Panel on The Intelligence-Policy Nexus
International Studies Association Annual Meeting
22-25.2.2017, Baltimore, Maryland

Assessing Change:
Intelligence Assessment in Support of Policy in Tumultuous Times

Yakov Ben-Haim[1]
Professor, Yitzhak Moda'i Chair in Technology and Economics
Technion – Israel Institute of Technology
Haifa, Israel
yakov@technion.ac.il

Abstract

We focus on the intelligence assessment of change, leading to support for policy decisions. Assessment must be based on addressing questions that produce substantively relevant understanding. However, severe uncertainty is the central challenge in highly dynamic environments. We explore Knight's concept of uncertainty, and the idea of indeterminism in human affairs as described by Shackle and Popper. This leads to the conclusion that intelligence assessment must be both substantively relevant and robust to uncertainty. However, we explain that relevance trades off against robustness. The strategy of robust-satisficing, based on info-gap decision theory, provides a conceptual framework for managing this trade off in assessment and policy selection. We illustrate the discussion by considering the conflict with the Islamic State.

1 Introduction

Intelligence assessment is challenging, especially in times of change. When analysts clamber on the slope of regional or global change, hard questions include: "Are new threats imminent?", "Are new actors emerging?", "Are we winning or losing?", "Is public support declining?" "How to support decision makers?"

We focus on assessing change in dynamic and uncertain environments, and using that assessment to support policy evaluation and selection. The issue is not primarily "Where are we on the slope of change?", but rather, "What is the gradient; what changes are taking place?"

Intelligence assessment is based on answering questions, like "Is violence increasing?" The choice of questions is critical. We claim that answers to these questions – and how one uses those answers – must have two conflicting properties: substantive relevance to decision makers, and robustness to error in one's knowledge. 'Substantive relevance' means enhancing understanding that leads to decisions for effectively achieving specified goals. 'Robustness to error' means that

[1] \papers\intl_stud-assoc2017\a-chg004.docx. 15.1.2017.

uncertainty of the answer does not alter the decision, that the decision would be the same even if the answer differed. Robustness is crucial because uncertainty is rampant, but robustness conflicts with relevance. We illustrate how to balance these considerations for counterinsurgency intelligence.

We begin with an example to set the stage, we then proceed to theoretical considerations of uncertainty, indeterminism, and robustness, and then return to apply these considerations to the example.

2 Example: Progress in fighting the Islamic State

Gibbons-Neff reported in the *Washington Post* on 26 April 2016 that "The flow of foreign fighters into Iraq and Syria has dropped from roughly 2,000 a month down to 200 within the past year, according to the Pentagon, which says the waning numbers are further proof of the Islamic State's declining stature."[2] This dramatic assessment has important strategic implications because "The declining number of fighters is a direct result of strikes that have targeted the terror group's infrastructure, [according to] Air Force Maj. Gen. Peter E. Gersten". In times of armed conflict, political and military leaders need to know if they are winning or losing, in order to maintain or modify current strategy or tactics.

Two difficult questions confound these assessments: what do the numbers mean, and is the assessment (the numbers and their meanings) robust against myriad uncertainties.

Does the 90% reduction reported by Gibbons-Neff really reflect diminished strength and stature, or does it reflect the rapid spread of the Islamic State to Libya, Tunisia, Sinai, Yemen and other regions, and diminished need (by the Islamic State) for fighters in Iraq and Syria? Has the 90% reduction actually reduced the Islamic State's military effectiveness, or does it reflect a diversion of non-military supporters whose service is more effective elsewhere?

Blanken and Lepore characterize the difficulty of such assessments by referring to the "cacophony of confounding variables introduced by the social context in which the operations take place."[3] A decline in *foreign* fighters does not rule out a rise in domestic support in the Arab region. Gibbons-Neff reports that "80 percent of Arab teens and young adults [aged 18 to 24] rule out any support for the Islamic State, a number that is up from 60 percent in 2015." However, the remaining 20% still represents millions of individuals whose support is significant. Furthermore, tacit or de facto support (perhaps due to fear) of the populations under Islamic State control is also important for the viability of that control. In other words, both the numbers,

[2] https://www.washingtonpost.com/news/checkpoint/wp/2016/04/26/number-of-foreign-fighters-entering-iraq-and-syria-drops-by-90-percent-pentagon-says, accessed 4.5.2016.

[3] Leo J. Blanken and Jason J. Lepore, Introduction: The challenge of wartime assessment, in Leo J. Blanken, Hy Rothstein and Jason J. Lepore, eds., 2015, *Assessing War: The Challenge of Measuring Success and Failure,* Georgetown University Press, pp.3-15.

and their meanings, could change substantially due to alternative interpretations or counting procedures. Without disputing the importance of a 90% reduction, both the number and its significance are vulnerable to substantial alteration given a different constellation of data and knowledge.

This points to the importance of multi-dimensional assessment that may be quantitative but that is strongly grounded in qualitative contextual understanding. For instance, Hix and Sepp discuss a 5-dimensional assessment of counterinsurgency in Iraq during 2004-2005, relevant also to the conflict with the Islamic State. They stress the need to address the following issues.[4] (1) Is the level of violence rising, and what is the nature of violent acts (e.g. coordinated or sporadic)? (2) Is the population freely giving information about insurgents to the police and armed forces? (3) Does insurgent propaganda substantially influence popular behavior? (4) Are the insurgents in control of their own casualties, that is, do the insurgents incur casualties in operations of their choice (e.g. ambushes or suicide bombings), or are their casualties primarily the result in US-coalition and Iraqi-government action? (5) Is there widespread popular support for the Iraqi government?

These five questions, with suitable semantic modifications, are all relevant to many assessment problems in counterintelligence and anti-terrorism. They are all grounded in historical or ethnological considerations as they bear on the viability of the insurgency. They relate both to the insurgents' efforts to rebel, and to coalition and government efforts to suppress the insurgency. The answers to those questions provide contextual understanding that is necessary in assessing change in the state of the conflict.

Assessment questions must be substantively relevant. However, while substantive contextual understanding is a necessary attribute of assessment questions, it is not sufficient. In addition, resulting decisions by political or military leaders must be robust against uncertainty in the answers. Before we can address issues of robustness to uncertainty we must first discuss Knightian uncertainty and Shackle-Popper indeterminism, and info-gap robust-satisficing, to which the next two sections are devoted. We then return to this example.

3 Knightian uncertainty and Shackle-Popper indeterminism

By 'uncertainty' we mean: ignorance or ambiguity or the potential for surprise. The concept of Knightian uncertainty is fundamental to our understanding of uncertainty in human affairs. Frank Knight's concept of 'true uncertainty' arises from innovation and initiative of entrepreneurs. In this connection, Knight asserts that "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

[4] William C. Hix and Kalev I. Sepp, Assessing Counterinsurgency: The Iraq War, 2004-5, in Blanken, Rothstein and Lepore, 2015, *op. cit.,* pp.197-213.

can be learned.[5] Knightian uncertainty arises from the unbounded potential for future innovation, or simply from ignorance of the vastly complex world. Uncertainty may also arise from deception or denial by an adversary.

G.L.S. Shackle[6] and, independently, Karl Popper,[7] explained a concept of indeterminism that derives from Knightian uncertainty.[8] 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 knowledge-based behavior will have an element of irreducible indeterminism today. Knightian uncertainty and Shackle-Popper indeterminism (SPI) imply a fundamental and irrevocable limitation in the ability to predict outcomes in human affairs.

We can understand SPI more thoroughly by contrasting it with the Newtonian paradigm. The physical sciences are grounded on the assumption that laws of nature exist and are discoverable, universal and stable in time. As our mastery of these laws improves, our ability to predict the trajectory of physical systems in space and time also improves. The methodology of science-based engineering, as it emerged in the 19th century, proceeds in two steps: first learn the laws that govern a system, and then use those laws to design or control the system. For instance, first learn Newton's equations of motion of physical objects, and then use that scientific understanding to control the motion of missiles, airplanes, and so on. We will call this methodology the Newtonian paradigm.

The Newtonian paradigm is widely endorsed in social sciences. For example, Ross and Makovsky write that "[w]hether trying to foster peace or liberalization in the Middle East or alter regime behaviors, we have to see the world as it is. … [O]ne cannot change an unacceptable reality before one understands it – and then it becomes possible to shape a strategy that produces change in stages."[9]

In contrast to the Newtonian paradigm, Shackle-Popper indeterminism implies that human affairs entail an element of behavior that does not follow stable universal discoverable laws. SPI implies an element of non-nomological behavior in human history that is fundamentally different from the realm of physics. The predictive power of the physical scientist depends on the ability to learn – at least to

---

5 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.

6 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.

7 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.

8 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).

9 Dennis Ross and David Makovsky, (2010), *Myths, Illusions, and Peace: Finding a New Direction for America in the Middle East,* Penguin Books, New York, p.315.

a high level of accuracy – the laws that govern the dynamics of the system of interest. Such laws exist for human systems only partially, so prediction in human affairs is likewise limited. While understanding and insight are important, it is equally important to recognize and account for Knightian uncertainty and Shackle-Popper indeterminism in making assessments and formulating policy.

We can get further insight into Knightian uncertainty and Shackle-Popper indeterminism by contrasting retrospective explanation with prospective prediction.

Historians can explain the origins of WW I or the fall of the Roman empire, though disputes will persist. This sort of explanation is not too different from the paleontologist's explanation of extinction of species.[10] *Ex post* explanation is useful and interesting, but one must beware not to conclude that stable, universal and discoverable laws underlie such explanations.[11] The ability to explain an event is very different from the ability to predict the outcome of an event. The difference is epistemological. We know many things after the fact that we do not know beforehand. Retrospectively, we know the outcome and we know a multitude of events, discoveries, innovations and inventions that accompanied the evolving outcome and led to the final result. Retrospective explanation benefits from this knowledge, while it is the unavoidable lack of this knowledge that causes Shackle-Popper indeterminism and that limits the power of prediction in human affairs.

We claim that the lack of crucial information and understanding limits the power of prediction. But more than that, we claim that comprehensive, stable and universal laws governing human behavior simply do not exist. Those laws evolve as events occur and as discoveries and inventions are made. This is the essence of SPI, and it means that prediction in human affairs cannot successfully or reliably imitate the Newtonian paradigm. It does not mean that nothing can be foreseen, or that all predictions will err. It does imply a limitation on the reliability of those predictions, and it has important methodological implications for intelligence assessment and policy prioritization in times of change, as we will see.

4 Optimizing and satisficing: An info-gap perspective

The decision methodology that could be called "outcome-optimization" begins by identifying the best available information, understanding, theoretical and contextual insight, including perhaps assessments of uncertainty. We will call this information our "knowledge". This knowledge entails information and understanding about friendly and adversarial capabilities, geopolitical constraints and opportunities, domestic politics, terrain, logistics, and whatever else is deemed relevant. Outcome-optimization chooses the option whose knowledge-based predicted outcome is best.

[10] Gould, Stephen Jay, 1989, *Wonderful Life: The Burgess Shale and the Nature of History*, W.W. Norton and Co., New York.
[11] Hegel and subsequently the Marxists would disagree.

Outcome-optimization is usually **unsatisfactory** for policy selection when facing change and uncertainty because our knowledge is likely wrong in important respects, as explained by Knightian uncertainty and by Shackle-Popper indeterminism. Instead, we advocate the decision methodology of robustly satisficing[12] outcome requirements.[13] The deficiency of outcome-optimization is important to the intelligence analyst because of the need to connect intelligence assessment to policy selection.

The basic idea of info-gap robust satisficing is to first identify outcomes that are essential – goals that must be achieved – and then to choose the decision that will achieve those critical outcomes over the greatest range of future surprise. We use our knowledge in two ways. First, to assess the putative desirability of the alternative decisions, and second, to evaluate the vulnerability of those alternatives to surprising future developments. A decision has high robustness if it achieves the critical goals despite large surprises or great errors in the knowledge. The robust-satisficing strategy is the one with maximal robustness against uncertainty while satisfying the critical requirements. In other words, what is optimized is not the predicted quality of the outcome, but rather the immunity to error and surprise. The outcome will be satisfactory, though not necessarily optimal, over the greatest range of future deviations from our current understanding. What constitutes a satisfactory outcome can be as modest or as ambitious as one wants, though the robustness varies accordingly.

We now return to the example that we introduced earlier in order to illustrate these ideas.

5 Example: Continuation

In discussing the assessment of change in the conflict with the Islamic State, we explained that assessment questions must be substantively relevant. We also emphasized that, while substantive contextual understanding is a necessary attribute of assessment questions, it is not sufficient. In addition, assessment supports decisions by political or military leaders, and these decisions must be robust against uncertainty in the assessment. If small errors in the answers to assessment questions could alter the leader's decision, then the robustness is low and the assessment questions are (at least in part) unsuitable, regardless of their

---

[12] To satisfice means "To decide on and pursue a course of action that will satisfy the minimum requirements necessary to achieve a particular goal." *Oxford English Dictionary,* online version accessed 7.4.2016.

[13] Further discussion of these ideas are found in Yakov Ben-Haim (2014) Strategy selection: An info-gap methodology, *Defense & Security Analysis,* 30(2): 106-119, and Yakov Ben-Haim (2015), Dealing with uncertainty in strategic decision-making, *Parameters,* the US Army War College Quarterly, 45(3) Autumn 2015. Robust-satisficing is central in info-gap decision theory. See Yakov Ben-Haim, *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edition (London: Academic Press, 2006). References to work of many scholars can be found at info-gap.com.

substantive relevance. Stated differently, we must ask: how wrong can our assessment be, and the leader's decision is still the same? Low robustness implies vulnerability (of the choice of action) to error in the assessment. Good use of assessment questions must lead to decisions that are insensitive to errors in the assessment. Equivalently, good use of assessment questions must generate robustness (of the decision) against uncertainty in the assessment. This is different from the substantive criteria for choosing the assessment questions. One should definitely use substantively relevant assessment questions, but that is not a sufficient criterion for selecting and using assessment questions. The questions and the way the answers are used should generate decisions that are robust against error in the assessment.[14]

It is important to emphasize that the decision-stability of a robust decision does not mean that the decision is rigid, unwavering or indifferent to existing knowledge. A robust decision is responsive to existing understanding, but is also consistent with a wide range of alternative possible knowledge that cannot be excluded on the basis of present understanding. This realism of the robustness is a result of the way that Knightian uncertainty and Shackle-Popper indeterminism are modeled and managed.

For example, the 5 questions posed by Hix and Sepp can all by answered in part either 'yes' or 'no'. Hence a voting mechanism can be used, at least initially, to determine whether the counterinsurgency (COIN) strategy needs to be changed.

In one voting mechanism, the policy decision (to maintain or to alter the current strategy) is determined by consensus.[15] We explain this with two illustrations.

Suppose the analyst faces the question: Is our COIN strategy faltering? An affirmative answer would imply the need for operational changes, while a negative answer would imply operational continuity. Using the 5 Hix-Sepp questions and the consensus rule, the answer would be affirmative if and only if all 5 answers indicate weakness or failure of the strategy. For example, an affirmative answer is supported by responses such as: "Is violence rising?" – Yes; "Does the population inform the police about the insurgents?" – No; and so on, with support for "COIN is faltering" also in the remaining three questions.

The consensus algorithm is highly vulnerable to uncertainty in the answers to the Hix-Sepp questions. A small error in assessment (a change of any one 'vote' among the 5 questions) would remove the consensus, reject the assessment that the strategy is faltering, and result in operational continuity. The consensus rule has low robustness to error in assessment. In other words, the consensus rule is vulnerable to uncertainty in the 5 answers. In contrast, a 4-out-of-5 rule would be less vulnerable to uncertainty in the answers.

[14] We are claiming that the intelligence analyst, in performing an assessment, must consider the subsequent decisions to be made by policy makers. This raises issues of policy neutrality of the analyst that are discussed elsewhere. See Yakov Ben-Haim, Policy neutrality and uncertainty: An info-gap perspective, *Intelligence and National Security,* vol. 31, #7, pp.978-992.

[15] Like the voting rule of the 5 permanent members of the U.N. Security Council.

A similar situation arises when the analyst faces the question: Is our COIN strategy succeeding? An affirmative answer implies operational continuity, while a negative answer implies operational change. The answer would be affirmative if and only if all 5 answers indicate success of the strategy: "Is violence rising?" – No; "Does the population inform the police about the insurgents?" – Yes; and so on. The analyst rejects the assessment that the COIN is succeeding unless consensus indicates its success. The decision rule again has low robustness to uncertainty because any single vote-change would lead to rejection of the conclusion that the COIN is succeeding.

We can now appreciate a fundamental dilemma that arises in managing Knightian uncertainty and Shackle-Popper indeterminism. On the one hand, we want strong substantive support for an assessment; consensus of answers to substantively relevant questions provides the strongest possible support. On the other hand, the consensus algorithm has low robustness to error in the answers: reversal (due to error or surprise) of any single response would change the assessment. This dilemma arises from an irrevocable trade off: requiring a more demanding outcome (consensus rather than, say, 4 out of 5), entails lower robustness to uncertainty. We want both strong substantive support and high robustness to uncertainty. However, these two attributes trade off, one against the other: greater substantive support for the assessment is obtained only by an algorithm that entails lower robustness to uncertainty. This trade off is generic, and is sometimes called the "pessimist's theorem" because any healthy pessimist knows that more exacting outcomes can be foiled by surprise in more ways and hence have lower robustness to uncertainty.

We have explained that the consensus rule – an assessment is rejected unless consensus on all questions is achieved – has low robustness to uncertainty in the answers to the questions. Consensus has substantive as well as political or institutional advantages. However, our discussion indicates that consensus rules lack robustness to uncertainty. Brams describes many political and social considerations in designing voting rules in democracies.[16] Our discussion of this example suggests an additional consideration: robustness of the voting outcome to uncertainty in the election results. Most importantly, one's ability to realistically integrate diverse information is enhanced by understanding the trade off between substantive support and robustness to uncertainty.

6 Conclusion

Intelligence assessment is a process of inference whereby information and understanding about the past and the present are used to draw conclusions about possible futures. Assessment is important because of its use in making decisions. We

[16] Brams, Steven J., 2008, *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*, Princeton University Press.

have described the method of robust satisficing, discussed its conceptual foundations and illustrated some of its implications for decision making. We conclude our discussion by suggesting that robust satisficing fits into a taxonomy of inferential methods. We will consider three categories of foresight or forecasting methods: early warning, scenario analysis, and robust satisficing. We will explain that these three methods can be understood as different types of inference: inductive, abductive, and deductive, respectively. There is important overlap between these methods, and the taxonomy does not comprehensively characterize these methods. Nonetheless it reveals some fundamental distinctions.

Inductive inference: "early warning" methods. The category of inferential methods that we are calling "early warning" includes all quantitative predictive methods (even when they are not intended to warn against anything) whereby data are used to infer something about the future. Trend analysis is in this category, as well as, for example, statistical prediction of state failure.[17] The early warning category also includes less quantitative "weak signal scanning,"[18] including qualitative and heuristic methods, that attempt to detect early signs of impending momentous change by parsing evidence of any sort.

One can characterize the early warning methods of assessment and prediction as inductive inference. As Peirce explained long ago, "Induction is where we generalize from a number of cases of which something is true, and infer that the same thing is true of a whole class."[19] In early warning methods, data or other perceptual evidence are used to infer that what is true of past and present observations will be true of a broader class including the future.

Abductive inference: scenario analysis. An example of abduction is: we observe that the lawn is wet, so we conclude that it rained last night. The conclusion is a simple explanation of the observation. The conclusion is not an induction because it is not a generalization based on accumulated evidence. The conclusion is neither necessarily true (the neighbor's sprinkler could have wetted the lawn) nor is it sufficient as an explanation (the lawn could have absorbed last night's rain and hence be dry now).[20] As Peirce wrote, an abduction is "the operation of adopting an explanatory hypothesis",[21] and "[t]he abductive suggestion comes to us like a flash. It

---

[17] G. King and L. Zeng, 2001, Improving forecasts of state failure, *World Politics,* Vol. 53, Issue 4, pp.623-658.

[18] Sandro Mendonça, Miguel Pina e Cunha, Frank Ruff and Jari Kaivo-oja, 2009, Venturing into the Wilderness: Preparing for Wild Cards in the Civil Aircraft and Asset-Management Industries, *Long Range Planning,* 42: 23-41. See p.27.

[19] Charles Sanders Peirce, Deduction, induction, and hypothesis, *Popular Science Monthly,* August 1878, reprinted in Charles Sanders Peirce, *Chance, Love and Logic: Philosophical Essays,* Edited by Morris R. Cohen, University of Nebraska Press, 1998, pp.131-149. See p.135.

[20] We say that proposition *A* is "necessary" for proposition *B* if *A* must be true if *B* is true. We say that proposition *A* is "sufficient" for proposition *B* if *B* is true whenever *A* is true.

[21] Charles Sanders Peirce, Abduction and induction, in *Philosophical Writings of Peirce,* Justus Buchler, editor, pp.150-156. See p.151. Dover reprint of *The Philosophy of Peirce: Selected Writings,* Routledge and Kegan Paul, 1940. The manuscript of this essay is from c. 1901.

is an act of *insight*, although of extremely fallible insight."[22] Peirce believed that abduction is a creative process that generates a new idea, a provisional hypothesized explanation. Halas demonstrates that abductive inference is not uncommon in the study of international relations, most notably in the study of agent-based models.[23]

Scenario analysis is not strictly an abductive inference from observation to hypothesized cause. Nonetheless, scenario analysis is the creative invention of hypothesized futures based on speculative contemplation of uncertain possible occurrences. The scenario is the hypothesized cause of an event that has not yet occurred but could perhaps occur. Scenario analysis is abduction in the future tense, abduction that is contingent on uncertain possibilities. Scenario analysis is a learning experience through which analysts extend their understanding. For example, consider a scenario of increasing Jihadist terrorism in Western countries. Within this scenario we can envision deliberate destruction of symbolic or historical sites, incitement by religious leaders, suicide bombers in movie theaters, and so on. The futuristic abduction is: observing (in our mind's eye) these consequences, we hypothesize the scenario of increasing Jihadist terrorism. Of course, scenario analysis goes further by exploring other possible implications of the scenario. What we have described as abduction is the process of inventing scenarios.

Deductive inference: Robust satisficing. A deduction is the use of a general rule and a specific case, to infer something about the specific case. Peirce's example is:[24]

"Rule. – All the beans from this bag are white.

" Case. – These beans are from this bag.

"∴ Result. – These beans are white."

Deduction is in a sense the reverse of induction, because induction goes from a collection of cases to the construction of a general rule. A deduction uncovers implications that are inherent in the rule (or principle or axiom), but does not invent anything new or outside of that axiomatic system.[25] In this way a deduction differs from an abduction which is a creative insight that extends the domain of thought.

We can now understand the deductive nature of an inference that is based on robust satisficing. Consider intelligence assessments such as "Situation *S* will arise within 1 year" or, "Policy *P* will achieve these specified goals". Suppose furthermore that these assertions are consistent with our current knowledge and understanding of the past and present. How much confidence can we have that they will hold true in the future?

---

[22] Charles Sanders Peirce, Perceptual judgments, in Buchler, *op. cit.*, pp.302-305. See p.304. The manuscript of this essay is from c. 1903.

[23] Matus Halas, 2015, In error we trust: an apology of abductive inference, *Cambridge Review of International Affairs,* Vol. 28, Issue 4, pp.701-720.

[24] Peirce, Deduction, induction, and hypothesis, *op. cit.* p.134.

[25] For Henri Poincaré this property of deduction seems, perplexingly, to reduce mathematics to "a gigantic tautology". The resolution, he suggested, was that "a single formula [contains] an infinite number of syllogisms". Henri Poincaré, 1905, *Science and Hypothesis,* English translation, republished by Dover Press in 1952, pp.1, 9.

We **cannot** know how large a surprise will occur, or how much our current understanding of the future errs due to Knightian uncertainty. More importantly, we do not confidently know the mechanisms that will govern the future evolution of the system in question because of Shackle-Popper indeterminism. However, we **can** deduce how much surprise or error we can tolerate. Confidence in the assertions about situation *S* or policy *P* is based on answers to the robustness question, which is: how much could we err, or how large a surprise could occur, and the assertion is still true? If the answer is that great error or surprise does not jeopardize the truth of the statement, then it robustly satisfies a requirement regarding time of occurrence (for situation *S*) or goal-achievement (for policy *P*).

The robustness question is not ontological. It does not ask “How wrong are we?” or “What is the largest surprise that could occur?”. We cannot answer these questions (though inductive early warning methods might help). But we can deduce an answer to the robustness question because it addresses what we do know. Stated differently, the answer to the robustness question does not tell us how much we actually err; it only tells us how much error we can tolerate. Whether or not our actual ignorance exceeds this limit remains unknown to us. The answer to the robustness question is deductive, and tells us nothing beyond what is inherent in our knowledge as it currently stands.

In summary, the info-gap method of robust satisficing begins by explicitly recognizing that we do not know how much we err, or how large a surprise might occur, or how bad things can get; there is no known worst case. The decision maker then specifies how good an outcome is required or, equivalently, how bad an outcome is tolerable. Finally, for any proposed intelligence assessment or policy recommendation, we determine the greatest horizon of uncertainty up to which that assertion is still valid. In particular, a robust assessment is one that remains valid despite great error or surprise. When the assertions entail policy options, the robustness analysis prioritizes those options based on confidence in achieving specified goals.

Acknowledgement This paper was written while the author was at the University of Oxford as a Distinguished Visiting Scholar, an Oxford-Martin Visiting Fellow and a Leverhulme Visiting Professor. The author is grateful for the support of these foundations.