

A new decision sciences for complex systems

Robert J. Lempert*

RAND, 1700 Main Street, Santa Monica, CA 90407

Models of complex systems can capture much useful information but can be difficult to apply to real-world decision-making because the type of information they contain is often inconsistent with that required for traditional decision analysis. New approaches, which use inductive reasoning over large ensembles of computational experiments, now make possible systematic comparison of alternative policy options using models of complex systems. This article describes Computer-Assisted Reasoning, an approach to decision-making under conditions of deep uncertainty that is ideally suited to applying complex systems to policy analysis. The article demonstrates the approach on the policy problem of global climate change, with a particular focus on the role of technology policies in a robust, adaptive strategy for greenhouse gas abatement.

robust adaptive planning | agent-based modeling | deep uncertainty | complex adaptive systems | computer-assisted reasoning

The study of complex systems provides powerful tools and concepts for capturing useful information about the world. For instance, the process of technology innovation and diffusion, crucial to many of today's most important policy decisions, displays emergence and a dynamics shaped by self-referential expectations in the face of imperfect information. Such properties are difficult to represent with traditional analytic approaches but spring naturally from the mathematics of complexity. Unfortunately, it has often proved difficult to apply this mathematics to policy problems. Complex systems often are characterized by uncertainty of a type that strains the traditional methods of decision analysis, vital to the systematic examination of policy alternatives.

Traditional decision analysis rests on key assumptions about the types of information available to the decision-maker. Powerful Bayesian methods produce a ranking of alternative strategies, and in particular, identify the optimum strategy, assuming the decision-maker has a well characterized system model and can represent uncertainty with probability distributions over the input parameters to that model. Such methods have proved extraordinarily useful for many problems because they help structure the extensive information decision-makers often have, offer a systematic, nonbiased treatment, and expose the numerous logical fallacies to which human reasoning is prone (1).

However, we often have information about complex systems different from that assumed by traditional decision analysis. For instance, complex systems often display regions of extreme sensitivity to the particular assumptions, while at the same time exhibit important regularities of macroscopic behavior. Such systems display one example of deep uncertainty, a situation where the system model and the input parameters to the system model are not known or widely agreed on by the stakeholders to the decision. Under such deep uncertainty, the standard tools of decision analysis are difficult to apply and may not accurately represent the goals of decision-makers. For instance, traditional decision analysis cannot easily address the types of adaptive, evolving strategies that decision-makers often employ when confronted by deep uncertainty.

As an alternative to systematic policy analysis, practitioners have often used a "flight simulator" approach to apply the mathematics of complex systems to policy problems. Decision-makers build insight by watching a variety of interesting model runs unfold. Such approaches are characteristic of the narrative, scenario-based methods of decision support which in recent years have come into vogue (2, 3). Scenario-based planning can help individuals recognize that the future may not be an extrapolation of the past and can help groups reach consensus on strategies even when the members cannot agree on the most likely future. However, scenario approaches do not offer two key benefits of quantitative decision analysis, particularly when applied to complex systems. They do not systematically use quantitative information to correct fallacies in human reasoning nor do they provide any systematic way to rank the desirability of alternative policy choices.

A New Decision Sciences for Complex Systems

In recent years, we have been developing new methods for assisting decision-making under conditions of deep uncertainty that are ideally suited to applying complex systems to policy problems. Called Computer-Assisted Reasoning (CAR; ref. 4), and previously, "exploratory modeling" (5), these methods exploit the qualitatively new capabilities of modern computation—pervasive computation, low-cost memory, and powerful, interactive visualizations—to combine the best features of traditional quantitative decision analysis with those of narrative, scenario-based planning. CAR is a general and powerful example of the emerging school of multiscenario simulation methods that are becoming increasingly common in the literature on decision-making under conditions of deep uncertainty (6–9).

CAR rests on the claim that, under conditions of deep uncertainty, an ensemble of plausible models, rather than any single model, best represents the available information about the future. Users can exploit the information in these ensembles with computer visualization and search. Visualizations, often generated by some algorithm or experimental design strategy, help users to generate hypotheses about desirable strategies. Computer search can then help test these hypotheses. CAR combines the human ability to intuit patterns and abstract a big-picture view with the computer's ability to test detailed implications of facts over a huge number of cases. The approach facilitates the assessment of alternative strategies with criteria such as robustness and satisficing rather than optimality. The former are particularly appropriate for situations of deep uncertainty.

Fig. 1 shows a typical flow of computational steps within CAR analysis. We begin by grouping the inputs into the analysis into policy *levers*, the various actions that the decision-maker can take to affect the world, and exogenous *uncertainties*, the major variables outside the control of the decision-maker, that could

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling," held October 4–6, 2001, at the Arnold and Mabel Beckman Center of the National Academies of Science and Engineering in Irvine, CA.

*E-mail: lempert@rand.org.




Fig. 1. Schematic of a CAR analysis.

affect outcomes of interest. By using these uncertainties and levers, we create a large ensemble of scenarios where each scenario is one guess about how the world works and one choice of actions to take. (Note that this is different from the common definition of scenario which is generally one unfolding of the external world, without consideration of the decision-makers' actions.)

This scenario ensemble becomes a computer-based object that we can view and manipulate to help inform our decisions, often employing a series of interactive computer visualizations, either individually or to support group processes. First, we create what we call *Landscapes of Plausible Futures*, visualizations which show a wide range of potential outcomes. Such visualizations help decision-makers to understand better the range of possibilities that confront them, to recognize the key driving forces that may cause the future to head in different directions, and to achieve buy-in to the analysis from stakeholders with different views of the future. Next, we compare the performance of alternative strategies across the landscape of plausible futures using *Robust Regions*, which is related to the policy region analysis of Watson and Buede (10). The Robust Regions help identify key strategies that perform relatively well compared with the alternatives over a wide range of scenarios. Under

conditions of deep uncertainty, decision-makers often seek such robust strategies (11). Often, such strategies are adaptive; that is, they evolve over time in response to new information.

The key to this CAR approach is an inductive, rather than deductive, approach to quantitative reasoning. In traditional decision analysis, users begin with specific assumptions about the system model and the likelihood of alternative input parameters to that model. Once specified, these assumptions propagate through the computer to produce a definitive ranking of alternative strategies. In contrast, CAR emphasizes the computer as an interactive tool to help users form and then examine hypotheses about the best actions to take in those situations poorly described by well known probability distributions and models. In the CAR approach, analysts and decision-makers begin with a wide range of plausible assumptions that capture the available information about a deeply uncertain future. Then, they use computer search and visualization to systematically apply this information to the comparison of policy options. For instance, they can identify strategies that are robust in the face of this deep uncertainty, find the scenarios most stressing for alternative robust strategies, and define any key tradeoffs among strategies.

These methods have been successfully used to illuminate a variety of notoriously intractable problems of strategic choice under deep uncertainty. These include procurement decisions for the U.S. Air Force (12) military strategy (13–15), science and technology planning decisions for Defense Advanced Research Projects Agency (DARPA; ref. 16), robust strategies for responding to global climate change (7, 17–19), long-range financial planning for the public university systems in several states (20), and in product planning, technology investment, and market-entry problems for the private sector (unpublished work). Here, we present an example application of these methods to the problem of climate change.

Fig. 2 shows an agent-based model of technology diffusion used to compare alternative climate change-abatement strategies. Economic agents choose among alternative technologies on the basis of forecasts of cost and performance, which, in turn, are influenced by learning among the agents and potential price decreases caused by increasing returns to scale. The agents have heterogeneous initial expectations about technology performance and heterogeneous preferences for technology cost/performance tradeoffs. The agents' choices influence the level of




Fig. 2. Agent-based model of technology diffusion used to compare alternative climate change-abatement strategies. [Reproduced with kind permission from figure 1 of ref. 18 (Copyright 2000, Kluwer Academic Publishers).]

energy prices and of greenhouse gas (GHG) emissions, which both influence the rate of economic growth. Policy decisions about the level of carbon taxes and technology incentives, which depend on observations of economic growth, damages and technology diffusion, also influence the agents' technology choices.

Robust Strategies for Climate Change

Global climate change is perhaps our most intractable environmental problem. Most economic activities, and in particular, the burning of fossil fuels, increase the concentration of long-lived GHGs in the atmosphere. The scientific evidence has become compelling that such emissions over the last 100 years have altered the earth's climate system. But, deep uncertainty remains about the future costs and benefits of any potential action to address the problem. It is clear, however, that if climate change is indeed a serious problem, society will have to make significant reductions in its GHG emissions, on the order of 80% below extrapolations of current trends, by the end of the 21st century. Technology innovation will likely play a major role in any changes of this scale.

This fact raises a key question: what should policy-makers do in the near-term to encourage such innovation? In practice, governments pursue a wide range of technology policies designed to improve the technology options for emissions reductions in the future. There is widespread agreement that government research and development can enhance future emissions-reducing innovations (although public-sector, energy-related research among most developed countries has been dropping in recent years). Research also supports the observation that private-sector firms will invest in new technologies when faced with the expectation of future regulations on GHG emissions. In addition, governments often pursue a wide range of technology incentive programs, such as tax credits, subsidies for low-emitting technologies, or market creation programs such as government procurements. These programs often appear attractive on both substantive and political grounds, but their record in practice is mixed. In addition, economic theory suggests that such policies often are inefficient compared with policies designed to "get the prices right," such as carbon taxes or tradable permits.

In our work, we have addressed the role of these technology incentives within the context of a robust response to climate change (18). The performance of alternative policies depends critically on the deeply uncertain characteristics of the dynamics of technology diffusion. Thus, we have used our CAR methods and an agent-based model of technology diffusion that focuses on the social and economic factors that influence how economic actors choose to adopt, or not to adopt, new emissions-reducing technologies. We link this agent model to a simple macro model of economic growth. The agent-based representation is particularly useful because it conveniently represents key factors influencing technology diffusion and, thus, policy choices such as the heterogeneity of technology preferences among economic actors and the flows of imperfect information that influence their decisions.

As shown in Fig. 2, each agent in our model represents a producer of a composite good that is aggregated as total Gross Domestic Product (GDP), using energy as one key input. During each time period, the agents first choose among several energy-generation technologies, and second, given their chosen technology, choose how much energy to consume. (That is, agents choose a production function and where to operate on that production function.) We assume that agents pick a technology to maximize their utility, which depends on each agent's expectations about the cost and performance of each technology. The agents have imperfect information about the current performance of new technologies but can improve their information

based on their own past experience, if any, with the technology and by querying other agents who have used it. The agents also are uncertain about the future costs of new technologies, which may or may not decline significantly because of increasing returns to scale. Agents estimate these future costs based on observations of the past rates of adoption and cost declines. Thus, the diffusion rate can depend reflexively on itself, because each user generates new information that can influence the adoption decisions of other potential users.

We use this model to compare the performance of alternative climate strategies. We posit two types of adaptive-decision strategies, "Limits-Only" and "Combined Strategy", both of which can evolve over time in response to new information. (We also considered no response and a strategy using only technology subsidies. Neither of these proved very attractive compared with the other two.) Our description of both strategies, shown in Fig. 3, is consistent with both the theoretical literature and qualitative understanding of the political conditions under which the policies would operate.

The "Limits-Only" strategy employs only a carbon tax (the arguments here also apply to emissions trading), whose level can change over time in response to observations of the rate of economic growth and the damages caused by climate change. The tax begins at some initial level per ton of emitted carbon and subsequently grows, remains constant, or returns to its initial level depending on observations of the damages caused by climate change and the economic growth rate. The "Combined Strategy" uses both the carbon tax and a technology subsidy, which can change over time in response to observations of the market share of low-emitting technologies. The subsidy begins at some initial level, expressed as a percent of the cost of the subsidized technology, and stays at a constant level over time until either the market share for low-emitting technologies goes above a threshold value or the market share fails to reach a minimum level after a certain number of years. If either of these conditions is met, the subsidy is permanently terminated.

The agent model of the effects of the alternative policies on the climate and economic systems presents deep uncertainty in over 30 different input parameters that represent a wide range of factors, including the macroeconomic effects of alternative policies on economic growth, the microeconomic preferences economic actors use to weigh the cost and performance of alternative technologies, the characteristics of new technologies, and the dynamics of the climate system. To compare policy choices in the face of this deep uncertainty, we use the process sketched in Fig. 1.

First, we generate a landscape of plausible futures by searching for the most diverse set of model inputs that plausibly represents the microlevel data and that gives plausible, macroscopic model outputs. In particular, we compare model outputs to constraints based on current market shares for energy technologies and aggregate levels of carbon emissions. We also constrain future technology diffusion rates to be no faster than the maximum historically observed rates for energy technologies. By using a genetic search algorithm, we generated a diverse ensemble of 1,611 different vectors of model inputs whose outputs meet these constraints. When used to project the model into the future, this ensemble gives forecasts of future GHG emissions that differ by over an order of magnitude, as shown in Fig. 4, a landscape of plausible futures. Combined with our alternative strategies, these futures give the scenario ensemble shown in Fig. 1.

Next, we use this scenario ensemble to generate hypotheses about robust strategies. There are too many dimensions of uncertainty in this model for an exhaustive search, so we first used statistical techniques (importance sampling) to find the input parameters most strongly correlated with a key model output of interest (GHG emissions after 50 years). These decisive uncertainties are the rate of cost reductions caused by

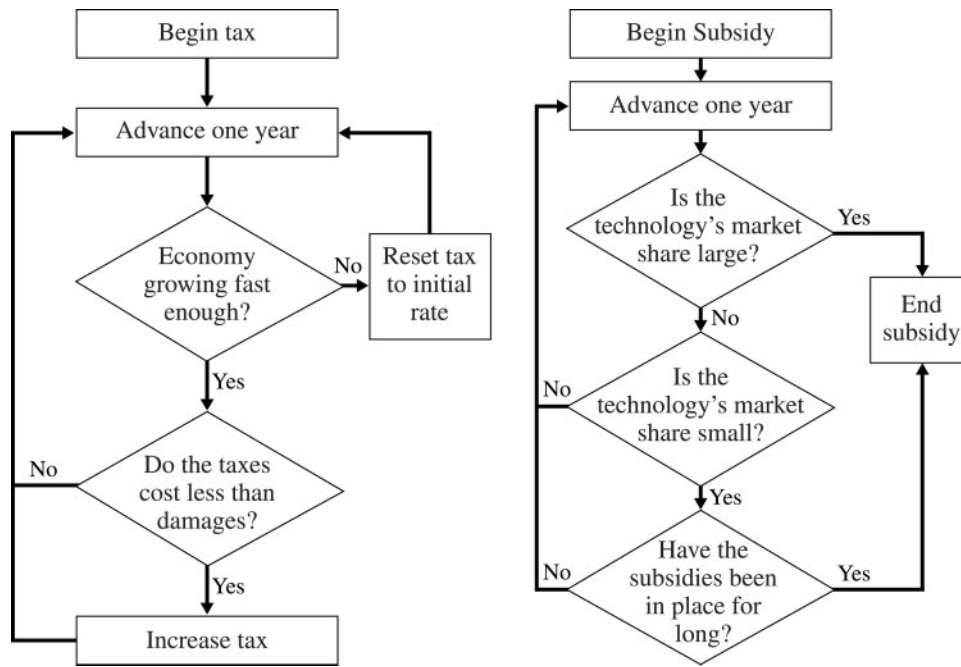


Fig. 3. Adaptive decision strategy for adjusting carbon taxes (Left) and technology incentives (Right) over time.

increasing returns to scale for the nonemitting technologies, the rate at which agents learn from one another about the performance of new technologies, the agents' risk aversion, and the heterogeneity of the agents' price-performance preferences for new technologies. Then, we viewed a series of interactive computer visualizations using different combinations of these four key uncertain inputs as independent variables (as well as a fifth uncertainty, the damages caused by climate change), each one comparing the performance of the "Limits-Only" and "Combined Strategy." Each visualization showed the performance of the two strategies as surface plots, measured as the present value of the GDP over the 21st century (reflecting both the costs and benefits of each strategy) as a function of two of the uncertainties, with the other inputs held constant at fixed values. A clear pattern emerged: the Limits-Only strategy is preferable in a

world where the agents' technology preferences are homogeneous, imperfect information effects are small, and the damages caused by climate change emerge slowly. When these conditions do not hold, the Combined (tax and subsidy) Strategy quickly becomes more attractive.

The robust region map in Fig. 5 summarizes these results. The figure shows the expectations about the future that should cause a decision-maker to prefer the Limits-Only strategy to the Combined Strategy. The horizontal axis represents the range of expectations a decision-maker might have for how likely it is—from very unlikely (Left) to very likely (Right)—that factors such as the potential number of early adopters and the amount of increasing returns to scale will significantly influence the diffusion of new technologies. The vertical axis represents the range of expectations a decision-maker might have that there




Fig. 4. A Landscape of Plausible Futures showing a wide range of future GHG emissions paths, all of which are consistent with available information. [Reproduced with kind permission from figure 3 of ref. 18 (Copyright 2000, Kluwer Academic Publishers).]




Fig. 5. A Robust Regions visualization showing regions in probability space where the expected gross world product resulting from the Limits-Only strategy is greater than that for the Combined Strategy as a function of the probability of high climate damages and of the potential for significant cost reductions for new technologies because of increasing returns to scale, reduction of uncertainty about the performance of new technologies because of learning among agents, and heterogeneity in the cost/performance preferences for new technologies among agents (“non-classical world”). [Reproduced with kind permission from figure 7 of ref. 18 (Copyright 2000, Kluwer Academic Publishers).]

will be significant impacts because of climate change (greater than 0.3% of the global economic product). The map shows that the Combined Strategy dominates even if decision-makers have only modest expectations that impacts from climate change will be significant and that information exchange and heterogeneity

among economic actors will be important to the diffusion of new, emissions-reducing technologies.

Testing the Hypotheses

Fig. 5 is based on an examination of only 6 of the 30 dimensions of uncertainty in the model. As a key final step, we tested the policy recommendations based on this reduced set of scenarios by launching a genetic search algorithm across the previously unexamined dimensions looking for counter examples to the conclusions summarized in Fig. 5. In this case, we found no plausible breaking scenarios. This analysis thus suggests that if decision-makers hold even modest expectations that market imperfections are likely to inhibit the diffusion of new, emissions-reducing technologies or that the impacts of climate change will turn out to be serious, then technology incentive programs may be a promising component of a robust response to climate change.

Conclusions

Models of complex systems capture much useful information about the world. However, such models can be difficult to apply to policy problems because the type of information they contain is often not that assumed by traditional methods of deductive decision analysis. New approaches, using inductive reasoning based on large ensembles of computational experiments, now make it possible to systematically compare policy alternatives using models of complex systems. The CAR methods shown here are widely applicable to policy problems characterized by complexity and deep uncertainty. These methods make it possible to capture more fully available information about complex systems, find policy responses that are robust against uncertainty, and bridge the gap between quantitative and qualitative approaches to decision-making.

This research was supported in part by the U.S. National Science Foundation under Grants ATM 95-22681 and BCS-9980337 and by the Integrated Assessment Program, Biological and Environmental Research, U.S. Department of Energy under Grant DE-FG03-96ER62262.

1. Dawes, R. (1988) *Rational Choice in an Uncertain World* (Harcourt Brace Jovanovich, San Francisco).
2. Schwartz, P. (1996) *The Art of the Long View* (Doubleday, Garden City, NY).
3. van der Heijden, K. (1996) *Scenarios: The Art of Strategic Conversation* (Wiley, New York).
4. Bankes, S. C., Lempert, R. J. & Popper, S. W. (2001) *Comput. Sci. Ind.* **3**, 71–77.
5. Bankes, S. C. (1993) *Oper. Res.* **41**, 435–449.
6. van Asselt, M. B. A. (2000) *Perspectives on Uncertainty and Risk* (Kluwer, Boston), p. 434.
7. Lempert, R. J. & Schlesinger, M. E. (2000) *Clim. Change* **45**, 387–401.
8. Caseman, E. A., Morgan, M. G. & Dowlatabadi, H. (1999) *Risk Anal.* **19**, 33–42.
9. Rotmans, J. & de Vries, H. J. M. (1997) *Perspectives on Global Change: The TARGETS Approach* (Cambridge Univ. Press, Cambridge, U.K.).
10. Watson, S. R. & Buede, D. M. (1987) *Decision Synthesis: The Principles and Practices of Decision Analysis* (Cambridge Univ. Press, New York).
11. Rosenhead, M. J., Elton, M. & Gupta, S. K. (1972) *Oper. Res. Q.* **23**, 413–430.
12. Brooks, A., Bankes, S. C. & Bennett, B. (1999) *Mil. Oper. Res.* **4**, 67–80.
13. Szayna, T. C., Byman, D. L., Bankes, S. C., Eaton, D., Jones, S. G., Mullins, R. E., Lesser, I. O. & Rosenau, W. (2001) *The Emergence of Peer Competitors: A Framework for Analysis* (RAND), MR-1346-A.
14. Dewar, J. A., Bankes, S. C., Edwards, S. J. A. & Wendt, J. C. (2000) *Expandability of the 21st Century Army* (RAND), MR-1190-A.
15. Davis, P. K. & Carrillo, M. J. (1997) *Exploratory Analysis of the Halt Problem: A Briefing on Methods and Initial Insights* (RAND), DB-232.
16. Lempert, R. J. & Bonomo, J. (1998) *New Methods for Robust Science and Technology Planning* (RAND), DB-238-DARPA.
17. Lempert, R. J., Schlesinger, M. E. & Bankes, S. C. (1996) *Clim. Change* **33**, 235–274.
18. Robalino, D. A. & Lempert, R. J. (2000) *Integr. Assess.* **1**, 1–19.
19. Lempert, R. J., Schlesinger, M. E., Bankes, S. C. & Andronova, N. G. (2000) *Clim. Change* **45**, 129–161.
20. Park, G. & Lempert, R. J. (1998) *The Class of 2014, Preserving Access to California Higher Education* (RAND), MR-971.