

The Evolution of Structured Analytic Techniques

By Richards J. Heuer, Jr.

Presentation to the National Academy of Science, National Research Council Committee on Behavioral and Social Science Research to Improve Intelligence Analysis for National Security, Washington, DC, December 8, 2009.

First, I want to thank the committee for the work you are doing toward our common goal of improving intelligence analysis, and for the opportunity to be here today to hopefully make some small contribution to this common effort.

I'll start with a little history. In 1975 I arranged to shift from the operations side of CIA to the analysis side in order to work in a new Analytic Methodology Division that had recently been created in direct response to criticism by prominent academics that the agency's analytic methods were way out of date. Our task was to examine quantitative methods that were developed in the 1960s during what was called the behavioral revolution in academic political science, and to test how these methods could be applied to intelligence analysis. I don't know the full story of the creation of this division, except that the order came down in a policy directive from the CIA Director, then Bill Colby. I can easily imagine that this recommendation originated with an NRC committee just like yours 35 years ago. In other words, I feel like I've been down this road before.

In 1977, at the annual International Studies Association convention, I chaired a panel called Quantitative Approaches to Political Intelligence: The CIA Experience. This was to show academia what we were doing in response to their criticism. Those panel presentations were subsequently published as a book that I edited, with that same title. This division did some very good and interesting quantitative work, but as far as I can tell, it did not lead to any long-term change in analytic methods. The only long term impact that came from the work of this division was my book on the Psychology of Intelligence Analysis. I'll tell you an interesting story of how that happened.

Immediately after that ISA presentation, I was approached by a man with a foreign accent who said, "Vee need to talk. Zee answer ees not in zee numbers, it's in zee head." That sounded interesting, so we had lunch, and he told me about Kahneman and Tversky's path breaking work in cognitive psychology. My new friend was an Israeli named Zvi Lanir. He was a senior officer in Israeli military intelligence at the time of the Yom Kippur War, when Israel was taken by surprise. It was a major Israeli intelligence failure. In the hope of avoiding such surprises in the future, the Israeli government sent him to the US for two years, working out of Columbia University, to study the new, quantitative approaches to political science that had been developed in the United States, and how they might be applied to intelligence analysis. So his mission for his government was exactly the same as our office's mission for the US government. However, he had quickly concluded, as he said, that the answer to good intelligence analysis is not in the numbers, it's in the head.

I then read Kahneman and Tversky's work, and that was the beginning of my continuing interest in cognitive psychology. I then began writing a series of papers and lecturing to CIA training courses on the cognitive problems and to a lesser extent the group process and organizational problems one encounters in doing analysis. One of the lessons I learned during that period is still very applicable today. If you want to change how analysis is done, you need to show analysts how you can help them. I could do that after my research in cognitive psychology taught me about how the mind works. You are not going to get very far by telling analysts that they are doing it all wrong and need to do it your way.

I retired from CIA in 1979 after 28 years of service, and moved to California with a contract to continue writing papers on cognitive issues, and this produced the bulk of the material that 15 to 20 years later was pulled together for publication in my book on the Psychology of Intelligence Analysis.

You may be interested in the origin of ACH, the Analysis of Competing Hypotheses methodology. Much of my operational experience and analytical experience at the agency was in counterintelligence. After retiring I published several pieces that were very well received on the psychology of deception. In 1984, the Reagan administration became concerned about Soviet deception in arms control agreements, and I was asked to develop a methodology for analyzing deception and to help develop and teach a new, interagency course in Deception Analysis. The problem we had to deal with was that when you asked analysts if they considered the possibility of deception, they would always respond: "Of course I did. I'm not stupid, but I didn't see any evidence of deception." Then you point out that if deception is well done, one should not expect to see evidence of it. Oops. The analysts then recognize that, well, maybe there is a problem you can help with.

So I developed ACH 25 years ago as the centerpiece for a new training course in Deception Analysis. The goal was to get analysts to seriously consider all the reasonable hypotheses, and to keep the deception hypothesis open until it can be refuted. I remember several times when students told us jokingly that they were lured to this course under false pretenses. They said: "This is not a course in deception analysis; it's a course in how to do good analysis of any type. I believe that course is still being taught, and the ACH method has been used for 25 years now by analysts working on deception issues, but it wasn't until my book came out in 1999 that it started getting broader usage. The automated version of ACH became widely available in 2005, and that has certainly boosted further usage.

My forthcoming book with Randy Pherson on Structured Analytic Techniques is scheduled to be published in mid-February. You've all seen several key chapters, so you should have a fairly good idea of what's in the book. We see this book not just as a collection of 50 techniques, but as one more step, and fairly a significant step, in the gradual evolution of how intelligence analysis is done.

The new insight helped to drive my work on this book is that structured analytic techniques are the *process* by which collaboration becomes most effective. I got the first

indication of this several years ago after CIA started using the ACH software, and we got feedback that what analysts like most about it is that they could get together and use ACH in a group, and that going through the step-by-step ACH process, identifying hypotheses, relevant evidence, and how each item of evidence relates to each hypothesis, helped the analysts identify early in any new project all the things they disagreed about and new information they hadn't known about. That happens because you have the whole analysis pictured there in the matrix. And being able to point to the matrix helped them depersonalize the disagreements.

Then writing the book with the additional techniques made me realize that this benefit is not something unique to ACH. All the structured techniques provide this same benefit. All these techniques guide the dialogue between analysts with common interests as they share evidence, share alternative perspectives, and discuss the meaning and significance of the evidence. Just as these techniques provide structure to our individual thought processes, they also structure the interaction of analysts within a small team or group. Because the thought process in these techniques is transparent, each step in the technique prompts relevant discussion within the team, and such discussion generates and evaluates substantially more divergent information and more new ideas than a group that does not use a structured process. Our experience is that this process helps to avoid the multiple pitfalls and pathologies that often degrade the performance of small teams or groups, and this experience is consistent with recent research on information sharing and team performance.¹

If there is one thing you take away from my presentation to day, please let it be that structured analytic techniques are *enablers* of collaboration. They are the *process* by which effective collaboration occurs. Structured techniques and collaboration fit together like hand in glove, and they need to be promoted and developed together. Unfortunately, the DNI leadership has not recognized this. For example, the DNI's *National Intelligence Strategy*, Enterprise Objective 4 on improving integration and sharing, makes no mention of improving analytic methods.²

The DNI gets credit for pushing hard for collaboration, and for developing the secure computer network that helps make it possible, but it does not mention structured techniques as a process for guiding that collaborative interaction. I see structured techniques as the final element needed to make all this work well together. At the National Intelligence Council, the use of structured techniques is the exception not the rule. Hopefully, our book, and this committee, through your recommendations, will help to change that. For the kinds of analysis the NIC does, the use of structured techniques should be the norm, not the exception.

CIA, not the DNI, is the agency that has been pushing structured analysis. One important innovation at CIA is the development in various analytic offices of what are called tradecraft cells. These are small groups of analysts whose job it is to help other analysts

¹ Jessica A. Mesmer-Magnus & Leslie A. DeChurch, "Information Sharing and Team Performance: A Meta-Analysis," *Journal of Applied Psychology*, 2009, Vol. 94, No. 2, 543-544.

² *The National Intelligence Strategy*, August 2009, available on the DNI website.

decide which techniques are most appropriate, help guide the use of such techniques by inexperienced analysts, and often serve as facilitators of group processes. These tradecraft cells are a very helpful innovation that should spread to the other agencies. They are a key element of any program to promote collaboration and improve the quality of analysis through the use of more structured thinking. CIA has learned from experience that they are an essential part of the process by which any change in how analysis is done gets implemented, and I certainly hope you will include that in your recommendations.

I understand you are all concerned about evaluating whether these structured techniques actually work. So am I. I'd love to see our methods tested, especially the structured analytic techniques Randy and I have written about. The only testing the Intelligence Community has done is through the experience of using them, and I think we all agree that's not adequate. The question that needs to be decided is how to go about the testing.

I understand that the committee has had considerable discussion and debate about this question. Some of you have emphasized the need to test the accuracy of these techniques. That would certainly be the ideal, but ideals are not always achievable. Unfortunately, there are major difficulties in testing structured techniques for accuracy, and the chances of such an approach having a significant favorable impact on how analysis is done are not very good. I see four reasons for this.

1. Testing for accuracy is difficult because it assumes that the accuracy of intelligence judgments can be measured. In practice, however, it is very difficult to assess the accuracy of intelligence judgments even after the fact when the results are known. That is because most judgments are conditional upon other things happening or not happening, or upon the validity of some assumption or some source. We say that X or Y very likely will happen if... or unless.... This is why people are still arguing about the accuracy of CIA's analysis of the Soviet Union 20 years after the end of the Cold War. People are still arguing about the bona fides of KGB defector Yuri Nosenko, a case that I was involved in 35 years ago. So the conclusions from this type of analysis are unlikely to be definitive.

2. There is a subset of analytic problems such as elections, when a definitive answer will be known in 6 or 12 months. Even in these cases there is a problem in measuring accuracy, because intelligence judgments are almost always probabilistic. We don't make predictions that something will happen or won't happen. We say it probably or very probably will or won't happen, or that it's highly unlikely. To test the relative accuracy of two different analytic methods, one would need to be able to distinguish an analytic technique that is accurate 80% of the time from one that is accurate 70% of the time. That would require quite a large number of experiments.

3. A third reason why a major effort to evaluate the accuracy of structured analytic techniques may not be feasible stems from our experience that these techniques are most effective when used as part of a group process. That means that the results of experiments using analysts working alone would not be generalizable to how the techniques are most commonly used in the Intel Community. It may not be practical to

test the effectiveness of a technique when each so called test subject is actually a small group, and, to make it worse, that is combined with the need to conduct a large number of experiments to account for the probabilistic nature of the results of any single experiment. The number of test subjects required to conduct experiments that are generalizable to how techniques are used in the Intelligence Community may make such experiments impractical.

4. If you are trying to change analysts' behavior, which has to be the goal of such research, you are starting with at least one strike against you, as much of your target audience already has a firm opinion, based on their personal experience that they believe is more trustworthy than your research. If your research confirms what they already think they have learned from their personal experience, they will accept it readily. If it does not agree with their personal experience, astute analysts will look at that experiment and ask: Was the technique used the same way that I use it? Is the issue analyzed comparable to the kinds of issues I analyze? Are the analysts as well trained and experienced in using that technique as I am? Are the test subjects as knowledgeable about that issue as I would be? If you can't answer yes to all those questions, good analysts will usually be guided by their personal experience in using a technique rather than by any research on the technique.

I believe the most feasible and effective approach for evaluating structured analytic techniques is to look at the purpose for which a technique is being used, and then to determine whether or not it actually achieves that purpose, or if there is some better way to achieve that same purpose.

Our book has eight chapters of techniques. The introduction to each chapter discusses the rationale for using the type of techniques in that chapter. How does using this type of technique help the analyst do a better job? And then for each individual technique the book describes not only how to use that technique, but also when to use it and the value added by using it. We need a program for testing and verifying whether or not those benefits are actually achieved, or whether we are deluding ourselves. And it is important that this testing be done in the natural environment where those techniques are actually being used.

For example, instead of trying to test the accuracy of conclusions drawn using ACH, which may vary greatly depending upon the nature of the problem being analyzed and the skill of the analyst, let's develop and test hypotheses about how the use of ACH and other techniques may or may not help an analyst do a better job. Here are three hypotheses that can be tested to evaluate whether ACH is an effective analytic technique.

1. Requiring that analysts start by developing a full set of alternative hypotheses reduces the risk of a common pitfall called satisficing, being satisfied with the first answer that seems to meet the need rather than looking for the best answer.
2. The process of trying to refute hypotheses rather than support a single favored hypothesis reduces what is called confirmation bias, thereby ensuring a more

careful consideration of other alternatives than typically occurs when ACH is not used.

3. The ACH process, as compared with traditional analysis, helps analysts present their conclusions in ways that are better organized and more transparent as to how these conclusions were reached.

I have personally made all those claims about ACH, but I don't expect everyone to believe them just because I say so. Some should be tested and documented. We also need research to find out what kinds of mistakes are made in using ACH, and how they can be avoided.

Procedures used for testing structured techniques can include simple empirical experiments, comparing analyses done with and without a specific technique against some common standard of analytic quality, structured interviews of analysts, managers, and customers; observing meetings of analysts as they use these techniques. All of those things can easily be done in house, but they're not being done now because no element of the Intelligence Community the charter and funding and staff to do that. That's the problem that I hope this committee will address.

I'll give you three additional examples of testing that I would like to see done. First, consider the issue of early warning indicators of a military attack. Unfortunately, some early warning indicators of military attack are ineffective at best, and misleading at worst, because they may also be consistent with a no attack scenario. For example, troops may have been mobilized in a threatening manner to create pressure for a political solution, but with no intent to actually attack. In other words, some indicators that seem logical actually have little diagnostic value. Randy Pherson developed a simple technique called the Indicators V alidator that takes an analyst through a series of steps that test the diagnosticity of indicators.

That technique was developed to meet a specific need. It could be tested easily by applying it to several established sets of indicators that are now being used, to see how many flunk the test. The Indicators V alidator can also be tested by asking two groups of test subjects to develop a set of indicators, with one group using the Indicators V alidator and the other not using it. Both of these tests should be easy to do, but again, where do you go in the Intel Community today to get that testing done?

The Key Assumptions Check is one of the most commonly used techniques, and I think we are all aware of the important role that assumptions play when analyzing any situation with multiple uncertainties. Actually assumptions may play such an important role that it would be advisable to recommend that a Key Assumptions Check should be done as part of all major analytic projects. Before making such a recommendation, however, I would want to do three different tests of the Key Assumptions Check technique:

1. Survey analysts to determine how often they use this technique, their criteria for when to use it, procedures for using it, what they see as the benefits gained from using

it, and when they have used it did it affect their judgment about the most likely explanation or outcome?

2. Compare the quality of a random sample of reports written without having done a Key Assumptions Check with reports written after such a check. How were the results of the Key Assumptions Check reflected in these products?

3. Determine whether a single analyst can effectively identify his or her own assumptions, or should this always be done as a small group process? The answer to this question seems obvious, but it is easy to document the magnitude of the difference with a before-and-after comparison. Analysts are commonly asked to develop their own set of assumptions before coming to a meeting to do the Key Assumptions Check, and these initial lists can be compared with the list developed as a result of the meeting. Documentation of the difference may be sufficient to warrant making the group process standard practice.

As with the other examples, however, the Intel Community has no organizational unit that is funded and qualified to do that sort of testing.

A final example is based on an article I saw in the Wall Street Journal several weeks ago. It was about the failures of economic analysis that led us to our current economic problems. It suggested that instead of estimating the likelihood that their plans will work, financial analysts should estimate the probability they might fail. That's a good idea that could also be applied to intelligence analysis.

I suspect that estimating the probability of your being wrong would substantially change the psychological context in which the probability estimate is made, and that it may provide a more objective/accurate estimate of probabilities than focusing only on the probability of being right. This could be used to help analysts better calibrate their judgments of probability and uncertainty. Obviously that would need some testing, but again, there is no place in the Intel Community where one can take that idea and get it tested.

There are differences of opinion about how and where such research should be done. Some want it done in some type of special institute with highly skilled methodologists so that it is totally scientific and objective, and not subject to any perverse influence by the analytic components defending the status quo. Experience tells us that will not work. With us today in the audience is Nancy Dixon, who has done several projects for the DIA Knowledge Lab that involve extensive interviewing of DIA analysts. One recent project was called "The Problem and the Fix for the U.S. Intelligence Agencies' Lessons Learned."³

It's about why so called lessons learned are so seldom implemented in the Intelligence Community. She concluded that "the likelihood that action will be taken based on the

³ Available on the Internet at <http://www.nancydixonblog.com/2009/07/the-problem-and-the-fix-for-the-us-intelligence-agencies-lessons-learned.html>

findings of a lessons learned study is greatly increased when those who are expected to implement the findings are involved in all phases of the process of creating the lessons learned." That makes a lot of sense to me, and that affects how and where I would like to see the research done.

The kind of research I just talked about can and should be done in-house with the assistance of those who are directly responsible for implementing the findings. It should not be contracted out, although contractors may be brought in house to provide special skills that are not available on the government staff. If the research is done in-house with real analysts working on real issues in the real environment in which analysis is actually done, the chances of lessons learned being accepted and followed should be pretty good.

Our forthcoming book has two major recommendations for DNI actions that we believe are needed to achieve the analytic transformation we would all like to see.

1. The DNI needs to require that the National Intelligence Council set an example about the importance of analytic tradecraft. NIC projects are exactly the kind of projects for which structured techniques should *always* be used, and this is not happening now. Their use is very much the exception, not the rule. The NIC projects are certainly important enough to warrant a little extra effort and a little extra time. And they usually involve the entire Intel Community, so this is a good way to spread the knowledge of these techniques throughout the Community. If the NIC isn't using structured techniques on a regular basis, how can we expect other analysts to do so? The NIC needs to be our model that sets the standards for others to follow!

We recommend that the DNI create a new position on the National Intelligence Council – a position that might be called vice chairman for analytic tradecraft. This new vice chairman should ensure that structured analytic techniques and appropriate collaboration procedures are used in all analytic projects for which the NIC has oversight. The vice chairman should also manage an after-action review of lessons learned after each major NIC project, based on input from all those who contributed to the project. The goal being to identify improvements in the NIC process. The case for such after-action review is also discussed in Nancy Dixon study. She argues that organizational learning needs to be embedded in the workflow process.

2. The second recommendation is that the DNI should create what might be called a center for analytic tradecraft. This center should be responsible for a program to evaluate existing analytic techniques and for development and testing of new techniques. It should also promote the development of tradecraft cells throughout the Community, and provide oversight, guidance, and assistance (as needed) to those cells. It should help and encourage analytic Communities of Interest, for example, by ensuring that they have appropriate technical support. It should be responsible for establishing and managing an IC-wide knowledge management system for the creation, capture, storage, and sharing of lessons learned about the use of structured analytic techniques.

Finally, it may be advisable for the NIC vice chairman for analytic tradecraft to be dual-hatted as deputy chief of the center for analytic tradecraft. This would ensure that the vice chairman of the NIC will have available whatever resources are needed to support the NIC projects. And it would also ensure that the work of the center for analytic tradecraft remains closely attuned to the needs of the NIC and the analytic community in general.

I hope this committee will support these recommendations, and I thank you again for the opportunity to be here and speak with you today.