Designing a Robust Decision–Based National Security Policy Process: Strategic Choices for Uncertain Times

Chapter Eleven

STEVEN W. POPPER, RAND CORPORATION

This chapter is extracted from *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, by Aaron B. Frank and Elizabeth M. Bartels, eds., RR-A1275-1, 2022 (available at www.rand.org/t/RRA1275-1).

Prepared for the Defense Advanced Research Projects Agency, Defense Sciences Office
Approved for public release; distribution unlimited

RAND
CORPORATION

For more information on this publication, visit **www.rand.org/t/RRA1275-1**.

About RAND

The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.
© 2022 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

CHAPTER ELEVEN

Designing a Robust Decision–Based National Security Policy Process: Strategic Choices for Uncertain Times

Steven W. Popper, RAND Corporation

The Pentagon. Whitehall. The Kremlin. Foggy Bottom. Horse Guards. Quai d'Orsay—these are how people once referred to the foreign policy and national security establishments of major nations. In the past, one could easily envision men (and it was only men who traveled in this world) arriving with briefcases to an office, sitting at desks, joining deliberative processes in meeting rooms, taking lunch in clubs or restaurants and then working often to a late hour illuminated by candle, gas, or incandescent lamps. During the days, officials would sift information, sort it in accordance with standard rules of thumb and established protocols, consider the information gained through their intelligence services' observation of their opposite numbers, debate policies and, over time, frame actions and responses that would then be commended as courses of action (COAs) to respective governments in minutes, memoranda, and white papers.

That was then, this is now. There are new players (some undetected for long periods); new arenas of competition; new stakes; an ever accelerating pace of communication and hence a decreasing time cycle of decision; widening variation of intentions, objectives, and strategies; and a vastly more voluminous information flow paradoxically accompanied by an alarming rise in fundamental uncertainties. The nature of international interaction has changed and, along with it, the processes and protocols of decisionmaking.

But have the processes of decisionmaking changed enough? In particular, has analytical support for decisionmaking made possible the type of transition that the new era calls for? In the face of fluid conditions, deep uncertainties, and changing relationships, what capabilities would be the most desirable in an apparatus for conducting deliberation and analysis of policy alternatives on national security, relations with allies or potential adversaries, and dealing with today's uncertainties and their challenges to U.S. long-term goals (as expressed in governing policy documents, such as the U.S national security strategy)?[1]

[1] White House, *National Security Strategy of the United States of America*, Washington, D.C., 2017.

A *value proposition* for innovative analytical support to meet contemporary needs should involve developing better methods for operating under uncertainty, integrating cross-agency processes, and enhancing the means for supporting organizational foresight. Although treated separately later, this chapter's thesis argues that the necessary transformations can be found through incorporating and exploiting the concept of robustness and in a refocus on decision support—rather than forecasting—as unifying principles within the National Security Enterprise (NSE). [2]

In this chapter, and in keeping with one of the major principles for enhancing policy robustness, we start by first identifying end goals. We enumerate six major difficulties that persist in the decisionmaking process now and identify potential means to overcome them. The next section considers the role that computation might play in analytical support to the NSE. The section following suggests that the very problem of deep uncertainty provides connective tissue for the integration of the planning, monitoring, and analytical systems that could well fulfill much of the agenda for contemporary analytical support to national security planning and decisionmaking. This is followed by a discussion of the role for adaptiveness and robustness that also lays out a design for analytical support to the NSE policy deliberation process—support that could better address deep uncertainties and, in so doing, go a long way toward dealing with the six obstacles. We close with a brief section of concluding thoughts.

Obstacles in NSE Decisionmaking to Be Overcome

Several aspects of contemporary NSE planning processes pose challenges for achieving the objective of becoming more flexible and adaptive in the face of greater uncertainty. We identify six obstacles and what we might do to overcome them in NSE planning.

Avoid Bifurcation in Focus by Evaluating Short-Term Actions in a Strategic Perspective

The United States has lost few battles in more than half a century. But there are few wars in which it has achieved the political outcomes it stated at the onset. This might partly stem from the strong distinction drawn in the United States between civilian policy and military operations and, thus, a certain bifurcation of focus. This exacerbates the fundamental difficulty of thinking along multiple timescales in parallel, especially in complicated or ill-defined arenas of conflict, such as undergoverned spaces. Knowing the next steps after securing battlefield victory is as important as achieving the victory itself. In the absence of the first,

[2] This chapter will shift across the components of the NSE, from policy deliberations in the U.S. Department of State, through military planning in the services and U.S. Department of Defense (DoD), to the offices involved in intelligence-gathering and all its forms. One of its central theses is the need for integrated assessment, planning, and implementation across these functions.

the second becomes hollow. The United States paid the price for that bifurcation in focus in the war in Iraq during which it easily won the military battle but did not consider sufficiently the consequences of victory. This led to the political vacuum in postwar Iraq. Being able to plan operations, evaluate intelligence, and trace operational pathways from means to ends within a consistent decisionmaking framework would support the difficult task of relating short-term actions to prospective long-term consequences.

Resistance to Uncertainty Absorption

The concept of *uncertainty absorption* formulated by Nobel laureate Herbert A. Simon captures the phenomenon of the lower levels within an organization being more cognizant of uncertainties pertaining to the sources, character, and quality of intelligence than are leaders in the higher levels.[3] What moves upward through organization channels is not raw intelligence gathered at the lower levels; rather, it is syntheses and interpretations based on such information. Nuanced understanding of the variation among sources is necessarily stripped away so as not to clog channels going upward and possibly compromise organizational function. Unfortunately, the ability to drill down and examine the foundations for an interpretation is also often stripped away or lost in transmission. What this loss leads to is less communication of subjective risk perceptions than might be purposeful within organizations and points to a need for processes better suited to enhancing the quantity and quality of information exchange. The quantification of intelligence, particularly the characterization of uncertainty by probabilities, might lead to similar effects. The capacity to convey messaging on COAs, means, and goals while retaining tools for interrogating an interpretation's underlying determinative factors would restore some of the nuance lost in the process of synthesis.

Penetrate Stovepipes

Establishing mission-oriented offices and agencies necessarily gives rise to inter-agency—and intra-agency—stovepipes of information and responsibility. The ideal would be to carry forward integrated discussions on policy objectives, intelligence, strategic concepts, operational requirements, mission requirements, existing and prospective capabilities, and mission-agency organizational goals. Although this is difficult enough to do within one mission agency, the difficulties of crafting comprehensive processes across those agencies are daunting and yet vital. The more complex the problem, the less reliable should be the confidence placed in any one organization to be uniquely authoritative and sufficiently expert.

Constrain Cognitive Dissonance

Victory disease—a term coined by Japanese officials to describe the state of mind of war staffs after early victories in World War II—equally applies to U.S. experience, such as the Iraq

---

[3] James G. March and Herbert A. Simon, *Organizations*, 2nd ed., Cambridge, Mass.: Blackwill, 1993.

example, the shared situational understanding prior to the Battle of the Bulge, or the march to the Yalu in 1950. The phrase rolls up aspects of the confirmation bias and communal reinforcement (groupthink) noted by psychologists. Any planning team may perceive in the evidence before them, particularly when questionable and less than clear, encouraging signs for the COA being advocated or pursued. The capacity to reformulate such evidence into configurations that might be equally plausible or depend on different assumptions that cannot be rejected based on information at hand is not only absent but consciously or unconsciously resisted. It is challenging to risk the ire of (and perhaps ostracism from) a planning group already under pressure for delivery and pleased with its own performance by raising late-stage doubts. Less disruptive embrace of *red teaming* within the flow of the planning process itself could shore up weak points of potential failure.

As a related matter, there is value in also constraining *proof by loud shouting* which is the ability of the most prevalent or well-articulated views to dominate. A decision process can often be gamed by adding redundant data and modeling runs that will, in effect, limit other voices from being heard and perhaps even lock them out entirely. There is value in an analytical approach that will limit domination by repetition and instead reward diversity in a systematic approach that is also purposeful and operationally meaningful.

Recapture Scenario Value

DoD has embraced scenario planning of future contingencies and the resulting potential demands. Integrated security constructs (ISCs) help enhance the joint understanding of potential operational and mission demands and so, working backward, the materiel, skills, and readiness posture required to prepare the forces needed to carry out missions. The process of generating ISCs generally enfolded into the Quadrennial Defense Review (QDR) process is complicated, costly, and subject to negotiation and approval by all the Services and offices of the defense establishment. Once the approved set of ISCs is compiled, all involved are both relieved and loathe to fiddle with them further. Worse, no one wishes an interservice contretemps to flare up by having new conditions supposed that would be viewed as invidious to policies and programs of one or another service once underway.

In short, the system manages to again remove uncertainties not explicitly captured during ISC development from further consideration. This is the antithesis of value that the ISCs were intended to provide. As with some of the dysfunctions already noted, the organizational psychology of dissent in planning is even more fraught than it previously was. There is a need to recapture the value intended to be gained by making possible scenario thinking within the national security planning apparatus.

Make Foresight Less Precarious

A precarious value within an organization or process is one that has not been sufficiently defined, is seen not to have received sufficient legitimacy through leadership support, or appears to be inimical to widely held understanding of vital missions.[4] Activities and functional teams seen as failing to, not contributing to, or perhaps even distracting from activities focused on the organization's bottom line will be sidelined—either consciously or unconsciously—and potentially ostracized functionally or eliminated entirely. Foresight activities generally fall within this category. Those activities by their nature look beyond the next production horizon, ask difficult questions of busy folks, and are carried out by people who are not seen as contributing directly to organizational goals.

Foresight offices or functions are created with much fanfare and then wither or disappear entirely in one organizational shuffle or another. Those that do not perish have mastered one simple task: They have caused internal demand for their output to grow. The value proposition for organizational foresight activities must be made sufficiently strongly that the line units of the organization perceive it.

Summary

Taken together, overcoming these six obstacles in planning form a tall order. They appear collectively to be a stretch well beyond the contemporary capability to fulfill them. Few would argue with their desirability, but practical considerations consign them to a realm of aspiration beyond realization. The very rise in uncertainty that calls out the need for change would appear to overwhelm the attempts to deal with these obstacles credibly in a meaningful analytical construct and organizational setting.

However, more explicit recognition of the same deep uncertainty that has so complicated the task of national security policy analysis may serve as the way to deliver on the value propositions outlined. The overwhelming problem of sheer numbers—so many uncertainties multiplying in interactions with each other—that manifests as myriad pathways from the present to the future gives pause. But it also raises a fundamental and perpetual concern of analysis: Are we asking the right questions given what we face? There is also the sensitive concern of whether the NSE as a whole is well served under new circumstances by the traditional separation between its intelligence-gathering, analysis, and knowledge creation component in the Intelligence Community (IC) and deliberations by the planners and decisionmakers (known as the *policy community*) on the other, with little comprehensive reference between them. If more direct interaction within the NSE is deemed potentially valuable, then how can the connection between the two be brokered without introducing another problem of policy contaminating intelligence—or vice versa?

[4] Burton R. Clark, "Organizational Adaptation and Precarious Values: A Case Study," *American Sociological Review*, Vol. 21, No. 3, 1956.

Three Pathways: The Role That Computation Might Play in Analytical Support to the NSE

Beginning with World War II, computation—or information machines—came to play a role for the military. From the U.S. Navy's development of combat information centers aboard its ships to the early computers used to crack the Enigma machine and reveal its coded information, this new type of machine became ubiquitous and preponderant. Command, control, communications, and information (C3I) had always been crucial in determining battle outcomes, but it was most conspicuous by the extreme difficulty in gaining anywhere near enough of each, to say nothing of their integration. Incremental gains in advantage were hard won. But now machines rather suddenly became capable of vastly enhancing these capacities.

We are still in the midst of this revolution. As technical means advance, computing moves from just providing information and communications to helping with decisions. The extent to which computing is brought into—or becomes in itself—a decision system deserves some consideration. There are three different but related channels for doing so—prediction, automated and expert systems, and adaptive planning—but they are sufficiently distinct that there is value in recognizing the differences among them. Each has a distinctive approach to dealing with the problem of making decisions in the presence of uncertainty.

Prediction

Military commands and national authorities have always sought to improve prediction, whether of adversary actions, external conditions affecting operations, or of likely outcomes from employing a COA. We have moved from Delphic oracles or auguries reading the entrails of sacrificed beasts to more-sophisticated techniques. Computer models and simulations are powerful tools for rigorously examining systems and outcomes. Such models are extraordinary human artifacts that encapsulate mountains of knowledge and experience gained in many fields. Models keep track of myriad causal relationships and compute complex interactions among and within various systems that influence one another. As much as any other technological advance, they extend the powers of human cognition and perception.

At the same time, model-based prediction also carries limitations and presents risks increasingly likely to be present the more we proceed in a direction that might convince us that we are actually getting somewhere. If so, enhanced awareness extended beyond the narrow purview of such simulation systems themselves might prevent us from falling into potential traps.

Predictive analytics necessarily take a narrow view of what constitutes models and what purpose they serve. The term *model* usually suggests a representation of reality in an artifact that, although limited, nevertheless seeks as much as possible (given resources and the specific analytical requirements being served) to portray the structure of the underlying true model of the system in question. Thus, the validity of such a model is determined by its predictive power, as in physical science or engineering, or postdictive capacity to explain variables' time series as in social science. The model not only becomes the central feature of the analysis, but alternative uses of models for nonpredictive exploratory or explanatory purposes are disparaged or ignored. The analytical enterprise becomes one of seeking to reduce uncertainty and verifying which of the many different sets of assumptions about causal relationships or the values of future model inputs can be shown most likely to prove true. The output from this research path focuses on the likely future states of the systems in question and their component elements; therefore, the path only indirectly touches on those questions raised by policy planners or decisionmakers.

This is not to gainsay the power of predictive analytics or disparage its achievements. But we need to recognize the point beyond which this method is no longer appropriate to the purpose when applied solely on its own. Unfortunately, the breakdown occurs precisely when it is most needed—when alternative choices for short-term action are many, systems are complex, and the uncertainties present are difficult to characterize probabilistically owing to either a deficit of information or fundamental disagreements about what the data we possess might mean.[5]

All too often, the analytical enterprise of model-based prediction focuses on determining which among a set of assumptions about an unknowable future is the most reliable. This has the potential of introducing a Red Queen's race and requiring increasingly exacting detail to enhance the fidelity of predictions.[6] (If a model generates measures of agricultural output in Ukraine, it is tempting to believe that distinguishing among wheat, rye, and buckwheat while further breaking them into several subclimate zones would increase model validity and forecast accuracy.) Perhaps quantum computing when it arrives can better support this mission creep, but it is unlikely to be resolved in the short or medium term.

Another concern is that the quest for more reliably predictive modeling could create an increasing black box problem. Fewer people become capable of comprehending what the model contains and so more are disenfranchised from supplying meaningful insights, expertise, or critiques. In particular, those assumptions made explicitly at the onset to allow for computational tractability will, over time, become implicit and less apparent even to those inside the group who really understand the model—to say nothing of the policy professionals who constitute the ultimate consumers of the output. Thus, potential points of departure from an unfolding reality may fail to be perceived.

The final concern is one of dependency and a false sense of confidence. The unknown raises anxiety. The greater the extent to which we feel in control over what the future may bring, the better we can anticipate it and the more confidence we feel. Predictive models can inculcate an illusion of control. Like a witch or sorcerer from a fairy tale who can exert

[5] The term *deep uncertainty* may refer to conditions under which we do not know (or cannot agree on) how best to characterize uncertainty about future values of key variables (parametric uncertainty) or the nature of causal mechanisms (structural uncertainty). Added to this is a normative component in which we do not know or cannot agree on the appropriate criteria, limit values, or priorities for assessing how well outcomes perform in achieving our goals.

[6] Derived from Lewis Carroll's *Through the Looking Glass*, a *Red Queen's race* is one in which "it takes all the running you can do to keep in the same place"—furious activity leads to no real advancement.

dominance by learning the true name of a thing or character, we feel that the more planners can learn the future's true name—that is, chart out estimates of the likelihoods of different outcomes—the greater the unspoken sense of control planners are likely to feel. We ease our anxiety as individuals and as professional planning teams at the potential cost of amplifying the opening for, and consequences of, surprise—precisely the opposite of the intended benefit from the resources devoted to predictive modeling of deeply uncertain decision spaces.

Automated and Expert Systems

The second path for bringing computing into planning and decisionmaking is a logical extension of trends elsewhere, although the path is perhaps less applicable to the planning problem in the short term. That path is to enhance our reliance on machine learning (ML) and the various approaches to artificial intelligence (AI). The machine becomes the decision system itself rather than supporting a human-moderated process. This path is far less advanced than that of enhanced prediction and remains more of a prospect than a tangible alternative. Expert systems do exist for remote medical diagnostic screenings and robotic surgeries too delicate for a human surgeon to perform reliably or achieve satisfactory outcomes. Expert systems also exist for assessing visual data. Journalism algorithms already generate business reports or sports coverage. But there is also increasing interest in using such autonomous systems in the military; if perfected, such expert systems also could conceivably be useful applications in national security policy planning.

As with predictive modeling, it is possible that future advances in computing will make this an increasingly tractable and accessible alternative. But this is unlikely to be the case in the foreseeable future. Therefore, this path will also have shortcomings that planners will need to be aware of. The first is that AI-based systems have substantial black box aspects. The entire purpose of a recursively trained system is to develop system-generated algorithms permitting it to achieve reliably positive outcomes against an increasingly complex set of challenges or indicators and to do so in a fraction of the time that humans would require. But by their very nature, such algorithms may be difficult for observers to fathom or document. Thus, it will require ceding a measure of human control. Added to this are the well-documented cases of algorithmic bias in which the repetitive reinforcement of algorithmic assessments might incline the system toward inferences and thus solutions based on unintended or even undesirable foundations. Researchers in this field have also found that unintended features in the data sets themselves can reinforce the tendency toward bias. Finally, there is a substantial ethical dimension in relying principally on the AI system's expertise, depending on the decisions being generated. Adding such automaticity to weapon or decision systems would bring us into a new world with uncertain prospects.

Adaptive Planning

The last of the three paths is the one we explore in the balance of this chapter: a human-in-the-loop approach to computer-based analytical support that creates adaptive planning systems within which human deliberation is supported by iterative analyses based on and feeding back into those deliberations. The approach represents a division of labor between computers and models doing what they are best at—tracking connections and generating cases based on data inputs, models and assumptions—and people doing what they are best at—discerning patterns, drawing inferences and, above all, posing more and different questions.

Adaptiveness: The Key for Moving Toward Decision Support

Before further considering the technical tools that may help fulfill the needs discussed in the first section of this chapter, it is worthwhile stepping back and looking at a critical juncture within the NSE. By focusing on it, we can motivate an approach toward change across the fullest set of NSE processes.

Making the Intelligence Community More Adaptive

A characteristic recent debate in the NSE has been how to improve the ability of the IC component to provide the policy community with reliable information (and the analyses used to convey information) when confronting the unknown. More generally, we can refer to this as the problem of knowledge-gathering and characterization—a function largely but not exclusively in the IC. This task is considered the IC's lane in the ideal—that of pure intelligence-gathering and exposition as opposed to the functions of policy planning, deliberation, and implementation. If the lines are blurrier in practice, there is nonetheless a wall intended to exist between intelligence and its knowledge formation activity and the policy deliberation process. The IC's straying beyond its limits would be viewed as running the risk of corrupting both the intelligence-gathering and policy deliberation processes—that is, politicizing the IC.

This ideal conception of the IC's role as focused solely on knowledge occasionally places it in the roles of either making forecasts or filling in blanks. At the same time, the prediction game has grown increasingly perilous, as shown by the IC's lapses in providing early warning of the Soviet Union's collapse; the Arab Spring; the September 11, 2001, terrorist attacks; and Iran's Islamic revolution. Perhaps the conception of the IC's modern role should be a bit different given the acceleration (along with decreased predictability) of change. Perhaps it always should have been so.

The IC must always be looking beyond the horizon. But for what purpose and as measured by what standards of performance? The same IC machinery used to gather information and frame analyses could be put toward endeavors with enhanced potential value within the NSE: providing analytical support to policy decisions. This line of effort would not be providing policy advice, but it would go beyond situation reporting as an end deliverable. This enhanced knowledge project can enhance policy actors' understanding of available COAs and the potential implications of those COAs across different plausible futures. And in the context of "wicked" problems with a plethora of variables—many hard or impossible to characterize by probabilities—the task would be to provide analytically informed assessments of which variables should weigh most prominently in the decision process so that short-term actions can closely conform with long-term policy objectives across many plausible futures.

This means a shift from trying to answer the question of "What will happen?" when it is objectively difficult to do so—that is, analysis as means to resolve which among the clashing assumptions will prove true in the future. Instead, the decision support focus would seek to provide input to illuminate questions more recognizably useful to the policy decision process:

- What future possibilities might affect the ability to achieve policy goals?
- How fragile is the current or intended strategy or COA to such changes?
- How could we modify our strategy to reduce exposure to such vulnerabilities?

This shift in posture and purpose for the data collection and analysis functions within the NSE would not require corresponding shifts in personnel, training, or modeling infrastructure, nor would it detract from reporting requirements. Rather, the same machinery can be leveraged to serve an enhanced rationale—one that is consonant with the historic tenets and purposes of IC activities but perhaps more suited to shifting and uncertain times. It would contribute to making the NSE more adaptive in two senses of the word. First, it would be better tuned and potentially responsive to the pace of external changes. Second, it could enable a posture for policy planning and deliberation that would be tuned to the need to strive for and embody robustness within the design of plans.

Moving Away from Optimality Toward Robustness

The term *robust* carries several denotations. Most relevant for this discussion is to define a plan as robust if it is one that *performs well, compared with the alternatives, over a wide variety of plausible futures*.[7] Robustness in this sense is a comparative quality. It emphasizes the decision support over the neutral knowledge-gathering aspect of the NSE. Robustness focuses on comparative advantages, disadvantages, and trade-offs among alternatives. A robust strategy might not be an optimizing strategy based on a specific set of presumed circumstances. It is more likely to do *well enough* across many plausible future circumstances than do as well as possible in many of them. But to the extent that it fails altogether in meeting objectives, it is likely to fail more *gracefully* (i.e., with less dire consequences) than might an optimizing strategy when it finds itself confronting similarly invidious circumstances compared with those it had been designed to operate within.

Performing well suggests a trade-off between meeting policy objectives and performing satisfactorily in many different futures. It also suggests the explicit establishment of criteria for assessing trade-offs. National security decisionmaking rarely has a single bottom line suf-

[7] Jonathan Rosenhead, Martin Elton, and Shiv K. Gupta, "Robustness and Optimality as Criteria for Strategic Decisions," *Journal of the Operational Research Society*, Vol. 23, No. 4, December 1, 1972; Robert J. Lempert, Steven W. Popper, and Steven C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative Long-Term Policy Analysis*, Santa Monica, Calif.: RAND Corporation, MR-1626-RPC, 2003.

ficient to stand in for all interests. For one thing, actions designed to achieve a short-term goal (e.g., disrupt adversary preparations) could be injurious to long-term objectives (e.g., returning adversary behavior to international norms). For another, any action has costs, whether weighed in terms of dollars, administrative attention, or political capital. And any action may well have direct or indirect influence on other interests. For all these reasons and more, the NSE decision process is a multi-attribute valuation problem requiring balance and nuance. In this context, performing well means meeting the several criteria established by policy leadership. This factor alone would mean leaving optimization behind and instead conducting analyses on a basis that comes closer to the satisficing approach that Herbert A. Simon suggested more realistically approximates the behavior of decision leaders within organizations: Their job is to find positions they judge to be well hedged between opportunity and threat.[8]

Adaptiveness is a characteristic that enhances a plan's robustness over a variety of plausible futures. Though this seems to be a truism, it too carries within it a subtle transformation in the conception of decisionmaking within the NSE. The professionals who staff these agencies are skilled and committed to excellent performance. They will naturally consider what might be Plan B if the future evolves deleteriously. But the adaptation will be situational, not a built-in characteristic of a truly adaptive plan that may be formally laid down as a series of "if–then" statements. A plan designed to be adaptive from the time of initial implementation will have determined in advance what signals should be observed to gain early warning of conditions as they develop and, based on those developments, adapt by shifting from one policy course to another that will have been previously judged to accord better with the emergent realities. Explicitly building capacity for adaptiveness may entail revising conceptions of what constitutes a plan and the analytical support required to both design it and expose it to policy deliberation.

Robustness and the capacity for adaptation that often is its motive force might be values used to distinguish and select among policy alternatives, but they are not ends in themselves. Robustness is a comparative value; there is need for a comparative yardstick that still centers on the goals set by policy leadership. So rather than measuring in absolute values, it is useful to employ the concept of regret to weigh choices. The regret of a robust strategy under a specific set of conditions is the difference between its performance along one or more attribute scales and that of a strategy that would have been optimized for those conditions. By definition, the latter strategy would have zero regret. The same is true for any two or more candidate robust strategies. By definition, one would have zero regret—among the available choices, that is the one that would do the best. Each of the others, if they did not perform the same according to the selected measure, would have some level of regret. This allows assessment of policy choices across many plausible futures and the several attributes for measurement that will have been selected by the decisionmakers.[9]

---

[8] Herbert A. Simon, "Rational Choice and the Structure of the Environment," *Psychological Review*, Vol. 63, No. 2, 1956.

[9] Note that a zero-regret strategy does not necessarily ensure a good outcome. That is, although the rank order of preference among available alternatives may lead to a preferred candidate relatively, it may be that

The remainder of this chapter will address the mechanisms for putting such decision support machinery in place to support analyses, policies, and decisions within the NSE.

The Technology of Complexity: Tooling for Robust Decision Support

NSE decisionmakers increasingly face growing uncertainties, dynamic links among complex systems, contention over assumptions and perceptions, divergent interests, and need for coordinated action even if consensus on trends and futures is elusive. Deep uncertainty exists when it is not possible to predict—or agree on—the probable values of future factors, competing models of causation cannot be rejected with the available evidence, or normative agreement on how to assess outcomes as successes or failures is contentious.[10]

There is an emerging field of theory, methods, and tools to provide analytical support for decisionmaking under deep uncertainty (DMDU).[11] DMDU methods share the general characteristic of shifting focus from optimization to instead seeking solutions that achieve robustness—the ability for a COA to perform well compared with the alternatives across plausible futures.[12] In particular, DMDU methods systematically explore and define adaptive rules and pathways to achieve robustness.

The hard challenge lies in the nontrivial effort of adapting the methods and conjectures of this nascent field and combining them with others in innovative ways to provide the envisioned capabilities to the NSE. Most renderings of the Heilmeier Catechism[13] used by the Defense Advanced Research Projects Agency to explore the merits of prospective research programs would ask, among other questions, what are the questions of interest that decisionmakers (for the purposes of this discussion, NSE decisionmakers) need to address, and what are the challenges to current methods for doing so? The prior portions of this chapter addressed these two points. The balance of this chapter is designed to answer the Heilmeier questions: What would a new approach look like, and how could a new program help to improve capabilities?

Although there is no example of a complete structure for robust decision–based analytical support within the existing NSE, we walk through how such a process might be conceived and conducted in the next section. The focus is on the capabilities of a reconceived decision process rather than the details of actual practices within the offices and agencies

all such candidates lead to disagreeable results in absolute terms when compared with policy objectives.

10 Decision Making Under Deep Uncertainty Society, "DMDU Society," webpage, undated.

11 Vincent A. W. J. Marchau, Warren E. Walker, Pieter J. T. M. Bloemen, and Steven W. Popper, eds., *Decision Making Under Deep Uncertainty: From Theory to Practice*, Cham, Switzerland: Springer, 2019.

12 Rosenhead, Elton, and Gupta, 1972.

13 Thomas Kalil, "Planning for US Science Policy in 2009," *Nature*, Vol. 443, No. 7113, October 2006.

that play roles within the NSE. Innovation would be required on the human and organizational aspects of implementation and process engineering, at least as much as for the technical aspects, for effective transitioning of this technology. Understanding how to do the least injury to existing working and reporting practices and being tuned to the differing formats of argumentation used within and between different offices (e.g., well-constructed narratives versus graphics or numerical tables) would be important values.

What follows is a high-level sketch of one possible robust decision–based architecture within the NSE. At several points, we will reference examples drawn from different subject areas. These serve solely to illustrate and provide more detail on process. The methods discussed are intended neither to be comprehensive nor exclusive of the possibility of other techniques. Because of this chapter's focus on entry points for incorporating advanced computing into NSE decisionmaking, we will base the discussion on the Robust Decision Making (RDM) method for creating human-in-the-loop adaptive reasoning systems.[14]

Deliberation with Robust Decisions Analytical Support

A National Research Council panel concluded that an effective approach to "wicked" problems,[15] such as those typified by climate policy but also manifestly present in the NSE, would be a process of "deliberation with analysis" that is an iterative interaction between the deliberating body and the analysts seeking to provide support to them.[16]

Figure 11.1 presents a concept for such an approach as part of an NSE policy deliberation process. In this conceptualization, the core engine for analysis is based on RDM principles portrayed in the center of the figure.

The model-based RDM method is not intended to produce better predictions but instead uses quantitative models and data to inform better decisions.[17] Although traditional analytical practice may focus on trying to derive consensus on assumptions about future states of the world and outcomes, RDM instead seeks to derive consensus on decisions—even when there may be considerable disagreement over assumptions or differing interests (Figure 11.2).[18] It does so through successive iterations of proposing different and successively more sophisti-

---

[14] Robert J. Lempert, "Robust Decision Making (RDM)," in Vincent A. W. J. Marchau, Warren E. Walker, Pieter J. T. M. Bloemen, and Steven W. Popper, *Decision Making Under Deep Uncertainty: From Theory to Practice*, Cham: Springer International Publishing, 2019.

[15] Horst W. J. Rittel and Melvin M. Webber, "Dilemmas in a General Theory of Planning," *Policy Sciences*, Vol. 4, No. 2, June 1, 1973.

[16] National Research Council, *Understanding Risk: Informing Decisions in a Democratic Society*, Washington, D.C.: The National Academies Press, 1996.

[17] Steven W. Popper, Robert J. Lempert, and Steven C. Bankes, "Shaping the Future," *Scientific American*, Vol. 292, No. 4, April 1, 2005.

[18] Nidhi Kalra, Stéphane Hallegatte, Robert Lempert, Casey Brown, Adrian Fozzard, Stuart Gill, and Ankur Shah, *Agreeing on Robust Decisions: New Processes for Decision Making Under Deep Uncertainty*, Washington, D.C.: World Bank, June 2014.

FIGURE 11.1
NSE Strategic Deliberation Supported by DMDU Robust Decisions Analysis

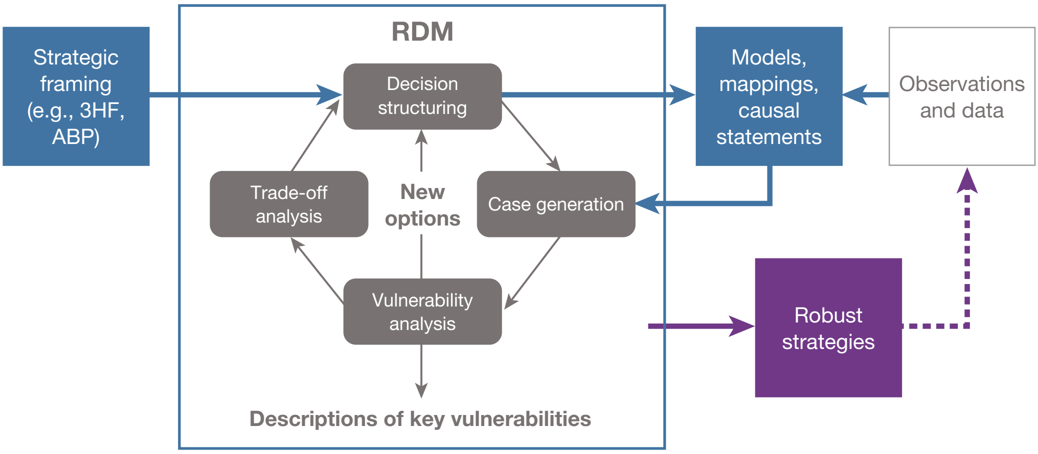

NOTE: 3HF = Three Horizon Foresight; ABP = assumption-based planning.

FIGURE 11.2
Contrasting the Analytical Strategies of Optimization (top) and RDM (bottom)

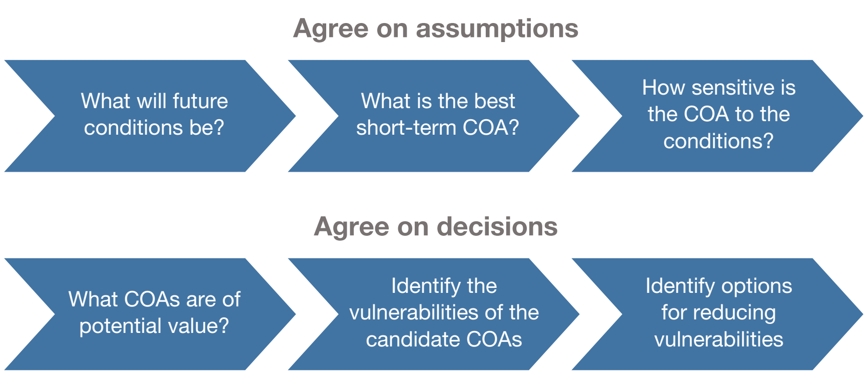

cated strategies, identifying vulnerabilities by exploring the effects on outcomes over varying assumptions, and then seeking options for reducing the revealed vulnerabilities.

Collaborative Framing of Objectives and Strategic Concepts

No serious discussion of strategy can occur without specifying the end objectives for which strategies are just the means. The parties to an NSE deliberation may be prepared to enter the decision structuring step immediately (center top of Figure 11.1). But particularly in an inter-agency process, specifying those goals may be difficult. Therefore, the left of Figure 11.1 also places an opportunity for beginning with a broader initial framing. A focus on robustness and adaptivity could usefully begin here. The effort could provide value in itself even if a more formal decision analysis is not feasible. Several nonquantitative DMDU methods can provide support. We will discuss two methods operated in parallel.

Standing in the present and looking forward is difficult, confusing, and contentious. It may prove easier for a mixed group to instead place itself in an explicitly defined future and look back. 3HF is a foresight technique that has been applied as a group tool for defining shared vision.[19] The focus is on establishing in some detail the characteristics of a desirable future (3rd Horizon) through a structured, qualitative process, characterizing the present (1st Horizon) condition in similar terms, contrasting the two, and identifying trends, obstacles, or trade-offs that might prevent the ideal 3rd Horizon condition from being realized. The heart of the exercise is to then identify and contrast alternative pathways for crossing the intervening period (2nd Horizon)—the “zone of conflict” in Curry and Hodgson’s parlance.[20] These different pathways or COAs can be viewed as transition trajectories corresponding to different candidate strategic concepts.

In extending the 3HF output, ABP may be used to examine different alternative COAs[21]—each based on one or more strategic concepts—to effect the transition across the 2nd Horizon.[22] ABP is a narrative technique originally intended as a way to discover important but implicit and potentially vulnerable assumptions within plans. When used at the outset of a policy deliberation, ABP framing can deconflict the more-usual advocacy behaviors found in policy discourse and enhance discussion of alternative choices.

ABP does so by comparing COAs side by side to elucidate for each their explicit and implicit *load-bearing assumptions*—those assumptions that, were they found to be invalid in the future, would then call into question the value of that COA as a vehicle for bringing about desirable outcomes. During this process, COAs may be modified or hybridized to shore up revealed weaknesses. After the winnowing process, each COA is then assessed for what *signals* might give early warning of impending vulnerabilities and what ancillary actions may be taken to either *shape circumstances* to be more conducive to the COA or *hedge* against its

[19] For more on 3HF, see Andrew Curry and Anthony Hodgson, “Seeing in Multiple Horizons: Connecting Futures to Strategy,” *Journal of Futures Studies*, Vol. 13, No. 1, August 2008; and Richard Silberglitt, Brian G. Chow, John S. Hollywood, Dulani Woods, Mikhail Zaydman, and Brian A. Jackson, *Visions of Law Enforcement Technology in the Period 2024–2034: Report of the Law Enforcement Futuring Workshop*, Santa Monica, Calif.: RAND Corporation, RR-908-NIJ, 2015.

[20] Curry and Hodgson, 2008. In the report, this is referred to as the “triangle” of conflict. It is conflictual because not only are the systems necessary to sustain the prospective 3rd Horizon coming into being, but also those of the 1st Horizon are resisting their loss of incumbency while the needs and goals of the 2nd Horizon time period also require contemporary support.

[21] James A. Dewar, Carl H. Builder, William M. Hix, and Morlie H. Levin, *Assumption-Based Planning: A Planning Tool for Very Uncertain Times*, Santa Monica, Calif.: RAND Corporation, MR-114-A, 1993; and James A. Dewar, *Assumption-Based Planning: A Tool for Reducing Avoidable Surprises*, New York: Cambridge University Press, 2002.

[22] For a discussion of this method, see Appendix A to Liisa Ecola, Steven W. Popper, Richard Silberglitt, and Laura Fraade-Blanar, *The Road to Zero: A Vision for Achieving Zero Roadway Deaths by 2050*, Santa Monica, Calif.: RAND Corporation, RR-2333-NSC, 2018.

potential failure. Figure 11.3 shows the idealized flow of an ABP process configured within a 3HF framing.

As the world changes, the systems and concepts well suited to current conditions (solid black line) may work increasingly less well in the future unless they also change. If they did so, they would be better suited to achieving objectives in the 2nd Horizon time frame (dashed black line). Because of fundamental uncertainty and despite what we may intend, the systems put in place to sustain the 3rd Horizon state may be more or less well suited to achieve the planning group's vision (dotted lines). Choice of COA, early warning, and hedging and shaping actions can affect actual outcomes.

Main Steps of RDM Process

Returning now to Figure 11.1, the main steps of an RDM policy analysis are seen in the center of the figure. The first and most crucial is for the parties to the deliberation, whether internal or cross-agency, to determine the decision structure framing the analysis. This effectively places the decisionmakers within the analytical process and so enhances focus on the decisions that require analytical support. This beginning at the "wrong end of the telescope" differs from more traditional approaches that might first seek to detail the system of interest without initial reference to the decision aspect. The latter would provide an objective view

FIGURE 11.3
Combining 3HF with ABP Weighs Both Future Visions and Alternative Strategic Pathways

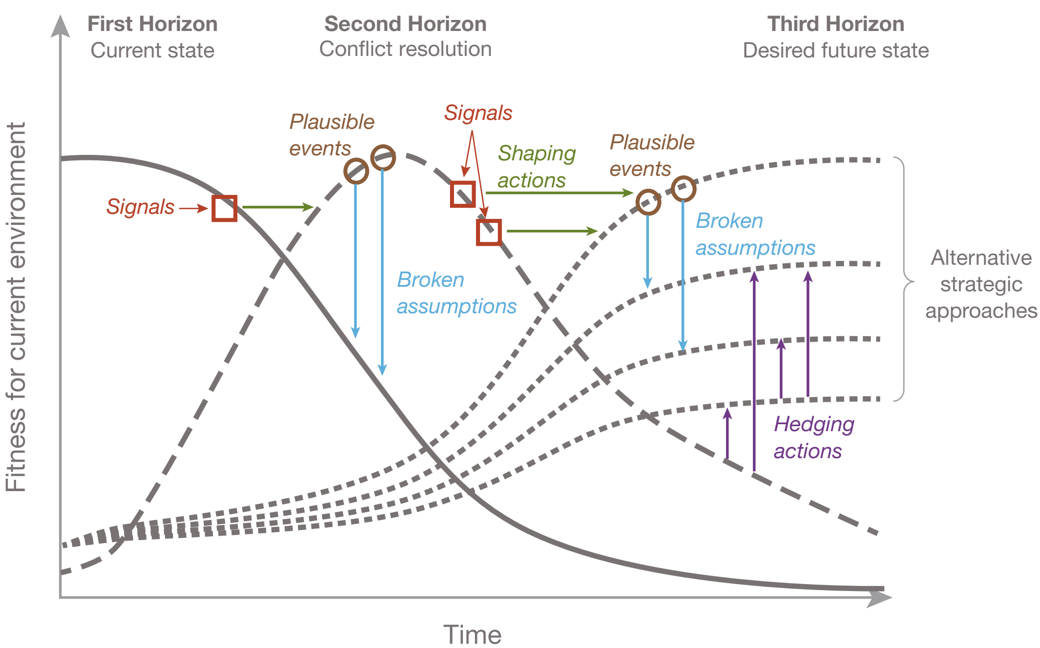

SOURCE: Ecola et al., 2018, Figure A.3.

that might leave to the planners the task of determining what the results may imply for the decisions at hand. The former brings the decision question into the heart of the analysis.

The decision structuring step gains its power through reviewing the factors of importance to a problem and placing them in one of four functional categories that will be explored in an RDM analysis as shown in the text box.[23]

The framing in the text box aids decision analysis in several ways. It provides a design specification for the model software in the RDM analysis. For example, the assignment of factors to the category of external uncertainties (X) versus that of policy levers (L) (which may be assembled as building blocks into a variety of candidate COAs) brings the actual policy process into the analysis; what is an external uncertainty to one group of policy actors is precisely the sphere of action of another. This reinforces a means and ends framing for the analysis versus detailed modeling of an entire system only to later find that only portions of that system are relevant to selecting among COAs (as opposed to predicting future outcomes.) This makes the modeling component more parsimonious of modeling resources.

Beyond the uncertainty that comes from not knowing values of important future factors, there is also structural uncertainty when there are alternative beliefs regarding causal relationships (R). This type of uncertainty means that two trained professionals, such as former United Nations Ambassadors John Bolton and Samantha Powers, can consider the same body of evidence and draw different policy conclusions, in part, because of different conceptions of underlying causal relationships. Finally, it is important to place exploration of measures of outcomes (M) at the same level of the other categories. Rather than just being the sausage that emerges from the NSE factory, policy choices will need to be made on the basis of trade-offs among prospective gains and losses. Thus, determining what set of outcomes will be deemed successful is also an active process of exploration. Creating the Table 11.1 structure at the onset provides a useful common vocabulary across disciplines, professions, and backgrounds to support deliberations and decisions over policy.

Categories of Factors to Be Explored in RDM

X: External Uncertainties (assumptions)	L: Policy Levers (choices)
Uncertain factors beyond control of planners that may affect their ability to reach goals	Actions that may combine into different alternative COAs in pursuit of goals
R: Relationships (models)	**M: Performance Metrics (outcomes)**
Alternative specifications of causal relationships among metrics, levers, and uncertainties	Metrics and associated satisficing criteria that reflect decisionmakers' goals

[23] Lempert, Popper, and Bankes, 2003.

The decision structuring process lays out the elements comprising the experimental design for the analytical machinery used to carry out the next RDM step, case generation (Figure 11.1). This supports the exploratory modeling at the heart of the method.[24] Instead of focusing on a small set of scenarios, exploratory modeling reasons over large ensembles of cases generated by simulating COA outcomes against several assumptions. Thus, the points of view of all sides are considered in the analysis, precluding the need for prior agreement on assumptions. The result of this step is a set of outcomes from the pairwise simulation of one COA (a defined set of policy levers [L]) played out across a selected test bed of alternative futures defined by different assumptions about the uncertainties (X) for all the COAs under study.

The resulting information is then evaluated in the discover vulnerabilities step. RDM invites an iterative process of discovery, reframing of questions, and COA refinement as indicated by the arrows showing recursive flow in the Figure 11.1 RDM box. The planning team and analysts draw inferences from reasoning over the ensemble of cases and pose new queries. Uncertainties are not characterized by assumed probabilities but are instead characterized by what information they convey about how to decide among alternative COAs. The analysis will not predict actual future outcomes. It rather provides better understanding of available alternative COAs and the potential implications of each and identifies those criteria on which a decision should be based to enhance robustness.

In particular, scenario discovery is a process of determining COA vulnerabilities. Because of the limitations of unaided human perception, such ML-based algorithms as the Patient Rule Induction Method (PRIM)[25] or Classification and Regression Tree (CART)[26]—perhaps combined with a principal component analysis[27]—are used to search for systematic successes or failures of COAs across many cases representing different assumptions about future states of the world. In discovering what is common across these cases, scenario discovery in effect proposes lower-dimension sets of factors—scenarios—that appear common across a class of outcomes. Meaningful scenarios are thus generated analytically rather than selected *ex ante*. These, in turn, convey important findings to the planning group, allowing them to reevaluate choices and better understand vulnerabilities among candidate COAs.

The implications of this capability for opening the aperture of NSE policy deliberation are potentially profound. Rather than the less tractable and often troubling question of "What will happen?" a new—and more operationally useful—one takes its place: "What would we need to assume or believe will be true to recommend selecting COA 1 instead of COA 2?"

[24] Steven C. Bankes, "Exploratory Modeling for Policy Analysis," *Operations Research*, Vol. 41, No. 3, June 1, 1993.

[25] Jerome H. Friedman and Nicholas I. Fisher, "Bump Hunting in High-Dimensional Data," *Statistics and Computing*, Vol. 9, No. 2, April 1, 1999.

[26] Leo Breiman, Jerome H. Friedman, Richard A. Olshen, and Charles J. Stone, *Classification and Regression Trees*, Boca Raton, Fla.: Routledge, 1984.

[27] Siddhartha Dalal, B. Han, Robert Lempert, A. Jaycocks, and A. Hackbarth, "Improving Scenario Discovery Using Orthogonal Rotations," *Environmental Modelling & Software*, Vol. 48, October 1, 2013.

This crucially changes the focus from outcome prediction to illumination of choices among alternatives and supports decisionmakers in hedging against risk.

Applying the prior steps (often several times) provides data for trade-off analyses to compare candidate COAs. Because the goal is not to forecast outcomes but to compare the robustness of alternative plans for meeting defined policy objectives across different future states of the world, the key is the concept of regret: How much would we regret (in terms of measurable objective value forgone) having chosen a particular COA compared with the alternative COA that would have been optimal for that set of conditions? When trade-off analysis occurs, it is likely that prior iterations of COA modification will have reduced many potential vulnerabilities within the remaining modified candidate COAs. Therefore, we are not interested (from the perspective of policy decisions) in any remaining uncertainties and vulnerabilities that do not change the preference rankings among alternatives. Rather, we are now able to identify and focus on those variations in possible external circumstances that would change that order of preference ranking.

Figure 11.4 illustrates the comparison of four candidate COAs.[28] Having discovered a stressful scenario in the prior step, the values being tested against are different assumptions

FIGURE 11.4

Trade-Offs in Regret Performance of Four COAs with Variations in Odds of Stressful Scenario

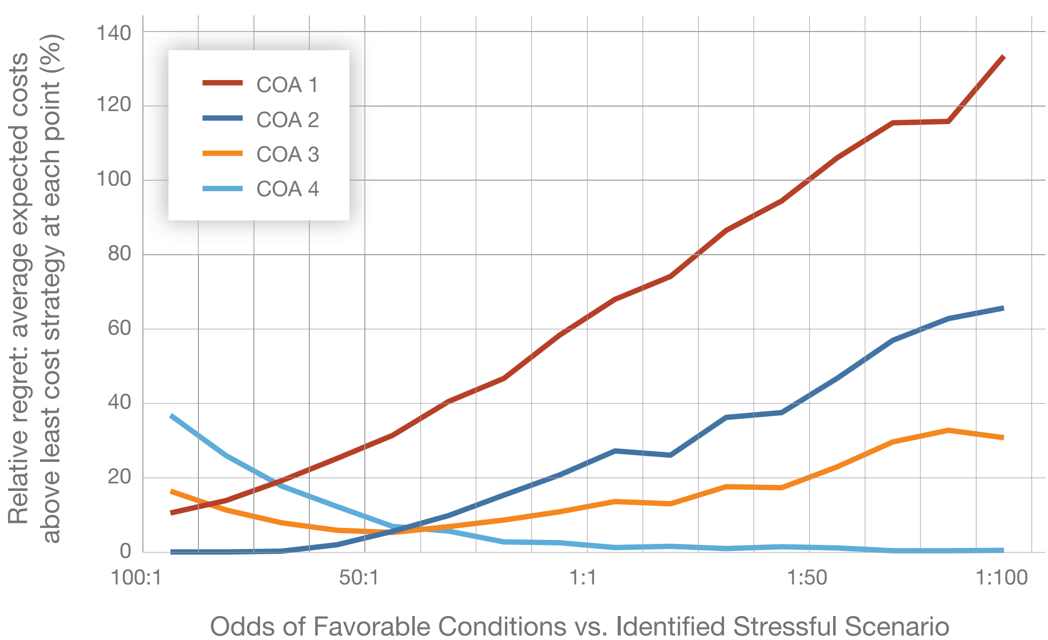

[28] This figure was adapted from Steven W. Popper, Claude Berrebi, James Griffin, Thomas Light, Endy M. Daehner, and Keith Crane, *Natural Gas and Israel's Energy Future: Near-Term Decisions from a Strategic Perspective*, Santa Monica, Calif.: RAND Corporation, MG-927-YSNFF, 2009.

of the odds of that scenario occurring. COA 1's expected outcome regret is higher at almost every point than all the others. COA 2 performs well when the assumed odds for the stressful scenario are low, while COA 4 does better in situations in which stressful conditions might be more expected. COA 3 would probably not be considered in an analysis that was keyed toward optimization or that looked only at two to three scenarios. It also is dominated at most points. However, perhaps it represents a hybrid version of COAs 2 and 4. It shows the characteristic of failing gracefully: Across the unknowns, its regret measure is at most points second best. Its performance appears robust to uncertainty as measured by the ratio shown on the horizontal axis.[29]

Output to Determine a Robust Strategy

The end process to RDM provides decisionmakers with strategy choices selected for robustness (the right side of Figure 11.1). An output, such as the one shown in Figure 11.4, presents several advantages to a senior decision group. It does not recommend policy but rather illustrates the basis for choice that would enhance the prospects for short-term actions to meet long-term goals across a rugged future landscape. It is an observational instrument rendering visible what was previously hard to perceive, in the same manner as a microscope or telescope. Senior leaders may bring their perspectives to what potential futures they find most worrying or credible or may bring additional external information to bear in selecting among choices. And to the extent there is true uncertainty with little prospect of assigning probabilities to future values of the uncertainties that RDM analysis shows affect preferences, the method discloses a robust candidate strategy (COA 3) that would form a hedged position.

Probabilities have not been brought into the analysis proper prior to the trade-off step, nor has the generation of the set of plausible future values that form the test screen for plan performance been done on a random basis.[30] This means that the analysis has a drill-down capacity; senior leadership might ask to see individual cases and outcomes that lead to the curves shown in Figure 11.4. It is possible to understand the cause and effect at play in each such case or set of cases. It is possible to see which changed assumptions might affect plan success or failure. Most importantly, graphics such as Figure 11.4 provide a basis for exchanging views, collectively recognizing which areas in a world of scarce intelligence or analytical

---

[29] We could alternatively have placed along the horizontal axis different assumed values for one of the key uncertain factors that vulnerability analysis will have disclosed does change preference rankings among choices—that is, as in Figure 11.4, regret values that cross over rather than maintain a consistent parallel valuation among COAs. (For an example of this approach, see Popper et al., 2009.)

[30] The Monte Carlo method generates an experimental design stochastically. Therefore, RDM instead frequently uses a Latin Hypercube experimental design which selects the desired number of test cases (100, 1,000 or 5,000,000) uniformly across all dimensions represented by the uncertain variables. This selection is more reproducible. Three such uncertainties would see points selected from within a cube; *N* uncertainties would be drawn from an *N*-dimensional hypercube. See M. D. McKay, R. J. Beckman, and W. J. Conover, "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code," *Technometrics*, Vol. 21, No. 2, 1979.

RDM in the Absence of Formal Models?

The NSE is broad and encompasses many issues that may be analyzed through formal computer modeling; these issues include the transition toward a modernized nuclear triad architecture while dealing with nuclear-armed adversaries, reduction of the vulnerability of the domestic defense industrial base to different types of shock or disruption, and strategies of assistance for developing countries' transition to non–fossil fuel energy. But there are also principal preoccupations in geopolitics and diplomacy for which no formal models exist and for which modeling is not part of analysts' training.[1] To what extent may the RDM process shown in Figure 11.1 be applied in such settings?

The word *model* usually implies a quantitative formalization. This fails to capture the concept's full potential.[2] In RDM, a model is regarded less as a detailed representation of reality than as a means for generating simulation outcomes consistent with known facts and relationships. This changes the model's role and cedes its traditional centrality in quantitative analysis to become instead subordinate to the main task of supporting deliberation and decision. What is required in RDM is an explicit, systematic mapping of cause and effect from applied action to the resulting outcomes within a setting described by the values assigned to environmental factors. In this sense, a model may consist of a set of explicit causal statements describing presumed relationships between inputs and outputs.

In this light, the world of foreign policy and diplomacy is rife with sophisticated, complicated models. However, they remain implicit, defined by the knowledge and accumulated experience of individuals, and not codified as formalized statements. A model sufficient to support RDM reasoning may be created through formal elicitation from subject-matter experts (or planners) of causal statements about the system or issue being considered. Differences between individual expressions of these causal constructs may be tested in precisely the same manner as any other exploration of differences within RDM.[3]

Elicited statements of causal relationships would then be compiled into a comprehensive structure. Different individuals might formulate such statements differently, ascribing greater or lesser influence of selected causes on the effects of interest. These differences would then be made explicit and thus testable and capable of serving as a basis for more focused policy reasoning. Within an RDM analytical structure, the potential exists for developing quantitative or ordinal characterization of uncertainties—for example, as

[1] A senior former State Department official once related to the author that the only times during his career that he attended a briefing for which the room's lighting was dimmed so that overhead or projected slides could be shown was when he was visiting in the Pentagon.

[2] Paul K. Davis and Steven W. Popper, "Confronting Model Uncertainty in Policy Analysis for Complex Systems: What Policymakers Should Demand, " *Journal on Policy and Complex Systems*, Vol. 5, No. 2, Fall 2019.

[3] For fuller development and application of this approach, see Steven W. Popper, "Robust Decision Making and Scenario Discovery in the Absence of Formal Models," *Futures & Foresight Science*, Vol. 1, No. 3–4, 2019.

formal mathematical models, decision trees operated as Bayesian networks, or a rules-based lexicographical mapping without being forced into the predictive role of assigning them probabilities.

A rebuttal to the argument that such models would be ad hoc and, therefore, potentially (or even necessarily) spurious, is the simple statement that such models exist and are already the implicit basis for argument and disagreement in the NSE deliberation process. The only difference with current practice is to elicit these tacit models, make them explicit, and render them in a form for comparison, testing, and leveraging their power for decision support in an innovative application. For an analysis conducted in this manner, the model is part of the analytical machinery but also an important output to be tuned and reevaluated as performance, outcomes, and data are subsequently gathered and assessed.

Using such a Robust Decision approach to NSE deliberations, even if formal models do not exist, would allow a realization of the ASDA vision inherent in Figure 11.1 of a common analytical process amalgamating knowledge-gathering, analysis, planning, and outcome evaluation into a unitary, operationally focused framing.

resources should be the focus for enhanced effort, and cooperatively finding ways of improving and adapting plans and implementation.

Any implementation of a resulting robust strategy itself has a recursive value in the Figure 11.1 schema. The activity of adaptation from which the property of robustness derives carries implicit within it a need for monitoring, evaluation, and course correction during implementation. The process outlined in Figure 11.1 may be seen as a means for fully implementing an Act-Sense-Decide-Adapt (ASDA) cycle for the NSE.[31] Not only does it formalize a systemic role and method for monitoring (sensing) and adapting functions, but it integrates both more organically with the systems for decision and action that more often receive emphasis.

Foresight activities in any organization can prove to be a precarious value in that looking forward beyond the institution's preoccupations may be considered, consciously or unconsciously, at best as a luxury or indulgence; at worst, it may be considered a drain on the time and resources of hardworking people whose task it is to get the immediate job done. The only way to escape this trap is for the functions of sensing and adaptation—the core of foresight activity—to be perceived by those within the organization as providing value. The Observations and Data loop in the far right of Figure 11.1 suggests that planning, analysis, implementation, and evaluation can be made part of a unified, recursive robust decision–based NSE policy architecture.

[31] The more-familiar Observe-Orient-Decide-Act loop (see John R. Boyd, The Essence of Winning and Losing," re-created briefing, Project on Government Oversight, August 2010) is based on a move-countermove paradigm that yields concepts of operations framed in this context. The ASDA cycle instead envisions operations within a complex adaptive systems paradigm and so prioritizes enhanced awareness to evolving situations and innovating new concepts of operations in response to them (see Justin Kelly and Mike Brennan, "OODA Versus ASDA: Metaphors at War," *Australian Army Journal*, Vol. 6, No. 3, Summer 2009).

Concluding Thoughts

Enhancing the capability for analytical decision support in NSE policy processes requires reexamining the role and orientation of the classic posture toward intelligence-gathering, analysis, and dissemination. Embracing the concept of robustness—with its prospect of allowing for more adaptive NSE policy in the uncertain times that lie ahead—suggests the value of similarly reexamining the optimization-based strategic approach that is implicit in traditional policy analyses. More often than we would care to admit, policies are selected with the unstated assumption of maintaining relevance from the time of their implementation until the time for evaluation arrives. Although few policies are truly *fire and forget* in the NSE policy realm and course correction will be required—particularly when human adversaries are involved—the structure of deliberations is framed in a way that carries little explicit recognition of that fact. Red teaming is an adversarial, post-processing step following plan formation rather than an essential component of its gestation.

The idealized robust decision–based framing of an NSE policy process presented in this chapter takes as one of its postulates that any policy implemented in a sufficiently complex setting, especially one that is adversarial, is in truth a policy experiment. If so, let the policy deliberation process reflect this reality in its fundamental framing. Doing so on a formal basis will be a great challenge. Yet, there is a profound asset buried within today's NSE that bodes well for eventual success: Many of the dedicated, talented, and creative professionals within that establishment make informal attempts as individuals to perform this reframing as best they can. It remains only to develop the structures that would permit these inclinations to become better supported. In addition to the tooling that already exists to support deductive reasoning within the NSE, there is also a need to institute a similar apparatus to assist in the inductive reasoning process ("What if . . . ?") that is its natural counterpart; this is a faculty that individuals routinely use but that has fewer means for expression within an agency hierarchy or inter-agency process.

Methods and application experience for each of the steps in the process in Figure 11.1 already exist. That does not make the goal of applying it in the NSE any less aspirational and challenging. The potential rewards for doing so are many. Recent events have shown that those at work in the NSE need "a prosthesis for the imagination."[32] NSE process and structure must better accommodate a more widely embracing phenomenology than it uses. More than ever, the NSE requires support for systematically exploring and formally employing the means for structured inductive and abductive reasoning through troubled times. The policy deliberation process would benefit from becoming more capable of eliciting different classes of knowledge and experience that might be brought to bear on the challenges of the future. There is a need to create an environment that encourages and can accommodate late-stage inferences and divergences of opinion without risking disruption of the larger policy deliberation process. And above all, the NSE as a whole would benefit from having the ability to

[32] Bankes, 1993.

incorporate the strategy development, analysis, planning, implementation, and monitoring and evaluation functions within an integrated, continuous process, such as in Figure 11.1. Bringing this aspiration into being within the NSE fully meets the criteria for being considered a worthy—"DARPA-hard"—challenge.

Acknowledgments

The author expresses his gratitude to the Acquisition and Technology Policy program of RAND's National Defense Research Institute, the Defense Advanced Research Projects Agency, and the editors of this report, Aaron B. Frank and Elizabeth M. Bartels, for the opportunity to contribute.

Abbreviations

3HF	Three Horizon Foresight
ABP	assumption-based planning
AI	artificial intelligence
ASDA	Act-Sense-Decide-Adapt
COA	course of action
DMDU	decisionmaking under deep uncertainty
DoD	U.S. Department of Defense
IC	Intelligence Community
ISC	integrated security construct
ML	machine learning
NSE	National Security Enterprise
RDM	Robust Decision Making

References

Bankes, Steven C., "Exploratory Modeling for Policy Analysis," *Operations Research*, Vol. 41, No. 3, June 1, 1993, pp. 435–449.

Boyd, John R., "The Essence of Winning and Losing," re-created briefing, Project on Government Oversight, August 2010. As of July 30, 2021:
http://pogoarchives.org/m/dni/john_boyd_compendium/essence_of_winning_losing.pdf

Breiman, Leo, Jerome H. Friedman, Richard A. Olshen, and Charles J. Stone, *Classification and Regression Trees*, Boca Raton, Fla.: Routledge, 1984.

Clark, Burton R., "Organizational Adaptation and Precarious Values: A Case Study," *American Sociological Review*, Vol. 21, No. 3, 1956, pp. 327–336.

Curry, Andrew, and Anthony Hodgson, "Seeing in Multiple Horizons: Connecting Futures to Strategy," *Journal of Futures Studies*, Vol. 13, No. 1, August 2008, pp. 1–20.

Dalal, Siddhartha, B. Han, Robert Lempert, A. Jaycocks, and A. Hackbarth, "Improving Scenario Discovery Using Orthogonal Rotations," *Environmental Modelling & Software*, Vol. 48, October 1, 2013, pp. 49–64.

Davis, Paul K., and Steven W. Popper, "Confronting Model Uncertainty in Policy Analysis for Complex Systems: What Policymakers Should Demand," *Journal on Policy and Complex Systems*, Vol. 5, No. 2, Fall 2019, pp. 181–201.

Decision Making Under Deep Uncertainty Society, "DMDU Society," webpage, undated. As of August 25, 2021:
https://www.deepuncertainty.org/

Dewar, James A., *Assumption-Based Planning: A Tool for Reducing Avoidable Surprises*, New York: Cambridge University Press, 2002.

Dewar, James A., Carl H. Builder, William M. Hix, and Morlie H. Levin, *Assumption-Based Planning: A Planning Tool for Very Uncertain Times*, Santa Monica, Calif.: RAND Corporation, MR-114-A, 1993. As of October 28, 2021:
https://www.rand.org/pubs/monograph_reports/MR114.html

Ecola, Liisa, Steven W. Popper, Richard Silberglitt, and Laura Fraade-Blanar, *The Road to Zero: A Vision for Achieving Zero Roadway Deaths by 2050*, Santa Monica, Calif.: RAND Corporation, RR-2333-NSC, 2018. As of January 26, 2022:
https://www.rand.org/pubs/research_reports/RR2333.html

Friedman, Jerome H., and Nicholas I. Fisher, "Bump Hunting in High-Dimensional Data," *Statistics and Computing*, Vol. 9, No. 2, April 1, 1999, pp. 123–143.

Kalil, Thomas, "Planning for US Science Policy in 2009," *Nature*, Vol. 443, No. 7113, October 2006, pp. 751–752.

Kalra, Nidhi, Stéphane Hallegatte, Robert Lempert, Casey Brown, Adrian Fozzard, Stuart Gill, and Ankur Shah, *Agreeing on Robust Decisions: New Processes for Decision Making Under Deep Uncertainty*, Washington, D.C.: World Bank, June 2014.

Kelly, Justin, and Mike Brennan, "OODA Versus ASDA: Metaphors at War," *Australian Army Journal*, Vol. 6, No. 3, Summer 2009, pp. 39–52.

Lempert, Robert J., "Robust Decision Making (RDM)," in Vincent A. W. J. Marchau, Warren E. Walker, Pieter J. T. M. Bloemen, and Steven W. Popper, eds., *Decision Making Under Deep Uncertainty: From Theory to Practice*, Cham, Switzerland: Springer International Publishing, 2019, pp. 23–51.

Lempert, Robert J., Steven W. Popper, and Steven C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative Long-Term Policy Analysis*, Santa Monica, Calif.: RAND Corporation, MR-1626-RPC, 2003. As of October 28, 2021:
https://www.rand.org/pubs/monograph_reports/MR1626.html

March, James G., and Herbert A. Simon, *Organizations*, 2nd ed., Cambridge, Mass.: Blackwill, 1993.

Marchau, Vincent A. W. J., Warren E. Walker, Pieter J. T. M. Bloemen, and Steven W. Popper, eds., *Decision Making Under Deep Uncertainty: From Theory to Practice*, Cham, Switzerland: Springer, 2019.

McKay, M. D., R. J. Beckman, and W. J. Conover, "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code," *Technometrics*, Vol. 21, No. 2, 1979, pp. 239–245.

National Research Council, *Understanding Risk: Informing Decisions in a Democratic Society*, Washington, D.C.: The National Academies Press, 1996.

Popper, Steven W., "Robust Decision Making and Scenario Discovery in the Absence of Formal Models," *Futures & Foresight Science*, Vol. 1, No. 3–4, 2019, pp. 1–13.

Popper, Steven W., Claude Berrebi, James Griffin, Thomas Light, Endy M. Daehner, and Keith Crane, *Natural Gas and Israel's Energy Future: Near-Term Decisions from a Strategic Perspective*, Santa Monica, Calif.: RAND Corporation, MG-927-YSNFF, 2009. As of January 26, 2022: https://www.rand.org/pubs/monographs/MG927.html

Popper, Steven W., Robert J. Lempert, and Steven C. Bankes, "Shaping the Future," *Scientific American*, Vol. 292, No. 4, April 1, 2005, pp. 66–71.

Rittel, Horst W. J., and Melvin M. Webber, "Dilemmas in a General Theory of Planning," *Policy Sciences*, Vol. 4, No. 2, June 1, 1973, pp. 155–169.

Rosenhead, Jonathan, Martin Elton, and Shiv K. Gupta, "Robustness and Optimality as Criteria for Strategic Decisions," *Journal of the Operational Research Society*, Vol. 23, No. 4, December 1, 1972, pp. 413–431.

Silberglitt, Richard, Brian G. Chow, John S. Hollywood, Dulani Woods, Mikhail Zaydman, and Brian A. Jackson, *Visions of Law Enforcement Technology in the Period 2024–2034: Report of the Law Enforcement Futuring Workshop*, Santa Monica, Calif.: RAND Corporation, RR-908-NIJ, 2015. As of October 28, 2021:
https://www.rand.org/pubs/research_reports/RR908.html

Simon, Herbert A., "Rational Choice and the Structure of the Environment," *Psychological Review*, Vol. 63, No. 2, 1956, pp. 129–138.

White House, *National Security Strategy of the United States of America*, Washington, D.C., December 2017.