

Many-Objective Robust Decision Making for Water Supply Portfolio Planning Under Deep Uncertainty

Joseph R. Kasprzyk^a, Shanthi Nataraj^b, Patrick M. Reed^a, Robert J. Lempert^b

^a*The Pennsylvania State University (jrk301@psu.edu, preed@engr.psu.edu)*

^b*RAND Corporation (snataraj@rand.org, lempert@rand.org)*

Abstract: Portfolios of market-based instruments have been shown to improve the reliability of water supplies, using simulations that utilize a single best estimate of distributions of data to evaluate performance. However, the estimates of problem information and likelihoods could be incorrect, especially when planning for climate change, which can modify streamflow availability, or projecting the trajectories of future water demands. These conditions are termed deep uncertainty, in which decision makers cannot fully conceptualize or agree upon the full range of risks to their system. This presentation will advance a new interactive framework that combines robust decision making (RDM) with many-objective optimization using evolutionary algorithms (MOEA) to confront deep uncertainty for water planning. The framework is demonstrated using a case study that examines a single city's water supply in the Lower Rio Grande Valley (LRGV) in Texas, USA. We use a MOEA to develop a tradeoff set of water supply portfolios for the LRGV, and develop a suite of values for key uncertainties using RDM that represent an ensemble of "states of the world". Each solution is tested under the ensemble of plausible future states of the world, with interactive visualizations being used to identify robust solutions for the system. Scenario discovery methods that use statistical data mining algorithms are then used to identify what assumptions and system conditions strongly control the cost-effectiveness, efficiency, and reliability of the robust alternatives. The results suggest that combining robust decision making, many-objective optimization, and visual analytics can dramatically improve risk-based planning decisions.

Keywords: robust decision making, many-objective optimization, water supply, interactive visual analytics

1 INTRODUCTION

Although climate change and urbanization pose serious threats to water management, new infrastructure projects are often not environmentally or politically desirable. In contrast, nonstructural approaches such as water marketing [Anderson and Hill, 1997] can help improve the reliability of municipal water supply by facilitating transfers of water between user sectors or regions. Recent studies [Kasprzyk et al. 2009; Kasprzyk et al. In-Press] have explored how a single city in the Lower Rio Grande Valley (LRGV) of Texas, USA can increase its supply reliability by augmenting traditional supply with transfers from a water market. A Monte Carlo simulation model of the LRGV uses expected value calculations of hydrology, demand, and lease pricing to evaluate each water portfolio design for the system. The simulation model is linked to a multiobjective evolutionary algorithm (MOEA), which generates Pareto-approximate alternatives for the LRGV's water system, balancing complex stochastic performance objectives, including cost, reliability, and the volume of wasted water specified by each portfolio.

However, calculating expected value performance objectives in this manner requires assumptions about the input data that could be violated in future planning periods. For example, there are concerns over reduced inflow from the Mexican tributaries feeding the LRGV's reservoir system, and climate change has the potential to increase the amount of evaporative energy in the system, exacerbating supply losses. These challenges show that using only the baseline historical data to select alternatives may fail to properly address unexpected system changes and shifts in estimated likelihoods.

Shifting likelihoods within the LRGV's supply system can be termed deeply uncertain [Knight, 1921], since decision makers cannot fully conceptualize or agree on the full range of risks to their system, or the likelihood of those risks [Langlois and Cosgel, 1993; Lempert, 2002]. Traditional scenario analysis [e.g., Schwartz, 1996] seeks to address deeply uncertain planning problems by building a small number of scenarios of the future. A significant issue with this approach, however, is that these scenarios do not link the assumptions about the system to the ultimate planning goal, such as how a system will respond to climate projections [Brown et al., 2011]. In general, scenario analysis fails to answer how changes in the assumptions of exogenous factors affect the likelihood of severe vulnerabilities of system performance. Robust Decision Making (RDM) [Lempert et al., 2006] has been proposed and demonstrated as an effective way to evaluate policy alternatives for their robustness to deeply uncertain planning conditions. RDM emphasizes simulating a broad range of plausible futures and identifying factors that shape future risks. However, to date RDM has not strongly emphasized the role of generating alternatives, as was demonstrated in the prior work with MOEAs in the LRGV.

The goal of this paper is to introduce *Many-Objective Robust Decision Making* and demonstrate it using the LRGV's water supply test case. Section 2 will present the Many-Objective RDM framework, and section 3 will review the prior work that used MOEAs to generate alternatives for the LRGV. Section 4 will summarize how the new framework negotiates robust planning alternatives, with conclusions given in section 5.

2. MANY-OBJECTIVE ROBUST DECISION MAKING

The many-objective robust decision making framework uses XLRM terminology [Lempert et al., 2003]. X refers to deep uncertainties for planning such as assumptions of input data distributions. Levers (L) are actions that decision makers can take to modify their system and are treated as decision variables in the optimization. The relationship, R, between actions and outcomes is coded in a quantitative simulation model. Finally, measures (M) are a set of quantitative outcomes of the system design expressed as objectives in the optimization.

Step 1 is *Problem Formulation*, a hypothesis of the best-known XLRM components for the system in an initial problem formulation. The problem formulation expresses multiple decision maker performance measures that will be used to define the Pareto optimal set of solutions: solutions are Pareto optimal if they are better than all feasible solutions in at least one objective. Note that the many-objective RDM framework is a type of a *posteriori* decision support, in that no weights or preferences between different objectives or measures are provided in the beginning of the analysis. The goal of the analysis is to provide a full set of alternatives that the decision makers can analyze, trading off various preferences after the alternative points have been generated.

Step 2, *Generating Alternatives*, uses a MOEA to generate alternatives for the system. MOEAs are heuristic, population-based search tools [Coello Coello et al., 2007; Nicklow et al. 2010] that use an iterative search of selecting high-quality population members and variation on those members to evolve a high-quality approximation to the Pareto optimal set. This process typically uses a baseline risk

simulation model with a single estimate of model parameters and of distributions of input data.

Step 3, *Uncertainty Analysis*, addresses the limitation of only using a single baseline risk simulation in the optimization. In collaboration with stakeholders, we identify key deep uncertainties and create a large ensemble of possible values for states of the world. This uncertainty ensemble is applied to each solution in the Pareto-approximate tradeoff set, and interactive visualizations are used to find a robust tradeoff solution. A robust solution is defined as one that has a low amount of deviation between its output measure values in the uncertainty ensemble relative to the original values from the baseline risk model.

Step 4, *Scenario Discovery*, uses statistical data-mining algorithms on the selected robust solution to identify values of the uncertainties that cause performance failures. These values of the uncertainties can motivate future data collection and monitoring in the system to plan for future risks.

3. GENERATING ALTERNATIVES FOR THE LRGV

The LRGV case study is used to demonstrate the framework, which is a risk-based water management problem where a single city tries to find the best combination of permanent rights and market-based options and leases for their water supply. We build off the analysis in Kasprzyk et al. [In-Press], in which six performance measures are used to quantify the performance of each supply portfolio. The original problem formulation uses eight decision variables to describe the city's design levers for their portfolio: a volume of permanent rights N_R , an adaptive options contract ($N_{O,low}$, $N_{O,high}$, and ξ), and supply/demand thresholds to control acquiring water on the market through options and leasing ($\alpha_{Jan-Apr}$, $\beta_{Jan-Apr}$, $\alpha_{May-Dec}$, $\beta_{May-Dec}$). In the set of supply/demand thresholds, alpha controls "when" to acquire water on the market: if the current supply is less than a ratio of alpha to the expected demand, the city must go to the market. Then, beta (constrained to be higher than alpha) controls "how much" the city buys. The original formulation contains one set of alpha/beta variables before May, and a separate set afterwards, for a total of four variables.

Kasprzyk et al. [In-Press] constructed 4 model cases to test the appropriate level of complexity for the LRGV's decision variables. The motivating question was: how many variables are necessary to generate high-quality planning alternatives? Case I uses a volume of permanent rights, a single-volume options contract, and one variable to determine both "when" to go to the market and "how much" to acquire (3 variables: N_R , N_O , and α). Case II varies the when/how much decision by the time of year, while keeping the single-volume options contract (4 variables: N_R , N_O , $\alpha_{Jan-Apr}$, and $\alpha_{May-Dec}$). Case III separates the "when" and "how much" decision (6 variables: N_R , N_O , $\alpha_{Jan-Apr}$, $\alpha_{May-Dec}$, $\beta_{Jan-Apr}$, $\beta_{May-Dec}$). Finally case IV adds the adaptive options contract for the full complexity of 8 decision variables. Kasprzyk et al. [In-Press] solved each of the four model cases using a MOEA, and compared their objective function performance to choose a preferred model case. Case III was chosen for further analysis, since it had preferred performance in cost, surplus water, and number of leases relative to the simpler model cases.

4. DISCOVERING ROBUST PLANNING ALTERNATIVES

In the analysis summarized in section 3, a single expected value calculation was used to generate the planning alternatives for the LRGV. This entails a single value for model parameters such as the demand growth rate, and a single distribution for each input variable such as inflow. The purpose of this section is to determine if our choice of case III was biased by the use of this single uncertainty estimate, hereafter termed the "baseline state of the world".

The third step in the framework is to test many plausible "states of the world" for the deep uncertainties. The uncertainties are sampled as dimensions of a Latin Hypercube Sample (LHS) of 10,000 samples. Two types of uncertainties are

sampled. The first type represents scaling factors for the input data; the lowest sampled inflows, for example, are scaled to be between 1 and 10 times likelier than in the baseline state of the world. The second type samples point-valued model parameters, such as the initial allocation to permanent rights. The ensemble of 10,000 samples is applied to each solution in the tradeoff set, and for each of the 10,000 “states of the world” we calculate a suite of output measures. The results are sorted, and we identify the most extreme 10% of the cases (i.e., the 10th percentile for maximization measures and the 90th percentile for minimization measures).

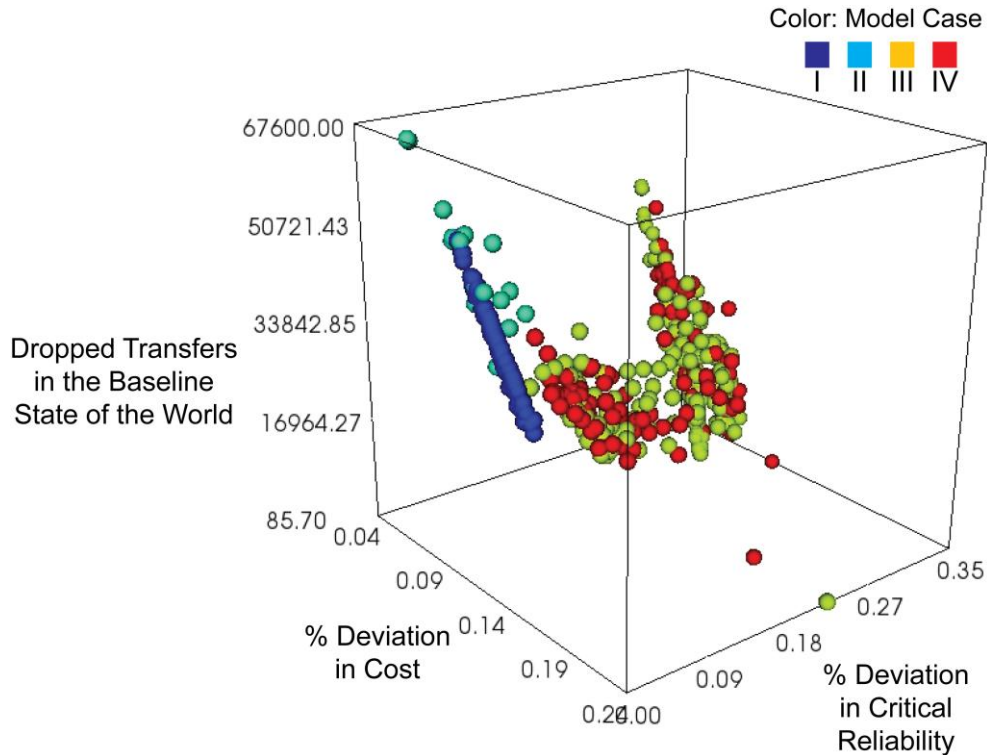


Figure 1. Glyph plot comparing dropped transfers in the baseline state of the world versus percent deviation for cost and critical reliability. Color shows the model case. The blue and cyan solutions have lower critical reliability deviation, meaning they are more robust than the more complex model cases.

Figure 1 shows a representative result from the third step of the framework. The vertical axis plots the dropped transfers objective from the baseline state of the world. The horizontal axes plot the percent deviation between the most extreme 10% of objective values in the ensemble of states of the world and the baseline condition, in cost and critical reliability. The color shows the model case, from the simplest case I formulation in blue to the most complicated formulation, case IV, in red. The figure shows that the solutions from case I have robust performance: near zero deviation in critical reliability with varying deviation in cost. The preferred solutions from case III, however, have higher deviation in critical reliability.

The fourth step of the framework uses scenario discovery [Bryant and Lempert, 2010] to identify values for the deep uncertainties that cause performance failure for the robust solution. Three groups of output measures were chosen to quantify the results. The cost group combines the cost and cost variability measures; the reliability group combines reliability, critical reliability, and drought reliability; and the market use group uses the number of leases. Similar to the percent deviation results shown in figure 1, we try to determine values of the uncertainties that cause performance measures to fall in the most extreme 10% of the distribution of outputs (the 90th percentile for minimization measures or the 10th percentile for maximization measures).

In preliminary work, we selected a solution from model case I termed the “robust” solution. Scenario discovery indicates that some dimensions are critically important for the more robust solution. If the likelihood of high losses, for example, is twice as high as in the baseline state of the world, all three groups of measures may have poor performance. Some dimensions must be scaled significantly higher, however; reliability does not have poor performance until low inflows are scaled 6.5 times higher than in the baseline state of the world. Looking across all dimensions, decision makers would have to be cautious about setting strict targets for the cost measures, since a large number of dimensions cause performance issues with cost, including high demand growth rates and low initial reservoir volumes.

5. CONCLUSION

This paper has demonstrated the value of combining many-objective optimization with robust decision making and interactive visualizations for finding high-quality robust planning alternatives. Prior work suggested that a moderate complexity set of decision variables termed “model case III” gave good performance for an urban water test case. However, the current analysis indicates that simpler decision variable formulations are able to provide more robust performance, maintaining high reliability under a large ensemble of plausible future states of the world. It is not the goal of this paper to claim that specific values of the uncertainty ensemble are likely to happen in the future. Instead, our approach addresses the pivotal question: “how wrong do our assumptions have to be to cause performance failure?”

By using the scenario discovery algorithm, we showed that the likelihood of high losses is critically important to avoiding high costs, low reliabilities, and high market use. This information can be used to promote monitoring in the water system; by monitoring evaporation rates in the system, decision makers can determine whether or not high losses are likely to cause vulnerabilities in their system. Furthermore, the many-objective robust decision making framework is iterative. Under new climate, land use, or water demand regimes, analysts can modify their modelling assumptions and “re-optimize” the system as they learn more about future challenges.

ACKNOWLEDGMENTS

The first author was funded by the US EPA STAR Graduate Fellowship program. The contents of the research are solely the responsibility of the grantee and do not necessarily represent the official views of the US EPA. Further, US EPA does not endorse the purchase of any commercial products or services mentioned in the publication. Also, this work was supported in part through instrumentation funded by the National Science Foundation through grant OCI-0821527.

REFERENCES

- Anderson, T.L. and P.J. Hill (eds) *Water Marketing: The Next Generation*. Rowman and Littlefield, Lanham, 1997.
- Brown, C., W. Werick, W. Leger, and D. Fay, A decision-analytic approach to managing climate risks: Application to the Upper Great Lakes, *Journal of the American Water Resources Association*, 47, 2011.
- Coello Coello, C., G.B. Lamont, and D.A. Van Veldhuizen (eds) *Evolutionary Algorithms for Solving Multi-Objective Problems*. Springer, New York, 2 ed, 2007.
- Kasprzyk, J.R., P.M. Reed, B.R. Kirsch, and G.W. Characklis, Managing population and drought risks using many-objective water portfolio planning under uncertainty, *Water Resources Research*, 45.
- Kasprzyk, J.R., P.M. Reed, G.W. Characklis, and B.R. Kirsch, Many-objective de Novo water supply portfolio planning under deep uncertainty, *Environmental Modelling and Software*, In-Press.
- Knight, F. *Risk, Uncertainty, and Profit*, Houghton Mifflin, Boston, 1921.

- Langlois, R.N., and M.M. Cosgel, Frank Night on risk, uncertainty, and the firm: A new interpretation, *Economic Inquiry*, 31, p.456-465, 1993.
- Lempert, R.J. A new decision sciences for complex systems, *Proceedings of the National Academy of Sciences*, 99, p. 7309-7313, 2002.
- Lempert, R. J., S. W. Popper, S. C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*, RAND, Santa Monica, 2003.
- Lempert, R. J., D. G. Groves, S. W. Popper, and S. C. Bankes, A general, analytic method for generating robust strategies and narrative scenarios, *Management Science*, 52, p. 514-528, 2006.
- Nicklow, J. et al. State of the art for genetic algorithms and beyond in water resources planning and management, *Journal of Water Resources Planning and Management*, 136, p. 412-432, 2010.
- Schwartz, P., *The Art of the Long View*, Doubleday, Garden City, 1996.