



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Asymmetry and Risk

B. T. Goodwin

July 6, 2011

Asymmetry and Risk
Washington DC, DC, United States
April 26, 2011 through April 26, 2011

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Asymmetry and Risk

Presented to the US State Dept Workshop on Advancing Technology To Support
Verification and Monitoring Challenges on the Road To Zero
April 26, 2011

Bruce T. Goodwin
Principal Associate Director Weapons and Complex Integration
Lawrence Livermore National Laboratory

The Administration has indicated its desire to advance nuclear weapons arms control treaties in the coming years. The nature of potential future agreements is likely to be very different from past experience. For example, in addition to seeking dramatically lower weapon numbers, counting methods may move from platform accounting and warhead attribution to a more detailed warhead accounting dealing with both operationally deployed and stockpile weapons. Regardless of one's opinion on the desirability of particular treaties or their attributes, technology can inform the uncertainties and thereby quantify risk and so enhance (or, depending upon outcome, diminish) confidence in provisions. For the US, these technological considerations will be informed and supported by four issues: The state of the US deterrent, extant bilateral asymmetries, mutual understanding and trust (or lack thereof), and verification capabilities. This paper will address these four issues and propose options to improve prospects in each technical area, with the goal of quantifying the risk in future agreements.

I. The potential fragility of the US deterrent

The current stockpile was built assuming that it would be large, diverse and changed out on an approximate ten-year cycle. It is now very old (average age greater than 25 years), with declining diversity, is static (i.e. no new systems since the late 1980's), designed for high yield to weight with relatively low margins to failure (who, back then, would know or care if a few out of many thousands failed to work), and shrinking in overall numbers. Yet it continues to be managed as if the stockpile of old existed today under the conditions of old. But conditions have changed. Edisonian* trial and error methods and rapid turnover production cycles become harder and harder to sustain as the stockpile ages and shrinks under today's conditions.

The intrinsic nature of a future, small deterrent leads to a few important considerations. A small stockpile must have few backups in order to be small. Therefore small changes will have large risk impacts. Current Edisonian, statistically based surveillance techniques (i.e. blindly pull a sample of weapons, cut them open and check to see if anything has failed) are unsustainable. This is because, as the stockpile shrinks, an increasingly larger proportion of deployed weapons must be pulled from service to maintain statistically based confidence. Finally, a future weapons infrastructure must be matched in size and capacity to the size of the deterrent. If the stockpile is small, so must be the infrastructure.

If not, the infrastructure will go idle. Because humans are not good at sustaining skills in inactivity, an infrastructure not sized to the stockpile it serves will become unaffordable, be significantly inactive and so, eventually, moribund.

For a deterrent to deter, both sides must be confident that it will work, not only under steady conditions, but also in the face of sudden change. The US must therefore be able to deal with the *"what ifs"*. Three leading examples are technical defects, geopolitical change, and scientific and/or technical disruptive technology (the "bolt out of the blue").

A technical defect (i.e. an aging or design problem found in surveillance) demands sufficient scientific/technical and infrastructure (i.e. manufacturing) responsiveness in order to understand the fix required and to deploy that fix quickly enough to sustain the deterrent posture. Such stockpile issues have, in fact, arisen relatively frequently.

Geopolitical change is a significant change in the threat posture of an existing or emergent adversary. Geopolitical change demands that the infrastructure (i.e. manufacturing base) be sufficiently responsive to counter the change in adversarial posture (e.g. treaty breakout) quickly enough to sustain the credibility of the deterrent.

Finally, the scientific and/or technical *"bolt out of the blue"* (for sudden) or *"frog in the pot"* (for gradual) is the emergence of a new, disruptive technology that diminishes or eliminates the credibility of the deterrent. By its very nature, such change is difficult to foresee and adapt to. Some long anticipated examples might be pure fusion weapons, long range, speed of light directed energy weapons or technology that "turns the oceans transparent". A pure fusion bomb would be disruptive because it would use no special nuclear materials and so have none of the production or tracking signatures of fission weapons (i.e. fissile material radiations) nor would it have the familiar fallout signature of fission weapons. Transparent oceans technology would make SLBM submarines no more defensible than the military dirigibles of the 1930's (i.e. large, slow and visible). The best defense against disruptive technology is a vital scientific/technical cadre and infrastructure to anticipate and respond to emergent threats.

Confidence then is derived from our ability to respond effectively to these and other possible futures. The 2010 Nuclear Posture Review (NPR) establishes requirements that seek to address these issues. The NPR defines a go forward stockpile that is effective, intrinsically safe and secure, maintainable and adaptable. Effective here means a stockpile that will work if required. This is a stockpile that has high margin to failure, so that aging or other flaws are unlikely to cause failure to function. An intrinsically safe and secure stockpile is one wherein safety and security (the shorthand "surety" is often used to capture both attributes) are engineered into the nuclear explosive rather than achieved through extrinsic, administrative procedures. The foundation of "surety" is insensitive high explosive (IHE). High explosive is the energy source that implodes the fissile material in a nuclear explosive leading to fission criticality and subsequent nuclear function. Unlike conventional chemical explosives, IHE is essentially impossible to detonate in accidents or by unauthorized human action. IHE is not just a little safer, it is really, fundamentally safe stuff. This

makes for both as smaller, more responsive infrastructure because more than one warhead can be safely worked on at a time. For related reasons, such an IHE stockpile is much easier to safely defend and so requires a much simpler, smaller, less expensive and more responsive protective posture.

Maintainability is key to a small, responsive infrastructure. Simple assembly operations are central to rapid maintenance and modern, simpler, less costly surveillance.

An adaptable stockpile is one wherein one nuclear explosive package can stand in for another on a different delivery platform. We do not have such capability today. An adaptive stockpile means that the deterrent can have many fewer non-deployed "warheads in waiting" to deal with failures in a given system. This means that the overall hedge can be reduced to as small as half the size of a non-adaptive hedge.

The problems with the existing stockpile present both issues and opportunities. Because the current stockpile is shrinking but was designed to be large, it is becoming hard to manage, especially in the face of aging. The opportunity exists today to use the life extension programs defined by the NPR to move to a deterrent that can be small yet manageable and responsive. As you will see further in this paper, we also can have the opportunity to include passive technology that could improve confidence in warhead accounting.

II. Asymmetry

The Russian and US deterrents are at equivalent levels of technical sophistication. Despite this, they are fundamentally different in design, size, manufacturing capacity, stewardship technology, delivery platforms, etc. This intrinsic asymmetry leads to very different drivers for deterrence confidence and creates very different risk perceptions for particular treaty features. For example, nuclear explosive design can make for divergent warhead verification technologies. The use of different fissile materials can make verification technology easier [e.g. plutonium) or more challenging (e.g. enriched uranium). Design and manufacturing technology may make for either very long or short stockpile life and so demand a corresponding smaller or larger manufacturing throughput. Platform design combined with nuclear explosive package design can make the strategic vs. tactical distinction artificial in terms of performance even if distinctions based upon range, mission and arms control counting remain. (I note here Russian public comments to the artificiality of this distinction). Finally, stockpile size and infrastructure capacity are inversely related. Small throughput demands a large non-deployed hedge to deal with the unexpected, whereas large throughput enables a small (or even no) non-deployed hedge.

These asymmetries make for divergent motivations for acceptable negotiating positions. As an example, in the area of latency, or the time required to up arm from a given deployment state. Russia and the US are in very different initial states and are asymmetrically latent. Latency depends upon the number of deployed warheads, reserve warheads, platforms and mounting points, available components, material reserves and individual national capabilities and capacities. The US and Russia are in different states with respect to virtually all of these factors. As an

example, the US has very few tactical weapons and Russia has many by comparison. If warheads are designed to be adaptable to different tactical or strategic platforms, the distinction is moot and Russia would have a potent latency advantage that they would be reluctant to yield. Russia appears to have a large manufacturing capacity, while the US has, by comparison, a much smaller capacity. This would drive the US to maintain a large reserve of stored, disassembled components and assembled, non-deployed warheads to compensate for the smaller infrastructure throughput. On the other hand, if Russian warheads are designed for a significantly shorter stockpile lifetime, they must have a large throughput to sustain a given number of warheads. These kinds of asymmetries can lead to errors in perception on both sides. For example, American engineered safety requirements in production facilities make these facilities very large for a relatively small throughput. Other nations use administrative controls to achieve facility safety and this makes for a much smaller footprint for a given throughput. One can foresee an asymmetry that could lead to distrust on the part of Russia because US claims of small throughput would ring hollow to them in the face of an apparently contradictory large facility footprint.

A straight-forward way to address problems of asymmetry is verification. If you know, there is less reason to worry. This raises the question: what do you need to know and how well you need to know it? Asymmetry and perception thereof are fundamental to this problem. Understanding the nature of an alien enterprise and how to account for mutual differences is key to solving this problem. If you know well enough, terms can be understood from a risk/confidence perspective and agreed upon (or not), by either side.

Ideally, the complete warhead cycle must be comprehensively thought about and understood. In this endeavor, asymmetry means that no mirror imaging can be allowed. In verification, it is not enough to just go after warheads - one need is to understand plants, reserves, transport, storage, components, material forms and quantities, platforms, design (at some level). It would be best to have a detailed understanding of the counterpart designs.

This leads to a fundamental need to understand the very different weapons enterprise of your counterpart. Capacities are driven by technical requirements and capabilities derived from the enterprise itself. Thus, a Russian weapons enterprise model is key to understanding the impacts of proposals for agreements.

The development of such a model has been surprisingly difficult to do even for the US's own enterprise, but one exists. Trust between US enterprise entities had to be developed to achieve the needed visibility into the system. So, the initial hard part is done. The new challenge is to develop an analogous model for the Russian enterprise. There is a major difference however. For the US model, we have the highest achievable certainty in inputs, inventories, facility capabilities and capacities, and output. We can test the model for validity against validated historical data. A Russian enterprise model would be like a US model *turned upside down*. Certainty now becomes uncertainty for inputs, inventories, facility capabilities and capacities, and output. Fortunately, the stockpile stewardship program has developed the ability to quantitatively evaluate uncertainty in modeling. The Quantification of Margins and Uncertainty methodology [or QMU] initially developed at Lawrence Livermore National Laboratory and extensively elaborated upon by both US nuclear laboratories over the last ten years is able to propagate uncertain inputs through

incomplete system models and yield reliable estimates of output values with quantified errors. Suffice it to say that the development of a good enough model will demand lots of help from the intelligence community both for inputs, models development and, especially for validation of outputs and associated uncertainties, and so risk evaluation.

III. Trust

Establishing trust is essential to any productive human activity. To increase Russian/US trust, establishing a program of nuclear weapons technical cooperation can have significant benefits. Mutual, technical visibility into each other's deterrent would improve understanding of risk in the technical aspects of any proposed future agreements.

The history of US/Russian technical cooperation to date has been mixed at best. Nuclear lab-to-lab cooperation from the '90's to the present has not been on equal terms nor mutually satisfying. Russian engagement began under economic duress, and they have resented their perceived (and real) junior status in these interactions. At this point in time, now that Russia is financially well off, cooperation has essentially ended. The Russian labs are no longer the worn looking places of fifteen years ago. The presence of modern, western automobiles and modern, private gasoline stations inside "the fence" testifies to this. Looking to the future, this must be a process between equals if we are to build a mutually useful relationship with some level of trust.

This begs the question "how?". What other cooperative experience can we learn from and so build upon? We have only two examples to use: the UK Mutual Defense Agreement and the cooperative agreement with France. The UK relationship is a less useful model. The UK/US special relationship has deep roots in the Manhattan Project in World War II. It benefits from the high degree of linguistic, cultural, and political alignment. Perhaps we may learn more from the less "aligned" French experience. While France is not an adversary, she is ardently independent and self-reliant. France defines her own path and cooperates for her own self-interest. The history of cooperation is long and slow. Mutual trust had to be achieved through hard work so that technical cooperation could advance. It was as Russian Director G. Rykovanov of the All-Russian Scientific Research Institute of Theoretical Physics (VNIITF) told me a decade ago, "We must get to know you in order to work with you."

With France, starting about four decades ago, initial cooperation was on safety and so unclassified for a long time. After decades, as the relationship and national interests developed, cooperation moved to limited classified subjects and after more decades, further expanded. Note that with a non-adversary, it has taken more than forty years to come to significant technical cooperation.

However, the "road to zero" in nuclear weapons is likely to be very long, winding and perhaps unending. So the sooner we start... But, why start? First, as already stated, the details of Russian technology very much inform US risks into the future [and vice versa]. Second, we know that they are our technical equals. While we may learn from each other, we (and they) are unlikely to significantly effect threats as a result of what is revealed. This is because our (and their) deterrent meets national needs. Third, from what we know to date, we could begin to discuss issues of nuclear weapons safety without too much risk and

work our way forward in a careful, measured way to increasing levels of technical sharing.

In the end, it is really important to know what we are dealing with so as to understand risk and how to verify with sufficient confidence. To repeat, it will take a very long time to advance mutual trust, so the sooner we start...

IV. Technical Verification

The goal here is to have full visibility into the weapon systems, platforms, components, materials and knowledge of whereabouts. This is not likely to happen soon. I do suspect that we (and they) have a pretty good general idea of what we all have. The problem with nuclear weapons is, what you don't know *can* hurt you. Therefore, increased visibility is a very good thing. With careful red teaming of risks, we should be able to develop technologies to verify a weapon and then tag it.

New LEP stewardship techniques for surveillance could be applied to tracking once a weapon or component is verified to be such (in or out of a storage can). Similarly, new techniques for measuring materials can give increased confidence that you know what's in a can. All of this must account for Russian/US classification asymmetries. For example, hypothetically, one side may protect details of high explosives technology that the other does not. The other may protect aspects of plutonium technology that the first does not. This is another reason why developing an "SRD" level (i.e. at the US secret restricted data level of classification) relationship with Russia could move verification forward.

Verifying and then tagging is the beginning. Tracking (i.e., continuity of knowledge) is the goal here. The key to this goal is to understand the sources and sinks of materials, components, warheads and platforms in the enterprise and to track them throughout. Hence, understanding how the enterprise works and moves makes a good enough enterprise model central to this process. We should therefore:

- 1 - Increase technical cooperation to gain sufficient technical understanding of potential risk from Russia;
- 2 - Deploy technology to verify (detect) and tag (track) materials, components and warheads;
- 3 - Develop a Russian enterprise model to evaluate their system and the evolution of risk to the US.

V. Conclusions

Whether one thinks that "zero" is or is not a good idea, a journey on "the road toward zero" will have many detours, hazards, opportunities and an unknown destination. One thing is clear. As stockpile numbers go down, each weapon takes on increased importance. In a "zero" nuclear weapons world, a single nuclear weapon changes everything. Therefore, any path we take must assure the stability of deterrence. A small number of technical objectives can help sustain stability as a stockpile shrinks. The NPR provides a solid path to a future smaller deterrent. The NPR calls for needed changes in technologies to enable a small deterrent (i.e. LEPs). Technically, we can achieve reduced risk levels if we think about what we will need and work to get prepared. Future negotiations must take account of asymmetries. A foundational level of trust could increase visibility enabling better risk quantification. Quantified risk increases understanding and will increase (or decrease as appropriate) confidence that can enable (or defer) future negotiation.

***Edisonian** refers to Edison's methods of invention, e.g. the "cut and try" method whereby he invented the light bulb by trying thousands of materials in series until one, carbonized cellulose, had sufficient lifetime to make a useful light bulb. This can be extended to all fields of technical endeavor. Surveillance can be considered Edisonian if the stockpile is randomly sampled to find unanticipated faults and the faults fixed when found. This is as opposed to an approach based upon understanding the basic mechanisms of operation and predictively designing from first principals, thereby reducing development time. Similar methods for surveillance endeavor to understand sources of aging and predictively sample for signs of anticipated degradation, thereby finding faults before they degrade performance.

Prepared by LLNL under Contract DE-AC52-07NA27344.