaps not. (Indeed, the new treaty allows bombers capable of carrying many warheads to be *counted* as though they bear only a single warhead. Whether the number of operationally deployed weapons will actually be much reduced by this

new agreement is not clear under such counting rules, and perhaps is not knowable at all; START counting rules seem to have been diluted.) This is where the *military significance* aspect of the Nitze-Baker test comes in: it allows for the possibility that if one's verification margin of error is sufficiently small, outcome divergences within that margin may be acceptable because they would not overturn the military balance that it is the fundamental ambition of the agreement to regulate. This is why effective verification does not require 100 percent certainty, and one of the reasons why it is possible to have arms control agreements between geopolitical rivals and ideological adversaries at all.

But let us now consider the challenge of "nuclear zero." The threshold of military significance in the context of complete disarmament may arise with the very first nuclear weapons developed. A sixteen-warhead uncertainty in the context of a Russo-American agreement in 2010 is one thing, but what if a state managed secretly to retain or develop sixteen weapons in a world in which no other states had nuclear weapons at all? Such an arsenal would have vast military import indeed. In fact, even a single weapon might be of great significance, especially if wedded to a long-range delivery system and thus potentially capable of incinerating any other country's capital city.

Context matters too, of course, and the threshold of significance would surely vary with the size and military power of such weapons' potential target. In the context of a global disarmament regime, the military significance thresholds of *all* relevant players would somehow have to be accommodated. The United States, large and powerful, might be willing to live with some uncertainty about other states' concealment of a few warheads. But how many nuclear weapons would it take to present a threat of "military significance" to a tiny country such as Israel? A global disarmament regime would thus create downward pressure on acceptable verification error margins. The presence or absence of missile (or other) defenses would also affect verifiability assessments, since robust defenses could help rob of significance small numbers of weapons sneaked past a verification regime. (This is why it is so perverse for disarmament advocates to oppose missile defense.)

In general, the smaller the arsenals, the more difficult it will be for verification to meet the Nitze-Baker standard. It is hard, therefore, to imagine a more difficult challenge for effective verification than a nuclear zero regime. Unless mechanisms were somehow developed to retard a violator's ability to quickly capitalize on his first few weapons as instruments of extortion or mass extermination—mechanisms such as widespread

and effective defenses, for instance, or a system whereby, in the event of violation, former nuclear weapons powers would be able to quickly reconstitute countervailing deterrent or reprisal forces—a nuclear zero regime would require a margin of error of, appropriately, nearly zero. Ensuring such high levels of certainty seems likely to be extremely challenging, making this question vital for disarmament research.

The Difficulties of Dismantlement

Dismantling a nuclear weapon under modern safety and security requirements is a rather difficult and demanding process. Nevertheless, it is a process that is well understood and that has for years been frequently practiced—at least in the United States, especially under the accelerated-dismantlement program directed by President George W. Bush and continued under President Obama. (The Pantex plant in Texas that handles this dismantlement in the United States, for instance, is said to be presently running at full capacity and faces a backlog that, without the construction of new facilities, could take fifteen years to exhaust.) And even more difficult than just dismantling the weapons is *verifying* their dismantlement.

To be sure, there are some efforts underway to address these verification issues. The British government, for instance, announced in 2007 its commencement of new detailed studies seeking to devise technical methods for verifying warhead elimination and establishing chains of custody for sensitive fissile materials derived from weapons. The U.K. has also called for expert-level discussions involving scientists from the nuclear laboratories of the five nuclear weapons states that are party to the Treaty on the Non-Proliferation of Nuclear Weapons (NPT), in order to explore possible routes toward verifiable elimination—though so far little has come of this; a preliminary meeting is said to have demonstrated how little the Russian and Chinese laboratories, in particular, are presently interested in transparency. Exercises conducted in 2007 by the United Kingdom and Norway—in which each party role-played the other, thus providing both sides an intriguing perspective on the tension between verification transparency and nuclear weapons security—have also explored questions related to managing challenge inspections of dismantlement activity. (Despite disarmament advocates' eagerness to cite the Anglo-Norwegian effort as proof that dismantlement verification will be possible, the project seems to have assumed complete good faith on both sides. Deception scenarios were apparently neglected, which hardly

seems wise in assessing the verifiability of disarmament by a potential adversary.)

American studies of transparency in warhead dismantlement began at the end of the Cold War under President George H.W. Bush, with particularly intense work being done in the late 1990s in anticipation of a possible START III. The U.S. Department of Defense began conducting Warhead Monitoring Technology Project exercises in 2001, for example, and U.S. experts have been doing some technological and operational development work on transparency measures for several years. The United States even conducted a fissile-material technology-transparency demonstration for a delegation of Russian scientists, and in the late 1990s the U.S. national laboratories began collaborating, at least fitfully, with their Russian counterparts on measures for verifiable warhead storage and transport tracking.

Nevertheless, there is much to do before even the simpler of the verification and chain-of-custody issues can be resolved. For instance, any answer to the problem of attribute monitoring—that is, of demonstrating the presence of weapons-grade plutonium in a declared container of ex-weapons material, yet without revealing information about that material (such as precise masses or isotopic ratios) that might be considered sensitive by its owner—would need to strike a precarious balance. It would need to absolutely protect sensitive host government information (about specific design secrets) yet also presumably provide enough data for diverse international observers to feel confident that real nuclear weapons were being dismantled, rather than just dummies, even dummies that contain nuclear material. To date, measures reportedly under examination by U.S. and U.K. laboratories have focused principally upon the first step of assuring the presence of nuclear material.

The various possessors of nuclear weapons today may each have different ideas of what constitutes sensitive information, and it still is far from certain at this point what data the rest of the international community would consider sufficient to give confidence in the reality of claimed dismantlement and to protect against "spoofing." For the dismantlement challenge to be overcome, nuclear weapons states would have to be willing to disclose enough information about designs and materials to satisfy the minimally-acceptable requirements for a credible verification data set.

Moreover, nuclear material derived from dismantled weapons would have to be handled in a way that both respects the sensitive information of its owner government and simultaneously provides an internationally-acceptable chain of custody and materials accountability. This could be

tricky, particularly if the objective is to verify the disposition of material corresponding to a specific *number* of warheads. To the extent that the masses of plutonium or uranium per warhead constitute sensitive information—which is certainly true in the United States—a verification system might face a dilemma: It could verify that a certain quantity of material had been presented and disposed of, *or* it could verify that a certain number of warheads-worth of such material had been processed, but it would be very hard to do *both* of these things at the same time, because permitting observers to derive mass-per-warhead numbers would itself compromise sensitive information.

In fact, monitoring disassembly might not really be worth the trouble, and conceivably might not be necessary if nuclear material were controlled in a way that steadily moved it out of a declared weapons stockpile into non-sensitive forms, making it accountable through more ordinary nuclear-safeguards methods. Ironically, this would require the disarmament verification regime to deliberately *refrain* from specific knowledge of warhead numbers—except at the end point, at which the certified absence of any non-safeguarded material would necessarily imply that the warhead count had reached zero. As we shall see, however, having any reasonable idea of when one has actually reached this point would require more detailed information about the dismantling country's total production than is presently available for anyone—and perhaps more detailed than it is possible to acquire.

Finding the Hidden

If an acceptable way of handling dismantlement verification could be worked out, global disarmament would still require solid assurance that what had been dismantled amounted to the *entirety* of all states' nuclear weapons arsenals—and that there was no fissile material tucked away somewhere that could secretly be used to make nuclear weapons.

The ability of any imaginable verification regime to search for hidden nuclear warheads or material beyond a state's declared stockpile would likely be quite limited. This limitation is partly technological in nature (sensors, for example, are very far from perfect) and partly procedural (complete certainty would require unbelievably intrusive inspections, and those inspections would be costly and personnel-intensive). It is quite easy to squirrel away nuclear weapons, and not much harder to conceal a small-scale production infrastructure. Nuclear material can be shielded and hidden handily; current technology can detect it only when within

mere feet or inches. Taking radiological samples from the surrounding environment is also of sharply limited use, especially in a former weapons production infrastructure: it can inform inspectors that nuclear material was formerly in a particular place, but cannot reveal where the material is now—and is in any case fairly unhelpful in a facility in which everything is already covered by isotopic traces from past weapons production.

The international community's track record in finding undeclared nuclear activities using longstanding International Atomic Energy Agency (IAEA) authorities is not very encouraging, and the agency itself has admitted that the tools at its disposal are insufficient in the face of determined concealment. The IAEA Additional Protocol—a voluntary measure that would expand the agency's inspection power—is still resisted by many governments, with countries such as Egypt, Brazil, Argentina, Venezuela, Syria, and Iran still holding out. If such countries consider even the limited inspection provisions of the Additional Protocol to be too onerous or expensive, it beggars belief that they would countenance an inspection and monitoring system capable of finding one or more carefully hidden nuclear weapons, even if it were possible to devise such a system. (When U.S. officials presented a robust and intrusive verification plan to North Korea in 2008 that was designed to provide assurances against Pyongyang's customary cheating on issues related to denuclearization, the North Koreans reacted with outrage and brought the talks to a standstill.) It might help ease such political and economic burdens if a system were implemented that focused verification efforts on governments considered particularly likely to cheat—as, in effect, the IAEA is attempting to do with safeguards that differentiate between investigative targets on the basis of state-specific information—but it would only address the underlying problem to the extent that it is possible to find hidden weapons at all.

And even if one could find assembled weapons that a sophisticated and determined cheater wished to hide—and with current technology and methods, the advantage still lies very much with the violator in this respect—it would be tremendously difficult to be confident in the absence of undeclared stocks of *unassembled* fissile material. Indeed, this would be the case in part because the nuclear weapons states themselves may be far from sure precisely how much nuclear material they have produced over the years. (The problem would be much worse, of course, if one could not entirely trust their honesty.) As a result, it would be difficult to be certain that the crucial final step in zero-based disarmament had in fact been achieved.

Proliferation of Fuel-Cycle Capabilities

There is also a disturbing nexus between the issue of future disarmament verification and the current global security challenge of stopping the proliferation of enrichment and reprocessing (ENR) technology. If ENR capabilities continue to spread, the corresponding proliferation of "latent" or "virtual" nuclear weapons programs would bode ill for the prospects of disarmament verification. The availability of fissile material has long been the primary choke point for nuclear weapons programs, and the easiest way to track the pace of potential programs—but for ENR-capable states, it would no longer be such a challenge. Future disarmament verifiers might have to find ways to detect aspects of nuclear weapons development not related to nuclear materials in order to guard against breakout from a disarmament regime.

Such non-nuclear work is likely to be harder to detect than fissile material activity, however, for it will lack the sort of telltale radiological signatures that inspectors currently look for in the context, for example, of IAEA safeguards inspections. Some non-nuclear work—such as the development of the specifically-shaped high explosive used in an implosion-type nuclear weapon—would surely represent a clear signal of weapons intentions. But other types of non-nuclear technology related to nuclear weaponry may be confusingly dual-use in nature. (The question of precisely what disarmament verifiers should be looking for in this regard could itself constitute sensitive information.) This is a conceptual nut that could prove very hard to crack.

As things stand today, a country operating a nuclear reactor, possessing nuclear material, or otherwise holding some such useful technology could always retain the option—however unlawful—of simply repudiating the verification system and moving forward quickly with the development of nuclear weapons. The possession of ENR technologies makes this rapid turnaround much easier. Stocks of low-enriched uranium usable in fueling light-water reactors can be enriched further to weapons-grade levels with relative ease. And plutonium can quickly and easily be chemically separated from spent reactor fuel—as North Korea did at Yongbyon—especially if one is not squeamish about radiation safety standards and is willing to expedite the process by using fuel not long removed from the reactor's core.

It is hard for a verification system to detect such activity if done clandestinely, but such actions can also be taken overtly by a violator who is not shy about being perceived as a dangerous scofflaw. These abuses cannot be precluded technologically; deterring or dealing with them takes

us out of the realm of technical verification and into the more political world of compliance enforcement. Suffice it to say that the less successful the world is in controlling the spread of ENR, the more difficult it will be to verify a future total disarmament regime. (Paradoxically, this difficulty has not stopped advocates of unrestricted ENR proliferation from propounding disarmament.)

Accounting for Nuclear Materials

Among the many challenges of disarmament verification will be devising new and improved methods to account for existing stocks of fissile material around the world—particularly, but by no means exclusively, in current weapons-possessing states. It would be difficult to feel confident that all nuclear weapons had been eliminated and that no country was preparing for breakout from a disarmament regime unless a detailed understanding had been developed about the production history and current fissile-material holdings of all countries with any production capability. Unfortunately, the methods now available may not be able to provide the necessary degree of certainty.

For example, Japanese officials admitted in 2003 that some 206 kilograms of weapons-usable plutonium had gone unaccounted for in the country's pilot plutonium-reprocessing plant over fifteen years. (This was in addition to 70 kilograms unaccounted for at a plutonium-based fuel fabrication plant.) The United Kingdom has also experienced such losses, with its Sellafield plant reporting 19 kilograms of unaccounted-for fissile material that same year.

This is not to single out Japan and the U.K., both of which have state-of-the-art nuclear industries. The United States, too, had an exceedingly difficult time accounting for plutonium at the now-dismantled Rocky Flats nuclear weapons pit manufacturing plant in Colorado. Even for sophisticated modern operators, it can be very hard to account for everything.

Moreover, such verification work would probably nowhere be harder than inside the nuclear weapons manufacturing industries of today's weapons possessors. The NPT's five acknowledged possessor states (the United States, Russia, France, China, and the United Kingdom) have been producing nuclear weapons for many years; their facilities were not designed to be subject to international safeguards, and their long years of operation under the less demanding safety and accountability standards of previous generations would vastly complicate efforts to ascertain production history and develop a detailed accounting of all materials. This complication is not

simply a problem for those five countries, however: the degree to which other weapons possessors could provide usefully detailed overall production data and account for all material ever produced is anybody's guess.

Unaccounted-for nuclear material presents a huge challenge for verification in a disarmament regime. It is hard to imagine that the world would feel terribly secure about the elimination of all nuclear weapons if fissile material sufficient for perhaps hundreds of weapons remained unaccounted for worldwide (and the estimates of unaccounted-for material grows with every year's operation of the global nuclear power industry). This uncertainty is another reason why the spread of ENR technology is so dangerous to the cause of disarmament. As more countries get into the business of creating fissile material, even very small margins of error in accounting for it will create large uncertainties about how many potential weapons-worth remain out there somewhere. Global nuclear disarmament would require far better—and to some extent even retroactive—materials accountability standards, technologies, and methodologies than exist today. It would also require the assent of governments around the world to costly and intrusive inspections, and a much more serious approach to the spread of ENR capabilities.

An Analogy That Should Give Pause

Several of the challenges raised by verification in a hypothetical global disarmament regime—particularly one that opted to emphasize nuclear-material controls instead of warhead-specific accountability—would be reminiscent of those that have arisen in debates over the proposed Fissile Material Cutoff Treaty. (The FMCT, which would ban the production of enriched uranium and plutonium for weapons purposes, was proposed by former president Bill Clinton at the U.N. in 1993 and has been in various stages of non-negotiation ever since.) From a disarmament perspective, comparisons to the FMCT ought to be worrisome, as the United States concluded in 2004 that no verification regime that could plausibly arise out of FMCT negotiations would be able to provide for the treaty's effective verification. The Obama administration reversed course on FMCT verification in early 2009, taking the position in international negotiations that the FMCT can and must be effectively verifiable. As yet, however, the Obama administration has offered no account of how FMCT verification would work, nor any rebuttal to Bush administration arguments that it would not. (This silence has not yet mattered much practically, because the Conference on Disarmament remains snarled in its usual procedural stalemate. If the

Conference deadlock were to end, or if FMCT talks were simply removed from the Conference forum, these questions would return to the fore.)

To be sure, a global nuclear disarmament regime would not face precisely the same tasks that would face FMCT enforcers. Presumably it would not face the same need to divine the purpose for which some quantity of discovered fissile material had been produced. (As mentioned, the FMCT would only ban the production of fissile material intended for weapons or other explosive purposes.) Nor, presumably, would a disarmament regime need to exclude from verification any material produced prior to a cutoff date, as would probably be necessary under the FMCT. Nonetheless, some disarmament challenges would be similar to those raised by the FMCT, including the question of how to manage verification with regard to material intended for "non-proscribed but sensitive" uses, the detection of fissile material production at clandestine facilities, and adequate monitoring of declared production. Some challenges, like accounting for fissile materials, would likely be much more difficult under total disarmament than under the FMCT.

In my previous job as U.S. Special Representative for Nuclear Nonproliferation, I heard a surprising number of diplomats at the Conference on Disarmament agree privately that the FMCT is indeed not really verifiable, arguing nonetheless that under the circumstances, simple agreement upon *some* international verification measures would be enough to declare success. (The Bush administration opposed this approach for fear that it would create a false sense of security: the treaty, it worried, would pretend to achieve effective verification while leaving dangerous scope for undetected violations.) Even if these diplomats' somewhat cynical approach were felt to be adequate for FMCT purposes, however—and the Obama administration seems to be inclining in this direction—it is a "solution" that would surely be harder to defend when applied to total nuclear disarmament. After all, a successfully negotiated FMCT would be born into a world in which nuclear weapons still exist in significant numbers, so producing more weapons-grade material would violate the treaty and could be destabilizing, but it would not necessarily radically change the global balance of power—at least not initially. But the stakes would be vastly higher in a world of *no* nuclear weapons because of the tiny margin of error that total nuclear disarmament would demand. It is hard to imagine countries laying their future security on the line in favor of a total disarmament regime that really only promised half-measures. Friends of total disarmament have much work to do in devising better answers to these problems.

Issues of Politics and International Dynamics

The verification of a total nuclear disarmament regime also poses several political challenges—including the problem of compliance enforcement, an endeavor that has not hitherto been the international community's strong suit. In this respect, U.S. Arms Control and Disarmament Agency Director Fred Iklé hit the nail on the head in a mere three words in the title of his 1961 *Foreign Affairs* article: "After Detection—What?"

Verification, after all, is only a means to an end: the correction of noncompliance and the deterrence of future violations. The point is not merely to document a violation for the sake of history, but rather to give the rest of the world a chance to respond to the breach. The IAEA already strives for "timely detection" of the diversion of nuclear material into nuclear weapons work, but as eminent strategists such as Iklé and Albert Wohlstetter understood years ago, timeliness is not a purely technical issue. It involves, as Wohlstetter put it a 1976 *Foreign Affairs* essay, "more than simply detecting a violation of an agreement. It means early detection of the approach by a government toward the making of a bomb in time for other governments to do something about it."

This latter aspect of timeliness—what one might call its political component, insofar as it entangles the effectiveness of verification with the responsiveness of real-world institutions—is logically inherent in any form of arms control, but it acquires particular salience in dealing with nuclear weapons because of the potential of such colossally powerful explosives to rapidly devastate a target nation. In the United States, the Acheson-Lilienthal Report of 1946, the first official study to call for comprehensive nuclear arms control, made clear that a workable verification system must provide "danger signals" that "flash early enough to leave time adequate to permit other nations—alone or in concert—to take appropriate action." (U.S. officials reaffirmed this view as recently as 2007.) In the nuclear arena, therefore, the warning required for effective verification has always encompassed a consideration of the time it would take for the international community to respond to a violation.

And even a timely warning is not the whole matter, for in order to forestall disaster, the recipients of the warning must then in fact mount an effective response. A verification system might be quite effective in providing timely warning, and yet the whole endeavor could come to naught if that opportunity is squandered. It is the policy challenge of compliance enforcement to ensure that warnings are followed by consequences sufficient to correct the problem—or at the very least to ensure that others

will not want to follow the violator's example. Without the meaningful prospect of effective compliance enforcement, a verification regime would be pointless.

Accordingly, even if effective means could be found to identify and eliminate all existing nuclear weapons, and to monitor relevant facilities and material stockpiles in such a way that cheating would be detected, actually reaching total disarmament would require giving today's possessor states confidence that any such detection would reliably result in effective compliance enforcement. This is but one among many reasons why the prospects for achieving full disarmament critically depend on whether the international community is able to meet the proliferation challenges presented by North Korea and Iran. If multilateral mechanisms cannot produce a reliable track record in addressing proliferation threats, which of today's nuclear states would be willing to give up its own nuclear weapons and entrust its future security to such feckless means?

By their very nature, these political challenges cannot be addressed simply by making available certain technologies or financial or human-capital resources, or by penning specific authorities into the text of a multilateral agreement. Meeting these challenges rests upon the collective political will of the international community—upon there being enough states sufficiently like-minded, diligent, trustworthy, and scrupulous to apply themselves consistently over time to ensuring that all parties are kept rigorously to the terms of an abolition regime.

Because this will be hard to guarantee in advance, advocates of full disarmament should not be surprised to find nuclear weapons states strategically "hedging" in ways that complicate the process of relinquishing their weapons. Such hedging would not necessarily preclude full disarmament, and indeed might well—from a political perspective—be essential to eliciting the participation of possessor states. But it would certainly make the process more complex and tricky. Disarmament advocates would do well to understand and assess the potential impact of various hedging options that might be available to weapons states, as some approaches may be greatly preferable to others from the perspective of global stability within a disarmament regime.

One such strategy, publicly discussed by U.S. officials in 2007 but in fact also having antecedents in the Acheson-Lilienthal Report and disarmament activist Jonathan Schell's 1984 book *The Abolition*, is the retention of a capability for "countervailing reconstitution"—that is, the ability to rebuild a nuclear arsenal in short order should some other party be caught attempting breakout. This, it was suggested, might make disarmament

more achievable by giving key nuclear weapons states a bit more confidence that they could eliminate existing weapons without imperiling their future security. It might also help to deter violations, insofar as a would-be violator would know that his violation would quickly be answered, and presumably his strategic gains as a nuclear weapons monopolist sharply undercut, by the reconstitution of countervailing deterrent arsenals. On the other hand, Nobel laureate and deterrence theorist Thomas Schelling has warned that in a world of "zero," any ability to reconstitute nuclear weapons—and in no imaginable world, he rightly points out, would all countries *entirely* lack the ability to build them—could lead to dangerous weapons-building "races" in times of war or crisis, and even create incentives for the preemptive use of nuclear weapons. At present, the Carnegie Endowment and the Hudson Institute are collaborating on a study of the various theoretical and programmatic issues raised by nuclear weapons reconstitution; it was also the subject of a September 2009 conference at Stanford's Hoover Institution. Now that even the Obama administration has found it necessary to increase spending on modernizing the U.S. nuclear weapons production infrastructure—and other weapons possessors continue their own longstanding modernization work—it is becoming clear that disarmament advocates must grapple with the challenge of hedging strategies such as countervailing reconstitution.

As difficult as it would be in a nuclear weapons-free world to deter regime breakout and manage other issues of post-nuclear stability, the dynamics of the *transitional* world through which we will necessarily have to pass en route to zero also present a challenge of their own. Unless one is to suppose that other nuclear weapons states would agree to abolish their weapons well before Washington and Moscow are prepared to take that step—which is theoretically possible, albeit quite unlikely—there will necessarily be a period between today's world and "nuclear zero" in which the United States and Russia have reduced their numbers to the point that all (or at least most) current weapons possessors have become, more or less, numerical peers.

If abolition is really to be the goal, at least passing through such a phase of near-parity seems inevitable. But the advocates of full disarmament seem entirely unprepared for such an environment of multiplayer nuclear deterrence between "near-peers." The world has some (infamously risky and worrisome) experience with nuclear deterrence on an essentially bipolar basis between the superpowers during the Cold War. It also has some (likewise worrisome) experience with multiplayer Great Power-balancing, for some centuries in Europe before 1939, as well as in various

periods of interdynastic Chinese history over the millennia. But mankind has no experience with multiplayer *nuclear-armed* deterrence, and it is currently difficult to say how this would prove different from such precedents. Since failures of deterrence in the nuclear context can so quickly result in catastrophe, this balancing is thus another crucial subject for future research, building on what piecemeal work has already been done.

Fundamental Ethical Questions

Considerations of the effectiveness of compliance enforcement, and of potential strategic hedging by weapons possessors contemplating disarmament, also point us toward the moral issues raised by the disarmament project. Ironically, given the sententiousness of so much disarmament advocacy, this is intellectual terrain on which the disarmament community is unused to engaging; all too many proponents seem to assume that their objective is of such unquestionable moral merit that it is unnecessary to do more than simply assert its virtues. If national leaderships are to be sold on the idea of full disarmament, however, it will have to be better defended in terms both political and ethical—and in a conceptual language consistent with, rather than self-consciously opposed to, the discourse of national security calculation in which such decision-making occurs.

In this regard, it is worth remembering that nuclear disarmament *per se* is not—or at least should not be—the ultimate objective of disarmament advocacy. The point is not to get rid of nuclear weapons just to get rid of nuclear weapons; it is, rather, to abolish them on account of the risks and dangers they present to global stability and even to civilization itself. This being the case, a genuinely ethical consideration of the subject also requires consideration of the risks and dangers presented by various alternative future security environments, both with and without nuclear weapons elimination. We cannot escape the necessity of weighing possible outcomes against each other.

It is very hard to defend the premise that there are *no* imaginable conditions under which a world free of nuclear weapons would be more dangerous and unstable than today's world. One foreign diplomat of my acquaintance likes to joke privately that the disarmament movement needs to be careful lest it "make the world safe again for large-scale conventional war." He has a point. We should remember that nuclear weapons helped end the bloodiest and most globe-convulsing conflict in human history, and it is not obvious that posterity would thank us for trading the world of 2010 for a world more reminiscent of 1914 or 1939.

This does not mean that we necessarily face such a stark choice between nuclear weapons possession and global war. Hopefully—and most likely—we do not. Nevertheless, rightly or wrongly, a number of governments clearly think that nuclear weapons contribute to their national security in important ways. Responsible present-day decision-makers need to have more confidence about their nations' security in a post-nuclear world than can be provided merely by the tautological observation that it would be a world free of nuclear weapons.

An Agenda for Disarmament Research

None of these points should be taken as an attack on the notion of nuclear disarmament. To the contrary, they are intended to provide an agenda that may help guide disarmament research and constructive thinking about international security dilemmas. Unless the tough questions can be persuasively addressed, it is very hard to imagine that real-world decision-makers will take total disarmament very seriously in the years ahead.

Advocates of total nuclear disarmament must begin to answer questions in seven essential areas. First, they must deal with the procedural and technical difficulties inherent in verifying a "nuclear zero" world: What is the minimum amount of information that would be required from one nuclear weapons state in order to give other nuclear weapons states confidence that it had dismantled real warheads, or at least that fissile material had really been removed from its weapons stocks? What is the minimum amount of information that would be required to elicit such confidence among non-weapons states? What is the maximum amount that weapons states would be able to reveal, consistent with their own national security requirements and their obligation not to contribute to nuclear weapons proliferation by passing sensitive weapons-related data to others? How could such approaches be made resistant to spoofing and deception? What levels of verification certainty are necessary in order to support disarmament; what levels are possible; and what happens if these do not coincide?

Second, they must explain what technologies and methodologies would be required not to simply hunt for undeclared nuclear activities (as the IAEA Additional Protocol is designed to do), but actually to find hidden nuclear weapons. Under conditions of maximal legal authority and permissive access, how low can verification mechanisms push uncertainty levels as a matter of technical capability? As a practical matter, what resource burdens would such verification mechanisms entail?

As a political matter, how—if at all—might states be persuaded to accept such costs and intrusiveness?

Third, they must explain how their long-term aim of total disarmament relates to the challenge of proliferating fuel-cycle capabilities. How specifically would the continued spread of ENR capabilities affect the technical challenges, resource burdens, and political costs of disarmament verification? What difference would it make if ENR proliferation continues principally as a matter of national government monopoly versus through increased reliance upon multilateral fuel-production capabilities? To what extent can our understanding of the interconnectedness of the ENR and disarmament issues help shape how the international community addresses (or fails to address) ENR challenges?

Fourth, disarmament proponents must address the problem of unaccounted-for nuclear materials. What can be done to improve the error margins inherent in present-day materials accountability? Is there any technical way to address the challenges of materials accountability and past production totals in the sometimes greatly contaminated environments of decades-old weapons-state production infrastructures? If so, what resource and political burdens would have to be borne in order to acquire the necessary data set?

Fifth, advocates of total nuclear disarmament must speak to the core political questions their position raises. Especially because timely warning depends in part upon the rapidity with which reactions to treaty violations can be mounted, what can be done—as a matter of law, institutional design, and (above all) politics—to maximize the chances that the international community will respond quickly and effectively to breaches of an abolition agreement, and is this enough? If such responses cannot be entirely assured, what hedging strategies might be adopted by states contemplating nuclear relinquishment, or simply worried that the system will be unable to prevent breakout? How might various alternative hedging strategies affect the stability of an abolition regime?

Sixth, they must also speak to the core questions of international dynamics. How would deterrence work in the multiparty world of numerical near-peers toward which we will most likely be moving as U.S. and Russian nuclear arms levels continue to fall? However long it lasts—whether it is merely a way station along the road to zero or in fact a new nuclear era in its own right—how would the dynamics of such a multiplayer security environment differ from the two-party nuclear deterrence to which strategists became accustomed during the Cold War? How would they differ from the non-nuclear balance-of-power systems seen in earlier eras?

And finally, they must address several fundamental ethical concerns. What strategies maximize the chances that a post-nuclear-weapons world will indeed be more able to manage international conflict than today's world? How can governments and international institutions best provide non-nuclear substitutes for whatever constructive role nuclear weapons may be felt to have played in the global security environment? In recent years, for instance, the United States has considered developing missile defenses for use against problem regimes, as well as new capabilities for deploying non-nuclear weapons anywhere around the globe within mere minutes. The Obama administration may also seek to bolster non-nuclear security relationships with allies presently dependent upon American nuclear weapons, so as to be able to provide strategic deterrence in some other fashion. Such post-nuclear security issues must be addressed if policymakers are ever to contemplate "zero" as a realistic possibility.

Taking such questions as a program for further research ought to unite both disarmament's most ardent advocates and its gloomiest critics. If the skeptics are right that complete nuclear disarmament is unachievable, or is perhaps even undesirable anyway, they should welcome a serious research agenda that would clearly expose any existing flaws. Disarmament advocates should also welcome this research, for it is hard to imagine reaching zero unless nuclear weapons states can be persuaded that satisfactory answers exist. Answering these questions will require much more rigorous thinking, intellectual and moral honesty, and sustained attention than has so far been in evidence. Accordingly, both camps—and everyone in between—ought to be able to agree that the time has come to face squarely the challenging questions presented by the notion of total nuclear disarmament. Failing to address these issues would be gravely irresponsible.