

To appear in *Intelligence and National Security*

Policy Neutrality and Uncertainty: An Info-Gap Perspective^o

YAKOV BEN-HAIM*

Abstract

Reducing uncertainty is a central goal of intelligence analysis. 'Reducing uncertainty' can mean (1) Reduce ignorance or ambiguity or potential for surprise in describing situations or intentions, or (2) Reduce adverse impacts of ignorance, ambiguity or surprise on decision outcomes.

We make two claims. First, the second meaning needs greater attention in intelligence analysis. Uncertainty itself isn't pernicious, but adverse impact of surprise is. Some policy options are less vulnerable to uncertainty than others. These less vulnerable (i.e. more robust) options can tolerate more uncertainty. Analysts should identify policy options that are robust to uncertainty.

Second, reducing the impact of uncertainty requires awareness of policymakers' goals. This needn't conflict with analysts' policy neutrality. Tension between neutrality and involvement arises in economics, engineering, and medicine. The method of info-gap robust-satisficing supports decision making under uncertainty in these and other disciplines. Implications for intelligence analysis are explored in this paper. We discuss the assessment of Iraqi WMD capability in 2002.

Introduction

Reducing uncertainty about situations that are relevant to policymakers is a major focus of intelligence analysis. There are two distinct meanings of 'reducing uncertainty':

- Reduce ignorance or ambiguity or potential for surprise in describing situations or protagonists' plans and intentions.
- Reduce vulnerability to ignorance, ambiguity or surprise. Reduce the potential for adverse impact of uncertainty on the outcome of a decision.

This dichotomy is a simplification of the intelligence analyst's mandate. Thomas Fingar explains that the "primary purpose ... is to reduce uncertainty [first meaning], identify risks

^o \papers\policy-neutral\pol-neut004.docx 15.11.2015.

* Corresponding author. yakov@technion.ac.il.

and opportunities, and, by doing so, deepen understanding so that those with policy-making responsibilities will make 'better' decisions."¹

Nonetheless, the distinction between these two challenges is important because it emphasizes that it is not uncertainty *per se* that is pernicious, but rather the possibility of adverse impact of surprise. For example, uncertainty about irrelevant situations needn't be reduced because it isn't pernicious. More importantly, some policy options are less vulnerable to uncertainty than others. These less vulnerable (equivalently, more robust) options can tolerate more uncertainty. The **first main claim** of this paper is that the second meaning needs greater attention in intelligence analysis, and that intelligence analysts should assist the policymaker to identify policy options that are robust to uncertainty. It is correct that if one drastically reduces uncertainty then it is very plausible that one has also drastically reduced the vulnerability to uncertainty. In the extreme, if uncertainty is eliminated, then vulnerability to surprise is also eliminated. However, intelligence analysis is usually done under constraints on time and other resources and major uncertainties will usually remain. The first main claim is that uncertainty should be managed by assessing and enhancing the robustness-benefit that accrues to the policy choice.

Reducing the impact of uncertainty requires awareness, by the analyst, of the policymaker's goals (avoiding war, promoting aid, balancing foreign powers, etc). Acquisition and possession of this awareness could tend to undermine the analyst's policy neutrality. Similar conflict between policy neutrality and policy involvement arises in advice by economists to policymakers, engineers to clients, and physicians to patients. Info-gap decision theory has proven useful in these and other disciplines for professional support of decision making under uncertainty.² The **second main claim** of this paper is that info-gap robust-satisficing provides a policy-neutral methodology for supporting policymakers' decisions. Info-gap theory is explained and its application to intelligence analysis is explored.

Attitudes to the Management of Uncertainty in Intelligence Analysis

The modern U.S. tradition in intelligence analysis, originating with Sherman Kent following WW II, is characterized by two linked conceptions. First, the analyst's task is to reduce

¹ Thomas Fingar, 2011, *Reducing Uncertainty: Intelligence Analysis and National Security*, Stanford Security Studies, p.25 See also pp.10, 50, 72.

² The main reference for info-gap theory is: Yakov Ben-Haim, *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edition, (London: Academic Press 2006). Citations to papers by many scholars from diverse disciplines can be found on info-gap.com.

uncertainty (in the first meaning) and to clearly present relevant data and knowledge to the policymaker. Second, the analyst must maintain strong policy neutrality: having no involvement in the policy implications of the intelligence product.

For instance, Fingar writes that "The purpose of intelligence since time immemorial has been to reduce uncertainty about the aspirations, intentions, capabilities, and actions of adversaries." He stresses later on that "I'm an analyst; ... I don't do policy."³ Likewise, James Bruce and Roger George write: "A key attribute of intelligence analysis is maintaining policy relevance while assiduously avoiding policy advocacy. This heritage of policy neutrality traces directly to Sherman Kent and is nearly hardwired in the culture of analysis."⁴ Similarly, Arthur Hulnick writes that "When the decision makers ask intelligence officers how they should respond to the crisis, typically intelligence officers decline to provide advice, thus staying clear of the policy process."⁵ Likewise James Wirtz explains that a normative theory "associated with the work of Sherman Kent, focuses on ensuring the independence of intelligence analysts when it comes to providing information to policymakers ... [and] identifies the importance of political and policy detachment".⁶ Similarly, Glenn Hastedt asserts that in "Conventional accounts ... [i]ntelligence is discussed as a value-neutral activity whose defining purpose is to help policymakers better understand a situation by furnishing them with analyzed information."⁷ Jeffrey Friedman and Richard Zeckhauser claim that "Much of the theory and practice of estimative intelligence aims to eliminate or reduce uncertainty".⁸ Stephen Marrin argues that "According to the standard model, intelligence organizations were created separate from decision-makers; to select facts and provide an objective assessment of them as an effort to get at one true answer which is then conveyed to consumers to be incorporated into their policy deliberations."⁹

There has, of course, been extensive critical discussion of the Kent tradition of uncertainty-reduction (first meaning) and strict policy neutrality of intelligence analysts.

³ Fingar, 2011, *op. cit.*, pp.6, 45-46.

⁴ James B. Bruce and Roger Z. George, Intelligence analysis: What is it - And what does it take?, in Roger Z. George and James B. Bruce, eds., *Analyzing Intelligence: National Security Practitioners' Perspectives*, 2nd edition, (Washington, DC: Georgetown University Press 2014), p.12.

⁵ Arthur S. Hulnick, What's wrong with the intelligence cycle?, in Loch Johnson, ed., *Essentials of Strategic Intelligence*, (Santa Barbara, CA: Praeger 2015) p.73.

⁶ James J. Wirtz, The intelligence-policy nexus, in Johnson, 2015, *op.cit.*, p.140.

⁷ Glenn Hastedt, Washington politics, homeland security, and the struggle against global terrorism, in Johnson, 2015, *op.cit.*, p.388. See also: Glenn Hastedt, The politics of intelligence and the politicization of intelligence: The American experience, *Intelligence and National Security*, 28(1) (2013) p.6.

⁸ Jeffrey A. Friedman and Richard Zeckhauser, Assessing uncertainty in intelligence, *Intelligence and National Security*, 27(6) (2012) p.824.

⁹ Stephen Marrin, Rethinking analytic politicization, *Intelligence and National Security*, 28(1) (2013) p.40.

Marrin presents a broad foundational critique of the "standard model" based on the tradition of Kent, and identifies "three significant problems". The first problem "is that analytic objectivity in an absolute sense does not exist. ... [I]ntelligence analysis ... contains a relatively high level of ambiguity in terms of the meaning of data and this causes relatively high levels of uncertainty and error in analytic conclusions and estimates." The second problem is that "the accuracy of any particular intelligence analysis may depend as much on the conceptual framework employed in a deductive way to derive meaning as from the inductive accumulation and assessment of the data itself. In this environment, multiple legitimate interpretations of the same situation are possible." The third problem is that there is "little evidence" to indicate that decision makers will change their opinions and policies when "confronted with assessments from experts that differ from their own".¹⁰

Peter Gill claims that "Intelligence seeks to improve knowledge by reducing uncertainty but is only ever partially successful."¹¹ Friedman and Zeckhauser claim that trying "to eliminate or reduce uncertainty ... is often impossible or infeasible." They advocate "that the goal of estimative intelligence should be to assess uncertainty" using probability distributions. They stress that "Using probability distributions removes the inclination to believe that alternative views are necessarily at odds with each other."¹² Similarly, Daniel Javorsek and John Schwitz write: "Only by explicitly identifying the probabilities of possible outcomes associated with a measure of significant consequence can we evaluate and improve our intelligence process while signaling the probability of catastrophic events."¹³

Not surprisingly, there is dispute among the critics of the Kent tradition. For instance not all authors agree that quantitative probability assessments are realistic. Javorsek and Schwitz continue their earlier discussion and explain that "the realities of the intelligence process preclude the specification of multiple future states with attendant probabilities".¹⁴ James Steinberg writes: "I am somewhat skeptical of what I believe is a false sense of concreteness implied in assigning numerical probabilities to individual events, particularly contingent outcomes that depend on choices others have yet to make."¹⁵ Nonetheless, Jack Davis emphasizes (without using probability) the importance of identifying a range of

¹⁰ Marrin, 2013, *op.cit.*, pp.41-43, 45.

¹¹ Peter Gill, 9/11 and Iraqi WMD, Johnson (2015) *op. cit.*, p.163.

¹² Friedman and Zeckhauser, 2012, *op.cit.*, pp.824, 832-833.

¹³ Daniel Javorsek II and John G. Schwitz, Probing uncertainty, complexity, and human agency in intelligence, *Intelligence and National Security*, 29(5) (2014) p.645.

¹⁴ *ibid*, p.648.

¹⁵ James B. Steinberg, The policymaker's perspective: Transparency and partnership, in George and Bruce (2014) *op.cit.*, p.96.

contingencies: "As a rule, the more important the intelligence issue and the greater the uncertainty and information gaps, the greater need for incorporating alternative explanations and projections into the text of an assessment."¹⁶

These and other similar critiques emphasize the limitation of presenting to the policymaker a single, perhaps even monolithic, portrayal of the situation.

Criticism of the Kent tradition of strict separation of analyst from policymaker is not new. Wirtz explains that a different operational framework, associated with "reforms instituted in the mid-1980s by then director of Central Intelligence Robert M. Gates, focuses on providing 'actionable' intelligence, information of immediate and direct use to policymakers." Wirtz quotes Gates as explaining "that the Intelligence community has to be right next to the policymaker, that [the analyst] has to be at his elbow - that he has to understand what is on his mind. He has to understand what his future concerns are. He has to understand what his agenda is. He has to understand some of the initiatives that he is thinking about taking."¹⁷ Davis elaborates that "if an analyst is not close enough to the process to feel the political pressures affecting policymaking, he or she probably is not close enough to produce professionally crafted deliverables that provide distinctive value added."¹⁸

Gates' challenge to Kent's strict separation between policymaker and analyst raises the issue of politicization, as opposed to policy neutrality, of intelligence analysis.

Hastedt claims that "politicizing an issue is not by definition corrupting it. ... It [politicization] says nothing about the agenda of those employing it."¹⁹ The politicization of analysis is pernicious if it entails distortion of the intelligence product to support or supplant specific policies. This can happen, for example, if policymakers "pose biased questions to the intelligence community to guarantee that analysis favorable to their positions will emerge".²⁰ Alternatively, if policymakers don't clarify their policy goals, then the intelligence community, in order to strengthen its relevance, may "establish intelligence priorities and policy requirements according to its own perception ... [and] will likely be accused of participating in policy by taking an undesirable active role in the sphere of politics."²¹ On the other hand, Woodard explains that "By explicitly stating a policy position, analysts would necessarily and more explicitly state their biases ... [and would add] more clarity to the

¹⁶ Jack Davis, Why bad things happen to good analysts, in George and Bruce (2014) *op.cit.*, p.132.

¹⁷ Wirtz (2015) *op. cit.*, pp.140, 143.

¹⁸ Jack Davis, Intelligence analysts and policy makers, in Johnson (2015) *op. cit.*, p.124.

¹⁹ Hastedt (2013) *op. cit.*, p.10.

²⁰ Wirtz (2015) *op. cit.*, p.144.

²¹ Marcos Degaut, Spies and policymakers: Intelligence in the information age, *Intelligence and National Security*, (2015) <http://dx.doi.org/10.1080/02684527.2015.1017931>, p.18.

situation." Having discussed evidence-based medicine, Nathan Woodard continues: "Evidence-based intelligence advocacy is a specific type of policy advocacy that seeks to be fair with all of the evidence at hand, and persuasively present the policy argument that is best supported by the evidence."²²

The practical meaning of policy neutrality, and whether neutrality prohibits explicit consideration of policy issues by analysts, and even whether strict abstinence from policy advocacy is possible, are all issues in scholarly dispute. An issue that has not received adequate attention in the Kent-Gates debate is that the importance of intelligence uncertainty depends on the impact of surprise on policy outcomes. Intelligence uncertainty is important if, and only if, corresponding surprises would seriously and adversely impact policy outcomes. This paper contributes a response to this challenge.

The first claim of this paper is that intelligence analysts need to manage the potential for adverse impact of uncertainty on the outcomes of policy choices. It is not uncertainty itself that is pernicious; bad outcomes are pernicious. This outcome-oriented analysis of uncertainty requires analyst familiarity with policymakers' goals. The second claim of this paper is that info-gap robust-satisficing provides a methodology for supporting policymakers' decisions that enables explicit inclusion of goals while avoiding analyst advocacy of alternative goals. Robust-satisficing facilitates deliberation and decision by policymakers; the analyst facilitates but does not advocate. These claims will be elaborated in subsequent sections.

Brenda Canton presents a detailed and explicit methodology for "the active management of uncertainty - what analysts do in response to what they don't know on key issues of policy interest."²³ A brief discussion of Canton's contribution will lead naturally to discussion of info-gap theory and the method of robust-satisficing.

The most salient aspect of Canton's methodology, for our discussion, is the 2nd of 4 components: "*Building Counterstrategies*. Analysts search for the knowable. If they cannot reduce information gaps, they seek better understanding of the nature and implication of what may be missing. Counterstrategies - action plans aimed at narrowing or eliminating information gaps - help analysts gather distinct-value-added information for policymakers." One might have the impression, from this passage, that Canton adheres to the traditional first meaning of the phrase 'reducing uncertainty'. However this is later clarified: "Analysts are accountable for their efforts to address information gaps identified in the drilling phase. This

²² Nathan Woodard, Tasting the forbidden fruit: Unlocking the potential of positive politicization, *Intelligence and National Security*, 28(1) (2013) pp.101-102.

²³ Belinda Canton, The active management of uncertainty, *International Journal of Intelligence and CounterIntelligence*, 21(3) (2008) p.487.

is the best way analysts can reduce substantive uncertainty and potential vulnerability, both for themselves and policy clients. In cases where an analyst cannot eliminate a gap - on a true mystery, for example - that officer is responsible for understanding and managing the consequences of missing information." Caution is required in interpreting this passage, but it seems to encompass both meanings of 'reducing uncertainty'. The first meaning is entailed in "reduce substantive uncertainty". However, the second meaning is suggested by "reduce ... potential vulnerability" and that the "officer is responsible for understanding and managing the consequences of missing information." If uncertainty is rampant, then the analyst does not present to the policymaker a single refined coherent picture, but must consider and manage the consequences of the information gaps: that which is unknown and relevant to the policy outcome.²⁴

What is an Info-Gap?

Information-gap decision theory originated in engineering design and safety analysis²⁵ and subsequently found application in biological conservation, economics, medicine, project management, national security and other areas.²⁶ What is an info-gap?

An info-gap is the disparity between what one *does know* and what one *needs to know* in order to make a responsible and reliable decision. An info-gap is not simply a gap in one's knowledge; it is a gap with significant consequences for the outcome of a decision. Two elements underlie the concept of an info-gap: uncertainty and consequence.

First is the element of **uncertainty**: ignorance or ambiguity or the potential for surprise. Frank Knight's concept of 'true uncertainty' arises from innovation and initiative of entrepreneurs. For Knightian uncertainty "there is no objective measure of the probability" because there is little or no experience with new innovations or initiatives from which frequencies or likelihoods can be learned.²⁷ Knightian uncertainty arises from the unbounded potential for future innovation, or simply from ignorance of the vastly complex world.

²⁴ Canton (2008) *op.cit.*, pp.490-491, 494.

²⁵ Yakov Ben-Haim, *Robust Reliability in the Mechanical Sciences* (Berlin: Springer-Verlag 1996).

²⁶ Ben-Haim 2006 *op.cit.* Yakov Ben-Haim, *Info-Gap Economics: An Operational Introduction*, (London: Palgrave-Macmillan 2010). See also info-gap.com

²⁷ Frank H. Knight, *Risk, Uncertainty and Profit*, (Hart, Schaffner and Marx, 1921, Re-issued by Harper Torchbooks, New York, 1965) pp.46, 120, 231-232. See also Frank H. Knight, *The Economic Organization*, (New York: Harper Torchbooks, 1933, 1951) p.120.

Uncertainty may also arise from deception or denial by an adversary, and Mark Phythian has pointed to the importance of Knightian uncertainty in intelligence analysis.²⁸

G.L.S. Shackle²⁹ and, independently, Karl Popper,³⁰ explained a concept of indeterminism that is related to Knightian uncertainty.³¹ Human behavior depends on what people (or groups) know: if you know it will rain then you'll take an umbrella; if you know the enemy has chemical weapons then you'll take a gas mask. However, what will be invented or discovered tomorrow cannot, by definition, be known today. Hence tomorrow's behavior will have an element of irreducible indeterminism today. Knightian uncertainty and Shackle-Popper indeterminism underlie the concept of an info-gap.

The second element underlying the concept of an info-gap is **consequence**: decisions or actions have consequences for which the policymaker is responsible. An info-gap is an uncertainty that can have significant impact on the outcome of a decision or policy. The policymaker is interested in info-gaps because the impact may be adverse.³² An info-gap is not simply a gap in one's information (or knowledge or understanding). An info-gap is the disparity between what is known and what needs to be known in order to reliably and responsibly make a decision.

In reducing an info-gap one is emphasizing both meanings of 'reducing uncertainty', as we noted earlier in connection with Canton (2008). We now examine what that entails.

What is Robust-Satisficing?

We mentioned earlier that the debate between the schools of Kent and Gates does not adequately address the fact that the importance of intelligence uncertainty depends on the impact of surprise on policy outcomes. We now explain the method of robust-satisficing and how it responds to this need.

Our first main claim is that intelligence analysts should reduce the vulnerability to uncertainty; they should reduce the potential for adverse impact on the policy outcome.

²⁸ Mark Phythian, Policing uncertainty: Intelligence, security and risk, *Intelligence and National Security*, 27(2) (2012) 187-205.

²⁹ G.L.S. Shackle, *Epistemics and Economics: A Critique of Economic Doctrines* (London: Transaction Publishers, 1992, originally published by Cambridge University Press, 1972) pp.3-4, 156, 239, 401-402.

³⁰ Karl Popper, 1982, *The Open Universe: An Argument for Indeterminism*. From the Postscript to *The Logic of Scientific Discovery* (London: Routledge 1982) pp.80-81, 109.

³¹ Yakov Ben-Haim, Peirce, Haack and info-gaps, pp.150-164 in *Susan Haack, A Lady of Distinctions: The Philosopher Responds to Her Critics*, edited by Cornelis de Waal, (New York: Prometheus Books 2007).

³² Info-gap theory also considers potentially propitious impacts of info-gaps - windfalls and favorable surprises - evaluated with the concept of opportuneness from uncertainty (Ben-Haim, 2006). We will focus on adverse impacts and consequently on robustness against uncertainty.

Canton rightly states that "Analysts search for the knowable."³³ Better knowledge can reduce one's vulnerability to an adversary's actions, especially in defensive situations as David Kahn has explained.³⁴ However, constraints on time and other resources preclude knowing everything. But even more fundamentally, one cannot know future innovations, inventions and discoveries. Consequently Knightian uncertainty and Shackle-Popper indeterminism can never be entirely eliminated. One will unavoidably face info-gaps about future threats. Our first claim is operationalized by enhancing the robustness against info-gaps. More precisely, the analyst needs to evaluate any proposed policy in terms of how large an info-gap it can tolerate and still achieve the policymaker's stated goals. The analyst must help the policymaker robustly satisfy the goals for which the policy is formulated. The concept of robustness identifies and manages important intelligence uncertainty.

So what is robust-satisficing?

Herbert Simon introduced the technical meaning of 'satisfice' as "To decide on and pursue a course of action that will satisfy the minimum requirements necessary to achieve a particular goal."³⁵ A policymaker who uses the concept of satisficing is not aiming at the best possible outcome. Instead of maximizing utility or minimizing damage, the policymaker aims to achieve an outcome that is *good enough*. For example, the policymaker tries to assure that loss is not greater than an acceptable level, or gain is no less than a required level. When choosing between alternative policies, the robust-satisficing policymaker will choose the policy that is expected to satisfy the critical requirements over the greatest spectrum of unknown future contingencies. Of course, what constitutes a good enough or an acceptable outcome, as specified by the policymaker, can be anywhere on the spectrum from very modest to highly demanding. Satisficing does not imply lack of ambition or demands. It simply means specifying how good an outcome is required, or how bad an outcome would still be acceptable.

It is the policymaker's task to specify the critical requirements. It is the analyst's task to participate in evaluating each proposed policy with respect to these requirements. (This separation of tasks underlies the analyst's policy neutrality.) Our first claim is that this evaluation is done by identifying, for each proposed policy, the largest info-gap that does not jeopardize achievement of the critical requirements. That is, any proposed policy is evaluated

³³ Canton (2008) *op.cit.*, 490.

³⁴ David Kahn, An historical theory of intelligence, *Intelligence and National Security*, 16(3) (2001) 85.

³⁵ *Oxford English Dictionary*, online edition accessed 15.6.2015. See also Herbert A. Simon, A behavioral model of rational choice, *Quarterly Journal of Economics*, 69 (1955) 174-183. Herbert A. Simon, *Models of Man*, (New York: John Wiley and Sons 1957).

by addressing the **robustness question**: how much can our current data, knowledge and understanding err, and the proposed policy will still satisfy the policymaker's stated outcome requirements? A policy is highly robust against uncertainty if the critical requirements are achieved even though great surprises occur. In contrast, a policy has low robustness, and is vulnerable to uncertainty, if even small info-gaps could result in failure to achieve some or all of the critical requirements. The analyst assists the policymaker in prioritizing the proposed policies according to their robustness for achieving the policymaker's critical requirements. The "bottom line" is that the policymaker chooses the policy that most robustly satisfies the policymaker's requirements, thus assuring the analyst's policy neutrality.

The robustness question is the info-gap analog of the question: Is this uncertainty important? Evaluation of the robustness entails an operational response to that question. Intelligence uncertainty is important if a corresponding surprise could seriously impair a policy outcome. Low robustness means that even minor surprise could be pernicious, while large robustness implies that only large surprise could be pernicious, so robustness helps to prioritize the options according to the ability to manage important uncertainty.

The robustness questions differs fundamentally from a question that one might think is central in policy selection. This different question is the **optimization question**: what policy achieves the best outcomes on the issues of importance to the policymaker? Policymakers have a moral (and perhaps statutory or political) obligation to do their best. The problem is that prioritization of policy alternatives based on their predicted outcomes is fraught with great uncertainty.³⁶ For example, the outcomes of innovative policies (that seem better than other policies) are usually less reliably predicted than the outcomes of more standard policies: a new and innovative policy is usually less well understood than a standard and familiar policy. Prioritization based on prediction-based optimization is highly uncertain and may be irresponsible.³⁷

When facing severe uncertainty, the policymakers' obligation to do their best is fulfilled by selecting the policy that achieves specified critical goals (that themselves may be highly demanding) over the greatest range of future surprises. Robust-satisficing is a responsible decision strategy in situations of severe uncertainty. It is the analysts' obligation to support this process by assisting policymakers to robustly satisfy their goals. Furthermore, because major uncertainties will always remain, analysts can augment the efficiency of their effort by

³⁶ Yakov Ben-Haim, Doing our best: Optimization and the management of risk, *Risk Analysis*, 32(8) (2012) 1326-1332.

³⁷ Yakov Ben-Haim, Craig D. Osteen and L. Joe Moffitt, Policy dilemma of innovation: An info-gap approach, *Ecological Economics*, 85 (2013) 130-138.

(1) reducing uncertainty (first meaning) when possible, and (2) enhancing robustness against persistent uncertainties. Thus both aspects of 'reducing uncertainty' are addressed.

The analyst is policy neutral in the sense that the analysis is contingent on goals chosen by the policymaker. In the spirit of 'what-if', the method of info-gap robust-satisficing allows the analyst to explain the impact of uncertainty on the outcome if the goal is this, or that, or something else. Robust-satisficing leads to policy prioritization that is contingent on the policymaker's goals and requirements. This obviates the analyst's tendency to advocate, without forgoing policy-relevance of the analytical product. The intention of the analyst's involvement in the policy process is not to influence the choice of goals. Rather, the analyst's relevance is assured by assisting policymakers to achieve goals that the policymakers themselves specify.

Example: 2002 Assessment of Iraqi WMD

Canton³⁸ presents a fascinating speculative reconstruction of how the intelligence assessment of Iraqi capabilities with chemical weapons (CW) could have been done differently in the months leading up to the 2003 US invasion of Iraq in "Operation Iraqi Freedom" (OIF). We will illustrate how an info-gap robust-satisficing analysis would contribute to the active management of uncertainty proposed by Canton. We will illustrate and elaborate on our two claims: (1) Attention needs to focus on enhancing robustness against uncertainty because the consideration of robustness can justifiably alter the policymaker's selection of policy. (2) This need not undermine the analyst's policy neutrality because the analyst's contribution is to assist the policymaker to explore the implications of uncertainty for achieving the policymaker's goals.

Canton relates that, after a structured procedure of review and analysis of diverse intelligence sources including signal and human intelligence, the team identifies three scenarios. Scenario 1: "Iraq destroyed stockpiles. Truck traffic [suspected of carrying CW] may be a denial and deception effort." Scenario 2: "Iraq has stockpiles but limited ability to produce more. Truck traffic likely military related." Scenario 3: "Iraq is replenishing existing stockpiles." (pp.510-511)

In the final analysis, Canton writes that "The lead analyst believes that scenario one is most likely but needs some qualification. He also believes that scenario three cannot be ruled out

³⁸ Canton (2008) *op.cit.*, pp.502-512.

entirely. ... The analyst states his reasons why scenario three cannot be ruled out: Intelligence collection on Iraq CW has been limited and weak to date, and volume is now growing but quality is doubtful. ... The issue [of scenario 3] must be raised in a final product but treated as a "worse case" scenario." (p.512) The team leader is now ready to write the intelligence assessment.

Info-gap robust-satisficing would carry the analysis one step further by assessing the robustness against uncertainty as a basis for the policymaker to choose between alternative actions for achieving specified goals. In the lead up to OIF the range of military and political policy options was large. For simplicity we will consider two prototypical options that in reality each entail a number of alternatives. The two options we consider are 'Initiation of War' (IW) and 'No Initiation of War' (NIW).

We proceed in three steps. (1) The team's final intelligence assessment underlies putative best-estimates of the cost (broadly construed) of the outcome of each policy option. This assessment would underlie policy selection that seeks to minimize the cost. (2) We evaluate the reliability and usefulness of these putative best estimates in light of the uncertainties. We do this by assessing the robustness against uncertainty, and the trade off between robustness and the cost of the policy outcome, for each policy option. (3) The putative optimum, identified in the first step, is not necessarily the most robust for satisfying the policymaker's goals. The analysis concludes by identifying the potential for a reversal of preference between the policy options.

Putative best-estimates of outcomes. The final conclusion of Canton's reconstruction of the intelligence analysis was that the "most likely" (p.512) contingency is scenario 1: "Iraq may possibly have some limited agent and precursor stockpiles ... [b]ut these are likely negligible ... We believe that Baghdad retains a chemical infrastructure capable of renewing CW agent production upon a decision to do so and may be considering such steps." (p.510) This assessment, if true, implies that initiation of war at this time (option IW) would entail many diverse and significant costs (human, political and monetary), while no initiation of war (option NIW) and continuing with the status quo would entail far lower costs on all dimensions of interest. Regardless of how costs are assessed, the most likely scenario implies that the putative best estimate of the cost of NIW is substantially less than the putative best estimate of the cost of IW.




Figure 1: Putative best-estimated outcome costs for two policy options.




Figure 2: Robustness curve for one policy option.

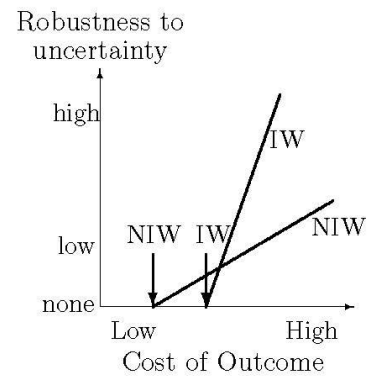


Figure 3: Robustness curves for two policy options.

This comparison of the outcome-cost of the two policy options, implied by scenario 1, is illustrated schematically in fig. 1. This figure is a **graphical metaphor**. The diverse costs are not evaluated numerically (though some costs may be quantitative). This figure is only intended to focus attention on the fact that the most likely contingency, scenario 1, predicts lower outcome-cost for option NIW than for option IW. The policymaker who asks the optimization question - what policy option is predicted to achieve the best outcome? - might tend to prefer NIW because it is predicted, by the most likely scenario, to have lower overall cost. The intelligence assessment does qualify the "most likely" status of scenario 1 by noting major uncertainties. Nonetheless, the outcome-optimizer could view scenarios 2 and 3 as outliers not justifying the greater predicted cost of IW.

Robustness against uncertainty. The robustness question is: what policy option is most tolerant to uncertainty in the assessment? Which option can absorb the greatest error without jeopardizing achievement of the critical goals? A policy option has low robustness if even small errors or surprises could result in failure to satisfy the outcome requirements. A policy has large robustness if only large error or surprise jeopardizes achievement of the requirements. We now explain how the robustness analysis considers both the uncertainties and the policymaker's goals, and how it does so without undermining the analyst's policy neutrality.

We begin by considering the robustness to uncertainty of option NIW. Suppose that the policymaker indicates to the analyst that only very low costs are tolerable, something around the predicted costs of the status quo or perhaps a bit more. Option NIW would, according to the best intelligence analysis, satisfy this requirement, but major uncertainties were identified

in the final intelligence assessment. This suggests that the actual costs of NIW may indeed exceed the predicted costs.

Fig. 2 is a schematic portrayal of the robustness idea, still considering only option NIW. We will explain two ideas. (1) Zeroing: predictions have no robustness against uncertainty. (2) Trade off: robustness increases as the outcome requirements become less demanding.

Zeroing. The horizontal axis of fig. 2 is the same as in fig. 1, and the vertical axis represents robustness against uncertainty. The predicted cost of NIW, indicated by the vertical arrow as in fig. 1, presumes the intelligence assessment of scenario 1 is correct. Even very small deviations of future reality from scenario 1 could cause the cost of NIW to exceed the predicted cost. In other words, the predicted outcome of NIW has no robustness against uncertainty. (Recall that an option is highly robust if only large errors can result in failure to achieve the specified goals. An option has low robustness if even small errors can lead to failure.)

The significance of the zeroing phenomenon is that prioritizing policy options according to their predicted outcomes, as suggested by fig. 1, is unwise (when facing radical uncertainty) because these predictions have no robustness against surprise.

Trade off. If the policymaker can tolerate only slightly greater costs than the scenario-1 prediction for NIW, then only small error can be tolerated and the robustness to uncertainty is low but not vanishing. The upward-sloping curve in fig. 2 is a graphical metaphor for the trade off between robustness and outcome requirement: as the policymaker relaxes the outcome requirement (accepting greater possible outcome cost), the robustness to uncertainty increases. This is not a surprising conclusion. A more demanding requirement (lower acceptable cost) is more vulnerable to failure resulting from pernicious surprise. The robustness trade off is an extension of the zeroing property. Predicted outcomes have no robustness, and worse-than-predicted outcomes are progressively more robust. The policymaker must explore the robustness-vs-outcome trade off in seeking an acceptable quality of outcome that can be confidently anticipated in light of the severe uncertainty.

Part of the procedural importance of figs. 1 and 2 is to make explicit the implications of the analysis of uncertainty. Fig. 1 depicts the predicted outcome of NIW based on the most likely scenario. Fig. 2 reminds the policymaker that such predictions are not robust against uncertainty, and thus are not a reliable basis for policy selection.

A further procedural implication represented by fig. 2 relates to feasible effort in reducing uncertainty (first meaning). The analyst's judgment, embodied in the gradual slope of the robustness curve in fig. 2, is that very high robustness for NIW is not plausible at the present

time. Recall that the intelligence report stated that "Intelligence collection on Iraq CW has been limited and weak to date, and volume is now growing but quality is doubtful."³⁹ Only major intelligence collection effort is likely to alter that judgment in the near term.

An additional procedural implication of the robustness analysis will be explained shortly in comparing alternative policy options.

A robustness curve such as in fig. 2 is schematic, not quantitative. It focuses attention on two orthogonal aspects of the policymaker's challenge: critical goals must be satisfied (horizontal axis), but major uncertainties must be managed robustly as well (vertical axis). A policy option cannot be evaluated solely in terms of its predicted outcome (the horizontal intercept in fig. 2) because that prediction has no robustness. Part of the analyst's task, in supporting the policymaker's policy choice, is to explain the irrevocable trade off between the robustness against uncertainty and the quality of the outcome. The analyst is not advocating or opposing the policy. The analyst is assisting the policymaker to explore the implications of uncertainty for reliably achieving the goals that the policymaker deems necessary.

Preference reversal among policy options. The robustness curve in fig. 2 is purely schematic: a graphical metaphor for the trade off between robustness and outcome. We do not evaluate the curve numerically, nor do we know the actual shape of the curve. The axes have no scale so it does not matter if we draw the curve with steep or gradual slope. However, what is true of all robustness curves is the upward slope, expressing the robustness trade off.

When we draw the robustness curves of two different options on the same graph, as in fig. 3, their relative slopes and intercepts express contextual understanding of the two options. This can reveal far reaching implications for the policymaker, as we now explain.

We know from fig. 1 that option IW has higher predicted cost (based on the most likely scenario) than option NIW. Likewise, in fig. 3 the horizontal intercept of the robustness curve for IW is to the right of the horizontal intercept for NIW. Both of these predicted outcomes have no robustness against uncertainty (this is the zeroing property).

Both robustness curves in fig. 3 have positive slope because the trade off between robustness and outcome is a property of all robustness curves. However, the analyst has drawn the IW robustness curve with steeper slope than for the NIW curve, reflecting different judgments about the uncertainties of these two options. While the initiation of war (IW) introduces many new uncertainties, it also takes the initiative and this makes the future somewhat more certain. Stated differently, one of the big uncertainties with not initiating war

³⁹ Canton (2008) *op.cit.*, pp.502-512.

(NIW) in the Iraqi situation was that scenario 1 could be fundamentally wrong, making war ultimately inevitable and possibly initiated surprisingly by the adversary. That major uncertainty about future war is removed with option IW. Stated yet another way, uncertainty tends to be more pernicious and dangerous in defense than in offense.⁴⁰

The steeper slope of IW reflects the somewhat lower outcome-uncertainty of IW over NIW. This means that a given increment of robustness is "purchased" at a lower increment of outcome-cost with IW than with NIW. Stated differently, the "cost of robustness" is lower with IW than with NIW due to the greater uncertainty surrounding the latter option.

The robustness curves in fig. 3 cross each other, revealing a fundamental dilemma facing the policymaker. On the one hand NIW is judged to entail lower cost based on the most likely scenario. On the other hand, NIW entails greater uncertainty than the IW option because the other side may initiate war in an unanticipated manner and time. NIW is putatively better but more uncertain than IW. The paradigm of this situation is the innovation dilemma, discussed earlier, because innovations are purportedly better than standard methods, but innovations are usually more uncertain precisely because they are new and less familiar.⁴¹ The innovation dilemma is revealed by the crossing of the robustness curves, and that crossing also assists the policymaker to resolve the dilemma, as we now explain.

The robust-satisficing methodology is to choose the option that achieves the policymaker's goals with greatest robustness against uncertainty. Suppose the policymaker judges that status quo no-war costs are tolerable but much greater costs are not. The policymaker's outcome requirement is at or slightly to the right of the NIW arrow in fig. 3. The IW option has no robustness in this range of outcome costs, while NIW has low but not vanishing robustness. NIW is the robust-satisficing policy choice, contingent on this specification of the policymaker's outcome requirement.

If, on the other hand, the policymaker envisions that the costs of war are tolerable, then the outcome requirement is at or to the right of the IW arrow in fig. 3. In this case IW is more robust than NIW over most of this range of outcomes, and IW is the robust-satisficing policy choice.

Either NIW or IW will be selected, depending on the policymaker's outcome requirement. The putative best choice is NIW, but the preference for this putative optimum may be reversed, depending on the policymaker's judgment regarding acceptable outcomes. The analyst's role in this deliberation is to assist the policymaker to understand the implications of

⁴⁰ Kahn (2001) *op.cit.*, p.85.

⁴¹ Ben-Haim, Osteen and Moffitt, *op.cit.*, 2013.

the putative best estimate (scenario 1) and of the attendant uncertainties. The policymaker brings outcome requirements (and other judgments) to the discussion; the analyst brings conceptual tools for assessing the confidence - robustness against uncertainty - of achieving these outcomes with the various policy alternatives.

The procedural importance of fig. 3 is in focusing attention on a common dilemma arising from uncertainty: NIW is putatively better but more uncertain than IW. The figure also supports the policymaker's deliberations in attempting to resolve this dilemma.

The potential for reversal of the policymaker's preference from NIW to IW bears some similarity to the precautionary principle. Michael Heazle has pointed out a basic asymmetry entailed in the precautionary principle whereby it shifts "the burden of proof away from those opposing some activity or substance to those who support it. This forces supporters of the activity or substance in question to demonstrate that it is not harmful while also decreasing or eliminating the need for opponents to show why, and how it might be considered dangerous".⁴²

The info-gap robust-satisficing analysis is more nuanced than the precautionary principle, and is linked closely to the policymaker's explicit judgments of acceptable outcomes. Scenario 1 is "most likely" and supports NIW, while scenario 3 is a highly unlikely "worst case" and supports IW. A precautionary approach would be to pre-empt the remote threat implied by scenario 3 because this scenario could be disastrous if it materialized. In the robust-satisficing analysis, the putative choice of NIW is rejected and IW is selected if the policymaker decides that IW more robustly achieves the required outcome. Robust-satisficing treats the two policy options, IW and NIW, equivalently. The selection between them hinges on the policymaker's requirements for the outcome.

Iraq 2002 Revisited

The robustness curves in fig. 3 for the two policy options - initiate or refrain from war - cross one another, expressing a common dilemma of uncertainty facing policymakers. While severe uncertainty is always challenging, the robustness analysis is sometimes less ambiguous because robustness curves do not necessarily intersect. We will illustrate this with an alternative speculative interpretation of the Iraqi situation in the run up to the war of 2003. We continue with the two broad policy options considered so far: IW or NIW.

⁴² Michael Heazle, Policy lessons from Iraq on managing uncertainty in intelligence assessment: Why the strategic/tactical distinction matters, *Intelligence and National Security*, 25(3) (2010) p.307.

Heazle writes: "The majority view across the US, UK, and Australian intelligence communities was that no tangible or immediate threat from Iraq existed - even if some WMD capability were assumed, an assumption broadly held at the time - and that containment appeared to be working."⁴³ While much analysis underlies and qualifies this conclusion, one might imagine that analysts could judge that the uncertainties of not initiating war (NIW) are no vaster, and plausibly even less so, than the uncertainties of initiating war (IW). Furthermore, the most likely scenario supports NIW.

A concise graphical metaphor for this judgment appears in fig. 4. The robustness curve for NIW hits the horizontal axis at lower outcome-cost than the curve for IW. The zeroing property reflects the judgment that the most likely scenario supports NIW, as before. The universal trade off between robustness and outcome-cost acts on both policy options. However, the trade off is less severe for NIW than for IW, unlike the assessment summarized in fig. 3. Whatever outcome cost the policymaker deems to be acceptable, provided that it is no less than the status quo, that cost is obtained over a wider range of surprise with NIW than with IW. The judgments summarized in fig. 4 implies that policy option NIW is more robust than IW at all outcome requirements greater than the predicted cost of NIW. Policy NIW is said to be 'robust dominant' over policy IW. There is no innovation dilemma in this assessment.

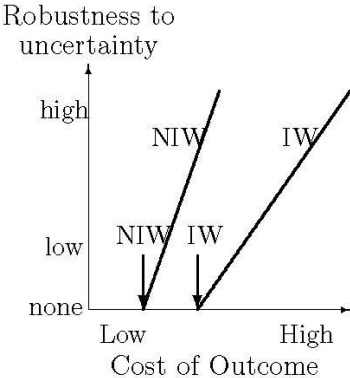


Figure 4: Robustness curves for two policy options.

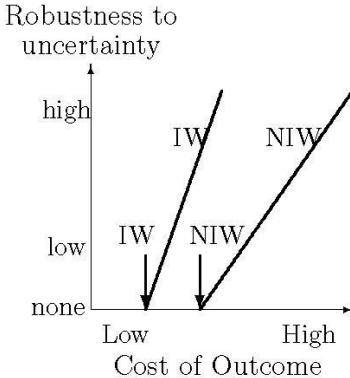


Figure 5: Robustness curves for two policy options.

Nonetheless, the robust dominance of NIW over IW does not necessarily imply that NIW is acceptable. It only implies that NIW is better (more robust against uncertainty) than IW over a particular range of outcomes. If the policymaker requires that the outcome cost be less, or at least no greater, than the predicted cost of NIW - more or less the cost of maintaining the

⁴³ Heazle (2010) *op.cit.*, p.305.

status quo - then NIW has no robustness against uncertainty and is therefore not an attractive option. The implication would be that the policymaker must "think outside the box" in search for some other policy alternative (perhaps a different refinement of the two classes of options we are considering).

Though the analyst has asserted the robust dominance of NIW over IW for a range of outcomes, the analyst has not intervened in the policy choice. The analyst's role is fulfilled by assisting the policymaker to explore the implications of uncertainty. The policymaker's role is to select a course of action in light of those (and other) considerations.

We now consider a different use of robustness curves: summarizing conflicting assessments. Robustness curves such as those in figs. 3 and 4 are useful if one recalls that they are intended as concise schematic summaries of underlying analytic interpretations. If that is kept in mind, then they can serve the additional purpose of highlighting the difference between alternative interpretations and supporting the deliberative process of policy analysis and selection.

We illustrate this by contrasting the interpretation embodied in fig. 4 with an interpretation that might seem to have guided US and UK policymakers. Our speculation is intended to illustrate an additional use of info-gap robustness curves. We are not embarking on an historical analysis of those governmental decision processes. As Heazle notes, "Debate over the respective policy motivations within the Bush administration and Blair and Howard governments continues and is unlikely to be conclusively resolved one way or another anytime soon."⁴⁴

Heazle writes that "the fact that the US was already committed to invading Iraq as early as mid-2002 is now generally accepted".⁴⁵ This could reflect an assessment by policymakers that the most likely scenario is that war with Iraq was inevitable and that the cost would be lower if the US takes the initiative and adopts policy IW.⁴⁶ Furthermore, the uncertainties accompanying IW are great but no greater, and plausibly less than, the uncertainties of leaving the initiative to Iraq or its allies and agents. In this situation the outcome-cost of IW is most likely lower than for NIW, and that IW is less vulnerable to surprise than NIW. In other words, this interpretation implies that IW is robust dominant over NIW. This assessment is summarized by the robustness curves of fig. 5.

⁴⁴ Heazle (2010) *op.cit.*, p.306.

⁴⁵ Heazle (2010) *op.cit.*, p.306.

⁴⁶ Daniel Bolger, *Why We Lost: A General's Inside Account of the Iraq and Afghanistan Wars*, (Boston: Houghton Mifflin 2014) p.71.

The schematic robustness curves in figs. 4 and 5 summarize conflicting assessments by different analysts. The purpose of juxtaposing these figures is to highlight the different policy implications of these competing interpretations. One might envision an analyst who has presented one of these interpretations to the policymaker who disagrees and responds with the other interpretation. Summarizing both positions with the robustness curves of figs. 4 and 5 facilitates the subsequent attempt to resolve the interpretive dispute. Is the dispute about the choice of the most plausible scenario? Or about the outcome-implication of that scenario (the horizontal intercept)? Is the negligible robustness to surprise of predicted outcomes fully appreciated (the zeroing property)? Does the policymaker appreciate the negligible robustness of trying to optimize the outcome by choosing the policy that is predicted to be best? Or perhaps the dispute arises from different assessments of the relative uncertainties of the two policy options (different slopes reflecting different trade offs between robustness and outcome)? These questions are made explicit in the difference between figs. 4 and 5.

Conclusion

Refined intelligence analysis is vital to policymakers and will command their attention if it convincingly assists them in facing their decision-making challenges. Furthermore, in the age of information and communication technologies, with easy access to vast information, intelligence services must heighten their usefulness more than ever.⁴⁷

The intelligence analyst's task "is to convert data into insight."⁴⁸ This entails, among other things, the task of reducing uncertainty. Within this latter task, we have claimed that it is not enough for analysts to reduce ignorance or ambiguity or potential for surprise in describing situations or protagonists' plans and intentions. Not all uncertainty is pernicious, and major uncertainties will always remain. We have claimed that analysts must also reduce policymakers' vulnerability to uncertainty.

In the extreme, if one eliminated all uncertainty, then one would have also eliminated the vulnerability to uncertainty. However, it is an "extrapolation error" to conclude that the analyst has fulfilled the mission of uncertainty-reduction by removing as much uncertainty as possible. Large uncertainty usually remains, so major vulnerabilities usually remain as well.

We have explained the info-gap robust-satisficing methodology by which policy options are prioritized according to their robustness against uncertainty while attempting to achieve

⁴⁷ Degaut (2015) *op.cit.*, pp. 2-3.

⁴⁸ Fingar, (2011) *op.cit.*, p.10.

specified policy goals. A policy is highly robust if only large uncertainty can jeopardize the attainment of the goals. Low robustness means that even small surprises can cause failure.

Evaluating the robustness of a policy requires the analyst to be familiar with the policymaker's goals, and this raises the spectre of policy advocacy by the analyst. All methods can be misused, but we have illustrated that the robust-satisficing method need not entail policy advocacy by the analyst. The analyst assists the policymaker to reliably achieve goals set by the policymaker. The policymaker's goals drive the analysis.

The robustness analysis supports deliberation and decision by the policymaker. The zeroing property of the info-gap robustness warns against prioritizing the policy options solely on the basis of their predicted outcomes because these predictions have no robustness against uncertainty. The trade off property of robustness vs. outcome helps identify reliable outcomes, assessed by their robustness against surprise. The robust-satisficing option is the most robust option, achieving the policymaker's specified goals over the widest range of uncertainty. We have illustrated that the most robust option may differ from the putatively optimal option identified by the predicted outcomes. The possibility for this reversal of preference, from the putative optimum to the robust-satisficing option, results from a subtle interaction between prediction and uncertainty that is succinctly portrayed by crossing robustness curves. Throughout the deliberation process, robustness curves serve as qualitative graphical metaphors for alternative interpretations and assumptions. The curves focus attention on disputes and assist in resolving them.

Notes on Contributor

Yakov Ben-Haim is a professor of mechanical engineering and holds the Yitzhak Moda'i Chair in Technology and Economics at the Technion - Israel Institute of Technology. He initiated and developed info-gap decision theory for modeling and managing severe uncertainty. Info-gap theory is applied around the world in engineering, biological conservation, economics, project management, climate change and natural hazard response, national security, medicine, and other areas (see info-gap.com).