Intelligence, Science and the Ignorance Hypothesis

David R. Mandel

Intelligence organizations perform many functions. For example, they conduct covert operations that would be too politically sensitive for militaries to undertake as explicit missions. They also collect vast amounts of data that are inaccessible to others, except perhaps other intelligence organizations. However, the principal purpose of information collection is to produce substantive intelligence that can inform policymakers, commanders and other decision-makers who guard their nation's interests and security. The fundamental premise of intelligence is that it serves to improve the planning and decision-making of these elite decision-makers by elevating the debate about policy options.¹ If this were not so, the multi-billion dollar annual budgets would not be justified. If the premise is justified, then clearly the intellectual rigor and accuracy with which intelligence assessments are produced is of paramount importance.

How do intelligence organizations ensure that the intelligence assessments they produce are as accurate and sound as they can be? In this chapter, I will propose that the short answer is "not very well at all." Methods and policies for ensuring analytic rigor have surely been implemented over the past several decades, but what has been endemic to these efforts is a rather pre-scientific, if not a fully anti-scientific, attitude towards their development and testing. After reviewing some examples of intelligence practices intended to ensure analytic rigor, I advance what I call the *ignorance hypothesis* to explain the absence of scientific attempts to discern what practices work from those that do not work so that intelligence organizations can effectively learn and adapt to the

challenges of the modern world. At face value, the ignorance hypothesis proposes that the absence of adequate scientific testing of analytic policies and practices is due primarily to widespread ignorance of scientific principles and values, within both intelligence and policy communities. At a deeper level, however, the ignorance hypothesis posits that there is also something special about the topic of analytic rigor that makes it especially impervious to scientific thinking. Before turning to these “why” questions, I must first at least sketch what intelligence communities are aware of vis-à-vis analytic rigor and what types of institutional responses they have offered. In so doing, I will focus on the US context, although the general points also apply to other intelligence communities.

Intelligence Analysis as Corruptible Human Judgment

In spite of the vast and impressive technologies brought to bear on collections challenges, intelligence organizations rely almost exclusively on human analysts to make the judgments that constitute finished intelligence. Likewise, the same analysts have considerable leeway in deciding how to express those judgments to their target audiences. As Sherman Kent aptly noted, substantive intelligence is largely human judgment made under conditions of uncertainty.² Among the most important assessments are those that not only concern unknowns but also potentially unknowables, such as the partially formed intentions of a leader in an adversarial state. In such cases, the primary task of the analyst is not to state what will happen but to accurately assess the probabilities of alternative possibilities as well as the degree of error in the assessments and to give clear explanations for the basis of such assessments.³

Intelligence communities are certainly not unaware of the problems inherent in human judgment under uncertainty. Maverick figures like Sherman Kent pioneered methods for improving the communication of uncertainty in assessments, and Richards Heuer Jr. not only summarized a great deal of relevant cognitive research for the US intelligence community in his 1999 book *Psychology of Intelligence Analysis*, he also pioneered several of the structured analytic techniques (SATs) that are still used by intelligence communities to this day.⁴ Indeed, it is common knowledge in intelligence organizations that humans are fallible and corruptible in many ways. They are often unreliable and/or systematically biased because the “fast” cognitive processes or “heuristics” they have been biologically adapted to use in prehistory are error-prone under given conditions.⁵ For example, when judging how frequent one event class is compared to another, humans often rely on the *availability heuristic*—namely, the ease with which instances in each event class come to mind.⁶ Thus, judgments of frequency will be influenced by factors affecting mental availability, such as advertising, social media and—yes—“fake news”, which were not biasing factors in human prehistory.

In a similar vein, intelligence communities are generally aware that probability, which is so central to intelligence assessments, is often judged using the *representativeness heuristic*—namely, the process of assigning probability based on how well individuating information seems to match alternative hypotheses.⁷ Thus, humans often fail to consider how influential the prior probabilities of the alternative hypotheses are on the posterior probabilities they are judging, except when the prior probabilities serve as anchors in which case humans tend to be overly “conservative,” which in the present context means they do not react fast enough to new and diagnostic information.⁸

This can cause them not only to be inaccurate, but also to be incoherent, such as when they judge a representative conjunction of two events, *A* and *B*, to be more likely than one alone—namely, what is known as the *conjunction fallacy*.⁹

What is more, however, the cognitive biases to which humans (including intelligence analysts!) are susceptible are not easily self-detected. When humans are overconfident, they tend to believe that they are right but they are unlikely to believe (accurately) that they are overconfident. Psychologists call the inaccessibility to one's own cognitive biases the *bias blind spot*.¹⁰ This “meta-bias” does not appear to be attenuated by greater cognitive sophistication and may in fact be correlated with intelligence, or at least markers of intelligence such as academic success.¹¹ Moreover, any benefit of training aimed at debiasing the bias blind spot seems to decay rapidly.¹² Stated plainly, people lack self-awareness of many, and probably most, of their cognitive biases. This ignorance state would be bad enough if it were a Rumsfeldian “unknown unknown”—namely, “mere” ignorance of the nature and severity of one's cognitive biases. However, cognition (i.e., reasoning and judgment), the central tool of the analyst, remains more akin to what Rumsfeld called the “unknown known” at the beginning of his interview with Errol Morris in Morris' 2013 documentary on the former US Secretary of State.¹³ The unknown known refers to something one thinks one knows to be true that turns out to be false. Not only do humans fail to detect their cognitive biases, they are quite convinced that such biases do not pose a problem for their thinking and deciding. Much as humans are adapted for cognitive bias, they are also biologically adapted for self-deception, which facilitates their ability to deceive others.¹⁴

The problems associated with “cold” cognitive biases are compounded by “hot” motivational biases,¹⁵ which can corrupt the integrity of intelligence production. Intelligence analysts might find it difficult to follow the mantra of “speaking truth to power” when the views of the powerful are known in advance and the powerful themselves are not known for their open accommodation of dissenting analysis. This applies to analysts in relation to intelligence directors and to analysts and directors in relation to their intelligence clients. Accountability pressures in response to career-influencing audiences can trigger defensive bolstering of assessments or pre-emptive self-criticism, both of which are extra-evidentiary psychological processes aimed at minding one’s reputation as an “intuitive politician” rather than focusing on achieving the most accurate, well-calibrated assessments possible under the circumstances.¹⁶ Perhaps this is why, although overconfidence is a well-documented bias,¹⁷ strategic intelligence forecasts systematically examined over several consecutive years have been found to be substantially *underconfident*.¹⁸ That is, even accurate forecasts tended to be communicated as a series of watered-down, hedge-filled estimates, and this was after a considerable proportion of the forecasts were excluded because of their reliance on unverifiable weasel words. Given that there is far more to lose by overconfidently asserting claims that prove to be false than by underconfidently making claims that prove to be true, intelligence organizations are likely motivated to make timid forecasts that water down information value to decision-makers—a play-it-safe strategy that anticipates unwelcome entry into the political blame games that punctuate history.¹⁹

Of course, the cold/hot dichotomy is closer to a rhetorical device than a scientific fact. The truth is that all products of judgment are the result of cognitive processes

undergirded by biochemical mechanisms. “Hot” and “cold” are useful metaphors insofar as they help us categorize human judgment and reasoning biases into “motivated” and “unmotivated” types. The types described as motivated, however, can vary greatly in their emotional intensity and goal directedness. For example, the underconfidence caused by strategically shrouding assessments in more uncertainty than required is “lukewarm” compared to the “fiery” political biases that some former intelligence directors and senior staff members exhibit towards their political opponents. As Gentry notes, the anti-Trump diatribes of former Central Intelligence Agency director, John Brennan, former Federal Bureau of Investigation director, James Comey, and former Office of the Director of National Intelligence (ODNI) director, James Clapper against President Donald J. Trump provide recent examples.²⁰ In such cases, the pursuit of truth and objectivity is sacrificed in favor of ends-justify-the-means tactics of “intuitive theologians”, who are often firmly convinced that they are on the right side of history. If the gravitational centers of power in such organizations hold overtly biased and uncertainty-intolerant views of the world, how can the organizations they lead possibly be expected to guide decision-makers to truthful, objective information about a complex world?

Untested Steps

Given that intelligence organizations principally rely on humans to produce substantive intelligence and given the predisposition of humans to exhibit a wide range of cognitive and motivational biases, what have such organizations done to mitigate bias and what have they done to verify whether their adopted methods are effective? In the US, the strategy taken to deal with this immense challenge has been to issue vague process directives that aim to promote the fuzzy concept of “analytic integrity”.

Specifically, ODNI's Intelligence Community Directive 203 (ICD203) admonishes all US analysts to provide assessments that are "objective", "independent of political consideration," "timely", "based on all available sources of intelligence information," and the analyst must further ensure that he or she "implements and exhibits Analytic Tradecraft Standards", which include (a) "properly describ[ing] quality and reliability of underlying sources," (b) "properly caveat[ing] and express[ing] uncertainties or confidence in analytic judgments," (c) "properly distinguish[ing] between underlying intelligence and analysts' assumptions and judgments," (d) "incorporat[ing] alternative analysis where appropriate," (e) "demonstrat[ing] relevance to U.S. national security," (f) "us[ing] logical argumentation," (g) "exhibiting consistency of analysis over time, or highlight[ing] changes and explain[ing] rationale," and (h) "mak[ing] accurate judgments and assessments".²¹

On the surface, who can object to such directives? Indeed, when a sample of over 100 Canadian intelligence professionals (who are not mandated to follow ICD203) was asked how strongly they valued such principles, each principle was strongly endorsed. Moreover, professionals who judged their organizations to be more compliant with these principles reported being more satisfied with their jobs as well as affectively and normatively committed to their organizations.²² This is all well and good, but how do intelligence organizations actually promote the achievement of these aspirational principles and how do they objectively measure how well each is met? Since the criteria are vague, it is unsurprising that individual raters' evaluations of how well an analytic product meets the ICD203 criteria are unreliable.²³ How, for instance, does ODNI measure objectivity? Is objectivity to be measured objectively through correspondence

measures such as proper scoring rules for forecasts, subjectively by analyst self-assessments or director evaluations, or inter-subjectively, for example, through consensus-seeking feedback from a group of peer experts? With rare exceptions,²⁴ intelligence organizations do not score their assessments using objective measures such as proper scoring rules (e.g. “Brier scores”). Therefore, the directive to “be objective” is either subjectively or inter-subjectively assessed to the extent that it is assessed at all. Self-reflection and peer-evaluation may each have beneficial effects on analytic practice, but neither constitutes a valid basis for gauging objectivity. One can hardly expect skeptical analysts to take objectivity seriously if the only way the requirement for objectivity is assessed is through subjective or inter-subjective measures. If the IC’s approach to objectivity is questionable, the directive to remain independent of political consideration is laughable. How can analysts or observers of the intelligence community take such statements seriously when the political biases of former directors are put on stark display?²⁵ Is it likely that the overt political biases they routinely exhibit upon taking up media commentator roles after their retirements were all but absent just a short time before when they were powerful directors? Would analysts working under them fail to notice?

As noted earlier, a related strategy for promoting analytic integrity is to advise analysts to use SATs. The “philosophy” behind SATs is that analysts cannot be trusted to reason in ways that would uphold the ICD203 criteria and, therefore, they need to be aided by techniques that should debias their judgments and improve the reliability and accuracy of their assessments. Most SATs, however, lack a sound philosophical or scientific basis in development and testing. Coveted SATs such as Heuer’s Analysis of

Competing Hypotheses (ACH),²⁶ which are taught to analysts in several countries, are dubious from the start. For instance, ACH guides analysts to evaluate each piece of evidence independently of all other pieces of evidence against each hypothesis. While “independent assessment” might sound positive, it militates against reasoning about the significance of configural fact patterns that provide keys to mysteries, which a process of evaluating each fact on its own would fail to reveal. Siloing evidence into distinct pieces to be evaluated separately can also militate against proper consideration of correlational structure in evidence. Karvetski, Mandel, and Irwin found that when two perfectly correlated sources of evidence presented in a hypothesis-testing task were made highly salient, participants who used ACH were more likely to double-count (rather than discount) the second, redundant source than when the redundancies were less salient.²⁷ Rather than correcting for redundancy, which a normative (e.g., Bayesian) approach would require, ACH amplified the mishandling of redundancy. ACH also assumes that evidence that is deemed to be consistent with a hypothesis should be given no weight in the calculation of support for alternative hypotheses. Rather, support is treated as an inverse function of the frequency of inconsistent evidence for a given hypothesis. As I have suggested elsewhere,²⁸ the rationale for this counter-normative information integration rule may have stemmed from Heuer’s interpretation of Sir Karl Popper’s theory of falsification.²⁹ However, Popper’s theory does not apply to non-universal hypotheses of the sort that intelligence analysts routinely deal with. “All swans are white” exemplifies the type of hypothesis Popper was addressing. No matter how many white swans one observes, it cannot prove the universal hypothesis to be true. However, finding a single non-white swan is sufficient to conclusively disprove the hypothesis. Hypotheses

in intelligence are neither about universals (i.e., “*All x are y.*”) nor is the evidence brought to bear on the evaluation of such hypotheses invariably unambiguous, such as whether a swan is white or not.

Given the conceptual failings of ACH, it is perhaps unsurprising that experiments testing its effectiveness have revealed disappointing results. Whitesmith found no evidence to suggest that ACH reduces cognitive biases or improves judgment accuracy.³⁰ Mandel, Karvetski, and Dhami found that intelligence analysts who were trained in and instructed to use ACH on a hypothesis-testing task were significantly less coherent in their judgments than analysts who were not trained in ACH and not asked to use any SAT in the task.³¹ In follow-up analyses from the same sample, the same authors found that analysts using ACH were less likely to use relevant base-rate information than analysts who did not use ACH.³² Karvetski et al. found that participants asked to judge the probability of alternative hypotheses being true were less accurate after using ACH than they were prior to using it.³³ Karvetski and Mandel found that participants who used ACH were, in multiple respects, less coherent in their judgments than participant who did not use the technique.

The preceding comments could easily lead one to believe that, while ACH is problematic, the overall SAT approach may be sound. Unfortunately, this is not so. Closer examination of other SATs reveals serious conceptual limitations. For instance, the SAT that Heuer and Pherson called Indicators Validator,³⁴ which is used for determining information collection priorities, has no basis in information theory and has been shown by Timms, Mandel, and Nelson to lead to collections recommendations that run counter to normative models of information utility, in part, due to its neglect of

relevant base-rate information.³⁵ More generally, the problem is that the intelligence community does not appear to have sufficiently drawn on relevant areas of science or philosophy in the idea generation phases of SAT development, nor has it utilized scientific methods to test whether the SATs they recommend that analysts use actually improve intelligence analysis one iota. This lack of scientific rigor persists despite many calls over the years for proper scientific attention to the development of methods for intelligence analysis.³⁶

To be fair, over the last decade, the US intelligence community, through ODNI's Intelligence Advanced Research Programs Activity (IARPA), has invested vast sums of money into programs aimed at improving intelligence analysis through funded research conducted by academic institutions and industries. This is a positive development and, in some sense, a valid response to calls for better scientific engagement. On smaller scales, some other intelligence communities are probably doing likewise. However, what remains unclear is how the new research is influencing intelligence training, current on-the-job practices, and executive-level plans to address the yawning gap between the status quo and optimal intelligence practice. From extensive discussions with tradecraft experts in several countries over more than a decade now, I have never gotten a sense that the status quo is being fundamentally transformed. Nor, as I have argued elsewhere, is the IARPA model proven to be a good model for the intelligence community's challenges.³⁷ While it beneficially shifts focus from over-reliance on "occasional mavericks" like Kent and Heuer to programmatic applied research, it is also limited by its high risk/high reward mandate and a contractual process that locks funded research teams into the equivalent of scientific straightjackets. In this sense, IARPA has not fulfilled Rieber and

Thomason's call for a "National Institute for Analytic Methods"³⁸ because it does not routinely empirically test in randomly controlled trials the methods that analysts in the intelligence community are instructed to use. Yet that is precisely what is needed to separate the wheat from the chaff—namely, to know what works and what doesn't. Intelligence organizations should be routinely performing such tests of their methods and adapting their analytic practices in response to the evidence they collect. They should also be tracking scientific research that directly tests the effectiveness of their methods, such as the several studies cited earlier, and they should be tracking developments in research fields, such as the cognitive and decision sciences, which have good prospects for informing the development of new analytic methods.³⁹

IARPA programs are also often launched as tournaments between competing "performer" teams that are scored roughly annually for the purpose of winnowing the performers. The tournament ends with a single winner, while the rest are presumably "losers." This process distorts the value of the funded research by making the winner's contributions appear more important than they actually are and the losers' contributions appear correspondingly less important. In reality, most teams will conduct research of value, and the differences in practical value to the intelligence community may be marginal, just as the difference between Olympic gold and silver medalists in a race can be a fraction of a second. Uncertainty is a key facet of all scientific research and scientists should highlight the limitations of their research, which might cast doubt on the reliability and validity of their findings. However, "performing" for funders incentivizes masking uncertainties and instead telling the most compelling narrative possible about the findings. Science aims to disentangle multiple causes of focal effects, yet it is in the

interest of performers to bundle as many possibly effective tactics into a brand in the hope that some subset of these tactics will eke out a victory. One could hardly blame the victors for this. It would be naïve to think that scientists do not act as intuitive politicians too.⁴⁰

The Ignorance Hypothesis Revisited

The intelligence community is aware that human judgment is central to intelligence production and that judgment is highly susceptible to cognitive and motivational biases, unreliability, imprecision and ultimately, error. Furthermore, it has taken steps to improve analytic rigor in day-to-day practices by implementing community-wide directives on analytic integrity and by training analysts to use SATs and then encouraging them to use them on the job. However, most of these endeavors are of questionable value because they have not been tested to verify if they work as intended. The same can be said of the intelligence community's methods for communicating uncertainty to end-users, which are based on the untenable assumption that the meaning of vague verbal probability terms can simply be stipulated by organizational fiat.⁴¹ Ample research,⁴² including on uncertainty communication methods used in intelligence production,⁴³ have shown that such attempts at establishing meaning by decree are ineffective, yet intelligence communities have doubled down on these methods in many cases, such as in the US and UK, turning them into community-wide reporting requirements.⁴⁴

Why have intelligence communities not done a better job of drawing on science to develop methods for effective intelligence production and why have they not also drawn

on well-established research methods for testing their best ideas in order to determine what works and what doesn't? As noted earlier, the ignorance hypothesis posits that the "experts" tasked with improving the quality of judgment in intelligence are not steeped deeply enough in scientific culture. They are unfamiliar with basic scientific concepts and lack sufficient knowledge about substantive areas of scientific research that are pertinent to improving judgment quality. They do not have the scientific tools at their disposal to approach the challenges they face with the requisite degree of skepticism. The experts tasked with making methodological improvements to intelligence practice typically don't want to get "down in the weeds" about research aimed at testing their preferred solutions. In all likelihood, their organizations provide them with no formal support mechanisms for such research even if they wanted it, and some may very well want it.

However, I doubt that most do. The likelier scenario is that most do not approach the task from a scientific frame of reference at all. Rather, they are likely to perceive their task as deciding what to borrow from existing methods and then deciding whether they want to adapt the methods in some number of ways to make it more sensible, useful or facile. In this process, they rely primarily, if not exclusively, on their "best judgment," which, in practice, means what they believe to be true or at least what they judge to be plausible. They may synthesize elaborate explanations of why a particular SAT is useful in training or why another isn't, but these explanations are meant to stand on their own; they are not treated as working hypotheses that give rise to necessary experimental tests of their validity. If challenged on the effectiveness of the methods they propose, they might cite anecdotal evidence to support their decisions, not realizing that such evidence is of low reliability, validity, and demonstrable generalizability.

Scientists also have beliefs about what is right and wrong. They formulate theories and hypotheses, but these must then be tested under controlled procedures that enable valid inference to be drawn from the test data. Moreover, if a test is designed in such a way that is biased in favor of the hypothesis, peer reviewers may challenge it. If the researchers nevertheless manage to publish it without correcting the flaws, other researchers may challenge it by designing better experiments and publishing their results, and these may cast doubt on the validity of the original research and its findings. Certainly, in practice, the system is not optimally self-correcting. However, it is the best system for knowledge generation that has ever been invented by our species.⁴⁵

The flaws in virtually all of the organizational interventions I have mentioned remain undetected by the intelligence community precisely because they do not apply a scientific lens to the problem or to their proposed solutions. For scientists, the reliability and validity of tests and measures is of paramount importance. Thus, it is natural for scientists such as Marcoci et al. to ask how reliable is the scoring of analytic products on ICD203—can it be done?⁴⁶ It is not merely that asking such questions is important. One must know how to translate those questions into experimental tests that are likely to yield test data from which one can draw valid inferences. Doing so requires considerable knowledge of research methodology, which most tradecraft experts I have encountered do not possess.

If those in charge of minding the system of intelligence production are not scientifically minded, as the ignorance hypothesis posits, then the current state of affairs

could have been anticipated. Why have those higher up the management and oversight chain not seen this coming? Is it because most directors, policymakers, and other elite decision-makers are also insufficiently steeped in scientific culture to have seen the problem? If so, then why does the intelligence community draw heavily on science in other areas of interest, such as advancing collections technology? The discrepancy seems to be related to the topic. Where intelligence methods concern human affairs such as judgment and decision-making and person-to-person communication, intelligence professionals seem to trust their intuitions about which methods will be beneficial and which won't. In other words, they lack skeptical prior probabilities about their skill in picking winners and avoiding dud methodologies. They will be the first to acknowledge that their solutions aren't perfect, but most I have spoken with accept the general premise that whatever they are doing is at least making the situation better. However, as noted earlier, this belief is contradicted by mounting research evidence that SATs can impede the quality of judgment.⁴⁷

Why don't analytic tradecraft methodologists assign nonzero probabilities to the possibilities that a proposed method has no appreciable benefit or a negative benefit? If, for instance, an analyst with an educational background in Russian studies were asked to design the next generation technology for encryption he would probably refuse the task. If given no choice, he would surely be aware that his best efforts probably would fail. However, if the same analyst moved into a tradecraft development position, it is likely that he would believe that he could think his way through the problem space and make valuable contributions. After all, he was an analyst and that would seem to confer a

nontrivial degree of expertise. I propose that this fictional analyst's relatively optimistic attitude in the second case owes greatly to the bias blind spot discussed earlier: the analyst over-estimates the degree to which his prior experience prepares him for the task of developing and testing analytic methods and, additionally, he fails to detect that he is over-estimating his pertinent knowledge.

Most people are not cryptographers and know their limits if asked to make contributions in that area. In contrast, everyone has experience of thinking and reasoning, judging, and deciding—and analysts have done so for a living, serving powerful clients. Accordingly, it may be extremely difficult for them to detect the limits of their knowledge. Those promoted to positions requiring oversight of analytic methods are likely to fall prey to what I have called the *goodness heuristic*, which specifies the following conditional rule: *if, upon mental inspection, an idea seems good, then act on it as if it were good because it probably is good*.⁴⁸ The goodness heuristic is a logical extension of Kahneman's WYSIATI principle (WYSIATI—pronounced *whiz-E-at-E*—stands for “what you see is all there is”).⁴⁹ When tradecraft developers try to think of what would help analysts they are likely to use a “positive-test strategy”,⁵⁰ recruiting ideas that form a positive match to the query, “what helps?” In WYSIATI terms, the positive attributes associated with one's idea will be “seen clearly”, but the drawbacks may remain imperceptible. The goodness heuristic takes this one step further by supporting the inference that if an idea “looks” or “sounds” good, then it probably *is* good and one should act upon it as if it is good.

However, for better or worse, reality does not bend to our misperceptions. When it comes to the promotion of analytic integrity, intelligence communities remain mired in Rumsfeldian unknown knowns, believing to be true that which is not true. Unfortunately, this had led to overly optimistic assessments of methodological success within intelligence communities and the misperception that they have ensured analytic rigor through their prescientific efforts.

Bibliography

Bessarabova, Elena, Cameron W. Piercy, Shawn King, Cindy Vincent, and Norah E. Dunbar. et al. "Mitigating Bias Blind Spot Via a Serious Video Game." *Computers in Human Behavior* 62 (2016): 452-66.

Budescu, David V., Han-Hui Por, Stephen B. Broomell, and Michael Smithson. "The Interpretation of IPCC Probabilistic Statements Around the World." *Nature Climate Change* 4 (2014): 508–512.

Chang, Welton, Elissabeth Berdini, David R. Mandel, and Philip E. Tetlock. "Restructuring Structured Analytic Techniques in Intelligence." *Intelligence and National Security* 33, no. 3 (2018): 337-56.

Dhami, Mandeep K., Ian K. Belton, and David R. Mandel. "The 'Analysis of Competing Hypotheses' in Intelligence Analysis." *Applied Cognitive Psychology* 33, no. 6 (2019):1080-90.

Dhami, Mandeep K. and David R. Mandel. "Words or Numbers? Communicating Probability in Intelligence Analysis." *American Psychologist* (2020). Advance online publication. <http://dx.doi.org/10.1037/amp0000637>

Dhami, Mandeep K., David R. Mandel, Barbara A. Mellers, and Philip E. Tetlock. "Improving Intelligence Analysis with Decision Science." *Perspectives on Psychological Science* 106, no. 6 (2015): 753–57.

Edwards, Ward. "Conservatism in Human Information Processing," in *Formal Representation of Human Judgment*, edited by Benjamin Kleinmuntz, 17–52. New York: Wiley, 1968.

Friedman, Jeffrey A. *War and Chance: Assessing Uncertainty in International Politics*. New York: Oxford University Press, 2019.

Friedman, Jeffrey A. and Richard Zeckhauser. "Assessing Uncertainty in Intelligence." *Intelligence and National Security* 27, no. 6 (2012): 824-47.

Gentry, John A. "'Truth' as a Tool of the Politicization of Intelligence." *International Journal of Intelligence and CounterIntelligence* 32, no. 2 (2019): 217-47.

Harari, Yuval Noah. *Sapiens: A Brief History of Humankind*. New York: Harper, 2014.

Heuer, Richards J. Jr. *Psychology of Intelligence Analysis*. Washington, DC: Center for the Study of Intelligence, 1999.

Heuer, Richards J. Jr. and Randolph H. Pherson. *Structured Analytic Techniques for Intelligence Analysis*. Washington, DC: CQ Press, 2014.

Ho, Emily, David V. Budescu, Mandeep K. Dhami, and David R. Mandel. "Improving the Communication of Uncertainty in Climate Science and Intelligence Analysis." *Behavioral Science & Policy* 1, no. 2 (2015): 43-55.

Hood, Christopher. *The Blame Game: Spin, Bureaucracy, and Self-Preservation in Government*. Princeton, NJ: Princeton University Press, 2011.

Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.

Karvetski, Christopher W. and David R. Mandel. "Coherence of Probability Judgments from Uncertain Evidence: Does ACH Help?" *Judgment and Decision Making* 15, no. 6 (2020): 939-58.

Karvetski, Christopher W., David R. Mandel, and Daniel Irwin. "Improving Probability Judgment in Intelligence Analysis: From Structured Analysis to Statistical Aggregation." *Risk Analysis* 40, no. 5 (2020): 1040-57.

Kent, Sherman. "The Need for an Intelligence Literature." *Studies in Intelligence*, 1, no. 1 (1955): 1-11.

Kent, Sherman. "Words of Estimative Probability." *Studies in Intelligence* 8, no. 4 (1964): 49-65.

Klayman, Joshua, and Young-won Ha. "Hypothesis Testing in Rule Discovery: Strategy, Structure, and Content." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15, no. 4 (1989): 596-604.

Kunda, Ziva. "The Case for Motivated Reasoning." *Psychological Bulletin* 108, no. 3 (1990): 480-98.

Lichtenstein, Sarah, Baruch Fischhoff, and Lawrence D. Phillips, Lawrence. "Calibration of Probabilities: The State of the Art to 1980," in *Judgment Under Uncertainty: Heurist-*

ics and Biases, edited by Daniel Kahneman, Paul Slovic, and Amos Tversky, 306-34. Cambridge: Cambridge University Press, 1982.

Mandel, David R. "Can Decision Science Improve Intelligence Analysis?" in *Researching National Security Intelligence: Multidisciplinary Approaches*, edited by Stephen M. Coulthart, Michael Landon-Murray, and Damien Van Puyvelde, 117-40. Washington, DC: Georgetown University Press, 2019.

Mandel, David R. "The Occasional Maverick of Analytic Tradecraft." *Intelligence and National Security* 35, no. 3 (2020): 438-43.

Mandel, David R. and Alan Barnes. "Accuracy of Forecasts in Strategic Intelligence." *Proceedings of the National Academy of Sciences* 111, no.30 (2014): 10984-89.

Mandel, David R. and Alan Barnes. "Geopolitical Forecasting Skill in Strategic Intelligence." *Journal of Behavioral Decision Making* 31, no. 1 (2018): 127-37.

Mandel, David R., Tonya L. Hendriks, and Daniel Irwin. *Policy for Promoting Analytic Rigor in Intelligence: Professionals' Views and Their Psychological Correlates*. (2021, January 20). Retrieved from psyarxiv.com/terux.

Mandel, David R. and Daniel Irwin. "Facilitating Sender-Receiver Agreement in Communicated Probabilities: Is It Best to Use Words, Numbers or Both?" *Judgment and Decision Making* 16, no. 2 (2021): 363-93.

Mandel, David R. And Daniel Irwin. "On Measuring Agreement with Numerically Bounded Linguistic Probability Schemes: A Re-analysis of Data from Wintle, Fraser, Wills, Nicholson, and Fidler (2019)." *PLoS ONE* 16, no. 3 (2021) e0248424. <https://doi.org/10.1371/journal.pone.0248424>

Mandel, David R. and Daniel Irwin. "Uncertainty, Intelligence, and National Security Decisionmaking." *International Journal of Intelligence and CounterIntelligence* (2020) <https://doi.org/10.1080/08850607.2020.1809056>

Mandel, David R., Christopher W. Karvetski, and Mandeep K. Dhimi. "Boosting Intelligence Analysts' Judgment Accuracy: What Works, What Fails?" *Judgment and Decision Making* 13, no. 6 (2018): 607-21.

Mandel, David R. and Philip E. Tetlock. "Debunking the Myth of Value-Neutral Virginity: Toward Truth in Scientific Advertising." *Frontiers in Psychology* 7, article 451 (2016): 1-5. <https://doi.org/10.3389/fpsyg.2016.00451>

Mandel, David R. and Philip E. Tetlock. "Correcting Judgment Correctives in National Security Intelligence." *Frontiers in Psychology* 9, article 2640 (2018): 1-5. <https://doi.org/10.3389/fpsyg.2018.02640>

Marcoci, Alexandru, Mark Burgman, Ariel Kruger, Elizabeth Silver, Marissa McBride et al. "Better Together: Reliable Application of the Post-9/11 and Post-Iraq US Intelligence Tradecraft Standards Requires Collective Analysis." *Frontiers in Psychology*, 9, article 2634 (2019):1-9, <https://doi.org/10.3389/fpsyg.2018.02634>

Marrin, Stephen. "Intelligence Analysis: Structured Methods or Intuition?" *American Intelligence Journal* 25, no. 1 (2007): 7-16.

Marrin, Stephen. "Training and Educating U.S. intelligence Analysts." *International Journal of Intelligence and CounterIntelligence* 22, no. 1 (2009):131-46.

Moore, Don. A. and Paul J. Healy. "The Trouble with Overconfidence." *Psychological Review* 115, no. 2 (2008): 502–17.

Office of the Director of National Intelligence. *Intelligence Community Directive 203: Analytic Standards*. Washington, DC, DNI, 2015. <https://fas.org/irp/dni/icd/icd-203.pdf>

Pool, Robert. *Field Evaluation in the Intelligence and Counterintelligence Context: Workshop Summary*. Washington, DC: The National Academies Press, 2010.

Popper, Karl. *The Logic of Scientific Discovery*. New York: Routledge, 2002. [Originally published in 1935 as *Logik der Forschung* Verlag von Julius Springer, Vienna, Austria]

Pronin, Emily, Daniel Y. Lin, and Lee Ross. "The Bias Blind Spot: Perceptions of Bias in Self Versus Others." *Personality and Social Psychology Bulletin* 28, no. 3 (2002): 369-81.

Rieber, Steven & Neil Thomason. "Creation of a National Institute for Analytic Methods." *Studies in Intelligence* 49, no. 4 (2005): 71–77.

Scopelliti, Irene, Carey K. Morewedge, Erin McCormick, H. Lauren Min, Sophie Lebrecht, and Karim S. Kassam. (2015). "Bias Blind Spot: Structure, Measurement, and Consequences." *Management Science* 61, no. 10 (2015): 2468-86.

Tetlock, Philip E. "Social Functionalist Frameworks for Judgment and Choice: Intuitive Politicians, Theologians, and Prosecutors." *Psychological Review* 109, no. 3 (2002) 451-71.

Tetlock, Philip E. and Barbara A. Mellers. "Intelligent Management of Intelligence Agencies: Beyond Accountability Ping-Pong." *American Psychologist* 66, (2011): 542–54.

Timms, Mark A. C., David R. Mandel, and Jonathan D. Nelson. "Applying Information Theory to Validate Commanders' Critical Information Requirements" in *Handbook of Military and Defence Operation Research*, edited by Natalie M. Scala and James P.

Howard, 331-44. Boca Raton, FL: CRC Press, 2020.

Tversky, Amos, and Daniel Kahneman. "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185 (1974): 1124-31.

Tversky, Amos and Daniel Kahneman. "Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment." *Psychological Review* 90, no. 4 (1983): 293-315.

von Hippel, William, and Robert Trivers. "The Evolution and Psychology of Self-Deception." *Behavioral and Brain Sciences* 34, no. 1 (2011): 1-16.

West, Richard F., Russell J. Meserve, and Keith E. Stanovich. "Cognitive Sophistication Does Not Attenuate the Bias Blind Spot." *Journal of Personality and Social Psychology* 103, no. 3 (2012): 506-519.

Whitesmith, Martha. "The Efficacy of ACH in Mitigating Serial Position Effects and Confirmation Bias in an Intelligence Analysis Scenario." *Intelligence and National Security* 34, no. 2 (2019): 225-42.

Wintle, Bonnie C., Hannah Fraser, Ben C. Wills, Ann E. Nicholson, and Fiona Fidler. "Verbal Probabilities: *Very Likely* to Be *Somewhat* More Confusing Than Numbers." *PLoS ONE* 14, no. 4 (2019) e0213522. <http://dx.doi.org/10.1371/journal.pone.0213522>

Notes

- ¹ Sherman Kent, "The Need for an Intelligence Literature," *Studies in Intelligence*, 1, no. 1 (1955): 1-11.
- ² Sherman Kent, "Words of Estimative Probability," *Studies in Intelligence* 8, no. 4 (1964): 49-65.
- ³ Friedman, Jeffrey A. Friedman and Richard Zeckhauser, "Assessing Uncertainty in Intelligence," *Intelligence and National Security* 27, no. 6 (2012): 824-47; David R. Mandel and Daniel Irwin, "Uncertainty, Intelligence, and National Security Decisionmaking," *International Journal of Intelligence and CounterIntelligence* (2020).
- ⁴ Richards J. Heuer Jr., *Psychology of Intelligence Analysis* (Washington, DC: Center for the Study of Intelligence, 1999).
- ⁵ Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011).
- ⁶ Amos Tversky and Daniel Kahneman, "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185 (1974): 1124-31.
- ⁷ Tversky and Kahneman, "Judgment Under Uncertainty."
- ⁸ Ward Edwards "Conservatism in Human Information Processing," in *Formal Representation of Human Judgment*, edited by Benjamin Kleinmuntz, 17-52 (New York: Wiley, 1968).
- ⁹ Amos Tversky and Daniel Kahneman, "Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment," *Psychological Review* 90, no. 4 (1983): 293-315.
- ¹⁰ Emily Pronin, Daniel Y. Lin, and Lee Ross, "The Bias Blind Spot: Perceptions of Bias in Self Versus Others," *Personality and Social Psychology Bulletin* 28, no. 3 (2002): 369-8; Irene Scopelliti, Carey K. Morewedge, Erin McCormick et al, "Bias Blind Spot: Structure, Measurement, and Consequences," *Management Science* 61, no. 10 (2015): 2468-86.
- ¹¹ Richard F. West, Russell J. Meserve, and Keith E. Stanovich, "Cognitive Sophistication Does Not Attenuate the Bias Blind Spot," *Journal of Personality and Social Psychology* 103, no. 3 (2012): 506-519.
- ¹² Elena Bessarabova, Cameron W. Piercy, Shawn King et al, "Mitigating Bias Blind Spot Via a Serious Video Game," *Computers in Human Behavior* 62 (2016): 452-66.
- ¹³ Errol Morris (Director), *The Unknown Known: What You Didn't Know You Didn't Know* (2013).
- ¹⁴ William von Hippel and Robert Trivers, "The Evolution and Psychology of Self-Deception," *Behavioral and Brain Sciences* 34, no. 1 (2011): 1-16.
- ¹⁵ Ziva Kunda, "The Case for Motivated Reasoning," *Psychological Bulletin* 108, no. 3 (1990): 480-98.
- ¹⁶ Philip E. Tetlock, "Social Functionalist Frameworks for Judgment and Choice: Intuitive Politicians, Theologians, and Prosecutors," *Psychological Review* 109, no. 3 (2002) 451-71.
- ¹⁷ Sarah Lichtenstein, Baruch Fischhoff, and Lawrence D. Phillips, Lawrence, "Calibration of Probabilities: The State of the Art to 1980," in *Judgment Under Uncertainty: Heuristics and Biases*, ed. Daniel Kahneman, Paul Slovic, and Amos Tversky (Cambridge: Cambridge University Press, 1982), 306-34; Don A. Moore and Paul J. Healy, "The Trouble with Overconfidence," *Psychological Review* 115, no. 2 (2008): 502-17.
- ¹⁸ David R. Mandel and Alan Barnes, "Accuracy of Forecasts in Strategic Intelligence," *Proceedings of the National Academy of Sciences* 111, no.30 (2014): 10984-89; Mandel and Barnes, "Geopolitical Forecasting Skill in Strategic Intelligence," *Journal of Behavioral Decision Making* 31, no. 1 (2018): 127-37.
- ¹⁹ Christopher Hood, *The Blame Game: Spin, Bureaucracy, and Self-Preservation in Government* (Princeton, NJ: Princeton University Press, 2011); Philip E. Tetlock and Barbara A. Mellers, "Intelligent Management of Intelligence Agencies: Beyond Accountability Ping-Pong," *American Psychologist* 66, (2011): 542-54.
- ²⁰ John A. Gentry, "'Truth' as a Tool of the Politicization of Intelligence," *International Journal of Intelligence and CounterIntelligence* 32, no. 2 (2019): 217-47.
- ²¹ Office of the Director of National Intelligence. *Intelligence Community Directive 203: Analytic Standards* (Washington, DC, DNI, 2015), 3-4.

- ²² David R. Mandel, Tonya L. Hendriks, and Daniel Irwin, *Policy for Promoting Analytic Rigor in Intelligence: Professionals' Views and Their Psychological Correlates*. (January 20, 2021).
- ²³ Alexandru Marcoci, Mark Burgman, Ariel Kruger, Elizabeth Silver, Marissa McBride et al, "Better Together: Reliable Application of the Post-9/11 and Post-Iraq US Intelligence Tradecraft Standards Requires Collective Analysis," *Frontiers in Psychology*, 9, article 2634 (2019): 1-9.
- ²⁴ Mandel and Barnes, "Accuracy of Forecasts"; Mandel and Barnes, "Geopolitical Forecasting Skill."
- ²⁵ Gentry, "'Truth' as a Tool."
- ²⁶ Heuer Jr., *Psychology*.
- ²⁷ Christopher W. Karvetski, David R. Mandel, and Daniel Irwin, "Improving Probability Judgment in Intelligence Analysis: From Structured Analysis to Statistical Aggregation," *Risk Analysis* 40, no. 5 (2020): 1040-57.
- ²⁸ David R. Mandel, "Can Decision Science Improve Intelligence Analysis?" in *Researching National Security Intelligence: Multidisciplinary Approaches*, eds. Stephen M. Coulthart, Michael Landon-Murray, and Damien Van Puyvelde (Washington, DC: Georgetown University Press, 2019), 117-40.
- ²⁹ Karl Popper, *The Logic of Scientific Discovery* (New York: Routledge, 2002, originally 1935).
- ³⁰ Martha Whitesmith, "The Efficacy of ACH in Mitigating Serial Position Effects and Confirmation Bias in an Intelligence Analysis Scenario," *Intelligence and National Security* 34, no. 2 (2019): 225-42.
- ³¹ David R. Mandel, Christopher W. Karvetski, and Mandeep K. Dhami, "Boosting Intelligence Analysts' Judgment Accuracy: What Works, What Fails?" *Judgment and Decision Making* 13, no. 6 (2018): 607-21.
- ³² Mandeep K. Dhami, Ian K. Belton, and David R. Mandel, "The 'Analysis of Competing Hypotheses' in Intelligence Analysis," *Applied Cognitive Psychology* 33, no. 6 (2019): 1080-90.
- ³³ Karvetski, Mandel, and Irwin, "Improving Probability Judgment."
- ³⁴ Richards J. Heuer, Jr. and Randolph H. Pherson, *Structured Analytic Techniques for Intelligence Analysis* (Washington, DC: CQ Press, 2014).
- ³⁵ Mark A. C. Timms, David R. Mandel, and Jonathan D. Nelson, "Applying Information Theory to Validate Commanders' Critical Information Requirements," in *Handbook of Military and Defence Operation Research*, eds. Natalie M. Scala and James P. Howard (Boca Raton, FL: CRC Press), 2020: 331-44.
- ³⁶ For example: Welton Chang, Elissabeth Berdini, David R. Mandel, and Philip E. Tetlock, "Restructuring Structured Analytic Techniques in Intelligence," *Intelligence and National Security* 33, no. 3 (2018): 337-56; Mandeep K. Dhami, David R. Mandel, Barbara A. Mellers, and Philip E. Tetlock, "Improving Intelligence Analysis with Decision Science," *Perspectives on Psychological Science* 106, no. 6 (2015): 753-57; David R. Mandel and Philip E. Tetlock, "Correcting Judgment Correctives in National Security Intelligence," *Frontiers in Psychology* 9, article 2640 (2018): 1-5; Stephen Marrin, "Intelligence Analysis: Structured Methods or Intuition?" *American Intelligence Journal* 25, no. 1 (2007): 7-16; Marrin, "Training and Educating U.S. intelligence Analysts," *International Journal of Intelligence and CounterIntelligence* 22, no. 1 (2009): 131-46; Robert Pool, *Field Evaluation in the Intelligence and Counterintelligence Context: Workshop Summary* (Washington, DC: The National Academies Press, 2010).
- ³⁷ David R. Mandel, "The Occasional Maverick of Analytic Tradecraft," *Intelligence and National Security* 35, no. 3 (2020): 438-43.
- ³⁸ Steven Rieber & Neil Thomason, "Creation of a National Institute for Analytic Methods," *Studies in Intelligence* 49, no. 4 (2005): 71-77.
- ³⁹ Mandel, "Can Decision Science."
- ⁴⁰ David R. Mandel and Philip E. Tetlock, "Debunking the Myth of Value-Neutral Virginity: Toward Truth in Scientific Advertising," *Frontiers in Psychology* 7, article 451 (2016): 1-5.

⁴¹ Mandeep K. Dhami and David R. Mandel, “Words or Numbers? Communicating Probability in Intelligence Analysis,” *American Psychologist* (2020); Jeffrey A. Friedman, *War and Chance: Assessing Uncertainty in International Politics* (New York: Oxford University Press, 2019).

⁴² David V. Budescu, Han-Hui Por, Stephen B. Broomell, and Michael Smithson, “The Interpretation of IPCC Probabilistic Statements Around the World,” *Nature Climate Change* 4 (2014): 508–512.

⁴³ Emily Ho, David V. Budescu, Mandeep K. Dhami, and David R. Mandel, “Improving the Communication of Uncertainty in Climate Science and Intelligence Analysis,” *Behavioral Science & Policy* 1, no. 2 (2015): 43-55; David R. Mandel and Daniel Irwin, “Facilitating Sender-Receiver Agreement in Communicated Probabilities: Is It Best to Use Words, Numbers or Both?,” *Judgment and Decision Making* 16, no. 2 (2021); Mandel, David R. And Daniel Irwin, “On Measuring Agreement with Numerically Bounded Linguistic Probability Schemes: A Re-analysis of Data from Wintle, Fraser, Wills, Nicholson, and Fidler (2019)” *PLoS ONE* 16, no. 3 (2021); Bonnie C. Wintle, Hannah Fraser, Ben C. Wills, Ann E. Nicholson, and Fiona Fidler, “Verbal Probabilities: *Very Likely* to Be *Somewhat* More Confusing Than Numbers,” *PLoS ONE* 14, no. 4 (2019).

⁴⁴ David R. Mandel and Daniel Irwin, “Uncertainty, Intelligence, and National Security Decisionmaking,” *International Journal of Intelligence and CounterIntelligence* (2020)

⁴⁵ Yuval Noah Harari, *Sapiens: A Brief History of Humankind* (New York: Harper, 2014).

⁴⁶ Marcoci et al, “Better Together.”

⁴⁷ For example, Mandel et al, “Boosting”; Karvetski et al, “Improving”; Christopher W. Karvetski, and David R. Mandel, “Coherence of Probability Judgments from Uncertain Evidence: Does ACH Help?” *Judgment and Decision Making* 15, no. 6 (2020): 939-58; Timms et al, “Applying.”

⁴⁸ Mandel, “Can Decision Science.”

⁴⁹ Kahneman, *Thinking*.

⁵⁰ Joshua Klayman, and Young-won Ha, “Hypothesis Testing in Rule Discovery: Strategy, Structure, and Content,” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15, no. 4 (1989): 596-604.