Revisiting the Uninvention Hypothesis: A Transactional View of Tacit Knowledge in Nuclear Weapons Design

Presentation to the Society for Social Studies of Science Annual Meeting, Montreal, Canada, October 10-13, 2007

Benjamin Sims
Sociologist
Statistical Sciences Group
Los Alamos National Laboratory

Introduction

This talk is about tacit knowledge and nuclear weapons. Some of you may be aware that this issue has been addressed in the past in a particularly excellent article by Donald MacKenzie and Graham Spinardi, *Tacit Knowledge, Weapons Design, and the Uninvention of Nuclear Weapons*. This article, published in 1995, has proven to be foundational both to the study of tacit knowledge in technological contexts, and to the cultural analysis of nuclear weapons production. My first goal in this talk is to update MacKenzie and Spinardi's analysis of nuclear weapons work, based on my own immersion in the nuclear weapons community, and in light of events since the 1990s.

Based on this updated analysis, I then suggest some new concepts that better capture the way tacit knowledge works in complex technological projects. In particular, I focus on how tacit knowledge is shared between distinct communities of practice, each of which has their own body of tacit knowledge that is inaccessible to outsiders. I argue that coordination between such communities is possible via a type of tacit knowledge I call *transactional knowledge*, which concerns the technical processes and social norms by which divergent bodies of tacit knowledge are coordinated to produce scientific results or working technological artifacts.

I then reassess the prospects for uninvention of nuclear weapons in light of this analysis. Specifically, based on an analysis of developments in the weapons community since the end of the Cold War (when MacKenzie and Spinardi wrote their article) I argue that straightforward uninvention is not a likely scenario for the future of nuclear weapons.

Tacit Knowledge and the Uninvention of Nuclear Weapons

Tacit knowledge has been a concern of science studies since the emergence of the field. The tacit knowledge focus was especially associated with the Sociology of Scientific Knowledge school of thought, and in particular the work of H.M. Collins. Collins' foundational article in this area, *The TEA Set: Tacit Knowledge and Scientific Networks*, was published in 1974. Since then, the idea of tacit knowledge has been taken up in numerous laboratory studies, and has been expanded into the technology and engineering arena as well.

The SSK account of tacit knowledge, as originally articulated by Collins, attributes several key characteristics to tacit knowledge. First, tacit knowledge is inextricably tied to a particular set of local practices. That is, it is embodied in groups of individuals who, by virtue of their participation in a common set of practices, come to share a particular "form of life," and therefore a common body of knowledge. The local nature of tacit knowledge makes its transmission across social and physical distances problematic. While some aspects of scientific knowledge can be captured explicitly in rules, algorithms, or other formalisms, tacit knowledge can only be transmitted by participation in shared practices. For example, Collins found that research groups that were able to build a TEA laser had almost always visited with another research group that had already been successful.

As a matter of methodology, the presence of tacit knowledge in a given social setting has commonly been inferred by one of two means in the science studies literature. First, through direct observation of technical experts and analysis of their own accounts of their work. If they are observed to employ a large amount of judgment in their work, and to be unable to enumerate formal rules underlying this judgment, this is an indication that they rely on tacit knowledge. Second, the presence of tacit knowledge has been deduced by studying patterns of replication in technical work – for example, replication of experiments or technological devices. If this cannot be accomplished through access to formalized knowledge alone, it is presumed to require some degree of tacit knowledge.

MacKenzie and Spinardi's article closely follows the outline of tacit knowledge I just gave. Methodologically, they first deduce the importance of tacit knowledge to nuclear weapons design and production by examining the replication of working nuclear weapons around the world. In every case, they find that the requisite knowledge was either transmitted through personal interaction or had to be painstakingly reinvented locally. Diagrams and written procedures were not, in themselves, sufficient to ensure transmission of weapons knowledge. They then further establish the relevance of tacit knowledge to nuclear weapons work through analysis of interviews with U.S. nuclear weapons designers. These designers report that their expertise involves a large element of judgment tied specifically tied to the context of weapons design and testing, and in particular to the use of simulation codes within this context. They worry that, if design and testing were to end, their knowledge of how to correctly use codes to predict weapon performance might not be transmissible to a new generation of designers. Worse, the new generation might continue to use codes without having the necessary judgment to understand their limitations, and might therefore lose touch with reality and become overconfident in predicting weapon performance.

From this analysis, MacKenzie and Spinardi conclude that:

> If there were a sufficiently long hiatus in their design and production (say a couple of generations), then that tacit knowledge might indeed vanish. Nuclear weapons could still be recreated, but not simply by copying from whatever artifacts, diagrams, and explicit instructions remained. In a sense, they would have to be *reinvented.* (47 AJS)

More specifically, they ask:

> As designers themselves age, leave, and die, the number who have first-hand experience of development through the point of full nuclear testing will steadily diminish; yet they will have to decide whether the inevitable changes in the arsenal matter. In such a situation, will explicit knowledge be enough? Will tacit knowledge and judgment survive adequately? For how long? (AJS, 92)

MacKenzie and Spinardi's account is methodologically thorough and very convincing in its demonstration of the relevance of tacit knowledge to weapons design and production. In particular, it accurately captures a set of concerns about knowledge preservation that weapon designers were beginning to articulate in 1990, which went on to become central in the debate over Stockpile Stewardship, the program that emerged to sustain the weapons laboratories after design and testing were finally halted in the early 1990s. However, I find some elements of their analysis to be inaccurate. In particular, they rely on a biased and incomplete view of the structure of the weapons community provided to them by the designers they interviewed. As a result, they attribute too much of the tacit knowledge needed to produce nuclear weapons to designers, while minimizing the complex division of labor between the numerous technical specialties involved in weapons design since the beginning of the Cold War. They don't entirely ignore the division of labor: there is discussion in their paper about the importance of skilled machinists, for example. What they ignore is the array of technical specialties that are integral to the modern weapons development process. These technical communities are mode up of many PhD-level scientists and engineers as well as skilled technicians, and are integral to the design and development process – they are not just passive followers of design directives. Going by common definitions of "design," much of the design of nuclear weapons is not done by designated "designers," although they do have a central role in the design process.

This observation has important implications for the "uninvention" of nuclear weapons. Whether the designer community, specifically, is able to maintain its tacit knowledge is only part of the story here. Other technical fields within the nuclear weapons complex must be considered as well. In addition, the ability of diverse technical fields to work together and share tacit knowledge effectively becomes a major concern.

Since MacKenzie and Spinardi's article appeared, the nuclear weapons community has undergone major changes in the way it functions, even though the weapons laboratories have not yet diminished in size or budgets. Since the end of the Cold War, when MacKenzie and Spinardi did their interviews, the scenario of a possible hiatus in design and testing has come true, and we are now at least half a generation into it. In its place, a Stockpile Stewardship program is underway that emphasizes development of advanced simulation codes, as well as the constant monitoring of the stockpile for unwanted changes, and minimal rebuilding of weapons to address any such changes. More recently, the concept of a Reliable Replacement Warhead, to be designed, produced, and fielded without full nuclear testing, has emerged as a possible path forward in maintaining a nuclear stockpile indefinitely. Analysis of these changes suggests that MacKenzie and Spinardi's scenario of tacit knowledge just vanishing in the absence of design and testing is probably an oversimplification.

Part of the reason why MacKenzie and Spinardi missed some complexities of nuclear weapons work, I think, lies in their reliance on tacit knowledge concepts more appropriate for studying small-scale science and the interaction of "core sets" of practitioners within a scientific field. Specifically, in their analysis of the knowledge of designers, they too easily fall into the assumption that the essential elements of weapons knowledge are embodied mainly within one unitary community with a common relationship to technology and practice. Based on this assumption, it becomes reasonable to argue that knowledge may be irrevocably lost if that single community becomes fragmented or its existing relationship to technology and practice is disrupted. In complex science and engineering projects, however, tacit knowledge is more distributed than this account suggests. The recent history of nuclear weapons work also suggests that tacit knowledge may be more malleable than MacKenzie and Spinardi propose, transforming rather than vanishing with changes in practice.

There is, in fact, a large body of research in science and technology studies that addresses how knowledge is transmitted and coordinated between scientific fields and across complex organizations. Peter Galison's work on trading zones in particle physics, Star and Griesemers' concept of boundary objects, and work by Joan Fujimura, Michael Lynch, and others on the circulation of standardized methods in the biosciences all address these issues. However, I am not aware than anyone tried to develop new ideas about tacit knowledge in light of these models of complexity and coordination in science and technology. H.M. Collins and Robert Evans have recently extended Collins' earlier work on tacit knowledge to encompass a detailed categorization of types of expertise, which includes a category of interactional expertise, in which an actor is able to successfully interact in a technical discipline without being a practicing member. This is an important element of what I call transactional knowledge, but it is only part of the story in the organizational settings I examine here.

Occupational Structure of the Weapons Community

The occupation of weapons designer, and the central role of the designer in the design and production process, was put in place during the Cold War. Each test device and each production weapon system had a lead designer, who was in charge of designing the weapon from a physics standpoint – determining the basic layout of materials and conducting the calculations to ensure that this layout would produce the desired yield. As the weapon moved into engineering and production phases of development, the designer still played a key role in determining whether engineering decisions and production methods were acceptable from the point of view of weapon physics and performance. As existing weapons age in the stockpile, there is also a designer in charge of the ongoing assessment and certification of each weapons system. Although designers are primarily concerned with the physics of nuclear weapons, some of them are actually engineers by training.

Designers are not a uniform occupational group. There is a distinction, for example, between designers who work mainly with the primary, or fission, component of the weapon, and those who work with the secondary, thermonuclear component. These designers work in different groups, and each design has a primary and a secondary designer. Another example of differentiation within the design community is the baseline code developer, who is in charge of integrating and stabilizing the suite of simulation codes used to assess and certify a particular weapon. Despite the title, baseline code developers are considered designers, since they do not actually make changes to the codes themselves.

Making changes to codes is the province of another group of specialists, called simply code developers. Code development was originally done by designers, but as computing capabilities grew, and codes grew more complex, code developers emerged as a distinct group. At the end of the Cold War, with the emergence of simulation as a key focus of the Stockpile Stewardship Program, code development became an organizationally distinct function, and more code developers were hired, many of whom were trained as computer scientists rather than engineers or physicists, solidifying the distinction between code developers and designers.

Code developers in turn rely on model developers, theoretical physicists who develop general models of the physics processes that occur during nuclear weapon detonation. The code developer is in charge of translating these theoretical models into working computer code. Model developers themselves often work with experimental physicists to refine and validate their models.

Engineers within the weapons laboratories (formally designated design agency engineers) are responsible for the detailed physical layout of the weapon and for documenting and assessing changes in weapon components over time. In the design process, weapon system engineers are responsible for translating the designers' physics work into a fully specified design that will integrate with other components within the delivery vehicle and can reasonably be produced by the production plants. System engineers are also attached to each system in the stockpile to monitor it as it ages. With the increased age of the stockpile since the end of the Cold War, a group of engineers specifically dedicated to tracking changes to the stockpile have emerged, called surveillance engineers. The engineering groups also have their own code developers that develop simulation tools for engineering purposes.

During the Cold War, diagnosticians were another important technical specialty. These were mainly physicists who were dedicated to developing the instruments used to measure the output of nuclear tests. Diagnosticians were also served by their own code development community.

As this breakdown of occupations suggests, the technical knowledge necessary to design and test nuclear weapons is not embodied solely in the role of the designer. Designers themselves are not a homogeneous social group, and do not possess a homogeneous body of tacit knowledge. Furthermore, the tacit knowledge they do possess, while of central importance to weapons design, exists in relation to a number of other technical fields. For this reason, I doubt that the tacit knowledge of weapons designers, in itself, would be sufficient to produce a working nuclear weapon of any sophistication. This lesson was learned as far back as the Manhattan Project, when physicists initially underestimated the scope of the project by several orders of magnitude, not realizing how much expertise beyond nuclear physics would be necessary to produce a weapon. A nation seeking to develop nuclear weapons has not only to cultivate a cadre of design experts, but to develop a whole set of distinct expert communities capable of working together.

Transactional Tacit Knowledge: Designers and Engineers

Designers and engineers I spoke to about the design process repeatedly mentioned the importance of compromise and negotiation in their interactions. The negotiations they describe engage issues of professional identity and power, but are also a deeply technical process. That is, rather than trying to simply force compromise by institutional means, participants are committed to mobilizing esoteric bodies of technical knowledge and making decisions directly justified by those bodies of knowledge. The sometimes adversarial character of the negotiations, however, suggests that the two bodies of knowledge have never really merged. Instead, the core knowledge of each group has remained distinct and apparently inaccessible to experts in other fields, further evidence that tacit knowledge is involved in weapons expertise. (In fact, this division of labor emerged during the Manhattan Project precisely because physicists quickly came up against the limits of their ability to master relevant engineering knowledge.)

Still, designers and engineers seem to have created a shared technical practice that involves its own knowledge of how to conduct technical transactions between the communities in a scientifically defensible way. I call the knowledge needed to effect this coordination *transactional knowledge*. Transactional knowledge encompasses Collins and Evans' interactional expertise – the ability to speak and understand the language of another community without being a practitioner – as well as knowledge specific to the shared context of practice, like understanding the allowable technical grounds for compromise and manipulating common formalisms.

A particularly technically challenging interaction between designers and engineers is assessing the impact of added engineering features (for example, a cable channel or structural support) on physics performance. As one designer explained, these impacts often come down to judgment rather than calculations:

[]

Here, the engineers need to have enough knowledge of weapons physics to identify components that could conceivably interfere with physics functionality, but at the same time rely on the expert judgment of designers to make the final determination on what configuration should be chosen.

A central feature of these negotiations, as described by both engineers and designers, is knowing when to insist on a particular configuration and when to back down, which requires some judgment about whether the other party's position is a matter of convenience or true technical necessity.

[]

However, the RRW design process provided an opportunity to hone these negotiation skills, and by the end designers and engineers appeared to have made significant progress in relearning how to make design compromises. This emphasis on being able to judge whether a person is taking a position for real technical reasons is a further indication that the knowledge of the two groups remains distinct. If both participants shared all of the same tacit knowledge, it would not be necessary to make such fine judgments about credibility.

Designers and engineers struggle to create a space of technical practice in which key aspects of each community's technical knowledge can coexist and influence one another without requiring complete integration between the communities. The transactional knowledge they deploy in this space simultaneously defines acceptable interactional styles and specifies the technical content of interaction. The fact that both the interactional norms and the technical knowledge appear to be difficult to acquire except through direct participation in the negotiating space suggests the relevant transactional knowledge is largely tacit.

Transforming Tacit Knowledge: Stockpile Surveillance

All of the scientists and engineers I talked to felt that lack of testing and design work made it difficult to successfully pass on their knowledge to a new generation. Designers were particularly concerned about this, and their concerns echoed those documented by MacKenzie and Spinardi in 1990. For example, one senior code developer, who had worked at the laboratory when design and testing were still going on, lamented that

[]

However, the younger designers (without test experience) I talked to appeared to be acutely aware of the difficulty of working in a post-design and testing environment, and of the limitations of codes. For example, one younger designer argued that

[]

Unlike older designers, however, some of the younger designers were willing to consider the possibility that some aspects of weapons knowledge had changed for the better since the end of the Cold War, despite the limitations imposed by lack of design and testing work:

[]

Both designers and engineers also felt that the working relationship between the two groups had fallen into decline since the end of testing. One engineer explained this in terms of loss of shared practice between the design division (X division) and the weapons engineering division (Engineering Sciences and Applications or ESA):

[]

Another engineer characterized the relationship in terms of a disconnect between designers and engineers:

[]

As the previous discussion of interactions between designers and engineers on the RRW project indicates, the transactional knowledge of engineers and designers was redeveloped, but only after a certain amount of trial and error in the context of a new design project.

Some tacit knowledge has surely been lost, and some working relationships broken, over 15 years without design and testing work. But this is not the whole story. Stockpile Stewardship also created new practices in the weapons community that developed new bodies of tacit knowledge and new contexts for interaction. In particular, I want to focus here on stockpile surveillance. During testing, surveillance was a tiny enterprise within the nuclear weapons complex, involving only one or two engineers at Los Alamos. Because weapon systems were typically retired before they began to age significantly, there was little need for surveillance data. As one designer explained, a couple of years after the end of design and testing, the surveillance program suddenly took on much greater significance:

[]

The surveillance program expanded until it employed about 15 full-time engineers at Los Alamos by 2000. Analysis of surveillance findings quickly became one of the main activities linking the design and the engineering communities, through an institutional mechanism known as the Significant Finding Investigation (SFI). In the context of these investigations, engineers carefully characterize any anomalies found in stockpile weapons and pass on this information to designers. Engineers, designers, material scientists, statisticians, and other technical experts then work together to analyze the anomaly, with designers making the ultimate determination whether it will have an impact on weapons performance. Any concerns identified through this process can then be addressed when the weapon undergoes a Life Extension Program, in which it is rebuilt for continued stockpile use. This process is similar to design and testing in the sense that it provides an alternative basis for assessing weapons performance through direct engagement with weapons as artifacts. As one engineer explained,

[]

Although this process is in no way seen as an adequate substitute for testing, it does at least provide a forum in which designers can interact with a wide range of other weapons experts to exercise their judgment about weapon performance. This process builds on an existing body of transactional tacit knowledge, although it also transforms it significantly. This has arguably resulted in the creation of a significant new body of tacit knowledge about nuclear weapons, focused on aging effects, and led to the emergence of surveillance engineering as a distinct technical field with its own esoteric body of knowledge.

The tacit knowledge involved in the analysis of surveillance findings could not have existed during design and testing because its object, the aging weapon, scarcely existed at that time. This new object, however, is only the result of gradual transformation, through aging, of the weapons stockpiled during the Cold War. As such, it provides a certain continuity in the knowledge of designers and engineers, who have been able to build on their Cold War practices to develop a new set of practices around the transformed object. The fact that this has occurred suggests that the likely fate of the tacit knowledge developed through the design and production of nuclear weapons is not simply to vanish, but to be transformed into something new that retains some relationship with the past. Whether this will eventually result in a body of tacit knowledge transformed beyond all recognition is a matter of judgment, but it may take more than one or two generations for this to occur. However, there is one extreme scenario in which MacKenzie and Spinardi's uninvention hypothesis could still hold: if all nuclear weapons work is shut down, the weapons design laboratories at Livermore and Los Alamos are closed or repurposed, and the production complex dismantled. Under such circumstances, which are politically improbable but not inconceivable, I believe that nuclear weapons would indeed be uninvented within a generation or two.

Conclusion

The cases presented in this paper suggest that new ideas about tacit knowledge are necessary if we are to understand the role tacit knowledge plays in complex, multi-disciplinary organizations. We already know quite a lot about how tacit knowledge is tied to particular technical communities and replicated across time and space. What is needed now is more emphasis on the interaction between distinct bodies of tacit knowledge, and on how these bodies of knowledge are coordinated to produce the kind of distributed, multi-disciplinary knowledge needed to execute large scientific and technical collaborations. Collins and Evans' recent work on developing categories of expertise, including interactional expertise, is an important effort in this regard, but in this work I try to extend tacit knowledge concepts in a more specifically organizational direction. I also place more emphasis on understanding not only how tacit knowledge appears and disappears over time, but how it can transform over time in a continuous way. These changes should make it possible to fully incorporate the concept of tacit knowledge into the literature on the large-scale structure and dynamics of scientific and technical work.

In this talk, I have tried to show that nuclear weapons knowledge is both more and less vulnerable to uninvention than MacKenzie and Spinardi and others have suggested. It is more vulnerable because it is distributed across so many technical specialties, and because it depends on maintaining the contexts for interaction among these fields as well as the core knowledge of each field. It is less vulnerable in the sense that it appears to be changing gradually rather than simply being preserved in static form until its practitioners leave this world.

This analysis may have implications for nuclear proliferation as well. When examining proliferation patterns, I would suggest it may be necessary to look beyond issues like whether a nation has the requisite technical experts and access to nuclear materials. We should also look at whether the country in question possesses the cultural outlook and experience with building interdisciplinary collaborations required to sustain a large-scale research and development program. Of course, the development of nuclear weapons could itself enable a country to develop these skills, in much the same way that the experience of the Manhattan Project set the stage for increasingly complex scientific and technological projects in the U.S. during the 20th century.