ADVANCING REGIONAL WATER SUPPLY PORTFOLIO MANAGEMENT AND COOPERATIVE INFRASTRUCTURE INVESTMENT PATHWAYS UNDER DEEP UNCERTAINTY

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Bernardo Carvalho Trindade

December 2019

ADVANCING REGIONAL WATER SUPPLY PORTFOLIO MANAGEMENT AND COOPERATIVE INFRASTRUCTURE INVESTMENT PATHWAYS UNDER DEEP UNCERTAINTY

Bernardo Carvalho Trindade, Ph.D.

Cornell University 2019

Water utilities globally are facing a great pressure to proactively address growing demands, higher levels of resource contention, and increasingly uncertain water availability. A tension exists between improving the efficiency and coordination of regional water supplies to delay or eliminate major infrastructure investments and the eventual effects of demand growth rates that necessitate water supply capacity expansions. The hard- and soft-path views for urban water portfolio planning, namely infrastructure construction and efficient regionally-coordinated management (e.g., restrictions, water transfers, financial instruments, etc.), should be integrated while minimizing costly and contentious infrastructure expansion. However, both the soft and hard approaches may be detrimental to utilitiesThis ’ finances. Integrating long-term water infrastructure investment and short-term management is a difficult task because of the potentially large number of decisions that must be considered over decadal timescales as well as the significant uncertainties inherent to the problem. This dissertation contributes a new uncertainty-sampling formulation for the stochastic simulation and optimization of water-infrastructure management policies (soft path alone) for regionally-coordinated water utilities. This sampling formulation accounts not only for well-characterized uncertainties, usually hydrological or climate extremes, but also for key urban-systems fac-

tors such as demand growth, effectiveness of water-use restrictions, construction costs, and financing uncertainties. This sampling scheme is later integrated into WaterPaths, a generalizable, cloud-compatible, open-source exploratory modeling framework designed to simulate and inform long-term regional investments in water infrastructure while simultaneously aiding regions to improve their short-term weekly management decisions. Uniquely, WaterPaths has the capability to identify the challenges and demonstrate the benefits of regionally-coordinated planning and management based on risk-of-failure decision metrics for groups of water utilities sharing water resources. The WaterPaths platform is introduced here through the hypothetical and realistic Sedento Valley test case, a test case designed to serve as a universal test case for decision-making frameworks and water-resources systems simulation software. Lastly, WaterPaths is integrated into the Deep Uncertainty (DU) Pathways framework. The DU Pathways framework is a decision-making framework developed for bridging long-term water supply infrastructure investments and improved short-term water portfolio management to yield robust regional water supply. The DU Pathways framework allows for elaborate uncertainty analysis to clarify how uncertainty may determine the appropriateness of infrastructure construction and the effectiveness of supply and financial drought-mitigation instruments used by regionally-coordinated water utilities. The ultimate goal of the DU Pathways is to aid stakeholders in discovering pathway policies that attain high performance levels across challenging and deeply uncertain futures and to guide robustness compromises that may be necessary between regional actors. As demonstrated in the Research Triangle region in North Carolina, DU Pathways clarifies how to identify robust infrastructure investment and management policies across the municipalities of Raleigh, Durham, Cary, and

Chapel Hill.

BIOGRAPHICAL SKETCH

Bernardo Trindade grew up in Brasília, the capital of Brazil. He attended the Federal University of Santa Catarina in the south of Brazil for one semester but then returned to Braslia to study at University of Brasília, where he received a B.S. in Civil and Environmental Engineering in 2009. After an six months as an intern on a construction site, Bernardo spend another six months designing sewage collection systems in Brazil before moving to the United States in 2010. Bernardo earned a M.S. in Civil and Environmental Engineering in 2012 from Auburn University (AL) with a concentration in computational pipe hydraulics, which led him to work at Bechtel Oil, Gas and Chemicals in Houston, TX, with natural-gas-liquefaction plants from 2012 to 2014. In 2014, he started started his Ph.D. at Cornell Environmental & Water Resources Systems program, after which he expects to apply his research in an industry position.

To my family, whose example illuminates the road leading me to ever greater success. *O que faz o homem o trabalho*.

ACKNOWLEDGEMENTS

I would not have succeeded in my Ph.D. without the intellectual and personal support of countless great people. I would like to first thank my advisor, Prof. Patrick Reed, for stirring my ideas in the right direction through his own while constantly providing me with the feedback and generous resources I needed to succeed. I am particularly grateful for him having brought me to this great university and encouraged me to persue my interests in topics ranging from machine learning to existentialism, Marxism, and wines, all of which changed my worldview and the professional life ahead of me. I also would like to thank John Foote, whose professional advice, friendship, introductions, inspiring example, and selfless commitment to my success have stirred my future career in a direction I could never have anticipated. I would also like to thank my current committee members: Professors Rich Geddes and Daniel Selva for the early and industry-focused advice, Prof. Rich Geddes and Prof. Oliver Gao for reviewing my five years worth of work, Prof. Geddes for introducing me for countless high-level professionals in the infrastructure industry, and Prof. Gao for graciously accepting the last-minute invitation to replace Prof. Daniel Selva as my committee member after the latter left Cornell. I would like to thank Prof. Pete Loucks for the constant mentorship about engineering and leadership, for the squash games, and for all the shared meals, Prof. Jery Stedinger for all conversations that made me see statistics as the fascinating discipline it is despite my initial prejudices, and Prof. David Bindel for helping me during the development of WaterPaths when he had no reason to do so. I would also like to thank my collaborators starting from Prof. Jon Herman for teaching me the rudiments of my future research, Dr. H. B. Zeff and Prof. Gregory Characklis for providing a pillar for my work at Cornell, David Gorelick for all the e-mail about bugs and improvements on WaterPaths, and David Gold for all the help close to the finish line. I would like to thank Dr. Jared Smith and Professors John Lamontagne and Julie Quinn for all the intellectual insight that so much improved the quality of my work and education. Lastly, I would like to thank all students I had the honor to help as a Graduate Teaching Assistant, who through intellectual rigor and while maintaining and friendly and fun environment pushed me to master the disciplines I was in charge of, and to Nadine Porter, Jeannette Little, Charissa King-O'Brien, Carl Cornell, Megan Keene, Tania Sharpsteen, Beth Korson, Joe Rowe, and Alita Howard for being such great department staff.

Besides great mentors and collaborators, Cornell has provided me with friends and colleagues of the highest quality and a fantastic and affectionate social life. I have been fortunate to spend time in and outside Hollister with and receive great friendship and support from Dr. Jared Smith, Prof. John Lamontagne, Prof. Julie Quinn, Prof. Jon Herman, David Gold, Jazmin Zatarain, Kwang-Bae Choi, Rohini Gupta, Peter Storm, Sarah Kayali El-Alem, Tori Ward, Dr. Min Pang, Dr. Tom Wild, Calvin Whealton, Candice Wu, Xiaotong Guo, Bernardo Casares, Justice Sefas, and Fábio Guilhon. I have also greatly appreciated the time spent with great music and swing dance with Renée Staffeld, my eternal swing-dance partner, Amanda Bartley, Furkat Mukhtarov, Leah Hoffmann, Aleksa Basara, Addie Lederman, Allison Stilin, Aimee Owens, Katie Spoth, Tres Reid, Sierra Cook, Chris Westersund, Renée Sifri, Yifei Zheng, YangYang Sun, Colin Bundschu, Mari Tader, Andrea Staffeld, Bill Staffeld, Anna-Katharina von Krauland, Irine Sarri, Cesare Paolucci, Céline Hollfelder, Tamilla Kassenova, Destiny Meador, Erica Watson, Gustavo Ribeiro, Erica Shaw, Howard Chong, Kevin Lin, Rachel Vanicek, Tommi Schieder, Schuyler Duffy, Vishnu Govindaraj, Yudi Pardo, Yupeng Lin. My adversaries inside the tennis court but friends and colleagues outside have also much supported me during my time at Cornell. They are Ian Bailey, Bruno Santarelli, Kevin O'Keefe, Andrew Fields, Evan Gadreau, Luca Moreschini, Chloé Albietz, Ricardo Paes, Chris Levine, Milo Johnson, Anda Perianu, Silviu Tanasoiu, Lev Kazakov, Caterina Kurz, Rachel Allen, Anita Jain, Allison Vizgaitis, Dina Rigas, and Rodrigo Viteri. I would like to also thank my great friends Ilse and Jon Jalvin for all the great wine and conversations, Sara Vandenbroek for her support during the long development of WaterPaths, and my great Katyusha Potemkina for all her support and patience as I try to close my work at Cornell.

This research was funded by the National Institute of Food and Agriculture, U.S. Department of Agriculture (WSC Agreement No. 2014-67003-22076). Additional support was provided by the U.S. National Science Foundation's Water, Sustainability, and Climate Program (Award No. 1360442). I also received funding from the Cornell School of Civil & Environmental Engineering as a Graduate Teaching Assistant. High-performance computing resources were provided by the Texas Advanced Computing Center (TACC), Blue Waters, and by the always-prompt support from the folks at the Cornell Center for Advanced Computing.

Finally, I would like to thank my parents and sister for providing me with wings as my duties tied my feet. They have always been a constant source of uplifting inspiration and love, making their presence felt at every hour from the other side of the equator. And none of my great achievements over the last ten years would have been possible were it not for the great country that is the United States of America, which has provided me with more love and opportunities than I will ever deserve.

TABLE OF CONTENTS

Biographical Sketch . . . iii
Dedication . . . iv
Acknowledgements . . . v
Table of Contents . . . viii
List of Tables . . . xii
List of Figures . . . xiv

1 Introduction **1**
1.1 Current Challenges for Water Infrastructure Planning and Management . . . 1
1.1.1 The Water Portfolio Planning and Management Problem . 3
1.1.2 Considering Uncertainty . . . 5
1.1.3 Existing Water-Infrastructure Planning and Management Software and Decision-Making Framworks and Their Limitations . . . 6
1.2 Dissertation Objectives . . . 11
1.3 Scope and Organization . . . 13
1.3.1 Chapter 2 — Background on Methodological Components 13
1.3.2 Chapter 3— Reducing Regional Drought Vulnerabilities and Multi-city Robustness Conflicts Using Many-objective Optimization Under Deep Uncertainty . . . 13
1.3.3 Chapter 4 — WaterPaths: A Stochastic Simulation Framework for Water Supply Infrastructure Investment Management . . . 14
1.3.4 Chapter 5 — Deeply Uncertain Pathways: Integrated Multi-City Regional Water Supply Infrastructure Investment and Portfolio Management . . . 16
1.3.5 Chapter 6 — Contributions and Future Work . . . 16
1.3.6 Author Contributions for Published Work . . . 17

2 Methodological Components **18**
2.1 Borg Multi-Objective Evolutionary Optimization Algorithm . . . 18
2.1.1 Adaptive Operator Selection . . . 20
2.1.2 Solution population . . . 22
2.1.3 Solution Archive and ϵ-Dominance . . . 23
2.1.4 Parallelization Strategies for the Borg MOEA . . . 25
2.2 Scenario Generation for Streamflow and Evaporation Rates Time Series . . . 27
2.3 Scenario Generation for Synthetic Demand Time-series . . . 30

3 Reducing regional drought vulnerabilities and multi-city robustness conflicts using many-objective optimization under deep uncertainty 32
3.1 Abstract . . . 32
3.2 Introduction . . . 33
3.3 Methodological Framework . . . 37
3.3.1 Overview . . . 37
3.3.2 States of the World . . . 39
3.3.3 Alternatives Generation . . . 40
3.3.4 Robustness Measures . . . 43
3.3.5 Robustness Controls . . . 44
3.4 The Research Triangle Test Case . . . 46
3.4.1 Triangle Region Management Model . . . 49
3.5 The Mathematical Formulation of the ROF Metric . . . 53
3.6 Computational Experiment . . . 60
3.6.1 Multiobjective Optimization Algorithm . . . 60
3.6.2 Hydrologic Scenarios and DU Optimization Runs . . . 60
3.7 Results and Discussion . . . 62
3.7.1 Optimization under Deep Uncertainties . . . 62
3.7.2 Robustness Comparison . . . 64
3.7.3 The Value of Regionally Coordinated Demand Management 70
3.7.4 Scenario Discovery . . . 72
3.8 Discussion . . . 75
3.9 Conclusions . . . 78

4 WaterPaths: An open source stochastic simulation framework for water supply portfolio management and infrastructure investment pathways 81
4.1 Abstract . . . 81
4.2 Software Availability . . . 82
4.3 Introduction . . . 83
4.4 WaterPaths Framework . . . 91
4.4.1 WaterPaths Overview . . . 91
4.4.2 ROF-based decision rules . . . 96
4.4.3 Mass-Balance Model . . . 99
4.4.4 Capturing Regional Decision Making . . . 100
4.4.5 Drought Mitigation Instruments . . . 101
4.4.6 Financial instruments . . . 102
4.4.7 Infrastructure Investment . . . 103
4.4.8 Flexible Representation and Evaluation of Uncertainties . 105
4.4.9 Currently Implemented Objective Functions . . . 107
4.4.10 Reducing Runtime With Local Parallel Computing . . . 109
4.4.11 Facilitating Decision Analytics . . . 112
4.4.12 WaterPaths Parallelization Strategy for Computational Policy Search . . . 114

4.5 Methodology 116
4.5.1 The Sedento Valley: An Illustrative Test Case 116
4.5.2 Problem Formulation 120
4.5.3 Discovering Robust Pathways for the Sedento Valley . . . 126
4.5.4 Identifying Tradeoffs 126
4.5.5 Evaluating Robustness 131
4.5.6 Pathway Analytics 133
4.5.7 Scaling Performance on Clusters and in the Cloud 135
4.6 Results 140
4.6.1 Performance and Robustness Tradeoffs 140
4.6.2 Policy Rules in the Robustness Compromises 145
4.6.3 Scenario Discovery for Compromise Policies 149
4.6.4 Infrastructure Ensemble Pathway Analysis 153
4.7 Conclusion 155

5 Deeply Uncertain Pathways: Integrated Multi-City Regional Water Supply Infrastructure Investment and Portfolio Management 157
5.1 Abstract 157
5.2 Introduction 158
5.3 The Long-Term Version of the Research Triangle Test Case 163
5.4 Methodology 167
5.4.1 Identifying Infrastructure Investment and Drought Management Policy Tradeoffs 168
5.4.2 Defining and Evaluating Robustness 181
5.4.3 Discovering Uncertain Scenarios that Control Robustness 184
5.4.4 Visual Infrastructure Pathway Analysis 187
5.5 Computational Experiment 188
5.6 Results 194
5.6.1 Comparing Tradeoffs After Re-evaluation 194
5.6.2 Scenario Discovery for Compromise Policies 199
5.6.3 Policy Rules in the Robustness Compromises 204
5.6.4 Infrastructure Pathway Ensemble 207
5.7 Conclusion 214

6 Contributions and Future Work 217
6.1 Contributions and Conclusions 217
6.2 Future Work 222
6.2.1 Strategies for Deep- and Well-Characterized-Uncertainty Sampling 222
6.2.2 Metrics for Long- and Short-Term Decision Making for Water Utilities 223
6.2.3 Design of New Financing Mechanisms and Insurance Instruments for Water Infrastructure Operators 224

6.2.4 Conflict Resolution Among Resource-Sharing Water Utilities 225

A Notation **227**

B Time-series Output for Compromise Solutions under selected scenarios **231**

C Deeply Uncertain Factor Sensitivities **235**

D Computational Experiment Hypervolume Plots **236**

E The Boosted Trees Algorithm **237**

E.1 Classification and Regression Trees 237
E.2 General Idea of Booosting 238
E.3 Geometric Intuition 239
E.4 Minimizing a Loss Function 241
E.5 AdaBoost 242
E.6 Sensitivities of Uncertainty Factors 244
E.7 Exponencial Convergence 245
E.8 Reduced Overfitting 245

LIST OF TABLES

3.1 Uncertainty Factors and Sampling Ranges for Many-Objective Robust Decision Making (vector Ψ in Equations 3.5 and 3.10). . . 42
3.2 Summary of reservoir storage capacities, demands, and allocations for the Research Triangle Region 48
3.3 Values used for ϵ-dominance. 62

4.1 Select modeling software for water resources planning and management 87
4.2 Select modeling software for water resources planning and management (continued) 88
4.3 Uncertainties included in extended version of the regional water utility planning and management problem presented in Zeff et al. (2016). Regional uncertainties are those for which the same value was used for all utilities, versus a value different value for each utility, as in the case uncertainties described as individual. . 106
4.4 Current infrastructure capacity of Sedento Valley Water Utilities. 118
4.5 Demand to Capacity Ratios for the Sedento Valley Water Utilities 118
4.6 Potential new infrastructure options in the Sedento Valley. 120
4.7 Decision variables pertaining to the short-term drought mitigation instruments: water-use restrictions, transfers, contingency fund and drought insurance. 124
4.8 Values of the long-term ROF trigger and daily demands that function as thresholds to trigger infrastructure construction by the utilities. 125
4.9 The ordinal variables below determine the infrastructure construction order and adjusts them for the total number of options available to each utility. The volumetric variable represents the volume of available storage within Falls Lake to be re-allocated from the water quality to Raleigh's municipal supply pool. . . . 125
4.10 Values used for ϵ-dominance. 125
4.11 Deep uncertainties considered for the Sedento Valley test problem. Unless specified otherwise the same minimum and maximum values for each uncertainty were applied for all utilities and infrastructure. 130
4.12 Architectures and configurations used for optimization scalability and cross-platform performance tests, followed by total run times and efficiencies for each optimization seed, ran on Stampede 2 and on the cluster The Cube with different parallel configurations. 138
4.13 Timings and efficiency of running WaterPaths with 1,000 realizations split among different numbers of cores on a 4-core Xeon Desktop and on a 48-core SkyLake node of Stampede 2. 139

5.1 Summary of proposed water supply infrastructure for the Research Triangle (Zeff et al., 2016; Triangle J, 2014). 165
5.2 Uncertainties included in extended version of the regional water utility planning and management problem presented in Zeff et al. (2016). Regional uncertainteis are those for which the same value was used for all utilities, versus a value different value for each utility, as in the case uncertainties described as individual. . 167
5.3 Decision variables pertaining to the short-term drought mitigation instruments: water-use restrictions, transfers, contingency fund and drought insurance. 189
5.4 Values of the long-term rof trigger and daily demands that function as thresholds to trigger infrastructure construction by the utilities. 191
5.5 The ordinal variables below determine the infrastructure construction order and adjusts them for the total number of options available to each utility. The volumetric variable represents the volume of available storage within Falls Lake to be re-allocated from the water quality to Raleigh's municipal supply pool. . . . 192
5.6 Values used for ϵ-dominance. 193

A.1 List of symbols. 230

C.1 Percent of the combined impurity the the entire boosted trees ensemble due to each uncertain factor for Durham, Cary and Raleigh under both buyer preferred policies D/B1 and R/B2. . . 235

LIST OF FIGURES

2.1 Flowchart of the Borg MOEA main loop adapted from (Hadka and Reed, 2013). First, one of the recombination operators is selected based on its performance. After, for an operator requiring k parents to generate an offspring, Borg selects $k-1$ from the population and one from the archive. The resulting offspring is evaluated and considered for inclusion in the population and archive. 20
2.2 Example distributions of offspring solutions (small dots) produced from parents solutions (large dots) by the six operators included in Borg MOEA. From Hadka and Reed (2014). 21
2.3 Illustration of two-objective optimization problem on an ϵ grid, adapted from Kasprzyk et al. (2013). Only one solution within each tolerance band ϵ_1 and ϵ_2 is kept in the Pareto set. Following the concept of ϵ-dominance, B, C, and E are the only non-dominated points. 24
2.4 Panel (2) shows a diagram of the master-worker implementation of the Borg MOEA. The master node maintains the ϵ-dominance archive and runs the main loop of the serial Borg MOEA shown in Figure 2.1. The decision variables generated through the six operators inplemented in the master node are transmitted to the worker nodes, and the evaluated objective function values and constraints are returned to the master node. Panel (b) shows a diagram of the multi-master implementation of the Borg MOEA. The multi-master Borg MOEA consists of two or more mastere-worker instances (the diagram depicts three) periodically communicating with one controller node, which maintains a global ϵ-archive. (1) Each master node periodically transmits its local ϵ-dominance archive to the controller, which updates the global archive. (2) When a master node is not advancing its pareto-approximate front, it sends a help message to the controller. (3) The controller responds with guidance, which includes the global ϵ-dominance archive and global operator probabilities. From Hadka and Reed (2014). 26
2.5 Panel (a) shows examples of flow duration curves for historical inflows, synthetically inflows with stationary log-mean, and synthetically inflows with non-stationary log-mean. Panel (b) shows all sinusoid-multiplier series of the log-mean of inflows and evaporation rates used to generate the panel (a). 29

3.1 Schematics of the deep uncertainty re-evaluation of candidate water portfolio solutions (termed DU re-evaluation). A set of objectives is calculated for each 1,000 SOW's that share the same $\mathbf{\Psi}$. These objectives are in turn used to calculate overall robustness and overall objectives for alternative θ. 41
3.2 WCU (a) and DU (b) optimization factor sampling schemes. The former creates states of the world by coupling synthetically generated set of streamflows with a base projection for 13 deeply uncertain factors, while the latter couples each set of synthetically generated streamflows with a different projection for the 13 deeply uncertain factors, shuffling the coupling at every function evaluation. 43
3.3 Map of the North Carolina Research Triangle region. Adapted from Zeff et al. (2014). 47
3.4 Average performance objective values for WCU and DU optimization alternatives over the 10 million SOWs from DU re-evaluation. The vertical axes denote average performance in each objective and each polyline represents an alternative. Brown tones represent alternatives from WCU optimization and blue tones represent alternatives from DU optimization. Transparency is applied to alternatives that do not meet a relaxed version of the performance criteria established by the studied utilities: Reliability $>$ 98.5% (relaxed from 99%), Restriction Frequency $<$ 20%, and Worst-Case Cost $<$ 5%). The color gradient is based on the restriction frequency objective (light for low frequency, dark otherwise, not applicable to highlighted solutions). None of the brown (WCU) alternatives meet the criteria, while 7 out of 2,954 DU solutions do. 63
3.5 Alternatives (x-axis) were rank-ordered according to robustness (y-axis) delivered to each utility. An alternative's robustness was defined as the percentage of the DU re-evaluation's 10,000 $\mathbf{\Psi}$ samples where an alternative meets the original robustness criteria elicited from water utilities. The brown bars represent the robustness of the original 84 WCU optimization's alternatives, while the blue bars represent the robustness of all 2,954 DU optimization's alternatives. Each group of bars is rank ordered by robustness. 66

3.6 Robustness tradeoffs among utilities. Each polyline represents an alternative, while the location where a polyline intersects an axes denotes the robustness value of the corresponding alternative for the corresponding utility. The blue lines indicate four illustrative compromise DU alternatives with fairly high robustness for all utilities, the brown line indicates the best compromise WCU alternative from prior work by Herman et al. (2014), while the gray lines represent the remaining DU alternatives. 67
3.7 Decision variables of the 4 compromise solutions for each utility. Different colors represent the different alternatives, and correspond to the ones in Figure 3.6. The "y" axis represent a scale from low use to high use. 68
3.8 Ranges of deeply uncertain factors which should be avoided in order to improve the robustness of each compromise solution. Demand growth (left box) and Falls Lake municipal supply allocation (right box) are the two most important sources of uncertainty for the four chosen compromise solutions. Utilities are encouraged to incorporate measures in their short-term management plan seeking to mitigate demand and be granted a greater allocation from Falls Lake. 71
3.9 Falls Lake Supply allocation vs. Demand growth pass/fail. A solutions performance is more impacted by variations in demand growth than in Falls Lake allocation. 72
3.10 Robustness improvement due to demand growth control. Demand growth management may lead to a robustness gain of up 30% for the four utilities studied here. Cary, Durham and OWASA have the potential to reach 100% calculated robustness, while Raleigh may reach over 90%. 74

4.1 WaterPaths main simulation loop. The main loop includes the (1) calculation of long-term ROF metric if at first week of the year, (2) calculation of the short-term ROF, (3) application of drought mitigation and financial policies, (4) system mass-balance modeling, and (5) system-state data collection. 95

4.2 For the calculation of the short-term ROF metric for the current week, 50 year-long simulations are ran for the coming year with all reservoir storage levels starting at their current levels, while the streamflows and evaporation rates for each simulation are those of one of the previous 50 years of recorded data and the demand are those of the previous year. For the calculation of the long-term ROF metric, the 50 year-long simulations are extended to one and a half years, with the last 6 months of the last simulation (the last year) are filled by synthetic data. The addition of 6 months to the 50 ROF simulations allows for capturing consecutive drought years in which the reservoirs are not filled during the wet season. 97
4.3 Schematics of (a) shared-memory parallelization and (b) hybrid shared- and distributed-memory parallelization for distributed heuristic optmization. Hyperthreading is not used in either panel. The collored and semi-transparent tiles represent threads, the dark solid tiles represent physical cores mounted on desktops or computing nodes (medium grey large and thin tiles serving as base for the cores), and the light grey sheet underneath the two nodes on panel (b) represents a cluster or cloud environment hosting the two depicted nodes. Blue represent the master threads, green the worker threads, and read the threads (one of which still being the master) of a master process. Threads connected by dashed black lines belong to the same process, with its own exclusive memory block, parallelized internally with OpenMP for shared-memory parallelization. Straight, full black two way arrows in panel (b) represent distributed-memory communications between four-core, four-thread master and worker processes, respectively represented by the red threads and the green threads in the yellow contour as an example. Panel (b) shows one master and seven worker processes performing a heuristic optimization run in parallel under the distributed-memory parallelization paradigm. 111
4.4 The Sedento Valley Regional Water Supply System. 117
4.5 Current Demand to Capacity ratios for the three utilities. A higher demand to capacity ratio indicates more stress on the system. 119
4.6 Policy design and analysis performed with WaterPaths within the DU Pathways framework. 127
4.7 Optimization loop and single-simulation input/output information flow. WaterPaths can be used to simulate a single policy, allowing its use for policy re-evaluation and robustness calculation, as demonstrated in Section 4.5.5, and for policy optimization adn tradeoff identification. 128

4.8 Evolution of the hypervolume for each of the nine sets. The overlap across the lines shows that the number of seeds was sufficient to assure consistency, while the gradual plateauing of the hypervolumes of all seeds shows that the number of function evaluations were enough to achieve maximum attainable convergence. 129
4.9 Performance and robustness tradeoffs in the DU Re-Evaluation space. Panel (a) illustrates the tradeoffs across the performance objectives quantified either using the mean or the worse first percentile across the tested SOWs, in which an ideal solution would be represented by a horizontal line at the bottom of the plot. Panel (b) illustrates the robustness tradeoffs across the three utilities. Here, robustness is calculated as the percent of RDM SOWs in which a policy met the performance criteria defined by the utilities and an ideal solution would be represented by a horizontal line at the top of panel (b). The highlighted policies represent the policies with best robustness for each utility (BRO-X), with the best robustness compromise across utilities (RoC), and best performance compromise (PC). 141
4.10 Robustness of pre-selected 229 policies for each utility. Each bar represents the attained robustness for one policy for each utility. The policies are sorted by robustness (bar height) for each utility. The policies with the highest value of the satisficing robustness for each utility, the policy with the best robustness compromise, and the policy with the best compromise of objectives values are highlighted in each subplot. 143
4.11 Parallel axes plots for the policy variables for the policies representing the best robustness for each utility, the best robustness compromise, and the best performance compromise. The table indicates the infrastructure construction order for each utility for each policy in the order shown in the legend. 145
4.12 Factor maps for the robustness compromise policy (RoC) showing the key uncertainty factors controlling the robustness the policy for all utilities. Success (grey dots and zone) required reliability $\geq$ 98%, restriction frequency $\leq$ 10% and annual financial worst 1st percentile cost $\leq$ 10%, while a failure to meet either criteria results in the scenario being considered a failure (red dots and region). White zones denote regions with a mixture of gray and red dots for which it is uncertain whether all criteria would be met. 150

4.13 Capacity expansion pathways and ROF dynamics for the most favorable scenario. Demand growth rate scenarios considered include the Most Favorable SOW, the SOW projected by the utilities, and the Least Favorable SOW. All infrastructure investments are illustrated as stacked color lines designating specific infrastructure investments across 1,000 hydrologic scenarios. . . 152

5.1 Map of the Research Triangle Region test case. Potential Infrastructure refers to infrastructure options presented in Triangle J (2014), that are being considered by the utilities. 164
5.2 The Research Triangle instance of the WaterPaths simulation model takes as inputs a policy to be evaluated, the system components and its connectivity, hydrological time series, infrastructure construction order, information about infrastructure to be built, different types of demands, and uncertainty. It outputs time series of system states, an objectives vector and pathways. . 166
5.3 Key steps of the DU Pathways approach and their key points . . 169
5.4 Illustrative portfolio of decisions related to short-term drought mitigation instruments and long-term priorities for infrastructure investments pathways that define a system's water supply policy. The optimization problem consists in finding the best position for each knob corresponding to each decision variable for each utility. 171
5.5 (a) In the WCU sampling scheme, sets of time-series samples of synthetically generated streamflows, evaporation rates and demands are generated for each realization and appended to the historical data. The best projected values (or other assumption) of the deeply uncertain sources of uncertainty are then coupled with the resulting sets of time series form all SOWs. (b) In DU sampling, the same sets of time-series are instead combined with samples of deep uncertainties to form the ensemble of SOWs. . . 175
5.6 Re-evaluation sampling scheme. The same sets of N_r hydrological and demand time series are combined with one vector sample of deeply uncertain factors one at a time, resulting in a number of SOWs equal to N_r times the number of samples of deeply uncertain factors. The resulting set of objectives vectors are used to back-calculate the objectives and robustness of the whole ensemble, while the pathways are used to analyze the dependency of infrastructure investments on deeply uncertainty factors. . . . 182

5.7 Comparative performance analyses based on the DU re-evaluation. Panel (a) illustrates the tradeoffs across the performance objectives quantified either using the mean or the worse first percentile across the tested SOWs. Panel (b) illustrates the robustness tradeoffs across the four utilities (satisficing metric). In both panels policies from the WCU formulation are brown and those from the DU optimization are blue. 195
5.8 Robustness of the 100 top most robust policies attained from the WCU and DU optimization schemes for each utility. Each bar represents the attained robustness for one policy for each utility. The policies are sorted by robustness (bar height) for each utility. The policies with the highest value for the satisficing robustness for each utility is highlighted each subplot. One of the most robust policy for Cary called the "supplier-preferred" policy while the most robust policies for Durham and Raleigh were called "buyer-preferred" 1 and 2. 196
5.9 Factor maps for two policies — buyer preferred 1 (Durham) and 2 (Raleigh) — showing the key factors controlling the robustness the solutions for all utilities but OWASA. Success required reliability $\geq$ 99%, restriction frequency $\leq$ 20% and annual financial worst 1st percentile cost $\leq$ 10%. A list with the combined impurity decreases for all factors can be found in Table C. 200
5.10 Parallel axis plots for the decision variables for the D/B1 and R/B2 buyer preferred policies. The red lines designate Durham's preferred buyer solution (D/B1). The green lines designate Raleigh's preferred buyer solution (R/B2). 205
5.11 Capacity expansion pathways and ROF dynamics for the Durham preferred policy DB/1. Demand growth rate scenarios considered include the Most Favorable SOW, the SOW with demand growth rates reduced to 80% of their original nominal values, and the Least Favorable SOW. All infrastructure investments are illustrated as stacked color lines designating specific infrastructure investments across 1,000 hydrologic scenarios. . . 208
5.12 Capacity expansion pathways and ROF dynamics for the Durham preferred policy DB/1. Demand growth rate scenarios considered include the Most Favorable SOW, the SOW with demand growth rates reduced to 80% of their original nominal values, and the Least Favorable SOW. All infrastructure investments are illustrated as stacked color lines designating specific infrastructure investments across 1,000 hydrologic scenarios. . . 211

B.1 System dynamics for Durham and Raleigh under policy D/B1 for multiple hydrologic realizations under Durham's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy. 231
B.2 System dynamics for Durham and Raleigh under policy D/B1 for multiple hydrologic realizations under Raleigh's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy. 232
B.3 System dynamics for Durham and Raleigh under policy R/B2 for multiple hydrologic realizations under Durham's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy. 233
B.4 System dynamics for Durham and Raleigh under policy R/B2 for multiple hydrologic realizations under Raleigh's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy. 234

D.1 Evolution of hypervolume for each master of each of two seeds ran for the WCU and DU searches. Since each seed had 4 island and each island ran 125,000 function evaluations, the total number of function evaluation for each of the 2 WCU and 2 DU seeds was 500,000. WCU seeds advanced faster at early NFE and converged slighly faster than the DU seeds, although the long plateaus at later NFE show that all seeds were given more than enough NFE to find the best Pareto approximate set of policies. . 236

E.1 Classification tree fit on illustrative data set. The first split to be foudn was the horizontal split, which applies to the entire data set. The three vertical splits were foudn branching off of the horizontal split and then off of each other. 237
E.2 Sequence of addition of single-split classification trees until all points are correctly classified. Panel (a) shows just the data, without trees, panel (b) shows data with one tree, panel (c) with two, and panel (d) with three. Variables x_1 and x_2 are two uncertain factors, the blue crosses represent a success scenario and a red circle represents a failure, and the blue and red regions correspond to regions classified as success or failure regions, with intermediate colors indicating uncertainty in the classification. . . 240

CHAPTER 1

INTRODUCTION

1.1 Current Challenges for Water Infrastructure Planning and Management

Water utilities globally are facing increasing pressure to proactively address growing demands, higher levels of resource contention, and increasingly uncertain water availability (Bonzanigo et al., 2018; Hall et al., 2014). A tension exists between improving the efficiency and coordination of regional water supplies to delay or eliminate major infrastructure investments and the eventual effects of demand growth rates that necessitate water supply capacity expansions. Soft water management strategies (conservation, inter-utility water transfers, demand management, etc.) are a key option to cope with growing domestic water demand (Gleick, 2002a, 2003). In the United States (US), a key driver increasing the importance of soft water management strategies is that most large projects of federal and/or state interest have already been built (Lund, 2013). However, growing urban water demands and increasingly uncertain climate conditions are motivating efforts on the hard side, meaning redesigning system capacity by upgrading existing and adding new infrastructure to maintain reliable water supply systems (Shafer and Fox, 2017).

The hard and soft path views for urban water portfolio planning should be treated as complementary, as even though infrastructure expansion may be needed, it is not desirable and may be mitigated by soft water management strategies. However, both the soft and hard approaches may be detrimental to utilities' finances. Together with outstanding debt, decreasing cash flows are significant factors that negatively impact utilities credit ratings, which can increase the cost debt when seeking to build new infrastructure (Moody's, 2017; Zeff et al., 2014; Zeff and Characklis, 2013). The financial impacts of drought mitigation strategies, the current negative financial outlook for the water utility sector (Moody's, 2019), and the lack of appropriate financing for water utilities (Gleick et al., 2014) all stress the need for new water infrastructure planning and management frameworks. More specifically, utilities and planners would benefit from a framework to facilitate improved water portfolio planning that effectively combines short-term drought mitigation, instruments for managing financial risks, and long-term infrastructure investment pathways while accounting for diverse sources of uncertainty.

Such framework needs to address the following key questions: (1) how should regional decision makers devise infrastructure investment and drought management strategies that adequately capture their interdependencies to improve water supply reliability and financial stability?, (2) how can stakeholders develop management and investment policies that are sufficiently adaptive and robust over decades given the impossibility of predicting climate, population shifts, political landscapes, and other sources of uncertainty?, (3) how do decision-makers know when to make substantial and irreversible investments in new infrastructure while avoiding stranded assets?, and (4) how should water utilities sharing water resources (e.g. watersheds, reservoirs, pipes) coordinate individual actions to avoid exacerbating vulnerabilities and conflicts?. These questions capture the defining challenges of the water utility planning and management problem.

1.1.1 The Water Portfolio Planning and Management Problem

Several studies have focused on integrating efficient short-term management strategies (soft path) into existing infrastructure to improve supply reliability at lower expected costs Helping to avoid unecessary infrastructure expansion (Lund and Israel, 1995; Wilchfort and Lund, 1997; Palmer and Characklis, 2009; Padula et al., 2013; Mozenter et al., 2018). Recent studies have explored demand management (Hall et al., 2019; Zeff et al., 2014; Borgomeo et al., 2018; Huskova et al., 2016; Haasnoot et al., 2013), treated water transfers (Caldwell and Characklis, 2014; Zeff et al., 2014; Watkins and McKinney, 1997; Hall, 2019), upstream reservoir releases (Borgomeo et al., 2018; Gorelick et al., 2019), and operation rules (Wu et al., 2017; Tian et al., 2018). Efforts over the last several years focused on regionalized management seek to identify coordinated management operations where the impacts from the measures taken by each utility has a predictable impact on the operations of the other utilities (Sheer, 2010; Zeff et al., 2014; Herman et al., 2014; Zeff et al., 2016; Borgomeo et al., 2018; Gorelick et al., 2019; Hall et al., 2019). For example, if multiple utilities rely on treated or raw water transfers as a drought mitigation measure for improving reliability, each transfer request has the potential to be directly impacted by constrained conveyance and water treatment capacities. The short-term water portfolio problem refers to how to identify the mixtures of actions that can be taken by the individual cooperating utilities that improve water-supply reliability and financial stability for the whole region.

Despite aiding utilities to attain higher levels of supply reliability, the use of drought management instruments such as water transfers or restrictions are likely be detrimental to utilities' finances due to potentially significant increases in the variability of their costs (Watkins and McKinney, 1997; Zeff et al., 2014). Together with outstanding debt, decreasing cash flows are significant factors that negatively impact utilities credit ratings, which can increase the cost of borrowing to build new infrastructure (Moody's, 2017; Zeff et al., 2014; Zeff and Characklis, 2013). This emphasizes the need to study existing and develop of new financial instruments, such as index insurance against droughts, contingency funds, and pricing structures to hedge water utilities' finances against droughts (Zeff and Characklis, 2013; Lopez-Nicolas et al., 2018). The financial impacts of drought mitigation strategies, the current negative financial outlook for the water utility sector (Moody's, 2019), the lack of appropriate financing for water utilities (Gleick et al., 2014), and the strong impact of operating decisions on the appropriateness of capacity expansion (Tian et al., 2018) stress the need for new frameworks that effectively combines short-term drought mitigation instruments backed by financial instruments and long-term infrastructure investment planning.

Under extreme demand increase or droughts, even well designed soft-path solutions and financial instruments may not eliminate the need to invest in the costly hard-path approach (Gleick, 2003; Hall et al., 2019). Furthermore, aging infrastructure may further exarcebate the need for costly investments in infrastructure (Hall et al., 2019), making even more pressing the need for a methodology focusing on minimizing investments in infrastructure without compromising water-supply services. Acknowledging this need, researchers have been developing methodologies to balance reliance on the soft- and hard-path approaches for addressing water supply planning and management. One example is the real option analysis (ROA) (Cox et al., 1979) literature of the last 15 years (Wang and de Neufville, 2005; Erfani et al., 2018; Fletcher et al., 2017; Hui et al., 2018; Fletcher et al., 2019), which is based on flexible decision-tree logic based on local conditions and found through single-objective optimization (normally cost) that emphasizes near-term information to inform decisions about infrastructure planning and policy implementation. Other researchers have presented frameworks that consider multiple objectives while seeking to recommend a sequence of infrastructure investments coupled with soft-path strategies (Zeff et al., 2016; Borgomeo et al., 2018; Huskova et al., 2016; Beh et al., 2015a; Mortazavi-Naeini et al., 2014). Alternatively, the Dynamic Adaptive Policy Pathways approach (DAPP) takes pre-specified adaptive pathways that combine management and infrastructure expansion actions based on monitored action triggers over time (Haasnoot et al., 2013; Kwakkel et al., 2014; Kingsborough et al., 2016, 2017; Zandvoort et al., 2017). Across all of these frameworks, one important and still unresolved issue is expanding their capacity to address a broad range of uncertainties.

1.1.2 Considering Uncertainty

Hydro-climatic uncertainties such as near-term streamflows and evaporation rates affect both short-term management and long-term planning problems, and have been extensively studied and characterized with well-established statistical techniques (Stedinger, 1993; Martins and Stedinger, 2000; Herman et al., 2016; Lall, 1995; Kirsch et al., 2013). However, water infrastructure investment and management problems include a multitude of uncertainties referred to hereafter as “deep uncertainties” (Kwakkel et al., 2016b; Knight, 1921). Deep uncertainties cannot, as described by Knight (1921, p. xiv), be treated as “a gamble on a known mathematical chance”. In other words, they lack consensus on their underlying probability distributions as well as on individual or collective outcomes. Such uncertainties have also been termed Knightian, or scenario uncertainties, by multiple authors, including Walker et al. (2013); Lempert et al. (2006) and Kwakkel et al. (2016b). Examples include regional demand growth, political and economic uncertainties, and climate change (Herman et al., 2014; Milly et al., 2008).

Uncertainty is often included in water utility planning and management exercises by sampling or designing representative states-of-the-world (SOWs) that shape water supply risks (Watson and Kasprzyk, 2016; Kwakkel et al., 2014), although how to perform such design or sampling is an important decision in any drought mitigation application, defining perceptions and preferences related to key performance objectives, vulnerabilities, and/or the robustness of a system. Fundamental to this challenge is the choice of what uncertain factors should be sampled or included in generating alternative scenarios. Next is a review of existing frameworks, how they solve the water utility planning and management problem, and how they handle the underlying uncertainties.

1.1.3 Existing Water-Infrastructure Planning and Management Software and Decision-Making Framworks and Their Limitations

There is an extensive body of literature dealing with water infrastructure planning and management under uncertainty. Broadly, there are three dominant classes of methodologies that are at present used in water supply planning contexts: (1) pre-specified deterministic scenarios/narratives, (2) scenario trees, and (3) globally sampled stochastic scenario analysis. As an example of deterministic scenario-based planning, Liu et al. (2008) analyzed four semi-arid river basins in the south-western US seeking to find water management strategies based on integrated modeling that would minimize lasting or irreversible environmental impacts (mostly on local vegetation) while allocating enough water for competing uses. The authors analyzed a pre-specified set of scenarios based on their *a priori* specified problem definition and input from stakeholders as well as water management experts. Scenario trees is used in ROA approaches such as Erfani et al. (2018), in which scenarios are created simulating a stochastic process so that as the simulation of future decisions progresses, all future possibilities share a common past. As an alternative example of stochastic scenarios, Brown et al. (2012), Moody and Brown (2013), and Zeff et al. (2016) use ensembles of reservoir inflows estimated using regression techniques to estimate and assess climate vulnerability. Climate vulnerability analysis focuses on reducing risks, or losses, calculated as the product of a loss function times subjective probabilities for the climate scenarios. The probabilities are elicited by contrasting a mixture of global climate model projections, paleodata (if available), and synthetically generated weather. These approaches are termed Decision Scaling, and have been recently expanded by Ghile et al. (2014) and Lownsbery (2014) to include hydroeconomic sensitivities.

The Decision Scaling example relates to a broader body of literature focused on planning under deep uncertainty. Exploratory modeling (Bankes et al., 2001), which involves running planning models across broad global samples of hypothesis, has emerged as a primary strategy for discovering which combinations of deeply uncertain factors are of high consequence in decision problems, as opposed to considering one most likely future projection (Walker et al., 2003; Lempert et al., 2006; Walker et al., 2013; Zeff et al., 2014; Herman et al., 2014). As noted by Dessai et al. (2009), the intent of exploratory modeling is to shift focus from predicting future conditions to understanding which conditions lead to decision relevant consequences. Given a better understanding of the consequences related to these discovered conditions, decision makers can then evaluate their relative importance or plausibility.

As reviewed by Herman et al. (2015), the selection of decision alternatives to analyze in exploratory modeling frameworks can either be done *a priori* or with computational search. In the case of decision alternatives that are specified *a priori*, the decision analysis then takes the form of a classical discrete choice among those alternatives based on their interpreted robustness. This approach has been commonly employed across the full range of decision support frameworks in this area (Brown et al., 2012; Moody and Brown, 2013; Hipel and Ben-Haim, 1999; Lempert et al., 2006; Bryant and Lempert, 2010). Conversely, a growing number of studies are using search strategies to discover high performing design alternatives in their initial stage of analysis and then evaluating their robustness to inform design selection (Korteling et al., 2013; Kasprzyk et al., 2013; Hamarat et al., 2014; Kwakkel and Haasnoot, 2015; Giuliani and Castelletti, 2016; Singh et al., 2015; Zeff et al., 2014; Herman et al., 2016). None of the questions presented thus far are new, although utilities and researchers still have difficulties in answering them through existing water-infrastructure planning and management frameworks (Bonzanigo et al., 2018; Paulson et al., 2018).

Considering the frameworks coupling soft- and hard-path approaches, the

ROA is limited in its accounting of stakeholders with diverse interests as a single-objective approach and the scope limiting challenges in addressing a broader suite of of uncertainties pose by its curse of dimensionality (Dittrich et al., 2016). Borgomeo et al. (2018) address some of these concerns by proposing a multi-objective evaluation of candidate water supply investments. They maximize robustness and minimize risk and cost of a 30-year long fixed construction schedule encompassing multiple infrastructure options. However, the derived sequence of investments is not adaptive and does not account for evolving information feedback. This limitations also exists in the recent studies by Beh et al. (2015b) and Huskova et al. (2016), although the former seeks to improve adaptivity by incorporating a similar emphasis on near term information as done in ROA. Alternatively, the DAPP approach (Haasnoot et al., 2013; Kwakkel et al., 2014; Kingsborough et al., 2017) adds flexibility to the planning process by making use of potentially diverse metrics (signposts) to inform adaptive action pathways. DAPP, however, still maintains a strong reliance on the use of limited numbers of predefined action sequences that require high levels of institutional stability and coordinated consensus over the long term (30 to 100 years in case studies).

The adaptive planning frameworks reviewed above all must confront issues related to institutional uncertainty. They all make the "omniscient and omnipotent lawgiver" assumption (Walker et al., 2001), ignoring all inherent "policy myopia" (Nair and Howlett, 2017). This is similar in effect to the "planning fallacy" described by Scruton (2013), in which a regimented plan prevents the necessary use of live information to effectively realize the needs that motivated its implementation in the first place. Additionally, as pointed by Scruton (2013) and Bosomworth et al. (2017), decades-long plans inherently assume that current goals and values are uncontested and will remain so for the duration of the plan, together with the required institutional context. These limitations, acknowledged by (Borgomeo et al., 2018; Beh et al., 2015b) though not addressed, make the case for a framework that maximizes the use of current data and minimizes the over-specification of decisions to be made in the future. Lastly, when commenting on the DAPP approach Bosomworth et al. (2017) makes the broadly valid point that the lack of flexibility in an approach may clash with case-specific institutional issues that may render the approach unfeasible. More specifically, planning approacha lacking adaptivity require higher degrees of abstraction of the particularities of individual and concrete institutional contexts, some of which in direct conflict with the structural assumptions that are core to planning frameworks. In light of these limitations, an approach that takes in current information at short time scales throughout the life of the plan and fits within the annual budget approval of most municipalities is desirable.

Zeff et al. (2016) introduced an integrated infrastructure planning and management approach, in which infrastructure investment and water portfolio management policies specify short- and long-term decision-making rules based on the risk-of-failure metric (ROF), as well as candidate orderings forsequencing of infrastructure investments. By combining infrastructure planning and management rules in a single policy, the infrastructure pathway approach aids the design of drought mitigation instruments that are carefully coupled with new infrastructure investments to minimize to help better balance system reliability and reduce debt burden. Despite the high adaptivity on both the short and long term, the infrastructure pathways framework introduced Zeff et al. (2016) is limited in its treatment of uncertainties. The limitation comes from its reliance on creating SOWs based only on well-characterized stationary hydro-climatic uncertainties while using expert-elicited assumed values for deep uncertainties (e.g., demand growth rates, pricing changes, behavioral responses to water restrictions, etc.), limiting the scope of the analysis and potentially limiting the robustness of the resulting policies. Furthermore, although the infrastructure pathways framework introduced by Zeff et al. (2016) proposes a fixed suite of major candidate investments over decades to come making it still subject to the same challenges that emerge changes in societal or institutional context, the plans proposed by this approach includes mechanisms for incorporating environment and economical information throughout its life, adding an opportunity for endogenous learning. Therefore, the joint-infrastructure pathways and water portfolio management introduced by Zeff et al. (2016) serves as the basis for the Deeply Uncertain (DU) Pathways methodology and for the WaterPaths simulation software, both developed in this thesis research.

1.2 Dissertation Objectives

This dissertation contributes the WaterPaths open-source stochastic simulation software and the DU Pathways methodology to reduce the computational and conceptual barriers to integrated infrastructure investment pathways and water portfolio management. In combination, WaterPaths and the DU Pathways framework introduced in this dissertation provide the first formal integration of dynamic and adaptive water infrastructure pathways, financial and supply- and reliability-focused drought-mitigation instruments, and the recent extensions of the Many-Objective Robust Decision-Making (MORDM) framework that incorporate deep uncertainties in search-based identification candidate actions (Kasprzyk et al., 2013; Kwakkel et al., 2014; Watson and Kasprzyk, 2016). For this, WaterPaths incorporates decision making at two time scales: weekly, for operational management decisions, and annual, for infrastructure investment decisions. Expanding on the approach recommended by Zeff et al. (2016), the integrated infrastructure investment and management policies designed with DU Pathways specify short- and long-term decision-making rules based on the risk-of-failure metric (ROF), as well as candidate orderings for infrastructure investments. By combining infrastructure planning and management rules in a single policy, the DU Pathways approach aids, here through its use of WaterPaths, the design of drought mitigation instruments tailored to new infrastructure to minimize the need of further investments.

To assure a broad uncertainty analysis, the DU Pathways methodology exploits WaterPaths to account for multiple well-characterized and deep uncertainties in each function evaluation during search, improving the search-based identification of candidate infrastructure investment and management policies. Additionally, the framework formalizes a detailed evaluation of the multi-city policies' robustness tradeoffs. DU Pathways aids stakeholders in navigating these tradeoffs through a combination of visual decision analytics (Woodruff et al., 2013) and enhanced scenario analysis in which boosted trees (Schapire, 1999) aid in the identification of the uncertainties most strongly shape the infrastructure investment and management policies' vulnerabilities. The application of the DU Pathways framework is demonstrated on the Research Triangle test case (Zeff et al., 2016; Herman et al., 2014; Zeff et al., 2014), where this dissertation will explore how the region's vulnerabilities evolve under alternative policies, identify significant interdependencies between the region's four utilities, and show the importance of carefully balancing supply and financial instruments.

1.3 Scope and Organization

1.3.1 Chapter 2 — Background on Methodological Components

This chapter presents an overview of computational methods that used in all of the subsequent chapters. These are the Borg multiobjective evolutionary algorithm (Borg MOEA) and the synthetic streamflow, evaporation rate, and demand time-series scenario generation technique. The Borg MOEA is used in all of the subsequent chapters to discover Pareto-efficient infrastructure investment and water portfolio management policy alternatives, while the time-series generation methods were used to generate SOW to evaluate performance.

1.3.2 Chapter 3— Reducing Regional Drought Vulnerabilities and Multi-city Robustness Conflicts Using Many-objective Optimization Under Deep Uncertainty

Chapter 3 introduces a multi-stakeholder extension of the MORDM framework that informed the formal model base development for WaterPaths (Chapter 4) and the broader DU Pathways framework (Chapter 5). This chapter focused on better accounting for deeply uncertain factors when identifying cooperative drought management strategies based on soft-path approaches for cooperative regional drought management. The analysis is applied in the Research Triangle area of North Carolina, with a focus on the water utilities within Raleigh, Durham, Cary and Chapel Hill. Prior analysis of this region through the year 2025 has identified significant regional vulnerabilities to volumetric shortfalls and financial losses. Moreover, efforts to maximize the individual robustness of any of the mentioned utilities also have the potential to strongly degrade the robustness of the others. The results in Chapter 3 show that appropriately designing adaptive and high-efficacy drought mitigation instruments required stressing them with a comprehensive sample of deeply uncertain factors in the computational search phase of MORDM. Search under the new ensemble of SOWs is shown to fundamentally change perceived performance tradeoffs and substantially improve the robustness of individual utilities as well as the overall region to water scarcity. The results in Chapter 3 show that search under deep uncertainty enhanced the discovery of how cooperative water transfers, financial risk mitigation tools. The results highlight that coordinated regional demand management employed jointly strongly improves regional robustness and decreases robustness conflicts between the utilities. Insights from this work have general merit for regions where adjacent municipalities can benefit from cooperative regional water portfolio planning.

1.3.3 Chapter 4 — WaterPaths: A Stochastic Simulation Framework for Water Supply Infrastructure Investment Management

Chapter 4 contributes the WaterPaths model: an open-source simulation software that implements, extends, and generalizes the short-term water portfolio management problem and extended uncertainty sampling presented in Chapter 3 and combines it with the integrated infrastructure pathways presented in Zeff et al. (2016). WaterPaths is designed to simulate stochastic risk and demand-based sequencing of long-term regional water infrastructure construction in parallel with short-term operational decisions (e.g., restrictions, transfers, financial instruments, etc), bridging the gap between weekly management and yearly planning for water utilities. The solution framework is designed to flexibly represent new conservation and financial mitigation strategies, such as schemes for allocating raw and treated transfer water among different utilities or different ways of structuring drought insurance policies. Moreover, Chapter 4 introduces a hypothetical Sedento Valley test case to demonstrate WaterPaths capabilities. The Sedento Valley test case is proposed as a new benchmark problem for testing optimization algorithms to be used in water system engineering and aid the development of improved planning methodolgies. The code, written in C++, makes use of parallelization strategies to perform scalable ensemble-based exploratory analyses ranging from small tests on desktops to larger runs on clusters or cloud computing platforms. It should be noted that the software development leading to WaterPaths occurred simultaneously with the formalization of the DU Pathways framework illustrated for the Research Triangle region in NC in Chapter 5, resulting in the formulations of both being partly indissociable. Therefore, a small amount of forward-referencing in Chapter 4 was innevitable.

1.3.4 Chapter 5 — Deeply Uncertain Pathways: Integrated Multi-City Regional Water Supply Infrastructure Investment and Portfolio Management

Chapter 5 contributes the Deep Uncertainty (DU) Pathways framework, in which WaterPaths is used to combine ROF decision rules presented in Chapter 3, dynamic adaptive policy pathways concepts, and a careful consideration of time-evolving information feedbacks to yield management-conditioned infrastructure pathways for regions. The DU Pathways' framework has been developed to carefully consider multi-actor regional contexts with the goal of aiding stakeholders in discovering pathway policies that attain high performance levels across challenging, deeply uncertain futures and to guide robustness compromises that may be necessary between regional actors. As demonstrated again in the Research Triangle region of North Carolina, DU Pathways clarifies how to identify robust infrastructure investment and management policies across the municipalities of Raleigh, Durham, Cary, and Chapel Hill. Our results provide insights about the most cost-effective infrastructure options to be pursued in the near-term, and clarify which sources of uncertainty drive the performance and robustness of the regional system and of the individual actors.

1.3.5 Chapter 6 — Contributions and Future Work

Finally, Chapter 6 summarizes the contributions of this dissertation to the body of existing knowledge in the field of water resources systems and planning. It also provides recommendations for future work, such as features to be added to WaterPaths and for different treatments of deep uncertainty for water resources planning and management.

1.3.6 Author Contributions for Published Work

Chapter 3: borrowing from previous studies by H.B. Zeff and J.D. Herman under the supervision of Professor Patrick Reed and Gregory Characklis, Bernardo Trindade conceived the studies, modified the inherited model, performed the optimization runs and data analysis, and wrote the paper under the supervision of Professors Patrick Reed and Gregory Characklis.

Chapters 4: borrowing from previous studies by H.B. Zeff and J.D. Herman under the supervision of Professor Patrick Reed and Gregory Characklis, Bernardo Trindade conceived the study, wrote a new model, performed the optimization runs and data analysis, and wrote the paper under the supervision of Professors Patrick Reed and Gregory Characklis.

Chapters 5: borrowing from previous studies by H.B. Zeff and J.D. Herman under the supervision of Professor Patrick Reed and Gregory Characklis, Bernardo Trindade conceived the study, implemented the test case conceived in partnership with David Gold, performed the optimization runs and data analysis, and wrote the paper together with David Gold under the supervision of Professors Patrick Reed and Gregory Characklis.

CHAPTER 2

METHODOLOGICAL COMPONENTS

2.1 Borg Multi-Objective Evolutionary Optimization Algorithm

In the recent literature, multiobjective evolutionary algorithms are emerging as a primary tool used for solving water supply infrastructure management and/or planning problems (Zeff et al., 2016; Quinn et al., 2018; Kwakkel et al., 2014; Borgomeo et al., 2016, 2018; Beh et al., 2015a). Zeff et al. (2016) employed multiobjective evolutionary algorithms in the context of the Research Triangle to identify water infrastructure management and planning policies effective across an ensemble of synthetically generated streamflows, evaporation rates, and demands (all well-characterized uncertainties). Such policies combine water transfers, financial risk instruments, demand management and infrastructure construction strategies. Seeking to find more robust candidate solutions, Kwakkel et al. (2014) and Watson and Kasprzyk (2016) explored simpler problem formulations while also sampling deep uncertainties in their search for optimal tradeoff policies. Combining the problem formulation presented in Zeff et al. (2016) and sampling strategies similar to Kwakkel et al. (2014) and Watson and Kasprzyk (2016) as desired in this dissertation creates a potential for more effective policies but in turn increases the complexity of the problem. Mathematically speaking, the policy structure, based mostly on discrete decisions and on permutations of infrastructure options, makes the resulting objectives (performance metrics) non-convex and discontinuous functions of the policy variables, while the uncertainty sampling adds a high degree of non-linearity to them.

In order to handle these difficulties and the high number of objective functions (up to six) of the the water portfolio planning problem, we used in the following chapters the master-worker and multi-master variants of the parallel Borg multi-objective evolutionary algorithm (MS- and MM-BorgMOEA) (Hadka and Reed, 2013; Hadka et al., 2013; Hadka and Reed, 2014). All versions of the Borg MOEA have demonstrated high levels of performance when applied to multiobjective problems in a variety of applications, including water supply portfolio planning, pollution control given ecological thresholds, groundwater monitoring design, and reservoir control (Bode et al., 2019; Reed et al., 2013; Hadka et al., 2012; Ward et al., 2015; Zatarain Salazar et al., 2016). Figure 2.1 shows a schematic flow chart of the main loop of the serial version (original) of the Borg MOEA.

Initially, a number of solutions (here infrastructure planning and management policies) are randomly generated and stored in the population (top-left box) after being checked for Pareto-dominance, described in Section 2.1.2. Next, an operator (dashed ellipses) is chosen to generate an offspring solution based on $k = 2, 3$ from the population and potentially of the archive, as explained in Section 2.1.1. This offspring is then evaluated with a model (here WaterPaths) to be potentially circled back to the population and/or archive, as described in Sections 2.1.2 and 2.1.3, closing the loop which is performed a pre-specified number of times. The following subsections will introduce some of the key concepts behind Borg MOEA highlighted in Figure 2.1 and both parallelization strategies.

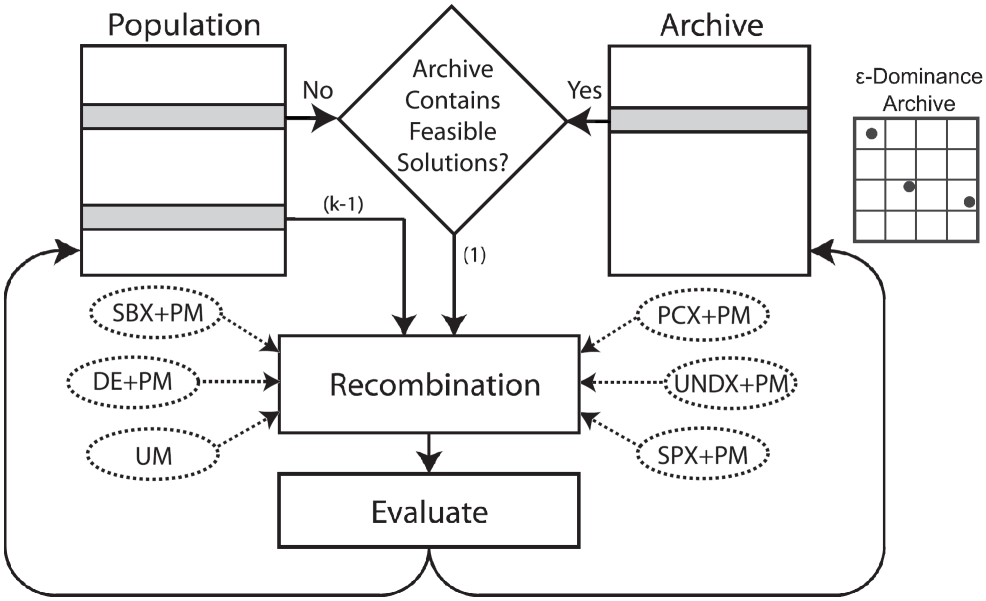

Figure 2.1: Flowchart of the Borg MOEA main loop adapted from (Hadka and Reed, 2013). First, one of the recombination operators is selected based on its performance. After, for an operator requiring k parents to generate an offspring, Borg selects $k - 1$ from the population and one from the archive. The resulting offspring is evaluated and considered for inclusion in the population and archive.

2.1.1 Adaptive Operator Selection

In order to be a suitable algorithm across a wide range of problems, Borg MOEA relies on adaptive operator selection (Vrugt and Robinson, 2007) based on probabilities calculated from solutions in its archive (Hadka et al., 2013) . An evolutionary optimization operator is a function that combines aspects of parent solutions to create an offspring solution in a semi-randomized fashion. The six operators implemented in Borg MOEA, shown in Figure 2.2, are: simulated binary crossover (Deb and Agrawal, 1994), differential evolution (Storn and Price, 1997), parent-centric crossover (Deb et al., 2002a), unimodal normal distribution crossover (KITA et al., 2000), simplex crossover (Tsutsui et al., 1999), polynomial mutation (PM), and uniform mutation (UM). Figure 2.2 shows the distributions of offspring solutions (small dots) produced from parent solutions (big dots) by all six operators in plemented in the Borg MOEA.

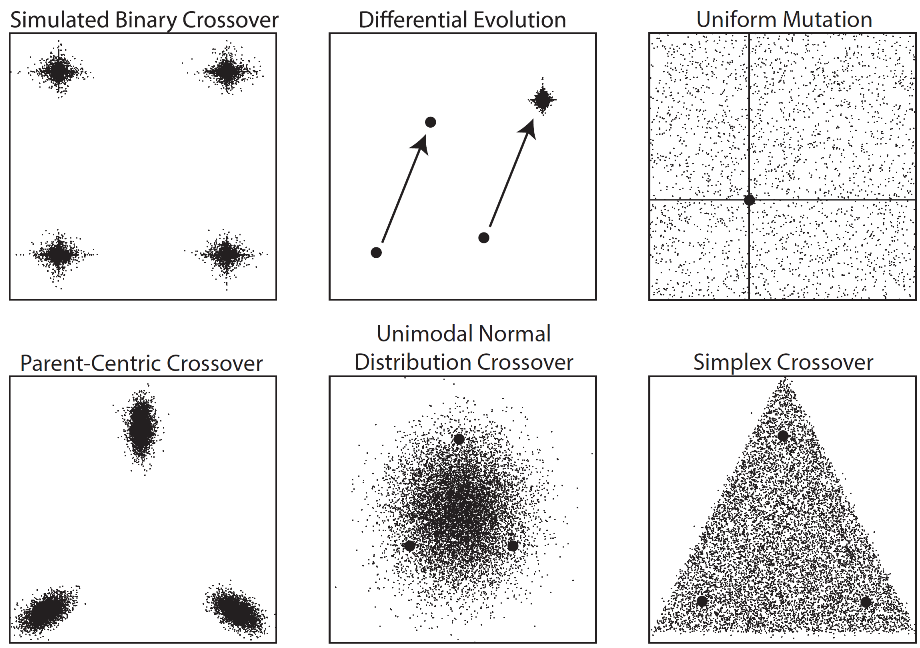

Figure 2.2: Example distributions of offspring solutions (small dots) produced from parents solutions (large dots) by the six operators included in Borg MOEA. From Hadka and Reed (2014).

At the beginning of an optimization run, all of Borg MOEA's K operators (here $K = 6$) are assigned equal probability $Q_i = 1/K$ of being selected to generate a new offspring solution. However, as the search progresses, meaning after a few iterations over the loop in Figure 2.1, the selection probabilities Q_i of each operator are adjusted according to how many Pareto-efficient solutions each operator has generated until then Hadka and Reed (2013); Hadka et al. (2013); Hadka and Reed (2014). More specifically, the probabilities Q_i are updated over time with a higher probability assigned to the operators with more solutions admitted into the archive at that stage of the search, as shown in Equation 2.1.

$$Q_i = \frac{C_i + \varsigma}{\sum_{j=1}^{K}(C_j + \varsigma)} \tag{2.1}$$

where C_i is the number of solutions generated by operator i that was admitted into the archive and ς is a constant to ensure even operators that are performing poorly at a given time have a small probability of being chosen—Hadka and Reed (2013) suggests $\varsigma = 1$. This allows the Borg MOEA to not only select operators that provide better offspring on each region of the solution space based on mathematical properties of both, but also to focus the later stages of the search on operators better at exploitation rather than exploration for the problem at hand, given sufficient search time. The offspring solutions generated by the operators may be discarded or stored in the population and/or archive, as described next.

2.1.2 Solution population

The Borg MOEA has two groups of solutions stored in memory at a given time: the population and the archive. The population is the most current pool of strongly non-dominated solutions, found by the Borg MOEA. A Pareto-efficient solution is one that perfoms better than all solutions in a set in a pairwise comparison in at least one objective. A set of Pareto-efficient solutions is called a Pareto set or Pareto front. Below is a definition of strong non-dominance:

$$F(\boldsymbol{a}) \succ F(\boldsymbol{b}) \iff F_i(\boldsymbol{a}) \leq F_i(\boldsymbol{b}) \;\; \forall i \in \{1, ..., D\} \;\wedge\; \exists i \; s.t. \; F_i(\boldsymbol{a}) < F_i(\boldsymbol{b}) \tag{2.2}$$

where F is a vector of D objectives in an all-minimization problem, and $\boldsymbol{a}$ and $\boldsymbol{b}$ are different vectors of decision variables. In other words, $\boldsymbol{a}$ strongly dominates $\boldsymbol{b}$ if all objective values in $F(\boldsymbol{a})$ are equal or smaller than in $F(\boldsymbol{b})$ and at least for one objective i, $F_i(\boldsymbol{a})$ is strictly smaller than $F_i(\boldsymbol{b})$ (Pareto, 1896; Coello et al., 2006).

A new offspring solution is admitted into the archive if it strongly dominates one or more population members, in which case replaces a randomly selected solution it dominates. If this offspring is dominated by at least one solution in the population (i.e. if it performs worse in every objective than at least one solution in the population), it is discarded. If it does not dominate nor is dominated by any solution in the population, it is either added to the population or replaces a member of the population at random so that the size of the population is maitained smaller than 125% of the size of the archive—and never smaller than 100% either.

2.1.3 Solution Archive and ϵ-Dominance

As explained in Section 2.1.2, Pareto-efficient solutions may be randomly deleted from the population, potentially resulting in the loss of a good solution. The goal of archive is to keep Pareto-efficient solutions on record in case of one is deletied from the population. The archive is a set of Pareto-efficient solutions found by the Borg MOEA from the start of the optimization process and kept on record until found to be dominated by a new solution. For an offspring solution to be included in the archive, it must not be dominated by any solution already in the archive by margins ϵ_i for all objectives $i \in \{1, ..., D\}$. This is because following the definition in Equation 2.2, it is possible for set of non-dominated solutions to contain dozens of thousands (up to an infinite number) of solutions, which is not desirable for decision-making purposes. Such stricter type of dominance, called ϵ-dominance (Laumanns et al., 2002), is used to keep the size of the archive resonably small (ideally a few hundreds) by preventing solutions non-dominated by a negligible fraction of the value of an objective from being included in the archive. In ϵ-dominance, solutions are placed on a grid as the one shown in Figure 2.3 and checked for dominance for each objective i according to ϵ_i tolerances. Laumanns et al. (2002) defined ϵ-dominance as:

$$F(\boldsymbol{a}) \succ_\epsilon F(\boldsymbol{b}) \iff (1-\epsilon_i) \cdot F_i(\boldsymbol{a}) \leq F_i(\boldsymbol{b}) \;\; \forall i \in \{1, ..., D\} \wedge \exists i \; s.t. \; (1-\epsilon_i) \cdot F_i(\boldsymbol{a}) < F_i(\boldsymbol{b}) \tag{2.3}$$

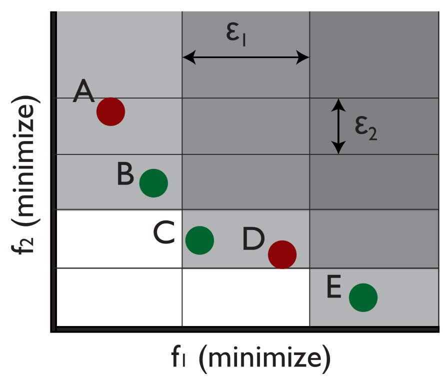

Figure 2.3: Illustration of two-objective optimization problem on an ϵ grid, adapted from Kasprzyk et al. (2013). Only one solution within each tolerance band ϵ_1 and ϵ_2 is kept in the Pareto set. Following the concept of ϵ-dominance, B, C, and E are the only non-dominated points.

The goal to be attained by running the Borg MOEA is therefore to find a best approximation of the true Pareto front for the problem, called the Pareto-approximate front. Beyond constraining the size of an Pareto-approximate solution set, the use of ϵ-dominance also guarantees two desirable properties of a multiobjective optimization algorithm given sufficient search time: convergence to the true Pareto-efficient solution set and diversity across the set. All ϵ_i values are specified by the user, observing the constraint $\epsilon_i > 0 \; \forall i \; \in \{1, ..., D\}$ (Laumanns et al., 2002).

2.1.4 Parallelization Strategies for the Borg MOEA

The Borg MOEA was originally developed as a serial search tool where function evaluation is done in the sequence for each member of its search population (Hadka and Reed, 2013). To make use of modern parallel computing clusters, two parallelization strategies have been developed for it in recent years: the Master-Worker Borg MOEA (MS-Borg MOEA) and the Multimaster Borg MOEA (MM-Borg MOEA), both of which are used in the subsequent chapters of this dissertations. In the MS-Borg MOEA algorithm, the master process distributes function evaluations across worker processes, each running on one or on a group of cores (processors), allowing for the search to be run in parallel. This distribution of work in MS-Borg MOEA is shown in Figure 2.1.4a.

On the other hand, the multi-master variant of the Borg MOEA generalizes the original Borg MOEA's adaptive operators across coordinating instances of the master process each with its own suite of workers, as shown in Figure 2.1.4b. This hierarchical parallelization scheme (Cantu-Paz, 2000; Hadka and Reed, 2014) takes advantage of both the multiple population and master-worker parallelization strategies to increase scalability and the difficulty of problems that can be addressed (Tang et al., 2007; Reed and Hadka, 2014). The MM-Borg MOEA's coordinating instances automatically detect search stall and share archived solutions and operator probabilities to enhance the efficiency and effectiveness of

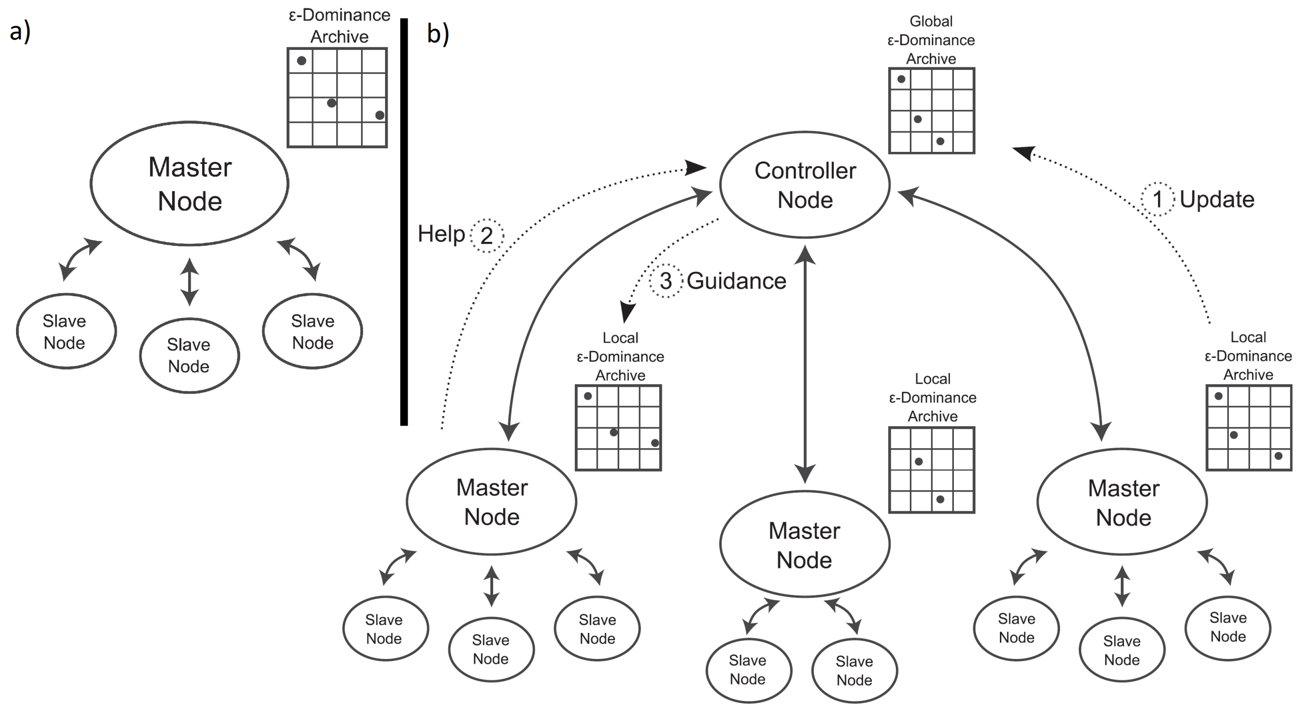

Figure 2.4: Panel (2) shows a diagram of the master-worker implementation of the Borg MOEA. The master node maintains the ϵ-dominance archive and runs the main loop of the serial Borg MOEA shown in Figure 2.1. The decision variables generated through the six operators inplemented in the master node are transmitted to the worker nodes, and the evaluated objective function values and constraints are returned to the master node. Panel (b) shows a diagram of the multi-master implementation of the Borg MOEA. The multi-master Borg MOEA consists of two or more mastere-worker instances (the diagram depicts three) periodically communicating with one controller node, which maintains a global ϵ-archive. (1) Each master node periodically transmits its local ϵ-dominance archive to the controller, which updates the global archive. (2) When a master node is not advancing its pareto-approximate front, it sends a help message to the controller. (3) The controller responds with guidance, which includes the global ϵ-dominance archive and global operator probabilities. From Hadka and Reed (2014).

search for difficult problems (Giuliani et al., 2018; Quinn et al., 2017, 2018). The next section explains the synthetic streamflow generation engine used to sample streamflows and evaporation rates for simulating the water portfolio planning problem.

2.2 Scenario Generation for Streamflow and Evaporation Rates Time Series

All synthetic inflows and evaporation rates for this work were generated using the methodology proposed by Kirsch et al. (2013). More specifically, this process was used initially to generate cross-site-correlated matrices with time series of log-normally distributed synthetic weekly inflows $\boldsymbol{NI}_{s_k} \in \mathbb{R}^{N_{ys} \cdot 52 \times N_r}$ for each reservoir and gauge k. Secondly, it was used to generate the corresponding matrices $\boldsymbol{E}_{s_k} \in \mathbb{R}^{N_{ys} \cdot 52 \times N_r}$ of evaporation rates for reservoirs only. Given the procedure for generating streamflow and evaporation series is nearly identical, matrices $\boldsymbol{NI}$ and $\boldsymbol{E}$ will be referred to as $\boldsymbol{Q}$ hereafter.

The first step is to create matrices of years of historical (subscript h) streamflows and evaporation rates $\boldsymbol{Q}_{h_k} \in \mathbb{R}^{N_{hr} \cdot 52 \times N_r}$, one of each for each water source k. Next, these matrices go through the transformation $\boldsymbol{Y}_{h_k}^{y,w} = ln\left(\boldsymbol{Q}_{h_k}^{y,w}\right)$ to create matrices $\boldsymbol{Y}_{h_s}$ for all sites. These are next converted into standard-normally-distributed matrices $\boldsymbol{Z}_{h_k} \in \mathbb{R}^{N_{hr} \times 52}$, in which:

$$\boldsymbol{Z}_h^{y,w} = \frac{\boldsymbol{Y}_h^{y,w} - \hat{\mu}}{\hat{\sigma}} \tag{2.4}$$

Matrix $\boldsymbol{Z}_{h_k}^{y,w}$ is next used to create matrix $\boldsymbol{C}_{s_k} \in \mathbb{R}^{N_{ys} \cdot 52 \times N_r}$. Each column of

$\boldsymbol{C}_{s_k}$ is filled by sampling with replacement from the corresponding week of the year (column) across all historical years in $\boldsymbol{Z}_{h_k}$. The column-wide sampling is performed by creating 52 N_{ys}-long vectors of bootstrap samples ranging from 1 to N_{hr}, to preserve cross-site correlation across reservoirs and gauges — the samples correspond to the index of the historical year from which the specific streamflow/evaporation sample is drawn. The samples Matrix $\boldsymbol{C}_{s_k}$ is, therefore, a matrix of uncorrelated standard-normally-distributed historical streamflows or evaporation rates. Temporal autocorrelation is then imposed on site k by using Choleskly decomposition to create an upper-triangular matrix $\boldsymbol{U}_k$ such that $corr(\boldsymbol{Z}_{h_k}) = \boldsymbol{U}_k^T \boldsymbol{U}_k$, and multiplying it by $\boldsymbol{C}_{s_k}$, or $\boldsymbol{S}_{s_k} = \boldsymbol{C}_{s_k}\boldsymbol{U}_k$, resulting in matrices $\boldsymbol{S}_{s_k}$ streamflows and evaporation rates. As a way of improving inter-annual correlation, the described process is performed on matrix $\boldsymbol{Q}^{'}_{s_k}$, a version of matrix $\boldsymbol{Q}_{s_k}$ shifted by 27 weeks Kirsch et al. (2013), resulting in matrices $\boldsymbol{S}^{'}_{s_k}$.

Matrix $\boldsymbol{S}^{'}_{s_k}$ has auto and cross-site correlated values of standard, log-normal streamflows with the same sample means and variances are the same as the historical data. The last step is therefore to transform matrices $\boldsymbol{S}^{'}_{s_k}$ from log back to real space using the log-means and log-variances of the historical data. Additionally, the log-mean of the historical data was scaled in the transformation to simulate climate change. The scaling was done by multiplying the historical log-mean and log-variance by weekly sinusoid-time-series, the parameters of which deeply uncertain themselves, before performing the reverse log transformation Quinn et al. (2018). The mathematical formulation of the described transformation is presented below:

$$\boldsymbol{Q}^{y,w}_{s_k,r} = e^{S^{'y,w}_{s_k,r}\hat{\sigma} + \hat{\mu}\cdot[A_r \cdot sin(\phi_r + f_r w)]} \tag{2.5}$$

where A_r is the amplitude assigned to that exploratory SOW, ϕ_r the phase, and f_r the frequency of the sinusoids. Examples of sinusoids used to create the streamflows and evaporation ensembles and the corresponding flow duration curves are presented in Figure 2.5.

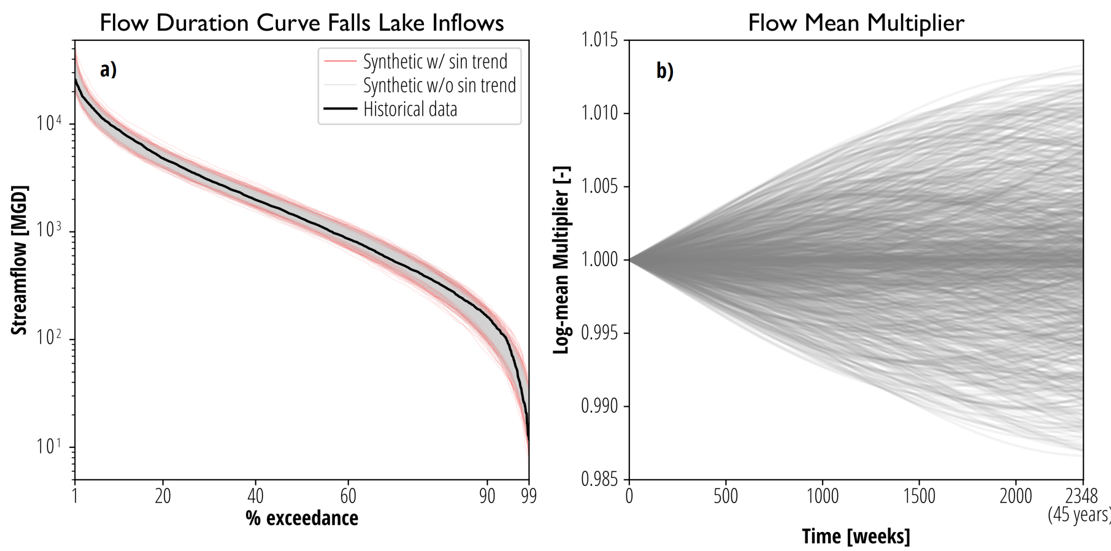

Figure 2.5: Panel (a) shows examples of flow duration curves for historical inflows, synthetically inflows with stationary log-mean, and synthetically inflows with non-stationary log-mean. Panel (b) shows all sinusoid-multiplier series of the log-mean of inflows and evaporation rates used to generate the panel (a).

The generation of evaporation rates time series happens concomitantly. Correlation between inflows and evaporation rates for a single reservoir is enforced by using the same years samples to build both inflow and evaporation matrices $\boldsymbol{C}_{s_k}$. The only difference in the synthetic time series generation processes for inflows and evaporation is that the historical evaporation rates used in this problem are normally distributed, which eliminates the need to work with log-transform the data.

The process for creating the re-evaluation series was nearly identical for optimization and re-evaluation series. The only differences are that that A_r, ϕ_r, and f_r are the same for all scenarios r for each of the 5,000 re-evaluation series sets. Therefore, all the variability across scenarios in each function evaluation comes from sampling process alone and not from the sinusoidal trends. Figure 2.5 shows the flow duration curves and sinusoid mean-multipliers for all exploratory scenarios.

2.3 Scenario Generation for Synthetic Demand Time-series

For generating stochastic demand time series correlated with the corresponding synthetic streamflows, the historical weekly deviations of demands and streamflows from the corresponding annual means are used to to build empirical joint probability distributions from which synthetic demands can be sampled based on future synthetic inflows. The first step is to whiten the historical inflows and demands. For this, the historical data is divided by the corresponding annual historic mean, as shown below for demand (the same logic applies to the log-streamflows):

$$D_{rmf}^{y,w} = \frac{D_h^{y,w}}{\bar{D}_h^y} \tag{2.6}$$

where $D_{mf}^{y,w}$ is the recorded demand as a fraction of the corresponding annual mean $\bar{D}^y$ for week w of historical year y, and D_r is the value of the corresponding recorded weekly demand. This process should eliminate changes in annual demands and inflows (mostly resulting from changes in population and local climate) so that all years of historical data are used to build empirical demand distributions for each week of the year, whose means and standard deviations are calculated as below (again, the same logic applied to log-streamflows):

$$\mu_{D_h}^{w} = \frac{\sum_{y=0}^{N_{hr}} D_{rmf}^{y,w}}{N_{hr}} \tag{2.7}$$

$$\sigma_{D_h}^{w} = \sqrt{\frac{\sum_{y=0}^{N_{hr}} \left(D_{rmf}^{y,w} - \mu_{D_h}^{w}\right)^2}{N_{hr} - 1}} \tag{2.8}$$

The historical demands and streamflows display a negative correlation during the months of April through September because of lawn irrigation and no correlation from October to March. The historical weekly streamflows are then split into sixteen bins for each week of the year based on their values. To generate a demand sample for a given future week during irrigation season, the generator selects the bin corresponding to the value of the whitened synthetic streamflow for that same week of the year, samples a whitened demand value from the corresponding historical whitened demands, and converts the value back to volume per time in real-space based on the corresponding projected annual mean. For non-irrigation season, sampling occurs across any of the historical whitened demands independently of the streamflows.

CHAPTER 3

REDUCING REGIONAL DROUGHT VULNERABILITIES AND MULTI-CITY ROBUSTNESS CONFLICTS USING MANY-OBJECTIVE OPTIMIZATION UNDER DEEP UNCERTAINTY

This chapter is drawn from the following peer-reviewed journal article: B. Trindade, P. Reed, J. Herman, H. Zeff, G. Characklis, Reducing regional drought vulnerabilities and multi-city robustness conflicts using many-objective optimization under deep uncertainty, Advances in Water Resources 104 (2017) 195-209. doi:10.1016/j.advwatres.2017.03.023. Funding for this work was provided by the National Institute of Food and Agriculture, U.S. Department of Agriculture (WSC Agreement No. 2014-67003-22076). Additional support was provided by the U.S. National Science Foundation's Water, Sustainability, and Climate Program (Award No. 1360442). The views expressed in this work represent those of the authors and do not necessarily reflect the views or policies of the NSF or the USDA.

3.1 Abstract

Emerging water scarcity concerns in many urban regions are associated with several deeply uncertain factors, including rapid population growth, limited coordination across adjacent municipalities and the increasing risks for sustained regional droughts. Managing these uncertainties will require that regional water utilities identify coordinated, scarcity-mitigating strategies that trigger the appropriate actions needed to avoid water shortages and financial instabilities. This chapter focuses on the Research Triangle area of North Carolina, seeking to engage the water utilities within Raleigh, Durham, Cary and Chapel Hill in cooperative and robust regional water portfolio planning. Prior analysis of this region through the year 2025 has identified significant regional vulnerabilities to volumetric shortfalls and financial losses. Moreover, efforts to maximize the individual robustness of any of the mentioned utilities also have the potential to strongly degrade the robustness of the others. This research advances a multi-stakeholder Many-Objective Robust Decision Making (MORDM) framework to better account for deeply uncertain factors when identifying cooperative drought management strategies. Our results show that appropriately designing adaptive risk-of-failure action triggers required stressing them with a comprehensive sample of deeply uncertain factors in the computational search phase of MORDM. Search under the new ensemble of states-of-the-world is shown to fundamentally change perceived performance tradeoffs and substantially improve the robustness of individual utilities as well as the overall region to water scarcity. Search under deep uncertainty enhanced the discovery of how cooperative water transfers, financial risk mitigation tools, and coordinated regional demand management must be employed jointly to improve regional robustness and decrease robustness conflicts between the utilities. Insights from this work have general merit for regions where adjacent municipalities can benefit from cooperative regional water portfolio planning.

3.2 Introduction

Understanding and exploring the combined states-of-the-world (SOWs) that shape water supply risks is a fundamental challenge in any drought mitigation application, defining perceptions and preferences related to key performance objectives, vulnerabilities, and/or the robustness of a system. Fundamental to this challenge is the choice of what uncertain factors should be sampled or included in generating alternative scenarios. Broadly, there are two dominant methodologies that are at present used in water supply planning contexts: (1) pre-specified deterministic scenarios/narratives and (2) globally sampled stochastic scenario analysis. As an example of deterministic scenario-based planning, Liu et al. (2008) analyzed four semi-arid river basins in the southwestern United States seeking to find water management strategies based on integrated modeling that would minimize lasting or irreversible environmental impacts (mostly on local vegetation) while allocating enough water for competing uses. The authors analyzed a pre-specified set of scenarios based on their *a priori* specified problem definition and input from stakeholders as well as water management experts. As an alternative example of stochastic scenarios, Brown et al. (2012) and Moody and Brown (2013) use ensembles of reservoir inflows estimated using regression techniques to estimate climate vulnerability. Climate vulnerability analysis focuses on reducing risks, or losses, calculated as the product of a loss function times subjective probabilities for the climate scenarios. The probabilities are elicited by contrasting a mixture of global climate model projections, paleodata (if available), and synthetically generated weather. These approaches are termed Decision Scaling, and have been recently expanded by Ghile et al. (2014) and Lownsbery (2014) to include hydroeconomic sensitivities.

The Decision Scaling example relates to a broader body of literature focused on planning under deep uncertainty. Deep uncertainties cannot, as described by Knight (1921, p. xiv), be treated as "a gamble on a known mathematical chance". In other words, they lack consensus on their underlying probability distributions as well as on individual or collective outcomes. Such uncertainties have also been termed Knightian, or scenario uncertainties, by multiple authors, including Walker et al. (2013); Lempert et al. (2006) and Kwakkel et al.

(2016b). Exploratory modeling (Bankes et al., 2001) — which involves running planning models across broad global samples of hypothesis — has emerged as a primary strategy for discovering which combinations of deeply uncertain factors are of high consequence in decision problems, as opposed to considering one most likely future projection (Walker et al., 2003; Lempert et al., 2006; Walker et al., 2013; Zeff et al., 2014; Herman et al., 2014). As noted by Dessai et al. (2009), the intent of exploratory modeling is to shift focus from predicting future conditions to understanding which conditions lead to decision relevant consequences. Given a better understanding of the consequences related to these discovered conditions, decision makers can then evaluate their relative importance or plausibility.

As reviewed by Herman et al. (2015), the selection of decision alternatives to analyze in exploratory modeling frameworks can either be done *a priori* or with computational search. In the case of decision alternatives that are specified *a priori*, the decision analysis then takes the form of a classical discrete choice among those alternatives based on their interpreted robustness. This approach has been commonly employed across the full range of decision support frameworks in this area (Brown et al., 2012; Moody and Brown, 2013; Hipel and Ben-Haim, 1999; Lempert et al., 2006; Bryant and Lempert, 2010). Conversely, a growing number of studies are using search strategies to discover high performing design alternatives in their initial stage of analysis and then evaluating their robustness to inform design selection (Korteling et al., 2013; Kasprzyk et al., 2013; Hamarat et al., 2014; Kwakkel and Haasnoot, 2015; Giuliani and Castelletti, 2016; Singh et al., 2015; Hadka et al., 2015).

This chapter advances the Many-Objective Robust Decision Making

(MORDM) framework initially introduced by Kasprzyk et al. (2013) by demonstrating the value of progressively transitioning from sampling well-characterized uncertainties only to also sampling deeply uncertain factors during the MORDM search phase. The inclusion of deeply uncertain factors in the search phase of the MORDM framework seeks to design portfolios of drought mitigation actions that perform well under a wide range of SOWs — where a SOW is defined here as a fully specified world, comprised of one fully specified sampled vector of deeply uncertain factors ($\mathbf{\Psi}$) and one streamflow time series for each reservoir. With this change, deep uncertainty sampling is performed for both generating candidate management alternatives and evaluating their robustness. This chapter builds on the hypothetical western water market test case explored by Watson and Kasprzyk (2016), which shows the value for increasing robustness by including deeply uncertain factors in the computational search and alternative generation phases of MORDM. Their study demonstrates that there were potential robustness benefits for a hypothetical water supply portfolio analysis for a single city in the Lower Rio Grande. In contrast, this chapter focuses on the Eastern U.S. using the Research Triangle region of North Carolina, engaging the water utilities serving Raleigh, Durham, Cary, Chapel Hill and Carborro in cooperative and robust regional water portfolio planning under deep uncertainty.

The Research Triangle region poses a multi-jurisdictional decision making context, where each municipality faces its own performance supply reliability and financial tradeoffs. Moreover, each of the utilities' water portfolio planning actions can have strong negative impacts on the neighboring systems. Efforts to maximize the individual robustness of any of the mentioned utilities have been shown to potentially significantly degrade the robustness of the others. Therefore, this chapter focuses on three core contributions: (1) demonstrate that adequately combining dynamic risk-triggered conservation and transfers with financial hedging requires the inclusion of deeply uncertain factors in the search phase of MORDM; (2) demonstrate that searching adaptive risk-based action triggers under a comprehensive sample of deeply uncertain factors may promote the discovery of a larger and more diverse suite of candidate solutions that meet stakeholder requirements in a broader variety of hypothetical future SOWs; and (3) illustrate how the number and diversity of robust water portfolios can aid in reducing planning conflicts among utilities (i.e., robustness tradeoffs across the municipalities).

3.3 Methodological Framework

3.3.1 Overview

As mentioned in Section 3.2, this chapter advances the MORDM framework initially introduced by Kasprzyk et al. (2013) by demonstrating the value of progressively transitioning from sampling well-characterized uncertainties only to also sampling deeply uncertain factors during search. The inclusion of deeply uncertain factors in the search phase of the MORDM framework seeks to identify alternatives that perform well under a wide range of SOWs. With this change, deep uncertainty sampling is performed for both generating and evaluating the robustness of candidate management alternatives. Following the taxonomy presented in Herman et al. (2015), the proposed MORDM methodology has the following four core component steps:

1. Sampling States of the World: Selection and definition of well-characterized and deeply uncertain factors that will be used to generate ensembles of SOWs.

2. Generation of alternatives: For one or more candidate problem formulations (i.e., objectives, decisions, constraints, and SOWs), exploiting many-objective evolutionary search under uncertainty to discover the Pareto approximate solutions.

3. Choosing Measures of Robustness: Evaluation of tradeoff alternatives' robustness across a broader ensemble of $\boldsymbol{\Psi}$ vectors using stakeholder elicited measures of acceptability (e.g., regret or satisficing-based robustness measures, as in Lempert et al. 2006 and Herman et al. 2014, respectively). Multiple candidate robustness measures may be elicited over iterations with stakeholders as part of a constructive learning feedback process (Tsoukias, 2008).

4. Discovering the Factors that Control Robustness: Use global sensitivity analysis and/or factor mapping to discover which uncertain factors control the robustness of Pareto approximate alternatives and warrant further potential monitoring or actions [e.g., scenario-discovery presented by Bryant and Lempert (2010)]. Sensitivity analysis methods can be used to rank order the most important factors (i.e., factor prioritization, as in Saltelli et al. 2008, 2006), or for factor mapping to find specific factors' value ranges that cause conditions of concern (i.e., categorical classification as introduced by Friedman and Fisher 1999; Breiman et al. 1984, respectively).

More detailed descriptions for each of these four component steps for the

MORDM framework are provided in Sections 3.3.2 to 3.3.5, below.

3.3.2 States of the World

Related to the present study, Zeff et al. (2014) accounted for a single source of well-characterized uncertainty (WCU) in their analysis of the Research Triangle's water supply tradeoffs through 2025, namely streamflows. Uncertain streamflows were included in the model using synthetically generated stochastic streamflow time series following the methodology presented by Kirsch et al. (2013). In Zeff et al. (2014), multiobjective search was used to identify candidate water portfolio designs. During the search phase of their study, all other factors such as regional demand trajectories, consumer response to water restrictions, prices for water transfers, etc., however, were treated deterministically. Transitioning the Research Triangle from planning solely based on deterministic historical streamflow records to synthetically generated streamflows showed that the region's planners were underestimating the severity and impacts of droughts. The initial use of the streamflow WCU served two benefits: (1) it highlighted that the region needs to more carefully consider transfers, financial risk mitigation, and conservation actions and (2) it illustrated consequential results that motivated the consideration of a broader range of uncertainties (e.g., demand growths, transfer pricing, imperfect responses to restrictions, demand hardening, etc.).

Building on the WCU optimization by Zeff et al. (2014), Herman et al. (2014) evaluated the robustness of the candidate water portfolios using a Latin Hypercube Sampling (LHS) of the 13 deeply uncertain parameters shown in Table

3.1, as a means of ensuring a good representation of possible futures. Figure 3.1 illustrates what we term the deep uncertainty re-evaluation (DU re-evaluation) used in the robustness calculations. For the Research Triangle test case, the DU re-evaluations pair the original synthetically generated stream flows, as used in Equation 3.17, with each LHS sampled set of deeply uncertain factors. For this chapter, the 1,000 original synthetic inflows used in Zeff et al. (2014) and Herman et al. (2014) are used in conjunction with 10,000 LHS sampled sets of deeply uncertain factors ($\mathbf{\Psi}$) to create 10 million SOWs. The 13 deeply uncertain factors in Table 3.1 can be grouped into four categories representing different aspects of the problem. The categories account for future hydrologic uncertainties, uncertainties in demand projections regarding timing and magnitude, legal and physical factors related to future storage capacities, and fluctuations in the water transfers' pricing. Ranges for deeply uncertain factors were chosen collaboratively with the Research Triangle utilities. As noted by Bryant and Lempert (2010), global sampling of deeply uncertain factors, with subsequent participatory assessments that facilitate the discovery of consequential combinations, allows for an improved *a posteriori* assessment by stakeholders of where to focus their attention and actions.

3.3.3 Alternatives Generation

In the context of the Research Triangle, Zeff et al. (2014) employed multiobjective evolutionary algorithms to identify water supply portfolios combining water transfers, financial risk instruments, and demand management. Each portfolio evaluation was based on multiple realizations of synthetically generated streamflows coupled with a set of deterministic base values for each of the 13

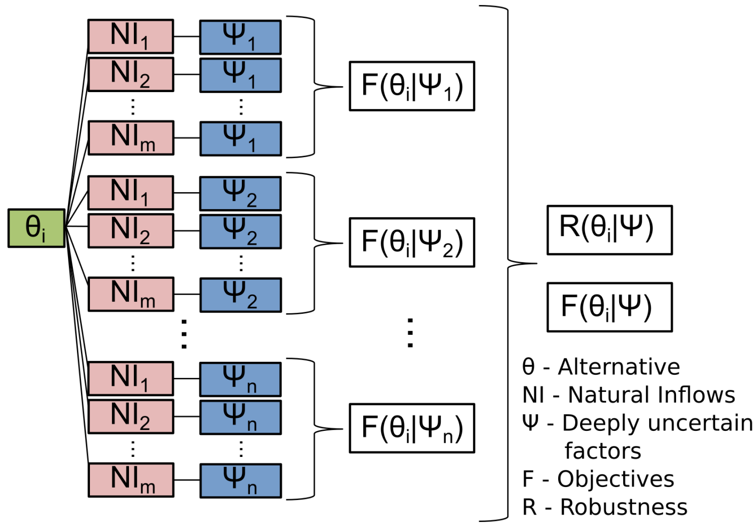

Figure 3.1: Schematics of the deep uncertainty re-evaluation of candidate water portfolio solutions (termed DU re-evaluation). A set of objectives is calculated for each 1,000 SOW's that share the same $\boldsymbol{\Psi}$. These objectives are in turn used to calculate overall robustness and overall objectives for alternative θ.

deeply uncertain factors presented in Table 3.1, representing what we term in this dissertation the WCU optimization was used initially to identify design alternatives (see Figure 3.2a). Transitioning the Research Triangle from planning solely based on the deterministic historical streamflow record to synthetic streamflows (Kirsch et al., 2013) as a well established uncertainty characterization tool showed that the region's planners were underestimating the severity and impacts of droughts.

The study transitions from the WCU optimization used in the prior published analysis of the Research Triangle to include deep uncertainties in the search and identification of candidate water portfolios. This transition poses a challenging multi-objective search problem that motivated our use of the MS-Borg algorithm, described in Section 2.1. The coupling of sampling of deeply uncertain factors with the computational search provided by the Borg MOEA

Table 3.1: Uncertainty Factors and Sampling Ranges for Many-Objective Robust Decision Making (vector Ψ in Equations 3.5 and 3.10).

Category	Name	Current Value	Lower Bound	Upper Bound
Climate	Inflows multiplier	1.0	0.8	1.2
	Evaporation multiplier	1.0	0.8	1.2
Demand	Consumer reductions multiplier	1.0	0.8	1.2
	Consumer reductions lag (weeks)	0	0	4
	Mean peaking factor	1.0	0.5	2.0
	Demand growth multiplier	1.0	0.5	2.0
	Standard deviation of demand variations	1.0	0.5	2.0
Capacity	Falls Lake municipal supply allocation	1.0	0.8	1.2
	Jordan Lake municipal supply allocation	1.0	0.8	1.2
	Cary treatment capacity multiplier	1.0	1.0	2.0
	Transfer connection capacity multiplier	1.0	1.0	2.0
Costs	Transfer cost ($/MG)	3000	2500	5000
	Insurance premium multiplier	1.2	1.1	1.5

provides direct feedbacks between the alternative generation and selection of SOWs steps of the MORDM framework. More specifically, it results in an alternative generation scheme that evaluates each candidate solutions against different SOWs, each with its own streamflow time series and vectors samples of 13 DU factors, for each function evaluation performed during the computational search. This alternative generation scheme, illustrated in Figure 2b, is termed here Deep Uncertainty (DU) optimization. In Figure 3.2b, the green boxes represent one alternative water portfolio, each red box represent a different set of

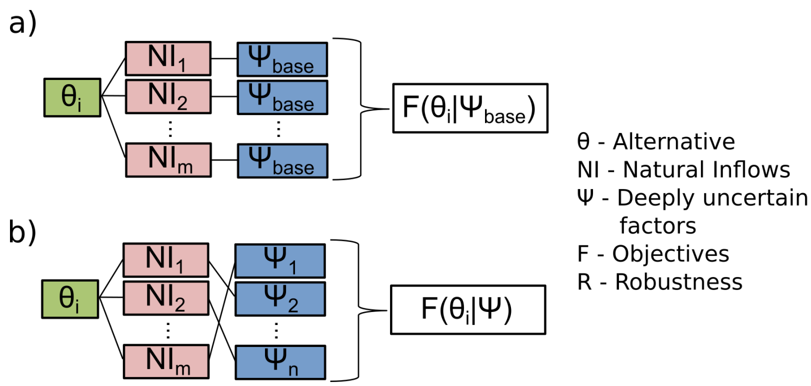

Figure 3.2: WCU (a) and DU (b) optimization factor sampling schemes. The former creates states of the world by coupling synthetically generated set of streamflows with a base projection for 13 deeply uncertain factors, while the latter couples each set of synthetically generated streamflows with a different projection for the 13 deeply uncertain factors, shuffling the coupling at every function evaluation.

synthetic reservoir inflows that are randomly paired with one of 1,000 samples of 13 DU factors. The new alternative generation scheme seeks to yield more robust alternatives, as Borg MOEA will search for alternatives whose performance should be acceptable under a broad range of SOWs. The DU optimization's sampling scheme shown in Figure 3.2b approximates the much broader and computationally demanding sampling used in the DU re-evaluation illustrated in Figure 3.1. It would be intractable to include all 10 million DU re-evaluation SOWs to evaluate each candidate water portfolio during search.

3.3.4 Robustness Measures

As reviewed by Herman et al. (2015), the three most common types of robustness measures are: expected value, regret and satisficing; all three can be employed in either univariate or multivariate contexts. The three measures are described in detail in Lempert and Collins (2007), the latter two being further developed and tested in the context of water supply portfolio analysis in Herman et al. (2015). Herman et al. (2014) have used the domain criterion robustness measure (i.e., percent of states of the world in which given performance criteria are met) on the Research Triangle problem, which was used in this work again because it best aligned with utilities preferences (Herman et al., 2015). As was discussed above in Section 3.3.2, the implementation of the domain criterion robustness measure for the Research Triangle problem is illustrated in the DU re-evaluation approach in Figure 3.1. This choice of robustness calculation was the motivation behind the computational search sampling scheme presented in this work (DU optimization, Figure 3.2b), which was intended to approximate the computationally intensive DU re-evaluation scheme (Figure 3.1) during the Borg MOEA search phase described in Section 3.3.3 for the sake of computational tractability.

3.3.5 Robustness Controls

Scenario discovery is a key step in many RDM frameworks and represents a factor-mapping sensitivity analysis used to determine under which ranges of the deep uncertain factors each alternative may be vulnerable to failures (Bryant and Lempert, 2010; Herman et al., 2014, 2015; Kasprzyk et al., 2013; Lempert et al., 2006; Halim et al., 2015; Kwakkel and Jaxa-Rozen, 2016; Groves and Lempert, 2007). Based on these ranges and on their experience with their system, decision makers can better understand where candidate actions perform acceptably and inform monitoring efforts to detect conditions of concern. PRIM and CART trees are suggested by Lempert et al. (2008), instead of more sophisticated statistical learning algorithms, for assessing scenario discovery due to their high interpretability.

PRIM works by gradually shrinking (peeling) prespecified ranges of factors according to a set confidence interval in order to find multi-dimensional boxes that encompass as high of a proportion of true positives as possible (Friedman and Fisher, 1999). In this application, positives are $\mathbf{\Psi}$ vectors for which an alternative fails to meet a given performance criteria. The algorithm returns a series of candidate boxes, each one devised after one peel and described by the following standard box metrics:

1. Density: the ratio of the number cases of interest (fails) inside that box to the total number of cases inside that box. A value of 1 is ideal to this metric.
2. Coverage: the ratio of the number of cases of interest inside that box to the total number of cases of interest in the entire domain. A value of 1 is ideal to this metric.
3. Equivalent Type II error: the ratio of the number of cases of interest outside the box to the total number of cases outside the box (ratio of false negatives). A value of 0 is ideal to this metric.

The analyst should choose a box according to stakeholders' interests. For example, relative to the water utilities in the Research Triangle, low type II error and high coverage are more important than high density, given the high risk aversion to any failures of the system. The type II error represents the relative perceived risk water utilities are willing to take of being wrong in their expected system performance. There is not a known universal density versus coverage ratio that can be used to all problems, as they are problem and stakeholder specific (Lempert et al., 2008). The described process of scenario discovery (robustness controls) is described in more detail in Bryant and Lempert (2010).

3.4 The Research Triangle Test Case

The Research Triangle region, as illustrated in Figure 3.3, is comprised of several municipalities, whose water supplies are provided by public utilities in Raleigh, Durham, Cary, and Chapel Hill/Carrboro (i.e., the Orange Water and Sewer Authority, or OWASA). Cary and Raleigh receive water from allocations of Jordan and Falls Lakes, respectively, two large flood control reservoirs operated by the United States Army Corps of Engineers. Volumetric allocations from each lake are designated by the state for municipal water supply. Raleigh has been granted 100% of the municipal water supply allocation to Falls Lake, and Cary holds 39% of the municipal water supply allocation to Jordan Lake. Although Cary is the only utility to operate an intake on Jordan Lake, other utilities also have allocations that can be accessed by purchasing water from Cary's water treatment plant during conditions of scarcity, subject to Cary's availability. Rapidly growing demands in the region have accelerated competition for the remaining unallocated portion of Jordan Lake's water supply pool. A number of municipalities, including Raleigh, the largest in the region, have requested new allocations, and it is likely that the pool will be completely allocated after the most recent round of allocation requests is approved by the state, even without Raleigh's request (Triangle J, 2014).

Table 3.4 summarizes the region's reservoir storage dedicated to municipal

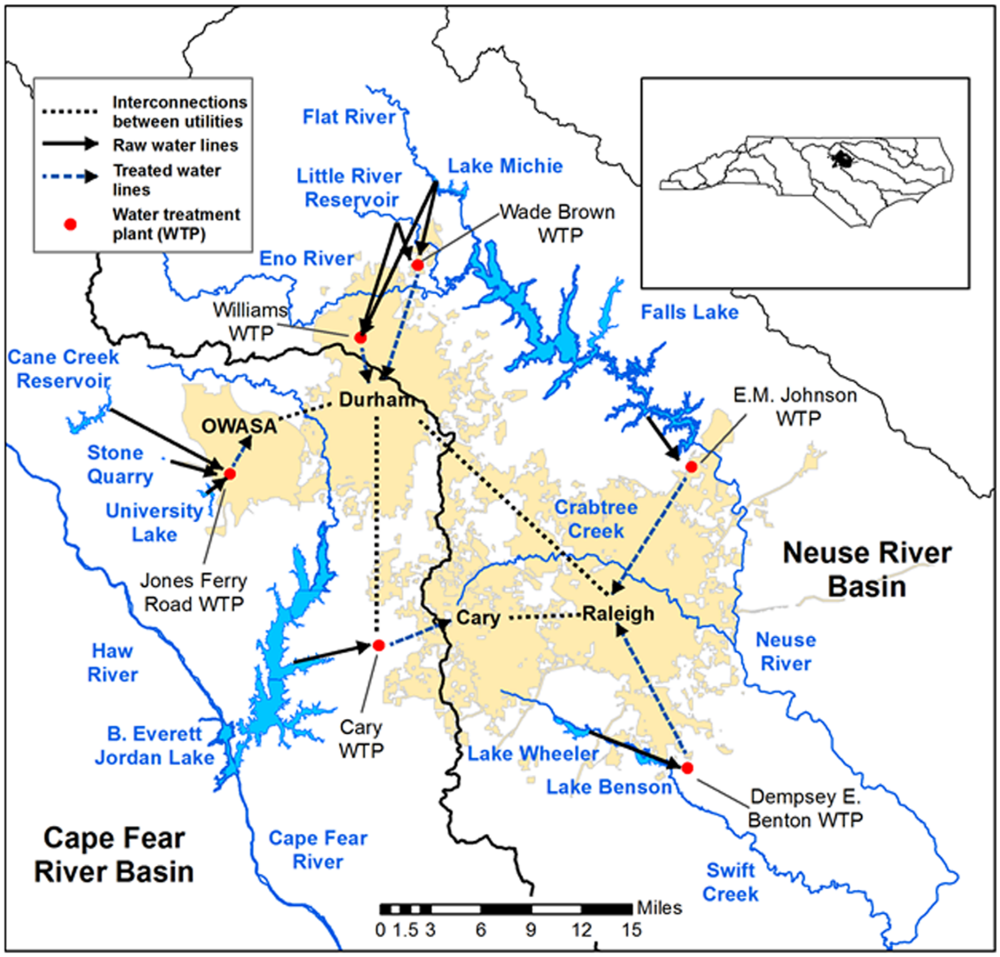

Figure 3.3: Map of the North Carolina Research Triangle region. Adapted from Zeff et al. (2014).

water supply as well as allocations for all four utilities. The Research Triangle is projected to experience a total demand growth of 33% through the planning horizon year of 2025 considered in this chapter (average of annual growth of 2.4%) (NCDENR, 2002). Given that the regional municipalities cannot bring new sources online before 2025, Jordan Lake represents the only major supply source that still has unallocated water for municipal use. Consequently, it is a current source of regional resource conflicts.

Although the municipalities are in close geographic proximity and share the

Table 3.2: Summary of reservoir storage capacities, demands, and allocations for the Research Triangle Region

Utility	Reservoir	Reservoir capacity reserved for municipal supply (Billion gallons)	Percentage of reservoir capacity reserved for municipal supply allocated to each utility currently in use	2013 Demand (Million gallons /day)	2025 Demand (Estimated) (Million gallons /day)	Demand/ Capacity Ratio (2025) ($10^{-3} year^{-1}$)
OWASA	Cane Creek	3.0	100%	8.0	9.0	8.2
	Stone Quarry		100%			
	University Lake		100%			
	Jordan Lake		5%			
Durham	Little River	6.4	100%	27.5	34.9	14.9
	Lake Michie		100%			
	Jordan Lake		10%			
Raleigh	Falls Lake	14.7	100%	57.4	76.2	14.2
Cary	Jordan Lake	14.9	39%	23.1	34.0	6.2

same hydroclimatic conditions, their different demand-to-capacity ratios yield highly variable drought impacts across the region's utilities. In severe droughts, transferring water from Jordan Lake via the existing Cary water treatment plant offers a means of more efficiently using the region's existing supply capacity to overcome water scarcity challenges (Caldwell and Characklis, 2014; Zeff et al., 2014). Transfers from existing sources can help to delay or avoid the costs of building new supply infrastructure, but they do create financial risks for utilities. Utilities typically set rates with the aim of returning revenues that cover their costs, which are dominantly fixed annual debt payments. The additional costs of water transfers during periods of drought can destabilize this budgetary balance. If they are large enough, these swings could harm the ability of a utility to repay their debt on schedule and increase future borrowing costs (Hughes and Leurig, 2013). The costs and financial risks associated with these transfer agreements for several types of financial instruments were studied in detail in Palmer and Characklis (2009), Zeff and Characklis (2013), Caldwell and Characklis (2014) and Zeff et al. (2014).

3.4.1 Triangle Region Management Model

Zeff et al. (2014) is the first study to present a fully detailed water supply simulation for the Research Triangle region's water utilities using drought mitigation instruments based on short term conservation, water transfers, and financial hedging. This work represents the culmination of several prior efforts (risk-of-failure triggers in Palmer and Characklis 2009, index insurance in Zeff and Characklis 2013, and water transfers in Caldwell and Characklis 2014). Zeff et al. (2014) used the model to find tradeoffs between reliability focused choices and their inherent financial risks when formulating drought management policies for the utilities in the Research Triangle. The WCU optimization formulation in that study considered hydrologic variability as the sole source of uncertainty. The policies consisted of different combinations of financial instruments and risk-of-failure (ROF, described in equations 3.8 below) triggered drought mitigation instruments. Herman et al. (2014) then used the same model to assess the robustness of the solutions found by Zeff et al. (2014) by testing them against a broader ensemble of uncertainties, some of them deep, finding that multiple utilities had significantly low robustness values. This chapter builds on these prior studies, seeking to discover if including deep uncertainties in the search phase of MORDM will result in improving the individual utilities as well as the whole Triangle region's robustness to drought. Given the growing complexity of the Research Triangle problem formulations as well as the evolving treatments of uncertainty, this section provides a careful mathematical description of the model linking the prior studies to the proposed DU optimization explored here.

Minimization Problem

Here, we define the optimization problem as finding the optimal policy θ^* that minimizes the objective function vector **F**, i.e.

$$\theta^* = \underset{\theta}{argmin}\ \mathbf{F} \tag{3.1}$$

where,

$$\mathbf{F} = \begin{bmatrix} -f_{REL}(\mathbf{x_s}, \theta_{\mathbf{rt}}, \theta_{\mathbf{tt}}, \theta_{\mathbf{jla}}) \\ f_{RF}(\mathbf{x_{rof}}, \theta_{\mathbf{rt}}, \theta_{\mathbf{tt}}, \theta_{\mathbf{jla}}) \\ f_{JLA}(\theta_{\mathbf{jla}}) \\ f_{AC}(\mathbf{x_{rof}}, \theta_{\mathbf{rt}}, \theta_{\mathbf{tt}}, \theta_{\mathbf{jla}}, \theta_{\mathbf{acfc}}, \theta_{\mathbf{irt}}) \\ f_{WCC}(\mathbf{x_{rof}}, \theta_{\mathbf{rt}}, \theta_{\mathbf{tt}}, \theta_{\mathbf{jla}}, \theta_{\mathbf{acfc}}, \theta_{\mathbf{irt}}) \end{bmatrix} \tag{3.2}$$

$$\mathbf{X} = \begin{bmatrix} \mathbf{x_{rof}} \\ \mathbf{x_s} \end{bmatrix} \tag{3.3}$$

$$\theta = [\theta_{\mathbf{rt}}, \theta_{\mathbf{tt}}, \theta_{\mathbf{jla}}, \theta_{\mathbf{acfc}}, \theta_{\mathbf{irt}}] \tag{3.4}$$

where $\mathbf{F}$ is the vector-based objective function, f_{REL} is the reliability objective, f_{RF} is the restriction frequency objective, f_{JLA} is the Jordan Lake allocation objective, f_{AC} is the average cost objective, f_{WCC} is the worse case cost objective, $\mathbf{X}$ is the state vector across the utilities and time, $\mathbf{x_{rof}}$ is a vector of approximately calculated risk of failures (Equation 3.8), $\mathbf{x_s}$ is a vector of combined utility storage, $\theta_{\mathbf{rt}}$ is a vector of restriction triggers, $\theta_{\mathbf{tt}}$ is a vector of transfer triggers, $\theta_{\mathbf{jla}}$ is a vector of Jordan Lake Allocations, $\theta_{\mathbf{acfc}}$ is a vector of annual contingency fund contributions, and $\theta_{\mathbf{irt}}$ is a vector of insurance restriction triggers. The storage state $\mathbf{x_s}$ can be defined as:

$$\mathbf{x_s}^w = f\left(\mathbf{x_s}^{w-1}, \mathbf{C}, \mathbf{D}^w, \mathbf{TF}^w, \mathbf{NI}^w, \mathbf{E}^w, \mathbf{W}^w | \mathbf{\Psi}\right) \tag{3.5}$$

$$\mathbf{D}^w = f(\mathbf{x_{rof}}^w, \theta_{\mathbf{rt}}) \tag{3.6}$$

$$\mathbf{TF}^{w} = f(\mathbf{x_{rof}}^{w}, \theta_{\mathbf{tt}}) \tag{3.7}$$

where $\mathbf{C}$ is a vector of reservoir capacities (7 reservoirs, some utilities have more than 1), $\mathbf{D}$ are the demands for each utility, $\mathbf{TF}$ are the transfer volumes for each utility, $\mathbf{NI}$ are the natural inflows in each reservoir, $\mathbf{E}$ the evapotranspiration in each reservoir, $\mathbf{W}$ the spillage of each reservoir, and $\boldsymbol{\Psi}$ is a vector of deeply uncertain factors such as population growth, future allocations yet to be negotiated, etc.

The approximate risk-of-failure (ROF) state x_{rof} is used in conjunction with decision variables θ_{rt} and θ_{tt} to trigger drought mitigation instruments. The primary drought mitigation instrument currently used by utilities is water use restrictions, which is triggered independently by each utility according to rules based on either days of supply remaining, minimum storage levels, or simplified risk calculations (Department of Water Management, 2009; The City of Raleigh Public Utilities Department, 2011; Goodwin, 2009; Orange Water and Sewer Authority, 2010). Despite being important system state variables, days of supply remaining and storage levels do not comprehensively represent the water supply systems' dynamics. Therefore, action triggers should be based on metrics that incorporate more representative dynamics in available capacity and evolving demands (Palmer and Characklis, 2009). The proposed approach based on ROF calculations allows instead for the inclusion of information from a wider range of state variables while making stakeholder decisions focus on risk-based action triggers. The scenarios for the ROF calculations are described by past recorded inflows and by assumptions about the deeply uncertain factors in $\boldsymbol{\Psi}$ presented in Table 3.1. The inflows are assumed to be statistically stationary, given this is a short-term planning study through 2025.

3.5 The Mathematical Formulation of the ROF Metric

The drought mitigation measures are triggered in a given week if the calculated value of the ROF state x_{rof} reaches the trigger values θ_{rt} and θ_{tt}. Equations 3.8 to 3.10 mathematically define how the ROF metric is computed:

$$x^{w}_{srof,j} = \frac{1}{N_{rof}} \sum_{y'=0}^{N_{rof}} f^{w}_{y',j}(\boldsymbol{NI}^{y'}, \boldsymbol{E}^{y'}) \tag{3.8}$$

where,

$$f^{w}_{y',j} = \begin{cases} 0 & \forall w' \in \{(y', w), ..., (y', w + T_{rof})\} \; : \; \frac{x^{y,w'}_{s',j}}{C_j} \geq s_c \\ 1 & \text{otherwise} \end{cases} \tag{3.9}$$

and,

$$x^{y',w'}_{s',j} = f\left(C_j, \boldsymbol{UD}^{w}_{j}, \boldsymbol{NI}^{y',w'}_{j}, \boldsymbol{E}^{y',w'}_{j}, \boldsymbol{W}^{y',w'}_{j} | \boldsymbol{\Psi}_s\right) \tag{3.10}$$

In Equations 3.8 to 3.10, w' and y' mean a week and a year simulated with past data for the calculation of the ROF. Variable $x^{w}_{rof,j}$ is the ROF for utility j in current week w, and $f_{y',j}$ is a binary variable for which 0 denotes a failure happened during the corresponding ROF simulation with data from past year y'. In the ROF metric, $\boldsymbol{NI}^{y'}$, $\boldsymbol{E}^{y'}$ and $\boldsymbol{W}^{y'}$ are the recorded natural reservoir inflows, evaporation rates and reservoir spillage, respectively, in year y' prior to current week w used in one of the N_{rof} simulations — note that in a simulation in which the ROF metric is calculated for a week 20 years from now, 20 of the N_{rof} years of hydrological data will belong to synthetically generated data. Variable T_{rof}

equals 52 weeks for this chapter, so that single-year droughts are captured. A failure is defined as the combined storage $x_{s',j}^{y',w'}$ for any realization y' divided by combined storage capacity C_j for utility j falling below critical storage s_c. Variable $x_{s',j}^{y',w'}$ is the vector of storage states calculated in one of the year-long ROF simulations using recorded hydrologic data from past year y'. In the calculations of $x_{s',j}^{y,w'}$, $\boldsymbol{UD}^w$ is the unrestricted demand (no restrictions enacted or transfers purchased) in week w, and $\boldsymbol{\Psi}_s$ is a matrix of sampled deeply uncertain factors obtained through WCU or DU optimization. Given a year has 52.178 weeks, every 6 years (when the rounded number of weeks in a year would reach 53) the first week of the following historical year will be used.

Financial Model

Financial losses are accumulated when a utility purchases transfers or enacts conservation, and are equal to the total transfer costs plus the revenue losses resulting from reduced water sales. Utilities can attempt to offset these losses through payouts from third party contracts, and/or annual contributions made to contingency funds (self-insurance). Both were jointly explored in this dissertation.

By purchasing third-party financial insurances, a utility pays an upfront fee — a premium — in exchange for a payout in the event of a drought. For contractual purposes, the drought conditions that trigger insurance payouts are defined as inflow values that would trigger a pre-determined demand restrictions schedule. To define these inflows, system simulations are run with projected demand values for the following year. The insurance triggers are set as the maximum cumulative values, through each week from the beginning of the simulation, where storage drops sufficiently (and, consequently, for the ROF to increase enough) for water use restrictions corresponding to θ_{irt} to be triggered in the simulations. Inflows rather than the enactment of restrictions are used as insurance criterion in order to avoid moral hazards, given that utilities would be able to artificially enact restrictions on times of financial need in order to get a payout. More details about the insurance model can be found in Zeff and Characklis (2013). The insurance payout (IPO) is set for each utility in the beginning of the simulation as the expected value of revenue losses if restriction stage θ_{irt} is enacted given the historical inflows:

$$\mathbf{IPO} = E\left[\mathbf{RL}^w\right] = E\left[\mathbf{UD}^w \cdot \mathbf{UWP} - \mathbf{D}^w \cdot \mathbf{WP}^w\right] \tag{3.11}$$

$$\mathbf{WP}_w = f(\mathbf{x_{rof}}^w, \theta_{\mathbf{rt}}) \tag{3.12}$$

where RL is the vector of revenue losses, UWP is the vector of unrestricted water prices (independent of time), and WP is the vector of water prices subject to restrictions.

The insurance premium (IP) is calculated based on the expected total payouts made during each historical year ($YIPO$) plus a loading (IPR) based on the variance of such total payouts to account for volatility.

$$\mathbf{IP} = E\left[\mathbf{YIPO}\right] + \mathbf{IPR} \tag{3.13}$$

$$\mathbf{IPR} = f\left(\sigma^2_{\mathbf{YIPO}}\right) \tag{3.14}$$

The total payouts $YIPO$ that inform the insurance price can be expressed as:

$$\mathbf{YIPO}^y = f(\theta_{\mathbf{irt}}, \mathbf{IPO}, \mathbf{NI}^y) \tag{3.15}$$

where y denotes a year.

The annual contingency fund contribution decision variable is used to set aside a percentage of the annual revenue into a contingency fund, which rolls over to future years indefinitely. The amount of money in a contingency fund is updated yearly according to the following expression:

$$\mathbf{CF}^{y+1} = \mathbf{CF}^y \cdot (1 + r) + \theta_{\mathbf{acfc}}{}^y \cdot \mathbf{ATR}^y \tag{3.16}$$

where CF is the total amount of money in the contingency fund, ATR is the annual total revenue, and r is the fund's interest rate. Contingency funds work as a utility's savings account for protection against financial instabilities that can arise from multiple causes. Utilities are expected to contribute annually to their own contingency funds with a percentage of the annual revenue. Although a contingency fund can be set large enough to attempt to provide financial protection against any drought or sequence of droughts, large funds would require high annual contributions resulting in opportunity costs relative to more efficient financial management strategies. Moreover, large contingency funds could be appropriated by municipal governments for other purposes. The annual contingency fund contribution is the percentage of annual revenue a utility should deposit in its contingency fund.

Objectives

1. *Reliability*: The reliability objective is an approximation of the system reliability calculated as the fraction of considered states of the world which may cause the storage level of any reservoir used by a given utility to drop below 20% of its maximum capacity in any given week (failure condition):
$$\text{maximize } f_{REL} = \underset{j}{min}\left[\underset{y}{min}\left(\frac{1}{N_r}\sum_{i=1}^{N_r} g_{i,j}^{y}\right)\right] \tag{3.17}$$
where,
$$g_{i,j}^{y} = \begin{cases} 0 & \forall w \ : \ \frac{x_{s,i,j}^{w,y}}{C_j} \geq S_c \\ 1 & \text{otherwise} \end{cases}$$
where $g_{i,j,y} = 0$ if there was a week in a given year of a particular realization where reservoir storage falls below S_c of capacity (20% in this chapter), and 1 otherwise, N_r is the number of realizations in one function evaluation, y is the simulation year, N_{ys} is the number of years in the

project horizon, i is the simulation realization index, and j is the utility index. Although the term reliability generally denotes a probability, this is not the case for f_{REL}, as the inclusion of the deep uncertainty factors in the calculation makes the SOWs not equally likely and not guaranteed to exhaust all possible future scenarios. However, f_{REL} can be used as a proxy for reliability for decision making purposes as a way of assessing how a utility can rely on each alternative relative to each other, given the uncertainty space defined by the utilities. Therefore, f_{REL} will be further referred to in this work as reliability for the sake of readability.

2. *Restriction Frequency*: Restriction frequency represents the fraction of years over the course of the planning horizon for which at least one week of restricted water use is expected:

$$\text{minimize } f_{RF} = \underset{j}{max} \left[\frac{1}{N_{ys} \cdot N_r} \sum_{i=1}^{N_r} \sum_{y=1}^{N_{ys}} h_{i,j}^{y} \right] \tag{3.18}$$

where,

$$h_{i,j}^{y} = \begin{cases} 0 & \forall w \ : \ x_{rof,i,j}^{y,w} \leq \theta_{rt,j} \\ 1 & \text{otherwise} \end{cases}$$

where $h_{i,j,y} = 0$ if there was a week in a given year of a particular realization where water use restrictions were enacted, and 1 otherwise.

3. *Drought Management Cost*: The average cost objective represents the expected yearly cost of all non-infrastructure water portfolio assets used to manage droughts over the planning horizon. These costs are revenue losses from restrictions, transfer costs, contingency fund contributions,

and third party insurance contract costs:

$$\text{minimize } f_{AC} = \max_{j}\left[\frac{1}{N_{ys}\cdot N_r}\sum_{i=1}^{N_r}\sum_{y=1}^{N_{ys}} SYC_{i,j}^{y}\right] \tag{3.19}$$

where,

$$SYC_{i,j}^{y} = \frac{\theta_{acfc,j}\cdot ATR_{i,j}^{y} + IP_{i,j}^{y} + max(RL_{i,j}^{y} + TC_{i,j}^{y} - YIPO_{i,j}^{y} - CF_{i,j}^{y}, 0)}{ATR_{i,j}^{y}}$$

where CFC is the one time contribution to a contingency fund in a given year y, IP is the insurance contract cost, RL is the revenue losses from water use restrictions, TC is the transfer costs, $YIPO$ is the total insurance payout over year y, CF is the available contingency funds, and ATR is the total annual volumetric revenue. All these variables are dollar values.

4. *Exposure to Financial Risk (worse case cost)*: The worse case cost objective represents the 1% highest single-year drought management costs observed across all analyzed SOWs over the planning horizon:

$$\text{minimize } f_{WCC} = \max_{j}\left\{\underset{i\in N_r}{quantile}(SYC_{i,j}, 0.99)\right\} \tag{3.20}$$

5. *Jordan Lake allocation*: The Jordan Lake allocation objective represents the allocation of the Jordan Lake water municipal supply pool to be granted to a utility, such that:

$$\text{minimize } f_{JLA} = \sum_{j=1}^{N_u} JLA_j \tag{3.21}$$

where JLA is the percent of the pool allocated to utility j, and N_u is the total number of utilities in the system.

The regional water supply portfolio problem formulated in equations 3.1 - 3.21 represents a challenging multi-jurisdictional planning problem, given its number of stakeholders and objectives, and its stochastic and non-linear nature. The transition to searching for appropriate decision triggers that effectively use a balanced portfolio of conservation, transfers, and financial risk measures has broad value beyond the Research Triangle region. Section 3.3.1 describes our extensions of the multi-stakeholder MORDM framework in more detail.

3.6 Computational Experiment

3.6.1 Multiobjective Optimization Algorithm

As mentioned in Section 2.1, the Borg MOEA (Hadka et al., 2013) has been shown to provide competitive or superior performance relative to other top performing MOEAs across a broad suite of challenging problems (Reed et al., 2013; Hadka and Reed, 2012a). The Borg MOEA combines ϵ-dominance based archiving (Laumanns et al., 2002) with an adaptive operator selection (Vrugt and Robinson, 2007) based on probabilities calculated from solutions in its archive (Hadka et al., 2013), making it suitable across a wide range of problems. The standard values of parameters of the Borg MOEA v1.4 Master-Slave were used in this work (see Hadka et al. 2013 for the specific values).

3.6.2 Hydrologic Scenarios and DU Optimization Runs

The number of Monte Carlo realizations used to estimate the objective functions' evaluations was 1,000 — empirical assessments with ensemble sizes varying from 100 to 5,000 showed that 1,000 evaluations per run is sufficient for approximating the mean and variances of the Monte Carlo distributions used for computing candidate solutions' objectives. This approach for determining sampling rates originates from early studies of metaheuristic search dynamics given a single noisy objective function (e.g., Miller and Goldberg 1996; Smalley et al. 2000), where it was shown that relatively small Monte Carlo samples per function evaluation can provide good approximations when verified with much larger verification samples post search. These early results were translated to multi-objective evolutionary algorithms in subsequent studies (see Goh and Tan 2007; Singh and Minsker 2008; Kasprzyk et al. 2012; and Reed et al. 2013), which empirically test a step-wise increase in Monte Carlo sampling rates per function evaluation with subsequent verification tests with much larger samples. The deeply uncertain factors were sampled using the Latin Hypercube Sampling technique, based on the ranges presented in Table 3.1, while the synthetic streamflows, evaporation rates, and demands were generated using the methods described in Section (2.2) to create stationary time series by not modifying the log-means with the sinusoid described in Equation 2.5.

Running 1,000 realizations per function evaluation, however, significantly increases computational burden, which in turn set a requirement for higher computing capacity. The optimization was performed on the Texas Advanced Computing Center's Stampede, where 25 random seed Borg MOEA trial runs each performed a total of 1,000,000 function evaluations while exploiting 2,000 cores (AMD Opteron Quad-Core processors @ 2.3 GHz). The Pareto approximate fronts of all seeds were then combined to form a best known approximation of the global Pareto front (i.e., the reference Pareto set). Table 3.3 shows the values of ϵ (i.e., significant precision) used for the optimization and for reference set generation, which are the same as in Zeff et al. (2014).

Table 3.3: Values used for ϵ-dominance.

Objective	Reliability	Restriction Frequency	Jordan Lake Allocation	Average Cost	Worse Case Cost
ϵ	0.005	0.02	0.002	0.05	0.01

3.7 Results and Discussion

3.7.1 Optimization under Deep Uncertainties

A key motivation of this dissertation is to use the search phase of the MORDM framework to identify more robust water portfolios for the Research Triangle region. The DU optimization scheme illustrated in Figure 3.2b was formulated to approximately sample the more complex DU re-evaluation space shown in Figure 3.1. As discussed in Section 3.3.3, it would not be computationally tractable to use the DU re-evaluation scheme within the optimization itself. Figure 3.4 shows a comparative verification of the tradeoffs attained by the WCU and DU optimization alternatives by showing their average performance across the 10 million SOWs that define the DU re-evaluation space.

In Figure 3.4, the points where each line (water portfolio alternative) intersects the vertical axes represents the average performance in each objective in the DU re-evaluation space. The plot is oriented such that the ideal alternative would be a horizontal line at the bottom of all of the axes. When comparing alternatives, a downward shift along the axes represents the direction of higher preference and diagonal lines capture pairwise objective tradeoffs. The ranges of the averaged objective values as well as their relative tradeoffs highlight the

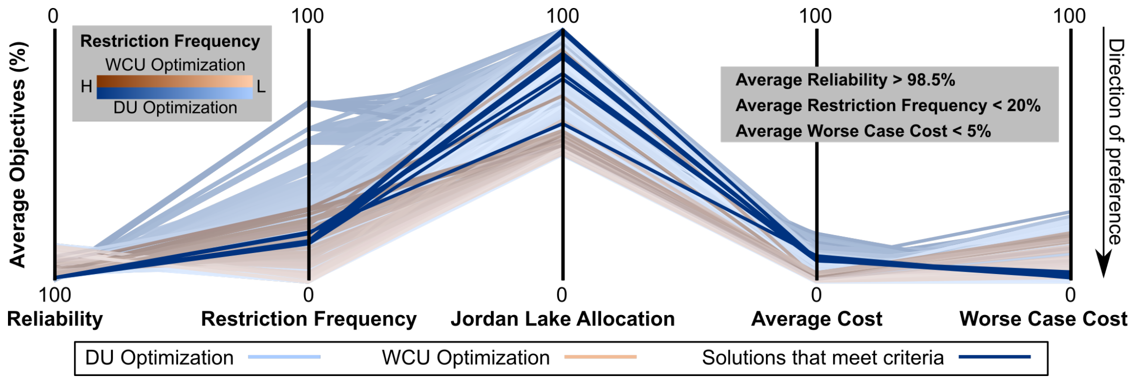

Figure 3.4: Average performance objective values for WCU and DU optimization alternatives over the 10 million SOWs from DU re-evaluation. The vertical axes denote average performance in each objective and each polyline represents an alternative. Brown tones represent alternatives from WCU optimization and blue tones represent alternatives from DU optimization. Transparency is applied to alternatives that do not meet a relaxed version of the performance criteria established by the studied utilities: Reliability > 98.5% (relaxed from 99%), Restriction Frequency < 20%, and Worst-Case Cost < 5%). The color gradient is based on the restriction frequency objective (light for low frequency, dark otherwise, not applicable to highlighted solutions). None of the brown (WCU) alternatives meet the criteria, while 7 out of 2,954 DU solutions do.

adaptivity of ROF-based transfers and restrictions in combination with the financial risk instruments in still attaining high levels of average performance despite the increased difficulty of the DU optimization's formulation. The most striking feature of Figure 3.4 is the increased use of Jordan Lake allocations, which indicates that the Research Triangle region's water utilities could face strong contention for this resource.

Figure 3.4 also illustrates which of the high performing DU optimization alternatives have high levels of performance on average when evaluated against the 10 million SOWs within the DU re-evaluation. In Figure 3.4, all alternatives whose values failed on average to meet the performance criteria established by the utilities — to attain at least 98.5% average reliability, average restriction frequencies no greater than 20% and average worst case costs no greater than 5% — were filtered out (transparent polylines). The reliability criteria set by the Research Triangle utilities was relaxed by 0.5% from the original 99% because none of the solutions from either the DU or the WCU optimization runs met this criterion without violating the restriction frequency and worse case cost criteria when evaluated against the broader ensemble of 10 million SOWs. Confirming the robustness assessment results in Herman et al. (2014), the alternatives from WCU optimization significantly deteriorated in their worse case costs, reliability and restriction frequencies; therefore all of the re-evaluated brown WCU solutions where filtered in Figure 3.4 (i.e., shown with a high level of transparency) — the re-evaluated averages for the reliability and worse case cost criteria would have to be relaxed to 95% and 25% respectively for any solutions from WCU optimization to avoid being brushed from Figure 3.4. The simplified but more comprehensive sampling of uncertainties used in the DU optimization proved to be helpful for generalizing to the DU re-evaluation space.

3.7.2 Robustness Comparison

Transitioning from the performance tradeoffs to satisficing frequencies, Figure 3.5 illustrates the significant differences in the robustness of the supply portfolios obtained with the WCU (brown) and DU (blue) optimization runs based on their calculated objectives values for each of the 10,000 samples of $\mathbf{\Psi}$ (deeply uncertain factors in Table 3.1). In this figure, the alternatives from both optimization schemes are rank-ordered (horizontal-axis) based on their robustness (vertical-axis). Figure 3.5 defined robustness as the percentage of the 10,000 $\mathbf{\Psi}$ samples where an alternative was able to meet all of the utilities' requirements as calculated with equations 3.17-3.21 for all 1,000 synthetically generated sets of natural inflows. The results show that DU optimization provides far more candidate solutions that are also significantly more robust. Both traits make it more likely that robust satisfactory compromise solutions can be found. However, it cannot be inferred from these results that DU optimization will always lead to more solutions, as the presented results may reflect particular characteristics of this problem's specific decision space.

The use of DU optimization improved the maximum attainable robustness, as calculated by the model, of all utilities but OWASA, who was already able to attain nearly 100% maximum attainable robustness in a majority of the WCU optimization solutions. Figure 3.5 shows that Cary, who had maximum attainable robustness of 86% for the WCU solutions reached 100% with the DU optimization's results. Likewise, Durham's maximum robustness went from 39% to 85%, and Raleigh's went from 18% to 59%. Although Durham's and Raleigh's robustness may not be as high as desired, their improvement from the DU optimization is substantial. It is important to notice that Figure 3.4 shows that such improvements in robustness comes at the expense of average expected performance, notably in average costs and Jordan Lake allocation — the transparent blue lines show consistently higher values for these two objectives than the transparent brown lines.

As noted by Herman et al. (2014), beyond their individual levels of attained robustness, it also important to understand and mitigate potential robustness conflicts where one of the utilities strongly benefits by degrading one or more

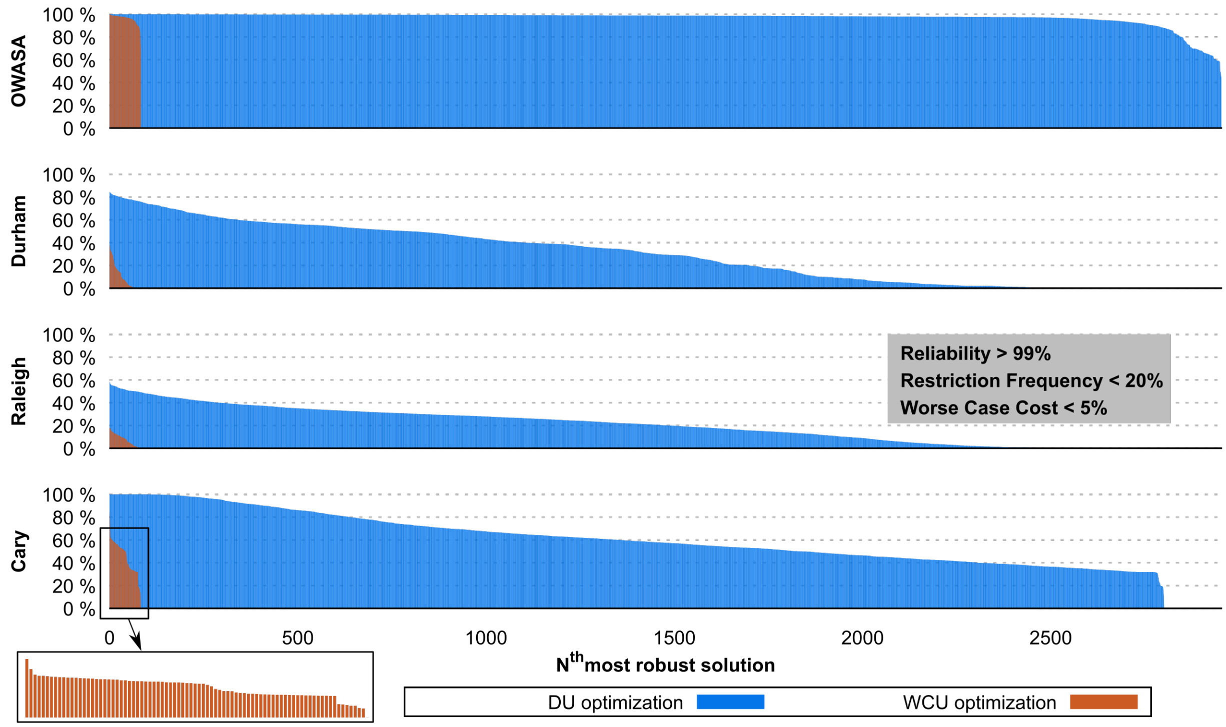

Figure 3.5: Alternatives (x-axis) were rank-ordered according to robustness (y-axis) delivered to each utility. An alternative's robustness was defined as the percentage of the DU re-evaluation's 10,000 $\boldsymbol{\Psi}$ samples where an alternative meets the original robustness criteria elicited from water utilities. The brown bars represent the robustness of the original 84 WCU optimization's alternatives, while the blue bars represent the robustness of all 2,954 DU optimization's alternatives. Each group of bars is rank ordered by robustness.

of the others. Figure 3.6 shows the robustness tradeoffs among all the utilities for the DU optimization alternatives. The ideal solution would be a horizontal line at 100% robustness for all of the utilities. The gray lines are all the solutions from the DU optimization run, while the blue and brown represent illustrative compromise solutions from the DU and WCU optimization runs, respectively, attained by filtering out solutions of high robustness sacrifice for any of the utilities. The gray diagonal crossing lines and color gradient for the DU optimization results suggest that increases in Raleigh's and Durham's robustness are often at the expense of Cary's, due to the amount of transfers that are being demanded. The brown line representing the candidate compromise alternative from the WCU optimization solutions shows low attainable system robustness, which confirms the value of DU optimization.

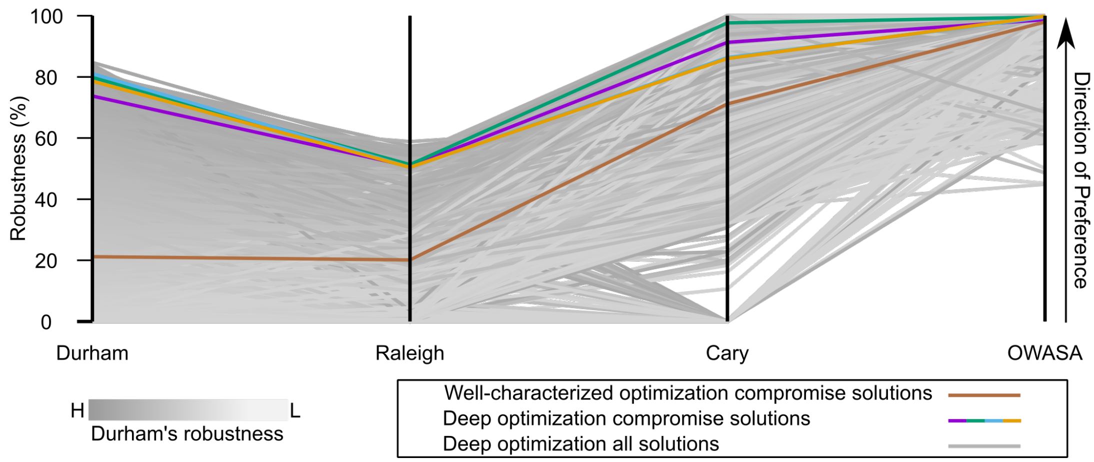

Figure 3.6: Robustness tradeoffs among utilities. Each polyline represents an alternative, while the location where a polyline intersects an axes denotes the robustness value of the corresponding alternative for the corresponding utility. The blue lines indicate four illustrative compromise DU alternatives with fairly high robustness for all utilities, the brown line indicates the best compromise WCU alternative from prior work by Herman et al. (2014), while the gray lines represent the remaining DU alternatives.

Figure 3.6 highlights four compromise alternatives from the DU optimization results that were able to maintain the utilities' robustness levels close to their maximum attained values. The DU optimization effectively exploited the flexibility of the ROF-based transfers and financial instruments under high uncertainty to identify compromise portfolios that strongly reduce regional robustness conflicts. This shows that a more comprehensive approach to uncertainty in the search phase has the potential to not only yield more robust alternatives, but also to decrease regional inter-utility conflicts during negotiations.

Figure 3.7 shows the decision variables for each utility for the four compromise alternatives. According to the figure, under a compromise scenario, Durham and Raleigh (panels a and c) have low restriction and transfer trigger values, which means they would often make use of these two instruments. In order to compensate for the financial losses due to restrictions and transfers, both utilities would have to rely on substantial contingency funds and possibly on third party insurances even for mild droughts.

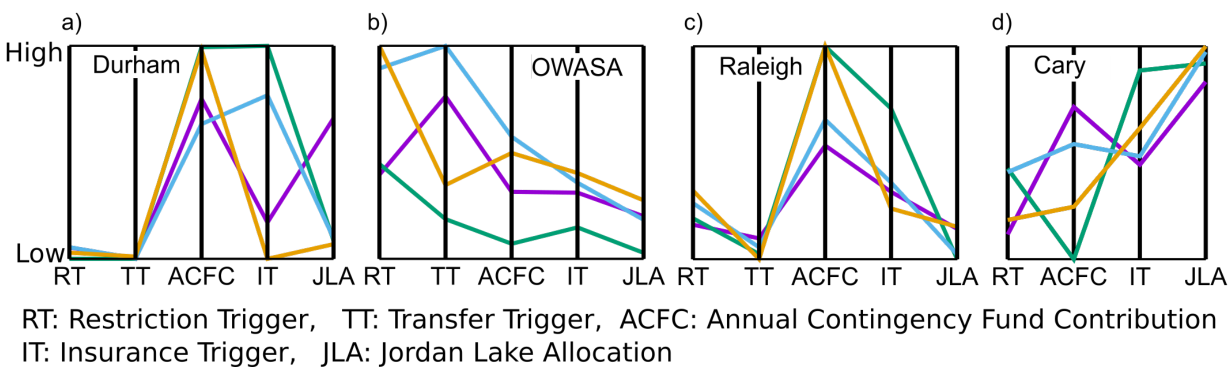

Figure 3.7: Decision variables of the 4 compromise solutions for each utility. Different colors represent the different alternatives, and correspond to the ones in Figure 3.6. The ”y” axis represent a scale from low use to high use.

Being a small utility in comparison with the others, OWASA (Figure 3.7d) is in a more comfortable position, not needing to rely often on restrictions or transfers. As a whole, OWASA seems to be largely insensitive to any of the instruments, as it performs well in most of the tested states of the world. However, due to its small budget and size, a more severe drought poses significant financial risk, which makes it have to rely on relatively low insurances trigger values.

Cary (Figure 3.7b) would require a significant Jordan Lake allocation to meet its demand because it relies on Jordan Lake as its primary water source. However, being the source of all transfers increases Cary's revenue — Durham and OWASA have allocations from Jordan Lake but have to pay a fee to Cary to use them, given Cary owns the Jordan Lake intake and treatment plant. This would put Cary in a relatively safe position against mild droughts, not strongly needing to set up a large contingency fund — which is mainly an instrument to mitigate financial impacts of mild droughts — or insurances that would be triggered often. However, Cary may not have enough revenue from water transfers in exceptionally dry years, given its water sales can be no greater than the allocations possessed by Durham and OWASA, making it have to rely on contingency funds as a complement to the insurance payouts. Under such scenario, the tradeoffs observed in Figure 3.6 between Raleigh and Durham would be more noticeable, given their low transfers trigger values and resulting in possible competition for transfers from Cary, which are limited by conveyance capacities. From these results, it is clear that an increased Falls Lake allocation would be paramount for maintaining Raleigh's reliability and regional robustness during low frequency extreme events, since Raleigh has the highest demands and Falls Lake is its only direct supply source (see Table 1).

The addition of the deeply uncertain factors in the search phase significantly impacted the underlying decisions that would result from the robust compromise policies. Both Raleigh and Durham are shown to require a substantially greater Jordan Lake allocation relative to the results of the WCU optimization shown in Herman et al. (2014). Additionally, the low values of the transfers ROF trigger for both utilities imply that the allocations are going to be frequently used (i.e., more risk averse ROFs). However, the increased transfers do not prevent both utilities from requiring relatively frequent restrictions when compared to previous results. Increased use of transfers and restrictions also serves to increase the required annual contingency fund contributions for Raleigh and Durham. Another important difference between WCU and DU optimization results is that OWASA transitioned from a very limited to a more significant use of restrictions, transfers, and the contingency funds.

3.7.3 The Value of Regionally Coordinated Demand Management

Beyond assessing the robustness of the utilities, it also important to understand the key deeply uncertain factors that control the success or failure of alternative regional water portfolios. Following prior MORDM studies (Kasprzyk et al., 2013; Herman et al., 2014), factor mapping sensitivity analysis, as described in section 3.3.5, was performed to identify the ranges deeply uncertain factors that were most related to system failures, and the results are shown in Figure 3.8. This is analogous to the scenario discovery approach as described by Bryant and Lempert (2010). A conservative approach was taken when selecting the PRIM box. Out of the over 80 boxes for each alternative, boxes with high coverage and low type II error were chosen at the expense of the density metric. The implication of this choice is that the bars on Figure 3.8 emphasize the ranges of parameters where the compromise alternatives are likely to perform satisfactorily.

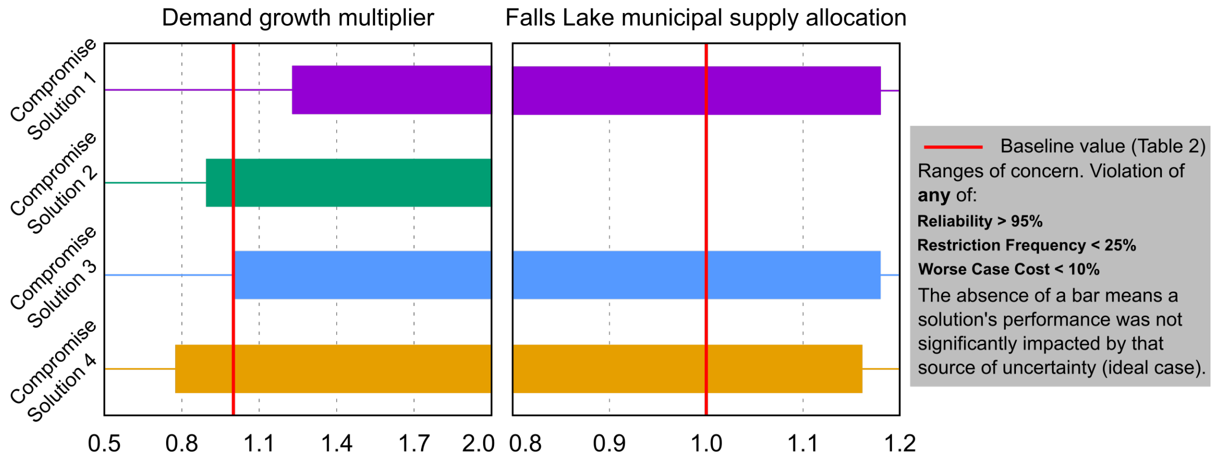

Figure 3.8: Ranges of deeply uncertain factors which should be avoided in order to improve the robustness of each compromise solution. Demand growth (left box) and Falls Lake municipal supply allocation (right box) are the two most important sources of uncertainty for the four chosen compromise solutions. Utilities are encouraged to incorporate measures in their short-term management plan seeking to mitigate demand and be granted a greater allocation from Falls Lake.

The two deeply uncertain factors that most strongly influence the robustness of the four compromise alternatives performances highlighted in Figure 3.8 are regional demand growth rates and the Falls Lake municipal supply allocation. The relative importance of Falls Lake municipal supply allocation in comparison to the deeply uncertain factors other than demand growth reemphasizes the importance of new water allocations/sources for the region, which is also expressed in Figure 3.7 by the high Jordan Lake allocation needed for Cary to be able to provide for frequent transfer requests from Raleigh and Durham. Generally speaking, these results show that utilities should carefully consider regional demand growth rates and possibly monitor the adequacy of the Falls Lake supply as well, depending on the chosen compromise solution. It is important to notice that demand growth multiplier and Falls Lake allocation had a much stronger effect on the system's performance than any climate or inflow factor through the 2025 planning horizon.

3.7.4 Scenario Discovery

Figure 3.9 provides a more illustrative view of a representative PRIM box with a low angle well-defined boundary for success versus failure in attaining the utilities performance requirements along the demand growth vertical axis. This clear performance boundary indicates that even though demand growth is not the only influential parameter determining the acceptability of an alternative's performance under a given $\boldsymbol{\Psi}$, it tends to be the most important one. Any further improvement in robustness for any of the compromise alternatives would then have to be accomplished by first trying to mitigate demand growth.

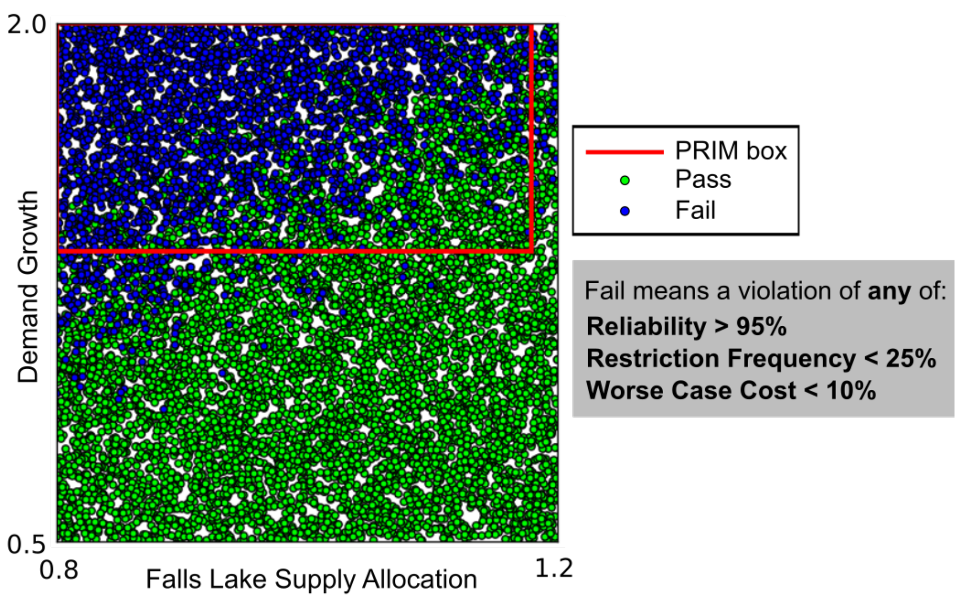

Figure 3.9: Falls Lake Supply allocation vs. Demand growth pass/fail. A solutions performance is more impacted by variations in demand growth than in Falls Lake allocation.

Following this result, a suite of robustness evaluations was performed in order to further understand the effects of changing the demand growth. Figure 3.10 shows the robustness results of robustness evaluations for new LHSs with the demand growth multiplier factor varying from 0.5 to 0.8, 1.0, 1.2, 1.4, 1.6, and 2.0 (results for 1.8 were similar to 2.0). The 0.8 multiplier represents a 20% reduction of the Research Triangle's demand growth rates relative to their assumed nominal value (i.e., the 1.0 multiplier). From a policy perspective, this would require that the utilities coordinate demand growth reductions from approximately 3% through 2025 to approximately 2.4% ($0.8 \cdot 3\%$) instead. As seen in Figure 3.10, demand growth management would be an effective means of mitigating droughts under various tested SOWs. Under scenarios of highly controlled demand growth, Cary, OWASA, and Durham could attain 100% robustness across several water portfolio options. In several of the candidate portfolios Raleigh's robustness exceeds 80%, which represents a drastic gain compared to allowing the demand to grow twice as much as the projected increase.

This work, as well as the prior motivating study (Herman et al., 2014) highlight the value of the water utilities monitoring and innovating their demand management schemes. A key contribution of this chapter relative to prior work is that the proposed DU optimization facilitated the discovery far more water portfolios with improved robustness. Our results highlight that a high level of adaptivity exists even when the system is confronted with the most challenging SOWs. A more detailed understanding of candidate demand management schemes holds significant value for future research. Cooperative water transfers, financial risk mitigation tools, and coordinated regional demand management must be explored jointly to decrease robustness conflicts between the utilities. The insights from this work have general merit for regions where adjacent municipalities can benefit from cooperative regional water portfolio planning.

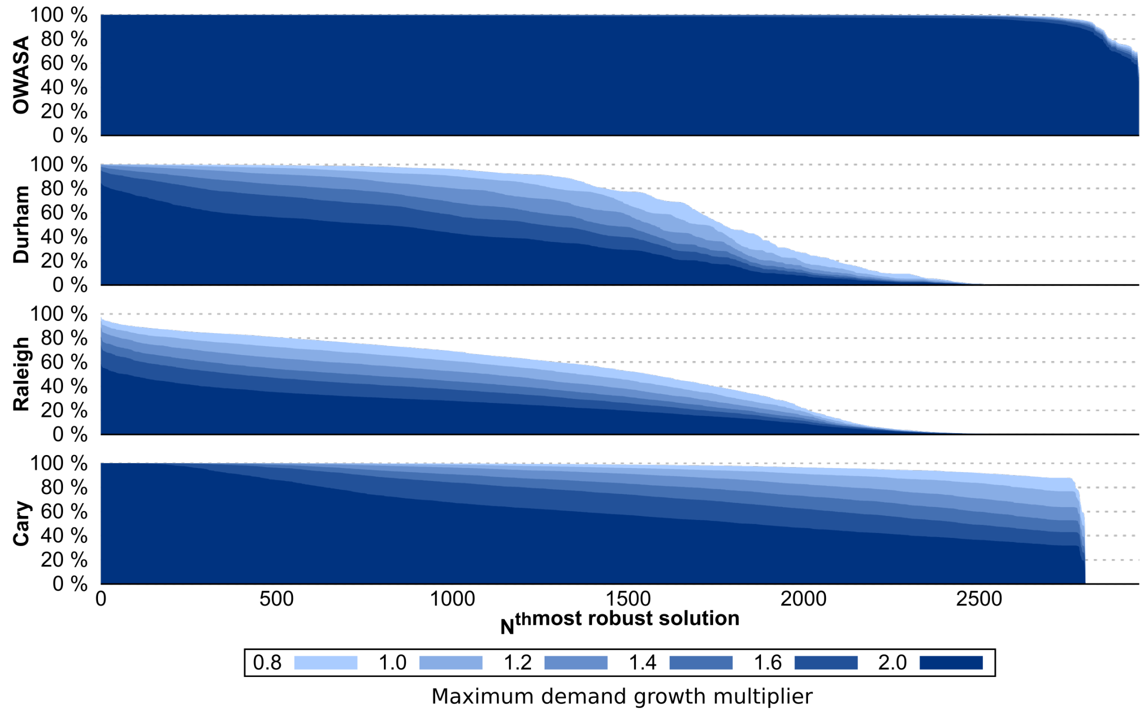

Figure 3.10: Robustness improvement due to demand growth control. Demand growth management may lead to a robustness gain of up 30% for the four utilities studied here. Cary, Durham and OWASA have the potential to reach 100% calculated robustness, while Raleigh may reach over 90%.

3.8 Discussion

While prior studies (e.g., Kwakkel et al. 2014, 2016a; Watson and Kasprzyk 2016) have also advocated for including deep uncertainties into the search for robust management actions, the core of the computational contributions in this work lie in the dynamic and adaptive nature of our ROF-triggered water supply portfolio instruments. Our DU optimization is seeking better risk-based management rules for the Research Triangle's interdependent regional actors. Each of the utilities is balancing supply reliability focused actions (i.e., transfers and restrictions) with their intrinsic financial risks (i.e., extreme tail costs). The combined use of ROF decision triggers and financial hedging is critical for DU optimization to realize its full value. In contrast, commonly employed average cost water supply formulations that focus on searching a space of non-adaptive decisions (e.g., see recent examples in Beh et al. 2015a; Borgomeo et al. 2016; Huskova et al. 2016) are strongly limited in their ability to respond to challenging combinations of deeply uncertain factors. More specifically, as formulated in equations 3.8 and 3.10, ROF is actually a dynamic representation of each utility's evolving capacity to meet their demand. Each utility's dynamic ROF state is subject to its ROF decision thresholds, the conditions of the specific SOW sampled, and the choices of the neighboring utilities. A key distinguishing feature of this portfolio formulation is that every sampled SOW yields a different suite of tailored actions taken in each management period subject to the ROF triggers that compose a given candidate solution. Akin to the concept of a closed loop feedback, each observed ROF state has a state-dependent action triggered. Put simply, this represents acting according to the world you are presently observing. In contrast, the more typical abstraction of drought decisions in the water supply literature (see literature in Beh et al. 2015a; Borgomeo et al. 2016; Huskova et al. 2016) is for each candidate solution to implement the same fixed set of actions in every sampled SOW and evaluate their expected performance measures (i.e., taking actions based on the average performance measures over possible future worlds). This difference is critical to the success of the DU optimization extension demonstrated in this chapter. For each candidate DU optimization solution, its ROF triggers yield water portfolio actions that are specific to the conditions observed in each sampled SOW. Consequently, a much higher degree of adaptivity and enhanced exploitation of information feedbacks are employed when the DU optimization confronts challenging SOWs. Moreover, robustness here requires the financial risk instruments to successfully mitigate extreme tail costs, making higher levels of restrictions and water transfers financially tenable in the most challenging SOWs encountered in the DU optimization (see Figures 3.4-3.6).

Successfully balancing robustness and regret is another potentially more subtle computational contribution in our transition from the WCU optimization to the DU optimization. In both cases, the same risk-based action rules where searched but the resulting action triggers are very different. Challenging SOWs in the DU optimization stressed the ROF-based rules to more fully exploit their potential adaptivity and yielded an strong increase in the number, diversity, and robustness of candidate solutions (see Figure 3.5). However, our results also highlight that increasing the robustness of individual Research Triangle utilities frequently increased the potential for regional resource conflicts (see Figure 3.6). If not properly understood and managed, these increased resource conflicts represent a significant unintended regret that emerges from the DU optimization. As highlighted by Giuliani and Castelletti (2016), robustness focused formulations abstract very high levels of risk aversion by designing for severe SOWs. Consequently, they can be mathematically prone to high levels of regret if solutions are not sufficiently adaptive to take advantage of less severe SOWs. Our results aid in better balancing these types of robustness and regret concerns. Although our most robust solutions are capable to confronting severe SOWs, their underlying ROF-based decision triggers would take actions relative to the world observed. Moreover, our scenario discovery shows that if the Research Triangle's utilities are willing to coordinate in modestly reducing their demand growth rates by 20%, their collective robustness increases dramatically while reducing large, potentially contentious Jordan Lake allocation requests.

Beyond the Research Triangle, our results have implications for rapidly growing urban regions that are now confronting increasing water scarcity concerns. Planning by nearly all water utilities at present, still relies solely on historical streamflow records, although it is common knowledge that even long streamflow records are poor representations of extreme drought quantiles (e.g., Fiering et al. 1971; Lettenmaier and Burges 1978; Potter 1976; Stedinger and Taylor 1982). Shifting from planning based solely on the historical streamflow to synthetic hydrology while also actively accounting for the diversity of factors included in our DU optimization holds significant promise as a means of improving the robustness of regional water supplies. That being said, this recommendation also highlights a rapidly growing gulf between emerging capabilities and the traditional planning frameworks that are currently commonly employed. This chapter also highlights that high density urban regions with multiple water supply utilities can and likely must carefully combine of highly adaptive reliability-driven actions that are supplemented with careful financial hedging. DU optimization was critical for discovering a soft path (Gleick, 2002a) mixture of conservation, transfers, and financial hedging in this chapter that could successfully overcome very severe combinations of stressors. This result has broad value for urban systems around the globe confronting climate change, rapid population growth, economic expansion, and limits on new source development.

3.9 Conclusions

This chapter contributes computational advances for a multi-stakeholder variant of the MORDM framework demonstrating how to include deep uncertainties in the generation of candidate policies to improve robustness and reduce regional robustness conflicts. In the water resources context, this work contributes a clear demonstration of the value of integrated regional water supply portfolio planning for mitigating drought risks, particularly for densely populated municipalities that are in close proximity to one another. The Research Triangle test case demonstrates that significant potential drought risks and inter-utility supply conflicts exist between utilities in the planning period of analysis from 2015 to 2025. This region currently encompasses more than 2 million users and trends of very rapid population growth.

This chapter highlights the need for flexible and scalable modeling frameworks to simulate the broad array of candidate management actions and uncertain stressors in coupled human-natural systems. The contributions of this work extend beyond computational algorithms. The Research Triangle management model that is core to this work requires detailed human system data (reservoir rules, finances, demands, rules for restriction, conditions for transfer, pricing, etc.) as well as natural system data (multiple ensembles of streamflows, evaporation, etc.). Having sufficient abstractions of the human systems strongly constrains the scope of candidate management actions (i.e., transfers, restrictions, and financial instruments) that can be simulated and explored under deep uncertainty. In terms of a broader guidance for future studies, densely populated urban regions with multiple utilities must be carefully abstracted as coupled humannatural systems. A multitude of factors fundamentally shape the degree to which droughts impact water supplies (e.g., magnitude and duration of rainfall deficits, coordination of crisis management, water treatment and conveyance capacities, population growth, financial stability, infrastructure maintenance, etc.). These interdependent factors yield complex dynamically evolving risks. Moreover, the overall systems' behavioral responses to these evolving risks are strongly shaped by asymmetries in regional impacts, highly diverse regional priorities, and conflicting stakeholder preferences. The DU optimization extensions of the MORDM framework contributed in this chapter act in combination with ROF-based decision triggers to provide a broad ability to explore the coupled humannatural systems' dynamics of the Research Triangle. More generally, this work demonstrates that understanding urban water systems' management tradeoffs, vulnerabilities, and dependencies requires modeling and analysis frameworks that are capable of capturing dynamic stocks and flows of risk itself.

The inclusion of deep uncertainties in the search phase of the MORDM framework made it possible to both discover a much more diverse range of drought mitigation actions that the municipalities could take while also improving their overall robustness. Robustness in this work is defined in collaboration with regional stakeholders and encompasses requirements for supply reliability, the frequency of demand restrictions, and the stability of financial risks. Across all of these concerns, our results show that regional demand growth rates dominate all other factors in controlling the robustness of alternative portfolios of drought mitigation actions. In combination, this chapter's improved search under deep uncertainty and coordinating regional demand management helps to eliminate many key tradeoffs and potential inter-utility conflicts. Future work should consider innovative strategies for coordinating individual and regional demand growth rates. In terms of the broader global water resources field, insights from this work have general merit for regions where adjacent municipalities can can introduce cooperative regional water portfolios.

CHAPTER 4

WATERPATHS: AN OPEN SOURCE STOCHASTIC SIMULATION FRAMEWORK FOR WATER SUPPLY PORTFOLIO MANAGEMENT AND INFRASTRUCTURE INVESTMENT PATHWAYS

Funding for this work was provided by the National Institute of Food and Agriculture, U.S. Department of Agriculture (WSC Agreement No. 2014-67003-22076). Additional support was provided by the U.S. National Science Foundation's Water, Sustainability, and Climate Program (Award No. 1360442). The views expressed in this work represent those of the authors and do not necessarily reflect the views or policies of the NSF or the USDA.

4.1 Abstract

Financial risk, access to capital, and regional competition for limited water sources represent dominant concerns in the US and global water supply sector. This chapter introduces the WaterPaths framework: a generalizable, cloud-compatible, open-source exploratory risk-modeling framework designed to inform long-term regional investments in water infrastructure while simultaneously aiding regions to improve their short-term weekly operational decisions. Uniquely, WaterPaths has the capability to identify the challenges and demonstrate the benefits of regionally coordinated planning and management for groups of water utilities sharing water resources. As a platform for decision making under deep uncertainty, WaterPaths accounts for uncertainties not only related to hydrological or climate extremes, but also to key urban systems factors such as demand growth, effectiveness of water-use restrictions, construction costs, and financing uncertainties. The WaterPaths platform is introduced

here through the fictional and realistic Sedento Valley test case, in which three resources-sharing water utilities improve their individual and joint infrastructure investments and weekly operations subject to various sources of deep uncertainty to attain higher supply and financial performance. The Sedento Valley test case was designed to serve as a universal test case for decision-making frameworks and water-resources systems simulation software.

4.2 Software Availability

- Name of Software: WaterPaths
- Description: WaterPaths is an open-source C++ model for the stochastic simulation of decision-making policies by water utilities concerning the use and upgrade of single- and jointly-owned infrastructure. It was designed to be easily customizable and to be used on high-performance-computing clusters and cloud for single and batch simulations and for policy optimization when coupled with a black-box multiobjective optimization algorithm. WaterPaths can export detailed time-series output of various system states and the values of objective functions (performance metrics) defined by the user. Included is also an example test case named Sedento Valley.
- Developer: B. Trindade (bct52@cornell.edu) with contributions by P. Reed. D. Gold contributed to the development of the Sedento Valley test case.
- Funding Source: Funding for this work was provided by the National Institute of Food and Agriculture, U.S. Department of Agriculture (WSC Agreement No. 2014-67003-22076). Additional support was provided by

the U.S. National Science Foundation's Water, Sustainability, and Climate Program (Award No. 1360442).

- Source Language: C++
- Supported Systems: Unix, Linux, Windows, Mac
- License: Apache 2.0
- Availability: https://github.com/bernardoct/WaterPaths

The code used for the WaterPaths optimization runs can be found in the git commit 1d07e944654569236dc496e5d6fa9933c75de9cd.

The code used for the WaterPaths re-evaluation runs can be found in the git commit 5b612cc8eae39acf4411d9674d897c03943f1f27.

4.3 Introduction

Recent projections estimate that the United States (US) will require over one trillion dollars of investment in water supply infrastructure in the next 20 years (ASCE, 2017). This investment represents a difficult challenge as water fees, which often fund water infrastructure projects, are already rising above the Consumer Price Index (USWA, 2019b,a; Hall et al., 2019) and initiatives such as the Water Infrastructure Finance and Innovation Act (commonly known as WIFIA) (Copeland, 2016) have only been able to address a fraction of the required investment (EPA, 2019). This challenge is mirrored around the world as water managers face the task of maintaining supply reliability under climatic, social, and financial uncertainties (Bonzanigo et al., 2018; AWWA, 2019). In many regions, rapid shifts in seasonal climate, drought patterns, and the frequency of extreme events threaten to overwhelm existing water supply capacity (Shafer and Fox, 2017). These threats are felt more acutely in large metropolitan areas, which have observed strong and sustained population growth in recent decades (McNabb, 2019; Census, 2019). To face these challenges, water managers must move beyond historic paradigms of water resources planning and management and transition to new methods of sustainable freshwater management (Gleick, 2018).

Historically, water managers have relied on infrastructure expansion to confront long-term supply risks and utilized water use restrictions to manage short-term drought crises (Gleick, 2002b). This approach often yields high financial burdens as utilities are forced to accumulate large amounts of debt. Revenue disruptions resulting from water use restrictions may cause utilities to miss payments on this debt, threatening their financial stability, and reducing their ability to make future investments (Hughes and Leurig, 2013; Moody's, 2017, 2019). Additionally, the land or water resources most suitable for new water supply infrastructure have largely been developed (Lund, 2013) and new water infrastructure projects (such as large dams) are often environmentally destructive and unpopular. These challenges have led to the proposal of soft path strategies to water resources management that seek to compliment centralized infrastructure with non-structural measures to improve efficiency and manage water demand (Gleick, 2002b).

Exploring synergies between long-term infrastructure investment pathways and non-structural drought mitigation instruments is one promising method for incorporating soft path strategies into water supply planning and management. Potential short-term management instruments include demand management

(Hall, 2019; Zeff et al., 2014; Borgomeo et al., 2018; Huskova et al., 2016; Haasnoot et al., 2013), treated water transfers through regional shared infrastructure (Caldwell and Characklis, 2014; Zeff et al., 2014; Watkins and McKinney, 1997; Hall, 2019), and raw water transfers through upstream reservoir releases (Borgomeo et al., 2018; Gorelick et al., 2019). Integrating these short-term management instruments with long-term infrastructure sequencing allows utilities to create adaptive management strategies that maintain reliability under drought conditions while minimizing the need for infrastructure expansion (Zeff et al., 2016).

The coupling of long-term infrastructure investment strategies with short-term drought mitigation necessitates innovations in decision support systems for water resources management problems. Broadly, decision support systems are composed of the suite of analytical, mathematical, and/or simulation-focused tools that aid planners in the design and operation of water resources systems (Loucks and Da Costa, 2013). There is large historical body of literature water resources planning and management where simulation models are a central focus for analysts developing insights on how a system behaves under varying system states or decision maker actions (see review in Loucks and van Beek 2017). Since the inception of the field in the early 1960s, these water systems models have typically consisted of networks of storage and junction nodes linked by conveyance structures such as pipelines, canals and river reaches (Maass et al., 1962; Harou et al., 2009). There a large range of modern example simulation software that have held substantial benefits for decision support applications including MODSIM (Labadie, 2011), Water Evaluation and Planning (WEAP) (Sieber, 2006), Interactive River-Aquifer Simulation (IRAS-2010) (Matrosov et al., 2011), WATHNET (Kuczera 1992), RIBASIM (Hydraulics, 2004), MIKE BASIN (Jha and Gupta, 2003), Source (Welsh et al., 2013), CALVIN (Draper et al., 2003) and OASIS (HydroLogics, 2009). For a summarized feature comparison between WaterPaths and other mentioned software/framework, see Tables 4.3 and 4.3.

More recently, there has been a growing interest in multiobjective simulationoptimization applications where models such as MODSIM, IRAS and WATHNET are coupled with multiobjective evolutionary algorithms (MOEAs) to aid in the discovery of key water supply planning and management tradeoffs (Matrosov et al., 2015; Borgomeo et al., 2018; Basdekas, 2014). MOEAs are global population-based search algorithms that evolve sets of Pareto approximate solutions to multiobjective problems through processes of mating, mutation and selection (for reviews see Nicklow et al. 2010, Maier et al. 2014 and Coello et al. 2006). However, analytical decision support tools such as MOEAs require the ability to develop an effective and efficient software coupling with simulation software. Moreover, typical applications substantially increase the computational demands of analyses because MOEA search typically requires thousands to tens of thousands of simulation-based evaluations of performance objectives to guide their search processes.

Two more nuanced concerns with regard to modern water systems simulation software is that (1) they often do not provide users with flexibility in representing complex state-aware actions as either an artifact of their approach to water balance modeling (e.g., optimization or rule-based allocations that cannot capture information feedbacks) and (2) they contain limitations that arise from a lack of access to their underlying source code bases (e.g., commercial software packages). Likewise, despite the growing recognition of that a broad

Table 4.1: Select modeling software for water resources planning and management

Framework	Purpose	Incorporated Uncertainty	Routing Method	Financial Modeling Capability	MOEA Ready	Represented Time-scales
IRAS-2010	Computationally efficient simulation model for flows and storages, multi-reservoir release, water consumption, hydropower and pumping energy use	Demand and inflow	Rule Based	Capital costs, operating costs, energy costs and hydropower revenue	Yes	Daily to decadal
OASIS	River Basin Management, Hydropower, Water Supply, Conflict Resolution	Inflow and demand	Optimization Driven	None	No	Minute to decadal
WEAP	Simulation of water supply systems through watershed-scale hydrologic processes, water demands and envrionmental requirements	Inflow, demand, land use	Optimization Driven	None	No	Daily to decadal
MODSIM	River Basin Management for analysis of long term planning, midterm management and short term operations	Inflow and demand	Optimization Driven	None	Yes	Minute to decadal
WATHNET	Simulation model for streamflow storage and transfer and water demand	Inflow and demand	Optimization Driven	None	Yes	Daily to decadal
RIBASIM	Multi sector planning to allocate water resources at the river basin level	Inflow and demand	Rule Based	Crop yeild and production costs	No	Daily to decadal
WaterPaths	Flexible simualtion and optimization to aid in long term infrastructure sequencing and short term drought mitigation	Inflow, demand, customizable suite of uncertainty	None	Water revenues, demand mangement and drought mitigation cost. Long term capital costs and infrastructure financing	Yes	Weekly to decadal

Table 4.2: Select modeling software for water resources planning and management (continued)

Framework	Batch processing	Endogenous Infrastructure Expansion	Represented decisions	GIS integration	Custom GUI	Extendable	Access	Illustrative citations
IRAS-2010	Yes	No	Static infrastructure and demand management	Yes, external	Link to hydroplatform	Yes, Fortran	Free	(Matrosov et al., 2011, 2015)
OASIS	No	No	Custom River Basin Operating Rules	No	Yes	No	Commercial	(HydroLogics, 2009)
WEAP	No	No	Custom River Basin Operating Rules	Yes	Yes	No	Free	(Sieber, 2006; Yates et al., 2005)
MODSIM	No	No	Custom River Basin Operating Rules	Yes	Yes	Yes, C# .NET or Visual Basic .NET	Free	(Labadie, 2011; Basdekas, 2014)
WATHNET	Yes	No	Static infrastructure and demand managment	Yes	Yes	Yes, custom scripting	Free	(Kuczera, 1992; Borgomeo et al., 2018)
RIBASIM	No	No	River Basin Operating Rules	Yes	Yes	No	Free	(Hydraulics, 2004)
WaterPaths	Yes	Yes	Dynamic state-based risk triggers, Custom metrics	No	No	Yes, C++	Free	This Manuscript

array of uncertainties strongly shape water resources systems (e.g., financial risks, behavioral responses, changing hydro-climatic conditions, shown Chapter 3), currently available water resources systems simulation software are again highly constrained in their representation and scalable computational support for including these concerns in water resources infrastructure investment and management applications.

Beyond limitations in current water systems simulation software, broader conceptual challenges must be considered in our decision support framework's treatment of uncertainties, in particular the presence of deep uncertainties (for recent reviews see Moallemi et al. 2019; Dittrich et al. 2016; Kwakkel and Haasnoot 2019; Herman et al. 2015). Deep uncertainty refers to conditions when decision makers do not know or cannot agree upon probability distributions of key system parameters and/or the system boundaries (Kwakkel et al., 2016a; Lempert, 2002). Planning and management under deep uncertainty shifts the task of discovering optimal management strategies for water resources systems, to crafting robust and adaptive strategies that maintain performance across a wide array of potential future conditions (Walker et al., 2013; Dittrich et al., 2016). Recent work in water resources systems planning and management has focused on bottom up decision. Frameworks such as Decision Scaling (Brown et al., 2012), Robust Decision Making (RDM) (Lempert et al., 2006), Many Objective Robust Decision Making (MORDM) (Kasprzyk et al., 2013) and Info-gap (Ben-Haim, 2006) provide methods to aid decision makers in the discovery of key uncertainties control system vulnerability. These methods are often computationally demanding exploratory Monte Carlo simulation software, that strongly benefit from highly scalable simulation software that can be adapted across a range of state-of-the-art computing architectures.

The challenges and needs summarized above motivated this chapter's contribution of the open-source WaterPaths is a stochastic simulation software that has been specifically develop to support applications of the DU Pathways decision support framework. WaterPaths allows for a flexible and efficient representation of multi-actor water resources systems while providing advanced computational support for multiobjective optimization algorithms and exploratory analyses of a broad range of uncertainties. WaterPaths has been specifically developed to focus on water supply infrastructure and water portfolio management applications for systems confronting water scarcity or increasingly severe droughts. WaterPaths provides decision makers with ability to flexibly abstract water infrastructure portfolio problems that contain measures for addressing short-term supply and financial impacts of droughts along with long-term capacity expansions. WaterPaths is developed to run on desktops, cloud computing systems, and on high performance computing resources, utilizing both shared memory and distributed memory parallelization to enable decision makers to efficiently design, optimize and evaluate candidate water supply portfolios over large ensembles of potential future states of the world.

We demonstrate WaterPaths by introducing the Sedento Valley Test Case, a new highly detailed and realistic hypothetical multi-actor regional water supply planning test case. WaterPaths is used to represent three hypothetical water utilities in the south eastern US that are facing the prospect of water shortage due to growing demand and changing climate. The utilities seek to craft both short-term drought mitigation responses and long-term infrastructure sequencing pathways that maintain reliable supply and financial stability. The utilities have the potential to cooperate using treated transfers through existing infrastructure and through shared infrastructure development. The DU Pathways framework is used to assess the utilities tradeoffs, robustness, and infrastructure investment pathways as a means of showcasing the key features of WaterPaths.

The framework is demonstrated on the Sedento Valley Test Case, a new generic multi-actor test bed for regional water supply planning. The model represents three hypothetical water utilities in the South Eastern United States facing the prospect of water shortage due to growing demand and changing climate. The utilities seek to craft both short-term drought mitigation responses and long-term infrastructure sequencing pathways that maintain reliable supply and financial stability. The utilities have the potential to cooperate using treated transfers through existing infrastructure and through shared infrastructure development. The DU pathways explored through this test case showcase the features of WaterPaths.

This chapter is organized as follows: Section 4.4 details the WaterPaths framework, Section 4.5 introduces the Sedento Valley test case, presents its development in WaterPaths and overviews the computational experiment, Section 4.6 presents results and discussion and details and Section 4.7 provides concluding thoughts.

4.4 WaterPaths Framework

4.4.1 WaterPaths Overview

Planned adaptation has been studies as an effective way to cope with deep uncertainty (Walker et al., 2003). One approach to adaptive infrastructure planning is Dynamic Adaptive Policy Pathways (DAPP). DAPP uses adaptation tipping points, system conditions that cause an alternative to fail to meet specified objectives (Kwadijk et al., 2010) to design signposts over 30 to 100 year-long schedules that trigger pre-specified adaptive actions(Haasnoot et al., 2013; Kwakkel et al., 2014; Kingsborough et al., 2016, 2017; Zandvoort et al., 2017). DAPP, however, still maintains a strong reliance on the use of limited numbers of predefined action sequences that require high levels of institutional stability and strong consensus over the long-term. Real Option Analysis (ROA) (Cox et al., 1979) is a probabilistic decision process for flexibility and adaptability in the context of irreversible decisions based on decision trees, lattices and Monte Carlo analysis (Erfani et al., 2018). Although recently applied in analysis of reservoir development and expansion (Erfani et al., 2018; Fletcher et al., 2019), drought planning (Fletcher et al., 2017) and flood protection (Hui et al., 2018), ROA is limited in its accounting of stakeholders with diverse interests for being a single-objective approach and its underlying mathematical tree logic quickly becomes computationally intractable when a large number of uncertainties are considered (Dittrich et al., 2016). Borgomeo et al. (2018) propose a multi-objective evaluation of candidate water supply investments as a way to consider diverse interests in which the authors maximize robustness and minimize risk and cost of a 30-year long fixed construction schedule for multiple infrastructure options. However, the resulting static sequence of investments is not adaptable to regional developments and to the local political process, a limitation only partly addressed by Beh et al. (2015b) and Huskova et al. (2016).

The recently proposed Deeply Uncertain Pathways (DU Pathways) framework, presented in detail in Chapter 5, utilizes state-adaptive portfolio based strategies that employ risk based rule systems to define infrastructure sequencing, drought mitigation response and financial instruments and tailor actions to observed system states. The framework builds upon the Many MORDM framework to include the design of rule systems that react to observed system states. DU Pathways provides a means of implementing closed loop control policies (Bertsekas et al., 1995) for infrastructure sequencing and drought mitigation. The simultaneous optimization of risk-based infrastructure triggers and drought-mitigation instruments effectively bridges the gap between long-term infrastructure investment and short-term water portfolio management. Chapter 5 demonstrates that the DU Pathways approach can produce robust and adaptive infrastructure pathways that perform well across broad sets of uncertainty and balance conflicting interests within a multi-actor systems. However, DU Pathways framework involves upgrades of the infrastructure system over time, is computationally intensive due to its required stochastic simulations, and requires large numbers of function evaluations during MOEA search, making it of difficult implemention with existing decision-support systems. To aid in the development and application of DU Pathways a new decision support framework is needed for their design and evaluation. This need motivated the development of WaterPaths.

WaterPaths is a open-source and modular model for the simulation and optimization of water infrastructure planning and management policies. WaterPaths designed to facilitate the application of the DU Pathways methodology and for easy addition of new features and other customization by an user with basic knowledge of C++. WaterPaths includes in one simulation stochastic risk calculations, stochastic uncertainty-based objective calculations, evolution of infrastructure system, and regionally coordinated planning operations of a group of water utilities. On a high level, the simulation of a system of cooperating water utilities and water infrastructure on WaterPaths is organized in terms of mass-balance performed for each piece of water infrastructure, utilities drawing water from and building new water infrastructure, and drought mitigation and financial policies acting on utilities to ensure safe operations from a supply and financial points of view. Each time it is run for a system, WaterPaths simulates a number ranging from hundreds to thousands of realizations (independent scenarios) and calculate the final performance metrics (objectives) based on averages and min-max rules of various system-state variables across all realizations. Each time step (day, week, or month) follows the following sequence, shown in Figure 4.4.1: (1) calculation of long-term ROF metric if at first week of the year, (2) calculation of the short-term ROF, (3) application of drought mitigation and financial policies, (4) system mass-balance modeling, and (5) system-state data collection.

The addition of new features such as new types of water infrastructure, customized reservoir control rules, and new drought mitigation policies is done by extending already implemented abstract classes. Similarly, any system metric for decision-making purposes can be implemented with the addition of one function to the existing code. The mass-balance model, ROF calculation model, and other aspects of the system-wide simulation are designed to work with any child class of the existing abstract classes, making it possible for a user to add various features to WaterPaths without venturing into the existing WaterPaths code.

Lastly, WaterPaths was designed for centralized system-state information recording to allow for easy implementation of custom objective functions. It currently has five objectives implemented, but more can be easily added to the

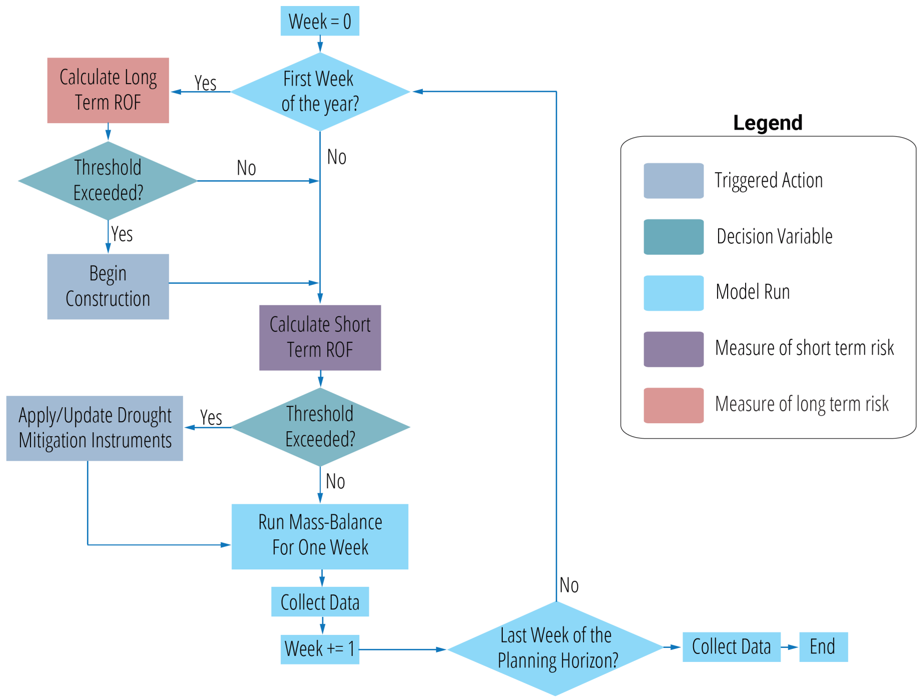

Figure 4.1: WaterPaths main simulation loop. The main loop includes the (1) calculation of long-term ROF metric if at first week of the year, (2) calculation of the short-term ROF, (3) application of drought mitigation and financial policies, (4) system mass-balance modeling, and (5) system-state data collection.

MasterDataCollector class. The currently implemented objectives are reliability, restriction frequency, net present value of infrastructure built during the life of the policy, maximum annual financial cost of drought mitigation and infrastructure construction, and worse-first-percentile annual financial cost of drought mitigation and infrastructure construction. Each of WaterPaths features are discussed in the subsequent sections.

4.4.2 ROF-based decision rules

WaterPaths formulates investment and management policies as state-aware rules that serve as action triggers based on short- and long-term ROF metrics (Palmer and Characklis, 2009; Zeff et al., 2016). More specifically, the default representation of decision policies in WaterPaths trigger each component action in a utility's investment and management portfolio in a given week if the calculated value of the ROF state x_{rof} reaches trigger values. These decision triggers are the core decision variables and allow for consistent and adaptive consideration of short-term temporary management actions or longer-term permanent water supply capacity expansions. As discussed in prior studies (Zeff et al., 2016) and in Chapter 3, the ROF-basis allows for low dimensional closed-loop rules that take actions tailored to the state of the world (SOW) being experienced (i.e., different action sequences for wet versus dry scenarios). Additionally, the ROF metric can be parametrized to capture only single-year droughts, the focus of short-term drought management instruments (here water use restrictions, transfers and drought insurance), and multiple-year drought, the focus of long-term actions (infrastructure construction). Equations 4.1 to 4.3, partly repeated here from Section 3.5 for completeness but now expandaded, mathematically define how the short- and long-term ROF metrics are computed. Figure 4.2 presents an equivalent graphical representation while making the distinction between the short- and long-term ROF metrics.

$$x^{w}_{srof,j} = \frac{1}{N_{rof}} \sum_{y'=0}^{N_{rof}} f^{w}_{y',j}(\boldsymbol{NI}^{y'}, \boldsymbol{E}^{y'}) \tag{4.1}$$

where,

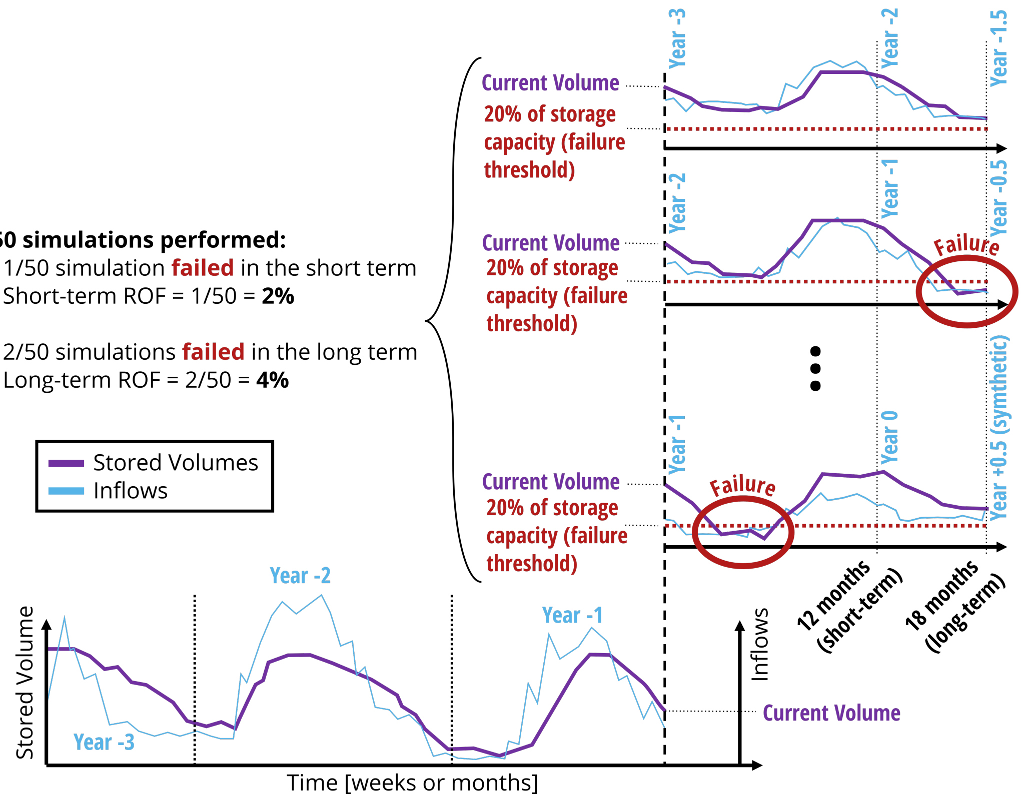

Figure 4.2: For the calculation of the short-term ROF metric for the current week, 50 year-long simulations are ran for the coming year with all reservoir storage levels starting at their current levels, while the streamflows and evaporation rates for each simulation are those of one of the previous 50 years of recorded data and the demand are those of the previous year. For the calculation of the long-term ROF metric, the 50 year-long simulations are extended to one and a half years, with the last 6 months of the last simulation (the last year) are filled by synthetic data. The addition of 6 months to the 50 ROF simulations allows for capturing consecutive drought years in which the reservoirs are not filled during the wet season.

$$f^{w}_{y',j} = \begin{cases} 0 & \forall w' \in \{(y',w), ..., (y', w + T_{rof})\} \ : \ \frac{x^{y,w'}_{s',j}}{C_j} \geq s_c \\ 1 & \text{otherwise} \end{cases} \tag{4.2}$$

and,

$$x^{y',w'}_{s',j} = f\left(C_j, \boldsymbol{UD}^{w}_{j}, \boldsymbol{NI}^{y',w'}_{j}, \boldsymbol{E}^{y',w'}_{j}, \boldsymbol{W}^{y',w'}_{j} | \boldsymbol{\Psi}_s\right) \tag{4.3}$$

In equations 4.1 to 4.3, w' and y' mean a week and a year simulated with past data for the calculation of the ROF. Variable $x^{w}_{rof,j}$ is the ROF for utility j in current week w, and $f_{y',j}$ is a binary variable for which 0 denotes a failure happened during the corresponding ROF simulation with data from past year y'. In the ROF metric, $\boldsymbol{NI}^{y'}$, $\boldsymbol{E}^{y'}$ and $\boldsymbol{W}^{y'}$ are the recorded natural reservoir inflows, evaporation rates and reservoir spillage, respectively, in year y' prior to current week w used in one of the N_{rof} simulations — note that in a simulation in which the ROF metric is calculated for a week 20 years from now, 20 of the N_{rof} years of hydrological data will belong to synthetically generated data. Variable T_{rof} equals 52 weeks for the short-term ROF, so that single-year droughts are captured, and 78 for the long-term ROF, so that each ROF year-long simulation potentially captures two dry seasons. A failure is defined as the combined storage $x^{y',w'}_{s',j}$ for any realization y' divided by combined storage capacity C_j for utility j falling below critical storage s_c. Variable $x^{y',w'}_{s',j}$ is the vector of storage states calculated in one of the year-long ROF simulations using recorded hydrologic data from past year y'. In the calculations of $x^{y,w'}_{s',j}$, $\boldsymbol{UD}^{w}$ is the unrestricted demand (no restrictions enacted or transfers purchased) in week w, and $\boldsymbol{\Psi}_s$ is a matrix of sampled deeply uncertain factors obtained through WCU or DU optimizatio. Given a year has 52.178 weeks (365.25 / 7), every 6 years (when the rounded number of weeks in a year would reach 53) the first week of the following historical year will be used. Although the ROF metric been proved to be an effective metric for water-systems decision making, other metrics such as the more conventional days-of-supply-remaining or new metrics can be implemented and tested with WaterPaths for the Sedento Valley test problem. The calculation of the ROF metric as well as the mass-balance simulation for the system at hand are dependent on WaterPath's mass-balance model, described next.

4.4.3 Mass-Balance Model

The core of WaterPaths is its mass-balance model. The mass-balance model solves mass-balance equations for all water infrastructure following its upstream-to-downstream order, found through topological sorting based on infrastructure connectivity information provided by the user as a directed graph. The mass-balance equation for reservoirs is shown in equation 4.4:

$$x_s^{w+1} = x_s^w + NI^w + SE^w + URO^w - ER^w \cdot RA(x_s^w) - EO^w - S^w - RD^w \tag{4.4}$$

where x_s^{w+1} is the volume of water stored in the reservoir at the week after the current week w, NI is the natural inflow into the reservoir from all its tributaries, SE is a treated sewage effluent discharged either on a tributary or directly on the reservoir, URO is the upstream reservoir total outflow (if such reservoir exists, mandated outflow plus spillage), ER is a non-dimensional evaporation rate, RA is the reservoir area as a function of stored volume, EO is the environmental outflow, RD^w is the total municipal demand drawn from that reservoir by one or more independently modeled utilities, and S is the reservoir spillage, which is set to zero unless the reservoir is completely full. If a reservoir is allocated for multiple uses (municipal supply to multiple utilities, water quality, etc.), each with a designated percentage of the total storage capacity, the inflow is split across all uses proportionally to their allocated volumes. Water in excess of any allocated capacity is redistributed among the others.

To allow for more realistic simulation of reservoirs, WaterPaths also provides an abstract class for reservoir control rules named MinEnvFlowControl and one for controls of other discharges such as treats waste water discharges (the latter flowing into a reservoir or stream), named ControlRules. WaterPaths also contemplates non-storage water infrastructure such as water intakes and water re-use stations, both of which already implemented, although others such as desalinization plants can be easily implemented by extending the abstract class WaterSources based on a custom mass-balance function. WaterPaths also has implemented upgrades of existing infrastructure, namely reservoir expansions and treatment capacity expansions, both meant to be potentially triggered throughout the simulation by one or multiple utilities based on the long-term ROF or a custom metric.

4.4.4 Capturing Regional Decision Making

Being designed for regional decision making, WaterPaths can simulate multiple water utilities, potentially with interconnected networks shared infrastructure, as separate elements within the system. Each utility in WaterPaths uses information about the allocated available volume of water from its infrastructure available (e.g. stored volumes, weekly volume allowed to be taken with an intake on a river, re-use station's treatment capacity in a given week) to split its demand among them, drawing as much from non-storage infrastructure before drawing from reservoirs. The volume of water in storage owned by a given utility is also used to calculate its short- and long-term ROF metrics described in Section 4.4.2, based on which each utility triggers the construction of new infrastructure and drought mitigation policies. The demand each water utility is required to fulfill may be mitigated or increased drought mitigation policies, described in Section 4.4.5, before it is drawn its infrastructure.

Additionally, water utilities in WaterPaths track their own finances. The finances as modeled as cost fluctuations due to drought mitigation and financial instruments and debt repayment, described respectively in Sections 4.4.5, 4.4.6, and 4.4.7, around operations and maintenance costs. The total annual revenue is also calculated based demand, consumer tiers, and corresponding tariffs, and is used to support objectives calculations and financial instruments.

4.4.5 Drought Mitigation Instruments

Building from prior works (Zeff and Characklis, 2013; Zeff et al., 2014, 2016; Palmer and Characklis, 2009; Caldwell and Characklis, 2014) and from Chapter 3, WaterPaths has implemented water-use restrictions and inter-utility, treated-water transfers, although other instruments can be easily added by creating a new child class of the DroughtMitigationInstrument class. Water-use restrictions as currently implemented in WaterPaths can be implemented on multiple tiers, each with different percentages of demand reduction based on stricter measures (reductions in lawn irrigation, car and sidewalk washing, etc.) and triggered by a different value of the ROF metric. As currently implemented, revenue losses due to restricted water sales can be mitigated by adjusted tariffs as provided by the user.

Treated-water transfers in WaterPaths are performed from a source utility to requesting utilities. As water use restrictions, requests for treated transfers are contingent on the current value of the short-term ROF metric for each requesting utility. However, treated-water transfers often suffer from two constraints: treatment capacity at the source utility and limited inter-utility conveyance capacity. To calculate the volume granted to each utility, WaterPaths solves a constrained allocation problem using a quadratic-programming algorithm (Goldfarb and Idnani, 1983; Gaspero, 2007). The problem is set up to minimize the mean-square error between the volumes requested by each utility adjusted for their ROF values (a utility with a higher ROF receives proportionally more water) and the volumes that can be transferred through the inter-utility network subject to conveyance constraints. Funds are then transferred from all requesting to the source utility. Drought mitigation instruments, however, can be costly, and utilities may benefit from setting in place financial instruments to absorb and stabilize such costs. More details about water-use restrictions and treated-water transfers are presented in Equations 3.6 and 3.7.

4.4.6 Financial instruments

WaterPaths allows utilities to hedge against the negative financial effects of the drought mitigation instruments by using financial instruments, such as the currently implemented drought insurance and contingency funds (Zeff et al., 2014). The drought insurance currently implemented in WaterPaths triggers a payout of a percentage of the previous year's annual revenue whenever the value of the ROF metric reaches a set value. The policy is updated, priced, and bought by the utilities every year based on the most current values of annual revenue and on a premium of 20% of the expected cost of the policy for the following year. Other types of insurance policy based on state variables other than the ROF can be easily added to WaterPaths by creating a child class of the DroughtMitigationPolicy class.

On the other hand, contingency funds are not dependent on the ROF metric, but on a fixed percentage of the annual revenue. On the last week of every year, each utility that has a contingency fund adds a fixed percentage of that year's total revenue to its contingency fund. These funds are then used during the following year to pay for drought mitigation policies, other financial policies and for debt issued to build new infrastructure. The drought mitigation and financial instruments, however, may not be enough to protect a utility against the unexpected negative financial effects of droughts, in which case construction of new infrastructure may be needed. More details and drought insurance and contingency funds are presented in Section 3.5.

4.4.7 Infrastructure Investment

As with the drought mitigation instruments and insurance, infrastructure construction is also triggered by an ROF value, although that of the long-term ROF metric. If a utility or region has several candidate infrastructure investment options, users or MOEA-based search must provide a prespecified construction sequence (i.e., a infrastructure pathway across time). To allow for infrastructure construction over the course of the simulated planning period, WaterPaths relies on information passed by the user or optimization algorithm about all infrastructure options (storage capacity, treatment capacity, cost, etc.) and about the sequence in which a utility is to build such options if the value of the long-term ROF metric reaches its trigger value. This is a unique feature of WaterPaths, which allows the user to prioritize options in the near term and define requirements.

Infrastructure options to be considered for construction can be located anywhere in and outside of the reservoir connectivity network. Infrastructure options in WaterPaths may take form of altogether new projects, such as new reservoirs, water intakes, treatment plants, and others, and of expansion of current infrastructure such as storage treatment capacity expansions and reservoir reallocations. This allows WaterPaths to account for flexible infrastructure development, justly emphasized in the Real Options Analysis literature as important elements in the long-term planning for water utilities (Fletcher et al., 2017, 2019; Erfani et al., 2018; Wang and de Neufville, 2005; Cox et al., 1979).

Infrastructure development is often financed over decades rather than paid upfront. WaterPaths allows for multiple types of bonds to be issued for financing new infrastructure, resulting in different possible debt repayment streams with different net present values for the same infrastructure option. Currently implemented in WaterPaths are level-debt-service, balloon-payment, and variable-interest bonds. However, WaterPaths allows for the design of creative finance mechanisms by allowing the user to create new types of bonds by creating children classes of the Bond class and including those in optimization and simulation exercises. Infrastructure development and policy design are highly dependent on uncertainty analysis, deemed as a central concern during the design of WaterPaths. The uncertainty handling within WaterPaths is described next.

4.4.8 Flexible Representation and Evaluation of Uncertainties

One of the key features that differentiate WaterPaths from existing simulation systems is its stochastic treatment of a broad array of uncertainties. As discussed in the Section 4.3, multiple commercial frameworks can only run one realization (a scenario fully specified by one time series of stream flows, evaporation rates and demands, when pertinent, and one value for other uncertainty factors) at a time, which implies in no explicit consideration of uncertainties (see Tables 4.3 and 4.3). Others, most notably the ROA frameworks, include uncertainty in policy optimization in a Bayesian fashion using decision trees and stochastic dynamic programming over expected costs (Fletcher et al., 2017, 2019; Hui et al., 2018). Lastly others consider uncertainty in a stochastic fashion by simulating a policy over hundreds or thousands of realizations every time the model is called, with policy optimization being performed by attaching the model to a black-box optimization algorithm (Zeff et al., 2014; Kwakkel et al., 2014; Watson and Kasprzyk, 2016; Zeff et al., 2016; Borgomeo et al., 2018). WaterPaths was designed based on the latter approach.

WaterPaths allows for both WCU and DU uncertainty sampling schemes (see Figure 3.2 and the broader Section 3.3.3). It currently requires as minimum uncertainty-related input one time series of inflows, evaporation rates, and demands for each piece of infrastructure and utility, when applicable, for each realization to be run. In addition, the user has the option of providing a series of multipliers representing deeply uncertainty factors, which if not provided will assume a default value of the unit. All time-series and multipliers are to be sampled externally from any desired distribution and passed to WaterPaths as input data. The user can add as many deeply uncertain factors as desired when setting up a problem or in new classes of water infrastructure, controls, and drought mitigation policies. The deep uncertainties currently included in WaterPaths are presented in Table 4.3. Simulating and optimizing a policy using a stochastic approach to modeling uncertainty is a mathematically complex and can be computationally costly, so care was taken to allow WaterPaths to both be computationally efficient and to make use of any scale of computational resources available to the analyst.

Type	**Uncertainty Source**	**Regional or Individual for Utility/ Infrastructure**
Supply/Demand	Growth of Mean Annual Demand	Regional
	Growth of Mean Annual Evaporation	Regional
Financial/ economics	Interest Rate	Regional
	Bond Term	Regional
	Discount Rate	Regional
Policy Effectiveness	Water Use Stage Restriction Efficacy	Individual
Infrastructure Construction	Permitting Time	Individual
	Construction Cost	Individual

Table 4.3: Uncertainties included in extended version of the regional water utility planning and management problem presented in Zeff et al. (2016). Regional uncertainties are those for which the same value was used for all utilities, versus a value different value for each utility, as in the case uncertainties described as individual.

4.4.9 Currently Implemented Objective Functions

WaterPaths has implemented by default five objective functions, which can be deleted, modified, and added for specific problem. The first two are the Reliability (f_{REL}), Restriction Frequency (f_{RF}), objectives presented in Section 3.5. However, a new Infrastructure Net Present Cost was added and the Drought Management Cost and Exposure to Financial Risk objectives presented in Section 3.5 were replaced by the Financial Cost and Worse First Percentile Cost below.

- *Infrastructure Net Present Cost* (f_{NPV}): The average net present cost of all new infrastructure build across all realizations:

$$\text{minimize } f_{NPV} = \frac{1}{N_r} \sum_{i=1}^{N_r} \sum_{y=1}^{BM} \frac{PMT}{(1+d)^y} \tag{4.5}$$

  where BM is the bond term, d is the discount rate (5%), y is the year of the debt service payment PMT since the bond was issued, with PMT being calculated as (assuming a level debt service bond):

$$PMT = \frac{P\left[BR(1+BR)^{BM}\right]}{\left[(1+BR)^{BM}-1\right]} \tag{4.6}$$

  where P is the principal (construction cost), BR is the interest rate to be paid to the lender BT is the bond term. The stream of payments is then discounted to present values.

- *Financial Cost* (f_{AC}): The financial cost objective represents the expected yearly cost of all water portfolio assets used to manage droughts over the planning horizon. These costs are revenue losses from restrictions, transfer costs, contingency fund contributions, third-party insurance contract costs, and debt repayment:

$$\text{minimize } f_{AC} = \max_{j} \left[\frac{1}{N_{ys} \cdot N_r} \sum_{i=1}^{N_r} \sum_{y=1}^{N_{ys}} SYC_{i,j}^{y} \right] \tag{4.7}$$

where,

$$SYC_{i,j}^{y} = \frac{\sum_{c \in C_j} PMT_{i,j,c} + \theta_{acfc,j} \cdot ATR_{i,j}^{y} + IP_{i,j}^{y}}{ATR_{i,j}^{y}}$$

where IP is the insurance contract cost in a given year y, $PMT_{i,j,c}$ is the debt payment for infrastructure option c if it belongs to the set C_j of infrastructure options to be built by utility j and is built in realization i, and ATR is the total annual volumetric revenue. All these variables are dollar values.

- *Worse First Percentile Cost* (f_{WFPC}): The worse case cost objective represents the 1% highest single-year drought management costs observed across all analyzed SOWs over the planning horizon:

$$SYC_{i,j}^{y} = \frac{max(RL_{i,j}^{y} + TC_{i,j}^{y} - \theta_{acfc,j} \cdot ATR_{i,j}^{y} - YIPO_{i,j}^{y}, 0)}{ATR_{i,j}^{y}} \tag{4.8}$$

where IP is the insurance contract cost in a given year y, RL is the revenue losses from water use restrictions, TC is the transfer costs, $YIPO$ is the total insurance payout over year y, CF is the available contingency funds, and ATR is the total annual volumetric revenue. All these variables are dollar values. The worse case cost objective is then:

$$\text{minimize } f_{WCC} = \max_{j} \left\{ \underset{i \in N_r}{quantile}(SYC_{i,j}, 0.99) \right\} \tag{4.9}$$

All six objectives are calculated based on the hundreds to thousands of realizations described in Section 4.4.8. Given the high computational demands of calculating such high numbers of realizations, WaterPaths implements concepts of parallel computing to decrease the simulation and policy optimization times. The focus of the remainded of this section is on the parallel structures of simulation and policy optimization with WaterPaths.

4.4.10 Reducing Runtime With Local Parallel Computing

WaterPaths' stochastic design allows for detailed uncertainty analysis but is computationally expensive due to the required high number of realizations and the stochasticity of the ROF calculations. However, given each realization can be simulated independently from all others, they can be easily distributed across multiple computational cores (CPUs) within a processor of a modern desktop, laptop, or cloud computing instance. This parallelization was achieved by using the shared-memory parallelization scheme (several cores working on the same process, meaning sharing the same block memory) implemented by means of OpenMP 5.0 code directives (Klemm et al., 2019). The OpenMP directives allow WaterPaths to create a given number of computational threads and distribute them across available cores. A thread is a set of computing instructions created to execute calculations independently of each other while accessing the same block of memory containing the problem's data. One of these threads is the master thread, which coordinates all worker threads among which it distributes all realizations, retaining some for itself to run, and runs the serial parts of the code (e.g. objectives calculation). WaterPaths creates by default a number of threads equal to the number of cores in the computer, although the user is given the option of manually setting the number of threads (a value of one means that all realizations will be run in serial). Figure 4.4.10a shows a WaterPaths run parallelized with shared-memory parallelization.

Manually setting a number of threads different than the number of cores may have four advantages. First, it allows the user to prevent WaterPaths from using cores already in use by other programs (or other WaterPaths instances) running simultaneously on the same computer, which may potentially significantly slow down all applications and maybe even freeze the computer. The second advantage is that although Hyperthreading[1] often results in performance gains for WaterPaths, it sometimes has the opposite effect, which and can be avoided by manually setting the number of used cores to the number of physical cores (e.g. four in case of a quad-core desktop). Thirdly, running one realization per core in some processors with higher cores counts (e.g. Intel's Knight's Landing series with its 68 cores) may reduce memory-access efficiency and negatively affect performance, in which case the smallest runtimes can be achieved by setting the number of used cores to a number slightly smaller than that of available physical cores. In fact, if running WaterPaths for batch simulations, it is often more efficient to run two to four WaterPaths instances simultaneously, each with fewer threads than the total number or cores and with all instances amounting to two threads per core when Hyperthreading is available. The last reason for manually setting the number of used cores is related to coupling WaterPaths to distributed optimization algorithms, in which multiple function evaluations (one fully stochastic WaterPaths run) are performed simultaneously. This use case is discussed next, following the exposition of a typical policy optimization exercise performed with WaterPaths.

[1]Hyperthreading allows for two threads to run simultaneously in one core by taking advantage of the high number of circuits within a core. It does so by feeding independent instructions from each thread to circuits not being utilized at a time by the other thread, potentially resulting in efficiency gains to the order of 30% although sometimes decreasing performance. The operational system of a computer with Hyperthreading enabled will count each physical core as two virtual cores.

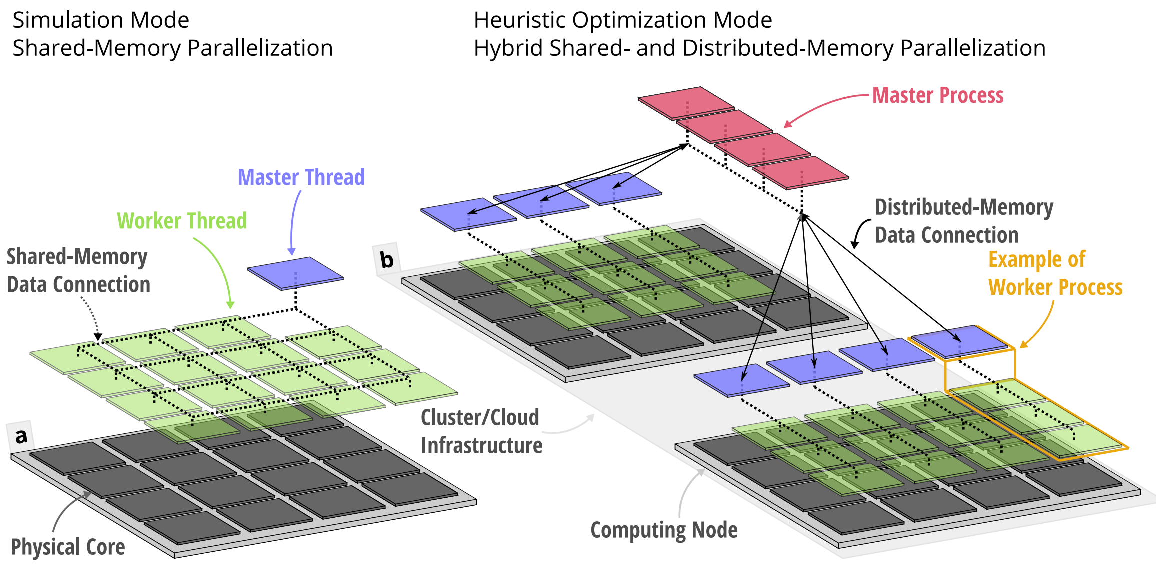

Figure 4.3: Schematics of (a) shared-memory parallelization and (b) hybrid shared- and distributed-memory parallelization for distributed heuristic optmization. Hyperthreading is not used in either panel. The collored and semi-transparent tiles represent threads, the dark solid tiles represent physical cores mounted on desktops or computing nodes (medium grey large and thin tiles serving as base for the cores), and the light grey sheet underneath the two nodes on panel (b) represents a cluster or cloud environment hosting the two depicted nodes. Blue represent the master threads, green the worker threads, and read the threads (one of which still being the master) of a master process. Threads connected by dashed black lines belong to the same process, with its own exclusive memory block, parallelized internally with OpenMP for shared-memory parallelization. Straight, full black two way arrows in panel (b) represent distributed-memory communications between four-core, four-thread master and worker processes, respectively represented by the red threads and the green threads in the yellow contour as an example. Panel (b) shows one master and seven worker processes performing a heuristic optimization run in parallel under the distributed-memory parallelization paradigm.

4.4.11 Facilitating Decision Analytics

Given the diversity and number of policy actions that can be evaluated with WaterPaths, the resulting water infrastructure investment and management pathways are typically composed of mixtures of discrete variables, continuous values, and permutations of infrastructure options. As a consequence, objective measuring system performance across different metrics are non-convex and discontinuous functions of the policy variables. Moreover, the uncertainty sampling enabled by WaterPaths has the potential to make the performance metric functions noisy and highly heterogeneous in their behavior across water supply or financial concerns. All these traits make the problem of manually or automatically designing policies for water infrastructure planning and management particularly difficult and attractive to researchers (Zeff et al., 2016; Fletcher et al., 2019; Borgomeo et al., 2018; Huskova et al., 2016; Kwakkel et al., 2014; Beh et al., 2017).

WaterPaths policy input and objectives output was designed for easy integration with any black-box multiobjective optimization algorithm and currently has out-of-the-box integration with the Master-Worker Borg Multiobjective Optimization Evolutionary Algorithm (MS Borg MOEA) (Hadka and Reed, 2013, 2014), although support other algorithms, written in C++ or other languages, can be easily added. The typical optimization problem formulation solved by WaterPaths coupled with a multiobjective optimization algorithm is shown below in Equations 4.10 through 4.15:

$$\theta^* = \operatorname{argmin}_\theta \mathbf{F} \tag{4.10}$$

s.t.

$$|\text{ME}| \leq 1 \; \forall \; \text{ME} \subseteq \; \text{BI} \tag{4.11}$$

Where:

$$\mathbf{F} = \begin{bmatrix} f_1\left(\mathbf{x}_{\text{srof}},\; \mathbf{\Psi_s},\; \mathbf{x}_{\text{lrof}},\; \theta_1,\; ...,\theta_m,\; \mathbf{OPV}\right) \\ f_2\left(\mathbf{x}_{\text{srof}},\; \mathbf{\Psi_s},\; \mathbf{x}_{\text{lrof}},\; \theta_1,\; ...,\theta_m,\; \mathbf{OPV}\right) \\ \vdots \\ f_n\left(\mathbf{x}_{\text{srof}},\; \mathbf{\Psi_s},\; \mathbf{x}_{\text{lrof}},\; \theta_1,\; ...,\theta_m,\; \mathbf{OPV}\right) \end{bmatrix} \tag{4.12}$$

$$\theta = [\theta_1,\; ...,\theta_m,\; \mathbf{OPV}] \tag{4.13}$$

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_{\text{srof}} \\ \mathbf{x}_{\text{lrof}} \\ \mathbf{x}_{\text{s}} \end{bmatrix} \tag{4.14}$$

$$\mathbf{\Psi}_s = \left\{ \begin{bmatrix} \psi_{\text{inflows, 1}} \\ \psi_{\text{inflows, 2}} \\ \vdots \\ \psi_{\text{inflows, } N_R} \end{bmatrix}, \begin{bmatrix} \psi_{\text{evaporation rates, 1}} \\ \psi_{\text{evaporation rates, 2}} \\ \vdots \\ \psi_{\text{evaporation rates, } N_R} \end{bmatrix}, \begin{bmatrix} \psi_{\text{demands, 1}} \\ \psi_{\text{demands, 2}} \\ \vdots \\ \psi_{\text{demands, } N_R} \end{bmatrix}, \underset{(Optional)}{\begin{bmatrix} \psi_{\text{DU,1}} \\ \psi_{\text{DU,2}} \\ \vdots \\ \psi_{\text{DU,}N_R} \end{bmatrix}} \right\} \tag{4.15}$$

Where **F** is a vector based objective function containing regional objectives f_1 through f_n. The management and investment policies are represented in θ, a vector containing all of the decision variables for all utilities. Decision variables are usually comprised of but not limited to decision ROF triggers θ_1 through θ_m plus other policy variables $\mathbf{OPV}$ such percentages of annual revenue to be paid by insurance in case of a drought. Matrix $\mathbf{X}$ has values of decision-relevant state variables for all utilities and is comprised of $\mathbf{x_{srof}}$, $\mathbf{x_{lrof}}$ and $\mathbf{x_s}$, vectors of short- and long-term values of the ROF metric (or any other custom decision metric), described in Section 4.4.2, and system states, respectively. Matrix $\mathbf{\Psi_s}$ contains vector samples of well-characterized and deeply uncertain time series and parameters. The matrix with samples deeply uncertain factors is an optional parameter. Matrix $\mathbf{X}$ has values of decision-relevant state variables for all utilities and is comprised of $\mathbf{x_{srof}}$, $\mathbf{x_{lrof}}$ and $\mathbf{x_s}$, vectors of short- and long-term values of the ROF metric, described in Section 4.4.2, and system states, respectively. Matrix $\mathbf{\Psi_s}$ contains vector samples of well-characterized and deeply uncertain time series and parameters. Well characterized uncertainties are contained in the vector $\mathbf{\Psi_{WCU}}$ and deep uncertainties are contained in $\mathbf{\Psi_{DU}}$.

4.4.12 WaterPaths Parallelization Strategy for Computational Policy Search

The optimization problem described in Equations 4.10 through 4.15 can be solved with various heuristics techniques that can in theory rely on a master process coordinating multiple worker processes running simultaneously and in charge of performing model evaluations. A process is an instance of the executable comprised of WaterPaths plus the optimization algorithm running on its own separate calculations and enclosed block of memory. After a worker process finishes performing its function evaluation of a policy generated by the master process, it returns the corresponding objective values to the master process. The master process then compares the objective values it just received to those of previously run policies, generates a new policy to be evaluated, and sends it to that worker process for evaluation, repeating the cycle. To allow for massive parallelization across hundreds to thousands of nodes, each with dozens of cores, this parallelization scheme must be developed with a distributed-memory parallelization paradigm, in which the master and each worker process have their own independent blocks of memory. An optimization algorithm written in C/C++ or Fortran for distributed-memory parallelization will likely be implemented with a library based on the MPI standard (Forum, 1994), such as OpenMPI (Gabriel et al., 2004) or Intel MPI (Intel, 2019), although other languages may have their own libraries.

Shared-memory parallelization, described in Section 4.4.10, plays an important role in enabling efficient use of parallel heuristic-optimization algorithms. The naïve parallelization strategy to make maximum use of all available cores would be to have one worker process per computing core, resulting in all cores being used simultaneously, each running one model evaluation. However, using this naïve strategy across tens of cores within a node or a desktop would likely result either on a freeze due to lack of RAM, as each WaterPaths function evaluation may require more RAM than available per core, or in inefficient memory access. To solve this problem, each computing node can have a small number of processes (distributed memory, each process has its own block of memory), each running one model evaluation at a time and distributing realizations across a smaller number of cores. Figure 4.4.10 illustrates this hybrid shared-/distributed-parallelization in action with WaterPaths.

4.5 Methodology

4.5.1 The Sedento Valley: An Illustrative Test Case

To demonstrate the functionality of WaterPaths, we have created a hypothetical test case that reflects many of the challenges facing regions with several independently operated urban water utilities. Furthermore, this test case was developed with other researchers in mind as a benchmark problem against which upcoming methodologies for water infrastructure planning and management can be tested. The test case consists of three utilities in close geographic proximity, facing increasing vulnerability to drought conditions due to climatic and population-based stressors. The close proximity of the three urban utilities in the test case provides an opportunity for regional cooperation and shared infrastructure development, but also creates the potential for inter-utility conflicts over scarce resources. This study demonstrates WaterPaths through a demonstrative implementation of the DU Pathways methodology (Chapter 5) to examine infrastructure investment and drought management strategies for the Sedento Valley. The DU Pathways analysis combines the key WaterPaths capabilities: (1) Monte Carlo evaluation of financial, institutional, and hydro-climatic uncertainties, (2) ROF-based adaptive infrastructure pathways (Zeff et al., 2016), and (3) a mix of short-term water supply portfolio instruments for drought mitigation (Characklis et al., 2006; Zeff et al., 2014).

The Sedento Valley, shown in Figure 4.4, is home to two medium sized cities, Dryville and Fallsland and a smaller municipality, Watertown. The population of the valley is near 1.5 million residents. The three municipalities are each supplied by independent water utility. The cities of Dryville and Fallsland currently receive water from the Autumn Lake reservoir, a large flood control reservoir owned and operated by the US Army Corps of Engineers (USACE). Watertown owns and operates College Rock Reservoir and receives water from Lake Michael, another large USACE operated reservoir. Dryville and Fallsland also have emergency allocations to Lake Michael to which they have access by purchasing water from Watertown's treatment plant and transferring via shared pipelines. Current supply capacities of each water utility can be found in Table 4.4.

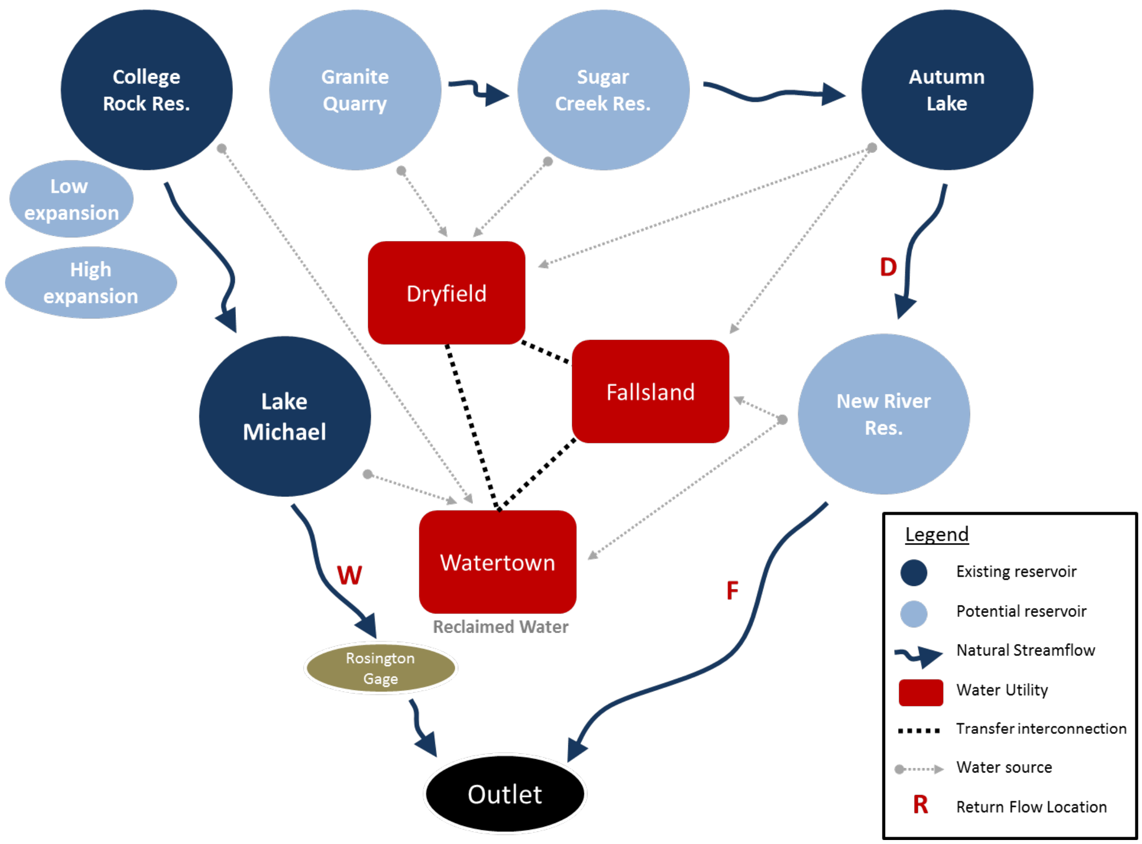

Figure 4.4: The Sedento Valley Regional Water Supply System.

A growing population has reduced each of the utilities' capacity-to-demand ratios, increasing their vulnerability to drought conditions. This problem is

Table 4.4: Current infrastructure capacity of Sedento Valley Water Utilities.

Utility	Infrastructure	Percent Allocation	Supply capacity (MG)
Watertown	College Rock Reservoir	100 %	1049
	Lake Michael	60 %	13,940
Dryville	Autumn Lake Reservoir	38 %	23,839
Fallsland	Autumn Lake Reservoir	62 %	23,839

Table 4.5: Demand to Capacity Ratios for the Sedento Valley Water Utilities

Utility	Baseline Capacity (MG)	2025 Demand (MG/Year)	Demand-Cap Ratio
Watertown	6629	12410	1.87
Dryville	9058.82	12738.5	1.41
Fallsland	14780.18	27813	1.88

compounded by the increased difficulty of large infrastructure investments (i.e., new reservoirs) due to scarcity in feasible sites, higher costs and reduced tolerance of environmental impacts. Increased hydro-climatic variability and uncertainties stemming from land use and land cover changes that effect reservoir inflows pose additional stresses to the region. Currently, the water utilities' drought mitigation strategies rely entirely on the water use restrictions, a measure that is expensive and deeply unpopular with local residents. While the three utilities are each facing increased water stress, their capacity-to-demand ratios and access to new supply options differ as shown in Figure 4.5 and Table 4.5.

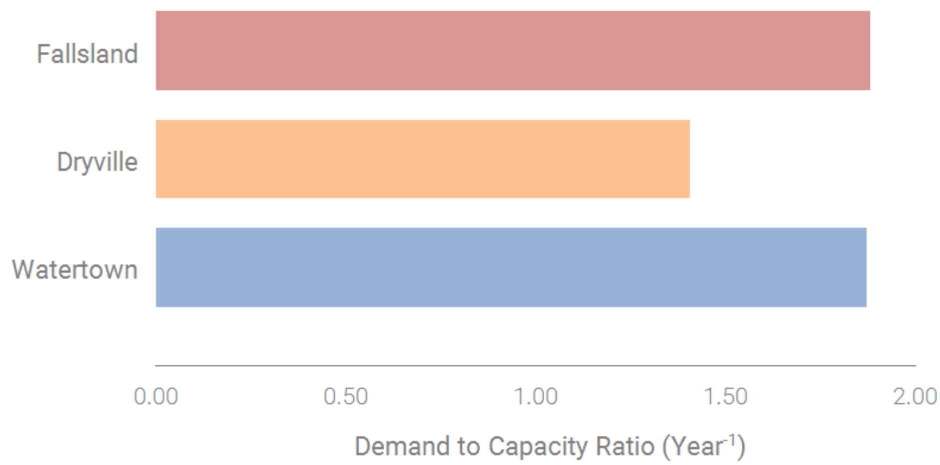

Figure 4.5: Current Demand to Capacity ratios for the three utilities. A higher demand to capacity ratio indicates more stress on the system.

The asymmetry across the utilities capacity-to-demand ratios, their close geographic proximity, and interconnected infrastructure represents an opportunity for cooperative regional water management strategies. There is an interest in cooperative infrastructure investment pathways and coordinated drought mitigation using a regionalized portfolio approach (Zeff et al., 2014) to improve regional and individual performance of water supply systems. Each of the three utilities has identified potential new infrastructure investments to improve their water supply reliability. Watertown has the option to expand college rock reservoir, this can either be a large expansion or a small expansion. Dryville has two small potential water supply options, the development of Granite Quarry into a reservoir and the construction of Sugar Creek Reservoir. Fallsland and Watertown have also been investigating a joint infrastructure investment in the construction of the New River Reservoir. A listing of potential infrastructure improvements can be found in Table 4.6. The three utilities are seeking to find a cooperative strategy that links infrastructure pathways planning with drought

Table 4.6: Potential new infrastructure options in the Sedento Valley.

Infrastructure	Utility (allocation %)	Cost ($MM)	Storage/ production	Permitting Period (years)
College Rock Reservoir expansion (Low)	Watertown	50	500	5
College Rock Reservoir expansion (high)	Watertown	100	1000	5
Watertown Reuse	Watertown	50	35 MGD	5
Granite Quarry	Dryville	22.6	200 MG	17
Sugar Creek Reservoir	Dryville	150	2909 MG	17
Dryville Reuse	Dryville	30	35 MGD	5
New River Reservoir	Fallsland (50%), Watertown (50%)	263	3700 MG	17
Fallsland Reuse	Fallsland	50	35 MGD	5

mitigation policy. As part of this policy, the utilities are interested in incorporating newly available financial tools to hedge against unexpected costs from drought mitigation.

4.5.2 Problem Formulation

The three utilities seek to solve the many objective optimization problem shown in Equations 4.16 through 4.21:

$$\theta^* = \operatorname{argmin}_{\theta} \mathbf{F} \tag{4.16}$$

s.t.

$$|\mathrm{ME}| \leq 1 \; \forall \; \mathrm{ME} \subseteq \; \mathrm{BI} \tag{4.17}$$

Where:

$$\mathbf{F} = \begin{bmatrix} -f_{\mathrm{REL}}\left(\mathbf{x_s},\ \theta_{\mathrm{rt}},\ \theta_{\mathrm{tt}},\ \theta_{\mathrm{lma}},\ \theta_{\mathrm{it}},\ \mathrm{ICO},\ \mathbf{\Psi_s}\right) \\ f_{\mathrm{RF}}\left(\mathbf{x}_{\mathrm{srof}},\ \theta_{\mathrm{rt}},\ \theta_{\mathrm{tt}},\ \theta_{\mathrm{lma}},\ \theta_{\mathrm{it}},\ \mathrm{ICO},\ \mathbf{\Psi_s}\right) \\ f_{\mathrm{NPV}}\left(\mathbf{x}_{\mathrm{lrof}},\ \mathrm{ICO},\ \mathbf{\Psi_s}\right) \\ f_{\mathrm{FC}}\left(\mathbf{x}_{\mathrm{srof}},\ \theta_{\mathrm{rt}},\ \theta_{\mathrm{lma}},\ \theta_{\mathrm{arfc}},\ \theta_{\mathrm{irt}},\ \theta_{\mathrm{it}},\ \mathrm{ICO},\ \mathbf{\Psi_s},\ \mathbf{x}_{\mathrm{lrof}}\right) \\ f_{\mathrm{WFPC}}\left(\mathbf{x}_{\mathrm{srof}},\ \theta_{\mathrm{rt}},\ \theta_{\mathrm{tt}},\ \theta_{\mathrm{lma}},\ \theta_{\mathrm{arfc}},\ \theta_{\mathrm{irt}},\ \theta_{\mathrm{it}},\ \mathrm{ICO},\ \mathbf{\Psi_s},\ \mathbf{x}_{\mathrm{lrof}}\right) \\ f_{\mathrm{lma}}\left(\theta_{\mathrm{lma}}\right) \end{bmatrix} \tag{4.18}$$

$$\theta = [\theta_{\mathrm{rt}},\ \theta_{\mathrm{tt}},\ \theta_{\mathrm{arfc}},\ \theta_{\mathrm{irt}},\ \theta_{\mathrm{it}},\ \mathrm{ICO},\ \theta_{\mathrm{lma}}] \tag{4.19}$$

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_{\mathrm{srof}} \\ \mathbf{x}_{\mathrm{lrof}} \\ \mathbf{x_s} \end{bmatrix} \tag{4.20}$$

$$\mathbf{\Psi}_s = \left\{ \begin{bmatrix} \psi_{\text{inflows, 1}} \\ \psi_{\text{inflows, 2}} \\ \vdots \\ \psi_{\text{inflows, } N_R} \end{bmatrix}, \begin{bmatrix} \psi_{\text{evaporation rates, 1}} \\ \psi_{\text{evaporation rates, 2}} \\ \vdots \\ \psi_{\text{evaporation rates, } N_R} \end{bmatrix}, \begin{bmatrix} \psi_{\text{demands, 1}} \\ \psi_{\text{demands, 2}} \\ \vdots \\ \psi_{\text{demands, } N_R} \end{bmatrix}, \underset{(Optional)}{\begin{bmatrix} \psi_{\text{DU,1}} \\ \psi_{\text{DU,2}} \\ \vdots \\ \psi_{\text{DU,}N_R} \end{bmatrix}} \right\} \tag{4.21}$$

Where $\mathbf{F}$ is a vector based objective function containing regional objectives f_{Rel}, reliability, f_{RF}, restriction frequency, f_{NPV}, net present value of infrastructure investment, f_{FC}, financial cost of drought mitigation and debt payment, f_{WFPC}, worst-first-percentile cost of the f_{FC} and f_{LMA}, Lake Michael allocation (detailed descriptions of the objectives are provided in Section 4.4.9 of the supplement). The management and investment policies are represented in θ, a vector containing all of the decision variables for the three utilities. Decision variables controlling supply stability are $\theta_{\mathbf{rt}}$, a vector of restriction triggers, and $\theta_{\mathbf{tt}}$, a vector of transfer triggers. Decision variable regulating financial stability are $\theta_{\mathbf{arfc}}$, a vector of annual reserve fund contributions, and $\theta_{\mathbf{irt}}$, a vector of insurance restriction triggers. Lastly, decision variables $\theta_{\mathbf{it}}$, a vector of long-term-ROF infrastructure construction triggers, ICO a matrix containing infrastructure construction ordering for each utility and $\theta_{\mathbf{lma}}$, a vector of Lake Michael allocations for the three utilities control infrastructure construction and allocations.

The bounds of the decision variables of the Sedento Valley problem are shown in Tables 4.7 through 4.9. The reader may find strange that values of 1.0 were accepted as the the ROF trigger values for restrictions and transfers for all utilities. The logic behind this choice is that more than providing utilities with precise values of ROF triggers to be implemented by the utilities, the goal of a study such as this one is to provide overarching strategies, in which a value close to 1.0 for restriction trigger would provide a clear signal that other alternatives should be explored besides restrictions so that utilities do not over-rely on it. Lastly, the ϵ-dominance precisions for the objectives of the Sedento Valley problem are shown in Table 4.10. The next subsection will describe how WaterPaths was designed not only for standalone simulations but also to facilitate decision making for regionally-coordinated water utilities.

Table 4.7: Decision variables pertaining to the short-term drought mitigation instruments: water-use restrictions, transfers, contingency fund and drought insurance.

Decision Variable	**Lower Bound**	**Upper Bound**
Dryville restriction ROF trigger	0%	100%
Fallsland restriction ROF trigger	0%	100%
Watertown restriction ROF trigger	0%	100%
Dryville transfer ROF trigger	0%	100%
Fallsland transfer ROF trigger	0%	100%
Fallsland Lake Michael allocation	5%	33.4%
Dryville Lake Michael allocation	5%	33.4%
Watertown Lake Michael allocation	33.4%	90%
Dryville annual contingency fund contribution as percentage of annual revenue	0%	10%
Fallsland annual contingency fund contribution as percentage of annual revenue	0%	10%
Watertown annual contingency fund contribution as percentage of annual revenue	0%	10%
Dryville insurance ROF trigger	0%	100%
Fallsland insurance ROF trigger	0%	100%
Watertown insurance ROF trigger	0%	100%
Dryville insurance payment as percentage of revenue	0%	2%
Fallsland insurance payment as percentage of revenue	0%	2%
Watertown insurance payment as percentage of revenue	0%	2%

Table 4.8: Values of the long-term ROF trigger and daily demands that function as thresholds to trigger infrastructure construction by the utilities.

Decision Variable	**Lower Bound**	**Upper Bound**
Dryville infrastructure construction long-term ROF trigger	0%	100%
Fallsland infrastructure construction long-term ROF trigger	0%	100%
Watertown infrastructure construction long-term ROF trigger	0%	100%

Table 4.9: The ordinal variables below determine the infrastructure construction order and adjusts them for the total number of options available to each utility. The volumetric variable represents the volume of available storage within Falls Lake to be re-allocated from the water quality to Raleigh's municipal supply pool.

Decision Variable	**Lower Bound**	**Upper Bound**
New River Reservoir ranking	1^{st}	8^{th}
Sugar Creek Reservoir ranking	1^{st}	8^{th}
College Rock Expansion Low ranking	1^{st}	8^{th}
College Rock Expansion High ranking	1^{st}	8^{th}
Watertown Reuse I ranking	1^{st}	8^{th}
Watertown Reuse II ranking	1^{st}	8^{th}
Dryville Reuse	1^{st}	8^{th}
Fallsland Reuse	1^{st}	8^{th}

Table 4.10: Values used for ϵ-dominance.

Objective	Reliability	Restriction Frequency	Infrastructure Net Present Cost	Financial Cost	Worse First Percentile Cost
Value	0.1%	2%	$10MM	2.5% of AR*	1% of AR*

4.5.3 Discovering Robust Pathways for the Sedento Valley

We will illustrate the core functionality of WaterPaths through an exercise for the design and analysis of a water-infrastructure planning and management policy. Such exercise consists in an application of the DU Pathways framework for the Sedento Valley example. The DU Pathways framework, as illustrated in Figure 4.5.3, has three core steps besides the problem formulation, all of which forming a loop repeated until a satisfactory solution has been found: (1) identify tradeoffs, (2) evaluating robustness, and (3) infrastructure pathway analytics (more details will be given on Chapter 5). Our application of these steps for the Sedento Valley test is described next.

4.5.4 Identifying Tradeoffs

The Sedento Valley water portfolio management and infrastructure investment pathway demonstration is a challenging high-dimensional (in terms of decisions and objectives) stochastic multiobjective problem. Our demonstration of WaterPaths capabilities and use for discoverying the Sedento Valley test case's tradeoffs exploits the MS Borg MOEA (Hadka et al., 2013). The MS Borg MOEA combines adaptive operator selection (Vrugt and Robinson, 2007) based on probabilities calculated from solutions in its ϵ-dominance archive (Laumanns et al., 2002) described in Equation 2.3 and Figure 2.3. This combination makes it suitable to solve problems with a wide range of mathematical characteristics. The Borg MOEA has demonstrated superior performance over a diverse set of multiobjective problems, such as benchmarking test problems, water supply portfolio planning, pollution control given ecological thresholds, groundwater

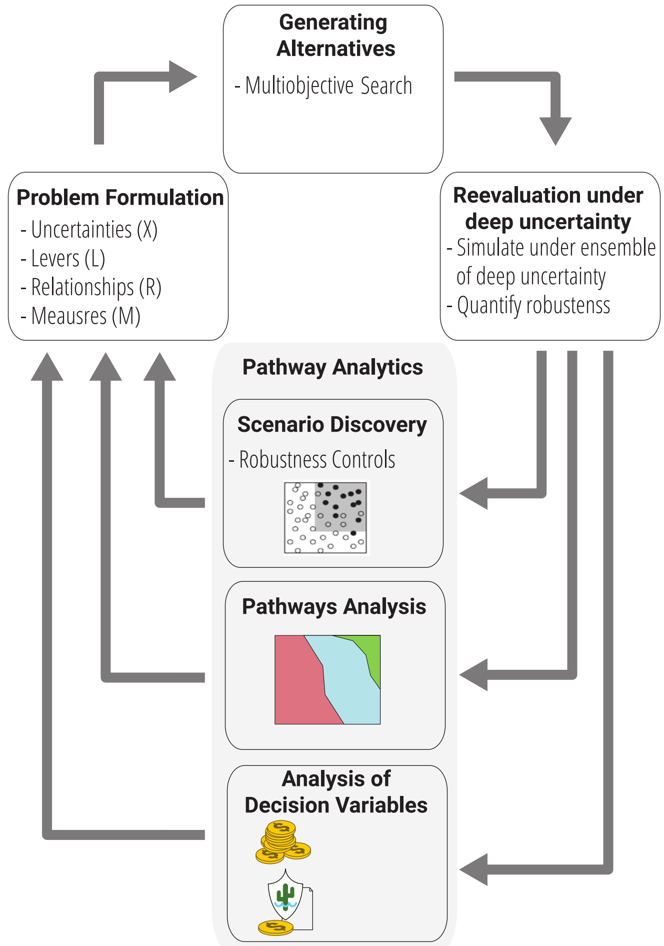

Figure 4.6: Policy design and analysis performed with WaterPaths within the DU Pathways framework.

monitoring design, and reservoir control (Hadka and Reed, 2012b; Reed et al., 2013; Zatarain Salazar et al., 2016; Ward et al., 2015). Furthermore, the MS Borg MOEA has demonstrated that it is capable of solving infrastructure pathway problems Zeff et al. (2016) without the need of parameter tuning, which makes it the ideal choice for engineers not familiar with the technical aspects of evolutionary optimization. The standard values of parameters of the Borg MOEA v1.8 Master-Worker were used in this work (see (Hadka et al., 2013) for the specific values). Although WaterPaths can be readily used with any available modern MOEA or other search tools, the MS-Borg MOEA parallel design facilitates our demonstration of the framework's ability to support state-of-the-art massively parallel analytical capabilities involving input and output of potentially high amounts of data. Figure 4.7 shows the optimization loop and the single-simulation input/output flow.

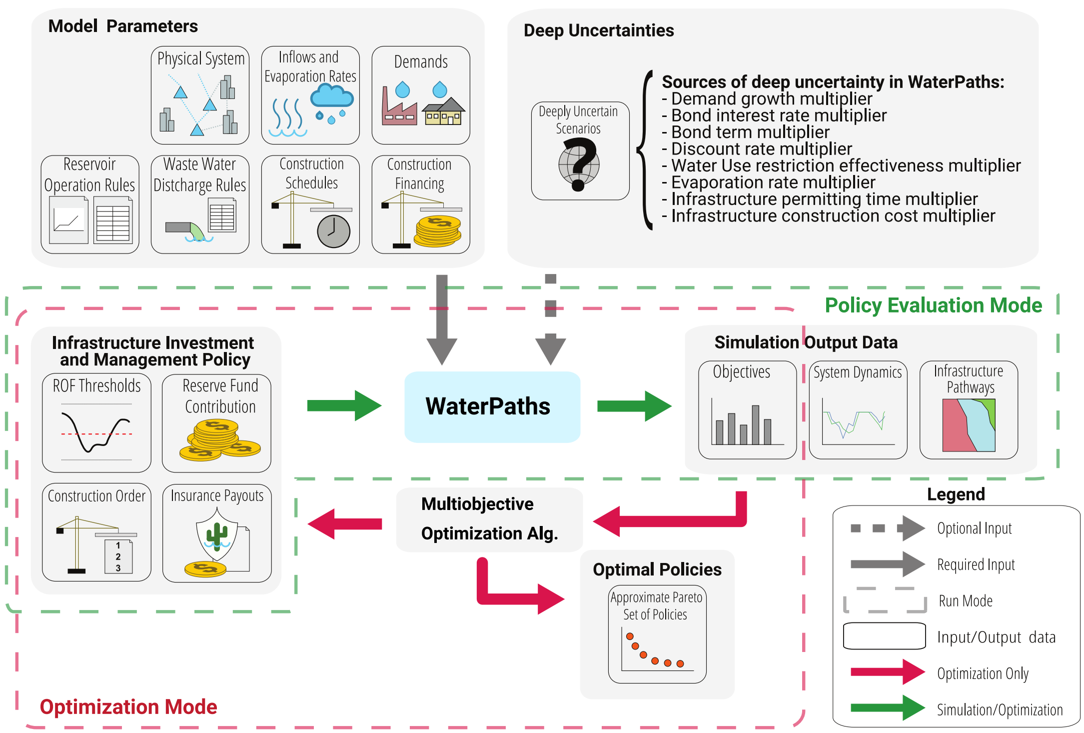

Figure 4.7: Optimization loop and single-simulation input/output information flow. WaterPaths can be used to simulate a single policy, allowing its use for policy re-evaluation and robustness calculation, as demonstrated in Section 4.5.5, and for policy optimization adn tradeoff identification.

For this work, we performed all optimization runs on Texas Advanced Computing Center's (TACC) Stampede 2. The runs were performed on the SKX computing nodes, comprised of two Intel Xeon Platinum 8160 ("Skylake"), with 48 cores @ 2.1 GHz and 192GHz of RAM. We found that nine different optimization seeds ran with MS-Borg MOEA, each with 125,000 function evaluations and a unique number of nodes, was sufficient to converge to the best attainable approximate Pareto set based on the evolution of the hypervolume metric, seen in Figure 4.5.4.

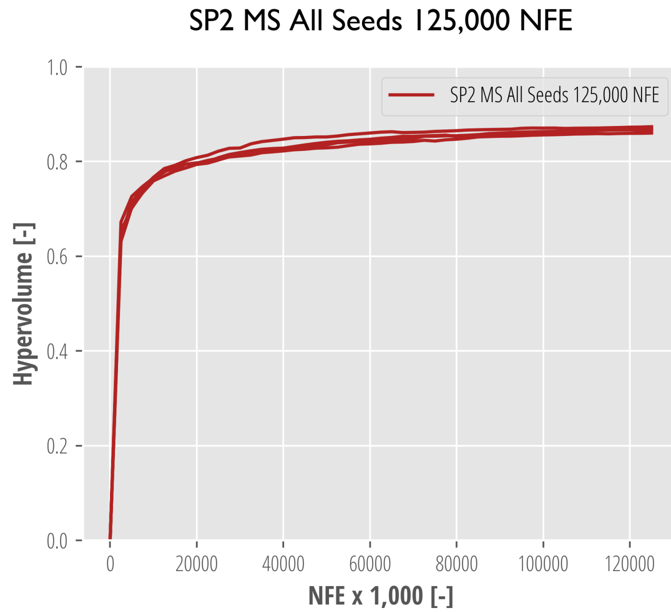

Figure 4.8: Evolution of the hypervolume for each of the nine sets. The overlap across the lines shows that the number of seeds was sufficient to assure consistency, while the gradual plateauing of the hypervolumes of all seeds shows that the number of function evaluations were enough to achieve maximum attainable convergence.

Each function evaluation was performed using the Monte Carlo functionality of WaterPaths by sampling 1,000 realizations. Each realizations fully defines one full draw of candidate values for all of the deeply uncertain factors of concern presented in Table 4.11. Each single draw of these DU values are then coupled to one realization of a 45-year synthetically generated record for streamflows, evaporation rates, and demands. Overall we augment the limits of the

Table 4.11: Deep uncertainties considered for the Sedento Valley test problem. Unless specified otherwise the same minimum and maximum values for each uncertainty were applied for all utilities and infrastructure.

Category	Factor name	Min [-]	Max [-]
Synthetic series trend	Sinusoid amplitude	0.8	1.2
	Sinusoid frequency	0.2	0.5
	Sinusoid phase	$-\pi/2$	$\pi/2$
Utilities	Demand growth multiplier	0.5	2.0
	Bond interest rate multiplier	1.0	1.2
	Bond term multiplier	0.6	1.0
	Discount rate multiplier	0.6	1.4
Drought mitigation instruments (restriction effectiveness multiplier)	Watertown	0.9	1.1
	Dryville	0.9	1.1
	Fallsland	0.9	1.1
New infrastructure	Permitting time multiplier	0.75	1.5
	Construction time multiplier	1.0	1.2

80-year historical record for the Sedento Valley system by creating 1,000 cross-correlated synthetic streamflows and evaporation-rate time series (realizations) for each reservoir and stream gauge. These natural times series are also used to inform the generation of 1,000 synthetic demand time-series for each utility, which account for correlations with changing levels of water scarcity for each utility. The simulation of all 1,000 fully specified realizations and subsequent calculation of the specified objectives constitutes one function evaluation. WaterPaths' software design permits significant flexibility for stochastic simulation as well as stochastic simulation-optimization.

The resulting output Pareto approximation sets from each of the MS Borg MOEA random seed trials on TACC Stampede 2 were combined and sorted using ϵ-dominance sorting to build a best-known reference set for the Sedento Valley demonstration problem. The use of ϵ-dominance provides a convenient means of preserving high-quality representations of key tradeoffs sets while reducing solution sets to a user preferred size to improve interpretability of tradeoff assessments and/or limit the computational demands of the robustness assessments described below (readers interested in more details can reference Reed et al. 2013; Kollat and Reed 2007). For the Sedento Valley example presented here, we filtered the reference set of solutions using two steps: (1) ϵ-dominance sorting of all solutions based with our specified ϵ values and (2) filtering solutions based onperformance goals for specific objectives.

4.5.5 Evaluating Robustness

The next step was to further stress test solutions in the final reference set, generated by ϵ-dominance sorting all solutions based on specified ϵ values, by re-evaluating them against a larger independent sampling of 5,000 vectors of the DU factors presented in Table 4.11. For each of the 5,000 re-evaluation runs, we created one set of 1,000 streamflows, evaporation rate, and demand time-series based on the first three deeply uncertain factor in Table 4.11 time series for each reservoir and gauge, resulting in 5,000,000 SOWs. The processes for the generation of inflows, evaporation rate, and demand time series is presented in the supplemental Section B.

Before proceeding with such computationally expensive re-evaluation and robustness evaluation, the initial reference set resulting in 830 non-dominated solutions was reduced in size to 229 solutions before being re-evaluated over 5,000 re-evaluation scenarios for computational tractability, as mentioned in Section 4.5.4. This was achieved by keeping in the reference set only solutions whose objective values as calculated during optimization displayed reliability greater than 98%, restriction frequency on less than 30% of the years, and worse first percentile cost smaller than %10 of the annual revenue. These criteria were chosen for (1) being similar to minimum performance standards deemed acceptable based on the authors' prior experience with real utilities, and (2) for reducing the number of solutions limit computational demands.

The resulting $229 \cdot 5,000 = 1,145,000$ simulations required for re-evaluation were split into 25 independent jobs on Stampede 2. Each job consisted in one Python script using MPI4Py (distributed memory parallelization) to use distribute blocks of (229 solutions) · (200 RDM scenarios) = (22,900 function evaluations) across 50 nodes, resulting in (229 solutions) · (4 RDM scenarios) = 916 function evaluations per node per job for each of the 25 jobs. The reason for the split was to improve the fault tolerance of the re-evaluation exercise by preventing a crashing function evaluation from crashing the entire re-evaluation exercise, as a crash in one function evaluation would crash the Python script and in turn all its function evaluations. Each of the 50 nodes ran one function evaluation at a time distributing its 1,000 realizations across all 48 cores with 2 realizations running simultaneously on each core.

The objective values for each of the 5,000 re-evaluation function evaluations for each of the policies in the abbreviated combined reference set were used to re-calculate the policies' objectives, this time based on $1,000 \cdot 5,000 = 5,000,000$

realizations rather than the initial 1,000 using during the MOEA search phase. Additionally, the 5,000 sets of objective values obtained for each policy were used to calculate the value of the satisficing robustness metric (Starr, 1962; Herman et al., 2014) for each utility under each policy, defined as the percentage of the 5,000 sets that met the performance criteria defined by the utilities.In short, the satisficing measure of robustness measures the percent of sample worlds where decision makers deem performance acceptable based specified performance requirements.

4.5.6 Pathway Analytics

Based on the robustness values, a solution was selected for a detail scenario discovery analysis (Bryant and Lempert 2010 and Section 3.3.5) performed with the Boosted Trees algorithm (Schapire, 1999), as further discussed in Section 5.4.3 and Appendix E. Scenario discovery is an effort to determine the uncertainties (factor prioritization) and/or combinations of their values (factor mapping) that most closely relate to performance and robustness. This information is presented in terms of a map of the space of uncertainty showing which combinations of values for uncertain factors would likely cause a policy to fail. Included in the axes of the maps is the uncertain factors that are most influential in that policy's performance, with their relative performances displayed in percentage metric analog to explained performance variance (see Section 5.4.3 and Appendix E for details). The Boosted Trees was used for scenario discovery for the Sedento Valley test case because of two advantages it holds over the other more easily interpretable but limited scenario discovery methods such as PRIM (see Section 3.3.5, Figure 3.8, and Bryant and Lempert 2010; Friedman and Fisher 1999), CART (Bryant and Lempert, 2010; Breiman et al., 1984), and logistic regression (Quinn et al., 2018; Gold et al., 2019): (1) Boosted Trees captures non-differentiable boundaries typical from threshold-based rules such as WaterPaths' ROF-based action triggers, (2) it captures non-linear boundaries without explicitly modeling variable interactions while being resistant to overfitting, helping assure scenario discovery maps that are as simple as possible to interpret.

To obtain the maps and factor priorities, a Python script provided with WaterPaths' source code reads all samples 5,000 of deeply uncertain factors and corresponding objectives, and for each solution fits a Boosted Trees classifier to the uncertain factors and objective values for each utility for each policy. The surface is then presented in maps indicating the regions of the space of uncertainties in which that policy is likely to fail for each utility, as well as the sources of uncertainty that most impact the performance of each utility under that policy. The script for fitting the Boosted Tree classifier and the plotting the maps is based on the Scikit-Learn machine-learning library (Pedregosa et al., 2011) and provided in the repository listed in Section 4.2.

Lastly, all infrastructure-pathway data output for each of the 5,000 function evaluations performed for the policy being analyzed is used to create plots representing infrastructure construction as a function of time. The data is comprised of the realization ID, the utility that triggered a project, the infrastructure option that was built, and the week in which the option became operational. Three pathway plots are created for each utility under the policy being analyzed: one presenting infrastructure construction over time for all realizations under the least favorable sample of deeply uncertain factors, one for the projected values for each deeply uncertain factor, and one for the most favorable. These plots can be used to understand which projects are likely to be suitable for construction in the near-term and how dependent on uncertainty the need for construction is. Next we provide details about the parallel architectures and configurations we exploited to provide an assessment of WaterPaths parallel performance.

4.5.7 Scaling Performance on Clusters and in the Cloud

In this comparative assessment, replicate optimization random seed trial runs were run on small and large traditional high-performance computing clusters as well as a virtual cloud cluster. In the scaling analysis, all WaterPaths model runs (or function evaluations) were distributed across computing nodes on Stampede 2 with Intel MPI for Linux v2019u5 (distributed memory parallelization library) (Intel, 2019) and on the Cornell Red Cloud and on the research cluster using OpenMPI v3.1.4 (Gabriel et al., 2004), while the sampled realizations that composed the function evaluations were distributed across cores within a node using OpenMP v5.0 (shared memory parallelization library) (Stallman and Community, 2017; Klemm et al., 2019). This hybrid parallelization scheme as discussed in Section 4.4.11 and in Figure 4.4.10b has the advantage of allowing the user to scale optimization runs up to as many computing nodes as are available while avoiding core idle time due to memory access and size limitations within a node. The within-node shared memory parallelization of realizations within a function evaluation also has the advantage of allowing the user to run more realizations simultaneously. This feature is particularly useful as the number of cores per processor is increasing and as cloud computing with nodes of various sizes are becoming more readily available.

In addition to optimization runs performed on Stampede 2 described above in Section 4.5.4, performance was assessed on a research cluster termed The Cube, whose nodes are comprised of two Intel Xeon E5-2680 (Sandy Bridge), @ 2.7 GHz and 128 GB of RAM (see Table 4.12). The cloud results were attained on Cornell's Red Cloud with nodes with 28 cores and variable architecture. The scaling analysis consists of two types of tests to evaluate search performance on each platform: search scalability and cross-platform wall clock search time comparisons. Our goal for the multi-architecture comparison tests is to benchmark how performance varies for typical small clusters, elite class high performance computing systems, and emerging cloud platforms. Our scalability benchmarking focuses on efficiency as a function of the number of nodes/cores as defined in equation 4.22 derived from Amdahl's law assuming negligible, necessarily-serial work. Our core demonstration results are based on nine random seed search trials on Stampede 2, where the seed trial count allowed our experiment to exploit the maximum allowed by the system administrators. The Cube and Red Cloud analyses are based on three random seed trial runs for the sole purpose of informing this computational scalability analysis. Each seed had one MS Borg MOEA master process and a number of worker processes based on the number of nodes and of processes per node.

$$E \approx \frac{N_0 \cdot T_{N_0}^{arch}}{N \cdot T_{N-nodes}^{arch}} \tag{4.22}$$

In equation 4.22 E is parallel efficiency, $T_{N_0}^{arch}$ is the total execution time on the base case (here, eight nodes) for a given architecture and N is the number of nodes to be compared against the base case. Equation 4.22 assumes the vast majority of the execution time is spent on tasks run in parallel. Table 4.12 shows which computer configurations were used for each test. The configurations of each seed can be found in Table 4.12 together with details about the architecture of the compute nodes used. Table 4.12 also shows that the exercise of optimizing a model built with WaterPaths using MS-Borg MOEA is nearly perfectly scalable when running function evaluations in parallel for optimization on either low or high numbers of nodes in high performance computing applications. These results highlight that when using WaterPaths that investments in high performance computing power will yield returns linearly proportional to the number of nodes, which corresponds to ideal scalability.

Table 4.12: Architectures and configurations used for optimization scalability and cross-platform performance tests, followed by total run times and efficiencies for each optimization seed, ran on Stampede 2 and on the cluster The Cube with different parallel configurations.

Platform	Node Type	# of Nodes	# of Cores/ Node	Total # of Cores	Processes per Node	Times [min]	Efficiency [-]
Stampede 2	2 x Intel Xeon Platinum 8160 (SkyLake), 2.1 GHz, 192 GB of RAM	8	48	384	4	1720	—
		16		768		882	98%
		32		1536		425	100%
		40		1920		337	100%
		48		2304		289	99%
		56		2688		245	100%
		64		3072		219	98%
		96		4608		143	100%
		128		6144		109	99%
Aristotle Cloud	Haswell (variable), 236 GB of RAM	8	28	224	2	3665	—
The Cube (Cornell Research Custer)	2 x Intel Xeon E5-2680 (Sandy Bridge), 2.7 GHz, 128 GB of RAM	8	16	128	3	8274	—
		16		256		4177	100%
		32		512		2020	98%

In addition to the results above, we also analyzed performance losses from distributing realizations of a single function evaluation across multiple cores, as opposed to the previously described scalability experiment which distributed multiple function evaluations across nodes for optimization. These tests considered 1, 2, and 4 cores within a single Intel Xeon CPU E5-2680 0 @ 2.70GHz standard desktop workstation and across 2, 4, 8, 16, 32, and 48 cores of a single Stampede 2 Skylake node. This scalability study is aimed at users running single Monte Carlo simulations or ensembles of Monte Carlo simulations in batch mode, as opposed to MOEA search runs. Table 4.13 shows the results of timing

WaterPaths running 1,000 realizations on different number of cores on a four-core Xeon desktop and on Stampede 2.

Table 4.13: Timings and efficiency of running WaterPaths with 1,000 realizations split among different numbers of cores on a 4-core Xeon Desktop and on a 48-core SkyLake node of Stampede 2.

—	Xeon Desktop		Stampede 2 Skylake	
Number of Cores	Time [s]	Efficiency	Time [s]	Efficiency
1	176	—	—	—
2	88	100%	128	—
4	44	100%	103	62%
8	—	—	63	51%
16	—	—	37	43%
32	—	—	20	40%
48	—	—	15	36%

Table 4.13 shows that WaterPaths also scales its realizations across cores on the Xeon desktop, which has a processor with few a few strong cores, at 100% efficiency. The same, however, does not happen on Stampede 2 Skylake nodes, which is a processor with several weaker cores, that become progressively more impacted with memory access issues with core count (i.e., less active memory per core for computations). Despite the losses in efficiency, a user can still achieve a 10-fold runtime reduction using all available cores, which is a much desirable feature to have at production scale.

4.6 Results

4.6.1 Performance and Robustness Tradeoffs

Figure 4.9 shows the objectives (panel "a") and robustness (panel "b") tradeoffs among Sedento Valley utilities following re-evaluation as output by WaterPaths. In Figure 4.9a, each axis represents an objective and each line represents a policy. The location where each line (policy) intersects a vertical axes represents it performance value in the corresponding objective. The highlighted policies represent the policies with best robustness for each utility (BRo-X), with the best robustness compromise across utilities (RoC), and best performance compromise (PC). The axes in Figure 4.9a are oriented such that the ideal policy would be a horizontal line at the bottom of all axes.

Figure 4.9a highlights moderately strong tensions between reliability and restriction frequency, which is expected given restrictions are a supply-reliability instrument. Likewise, significant tradeoffs exist between Infrastructure NPV and Financial Cost for higher reliabilities, indicating that the construction of infrastructure successfully offsets the need (and therefore the cost) of costly drought mitigation and financial instruments. Figure 4.9a also shows that the low-reliability policies tend to generally have between 5% and 15% of restriction frequency and that policies with low restriction use rank high in infrastructure NPV, further suggesting that building the suggested infrastructure successfully offsets the use of drought mitigation instruments. However, despite the general trends, a small group of policies with the highest reliability still make use of moderate restrictions and investments in new infrastructure. Such conclusions would be difficult to reach with existing frameworks and were made possible

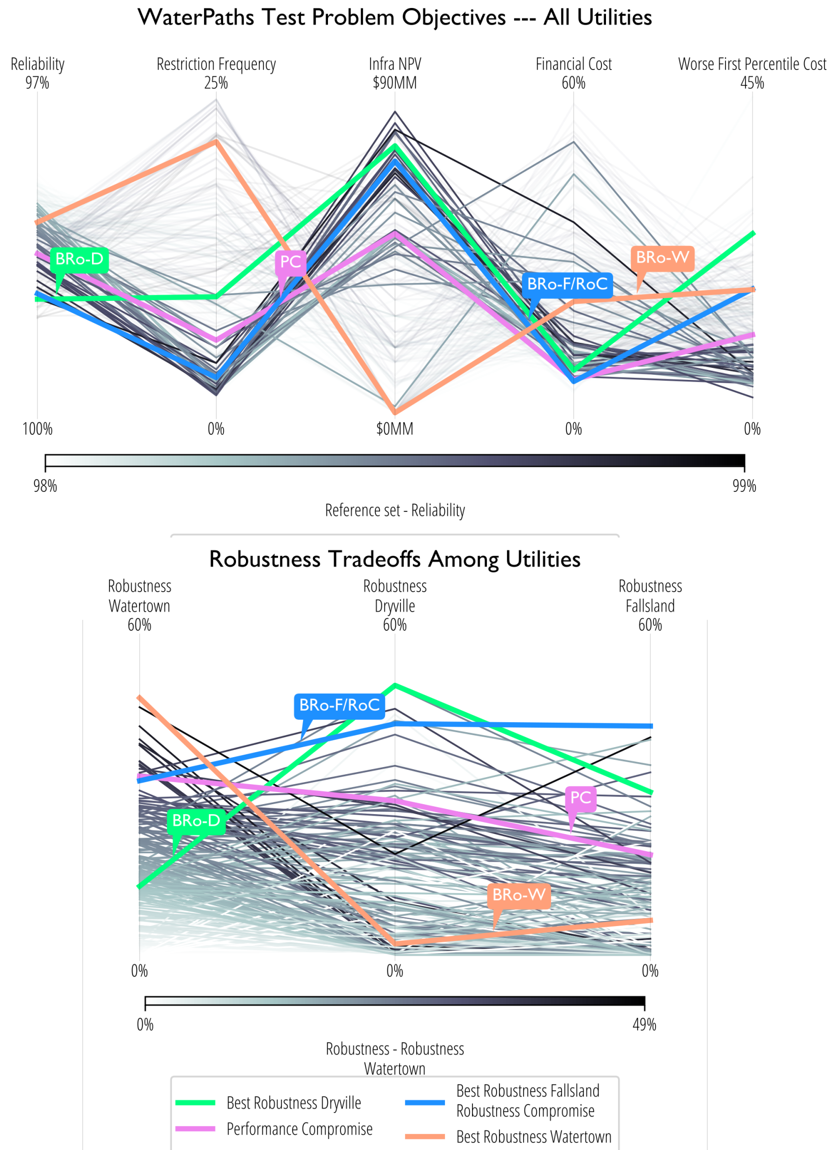

Figure 4.9: Performance and robustness tradeoffs in the DU Re-Evaluation space. Panel (a) illustrates the tradeoffs across the performance objectives quantified either using the mean or the worse first percentile across the tested SOWs, in which an ideal solution would be represented by a horizontal line at the bottom of the plot. Panel (b) illustrates the robustness tradeoffs across the three utilities. Here, robustness is calculated as the percent of RDM SOWs in which a policy met the performance criteria defined by the utilities and an ideal solution would be represented by a horizontal line at the top of panel (b). The highlighted policies represent the policies with best robustness for each utility (BRO-X), with the best robustness compromise across utilities (RoC), and best performance compromise (PC).

by WaterPaths' unique combination of supply and financial models, policies implemented such that their ideal operation is discovered instead of prespacified, and possibility of infrastructure construction during the course of a simulation.

Transitioning to robustness tradeoffs, Figure 4.9b uses a similar parallel axes plot to display the percentages of SOWs where each utility meets their goal performance requirements: reliability $> 98\%$, restriction frequency $< 10\%$ and annual worst-first-percentile cost $< 10\%$. In Figure 4.9b, the direction of preference is upward, so that the ideal policy would be represented by a horizontal line at the top of the axes (i.e., 100% for all utilities). The most robust policies for each utility and the best robustness and objectives-performance compromises across all utilities are again highlighted in different colors in Figure 4.9b. The highlighted best-robustness policies in Figure 4.9b suggest inter-utility robustness tradeoffs for the highest levels of attained robustness, which indicate complex interdependencies between the utilities caused by utilities attempting to simultaneously use constrained regional resources. Furthermore, performance compromise solution shows that focusing on performance levels may degrade robustness for all utilities, which would likely go against the utilities' risk aversion. The capability of simulating resources shared by utilities to allow for regional conflict analysis is increasingly important given how the distance between areas serviced by water utilities is decreasing, and is a distinct capability of WaterPaths.

Figure 4.10 supplements Figure 4.9b by comparing utilities' robustness profiles across all of the policies found through optimization. One important result highlighted in Figures 4.10 and 4.9b is the strong source of potential regional tension depending on if the utilities seek to focus on overall performance trade-

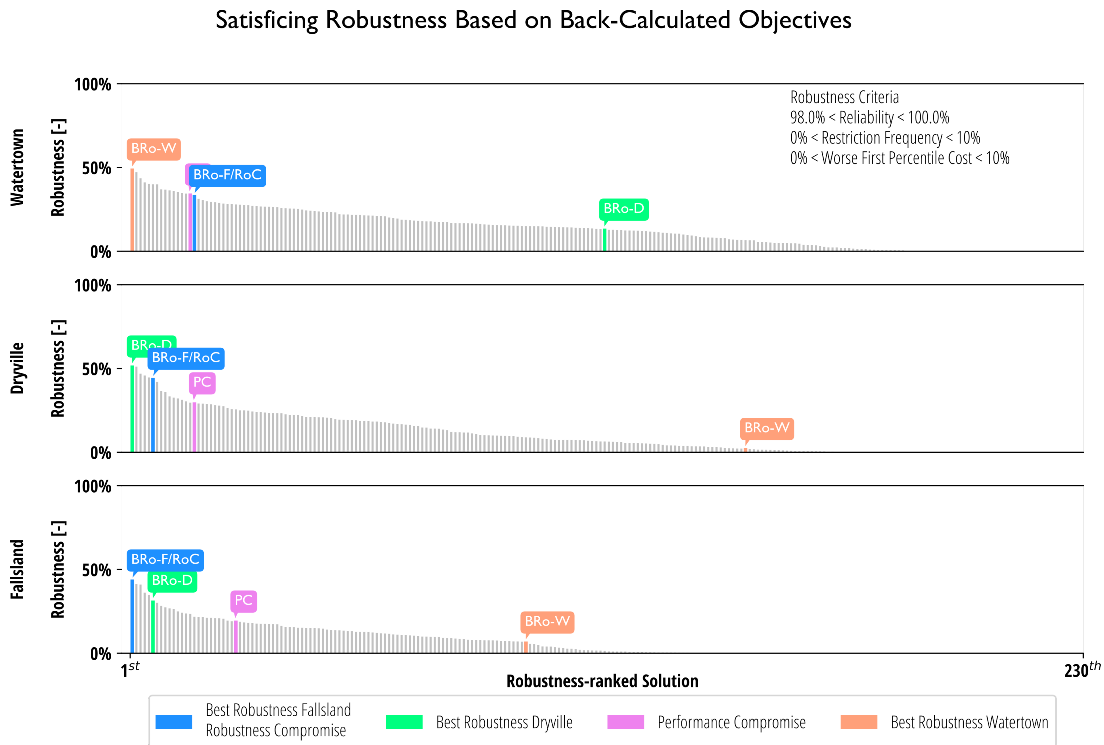

Figure 4.10: Robustness of pre-selected 229 policies for each utility. Each bar represents the attained robustness for one policy for each utility. The policies are sorted by robustness (bar height) for each utility. The policies with the highest value of the satisficing robustness for each utility, the policy with the best robustness compromise, and the policy with the best compromise of objectives values are highlighted in each subplot.

offs versus individual robustness tradeoffs. For Dryville and Fallsland, the illustrated distances between the blue (robustness compromise) and the fuchsia (performance compromise) solutions in both figures implies that seeking a regional management and investment policy compromise for performance tradeoffs could make Dryville and Fallsland more vulnerable individually. Figure 4.10 highlights that there are very few solutions that are simultaneously close to the highest attainable robustness for each individual utility, as illustrated by the steep downward slope in bar heights on the leftmost region of all three bar charts, which is concern that should be considered explicitly in any regional water portfolio management and infrastructure investment policy negotiation. WaterPaths provides detailed analyses that can be helpful for these types of regional decision-making contexts. The simulation framework allows stakeholders to explore how the component policy decision variables themselves are driving these tradeoffs.

4.6.2 Policy Rules in the Robustness Compromises

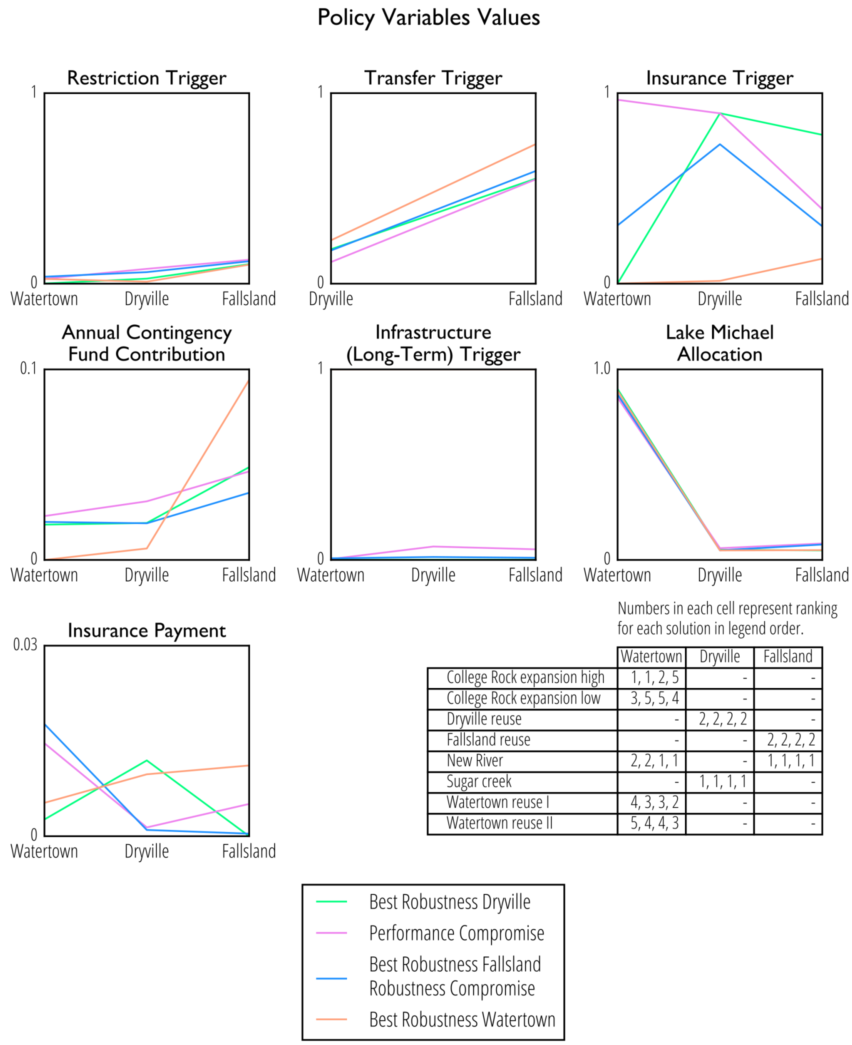

	Watertown	Dryville	Fallsland
College Rock expansion high	1, 1, 2, 5	-	-
College Rock expansion low	3, 5, 5, 4	-	-
Dryville reuse	-	2, 2, 2, 2	-
Fallsland reuse	-	-	2, 2, 2, 2
New River	2, 2, 1, 1	-	1, 1, 1, 1
Sugar creek	-	1, 1, 1, 1	-
Watertown reuse I	4, 3, 3, 2	-	-
Watertown reuse II	5, 4, 4, 3	-	-

Figure 4.11: Parallel axes plots for the policy variables for the policies representing the best robustness for each utility, the best robustness compromise, and the best performance compromise. The table indicates the infrastructure construction order for each utility for each policy in the order shown in the legend.

Figure 4.11 detailed illustration of the rules that compose the policies highlighted in Figures 4.9 and 4.10. Recall that each policy is comprised of: (1) ROF triggers for short-term mitigation instruments [Figures 4.11a, 4.11b, and 4.11c] and for long-term infrastructure investments [Figure 4.11e], (2) allocations to the water supply pool in Lake Michael allocation [Figure 4.11f], (3) annual contingency fund contributions and insurance payouts defined as percentages of annual revenue [Figures 4.11d and 4.11g], and (4) infrastructure construction order [table within Figure 4.11]. The first striking feature of Figure 4.11 is the similarities between all highlighted solutions regarding the use of supply-reliability tools, namely restrictions, transfers, infrastructure construction, Lake Michael Allocation, and infrastructure construction order, the latter for Dryville and Fallsland. In all policies, all utilities make extensive use of restrictions due to its low trigger value, build considerable infrastructure, and almost the entirety of the Lake Michael is allocated to Watertown. In addition, Dryville also has high reliance on transfers from Watertown. In comparing Figure 4.11b to Figure 4.10, it is possible to notice a possible relation between transfer obligations for Watertown and its robustness, as the most robust policy for Watertown implies less treated water transfers and robustness for the other two utilities. The variability in all utilities' financial strategies across solutions, however, requires better scrutiny. All policies except for Watertown's most robust have values between 2.5% and 5% for the contingency fund variable, suggesting this is an instrument to be explored by all utilities. In addition, Dryville insurance strategy seems to gravitate around infrequent use of insurance with low payments, which effectively means insurance is not of use for Dryville. Watertown's and Fallsland's financial strategies varied across policies with no discernible pattern.

Figure 4.11 provides a detailed illustration of the rules that compose the water portfolio management and infrastructure investment policies highlighted in Figures 4.9 and 4.10. Recall that each policy is comprised of: (1) ROF triggers for short-term mitigation instruments [Figures 4.11a, 4.11b, and 4.11c] and for long-term infrastructure investments [Figure 4.11e], (2) Lake Michael allocations [Figure 4.11f], (3) annual contingency fund contributions and insurance payouts defined as percentages of annual revenue [Figures 4.11d and 4.11g], and (4) infrastructure construction order [table within Figure 4.11]. The first striking feature of Figure 4.11 are the similarities between all of the highlighted solutions regarding the use of supply focused instruments that are strongly influential for reliability, namely restrictions, transfers, infrastructure construction, Lake Michael Allocation, and infrastructure construction order, the latter for Dryville and Fallsland. In all policies, all utilities make extensive use of restrictions due to its low trigger value, build considerable infrastructure, and almost the entirety of the Lake Michael is allocated to Watertown. In addition, Dryville also has high reliance on transfers from Watertown. Comparison of Figure 4.11b to Figure 4.10 highlights a possible relation between transfer obligations for Watertown and its robustness, as the most robust policy for Watertown implies less treated water transfers [higher triggers in Figure 4.11] and robustness for the other two utilities [smaller robustness for Dryville and Fallsland in Figure 4.10]. Transitioning to the utilities use of financial instrument, there are significant differences in their financial decisions across the highlighted solutions. Figure 4.10 shows that all of the highlighted policies except for Watertown's most robust have values between 2.5% and 5% for the contingency fund variable, suggesting this is an instrument to be explored by all utilities. In addition, Figure 4.11 shows that Dryville's insurance strategy tends to gravitate towards either unusually infrequent use of insurance and/or low payments, both effectively meaning insurance is less useful for Dryville. Watertown's and Fallsland's financial strategies, on the other hand, varied across policies depending their focus on performance, robustness, and alternative compromises.

A key feature of WaterPaths is providing a self-consistent means of simulating and evaluating how short-term management actions shape sequences of infrastructure investments. In Figure 4.11 the ROF-based management and investment action triggers interact with the ordering or prioritization of infrastructure sequences for each of the utilities. As Figure 4.11 shows, the infrastructure construction order is stable for Dryville and Fallsland across solutions but varies for Watertown. In all policies, Watertown builds early on the New River reservoir, which is shared with Fallsland and for whom it is also the first option to be built, indicating the importance of this joint project. Also, all utilities in all solutions prioritize the construction of added storage as opposed to water reuse stations except for Watertown in its most robust policy, which indicates high storage's higher efficacy in improving supply and financial performance. The inversion in priority between expanding the College Rock Reservoir and building reuse stations for Watertown in Watertown's most robust policy may also been connected to its high robustness in this policy and to the decrease of Fallsland robustness, given the College Rock Reservoir helps stabilize the flows into the New River Reservoir. The similarities in the use of reliability-focused instruments across utilities also suggests that their differences in robustness may be substantially dependent on the utilities' differences in their use of their financial instruments. This leaves room for fine-tuning of policies using WaterPaths, given that financial instruments are part of individual (as opposed to regional) planning and that the during optimization the financial instruments were optimized only for the worse performing utility for each objective. With an understanding of the decisions being made by the utilities under each policy, the next step is to understand which uncertainty factors and their values drive robustness.

4.6.3 Scenario Discovery for Compromise Policies

WaterPaths as a flexible Monte Carlo simulation framework enables careful exploratory modeling to better understand what deeply uncertain factors most strongly influence robustness. Each map in Figure 4.12 shows the performance attained across scenarios for the two most important uncertainties for each of the three utilities when the Sedento Valley implements the policy with the best robustness compromise, highlighted in Figures 4.9 and 4.10. In each panel of Figure 4.12, the percentages shown in the axes denote the importance of each source of uncertainty in determining whether the mentioned policy would or would not meet the utilities performance criteria in the samples scenarios, as identified by the boosted trees algorithm as mentioned in Section 4.5.3. More specifically, the percentages represent the decrease in impurity of the tree ensemble from splits on that factor, with a higher percentage indicating a higher importance of that factor. The red dots denote scenarios in which the utility fails to meet its performance targets and gray dots represent scenarios a utility performs as or better than required. Red shading represents regions of the uncertainty space where a utility under the mentioned policy will likely not meet the performance criteria with P(meeting criteria|scenario) ≤ 0.5, white regions denote inconclusive regions in which $0.98 \geq$ P(meeting criteria|scenario) ≥ 0.5 (inconclusive regions), and gray regions represent scenario regions in which the utility will likely meet the performance criteria with P(meeting criteria|scenario)

≥ 0.98. All probabilities were calculated using Boosted Trees with 500 trees of depth 4. In each panel of Figure 4.12, the stars represent the most favorable scenario (Gray star), the least favorable (red), and the future according to projections (blue star).

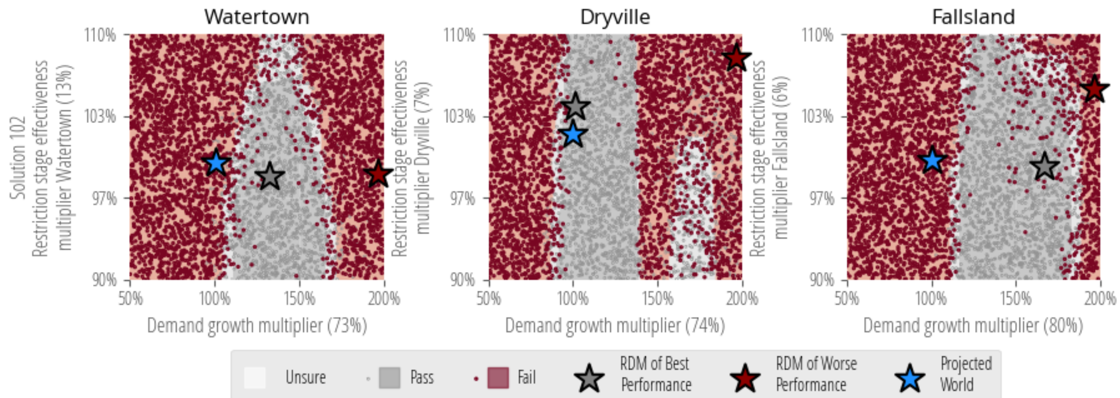

Figure 4.12: Factor maps for the robustness compromise policy (RoC) showing the key uncertainty factors controlling the robustness the policy for all utilities. Success (grey dots and zone) required reliability $\geq 98\%$, restriction frequency $\leq 10\%$ and annual financial worst 1st percentile cost $\leq 10\%$, while a failure to meet either criteria results in the scenario being considered a failure (red dots and region). White zones denote regions with a mixture of gray and red dots for which it is uncertain whether all criteria would be met.

Figure 4.12 shows that all of the utilities' robustness are mostly contingent on demand growth, but that contrary to expectations both low and high values of demand growth may cause the three utilities to fail. However, low- and high-demand failures happen for very different reasons. Failures due to low demand result for all three utilities when low long-term ROF values do not trigger infrastructure construction leading to over-dependence on financially challenging restrictions (i.e., high restriction frequencies and volatile cost swings). The low-demand failures result when restrictions and short-term transfers are not sufficient for preventing supply failures, causing low reliabilities and high restriction frequencies. In low-demand scenarios, all of the utilities struggle with financial stability and fail to meet the robustness criterion for the worse first percentile cost objective, which is worse first percentile cost smaller $\leq 10\%$.

Overall, Figure 4.12 provides some important regional insights for the Sedento Valley. The dependence on failures in all three objectives on demand growth rates indicate an intimate relationship between supply reliability and financial stability. Overall, keeping demand growth rates at 20% above the utilities' projected value would be an important step for all of the utilities to ensure satisfactory regional performance while avoiding stranded assets, which requires high demands because of their high costs. The next section will explore how the scenarios highlighted above by the stars drive infrastructure construction, as a way of understanding the relationship between infrastructure construction and robustness, performance, uncertainty factors.

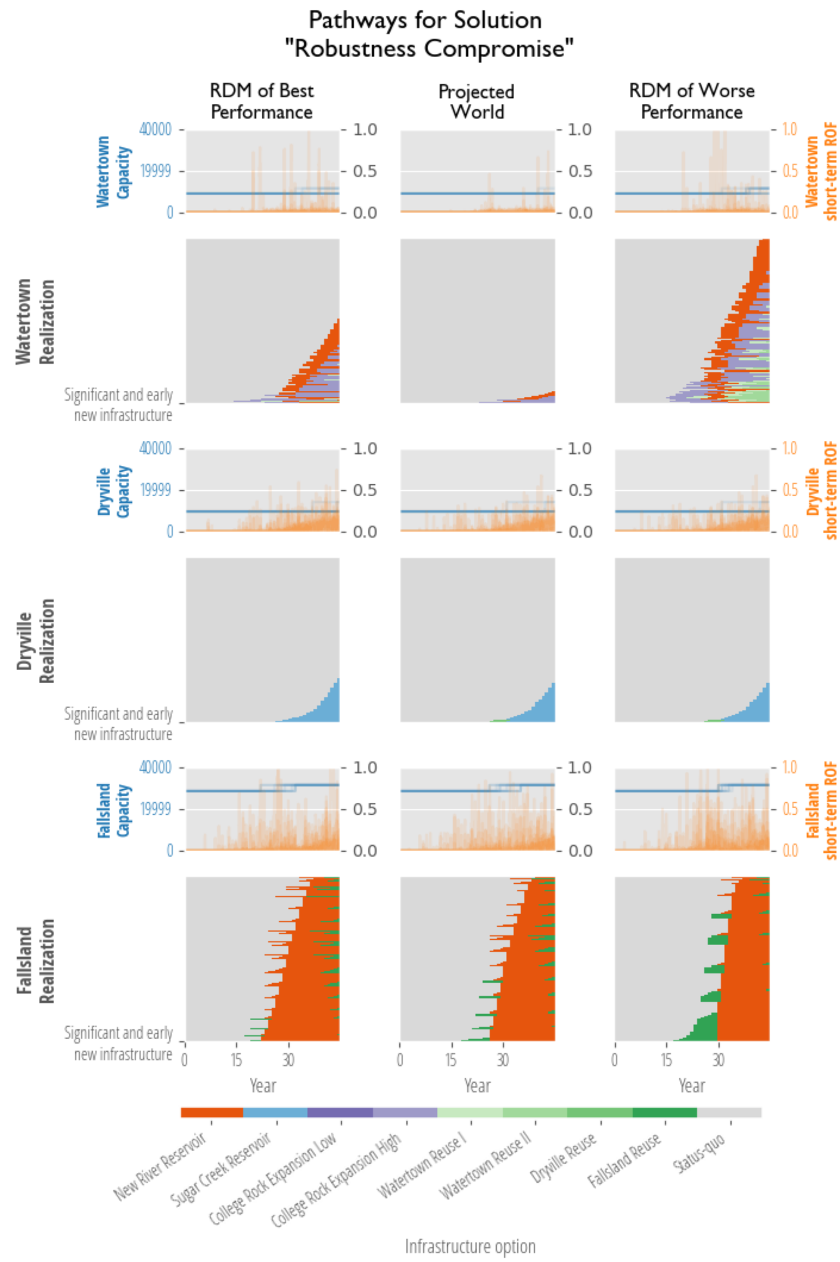

Figure 4.13: Capacity expansion pathways and ROF dynamics for the most favorable scenario. Demand growth rate scenarios considered include the Most Favorable SOW, the SOW projected by the utilities, and the Least Favorable SOW. All infrastructure investments are illustrated as stacked color lines designating specific infrastructure investments across 1,000 hydrologic scenarios.

4.6.4 Infrastructure Ensemble Pathway Analysis

Figure 4.13 illustrates the detailed analyses that WaterPaths enables for each of the utilities, which includes the joint outputs for capacity expansions, short-term ROF dynamics, and pathways for each of the utilities under their individual most favorable, projected, and least favorable scenarios highlighted by starts in Figure 4.12. Figures 4.13a - 4.13c plot Watertown's storage capacity and short-term ROF dynamics for each of the three scenarios. Likewise, Figures 4.13d - 4.13f show Watertown's resulting ensemble of infrastructure pathways across the three highlighted demand growth scenarios. The same panel layout is repeated for the Dryville and Fallsland pathways. The horizontal axes in the pathway panels represent time in weeks from 2015 to 2060, the planning horizon. Each infrastructure pathway plot displays a stack of results attained for 1,000 SOWs and the horizontal multicolored lines, each represent the construction sequence that would be triggered under a specific hydrological scenario on the vertical axis. For each horizontal line in the stack, the color of a segment represents the infrastructure option built last in time. That being the case, a transition in color indicates the time when an infrastructure option becomes operational. For example, the orange horizontal segment starting at the bottom-center (significant and early infrastructure, year 26) of the panel "p" indicates that in that realization, one with more build infrastructure options, Fallsland finished the construction of the New River Reservoir 9 years after finishing the construction of its reuse station, brought online in year 17.

Figure 4.13, panels "d" through "f" indicate that the amount of infrastructure built by Watertown varies significantly with demand growth. While the scenario with the highest demand growth triggers infrastructure investments for all hydrological realizations and still fails to minimize short-term ROF values, the projected future triggers infrastructure investments in less than 10% of the hydrological realizations. However, the low investments in infrastructure observed in the projected scenario, shown in panel "e," makes Watertown more susceptible to swings in revenue in more challenging realizations due to the need for frequent restrictions. Such financial swings cause Watertown to fail the worse-first-percentile-cost target, resulting in a red zone around the blue star in Figure 4.12a. On the other hand, all three scenarios for Dryville and Fallsland see the same infrastructure-investment patterns with small timing variations, which combined with the short-term drought mitigation and financial instruments are generally enough to make utilities meet their robustness performance criteria in the regions highlighted in gray in Figure 4.12a and (b).

An important aspect of the infrastructure investment problem is the effect of permitting times on the scheduled capacity expansions, which may cause the construction of an infrastructure option with high permitting time to be postponed in favor of another option with already granted permits. In the pathways illustrated in Figures 4.13 panels "d" through "f" and "p" through "r" for Watertown and Fallsland, the occurrence of purple (high-capacity College Rock expansion) and green lines (Fallsland reuse) segments, respectively, before the orange regions (New River reservoir) shows that this issue often happen to the jointly-built New River reservoir, resulting in an observed infrastructure construction order that does not match the order specified in the corresponding policy (robustness compromise policy) in Figure 4.11. More specifically, in scenarios of higher stress, when either Watertown or Fallsland trigger the News River reservoir its permit may be strongly delayed by the extended permitting period, so Watertown may trigger instead the College Rock reservoir expansion and Fallsland may trigger the construction of its water reuse station, both with substantially reduced permitting times.

4.7 Conclusion

Recent projections estimate that the United States (US) will require over one trillion dollars of investment in water supply infrastructure in the next 20 years, which can potentially be minimized by exploring synergies between infrastructure investment and non-structural drought mitigation instruments, both shared among regional actors. The coupling of long-term infrastructure investment strategies with short-term drought mitigation necessitates a model with capabilities not all found in a single currently available software: joint simulation of multiple utilities involved in joint planning and operations of shared water resources and infrastructure, joint financial-supply simulation, broad array of uncertainties within a single simulation, easy customizability, and ready integration with multiobjective optimization algorithms for use with high performance computing. WaterPaths was designed and shown to fill these gaps.

We present the design of WaterPaths, an open-source model written in the C++ programming language and designed to fit the DU Pathways. In addition, we demonstrate the application of WaterPaths on a hypothetical water regionally-coordinated infrastructure planning and management problem, called the the Sedento Valley. Using WaterPaths, we modeled the decision making and system upgrades of the regionally-coordinating water utilities in the Sedento Valley under well-characterized and deep uncertainty to explore policy options, discovered inter-dependencies between utilities, and discovered vulnerabilities revealed under deep uncertainty. Additionally, WaterPaths allowed us to discover policy alternatives for all utilities that land in different location of the objectives-tradeoff curve by integrating the MS Borg MOEA and running the optimization exercises in different high-performance-computing platforms.

The Sedento Valley test case was designed to be itself a contribution to the field of water systems engineering. It was designed to be a minimalist test case for the direct comparison of competing regional-decision-making frameworks, for the design and comparison of drought mitigation strategies, and for the evaluation of decision-making metrics. While minimalist for computational tractability and interpretability, the Sedento Valley test case maintains a high degree of complexity due to the number of utilities involved, to the unbalanced distribution of resources and liabilities across utilities, and to the utilities locations in relation to each other in the local shared basins. The code repository provided in WaterPaths' code repository contains all the data needed to reproduce the test case in WaterPaths or other frameworks and to compare new methodologies and policy instruments against the ones presented this paper.

CHAPTER 5

DEEPLY UNCERTAIN PATHWAYS: INTEGRATED MULTI-CITY REGIONAL WATER SUPPLY INFRASTRUCTURE INVESTMENT AND PORTFOLIO MANAGEMENT

This chapter is drawn from the following peer-reviewed journal article: Trindade, B., Reed, P. M., and Characklis, G., "Deeply Uncertain Pathways: Integrated Multi-City Regional Water Supply Infrastructure Investment and Portfolio Management.", Advances in Water Resources, v134, DOI:10.1016/j.advwatres.2019.103442, (2019). Funding for this work was provided by the National Institute of Food and Agriculture, U.S. Department of Agriculture (WSC Agreement No. 2014-67003-22076). Additional support was provided by the U.S. National Science Foundation's Water, Sustainability, and Climate Program (Award No. 1360442). The views expressed in this work represent those of the authors and do not necessarily reflect the views or policies of the NSF or the USDA.

5.1 Abstract

This chapter contributes the Deep Uncertainty (DU) Pathways framework for bridging long-term water supply infrastructure investments and improved short-term water portfolio management (e.g., restrictions, water transfers, financial instruments, etc.) to yield robust regional water supply. The DU Pathways framework combines flexibility-providing risk-of-failure (ROF) decision rules presented in Chapter 3.5, dynamic adaptive policy pathways concepts, and a careful consideration of time-evolving information feedbacks to yield management-conditioned infrastructure pathways for regions. The DU Pathways' framework has been developed to carefully consider multi-actor regional

contexts with the goal of aiding stakeholders in discovering pathway policies that attain high performance levels across challenging, deeply uncertain futures and to guide robustness compromises that may be necessary between regional actors. As demonstrated in the Research Triangle region of North Carolina, DU Pathways clarifies how to identify robust infrastructure investment and management policies across the municipalities of Raleigh, Durham, Cary, and Chapel Hill. Our results provide insights about the most cost-effective infrastructure options to be pursued in the near-term, and clarify which sources of uncertainty drive the performance and robustness of the regional system and of the individual actors.

5.2 Introduction

Water utilities globally are facing a growing pressure to proactively address growing demands, higher levels of resource contention, and increasingly uncertain water availability (Bonzanigo et al., 2018; Hall et al., 2014). A tension exists between improving the efficiency and coordination of regional water supplies to delay or eliminate major infrastructure investments and the eventual effects of demand growth rates that necessitate water supply capacity expansions. The soft water management strategies presented in Chapter 3 (conservation, inter-utility water transfers, demand management, etc.) are a key option to cope with such growing domestic water demand (Gleick, 2002a, 2003). In the United States (US), a key driver increasing the importance of soft water management strategies is that most large projects of federal and/or state interest have already been built (Lund, 2013). However, growing urban water demands and increasingly uncertain climate conditions are motivating efforts for redesigning system capacity by upgrading existing and adding new infrastructure to maintain reliable water supply systems (Shafer and Fox, 2017).

The hard and soft path views for urban water portfolio planning should be treated as complementary, as even though infrastructure expansion may be needed, it is not desirable and may be mitigated by soft water management strategies. However, both the soft and hard approaches may be detrimental to utilities' finances. Together with outstanding debt, decreasing cash flows are significant factors that negatively impact utilities credit ratings, which can increase the cost of borrowing to build new infrastructure (Moody's, 2017; Zeff et al., 2014; Zeff and Characklis, 2013; Trindade et al., 2017). The financial impacts of drought mitigation strategies studied in Chapter 3, the current negative financial outlook for the water utility sector (Moody's, 2019), the lack of appropriate financing for water utilities (Gleick et al., 2014) all stress the need for new frameworks that facilitate improved water portfolio planning that effectively combines short-term drought mitigation instruments for managing financial risks and long-term infrastructure investment pathways while accounting diverse sources of uncertainty. These needs motivated the creation of the Deep Uncertainty Pathways (DU Pathways) framework proposed and demonstrated in this study.

Integrating long-term water infrastructure investment and short-term management is a difficult task because of the potentially large number of decisions that must be considered over decadal timescales as well as the significant uncertainties inherent to the problem. Hydro-climatic uncertainties such as near-term streamflows and evaporation rates have been extensively studied and characterized with well-established statistical techniques (Stedinger, 1993; Martins and

Stedinger, 2000; Herman et al., 2016; Lall, 1995; Kirsch et al., 2013). However, water infrastructure investment and management problems include a multitude of deep uncertainties — uncertainties for which decision makers cannot agree on their boundaries, importance, and probability distributions (Kwakkel et al., 2016b; Knight, 1921), see Section 1.1.2. Examples include regional demand growth, political and economic uncertainties, and climate change (Herman et al., 2014; Trindade et al., 2017; Milly et al., 2008).

Although these are not new issues, utilities are still struggling with them (Bonzanigo et al., 2018; Paulson et al., 2018). Although WaterPaths was developped with these issues in mind, existing commercial frameworks for water infrastructure management modeling (e.g., HydroLogics, 2009; Sieber and Purkey, 2015; Labadie, 2011) focus on existing infrastructure, have strong limitations on their inclusion of uncertainties, and do not connect to key concepts from the infrastructure investment literature

Our DU Pathways framework broadens the water portfolio problem to also consider long-term planning and the construction of new infrastructure into the short-term management framework presented in Chapter 3. The DU Pathways framework uses some features of ROA (Wang and de Neufville, 2005; Erfani et al., 2018; Fletcher et al., 2017; Hui et al., 2018) such as a focus on flexible decision rules and a careful consideration of short-term versus long-term information. However, ROA as a single-objective approach is limited in its accounting of stakeholders with diverse interests and its underlying mathematical tree logic becomes computationally intractable when large number of uncertainties are considered [curse of dimensionality, (Dittrich et al., 2016)]. Borgomeo et al. (2018) address some of these concerns by proposing a multi-objective evaluation of candidate water supply investments. They maximize robustness and minimize risk and cost of a 30-year long fixed construction schedule encompassing multiple infrastructure options. It should be noted that the derived sequence of investments is not adaptive and does not account for evolving information feedbacks. These limitations also exist in the recent studies by Beh et al. (2015b) and Huskova et al. (2016), although the former seek to improve adaptivity by incorporating a similar emphasis on near term information as done in ROA. Alternatively, the recently introduced dynamic adaptive pathways policy (DAPP) approach (Haasnoot et al., 2013; Kwakkel et al., 2014) adds flexibility to the planning process by making use of potentially diverse sets metrics (signposts). DAPP, however, still maintains a strong reliance on the use of limited numbers of predefined action sequences that require high levels of institutional stability and coordinated consensus over the long term.

The DU Pathways framework introduced in this work study is first formal integration of dynamic and adaptive water infrastructure pathways (Zeff et al., 2016; Haasnoot et al., 2013) (annual, long-term planning), financial and supply-reliability drought mitigation instruments [weekly, short-term decisions discussed in Chapter 3 and in Zeff et al. 2014 and Trindade et al. 2017], and the recent extensions of the MORDM framework that incorporate deep uncertainties in search-based identification candidate actions (Kasprzyk et al., 2013; Trindade et al., 2017; Kwakkel et al., 2014; Watson and Kasprzyk, 2016). DU Pathways addresses decision making at two time scales: weekly, for operational management decisions, and annual, for infrastructure investment decisions. Following the approach recommended by Zeff et al. (2016), the integrated infrastructure investment and management policies specify short- and long-term decision-making rules based on the risk-of-failure metric (ROF), as well as candidate orderings for infrastructure investments. By combining infrastructure planning and management rules in a single policy, the water portfolio planning approach aids the design of drought mitigation instruments tailored to new infrastructure to minimize the need of further investments. The main limitation of the infrastructure pathways work by Zeff et al. (2016) is its limited treatment of uncertainties. The DU Pathways framework contributed here significantly expands the suite of deep uncertainties considered in the search-based identification of candidate infrastructure investment and management policies. Additionally, the framework formalizes a detailed evaluation of the multi-city policies' robustness tradeoffs. DU Pathways aids stakeholders in navigating these tradeoffs through a combination of visual decision analytics Woodruff et al. (2013) and enhanced scenario discovery in which boosted trees Schapire (1999) aid in the identification of the uncertainties most strongly shape the infrastructure investment and management policies' vulnerabilities. The application of the DU Pathways framework is demonstrated the Zeff et al. (2016) long-term version of Research Triangle test case described in Section 3.4 and prior works (Trindade et al., 2017; Herman et al., 2014; Zeff et al., 2014), where we have discovered how the region's vulnerabilities evolve under alternative policies, clarified significant interdependencies between the region's four utilities, and show the importance of carefully balancing supply and financial instruments. Furthermore, this is chapter also serves to test the WaterPaths framework in a challenging real-world regional decision context.

5.3 The Long-Term Version of the Research Triangle Test Case

The DU Pathways framework introduced in this study is demonstrated on the long-term version of Research Triangle region in North Carolina presented in Section 3.4. As mentioned in Section 3.4, the present test case is comprised of the four inter water utilities (Figure 5.1) interconnected through water mains and currently responsible for meeting the majority of the region's growing demand in the Research Triangle portion of the Neuse and Cape Fear river basins: Raleigh, Durham, Cary and Orange Water and Sewer Authority (OWASA, representing Carrboro and Chapel Hill). The Triangle Water Supply Partnership (TWSP) of which the four utilities participate has focused on developing collaborative short-term drought management and infrastructure planning solutions for its members with specific focus on meeting projected 2060 water supply demands (Triangle J, 2014). The two key challenges for the Research Triangle Region are to develop (1) collaborative plans for short-term regional management of the existing water infrastructure, as mentioned in Section 3 and (2) plan long-term infrastructure investments to meet future demands.

The the long-term version of the Research Triangle test case includes all current infrastructure systems providing water to the four utilities presented in Section 3.4 (Table 3.4) as well as recently proposed candidate investments (Table 5.1) including additional storage, conveyance and treatment projects across the four utilities laid out by the JLP in Triangle J (2014). OWASA and Durham own allocations in the Jordan Lake that can currently only be accessed through Cary's water treatment plant on eastern shore of Jordan Lake, subject to treatment and inter-utility conveyance availabilities (Figure 5.1). Currently, Jordan Lake's water supply storage is still not fully allocated and several utilities in the

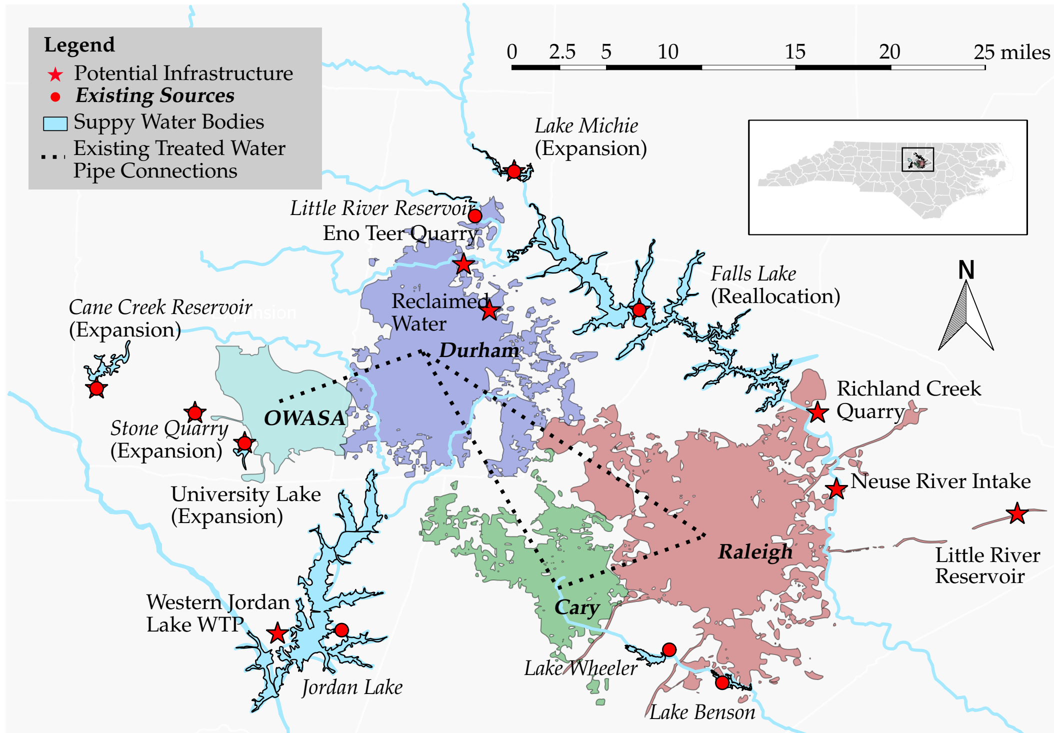

Figure 5.1: Map of the Research Triangle Region test case. Potential Infrastructure refers to infrastructure options presented in Triangle J (2014), that are being considered by the utilities.

Research Triangle are planning on submitting new allocation requests, which gives rise to potential resource contention that the the four studied water utilities want to avoid.

Key short-term (i.e. drought) regional management instruments include water use restrictions and treated water transfers that are employed using weekly operational rules with the goal of making better use of the Research Triangle region's shared resources. Chapter 3 focused on defining management operations collaboratively so that the impacts from the measures taken by each utility has a predictable impact on the operations of the other utilities, while this chapter will solve the same problem but considering added infrastructure and over a longer period of time (through 2060).

Utility	Infrastructure Option	Storage/ Production	Cost ($MM)
OWASA	University Lake expansion	2550 MG	107
OWASA	Cane Creek Reservoir expansion	3000 MG	127
OWASA	Stone Quarry shallow expansion	1500 MG	1.4
OWASA	Stone Quarry deep expansion	2200 MG	64.6
Durham	Teer Quarry expansion	1315 MG	22.6
Durham	Reclaimed water (low)	2.2 MGD	27.5
Durham	Reclaimed water (high)	11.3 MGD	104.4
Durham	Lake Michie expansion (low)	2500 MG	158.3
Durham	Lake Michie expansion (high)	7700 MG	203.3
Raleigh	Little River Reservoir	3700 MG	263
Raleigh	Falls Lake WQ Pool Reallocation	4100 MG	68.2
Raleigh	Neuse River Intake	16 MGD	225.5
Regional	Western Jordan WTP (low capacity)	33 MGD	243.3
Regional	Western Jordan WTP (high capacity)	54 MGD	73.5

Table 5.1: Summary of proposed water supply infrastructure for the Research Triangle (Zeff et al., 2016; Triangle J, 2014).

As a first step in our development of the DU Pathways framework contributed in this study, we borrowed from the model presented in Section 3.4.1 and developed a flexible regional model using WaterPaths that can encompass a breadth of system uncertainties and support our evaluation of the diverse portfolios of short-term management actions as well as long-term infrastructure investment pathways. At its most basic core, our simulation of the Research Triangle is a water mass-balance model with the inputs and outputs illustrated in Figure 5.2 below.

The model is used in a set of optimization runs whose results were further assessed to identify robust infrastructure management and investment policies for each of the four Research Triangle utilities. Beyond management actions, made at a weekly timestep, the utilities must also define infrastructure investment pathways that specifically prioritize and sequence major projects (among

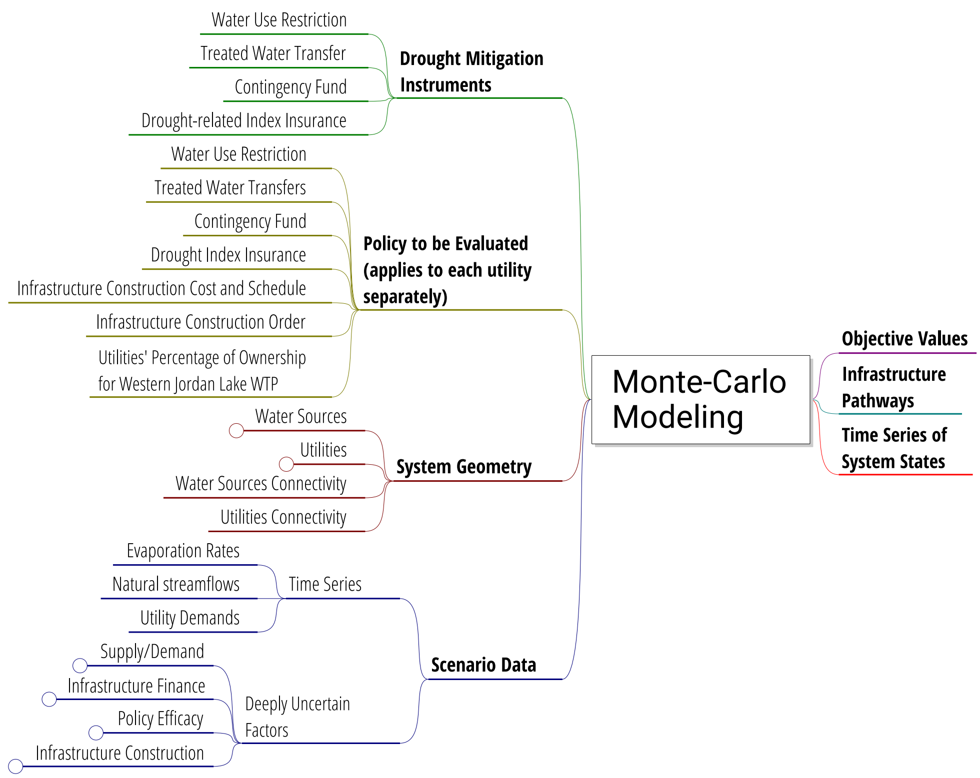

Figure 5.2: The Research Triangle instance of the WaterPaths simulation model takes as inputs a policy to be evaluated, the system components and its connectivity, hydrological time series, infrastructure construction order, information about infrastructure to be built, different types of demands, and uncertainty. It outputs time series of system states, an objectives vector and pathways.

options specified in Triangle J, 2014). As with the short-term drought mitigation instruments (restrictions and transfers), such infrastructure sequences and corresponding decision-making rules are sought that will positively impact the Research Triangle region's supply reliability and financial stability. The management and investment actions are critically interdependent, given that effective short-term drought mitigation instruments can serve to dramatically delay and/or reduce capital investments (Zeff et al., 2016).

Type	Uncertainty Source	Regional or Individual for Utility/ Infrastructure
Supply/Demand	Growth of Mean Annual Demand	Regional
	Growth of Mean Annual Evaporation	Regional
Financial/ economics	Interest Rate	Regional
	Bond Term	Regional
	Discount Rate	Regional
Policy Effectiveness	Water Use Stage Restriction Efficacy	Individual
Infrastructure Construction	Permitting Time	Individual
	Construction Cost	Individual

Table 5.2: Uncertainties included in extended version of the regional water utility planning and management problem presented in Zeff et al. (2016). Regional uncertainteis are those for which the same value was used for all utilities, versus a value different value for each utility, as in the case uncertainties described as individual.

5.4 Methodology

Expanding the methodology presented in Section 3.3, the major steps of the DU Pathways methodology as illustrated in Figure 5.3 are as follows:

1. Identifying Infrastructure Investment and Drought Management Policy Tradeoffs: Candidate Pareto-approximate infrastructure investment and water portfolio management policies are identified using robust multiobjective optimization (Reed et al., 2013; Hadka and Reed, 2014; Deb et al., 2002b; Shah and Ghahramani, 2016; Hernández-Lobato et al., 2016).

2. Defining and Evaluating Robustness: Exploiting a more expansive sampling of the potential deep uncertainties that could pose challenging states-of-the-world (SOWs) in which robustness metrics can be used to

evaluate and rank-order solutions for their robustness. For reviews and comparisons of alternative definitions of robustness, see (McPhail et al., 2018; Herman et al., 2015; Giuliani and Castelletti, 2016).

3. Discovering Uncertain Scenarios that Control Robustness: The Pareto approximate of policies of interest once evaluated against a broader sample of SOWs undergo further diagnostic assessment by means of sensitivity analysis (Sobol, 2001; Saltelli, 2002; Saltelli et al., 2008; Morris, 1991; Campolongo et al., 2007; Herman and Usher, 2017; Jaxa-Rozen and Kwakkel, 2018) and/or scenario discovery (Bryant and Lempert, 2010; Quinn et al., 2018; Kwakkel, 2017) to determine which specific uncertainties dominantly impact system's performance (factor prioritization) for specific combinations of values (factor mapping).

4. Visual Infrastructure Pathway Analysis: Use of visual analytics to determine how infrastructure investment decisions are impacted by deviations from base projections for the most important sources of uncertainty.

The four DU Pathways framework steps listed above and summarized graphically in Figure 5.3 are described in more detail in the remainder of this section.

5.4.1 Identifying Infrastructure Investment and Drought Management Policy Tradeoffs

As mentioned in Section 5.2, DU Pathways framework bridges three active areas of research in recent literature: (1) dynamic adaptive pathway planning (Haas-

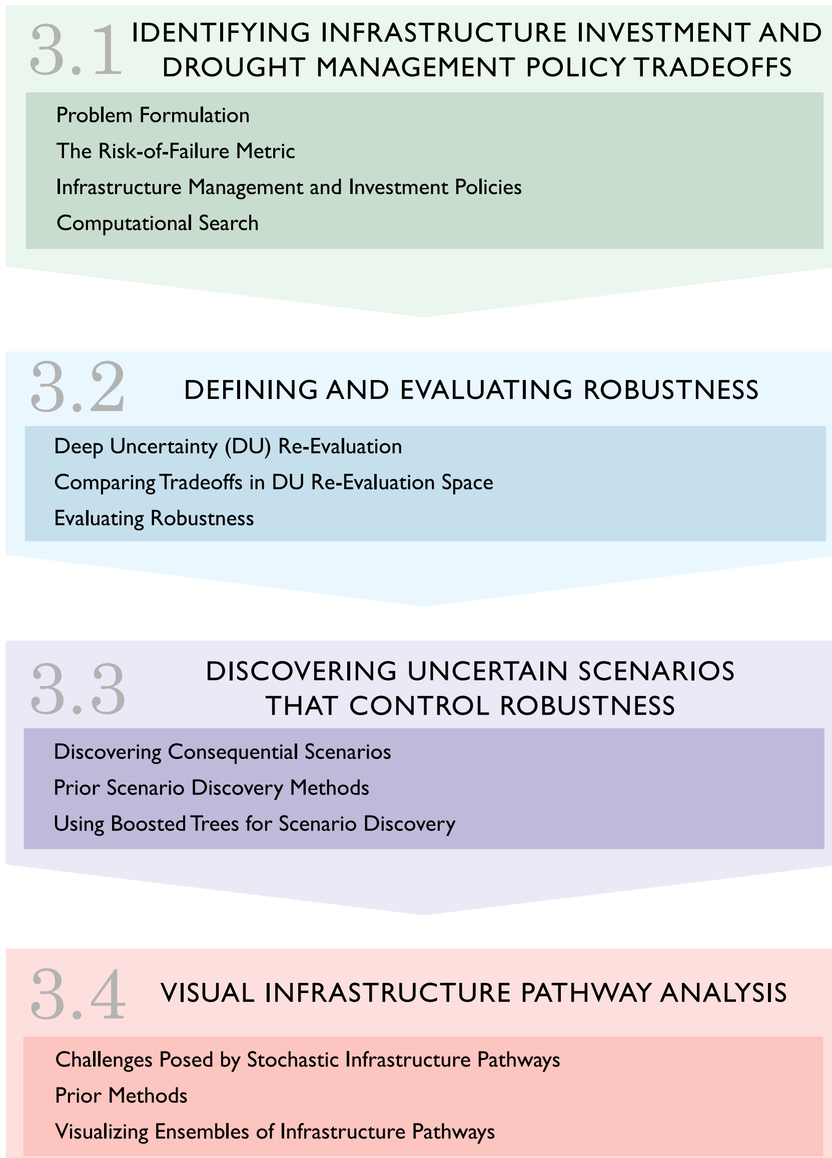

Figure 5.3: Key steps of the DU Pathways approach and their key points

noot et al., 2013; Kwakkel et al., 2014; Zandvoort et al., 2017), (2) ROF-based urban water portfolio planning (Palmer and Characklis, 2009; Zeff et al., 2014, 2016; Herman et al., 2014), and (3) multi-objective optimization under deep uncertainty (Kwakkel et al., 2014; Watson and Kasprzyk, 2016). In DU Pathways, action triggers based on the ROF dynamics are the basis of infrastructure investment and management policies. These policies are in turn what determine the infrastructure investment pathways. The ROF-based action triggers (or decision variables) are based on the premise of monitoring storage-to-demand dynamics to yield a time-continuous assessment of when the risks of unacceptably low water supply capacity requires action. As demonstrated by Zeff et al. (2016), the ROF-based action triggers have the advantage of providing system adaptability across planning time-scale decisions (i.e., annual infrastructure construction) and short-term management time-scale decision (i.e., the weekly use of drought mitigation instruments). In either case, actions are triggered when the given ROF metric (Palmer and Characklis, 2009) crosses ROF threshold values stated in a fully specified infrastructure investment and management policy (Zeff et al., 2016). The set of all regional ROF-triggers, other financial policy variables and infrastructure construction order for all utilities in the problem defines the operations of all drought mitigation instruments and the strategy for the construction of infrastructure. This set is what is in this work called an infrastructure investment and management policy, shown in Figure 5.4. With that, the minimization problem to be solved then becomes a matter of finding positions for the decision levers in Figure 5.4, where the levers represent the decision/policy variables, that lead to Pareto-approximate policies.

Problem Formulation

The multi-objective minimization problem formulation presented below in Equations 5.1 - 5.6 is a general mathematical abstraction for search-based identification of regional water infrastructure investment and management policies, denoted by θ, for all members of a cooperative group of water utilities. The optimization problem focuses on finding Pareto-optimal policies θ^* that minimize

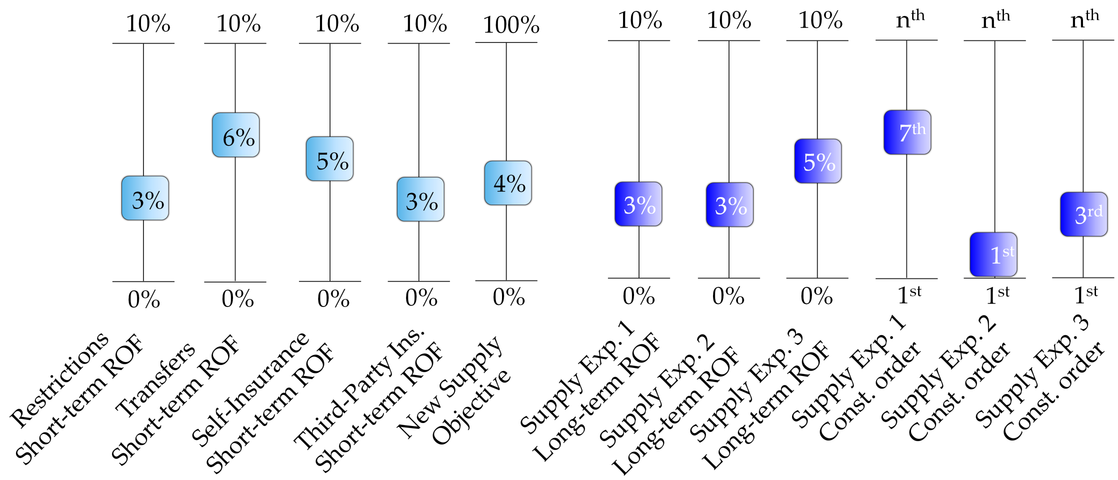

Figure 5.4: Illustrative portfolio of decisions related to short-term drought mitigation instruments and long-term priorities for infrastructure investments pathways that define a system's water supply policy. The optimization problem consists in finding the best position for each knob corresponding to each decision variable for each utility.

the objective function vector $\boldsymbol{F}$:

$$\theta^* = \underset{\theta}{argmin}\ \boldsymbol{F} \tag{5.1}$$

where,

$$\boldsymbol{F} = \begin{bmatrix} -\boldsymbol{f_{REL}}(\boldsymbol{x}_s, \boldsymbol{\theta}_{rt}, \boldsymbol{\theta}_{tt}, \boldsymbol{\theta}_{jla}, \boldsymbol{\theta}_{it}, \boldsymbol{ICO}, \boldsymbol{\Psi_s}) \\ \boldsymbol{f_{RF}}(\boldsymbol{x}_{srof}, \boldsymbol{\theta}_{rt}, \boldsymbol{\theta}_{tt}, \boldsymbol{\theta}_{jla}, \boldsymbol{\theta}_{it}, \boldsymbol{ICO}, \boldsymbol{\Psi_s}) \\ \boldsymbol{f_{NPV}}(\boldsymbol{x}_{lrof}, \boldsymbol{ICO}, \boldsymbol{\Psi_s}) \\ \boldsymbol{f_{FC}}(\boldsymbol{x}_{srof}, \boldsymbol{\theta}_{rt}, \boldsymbol{\theta}_{tt}, \boldsymbol{\theta}_{jla}, \boldsymbol{\theta}_{acfc}, \boldsymbol{\theta}_{irt}, \boldsymbol{\theta}_{it}, \boldsymbol{ICO}, \boldsymbol{\Psi_s}, \boldsymbol{x}_{lrof}) \\ \boldsymbol{f_{WFPC}}(\boldsymbol{x}_{srof}, \boldsymbol{\theta}_{rt}, \boldsymbol{\theta}_{tt}, \boldsymbol{\theta}_{jla}, \boldsymbol{\theta}_{acfc}, \boldsymbol{\theta}_{irt}, \boldsymbol{\theta}_{it}, \boldsymbol{ICO}, \boldsymbol{\Psi_s}, \boldsymbol{x}_{lrof}) \\ \boldsymbol{f_{JLA}}(\boldsymbol{\theta}_{jla}) \end{bmatrix} \tag{5.2}$$

s.t.

$$|\boldsymbol{ME}| \leq 1\ \forall\ \boldsymbol{ME} \subseteq \boldsymbol{BI} \tag{5.3}$$

where

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_{srof} \\ \boldsymbol{x}_{lrof} \\ \boldsymbol{x}_s \end{bmatrix} \tag{5.4}$$

$$\theta = [\boldsymbol{\theta}_{rt}, \boldsymbol{\theta}_{tt}, \boldsymbol{\theta}_{acfc}, \boldsymbol{\theta}_{irt}, \boldsymbol{\theta}_{it}, \boldsymbol{ICO}, \boldsymbol{\theta}_{jla}] \tag{5.5}$$

$$\boldsymbol{\Psi_s} = \begin{cases} \boldsymbol{\Psi}_{WCU} = \begin{bmatrix} - & \boldsymbol{\psi}_{WCU} & - \\ - & \boldsymbol{\psi}_{WCU} & - \\ & \vdots & \\ - & \boldsymbol{\psi}_{WCU} & - \end{bmatrix} \\ \boldsymbol{\Psi}_{DU} = \begin{bmatrix} - & \boldsymbol{\psi}_{DU,1} & - \\ - & \boldsymbol{\psi}_{DU,2} & - \\ & \vdots & \\ - & \boldsymbol{\psi}_{DU,N_R} & - \end{bmatrix} \end{cases} \tag{5.6}$$

In the equations above, $\boldsymbol{X}$ is the time-varying state matrix across all of the utilities, where the components $\boldsymbol{x}_{srof}$ and $\boldsymbol{x}_{lrof}$ are vectors of short- and long-term values for the dynamic tracking of ROF. This distinction refers to the short-term ROF metric calculated over 52 weeks and used for the short-term instruments, as presented in Section 3, and the long-term version of the same metric using $T_{ROF} = 78$ weeks in Equation 3.8 used for infrastructure triggering. The decision variable $\boldsymbol{\theta}_{acfc}$ is a vector of annual contingency fund contributions, which are percentages off annual revenue saved in a utility's drought mitigation fund. As mentioned above, restrictions, transfers, insurance, and infrastructure investment decisions are all formulated consistently through time using ROF triggers. The overall regional policies are composed of the following decisions for all utilities: the $\boldsymbol{\theta}_{rt}$ vector of restriction triggers, the $\boldsymbol{\theta}_{tt}$ vector of transfer triggers, the $\boldsymbol{\theta}_{jla}$ vector of Jordan Lake Allocations, the $\boldsymbol{\theta}_{irt}$ vector of insurance restriction triggers, the $\boldsymbol{\theta}_{it}$ vector of long-term-ROF infrastructure construction triggers, and the $\boldsymbol{ICO}$ vectors with the infrastructure construction ordering for each utility. Overall, all decision variables are encompassed by the triggers $\boldsymbol{\theta}$ plus the infrastructure construction order $\boldsymbol{ICO}$. A Pareto optimal management and investment policy is denoted by $\boldsymbol{\theta}^*$ in Equation 5.1. Equation 5.3 enforces the requirement that two infrastructure options, such as high and low capacity expansions of a reservoir, cannot be built in the same realization (Zeff et al., 2016; Erfani et al., 2018). In Equation 5.3, ME represents a generic subset of mutually exclusive infrastructure options within the set of built or prospective infrastructure BI. Matrix $\boldsymbol{\Psi}_s$ is comprised of the vector samples of the deeply uncertain variables of concern where each element corresponds to a specific uncertainty. $\boldsymbol{F}$ is the vector-values objective function where f_{REL} is the supply reliability objective, f_{RF} is the restriction frequency, f_{NPV} is the net present cost of infrastructure, f_{FC} is the total cost of financial drought mitigation instruments (contingency fund and insurance loading) plus debt repayment, f_{WFPC} is the worse-first-percentile of financial variability caused by the supply drought mitigation instruments (restrictions and transfers), and f_{JLA} is the combined Jordan Lake allocation objective, representing the allocation of the municipal supply pool of the Jordan Lake, in the Research Triangle, NC, to be granted to a utility, such that:

$$\text{minimize } f_{JLA} = \sum_{j=1}^{N_u} JLA_j \tag{5.7}$$

where JLA is the percent of the pool allocated to utility j, and N_u is the total number of utilities in the system. Utilities want to minimized requested allocation due to local political contention surrounding the Jordan Lake allocation requests. Such objective is specific to the test case presented in this study, although groups of utilities in different regions may go through similar conflicts. A detailed formulation of the other five objectives is presented in Sections 3.5 and 4.4.9. Both sampling processes for matrix $\boldsymbol{\Psi}_s$, well-characterized uncertainty (WCU) sampling ($\boldsymbol{\Psi}_{WCU}$) and deep uncertainty (DU) sampling ($\boldsymbol{\Psi}_{DU}$), are described in detail in the next section.

Sampling States-of-the-World for Alternatives Generation

All of the objectives presented above in Equation 5.2 except for the Jordan Lake allocation are stochastically evaluated within a Monte Carlo simulation framework across an ensemble of candidate SOWs comprised of time series of inflows, eraporation rates, and demands, and potentially several deeply uncertain factors. The DU Pathways framework distinguishes between two formulations of the water portfolio management and infrastructure investment problem. These formulations differ in their forward Monte Carlo-based evaluations based on sampling the matrices of natural inflows ($\boldsymbol{NI}$), evaporation series ($\boldsymbol{E}$), unrestricted demands ($\boldsymbol{UD}$), and samples of deeply uncertain factors ($\boldsymbol{\Psi}_s$). The Well-Characterized Uncertainty (WCU) and Deep Uncertainty (DU) sampling schemes employed in this study are illustrated in Figure 5.5 and were imported from Section 5.4.1.

Figure 5.5a shows that the WCU sampling scheme specifies SOWs (rows of rectangles) by sampling uncertainties for which historical-data-driven distributions are available from their historical observations, represented by colored rectangles (e.g., natural stream inflows, evaporation rates, and demand fluctu-

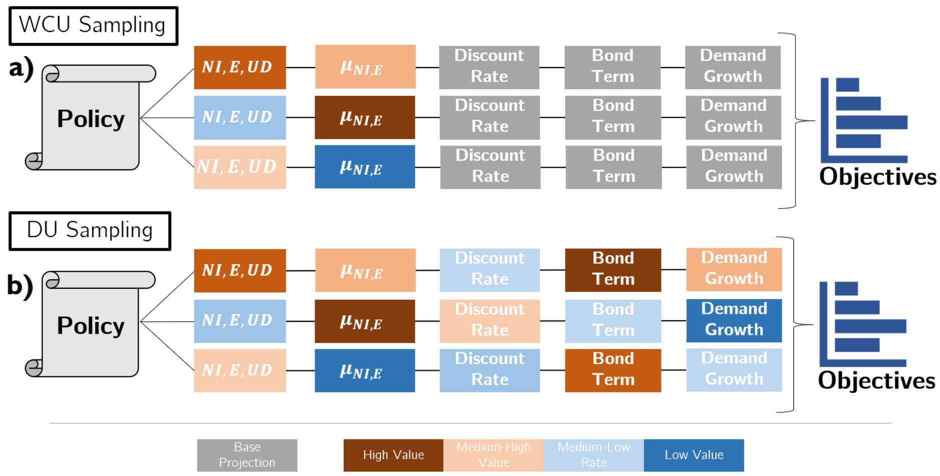

Figure 5.5: (a) In the WCU sampling scheme, sets of time-series samples of synthetically generated streamflows, evaporation rates and demands are generated for each realization and appended to the historical data. The best projected values (or other assumption) of the deeply uncertain sources of uncertainty are then coupled with the resulting sets of time series form all SOWs. (b) In DU sampling, the same sets of time-series are instead combined with samples of deep uncertainties to form the ensemble of SOWs.

ations around projected mean demands). For all other factors, represented by grey rectangles in Figure 5.5a, deterministic best projections are used. Alternatively, the DU sampling strategy shown in Figure 5.5b increases the number of uncertain factors that are varied across SOWs by sampling uncertainties not only from data-derived distributions but also from expert elicitation. More formally, the WCU and DU sampling schemes can be abstracted mathematically as a forward Monte Carlo approximation of the system performance (see Equation 5.8). Table 5.2 shows the deep uncertainties considered in this study, and as represented in Equation 5.8 are assumed to be uncorrelated with the well-characterized uncertainties.

$$\boldsymbol{NI}_s, \boldsymbol{E}_s, \boldsymbol{UD}_s, \boldsymbol{\Psi}_s \sim f(\boldsymbol{NI}_h, \boldsymbol{E}_h, \boldsymbol{UD}_h, \boldsymbol{S}, \boldsymbol{R}) \cdot \mathcal{D}(\cdot) \tag{5.8}$$

Within Equation 5.8 the subscript s denotes synthetic and h denotes historical time series. Synthetically generated series compose matrices for natural inflows ($\boldsymbol{NI}$), evaporation ($\boldsymbol{E}$), and unrestricted demands ($\boldsymbol{UD}$) where their dimensions are the number of synthetic series samples (N_r) by the total number of weeks in the full planning horizon (N_w). The historical data used to derive the synthetic series (subscript h) are represented in respective matrices in equation 5.8 of dimension of the number years in the historical record (N_{hr}) and the number of weeks in a year (N_w). Cross-site correlation and temporal auto-correlation in the historical natural inflows and evaporation time series are captured with the $\boldsymbol{S}$ and $\boldsymbol{R}$ matrices, respectively. The joint-pdf represented by f characterizes the inflows, evaporation rates, and demand fluctuations around annual means, all well-characterized uncertainties, while $\mathcal{D}(\cdot)$ is an assumed pdf for the deeply uncertain factors that can take the following forms, depending on the sampling scheme:

$$\mathcal{D}_{WCU}(\boldsymbol{\mathcal{X}}) = \prod_{i=0}^{N_f} \delta(\mathcal{X}_i) \tag{5.9}$$

$$\mathcal{D}_{DU}(\boldsymbol{l}, \boldsymbol{u}) = \prod_{i=0}^{N_f} \mathcal{U}(l_i, u_i) \tag{5.10}$$

Within Equation 5.9, $\delta(\cdot)$ is a Dirac function, $\mathcal{X}$ is a vector of best projections for the deeply uncertain factors considered in the problem, i is the index of a deeply uncertain factor, and N_f is the total number of deeply uncertain factors. Within Equation 5.10, $\boldsymbol{l}$ and $\boldsymbol{u}$ are sets of lower and upper bounds for the deeply uncertain factors, and $\mathcal{U}$ denotes a uniform distribution. The initial probabilistic management and infrastructure investment pathways concepts introduced by Zeff et al. (2016) used only the WCU sampling strategy. The DU Pathways framework significantly expands the sampling scope to consider deep uncertainties and increase their role in the discovery of alternatives. Non-stationary hydro-climatology trends are accounted for in this work by applying a weekly series of multipliers to inflow and evaporation log-mean and log-variance multipliers, increasing and/or decreasing them over time (Quinn et al., 2018). In this work, the series of multipliers was created from sinusoid functions, to assure the inclusion of scenarios in which trends go up and/or downwards (more details in Section (2.2)). The ROF metric, used as a basis of the infrastructure investment and management decision triggers, is presented next.

Tracking Risk-of-Failure Dynamics

The ROF dynamics for each utility are tracked for each week by following the procedure described in Section 4.4.2. The ROF metric presented in Section 4.4.2 simultaneously encompasses information about current demands and storages, and of historical inflows and evaporation rates, improving therefore on information use relative to traditional take-or-pay agreements and days-of-supply-remaining metric (Palmer and Characklis, 2009) and providing the system with more adaptability. We recommend an N_{rof} (one-year long simulations used in the ROF calculation for one week) of at least 50 ROF simulations so that the ROF metric has a precision of at least $1/50 = 2\%$.

Drought Mitigation Instruments and Financial Model

The drought mitigation instruments explored in this study build from Chapter 3 and are focused on water use restrictions and treated inter-utility water transfers

(Zeff et al., 2014, 2016; Caldwell and Characklis, 2014; Palmer and Characklis, 2009; Zeff et al., 2014). Obviously, the overall DU Pathways framework could consider other short-term management actions as would be appropriate in other regions of application. The use of restrictions and transfers are contingent on the value of the short-term ROF metric, as shown in Equations 3.6 and 3.7. Restrictions are tiered based on the progressive severity of a drought, so each tier may have its own short-term ROF trigger value. The transferred volumes requested by utilities are capped for two reasons: (1) limited conveyance capacity and (2) competition between two or more utilities for available volume for transfer, in which case each utility receives a volume proportional to its short-term ROF.

Inter-utility water transfers in combination with water use restrictions have negative financial impacts on the water utilities by reducing revenues and increasing costs. Drought-focused index insurance instruments provide a means to hedge utilities against the financial risks associated with drought mitigation measures. Index insurance, in which payouts depend on an index value, have been studied in literature to hedge utilities against droughts (Zeff and Characklis, 2013; Zeff et al., 2014), hydropower generators against climate risks (Foster et al., 2015), climate-related government regulation (Denaro et al., 2018), and groundwater irrigated agriculture against government groundwater pumping regulations during droughts (Blanco and Gómez, 2013). The ROF-triggered drought index insurance structure used to hedge utilities against financial losses in this study exploits a binary payout structure (Foster et al., 2015; Zeff et al., 2016). Details about the insurance and contingency models used in this work can be found in Section 3.5.

Prioritized Sequencing of Supply Infrastructure Investment Pathways

As with drought mitigation and financial instruments, infrastructure construction or permitting in the DU Pathways framework is triggered rather than pre-scheduled in time (Walker et al., 2001). Triggering occurs in years when the long-term ROF metric, defined in Equations 3.8 to 3.10 as an annual calculation using $T_{ROF} = 78$, crosses a set threshold (Zeff et al., 2016). The DU Pathways framework specifies that When infrastructure triggering occurs for a given utility, construction begins for the next infrastructure option in the construction sequence specified in the policy.

When infrastructure is triggered, utilities issue debt to finance the new infrastructure. The infrastructure and debt service payments stream for all utilities are then updated every year according to Equation 5.11.

$$\boldsymbol{OI}^y, \boldsymbol{DS}^y = f(\boldsymbol{x}_{lrof}, \boldsymbol{\theta}_{irof}, \boldsymbol{ICO}, \boldsymbol{OI}^{y-1}, \boldsymbol{DS}^{y-1}) \tag{5.11}$$

In Equation 5.11, $\boldsymbol{OI}^y$ is the vector with the IDs of all online supply infrastructure, $\boldsymbol{DS}$ is a matrix with the streams of debt service payments for all utilities, $\boldsymbol{x}_{lrof}$ is the vector of long-term ROFs for all utilities, $\boldsymbol{\theta}_{irof}$ is the vector of long-term ROF triggers for all utilities, and $\boldsymbol{ICO}$ is a matrix with the infrastructure construction order for each utility.

Another distinguishing feature of the DU Pathways formulation approach is the direct accounting for the possibility that infrastructure may be built and operated jointly by multiple utilities (Zeff et al., 2016). If triggered by one utility, such projects are built by all utilities who have pledged to be part of the project with the cost proportionally split, meaning that associated risks would be shared among utilities. The supply infrastructure construction logic, together with the short-term drought mitigation instruments, are integrated with a mass-balance model and the WCU and DU uncertainty sampling into a combined framework to allow for the evaluation of drought mitigation policies, as described next.

Evolutionary Multi-Objective Search

The complexity of the DU pathways decisions, the large number of objectives, and the breadth of modeled uncertainties motivated our use of the multi-master parallel Borg multi-objective evolutionary algorithm (MM-BorgMOEA) Hadka and Reed (2013, 2014) described in Section 2.1. Given that the Borg MOEA is a stochastic global optimizer, in both the WCU and DU optimization formulations, multiple random seed search trials of the MM-Borg MOEA are used to attain Pareto-approximate sets of policies (see equation 5.1 in Section 5.4.1). After running all the seeds, we then identify the best known Pareto-approximate sets (also termed the reference sets) by epsilon non-dominated sorting (Hadka and Reed, 2013; Laumanns et al., 2002) across all of the attained approximate sets. More details on the specific MM-Borg MOEA parametrization and parallel configuration used in this work are provided subsequently in Section 5.5. All of our subsequent results, including our more detailed robustness assessments, focus on the best known reference sets obtained for the WCU and DU sampling schemes optimization formulations.

5.4.2 Defining and Evaluating Robustness

Deep Uncertainty Re-Evaluation

The DU Re-Evaluation illustrated in Figure 5.6 defines a broader and more challenging set of SOWs that are used to assess the robustness of candidate infrastructure investment and management policies. As noted in prior studies Groves and Lempert (2007); Bankes (1993), the intent of this exploratory modeling and analysis is to shift away from predicting to instead refocus on discovering consequential scenarios. The DU Re-Evaluation sampling scheme illustrated in Figure 5.6 also provides a challenging but consistent space for comparing the performance objectives and robustness of policies found with WCU and DU optimization formulations. Scenario generation for the DU Re-Evaluation sampling scheme consists of two steps. The first step is to generate a number N_r of new synthetic inflows and evaporation-rates annual time series $(\boldsymbol{NI}_s, \boldsymbol{E}_s, \boldsymbol{UD}_s)$ following the approach presented in Kirsch et al. (2013) and correlated synthetic demands — both procedures are described in Section (2.2). The second step is to sample a matrix $\boldsymbol{\Psi}_s$ of vectors of deeply uncertain factors $\boldsymbol{\Psi}$ from distribution $\mathcal{D}_{DU}(\cdot)$. Lastly, all of time series are coupled with one vector of deeply uncertain factors $\boldsymbol{\Psi}_s$ at a time for all vectors in matrix $\boldsymbol{\Psi}_s$, as illustrated in Figure 5.6. Each Monte Carlo evaluation (N_r sets of inflow, evaporation rate and demand time series paired with the same sample of deeply uncertain factors, or one row $\boldsymbol{\Psi}$ of matrix $\boldsymbol{\Psi}_s$) returns one value of each objective and N_r pathways. The objective values of all Monte Carlo evaluations are later used to back-calculate the objective values of the entire re-evaluation exercise for a given solution, as described in the subsection below.

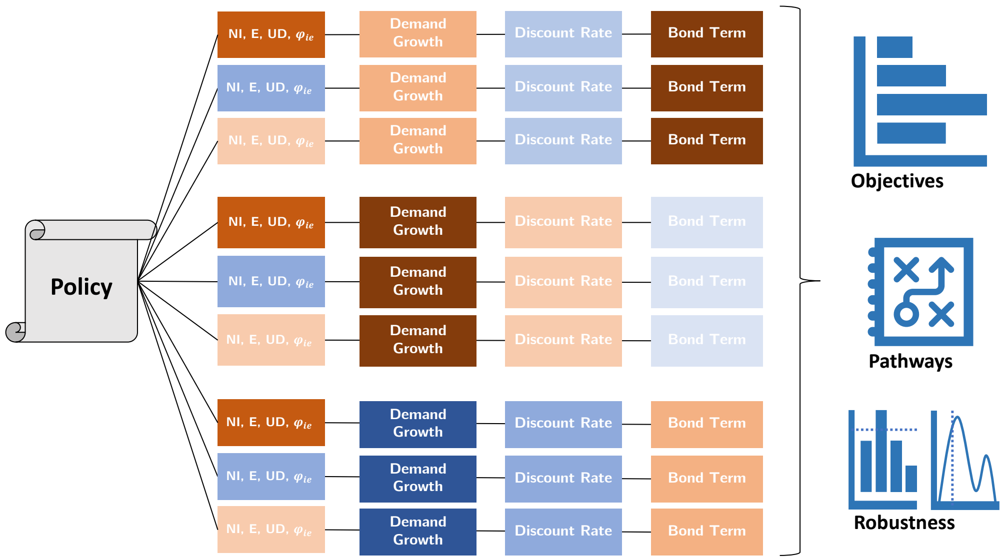

Figure 5.6: Re-evaluation sampling scheme. The same sets of N_r hydrological and demand time series are combined with one vector sample of deeply uncertain factors one at a time, resulting in a number of SOWs equal to N_r times the number of samples of deeply uncertain factors. The resulting set of objectives vectors are used to back-calculate the objectives and robustness of the whole ensemble, while the pathways are used to analyze the dependency of infrastructure investments on deeply uncertainty factors.

Objectives Estimation Based on the DU Re-Evaluation

The re-evaluation samples provide means of verifying performance by recomputing the WCU or DU optimization solutions' objectives attained from the reduced form Monte Carlo sampling schemes during search (see Figure 5.5). It should also be noted the WCU and DU optimization formulations are inconsistent with one another in terms of the number and severity of uncertainties that are included in search. We overcome this inconsistency by using the DU Re-Evaluation to verify and compare performance tradeoffs when evaluated under a consistent suite uncertainties. The DU Re-Evaluation yields an ensemble for the objectives, which are aggregated as follows:

- Reliability, Restriction Frequency, Net Present Cost of infrastructure and total Annual Financial Cost of drought mitigation costs are calculated as their means across DU re-evaluation samples.
- the Worse First Percentile Cost is calculated as the worse first percentile across the samples, mimicking the original objective as a measure of supply-enabled financial risk.
- the Jordan Lake Allocation objective does not change across samples.

The next step is to use the same ensemble of objectives used to back-calculate the overall objective values for a single policy evaluation and use it to evaluate its robustness.

Evaluating Robustness

A myriad of robustness metrics have been developed for different applications while being general enough to be applied in water systems engineering. In this chapter, we again exploit the domain criterion satisficing measure (introduced in Starr (1962), formalized in Schneller and Sphicas (1983) and used in several prior studies, such as Herman et al. (2014), Herman et al. (2015) and Lempert and Collins (2007)) used in Chapter 3. Satisficing metrics are preferred by stakeholders whose primary concern is to not violate performance constraints, despite losses of attainable performance (regret) (McPhail et al., 2018; Simon, 1959). The reader is encouraged to refer to recent works comparing approaches to robustness (Lempert and Collins, 2007; McPhail et al., 2018; Dittrich et al., 2016; Herman et al., 2015; Giuliani and Castelletti, 2016) if interested in more details.

5.4.3 Discovering Uncertain Scenarios that Control Robustness

The SOWs sampled in the DU Re-Evaluation sampling scheme illustrated in Figure 5.6 are next used for scenario discovery (Groves and Lempert, 2007; Bryant and Lempert, 2010). In scenario discovery, the analyst uses the objectives of a given policy calculated for the policy re-evaluation under different scenarios (Figure 5.6) to map regions of the space of uncertainties (or scenarios) to be sought or avoided should that policy be implemented. A classification of regions of concern would ideally include most scenarios in which the policy in question meets the conditions of performance measures of focus (high coverage) and avoid mixtures with other scenarios (high density). Scenario discovery is the step of the DU Pathways framework that allows analysts to switch from asking "what should decision-makers do?" to "what scenarios should decision-makers be concerned by if they implement candidate actions?".

Seminal studies introducing scenario discovery (Lempert et al., 2008) suggested the use of the Patient Rule Induction Method, or PRIM (Friedman and Fisher, 1999), used in Chapter 3, and Classification and Regression Trees, or CART (Breiman et al., 1984). Pass/fail boundaries found by both methodologies are hypercubes orthogonal to all the sources of uncertainties (Dalal et al., 2013). The advantage of both methods is that when mathematically appropriate they yield interpretable and relevant scenarios, as in Section 3.7.4. The main disadvantage of both methods is that they do not capture interactions between sources of uncertainties, which is often significant, preventing boundaries expressed as combinations of values for multiple sources of uncertainty. The issue of capturing interactions between uncertainties was addressed in Quinn et al. (2018) by using logistic regression to find the pass/fail boundaries. Logistic regression is a linear classifier, and therefore will return pass/fail boundaries that are straight but not necessarily orthogonal to the uncertainty axes, accounting therefore for interactions between sources of uncertainty. Logistic regression can be used to find non-linear boundaries by adding features to the original data. For example, if a data set of scenarios has two sources of uncertainty, x_1 and x_2, adding an independent term $x_1 \cdot x_2$ and training a logistic regression model on all three terms will return a non-linear model. However, there are three disadvantages for this approach: (1) output rules are not as easy to interpret as those from PRIM and CART, (2) the approximate mathematical representation of the interactions between sources of uncertainty must be pre-specified, and (3) the boundaries are necessarily smooth.

In the DU pathways framework presented in this study, the mixture of ROF-based rules for short-term management actions and long-term large discrete investments in infrastructure yields failure regions in the uncertainty space delimited by mixtures of sampled SOWs that are non-linear and multimodal. These complexities cannot be handled by logistic regression without the manual addition of several interaction terms — kernel logistic regression (Zhu and Hastie, 2005) can be used to fit such complex boundaries, but given kernel methods are non-parametric, no information about factors' sensitivities would be available. Alternatively, we propose the use of boosted trees (Drucker and Cortes, 1996; Freund et al., 1999; Schapire, 1990). Boosting is a distribution-free strategy to turn a weak learning model — models whose accuracy is only slightly better than random guessing — into a strong learning model — models that can achieve arbitrarily high accuracy (Schapire, 1990). It works by creating and ensemble of weak classifiers and forcing part of the ensemble to focus on the hard-to-learn parts of the problem, while other parts focus on the easy-to-learn parts. A boosted ensemble model therefore has the following shape:

$$H_T(\boldsymbol{x}) = \sum_{t=1}^{T} \alpha_t h_t(\boldsymbol{x}) \tag{5.12}$$

where H is the ensemble model, T is the total number of models in the ensemble, α_t is the boosting multiplier corresponding to model h_t where t is the model index, and x is the vector of explanatory variables. Model complexity can be adjusted by setting the number of models (T) in the ensemble. The boosting algorithm is described in more detail in Appendix E.

Boosting is often applied to Classification and Regression Trees (CART) trees (Breiman et al., 1984) with small maximum depth (e.g., 3 or 4), that are high bias (weak) classifiers, yielding an ensemble of trees. CART trees are comprised of nested orthogonal partitions within the explanatory variables space. The partitions are found one at a time by finding the variable and its value where a split would decrease some loss function (or leaf impurity) the most. Since in scenario discovery the analyst wants to find regions of the deep uncertainty space in which a policy will fail to reach pre-established performance criteria by any amount, classification trees are used here. Boosted trees have three advantages over the previously mentioned scenario discovery methods: (1) the approach captures non-differentiable boundaries that typical arise from threshold-based rules (e.g., ROF-based action triggers), (2) it captures non-linearity without explicitly modeling variable interactions, and (3) it is resistant to overfitting, helping assure scenario discovery maps that are simple to interpret. The use of boosted trees as opposed to random forests was motivated by the preference for simple boundaries between success/fail regions to be achieved with trees as shallow as possible. An important concept for understanding the impact of a particular uncertain factor on the performance of a policy is that of leaf impurity. A impure leaf has a mix of success and fail scenarios, as opposed to all scenarios within it corresponding to either one or the other. As mentioned, every time a partition is added to a tree by splitting a leaf into two based on an uncertainty factor at a time, it is done so decrease the impurity of the resulting leafs as much as possible. The percentage total impurity decrease across all trees due to splits on each uncertainty factor can be used as a measure of the impact of that deep uncertainty factor in the performance of a policy.

5.4.4 Visual Infrastructure Pathway Analysis

Lastly in the DU Pathways framework, after key uncertainties have been identified, we suggest detailed visual diagnostic assessments of key pathway solutions of interest to decision makers. In this study, we demonstrate the value of considering how the infrastructure investment and management pathways evolve for three scenarios: the most favorable values for the most influential uncertainties, an intermediate case of interest, and one with the least favorable values for the most influential uncertainties. The challenging part of this is step is the design of a plot that displays the pathways of all hydrologic realizations (e.g., 1,000 realizations) for a policy, while clearly showing the times at which each infrastructure option is built.

There are in the literature examples of visualizations of infrastructure pathways. The first and most well-known is the "metro maps" (Haasnoot et al., 2013; Kwakkel et al., 2016a; Zandvoort et al., 2017). The metro maps have the advantages of being easy to interpret and to present to audiences. However, they allow for the representation of only a handful of pathways before the pathways become difficult to distinguish. Another option is the probabilistic pathways Zeff et al. (2016), which was derived from the concept of metro maps (Haasnoot et al., 2013). The probabilistic pathways visualization makes use of transparency to indicate the timing and frequency that each candidate infrastructure option is implemented over an ensemble hydro-climatic scenarios. However, distinguishing the pathways differences in each of the ensemble realizations is challenging. The pathways visualization proposed in this work displays the pathways explicit temporal evolution for major investments for all hydrological realizations as a stacked ensemble plot. The resulting plots provide a visual indication of what new water supply infrastructure would be needed as function of key hydrological or other exogenous scenario gradients. This can help planners provide resources for near-future construction and for permitting, detailed design, public consultations, etc.

5.5 Computational Experiment

The computational experiment was divided into two phases: optimization and re-evaluation. Given the runtime of the Research Triangle model, we used the Texas Advanced Computing Center's Stampede 2 for both optimization and re-evaluation phases. The optimization phase was comprised of two MM-Borg MOEA (see Chapter 2.1.4) seeds trials for the WCU formulation and two random seed trials for the DU optimization. Each optimization trial was run on 256 MPI processes (1 controller, 4 master, and 251 workers) with 24 threads each on 128 2 x Intel Xeon Platinum 8160 (Skylake) nodes (48 cores per node) with 1,000 realizations per function evaluation — see schematic or realization sampling for each type of search in Figure 5.5. A total of 500,000 function evaluations were used for each trial, 125,000 per master. Our runs were evaluated for their search solution quality by confirming that further search had minimal hypervolume benefits (see Figure D.1). The resulting solutions were each re-evaluated on nodes of the same type over 1,000 sets of hydrological/demand series against 2,000 SOWs, resulting in 2,000 re-evaluation function evaluations for each re-evaluated solution—see Figure 5.6, but instead of 3 SOWs, as in the figure, 2,000 were used. All time series were generated using the methods described in Section 2.2.

The ranges of values considered for the decision variables are presented in Tables 5.3-5.5 and the ϵ values (i.e., significant precision) for all objectives are presented in Table 5.6. The MM-Borg MOEA was parameterized following prior published recommendations Hadka and Reed (2014).

Table 5.3: Decision variables pertaining to the short-term drought mitigation instruments: water-use restrictions, transfers, contingency fund and drought insurance.

Decision Variable	**Lower Bound**	**Upper Bound**
Durham restriction ROF trigger	0%	100%
OWASA restriction ROF trigger	0%	100%
Raleigh restriction ROF trigger	0%	100%
Cary restriction ROF trigger	0%	100%
Durham transfer ROF trigger	0%	100%
OWASA transfer ROF trigger	0%	100%
Raleigh transfer ROF trigger	0%	100%

Continued on next page

Table 5.3 – continued from previous page

Decision Variable	Lower Bound	Upper Bound
OWASA Jordan Lake allocation	5%	47%
Raleigh Jordan Lake allocation	0%	37%
Durham Jordan Lake allocation	5%	42%
Cary Jordan Lake allocation	35%	72%
Durham annual contingency fund contribution as percentage of annual revenue	0%	10%
OWASA annual contingency fund contribution as percentage of annual revenue	0%	10%
Raleigh annual contingency fund contribution as percentage of annual revenue	0%	10%
Cary annual contingency fund contribution as percentage of annual revenue	0%	10%
Durham insurance ROF trigger	0%	100%
OWASA insurance ROF trigger	0%	100%
Raleigh insurance ROF trigger	0%	100%
Cary insurance ROF trigger	0%	100%
Durham insurance payment as percentage of revenue	0%	2%
OWASA insurance payment as percentage of revenue	0%	2%
Raleigh insurance payment as percentage of revenue	0%	2%

Continued on next page

Table 5.3 – continued from previous page

Decision Variable	**Lower Bound**	**Upper Bound**
Cary insurance payment as percentage of revenue	0%	2%

Table 5.4: Values of the long-term rof trigger and daily demands that function as thresholds to trigger infrastructure construction by the utilities.

Decision Variable	**Lower Bound**	**Upper Bound**
Durham infrastructure construction long-term ROF trigger	0%	100%
OWASA infrastructure construction long-term ROF trigger	0%	100%
Raleigh infrastructure construction long-term ROF trigger	0%	100%
Cary infrastructure construction long-term ROF trigger	0%	100%
Daily demand Infrastructure Trigger for Cary WTP Expansion 1	0	100 MGD
Daily demand Infrastructure Trigger for Cary WTP Expansion 2	0	100 MGD

Table 5.5: The ordinal variables below determine the infrastructure construction order and adjusts them for the total number of options available to each utility. The volumetric variable represents the volume of available storage within Falls Lake to be re-allocated from the water quality to Raleigh's municipal supply pool.

Decision Variable	**Lower Bound**	**Upper Bound**
University Lake expansion ranking	1^{st}	22^{nd}
Cane creek expansion ranking	1^{st}	22^{nd}
Stone quarry reservoir expansion shallow ranking	1^{st}	22^{nd}
Stone quarry reservoir expansion deep ranking	1^{st}	22^{nd}
Teer quarry expansion ranking	1^{st}	22^{nd}
Reclaimed water ranking low	1^{st}	22^{nd}
Reclaimed water high	1^{st}	22^{nd}
Lake Michie expansion ranking low	1^{st}	22^{nd}
Lake Michie expansion ranking high	1^{st}	22^{nd}
Little River reservoir ranking	1^{st}	22^{nd}
Richland creek quarry rank	1^{st}	22^{nd}
Neuse river intake rank	1^{st}	22^{nd}
Re-allocate Falls Lake rank	1^{st}	22^{nd}
Western Jordan Lake treatment plant rank OWASA low	1^{st}	22^{nd}
Western Jordan Lake treatment plant rank OWASA high	1^{st}	22^{nd}
Western Jordan Lake treatment plant rank Durham low	1^{st}	22^{nd}

Continued on next page

Table 5.5 – continued from previous page

Decision Variable	**Lower Bound**	**Upper Bound**
Western Jordan Lake treatment plant rank Durham high	1^{st}	22^{nd}
Western Jordan Lake treatment plant rank Raleigh low	1^{st}	22^{nd}
Western Jordan Lake treatment plant rank Raleigh high	1^{st}	22^{nd}
Western Jordan Lake treatment plant OWASA fraction	1^{st}	22^{nd}
Western Jordan Lake treatment plant Durham fraction	1^{st}	22^{nd}
Western Jordan Lake treatment plant Raleigh fraction	1^{st}	22^{nd}
Volume to be re-allocated within Falls Lake for Raleigh	0MG	10,000MG

Table 5.6: Values used for ϵ-dominance.

Objective	Reliability	Restriction Frequency	Infrastructure Net Present Cost	Financial Cost	Worse First Percentile Cost	Jordan Lake Allocation
Value	0.2%	2%	$10MM	2% of AR*	1% of AR*	2.5%

5.6 Results

5.6.1 Comparing Tradeoffs After Re-evaluation

Our results compare the impacts of changing the SOW sampling strategies (Figure 5.7a) across the WCU and DU optimization formulations. The two core advantages of carefully tracking the impacts of including more uncertainties in the multiobjective search are (1) an ability to distinguish how the tradeoffs for the Research Triangle infrastructure investment and management pathways change when stressed with more diverse scenarios for the future and (2) to confirm if switching from the more traditional WCU to the approximate SOW sampling used in the DU optimization significantly impacts the robustness attained by the individual utilities as expected. The DU optimization's approximate sampling avoids the full computational cost of the comprehensive DU re-evaluation (Figure 5.6) during search. For a fair comparison, the resulting Pareto-approximate infrastructure investment and management policies found with the WCU and DU optimization schemes were each fully re-evaluated to assess their expected tradeoffs and robustness using the more comprehensive suite of SOWs captured by the DU re-evaluation ensemble.

Figure 5.7a shows the post-re-evaluation objective tradeoffs (panel "a") and the utilities relative robustness tradeoffs (panel "b") for all policies obtained through the DU and WCU optimization formulations. In Figure 5.7a, each axis represents an objective, while each line represents a regional infrastructure investment and management policy. The points where each line (policy) intersects the vertical axes represents the performance value of the metric corresponding to that axis. In Figure 5.7a, the parallel axes plot is oriented such that the ideal policy would be a horizontal line at the bottom of all axes.

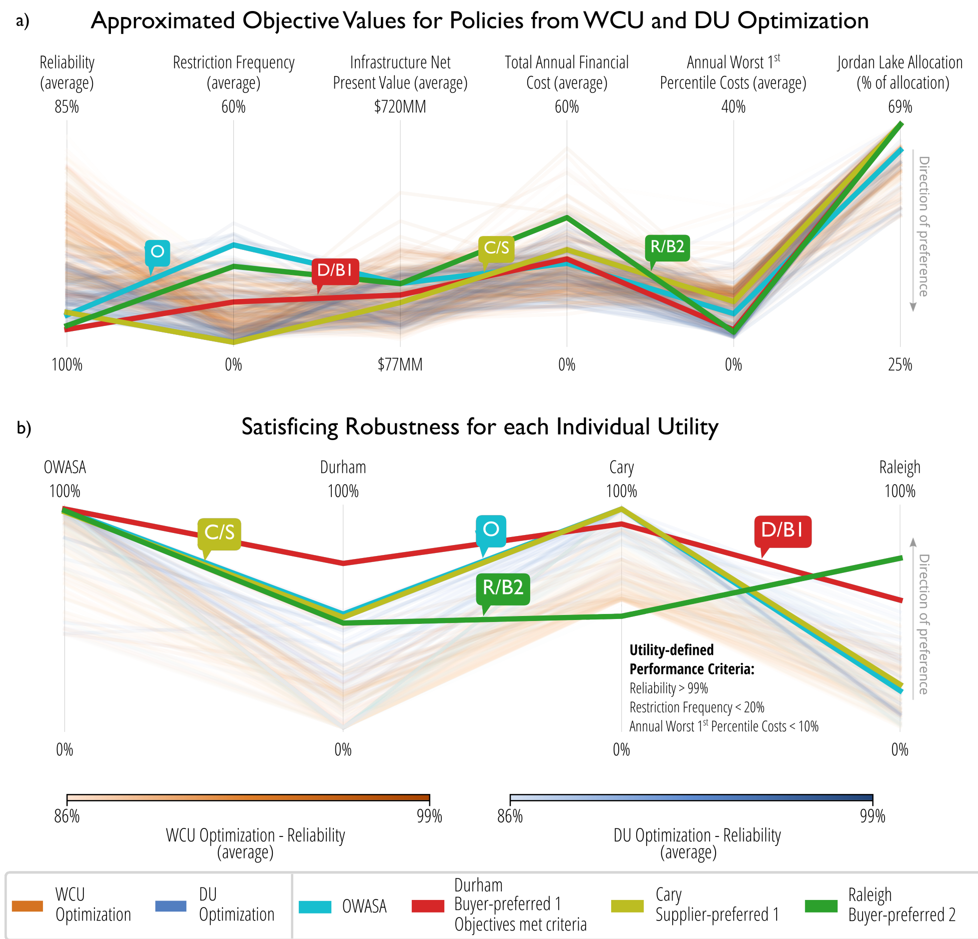

Figure 5.7: Comparative performance analyses based on the DU re-evaluation. Panel (a) illustrates the tradeoffs across the performance objectives quantified either using the mean or the worse first percentile across the tested SOWs. Panel (b) illustrates the robustness tradeoffs across the four utilities (satisficing metric). In both panels policies from the WCU formulation are brown and those from the DU optimization are blue.

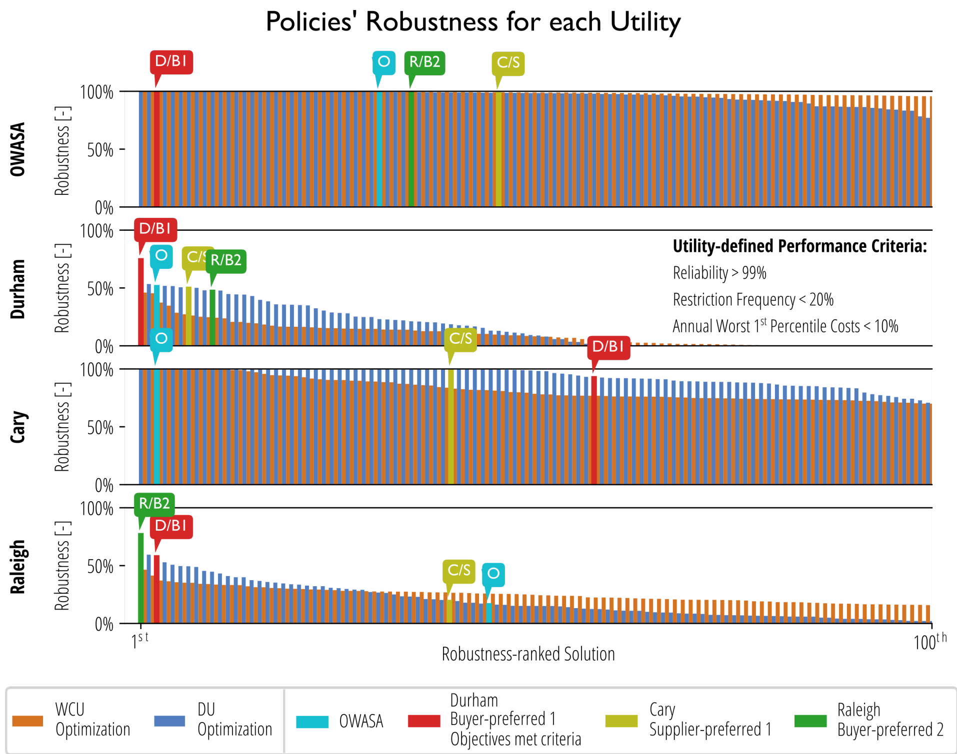

Figure 5.8: Robustness of the 100 top most robust policies attained from the WCU and DU optimization schemes for each utility. Each bar represents the attained robustness for one policy for each utility. The policies are sorted by robustness (bar height) for each utility. The policies with the highest value for the satisficing robustness for each utility is highlighted each subplot. One of the most robust policy for Cary called the "supplier-preferred" policy while the most robust policies for Durham and Raleigh were called "buyer-preferred" 1 and 2.

Figures 5.7 shows that the policies from DU optimization were on average cheaper to operate and less infrastructure intensive than their WCU optimization counterparts despite the comparable restriction frequency. The attained values of the ROF triggers for the policies discovered using the DU optimization better balance the reliability benefits and financial impacts of the short-term drought mitigation, financial instruments, and infrastructure investments while avoiding stranded infrastructure assets. Figure 5.7a shows that search under deep uncertainty served to shift policies towards higher levels of performance.

The more diverse SOWs in the DU optimization appear to sufficiently stress the management and investment policies to force them to more effectively use the full suite of water supply portfolio options at the utilities' disposal. Figure 5.7a also highlights strong tensions between reliability and restriction frequency across the WCU and DU optimization formulations. Likewise, significant tradeoffs exist between Reliability and Total Financial Cost, meaning that if a utility wants high reliability it has to pay for it through a contingency fund and/or drought insurance. Lastly, Figure 5.7a highlights a tension between average annual cost and annual worst-first-percentile cost, which is important as the utilities often struggle when confronted with significant variability in their annual costs and are often willing to pay a premium for more predictable costs.

Transitioning to robustness tradeoffs, Figure 5.7b uses a similar parallel axes plot to display the percentages of DU re-evaluation SOWs where each utility meets their goal performance requirements: reliability $> 99\%$, restriction frequency $< 20\%$ and annual worst-first-percentile cost $< 10\%$. In terms of ideal performance for robustness in Figure 5.7b, the direction of preference is upward, where the ideal policy would be represented to a line horizontally intersecting each of the utilities robustness axes at their top maximum values (i.e., 100%). Figure 5.7b shows that strong robustness tradeoffs exist between Raleigh, Cary, and Durham. In contrast, OWASA attains close to maximum robustness for almost all policies. The specific policies that represent the most robust policies for each utility are highlighted in different colors in Figure 5.7b. Figure 5.8 supplements Figure 5.7b by more directly comparing the utilities' satisficing robustness for each of the policies found by the WCU and DU optimization formulations. To clarify our analysis, we designate the most robust policies for Durham and Raleigh, respectively, as buyer-preferred policies D/B1 and R/B2, given their dependence on buying water transfers from Cary. The most robust policy for Cary is designated the supplier-preferred policy, or C/S in Figures 5.7 and 5.8. Figure 5.7a shows that the most robust policy for Durham (D/B1) was the only policy whose DU re-evaluation objectives met the performance criteria established by the utilities: reliability $>$ 99%, restriction frequency $<$ 20% and annual worst-first-percentile cost $<$ 10%. Even though WCU optimization found 246 Pareto approximate policies versus 112 from DU optimization, only those policies attained by including deeply uncertain factors in search attained both high levels of objective performance and robustness against deeply uncertain scenarios (DU Re-evaluation). Also, the most robust management and investment policies for each individual utility in the Research Triangle were obtained through DU optimization.

Overall, Figures 5.7 and 5.8 show that the policies from DU optimization were on average more reliable, less infrastructure intensive, financially cheaper and less risky despite slightly higher restriction frequency than their WCU optimization counterparts when both sets of policies are exposed to broader uncertainty. The policies resulting from the DU optimization formulation have combinations of ROF triggers that allow the system to respond faster and more cost-efficiently to drought events, avoiding costly stranded infrastructure and high capital debt levels while maintaining reliability. Even if the reliability performance requirement is relaxed to 97%, only 4 policies from WCU optimization would meet the criteria versus 34 from DU optimization.

5.6.2 Scenario Discovery for Compromise Policies

Building on a better understanding of the performance tradeoffs and robustness conflicts as illustrated in Figures 5.7 and 5.8, it is also important to understand what deeply uncertain factors most strongly influence robustness. Figure 5.9 shows scenario discovery maps for Durham, Cary and Raleigh. Each map shows the performance attained across sampled scenarios for the two most important uncertainties for each utility when the Research Triangle implements one of the two buyer preferred policies highlighted in Figures 5.7b and 5.8. In each panel of Figure 5.9, the two deep uncertain factors plotted on the horizontal and vertical axes are the factors identified by the boosted trees algorithm to be the most dominant in shaping robustness performance for a given policy. Such dominance is quantified by the percentages in parentheses next to each factor label, which indicate the decrease in impurity of the tree ensemble from splits on that factor. As summarized in section 5.4.3, significant changes in the percentage of impurities from tree-based splits of a factor represents the direct reduction of classification errors in terms of which sampled SOWs attain success or failure. High percentage changes in impurities consequently can be interpreted as a measure of utilities' performance sensitivity to a given factor. Table C in the Supplement Section C reports the impurity decreases (or sensitivities) for all of the sampled factors for Durham, Cary and Raleigh under both buyer preferred policies D/B1 and R/B2. Red regions in the factor maps of Figure 5.9 represent regions in which a utility will likely not meet its performance targets, grey regions represent high performance scenarios, and white regions are deemed inconclusive due to close proximity mixtures of failures and successes in a given region of the scenario space. In each of the six panels of Figure 5.9, the stars represent the most favorable scenario (grey star), the least favorable (red), and a projected future where demand growth rates in the Research Triangle region are 80% of the nominal values assumed in the WCU optimization values (blue star).

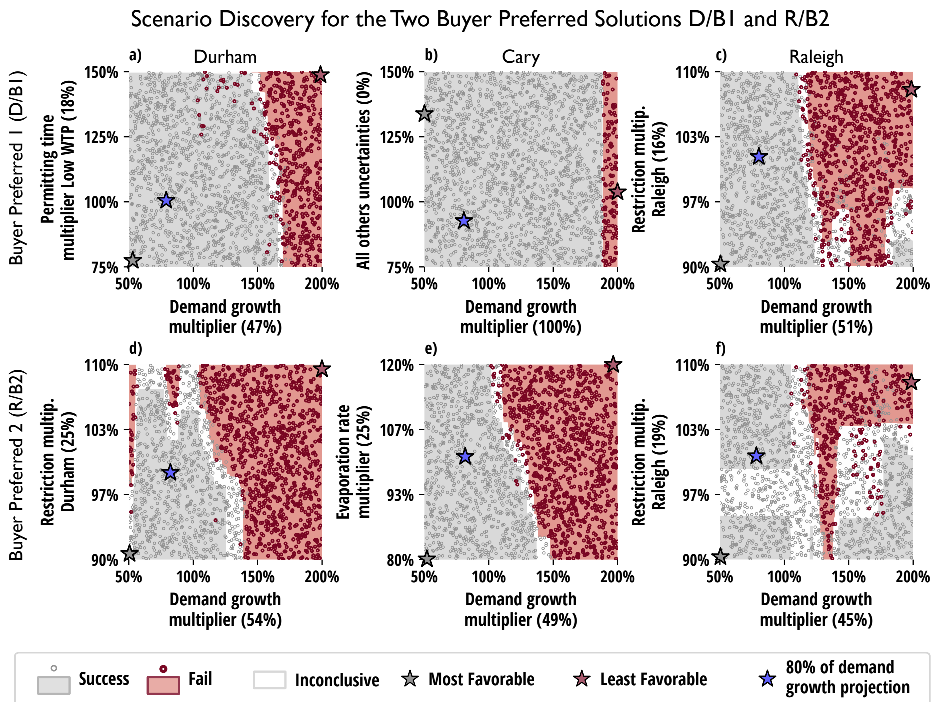

Figure 5.9: Factor maps for two policies — buyer preferred 1 (Durham) and 2 (Raleigh) — showing the key factors controlling the robustness the solutions for all utilities but OWASA. Success required reliability $\geq$ 99%, restriction frequency $\leq$ 20% and annual financial worst 1st percentile cost $\leq$ 10%. A list with the combined impurity decreases for all factors can be found in Table C.

Figures 5.9a-5.9c show the scenario discovery maps obtained for Durham, Cary and Raleigh, respectively, under Durham's preferred policy D/B1. Figure 5.9a shows that Durham's failures are mostly contingent on demand growth. The same figure also shows that the D/B1 management and investment policy effectively hedges Durham's robustness against demand growth, requiring a demand growth rate more than 50% over the nominal rate assumed by Durham before failing to meet the utility's performance requirements. However, if demand growth rates are 50% to 70% above the projected values, Durham may still meet the performance criteria depending on how many years the utilities take to get the permits for the low capacity Joint WTP. Figure 5.9b shows that Cary's failures under policy D/B1 are contingent on extreme demand growth rates exceeding their nominal projections by 85%. Cary's failures are due to a hard treatment capacity boundary where demand would surpass the maximum upgraded treatment capacity for Cary's own water treatment plant on the Jordan Lake.

Figure 5.9c, on the other hand, shows that Raleigh has a far more complex failure dynamic. Figure 5.9c shows that the most dominant factor influencing Raleigh robustness is demand growth, though Raleigh's performance does not decrease monotonically as demand growth increases. Rather, Raleigh's displays poor performance at demand growths between 50% and 80% higher than the nominal projected values, but at twice the projected values Raleigh meet the performance criteria again. The reason for the stratification of failure scenarios in Figure 5.9c is because scenarios in the red vertical strips emerge due to poor coordination between infrastructure construction and the effective use of short-term drought mitigation instruments. Figure 5.9c shows that in low demand growth scenarios, Raleigh builds storage infrastructure slowly to avoid costly stranded assets while supply risks are managed effectively by short-term mitigation instruments. On demand growth scenarios greater than Raleigh's nominal projections, Raleigh fails to meet the performance goals (red dots) when new storage infrastructure is not built fast enough to keep up with demand growth rate, straining the short-term drought mitigation instruments. In Figure 5.9c, the complex non-convex fingering of failure or inconclusive scenarios result when the D/B1 investment long-term ROF rules do not trigger sufficiently rapid capacity expansion and the short-term ROFs are not capable of mitigating the demand growth rates (i.e., cases of slow recognition of growing risks). The complex non-linear fingering of the failures captured in several panels of Figure 5.9 are an interesting result that provides a general insight, as this behavior could occur for any kind of decision-making metric based on thresholds, such as the ROF or days-of-supply-remaining. Geometrically, this allows a utility to more explicitly recognize the scenarios their investment and management rules do not perform as wished.

Transitioning to the Raleigh preferred buyer preferred solution (R/B2) in Figures 5.9d - 5.9f, the demand growth rate remains a dominant factor shaping the robustness performance for all of the utilities. However, Figure 5.9d shows that the effectiveness of Durham's water use restrictions emerges as the second most important factor influencing Durham's performance. Switching to the R/B2 investment and management policy yields an appreciable increase in the number and complexity of failure scenarios for Durham. The underlying mechanisms that yield the complex fingering of failure scenarios in Figure 5.9d for Durham are similar to those for Raleigh under the D/B1 policy in Figure 5.9d: poor coordination between infrastructure construction and the short-term use of drought mitigation instruments. For the R/B2 policy, Durham has far less hedging to robustness failures and is specifically more sensitive to the effectiveness of their populations response to restrictions.

Transitioning to Cary, Figure 5.9e shows that Cary's failure region is far larger under policy R/B2 than under D/B1. Interestingly, evaporation emerges as an important uncertain factor for Cary instead of for Raleigh, despite Falls lake (Raleigh's main supply source) being rather shallow. The emergence of evaporation as key concern for Cary would imply that longer term warming from climate change could pose a severe challenge to the R/B2 policy. To our knowledge, these results are the first to show the complex emergence of vastly different failure modes and asymmetric sensitivities that can emerge for multi-utility regionalization of investment and management pathways.

Lastly, Figure 5.9f shows that the R/B2 policy strongly hedges Raleigh while degrading the robustness of both Cary and Durham as shown for the D/B1 policy. Regional demand growth rate and the response of Raleigh's customers to restrictions dominate its robustness performance. Figure 5.9f again shows a complex fingering of failure scenarios as has been noted above where slow recognition of the need to expand capacity when demand growth rates exceed nominal projections by approximately 40%. Moderate losses in the effectiveness of restrictions serves to compound Raleigh's robustness failures. The consistent importance of Raleigh's water use restriction effectiveness across both the D/B1 and R/B2 policies stresses the importance of Raleigh's readiness to effectively implement water-use restrictions and make sure their population will be responsive.

Overall, Figure 5.9 provides some important regional insights for the Research Triangle. The dominant impact of demand growth rates emphasizes that supply reliability is harder to achieve than financial goals for this system. Overall, keeping demand growth rates at or below the utilities' projected value would be an important step for all of the utilities to ensure satisfactory regional performance. Furthermore, in most cases the reductions by 20% of the projected growth (suggested by the utilities themselves) would significantly improve satisfactory performance while also adding a safety margin that decreases the region's need for new infrastructure.

Figure 5.9 also contributes some general insights regarding the value of using boosted trees to guide the scenario discovery process for complex infrastructure pathway problems. Boosted trees succeeded in capturing complex and distinctive features of the factor mapping plots without inherent overfitting. Most non-linearities and fingering within complex transition zones between successful and failed scenarios were identified by the boosted trees approach. In addition, the identification of inconclusive regions between passing and failing scenario regions are helpful for avoiding errors in complex scenario discovery contexts.

5.6.3 Policy Rules in the Robustness Compromises

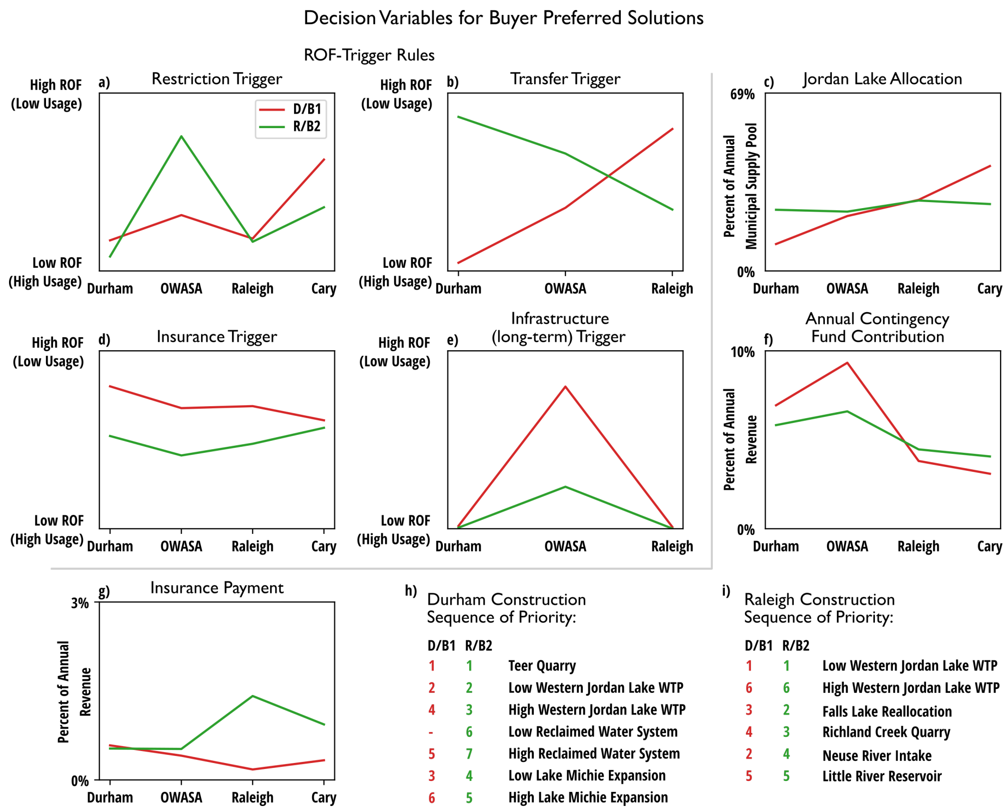

Figure 5.10: Parallel axis plots for the decision variables for the D/B1 and R/B2 buyer preferred policies. The red lines designate Durham's preferred buyer solution (D/B1). The green lines designate Raleigh's preferred buyer solution (R/B2).

Figure 5.10 provides a more detailed illustration of the underlying decisions or rules that compose the Durham (D/B1) and Raleigh (R/B2) preferred policies. Recall that the investment and management policies are comprised of: (1) ROF triggers for short-term mitigation instruments (Figures 5.10a, 5.10b, and 5.10d) and for long-term infrastructure investments (Figure 5.10e), (2) allocations to the water supply pool in Jordan Lake allocation (Figure 5.10c), (3) annual contingency fund contributions and insurance payouts defined as percentages of annual revenue (Figure 5.10f and 5.10g), and (4) infrastructure construction order (Figure 5.10h and 5.10i). In both the D/B1 and the R/B2 policies, there is significant similarity between Durham and Raleigh in their high reliance on restrictions (Figure 5.10a), on their prompt triggering of new infrastructure by utilizing low long-term ROF triggers (Figure 5.10d), and in their moderate-to-high use of contingency funds ((Figure 5.10f). Figure 5.10a shows that Durham makes consistently high use of transfers under Durham's preferred policy (D/B1) by setting the transfer trigger value low. This in turn requires a high annual contingency fund contribution as highlighted in Figure 5.10f. Figures 5.10g and 5.10e show that Durham uses small insurance payouts in both policies at moderate and high ROFs, which emphasizes the preferred strategy of relying dominantly on a contingency fund.

In comparing Raleigh versus Durham's preferred policies, Figure 5.10b highlights competition for transfers. This partly explains the robustness conflicts between Durham and Raleigh, given their shared use of limited conveyance capacity for transfers from Cary. In Durham's preferred robustness compromise policy (D/B1), Durham consistently uses the full capacity of the intra-utility transfer pipes from 2015 until 2027. Alternatively, in Raleigh's preferred solution (R/B2), Durham's limited use of transfers allowed Raleigh to request medium-to-high volumes of transfers whenever needed (see supplementary Figures B.1 to B.4). These differences are reflected in each utility's approach to their finances because although transfers are an effective way of decreasing the need for restrictions, they can be costly. Under policy R/B2 Raleigh opts for moderate payouts at moderate ROFs to hedge against the higher use of transfers because of the lower associated ROF trigger (Figure 5.10b). In fact, Raleigh's sparse use of transfers under policy R/B2 warrant the use of drought insurance in contrast to policy D/B1. Most insurance payouts for Raleigh under R/B2 happen during periods of high use of transfers (see supplementary Figures B.3 and B.4). Raleigh's use of financial insurance in the R/B2 policy serves to reduce the political risks and opportunity costs associated with a large and rarely used contingency fund.

5.6.4 Infrastructure Pathway Ensemble

The Durham preferred policy (D/B1) and the Raleigh preferred policy (R/B2) capture alternative strategies for navigating the performance and robustness conflicts highlighted in Figures 5.7 and 5.8 and each lead to distinctly different future pathways for the Research Triangle. Figure 5.11 contributes a detailed analysis of the capacity expansions and ROF dynamics for Durham and Raleigh that emerge from the D/B1 policy. The plots in Figure 5.11 focus on the resultant capacity expansion pathways and ROF time series corresponding to the three highlighted scenarios from Figure 5.9: the Most Favorable SOW (left column), the SOW with 80% of the Nominal Demand Growth Rate (middle column), and the Least Favorable SOW (right column) for Durham and Raleigh. Figures 5.11a - 5.11c plot Durham's storage capacity and short-term ROF dynamics for each of the three demand growth rate scenarios. Likewise, Figures 5.11d - 5.11f plot Durham's resulting ensemble of infrastructure pathways across the demand growth scenarios. Figures 5.11d - 5.11f show the infrastructure pathways for the same scenarios. The horizontal axes of each pathway panel represents time from 2015 to 2060, encompassing planning horizon of 45 years. All of our illustrated infrastructure pathway plots display stacks of 1,000 horizontal multicolored lines, each line corresponding to the investment sequences for one hydrological realization. For each realization line, the colors represent the infrastructure project that was built last in time. For example, the dark green horizontal line segment close to the center of the plot indicates that in that real-

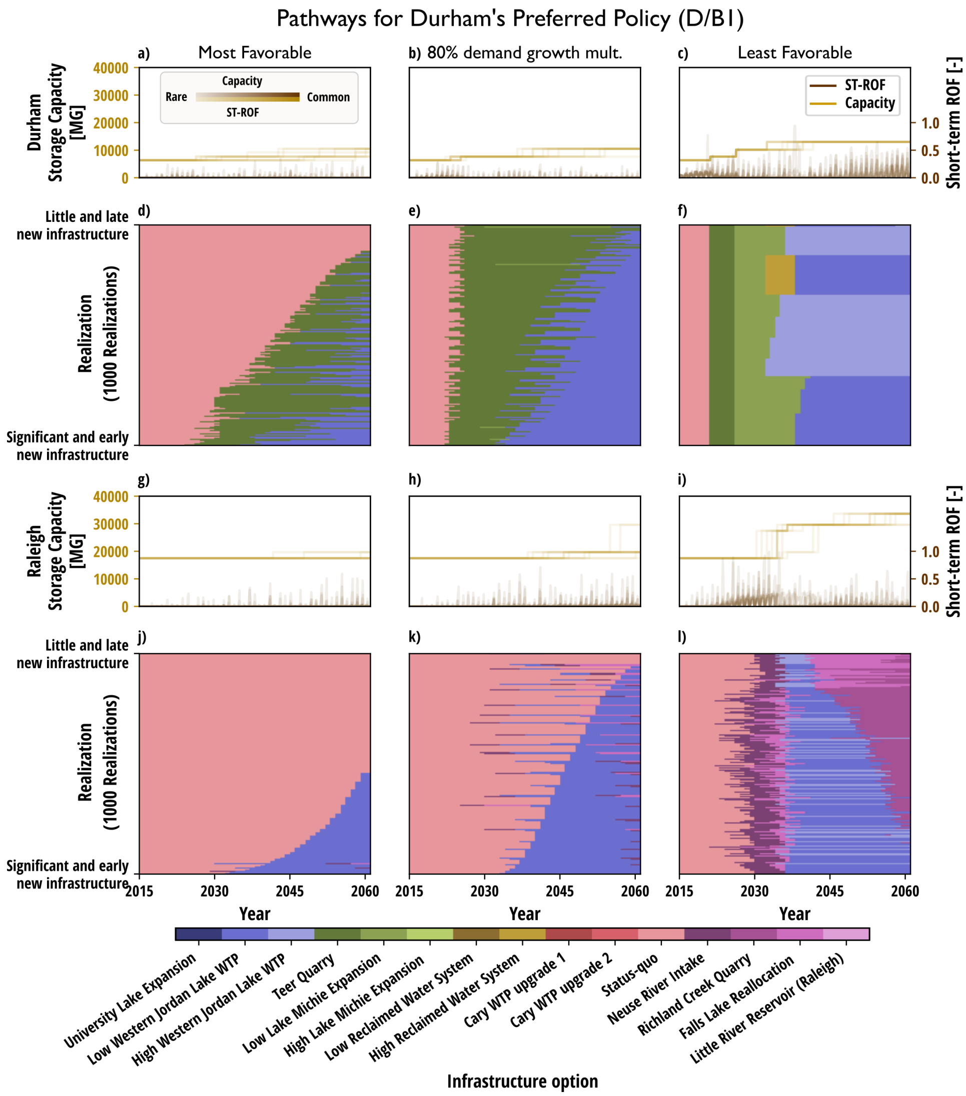

Figure 5.11: Capacity expansion pathways and ROF dynamics for the Durham preferred policy DB/1. Demand growth rate scenarios considered include the Most Favorable SOW, the SOW with demand growth rates reduced to 80% of their original nominal values, and the Least Favorable SOW. All infrastructure investments are illustrated as stacked color lines designating specific infrastructure investments across 1,000 hydrologic scenarios.

ization indicates that Durham builds the Teer Quarry typically between 2040 to 2050.

An important concern captured in this study relates to the importance of permitting times, which vary significantly across different infrastructure investments. In the pathways illustrated in Figure 5.11, the infrastructure construction order for a policy may not match the order in which infrastructure is built over the course of a simulation due to permitting times. This issue frequently emerged when considering investments in the Joint WTP for Durham under all three scenarios (Figures 5.11d - 5.11f) and for Raleigh under Raleigh's least favorable scenario for policy D/B1 (Figure 5.11l). The infrastructure construction order for policy D/B1 in Figure 5.10h shows that the low capacity version of the Joint WTP was ranked first and second in the infrastructure construction queues of Raleigh and Durham under policy D/B1. However, the colors between the salmon (status-quo) and light or medium blue (low and high capacity Joint WTP) in Figures 5.11d through 5.11f and 5.11l show that the Joint WTP was not the first option built by either utility under the considered demand growth rate and hydrologic scenarios. The reason for the apparent incongruence is that the permitting and planning period of the Joint WTP is sufficiently long as to allow other options — namely Teer Quarry Reservoir, Low Capacity Lake Michie expansion, Falls Lake reallocation and in a few cases Richland Creek quarry and reclaimed water — to be built first by Durham and Raleigh.

A more detailed analysis of the pathways resulting from Durham's preferred policy D/B1 shows that the Joint WTP expansion and co-investment is important and dominantly triggered by Durham. Figures 5.11 and B.1 shows for even the Least Favorable SOW that Durham does not need to invest in any other projects in the queue after either version of the Joint WTP is built. The water treatment plant investment successfully reduces Durham's long-term ROF to 0 after its construction (see supplementary Figures B.1 and B.2). Transitioning to Raleigh, despite having to still build one more project in the Least Favorable SOW (Figure (5)l), the short-term ROF in Figure 5.11i noticeably decreases with the construction of the Joint WTP. This reduces the pressure on Raleigh's short-term drought mitigation instruments. Despite the Joint WTP being a fairly expensive project when compared to other infrastructure alternatives, its shared cost makes it effective reliability- and financially-wise for both Durham and Raleigh because under both policies both utilities have a substantial allocation at Jordan Lake (see Figure 5.10c).

Aside from the Joint WTP, most of the focus of Durham and Raleigh under policy D/B1 is on storage infrastructure. This conclusion follows from the high prevalence of storage infrastructure seen in all of the pathways (Durham — Figures 5.11d - 5.11f and Raleigh — Figures 5.11j - 5.11l). The investments in storage come at the expense of investments in reclaimed water and intakes, which are comparatively cheaper but not as effective at mitigating supply shortfalls. The only appearance of reclaimed water infrastructure in the pathways occurred for Durham under their Least Favorable SOW (Figure 5.11f) largely as a consequence of construction times for the other water sources. Lastly, it is also interesting to observe that even in their Least Favorable SOW (Figure 5.11l), Raleigh goes through substantial storage capacity growth without building a single reservoir or reservoir expansion. Instead, increases in their supply capacity result from direct access to its Jordan lake water supply pool via the Joint WTP and through the Falls Lake's municipal pool relocation, which are cheaper and less disruptive projects than storage infrastructure.

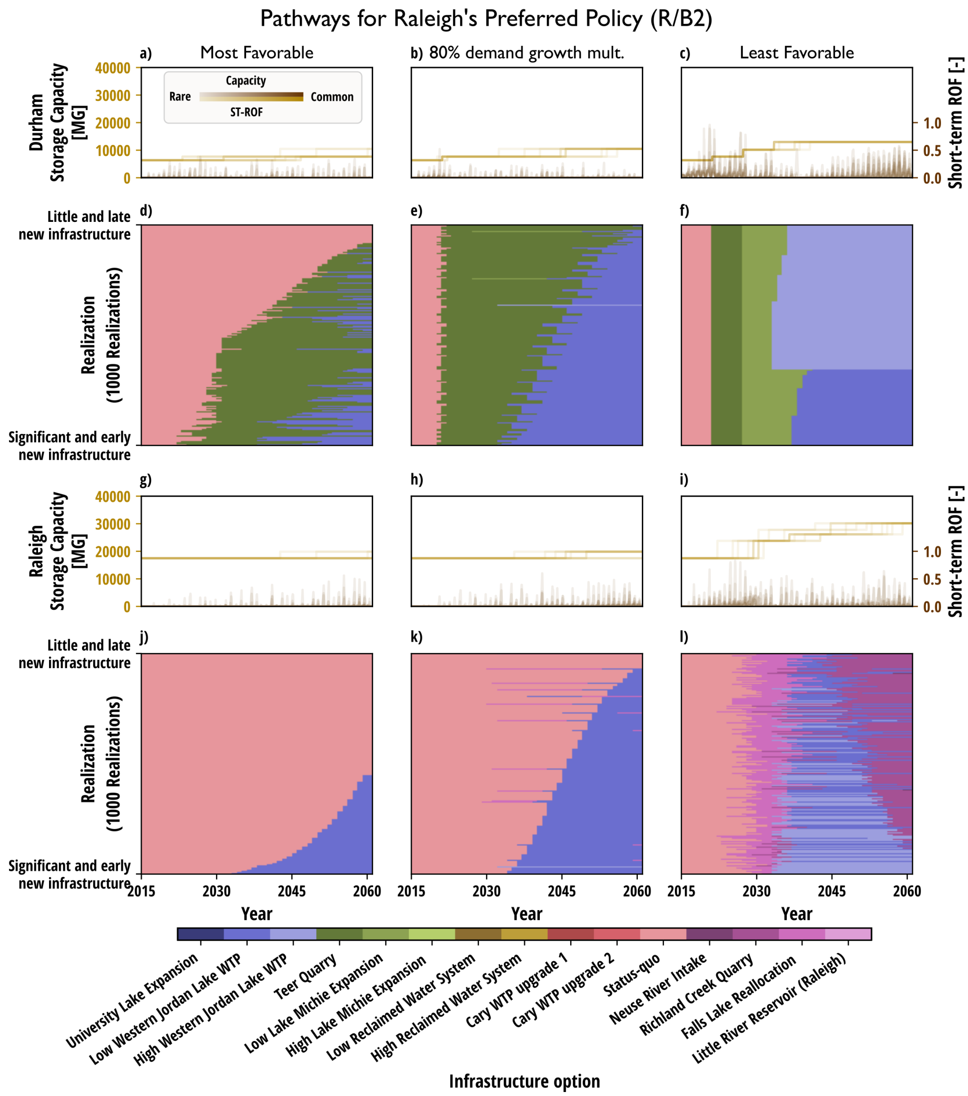

Figure 5.12: Capacity expansion pathways and ROF dynamics for the Durham preferred policy DB/1. Demand growth rate scenarios considered include the Most Favorable SOW, the SOW with demand growth rates reduced to 80% of their original nominal values, and the Least Favorable SOW. All infrastructure investments are illustrated as stacked color lines designating specific infrastructure investments across 1,000 hydrologic scenarios.

Transitioning to Raleigh's preferred buyer policy R/B2, Figure 5.12 shows the ROF dynamics and infrastructure pathways for Durham and Raleigh under the same three demand growth rate scenarios displayed in Figure 5.11. Although Figures 5.11 and 5.12 show many similarities in the resulting infrastructure investment pathways, key differences that favor Raleigh do emerge for the R/B2 policy beyond the utilities competition for transfer water discussed above. Under policy D/B1, as illustrated in Figure 5.11, it is always Durham who triggers either version of the Joint WTP. However, for the policy R/B2 under Raleigh's Least Favorable SOW, the low capacity version of the Joint WTP is triggered 60% of the time by Raleigh. This typically happens when Raleigh finishes implementing the Falls lake reallocation and Durham has the low capacity version of the Lake Michie expansion under construction. The Joint WTP remains as an effective investment to reduce the ROFs of both Durham and Raleigh under the R/B2 policy.

Differences in the D/B1 and R/B2 policies are most pronounced for high demand growth scenarios. Recall from our scenario discovery results in Figure 5.9 that short-term use and competition for water transfers yields strong differences in Durham's and Raleigh's robustness (and failure mechanisms). This results despite the similarity of the investment pathways across both policies shown in Figures 5.11 and 5.12. This further stresses that sound infrastructure investment strategies must be complemented by a sound short-term drought mitigation instruments for performance levels to be maintained under extraneous scenarios. More specifically, effective use of restrictions and transfers strategies coupled with financial instruments to hedge utilities serve to reduce and delay for years the need for building new infrastructure projects. Still, despite a loss in performance, under policy R/B2 a combination of an earlier Falls Lake reallocation and effective use of transfers before the construction of the second infrastructure option allowed for a more gradual system expansion, resulting in less new built infrastructure (i.e., the Neuse River intake was not built and the final storage capacity for the system was about 3,000 MG smaller).

Despite the mentioned differences in the pathways between both solutions shown in Figures 5.11 and 5.12, their similarities (similar timings and options built) imply general insights for the Research Triangle region. Short-term drought mitigation instruments are critical to the evolution of the system's infrastructure investment pathways. Regionally coordinated strategies for managing demand growth rates are an important risk hedge for all four of the Research Triangle utilities. Overall, an important take-away point from the pathways plots for Raleigh and Durham is that the Teer Quarry and the low-capacity Joint WTP projects are very likely to be needed in the foreseeable future. Therefore, Durham, Raleigh and Cary should include the Joint WTP in their budgeting process, perform more detailed design analysis and contemplate their permitting requirements. The main recommendation for Raleigh would be to initiate an application for a reallocation of the Falls Lake's water quality pool to its own supply pool. More broadly, DU optimization significantly reshaped system objectives tradeoffs, robustness, scenario discovery results, portfolio mix of decisions, and the corresponding pathways. The analyses framework as illustrated in our results provides a rich context for an understanding of the system.

5.7 Conclusion

This work contributes the DU Pathways framework, which includes deep uncertainty in the search phase of the infrastructure pathways problem. In the context of water resources systems, this work demonstrates the value of integrated regional infrastructure planning and management coupled with deep uncertainty analysis for mitigating drought risks and their corresponding financial impacts. Such risks are even more prevalent in areas of high population growth, such as the Research Triangle region in North Carolina. The Research Triangle regional system demonstrates how performance and robustness conflicts can emerge between geographically close utilities in their long planning.

The DU Pathways framework's flexibility in incorporating and exploring uncertainties as demonstrated in this study highlights some general insights and benefits for regional water utilities seeking to coordinate their short-term management actions (weekly) and long-term (annual) investment decisions. A key benefit is the ability to create more flexibility and reduce required infrastructure debt burden by carefully balancing short-term decision making and long-term infrastructure investment decisions. The consistent integration of both management and investment decisions using ROF-based rules provides operators with ways to account for information beyond measurements of reservoir levels and past data, allowing for decisions based on possible future conditions as well. Another major benefit is the possibility of performing uncertainty analysis on the effectiveness of alternative infrastructure investment pathways for individual utilities as well as broader cooperating regional coalitions. This analysis not only informs utilities about what uncertainties to monitor when planning the early studies preceding the triggering of new infrastructure, but also about what investments would be more effective as the results from monitoring emerge. Lastly, the DU Pathways framework allowed for a clear understanding of strong interdependencies between short-term drought mitigation and financial instruments as well as their potential to determine the proper timing of investments in new infrastructure. The DU Pathways framework centralizes and promotes all these benefits for single and regional water utility infrastructure planning and management.

The inclusion of deep uncertainties in the search-based identification water supply management and investment portfolios yielded pathway policies that are more robust (according to performance criteria defined by the utilities themselves) and that also provide better performance and robustness compromises between the four utilities in the Research Triangle Region. The scenario discovery analysis showed that the key driver of system performance and robustness is demand growth, which utilities can act on through coordinated demand management. The infrastructure pathways analysis also showed that demand growth may determine whether specific investments in infrastructure will be needed by a given point in time (should the right conditions be observed), how key uncertainties may drive the need for investments to be sooner or later, or if some candidate projects would never be prioritized at all even under the most adverse futures. This information is crucial for utilities to design their rate structure to accommodate for debt servicing for the maintenance of their credit rating when issuing the required bonds. Future work could consider a wider variety of drought mitigation instruments, new strategies for sampling and potentially screening deep uncertainties, and short-term drought mitigation policies that adapt their triggers with the construction of new infrastructure (a meta-policy, so to speak). The insights of this work have broad merit for water utilities facing growing pressures to better coordinate management and investment decisions to confront resource contention, growing urban demands, and increasingly extreme droughts.

CHAPTER 6

CONTRIBUTIONS AND FUTURE WORK

6.1 Contributions and Conclusions

As discussed in Chapters 3 and 4, water utilities are increasingly challenged to balance the burdens from required new investments in infrastructure with innovating their management of existing resources to confront pressures from growing populations, changes in climate, evolving environmental regulations, growing financial risks, and many other factors. This dissertation seeks to aid utilities in confronting these challenges through three key contributions. First, the Deep Uncertainty (DU) Pathways framework formalizes uncertainty-driven infrastructure planning and management by bridging stochastic adaptive infrastructure pathways and multi-objective robust decision making (MORDM). Secondly, this dissertation contributes the open source WaterPaths model for simulating planning and management decision making by water utilities and designed to enable the seamless application of the DU Pathways approach to any single or regionally coordinated group of water utilities. The third major contribution of this dissertation is the Sedento Valley hypothetical test case that can serve as a shared training and diagnostic test bed to promote future innovations in water infrastructure planning and management frameworks. These three contributions were motivated under the unifying and specific goal of improving decision-making and policy design for water-infrastructure systems with multiple stakeholders and conflicting objectives. These contributions are discussed in more detail below.

(1) Enabling a more sophisticated incorporation of uncertainty in infrastructure invest-

ment pathways that are carefully coordinated with short-term water portfolio management instruments yields regional actions that are more robust and financially stable.

Given the long lives, high costs, and irreversibility of investments in public infrastructure assets, decisions to pursue specific options require careful thought to assure the population's needs are met and that the new asset does not become stranded (Grant et al., 2013). These risks can be mitigated by (1) performing careful uncertainty analysis (Nair and Howlett, 2017), (2) by presenting all stakeholders with multiple options, all in accordance with the specific institutional backgrounds, to foster a high-level political debate (Bosomworth et al., 2017), by (3) using short-term management strategies tailored to the resulting infrastructure system and vice-versa (Zeff et al., 2016), and by (4) using regional planning including various municipalities impacted and benefitted by the asset. This dissertation contributes the DU Pathways framework, an extension the MORDM framework that includes strategies for sampling deeply uncertain factors during the computational policy search as a way of assuring the discovery of regional infrastructure planning and management policies that are robust to unforeseen conditions (Chapters 3 and 5). To improve compliance with the timelines of utilities' operations and local political processes, the DU Pathways framework simultaneously employs weekly and annual rule-based decision-making based on the ROF metric for decisions regarding short-term drought management and infrastructure planning, respectively.

This dissertation has broadened the scope of uncertainties that can be sampled during automated infrastructure planning and management policy search as well as the robustness assessments of the discovered policies. Chapter 3 introduced the concept of DU uncertainty sampling and used it to find a drought-management policy for the Research Triangle test case. Such policies have higher robustness to deeply uncertain factors for all stakeholders, fewer conflicts over robustness tradeoffs among stakeholders, and that nonetheless exceeds performance expectations under projected values for all analyzed deeply uncertain factors without being explicitly optimized to do so. Chapter 5 built on these findings to extend the framework proposed by (Zeff et al., 2016) and achieve the goal of finding infrastructure planning and management policies for the Research Triangle test case that are robust and adaptable to a wide variety of uncertainties, minimizing the risk of stranded-assets and capacity shortfalls. Furthermore, the resulting policies are designed to work with drought mitigation measures trusted by utilities and accepted by the public, as well as to respect the nature and pace of local political and institutional processes. In addition, the Boosted-Trees-based scenario discovery methodology introduced in Chapter 5 filled the last gap regarding the applicability of scenario discovery theory, namely that in which the uncertainty-performance relationships cannot be fully explained by models with pre-specified functional forms. By incorporating uncertainty sampling in all stages of the MORDM framework and improving later uncertainty analysis, this dissertation seeks to promote water-infrastructure supply and financial modeling under uncertainty as a way to plan and operate water infrastructure with greater supply reliability and financial stability.

(2) The WaterPaths model addresses limitations in existing software for modeling water infrastructure by incorporating stochastic uncertainty analysis, infrastructure expansions, and regional infrastructure planning and management.

Despite the demonstrated benefits of the DU Pathways framework for the formulation of regional water infrastructure planning and management policies, performing all its steps with existing simulations models would not be possible because the existing simulation tools lack needed features. More specifically, the existing models are not designed for stochastic uncertainty analysis or for evaluating broad categories of objectives that bridge water supply reliability as well as financial concerns. Existing simulation frameworks lack out-of-the-box integration with multiobjective optimization algorithms and would not be capable of representing diverse infrastructure investment or water portfolio management actions. These concerns are particularly acute for multi-stakeholder regionally coordinated operations and planning and for joint supply and financial modeling. Lastly, large scale computational search or ensemble-based exploratory modeling would not be tractable with current simulation frameworks as they lack functionality on cloud and high-performance-computing platforms and often do not provide open-source access to their underlying codes. Chapter 4 demonstrates how WaterPaths fills these capability gaps with a risk-based planning and management methodology, while Chapter 5 shows its first successful real-world application in the Research Triangle, NC. WaterPaths offers regions with interconnected utilities the possibility of a full DU Pathways analysis with minimal model setup time. Furthermore, as an open-source code WaterPaths is designed for ease of modification and extension. The simulation framework provides researchers and practitioners with a test bed for novel drought mitigation instruments, types of infrastructure projects, decision-making metrics, and uncertainty sampling techniques. Currently, multiple researchers are employing and extending WaterPaths to study systems in the US, Brazil, and South Korea.

(3) The Sedento Valley test case contributes a highly realistic and challenging hypothet-

ical system that can support educational training efforts, algorithmic benchmarking for emerging search or scenario discovery algorithm innovations, for testing the efficacy of new water management and planning methodologies.

In Chapter 4 the contribution of the Sedento Valley test case as a new generic multi-actor test bed for regional water supply planning has broad value for disseminating and extending the innovations presented in this dissertation. The model represents three hypothetical water utilities in the South Eastern US facing the prospect of water shortage due to growing demand and changing climate. The utilities seek to craft both short-term drought mitigation responses and long-term infrastructure sequencing pathways that maintain reliable supply and financial stability. The utilities have the potential to cooperate using treated transfers and through shared infrastructure development. The Sedento Valley test case was designed to be a realistic but computationally tractable tool for the direct comparison of competing regional-decision-making frameworks, for the design and comparison of drought mitigation strategies, and for the evaluation of decision-making metrics. While care was taken to limit computational demands and maintain clear interpretability, the Sedento Valley test case maintains a high degree of complexity due to the number of utilities involved, to the unbalanced distribution of resources and liabilities across utilities, and the utilities locations in relation to each other in the local shared river basins. The code repository provided with WaterPaths in Chapter 4 contains all the data needed to reproduce the test case in WaterPaths or other frameworks and to compare methodologies and policy instruments against what was presented in Chapters 5 and 4.

6.2 Future Work

This dissertation provides a foundational basis for the study of several challenges in water systems planning. This section highlights four suggested extensions: strategies for deep- and well-characterized-uncertainty sampling, metrics for long- and short-term decision making for water utilities, design of new financing and insurance instruments for water infrastructure operators, and conflict resolution among resource-sharing water utilities.

6.2.1 Strategies for Deep- and Well-Characterized-Uncertainty Sampling

Uncertainty factors are called deep when there is no explanatory probability distribution accepted by the various parties to a decision (Kwakkel et al., 2016b). Innovations in exploratory modeling frameworks that address deeply uncertainty factors and their impacts on a water infrastructure systems should focus on more efficient computational sampling strategies to inform, the calculation of robustness metrics or stochastic objectives. Computational parallelization and careful data management workflows would be highly beneficial for expanding analysts ability to parse through large multi-objective solution sets of planning policies and chose a limited number of them for scenario discovery, as performed in Chapters 3, 4 and 5 of this dissertation. The calculation of such robustness metrics and objectives requires in turn that values for deeply-uncertainty factors be assigned a distribution, and this choice may potentially impact the selection of a policy for scenario discovery (Reis and Shortridge, 2019). However, the topic of deep-uncertainty sampling has received little attention despite its potential to significantly impact decision making.

WaterPaths provides the ideal platform for the study of new workflows and new innovative stochastic sampling techniques when evaluating objectives and system robustness. Its stochastic design for objective calculations and batch mode provide an easy way to run sampling studies without the need for codewrappers or other workarounds as is typical current models to run all needed simulations. Furthermore, researchers interested in the topic would benefit from the Sedento Valley test case, which is already implemented in WaterPaths and has been extensively analyzed in this dissertation with the standard uniform-distribution approach.

6.2.2 Metrics for Long- and Short-Term Decision Making for Water Utilities

Various studies have analyzed metrics to inform short-term decision making by water utilities, such as the days-of-supply-remaining metric (Fisher, 1993; Fisher and Palmer, 1997) and the ROF metric exploited in this dissertation (Palmer and Characklis, 2009; Zeff et al., 2014) However, the field would benefit from a unified study of metrics for decision making. Researchers interested in the topic could extend WaterPaths to include a variety of decision-making metrics and perform MORDM analyses for the Sedento Valley test case with different metrics. The MORDM analyses could be followed by a detailed analysis of time series of system states for a granular understanding of how each metric affects decision making over time and for different conditions for water utilities. Such study would have great practical impact on the field of water infrastructure planning and management, demonstrating to water utilities the value of potentially switching metrics used for their operations and planning.

6.2.3 Design of New Financing Mechanisms and Insurance Instruments for Water Infrastructure Operators

As shown in Chapters 3 and 5 and in other works (Zeff et al., 2014, 2016; Herman et al., 2014), analyzing the financial component of water infrastructure planning and management is as important as the supply component to water utilities. Increasing pressure on water utilities to operate on ever smaller margins, the need for infrastructure refurbishment and capacity expansions (AWWA, 2019), and of credit rating standards for water utilities (Moody's, 2017, 2019), the municipal water sector would benefit from a continued effort to design innovative financial instruments and mechanisms for structuring dept. Although there are recent studies that are more directly engaging with the challenges associated with designing financial insurance programs for water utilities (Baum et al., 2018; Zeff and Characklis, 2013) and the existing Water Infrastructure Finance and Innovation Act (WIFIA) program for financing water infrastructure (Copeland, 2016). The area still appears to be in its technical infancy and is not at present capable of effectively addressing the global and US needs for massive investments in water infrastructure as well as sustaining the systems with well-funded operations and maintenance programs (Hallegatte et al., 2019). WaterPaths can be an effective tool for testing financial instruments and debt structure. It can be used by researchers and single or groups of municipalities to test general or tailored financial instruments and debt structure before committing to them in practice. Such proof-of-concepts may encourage case-specific implementation of such financial strategies, allowing for cost savings and improved financial stability across the municipal water industry.

6.2.4 Conflict Resolution Among Resource-Sharing Water Utilities

The global densification of urban systems means that water utilities that were once isolated are becoming more dependent on water resources shared with neighboring utilities. In such an environment, water transfers allow a group of water utilities to share capacity in ways that are more efficient than relying on independent sources (Vaux Jr and Howitt, 1984; Lund and Israel, 1995; Green and Hamilton, 2000; Olmstead, 2010). The trend of cooperation among utilities has also been encouraging the joint development and operation of shared infrastructure (see Zeff et al. 2016; Triangle J 2014 and Chapter 5). Such use and development of resources shared by multiple municipalities is explored in Chapters 3 and 5 of this dissertation and in other works (Zeff et al., 2014; Gorelick et al., 2019; Caldwell and Characklis, 2014; Borgomeo et al., 2018), although with a major simplification: the assumption that perfect contracts will be created and followed through the lifespan of a project. Gold et al. (2019) take a promising step by exploring how deviations in implementation of a plan can fundamentally shift perceived tradeoffs and robustness. They touch on the potential for game theory informed strategies for navigating the resulting conflicts among utilities in the Research Triangle region. Overall the body literature on the topic the stability of regional coalitions of utilities with conflicting interests and increasingly severe deeply uncertain stressors is underdeveloped. WaterPaths provides interested researchers with the unique ability to jointly simulate the behavior of multiple utilities in a regional system that can be abstracted with a broad array policy rule behaviors in highly dynamic and stochastic contexts (e.g., stochastic dynamic mechanism design or evolutionary dynamic games). WaterPaths can be used for the design of contracts to simulate the effects of utilities breaching or opting out of them, allowing for the design of policies with more safeguards and in more agreement with local institutional and political processes. These extensions would provide a deep and rich area of research that could be of immediate value for aiding regions globally that are straining to maintain the reliability and financial stability of their water supply services.

APPENDIX A

NOTATION

Symbol	Description
ATR	annual total revenue
CF	total amount of money in the contingency fund
D	actual (restricted) demands
D^w	actual water demand (as opposed to unrestricted) on week w
$\boldsymbol{DS}$	matrix with the streams of debt service payments for all utilities
$\boldsymbol{E}$	evaporation series matrix
f	joint-pdf of inflows, evaporation rates, and demand fluctuations around annual means, all well-characterized uncertainties
$\boldsymbol{F}$	vector-values objective function
f_{FC}	financial cost of drought mitigation instruments plus debt repayment
f_{JLA}	combined Jordan Lake allocation objective
f_{NPV}	net present cost of infrastructure
f_{REL}	supply reliability objective
f_{RF}	restriction frequency
f_{WFPC}	worse-first-percentile of the financial-cost objective
$\cdot_h$	historical data used to derive the synthetic series
H	ensemble model
h_t	model with index t of the ensemble
$\cdot_i$	index of a deeply uncertain factor

$\boldsymbol{ICO}$	matrix with vectors with the infrastructure construction ordering for each utility
$\boldsymbol{l}$	sets of lower bounds for the deeply uncertain factors
IPO^w	payout received by a utility in a given week
IPR	insurance loading (1.2 for this work, representing a loading of 20%)
N_{hr}	number years in the historical record
$\boldsymbol{NI}$	natural inflows matrix
N_{rof}	number of one-year long simulations used to calculate the ROF metric
N_f	total number of deeply uncertain factors
N_r	number of synthetic series samples
N_w	total number of weeks in the full planning horizon or number of weeks in a year
$\boldsymbol{OI}^y$	vector with the IDs of all online supply infrastructure
PO^y	insurance payout fixed on year's y and insurance contract
$\boldsymbol{R}$	temporal auto-correlation in the historical natural inflows and evaporation time series
$\cdot_s$	denotes synthetic time series
$\boldsymbol{S}$	cross-site correlation matrix
$\cdot_t$	index of model of boosting multiplier or demand tier (type of consumer)
T	total number of models in the ensemble
TF^w	volume of treated water to be transferred in week w

$\boldsymbol{u}$	sets of upper bounds for the deeply uncertain factors
$\boldsymbol{UD}$	matrix with unrestricted demands
UWP	unrestricted water prices
WP	restricted water prices
x	vector of explanatory variables
$\boldsymbol{X}$	time-varying state matrix across all of the utilities
$\boldsymbol{x}_{lrof}$	long-term ROF values
$\boldsymbol{x}_{lrof}$	vector of long-term ROFs for all utilities
$\boldsymbol{x}_{srof}$	short-term ROF values
$\cdot^{w}$	week number
$\cdot^{y}$	year number
α_t	boosting multiplier corresponding to model h_t
$\delta(\cdot)$	Dirac (delta) function
$\boldsymbol{\theta}_{acfc}$	vector of annual contingency fund contributions, which are percentages off annual revenue saved in a utility's drought mitigation fund.
θ_{ins}	ROF trigger for the index insurance
θ_{ipo}	percent of the previous year's revenue used to calculate the insurance payout
$\boldsymbol{\theta}_{irt}$	vector of insurance restriction triggers
$\boldsymbol{\theta}_{it}$	vector of long-term-ROF infrastructure construction triggers
$\boldsymbol{\theta}_{jla}$	vector of Jordan Lake Allocations
$\boldsymbol{\theta}_{rt}$	vector of restriction triggers
$\boldsymbol{\theta}_{tt}$	vector of transfer triggers

$\boldsymbol{\theta}_{irof}$	vector of long-term ROF triggers for all utilities
$\boldsymbol{\theta}^{*}$	Pareto optimal management and investment policy
$\boldsymbol{\Psi}_{DU}$	matrix of DU vector samples of the deeply uncertain variables of concern where each element corresponds to a specific uncertainty.
$\boldsymbol{\Psi}_{s}$	matrix of vector samples of the deeply uncertain variables of concern where each element corresponds to a specific uncertainty.
$\boldsymbol{\Psi}_{WCU}$	matrix of WCU vector samples of the deeply uncertain variables of concern where each element corresponds to a specific uncertainty.
$\mathcal{D}(\cdot)$	assumed pdf for the deeply uncertain factors
$\mathcal{U}$	uniform distribution
$\mathcal{X}$	vector of best projections for the deeply uncertain factors considered in the problem

Table A.1: List of symbols.

APPENDIX B

TIME-SERIES OUTPUT FOR COMPROMISE SOLUTIONS UNDER SELECTED SCENARIOS

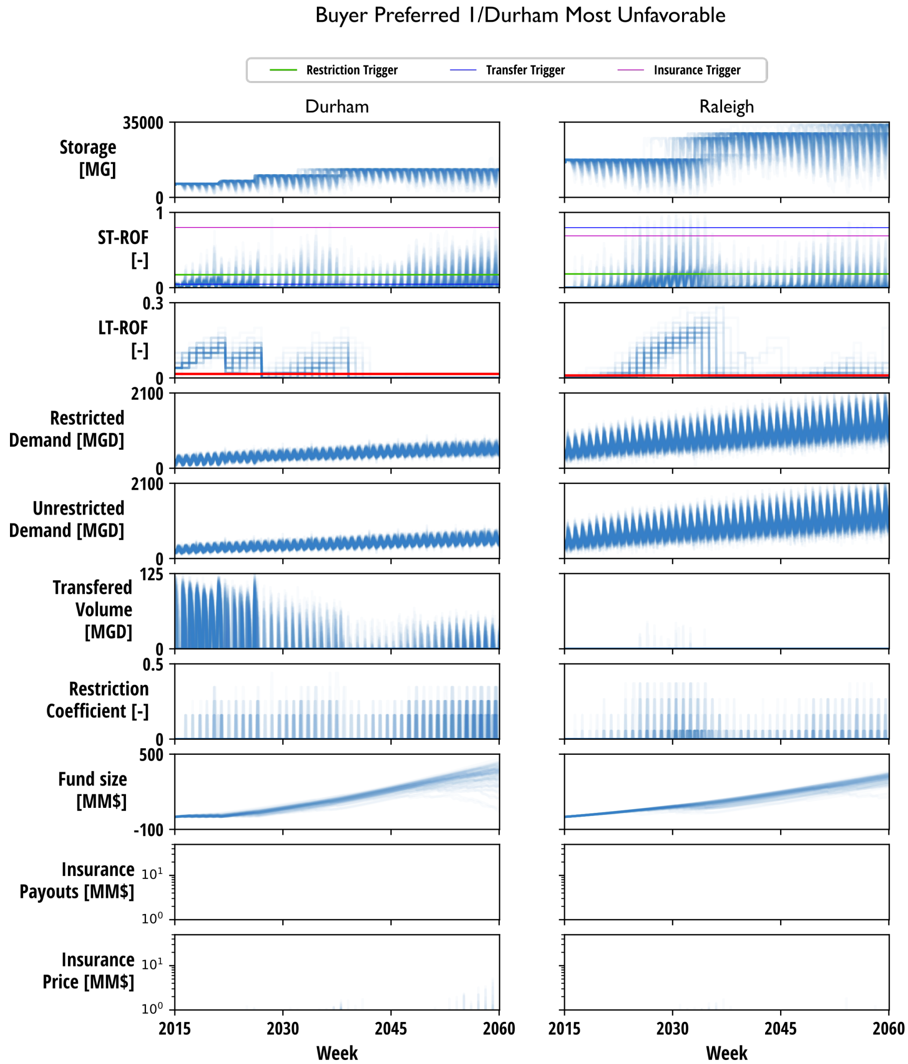

Figure B.1: System dynamics for Durham and Raleigh under policy D/B1 for multiple hydrologic realizations under Durham's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy.

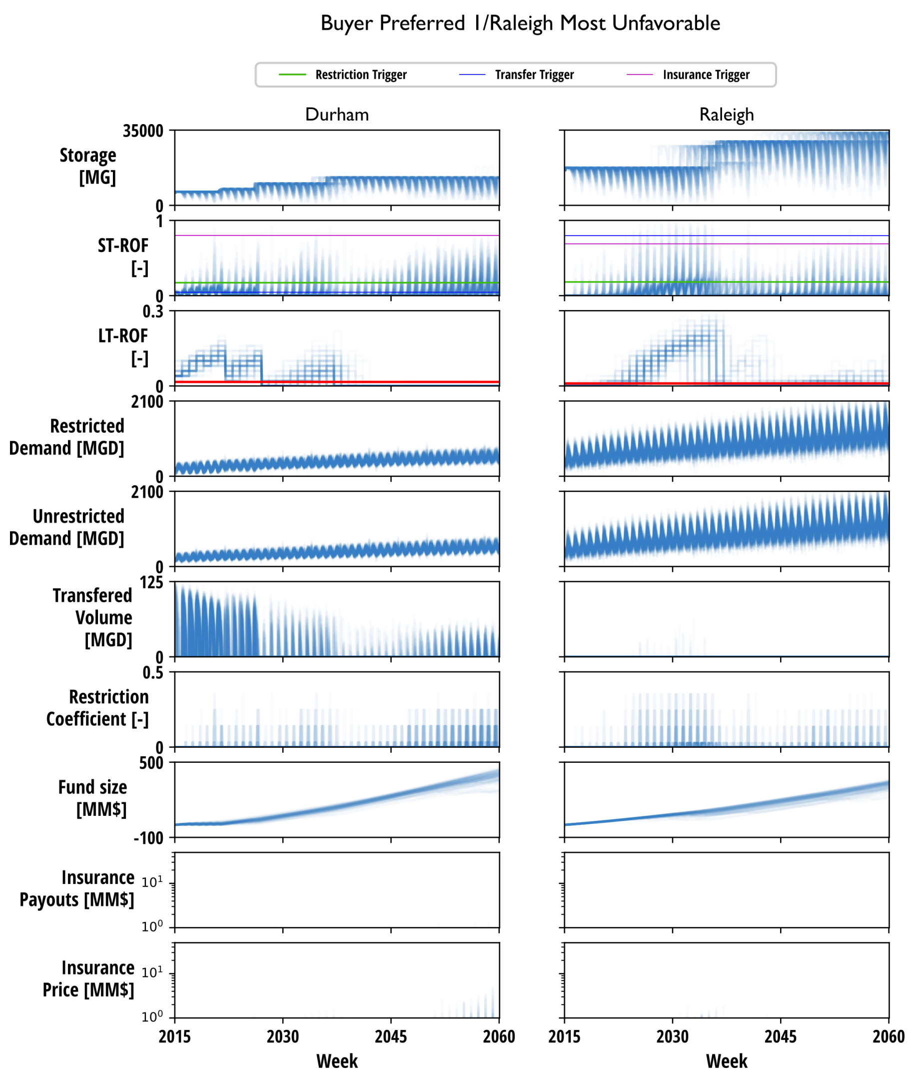

Figure B.2: System dynamics for Durham and Raleigh under policy D/B1 for multiple hydrologic realizations under Raleigh's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy.

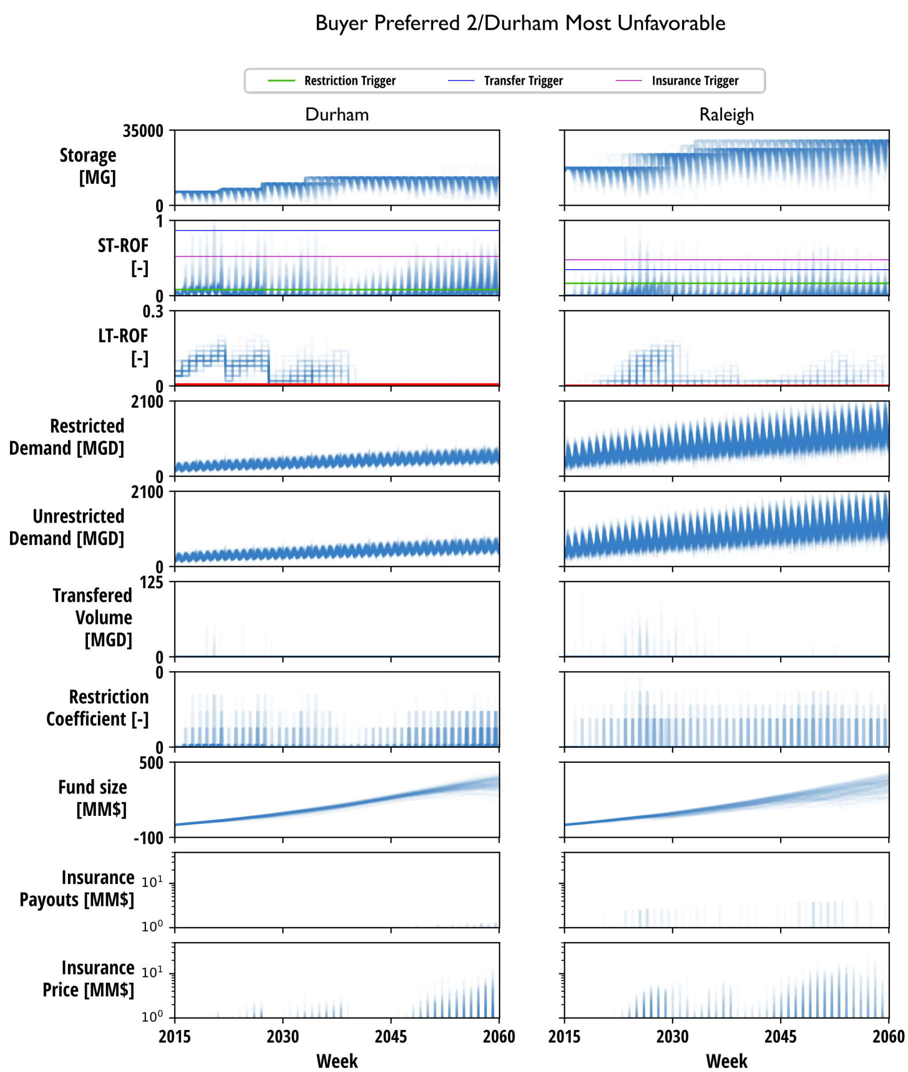

Figure B.3: System dynamics for Durham and Raleigh under policy R/B2 for multiple hydrologic realizations under Durham's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy.

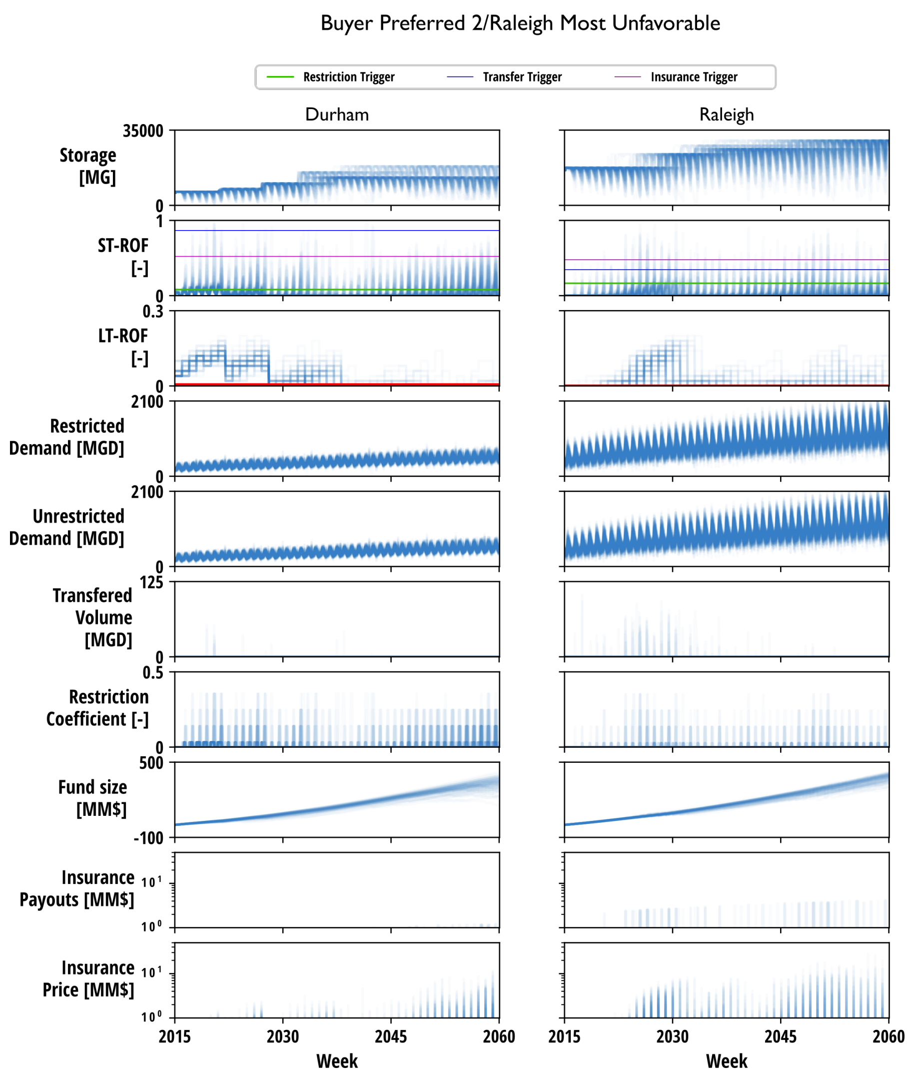

Figure B.4: System dynamics for Durham and Raleigh under policy R/B2 for multiple hydrologic realizations under Raleigh's most unfavorable scenario. Transparent lines indicate that few realizations had similar dynamics and the horizontal colored lines represent ROF trigger values for the corresponding policy.

APPENDIX C

DEEPLY UNCERTAIN FACTOR SENSITIVITIES

	D/B1			**R/B2**		
Uncertainty	Durham	Cary	Raleigh	Durham	Cary	Raleigh
Regional Demand Growth Multiplier	47 %	100 %	51 %	53 %	49 %	45 %
Regional Bond Interest Rate Multiplier	1 %	0 %	0 %	0 %	1 %	1 %
Regional Bond Term Multiplier	0 %	0 %	0 %	0 %	0 %	0 %
Regional Discount Rate Multiplier	1 %	0 %	0 %	0 %	1 %	1 %
Restriction Efficacy Multiplier OWASA	1 %	0 %	0 %	2 %	0 %	0 %
Restriction Efficacy Multiplier Durham	14 %	0 %	0 %	25 %	0 %	0 %
Restriction Efficacy Multiplier Cary	0 %	0 %	1 %	0 %	0 %	0 %
Restriction Efficacy Multiplier Raleigh	1 %	0 %	16 %	0 %	0 %	19 %
Regional Evaporation Rate Multiplier	1 %	0 %	1 %	0 %	25 %	1 %
Permitting Time Low WJLWTP	18 %	0 %	10 %	0 %	4 %	3 %
Construction Time Multiplier Low WJLWTP	2 %	0 %	1 %	1 %	1 %	0 %
Permitting Time High WJLWTP	0 %	0 %	0 %	1 %	0 %	3 %
Construction Time Multiplier Low WJLWTP	1 %	0 %	0 %	0 %	0 %	1 %

Table C.1: Percent of the combined impurity the the entire boosted trees ensemble due to each uncertain factor for Durham, Cary and Raleigh under both buyer preferred policies D/B1 and R/B2.

APPENDIX D

COMPUTATIONAL EXPERIMENT HYPERVOLUME PLOTS

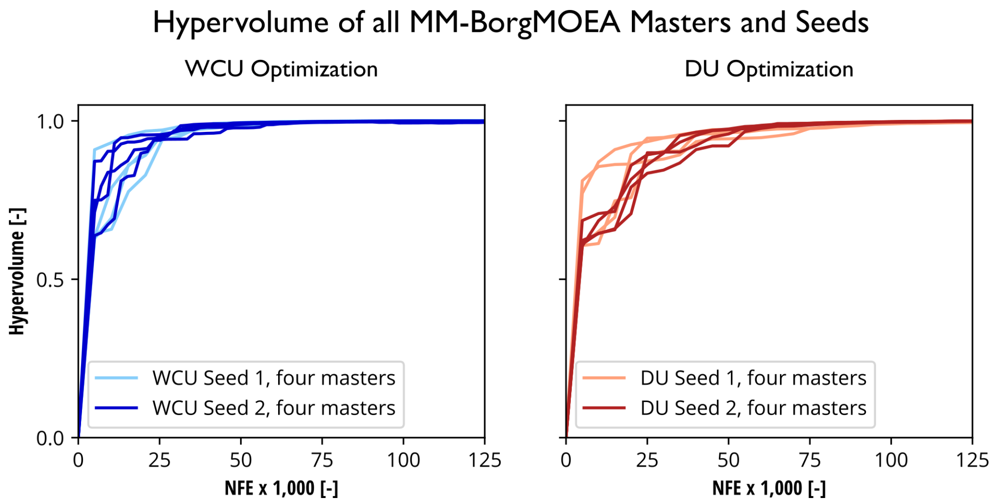

Figure D.1: Evolution of hypervolume for each master of each of two seeds ran for the WCU and DU searches. Since each seed had 4 island and each island ran 125,000 function evaluations, the total number of function evaluation for each of the 2 WCU and 2 DU seeds was 500,000. WCU seeds advanced faster at early NFE and converged slighly faster than the DU seeds, although the long plateaus at later NFE show that all seeds were given more than enough NFE to find the best Pareto approximate set of policies.

APPENDIX E

THE BOOSTED TREES ALGORITHM

E.1 Classification and Regression Trees

Classification and Regression Trees (CART) works by orthogonally dividing a space into regions and attaching a label or number to all points inside each region (Breiman et al., 1984). Each tree may sequentially divide the space multiple times, with each split being the one that minimizes the classification error on the overal tree at the time it is found. Take as an example the tree three shown in Figure E.2a, where the blue crosses represent a success scenario and a red circle represents a failure. The first split to be found, the one that results in the highest number of correctly-classified points among all possible orthogonal splits on both axes, would be the horizontal. The second split, a vertical one, would be found next and branching off of the horizontal split, the third (also vertical) off of the second (vertical as well), and so on.

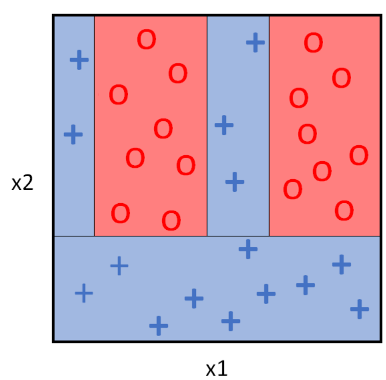

Figure E.1: Classification tree fit on illustrative data set. The first split to be foudn was the horizontal split, which applies to the entire data set. The three vertical splits were foudn branching off of the horizontal split and then off of each other.

More previsely, the ideal split at each iteration for a classification problem (e.g., scenario discovery) can be found by checking all possible splits for the one that minimizes one of a variety of metrics, such as the Gini impurity calculated for a tree (Breiman et al., 1984), shown in Equation E.3:

$$p_k = \frac{|S_k|}{|S|} \leftarrow \text{fraction of scenarios in leaf } S_n \text{ with success/fail label } k \tag{E.1}$$

$$G(S_n) = \sum_{k=\{\text{success, fail}\}} p_k(1-p_k) \tag{E.2}$$

with the Gini impurity after the split being:

$$G^T(S) = \frac{|S_1|}{|S|}G^T(S_1) + \frac{|S_2|}{|S|}G^T(S_2) + ... + \frac{|S_N|}{|S|}G^T(S_N) \tag{E.3}$$

for data set S, $S_n \subseteq S$, $S_{n,k} \subseteq S_n$ where $S_{n,k} = \{(\boldsymbol{x}, y) \in S_n : y = k\}$ (k being either success or fail), and $S = S_{success} \cup S_{fail}$. In Equations E.1-E.3, S_n is the data in a leaf, p_k is the fraction of points belonging to class k within leaf n, $G(S_n)$ is the Gini impurity within a leaf, $G^T(S)$ is the Gini impurity for the entire tree. The change in the total impurity across a tree due to all splits on a particular uncertain factor (dependent variable) is also an indicator of how sensitive to that factor performance is.

E.2 General Idea of Booosting

The idea of boosting is to train an ensemble of shallow classification trees (or any weak classifier), each of which an expert in a specific region of the uncertainty space. The ensemble of trees has the form below:

$$H(\boldsymbol{x}) = \sum_{t=1}^{T} \alpha_t h_t(\boldsymbol{x}) \tag{E.4}$$

where H is the ensemble of trees, α_t is the weight assigned to weak classifier h_t around samples $\boldsymbol{x}$ at iteration t, and T is the number of trees in the ensemble.

Boosting creates such an ensemble in a similar fashion to gradient descent. However, instead of:

$$\boldsymbol{x}_{t+1} = \boldsymbol{x}_t - \alpha \nabla \ell(\boldsymbol{x}_t) \tag{E.5}$$

as in gradient descent in real space, where ℓ is a loss function, Boosting is trained via gradient descent in functional space, so that:

$$H_{t+1}(\boldsymbol{X}) = H_t(\boldsymbol{X}) + \alpha_t \nabla \ell(h_t(\boldsymbol{X})) \tag{E.6}$$

The question then becomes how to find the α_t's and the trees h_t. Before answering these questions, we should get a geometric intuition for Boosting first.

E.3 Geometric Intuition

Figure E.2a shows a set of blue crosses and red circles denoting success and failure scenarios, respectively, we would like our ensemble of classification trees with depth of one split (hereafter called a tree stump) to correctly classify. For simplicity, let's assume all trees will have the same weight α_t in the final ensemble.

The first tree stump, a horizontal divide in panel "b," classified ten out of thirteen points correctly but failed to classify the remaining three. Since it incorrectly classified a few points in the last attempt, we would like the next classifier correctly classify these points. To increase the chances of that being the case,

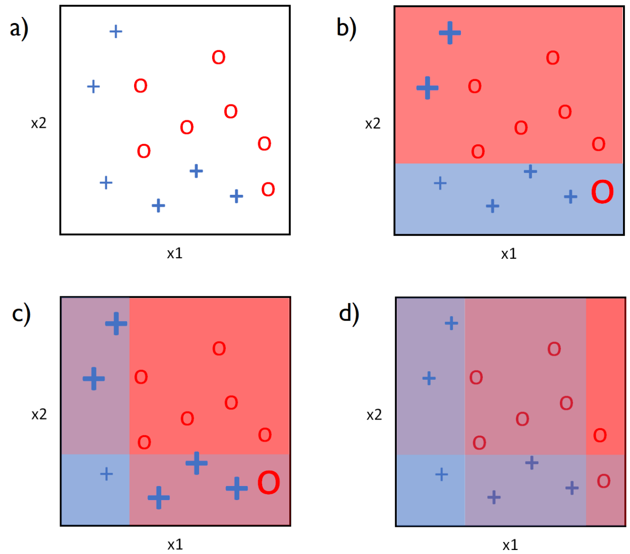

Figure E.2: Sequence of addition of single-split classification trees until all points are correctly classified. Panel (a) shows just the data, without trees, panel (b) shows data with one tree, panel (c) with two, and panel (d) with three. Variables x_1 and x_2 are two uncertain factors, the blue crosses represent a success scenario and a red circle represents a failure, and the blue and red regions correspond to regions classified as success or failure regions, with intermediate colors indicating uncertainty in the classification.

boosting will increase the weight w of the points that were misclassified earlier before training the new tree stump. The second tree stump, a vertical line on panel "c," correctly classifies the two blue crosses that were originally incorrectly classified, although it incorrectly classifies three crosses that were originally correctly classified. For the third classifier, Boosting will now increase the weight of the three bottom misclassified crosses as well of the other misclassified two crosses and circle because they are still not correctly classified – technically speaking they are tied, each of the two classifiers classifies them in a different way, but we are here considering this a wrong classification. The third iteration will then prioritize correcting the high-weight points again, and will end up as the vertical line on the right of panel "d." Now, all points are correctly classified.

There are different ways to mathematically approach Boosting. But before getting to Boosting, it is a good idea to go over gradient descent, which is a basis of Boosting. Following that, the AdaBoost algorithm will be presented, which is a boosting algorithm that assumes an exponential loss function.

E.4 Minimizing a Loss Function

Gradient Descent in Functional Space

In real-space gradient descent function minimization algorithm, the goal is to start from an initial guess x_0 and iteratively find the value of x corresponding to the minimum value of $f(x)$, which in machine learning is a loss function, henceforth called $\ell(x)$. Gradient descent does that by moving one step s at a time starting from x^0 in the direction of the steeped downhill slope at location x_t, the value of x at the t^{th} iteration (Murphy, 2012; Hastie et al., 2009).

$$\ell\left[x_{t+1} - \alpha g(x_t)\right] \approx \ell(x^t) - \alpha g(x_t)^T g(x_t) \tag{E.7}$$

where α, called the learning rate, must be positive and can be set as a fixed parameter. The dot product on the last term $g(x_t)^T g(x_t)$ will also always be positive, which means that the loss should always decrease.

In gradient descent for functional space, x is a function instead of a real number. This means that the loss function $\ell(\cdot)$ is a function of a function, say $\ell(H(\boldsymbol{x}))$ instead of a real number x. By analogy, gradient descent in functional space applied to classification trees works by adding trees to an ever growing ensemble of trees. Using the definition of a functional gradient, which is beyond the scope of the this post, this leads us to (Murphy, 2012):

$$\ell(H + \alpha h) \approx \ell(H) + \alpha{<}\nabla\ell(H), h{>} \tag{E.8}$$

where H is an ensemble of trees, h is a single tree, and the $<f, g>$ notation denotes a dot product between f and g. Gradient descent in function space is an important component of Boosting. Based on that, the next section will talk about AdaBoost, the Boosting algorithm used in this study.

E.5 AdaBoost

Basic Definitions

AdaBoost is a Boosting algorithm, whose goal is to find an ensemble function H of weak classifiers h, here shallow classification trees, that minimize an exponential loss function below for a binary classification problem (Murphy, 2012):

$$\ell(H) = \sum_{i=1}^{n} e^{-y(x_i)H(x_i)} \tag{E.9}$$

where $x_i, y(x_i)$ is the i^{th} data point in the training set. The step size α can be interpreted as the weight of each classifier in the ensemble, which optimized for each function h added to the ensemble. The uncertain factors independent variables $\boldsymbol{x}$ have corresponding vector of dependent variables $\boldsymbol{y}(x_i)$, in which each $y(x_i) \in \{-1, 1\}$ is a vector with the classification of each point, with -1 and 1 representing the two classes to which a point may belong (say, -1 for red circles and 1 for blue crosses). The weak classifiers h in AdaBoost also return $h(x) \in \{-1, 1\}$.

The weight α_t of each weak classifier h is calculated as:

$$\alpha = \frac{1}{2} ln \frac{1-\epsilon}{\epsilon} \tag{E.10}$$

where ϵ is the classification error of weak classifier h_t. The error ϵ for each iteration is:

$$h(\boldsymbol{x}_i) = argmin_h \sum_{i:h(\boldsymbol{x}_i) \neq y(\boldsymbol{x}_i)} w_i \tag{E.11}$$

which is the summation of the weights of the points misclassified by $h(\boldsymbol{x}_i)_t$ – e.g., in panel "b" the error would be summation the of the weights of the two crosses on the upper left corner and of the circle at the bottom right corner. Now let's get to these weights.

There are multiple ways we can think of for setting the weights of each data point at each iteration, corresponding to the sizes of the points in Figure E.2. The way presented in Schapire (1990) is a common choice:

$$w_{t+1} = w_t \frac{e^{-\alpha_t h_t(\boldsymbol{x}) y(\boldsymbol{x})}}{2\sqrt{\epsilon(1-\epsilon)}} \tag{E.12}$$

Algorithm 1 shows a pseudo-code of AdaBoost.

Algorithm 1: AdaBoost pseudo-code.

Input : $(\boldsymbol{X}, \boldsymbol{y})$, $\mathbb{H}$ (classification tree)

Output: Ensemble classifier $H(\boldsymbol{x})$

```
H=0
w⃗ = {1/n, ..., 1/n}
for t = 0 : T − 1 do
    find h^t by training classification tree on weighed scenarios
    calculate error: ∑_{i:h(x_i)≠y(x_i)} w_i
    if ϵ < 1/2 then
        calculate iteration classification tree weight: α = 1/2 ln (1−ϵ)/ϵ
        add classification to ensemble: H^{t+1} = H^t + α^t h^t
        update scenarios' weights: w^{t+1} = w^t e^{−α^t h^t(x⃗) y(x⃗)} / (2 sqrt(ϵ(1−ϵ)))
    else
        return H;
    end
end
```

E.6 Sensitivities of Uncertainty Factors

Each tree added to the ensemble has its own values of impurity decrease for each uncertain factor. The average total impurity decrease for an uncertain factor x across all trees h in the ensemble H can be used as a measure of how significant that factor is in determining the system's performance under a given scenario (Hastie et al., 2009), as shown in Equation E.13.

$$s_d = 100\% \cdot \frac{1}{T} \sum_{t=1}^{T} \frac{\sum_{d=1}^{D} \mathbb{I}_d \cdot (G_{d-1}^T(S) - G_d^T(S))}{G_0^T(S)} \tag{E.13}$$

where s_k is the sensitivity index for factor k, d in the split index (first split, second split, etc.), D is the depth of the trees, and $\mathbb{I}_{d,k}$ is a binary variable assuming the value of 1 if split d was performed on k and 0 otherwise. Although deeper trees will more likely report scores for more factors, the scores for factors that only appear on deeper trees are likely to be negligible. Furthermore, Boosting is designed for high bias classifiers, meaning shallow trees (depths between 2 and 4), so a deeper tree may hamper its performance.

E.7 Exponencial Convergence

Boosted algorithms, not only AdaBoost and independently of which weak classified is chosen, are guaranteed to converge. It is possible to derive an expression for the upper bound of the error, resulting in:

$$\ell(H) \leq n(1 - 4\gamma^2)^{T/2} \tag{E.14}$$

which means that the training error is bound by an exponential decay as you add classifiers to the ensemble. This result applies to any boosted algorithm.

E.8 Reduced Overfitting

Lastly, Boosted algorithms are remarkably resistant to overfitting. According to Murphy (2012), a possible reason is that Boosting can be seen as a form of ℓ_1 regularization, which is prone to eliminate irrelevant features and thus reduce overfitting. Another explanation is related to the concept of margins, so that at least certain boosting algorithms force a classification on a point only if possible by a certain margin, thus also preventing overfitting.

BIBLIOGRAPHY

ASCE (2017). 2017 infrastructure report card.

AWWA, A. W. W. A. (2019). State of the water industry report.

Bankes, S. (1993). Exploratory modeling for policy analysis. *Operations Research*, 41(3):435–449.

Bankes, S. C., Lempert, R. J., and Popper, S. W. (2001). Computer-Assisted Reasoning. *Computing in Science and Engineering*, 3(2):71–76.

Basdekas, L. (2014). Is multiobjective optimization ready for water resources practitioners? utilitys drought policy investigation.

Baum, R., Characklis, G. W., and Serre, M. L. (2018). Effects of geographic diversification on risk pooling to mitigate drought-related financial losses for water utilities. *Water Resources Research*, 54(4):2561–2579.

Beh, E. H., Maier, H. R., and Dandy, G. C. (2015a). Scenario Driven Optimal Sequencing Under Deep Uncertainty. *Environ. Model. Softw.*, 68(C):181–195.

Beh, E. H., Zheng, F., Dandy, G. C., Maier, H. R., and Kapelan, Z. (2017). Robust optimization of water infrastructure planning under deep uncertainty using metamodels. *Environmental Modelling & Software*, 93:92–105.

Beh, E. H. Y., Maier, H. R., and Dandy, G. C. (2015b). Adaptive, multiobjective optimal sequencing approach for urban water supply augmentation under deep uncertainty. *Water Resources Research*, 51(3):1529–1551.

Ben-Haim, Y. (2006). *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*. Elsevier. Google-Books-ID: yR9H_WbkIHkC.

Bertsekas, D. P., Bertsekas, D. P., Bertsekas, D. P., and Bertsekas, D. P. (1995). *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA.

Blanco, C. D. P. and Gómez, C. M. G. (2013). Designing optimum insurance schemes to reduce water overexploitation during drought events: a case study of la campiña, guadalquivir river basin, spain. *Journal of Environmental Economics and Policy*, 2(1):1–15.

Bode, F., Reed, P., Reuschen, S., and Nowak, W. (2019). Search space representation and reduction methods to enhance multiobjective water supply monitoring design. *Water Resources Research*.

Bonzanigo, L., Rozenberg, J., Felter, G. C., Lempert, R. J., and Reed, P. M. (2018). Building the Resilience of WSS Utilities to Climate Change and Other Threats : A Road Map. Technical Report 133227, The World Bank.

Borgomeo, E., Mortazavi-Naeini, M., Hall, J. W., and Guillod, B. P. (2018). Risk, robustness and water resources planning under uncertainty. *Earth's Future*, 6(3):468–487.

Borgomeo, E., Mortazavi-Naeini, M., Hall, J. W., O'Sullivan, M. J., and Watson, T. (2016). Trading-off tolerable risk with climate change adaptation costs in water supply systems. *Water Resources Research*, 52(2):622–643.

Bosomworth, K., Leith, P., Harwood, A., and Wallis, P. J. (2017). What's the problem in adaptation pathways planning? the potential of a diagnostic problem-structuring approach. *Environmental Science & Policy*, 76:23–28.

Breiman, L., Friedman, J., Stone, C., and Olshen, R. (1984). *Classification and*

Regression Trees. The Wadsworth and Brooks-Cole statistics-probability series. Taylor and Francis.

Brown, C., Ghile, Y., Laverty, M., and Li, K. (2012). Decision Scaling: Linking bottom-up vulnerability analysis with climate projections in the water sector. *Water Resources Research*, 48(9):W09537.

Bryant, B. P. and Lempert, R. J. (2010). Thinking inside the box: A participatory, computer-assisted approach to scenario discovery. *Technological Forecasting and Social Change*, 77(1):34–49.

Caldwell, C. and Characklis, G. W. (2014). Impact of Contract Structure and Risk Aversion on Inter-utility Water Transfer Agreements. *Journal of Water Resources Planning and Management*, 140(1):100–111.

Campolongo, F., Cariboni, J., and Saltelli, A. (2007). An effective screening design for sensitivity analysis of large models. *Environmental Modelling & Software*, 22(10):1509–1518.

Cantu-Paz, E. (2000). *Efficient and Accurate Parallel Genetic Algorithms*. SPRINGER NATURE.

Census, U. S. C. B. (2019). Numeric population change by metropolitan/micropolitan statistical area, 2010-2018.

Characklis, G. W., Kirsch, B. R., Ramsey, J., Dillard, K. E. M., and Kelley, C. T. (2006). Developing portfolios of water supply transfers. *Water Resources Research*, 42(5).

Coello, C. A. C., Lamont, G. B., and Veldhuizen, D. A. V. (2006). *Evolutionary Algorithms for Solving Multi-Objective Problems (Genetic and Evolutionary Computation)*. Springer-Verlag, Berlin, Heidelberg.

Copeland, C. (2016). *Water infrastructure financing: the water infrastructure finance and innovation act (WIFIA) program*. Congressional Research Service Washington, DC.

Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263.

Dalal, S., Han, B., Lempert, R., Jaycocks, A., and Hackbarth, A. (2013). Improving scenario discovery using orthogonal rotations. *Environmental Modelling & Software*, 48:49–64.

Deb, K. and Agrawal, B. (1994). Simulated binary crossover for continuous search space (technical reports iitk/me/smd-94027). *Convenor: Indian Institute of Technology, Department of Mechanical Engineering*.

Deb, K., Joshi, D., and Anand, A. (2002a). Real-coded evolutionary algorithms with parent-centric recombination. In *Proceedings of the 2002 Congress on Evolutionary Computation. 02 (Cat. No.02TH8600)*. IEEE.

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002b). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197.

Denaro, S., Castelletti, A., Giuliani, M., and Characklis, G. W. (2018). Fostering cooperation in power asymmetrical water systems by the use of direct release rules and index-based insurance schemes. *Advances in Water Resources*, 115:301–314.

Department of Water Management (2009). City of Durham Water Shortage Response Plan.

Dessai, S., Hulme, M., Lempert, R., and Pielke, R. (2009). Do We Need Better Predictions to Adapt to a Changing Climate? *Eos, Transactions American Geophysical Union*, 90(13):111–112.

Dittrich, R., Wreford, A., and Moran, D. (2016). A survey of decision-making approaches for climate change adaptation: Are robust methods the way forward? *Ecological Economics*, 122:79–89.

Draper, A. J., Jenkins, M. W., Kirby, K. W., Lund, J. R., and Howitt, R. E. (2003). Economic-engineering optimization for california water management. *Journal of water resources planning and management*, 129(3):155–164.

Drucker, H. and Cortes, C. (1996). Boosting decision trees. In *Advances in neural information processing systems*, pages 479–485.

EPA, P. O. (2019). Epa receives 51 requests totaling over $6 billion for third round of wifia funding. *EPA Press Releases*.

Erfani, T., Pachos, K., and Harou, J. J. (2018). Real-options water supply planning: Multistage scenario trees for adaptive and flexible capacity expansion under probabilistic climate change uncertainty. *Water Resources Research*.

Fiering, M. B., Bund, B., and Jackson, B. B. (1971). *Synthetic streamflows*, volume 1. American Geophysical Union.

Fisher, S. (1993). Days-of-supply-remaining as an indicator of drought severity in water supply planning and management. Master's thesis, University of Washington.

Fisher, S. and Palmer, R. N. (1997). Managing water supplies during drought: triggers for operational responses. *Water Resources Update*, 3(108):14–31.

Fletcher, S., Lickley, M., and Strzepek, K. (2019). Learning about climate change uncertainty enables flexible water infrastructure planning. *Nature Communications*, 10(1).

Fletcher, S. M., Miotti, M., Swaminathan, J., Klemun, M. M., Strzepek, K., and Siddiqi, A. (2017). Water supply infrastructure planning: Decision-making framework to classify multiple uncertainties and evaluate flexible design. *Journal of Water Resources Planning and Management*, 143(10):04017061.

Forum, M. P. (1994). Mpi: A message-passing interface standard version 3.0. Technical report, Message Passing Interface Forum, Knoxville, TN, USA.

Foster, B. T., Kern, J. D., and Characklis, G. W. (2015). Mitigating hydrologic financial risk in hydropower generation using index-based financial instruments. *Water Resources and Economics*, 10:45–67.

Freund, Y., Schapire, R., and Abe, N. (1999). A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612.

Friedman, J. H. and Fisher, N. (1999). Bump hunting in high-dimensional data. *Stat. Comput.*, 9(2):123–143.

Gabriel, E., Fagg, G. E., Bosilca, G., Angskun, T., Dongarra, J. J., Squyres, J. M., Sahay, V., Kambadur, P., Barrett, B., Lumsdaine, A., Castain, R. H., Daniel, D. J., Graham, R. L., and Woodall, T. S. (2004). Open MPI: Goals, concept, and design of a next generation MPI implementation. In *Proceedings, 11th European PVM/MPI Users' Group Meeting*, pages 97–104, Budapest, Hungary.

Gaspero, L. D. (2007). Quadprogpp, a quadratic programming c++ library based on the goldfarb-idnani active-set dual method. Technical report, University of Udine.

Ghile, Y., Taner, M., Brown, C., Grijsen, J., and Talbi, A. (2014). Bottom-up climate risk assessment of infrastructure investment in the Niger River Basin. *Climatic Change*, 122(1-2):97–110.

Giuliani, M. and Castelletti, A. (2016). Is robustness really robust? How different definitions of robustness impact decision-making under climate change. *Climatic Change*, 135(3):409–424.

Giuliani, M., Quinn, J. D., Herman, J. D., Castelletti, A., and Reed, P. M. (2018). Scalable multiobjective control for large-scale water resources systems under uncertainty. *IEEE Transactions on Control Systems Technology*, 26(4):1492–1499.

Gleick, P. H. (2002a). Soft water paths. *Nature*, 418:373.

Gleick, P. H. (2002b). Water management: Soft water paths. *Nature*, 418:373.

Gleick, P. H. (2003). Global freshwater resources: Soft-path solutions for the 21st century. *Science*, 302(5650):1524–1528.

Gleick, P. H. (2018). Transitions to freshwater sustainability. *Proceedings of the National Academy of Sciences*, 115(36):8863–8871.

Gleick, P. H., Institute, P., and Ajami, N. (2014). *The World's Water, Volume 8: The Biennial Report on Freshwater Resources*. ISLAND PR.

Goh, C. K. and Tan, K. C. (2007). An Investigation on Noisy Environments in Evolutionary Multiobjective Optimization. *IEEE Transactions on Evolutionary Computation*, 11(3):354–381.

Gold, D. F., Reed, P. M., Trindade, B. C., and Characklis, G. W. (2019). Identifying actionable compromises: Navigating multi-city robustness conflicts to discover cooperative safe operating spaces for regional water supply portfolios. *Water Resources Research*. (In-Press).

Goldfarb, D. and Idnani, A. U. (1983). A numerically stable dual method for solving strictly convex quadratic programs. *Mathematical Programming*, 27:1–33.

Goodwin, L. (2009). Water Shortage Response Plan.

Gorelick, D. E., Zeff, H. B., Hughes, J., Eskaf, S., and Characklis, G. W. (2019). Exploring treatment and capacity-sharing agreements between water utilities. *Journal - American Water Works Association*, 111(9):26–40.

Grant, S. B., Fletcher, T. D., Feldman, D., Saphores, J.-D., Cook, P. L. M., Stewardson, M., Low, K., Burry, K., and Hamilton, A. J. (2013). Adapting urban water systems to a changing climate: Lessons from the millennium drought in southeast australia. *Environmental Science & Technology*, 47(19):10727–10734.

Green, G. P. and Hamilton, J. R. (2000). Water allocation, transfers and conservation: Links between policy and hydrology. *International Journal of Water Resources Development*, 16(2):197–208.

Groves, D. G. and Lempert, R. J. (2007). A new analytic method for finding policy-relevant scenarios. *Global Environmental Change*, 17(1):73–85. Uncertainty and Climate Change Adaptation and Mitigation.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., and ter Maat, J. (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change*, 23(2):485–498.

Hadka, D., Madduri, K., and Reed, P. (2013). Scalability Analysis of the Asynchronous Master-Slave Borg Multiobjective Evolutionary Algorithm. In *The*

16th International Workshop on Nature Inspired Distributed Computing (NIDISC) at the 27th IEEE/ACM International Parallel and Distributed Processing Symposium (IPDPS), Boston, MA.

Hadka, D. and Reed, P. (2012a). Diagnostic Assessment of Search Controls and Failure Modes in Many-Objective Evolutionary Optimization. *Evolutionary Computation*, 20(3):423–452.

Hadka, D. and Reed, P. (2012b). Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. *Evolutionary Computation*, 20(3):423–452.

Hadka, D. and Reed, P. (2013). Borg: An Auto-Adaptive Many-Objective Evolutionary Computing Framework. *Evolutionary Computation*, 21(2):231–259.

Hadka, D. and Reed, P. (2014). Large-scale Parallelization of the Borg Multiobjective Evolutionary Algorithm to Enhance the Management of Complex Environmental Systems. *Environmental Modelling and Software*.

Hadka, D., Reed, P., and Simpson, T. (2012). Diagnostic Assessment of the Borg MOEA for Many-Objective Product Family Design Problems. In *WCCI 2012 World Congress on Computational Intelligence, Congress on Evolutionary Computation*, pages 986–995, Brisbane, Australia.

Hadka, D. M., Herman, J. D., Reed, P. M., and Keller, K. (2015). An open source framework for many-objective robust decision making. *Environmental Modelling and Software*, 74:114–129.

Halim, R. A., Kwakkel, J. H., and Tavasszy, L. A. (2015). A scenario discovery study of the impact of uncertainties in the global container transport system on European ports. *Futures*, pages –.

Hall, J. (2019). A simulation tool to guide infrastructure decisions: System-of-systems modeling aids prioritization and uncertainty planning. *IEEE Systems, Man, and Cybernetics Magazine*, 5(3):10–20.

Hall, J. W., Grey, D., Garrick, D., Fung, F., Brown, C., Dadson, S. J., and Sadoff, C. W. (2014). Coping with the curse of freshwater variability. *Science*, 346(6208):429–430.

Hall, J. W., Mortazavi-Naeini, M., Borgomeo, E., Baker, B., Gavin, H., Gough, M., Harou, J. J., Hunt, D., Lambert, C., Piper, B., Richardson, N., and Watts, G. (2019). Risk-based water resources planning in practice: a blueprint for the water industry in england. *Water and Environment Journal*.

Hallegatte, S., Maruyama, R. J. E., and Rozenberg, J. (2019). Lifelines : The resilient infrastructure opportunity : Overview. Technical report, The World Bank, Washin.

Hamarat, C., Kwakkel, J. H., Pruyt, E., and Loonen, E. T. (2014). An exploratory approach for adaptive policymaking by using multi-objective robust optimization. *Simulation Modelling Practice and Theory*, 46:25–39.

Harou, J. J., Pulido-Velazquez, M., Rosenberg, D. E., Medellín-Azuara, J., Lund, J. R., and Howitt, R. E. (2009). Hydro-economic models: Concepts, design, applications, and future prospects. *Journal of Hydrology*, 375(3-4):627–643.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. Springer-Verlag New York Inc.

Herman, J. and Usher, W. (2017). SALib: An open-source python library for sensitivity analysis. *The Journal of Open Source Software*, 2(9):97.

Herman, J. D., Reed, P. M., Zeff, H. B., and Characklis, G. W. (2015). How should robustness be defined for water systems planning under change? *Journal of Water Resources Planning and Management*, page In Press.

Herman, J. D., Zeff, H. B., Lamontagne, J. R., Reed, P. M., and Characklis, G. W. (2016). Synthetic drought scenario generation to support bottom-up water supply vulnerability assessments. *Journal of Water Resources Planning and Management*, 142(11):04016050.

Herman, J. D., Zeff, H. B., Reed, P. M., and Characklis, G. W. (2014). Beyond optimality: Multistakeholder robustness tradeoffs for regional water portfolio planning under deep uncertainty. *Water Resources Research*, 50(10):7692–7713.

Hernández-Lobato, D., Hernandez-Lobato, J., Shah, A., and Adams, R. (2016). Predictive entropy search for multi-objective bayesian optimization. In *International Conference on Machine Learning*, pages 1492–1501.

Hipel, K. W. and Ben-Haim, Y. (1999). Decision making in an uncertain world: Information-gap modeling in water resources management. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 29(4):506–517.

Hughes, J. and Leurig, S. (2013). Assessing Water System Revenue Risk: Considerations for Market Analysts. Technical report, Ceres.

Hui, R., Herman, J., Lund, J., and Madani, K. (2018). Adaptive water infrastructure planning for nonstationary hydrology. *Advances in Water Resources*, 118:83–94.

Huskova, I., Matrosov, E. S., Harou, J. J., Kasprzyk, J. R., and Lambert, C. (2016).

Screening robust water infrastructure investments and their trade-offs under global change: A london example. *Global Environmental Change*, 41:216–227.

Hydraulics, D. (2004). Ribasim: River basin simulation program operating manual and description. *Netherlands, Delft*.

HydroLogics (2009). *User Manual for OASIS With OCL*. HydroLogics, 10440 Shaker Drive, Suite 104, Columbia, MD, 21046, model version 3.10.8, gui version 4.6.16 edition.

Intel (2019). *Intel MPI Library Developer Reference for Linux OS*. Intel, 2019u5 edition.

Jaxa-Rozen, M. and Kwakkel, J. (2018). Tree-based ensemble methods for sensitivity analysis of environmental models: A performance comparison with sobol and morris techniques. *Environmental Modelling & Software*, 107:245–266.

Jha, M. K. and Gupta, A. D. (2003). Application of mike basin for water management strategies in a watershed. *Water International*, 28(1):27–35.

Kasprzyk, J. R., Nataraj, S., Reed, P. M., and Lempert, R. J. (2013). Many Objective Robust Decision Making for Complex Environmental Systems Undergoing Change. *Environmental Modelling and Software*, 42:55–71.

Kasprzyk, J. R., Reed, P. M., Characklis, G. W., and Kirsch, B. R. (2012). Many-objective de Novo water supply portfolio planning under deep uncertainty. *Environmental Modelling and Software*, 34:87–104. Emulation techniques for the reduction and sensitivity analysis of complex environmental models.

Kingsborough, A., Borgomeo, E., and Hall, J. W. (2016). Adaptation pathways in practice: Mapping options and trade-offs for london's water resources. *Sustainable Cities and Society*, 27:386–397.

Kingsborough, A., Jenkins, K., and Hall, J. W. (2017). Development and appraisal of long-term adaptation pathways for managing heat-risk in london. *Climate Risk Management*, 16:73–92.

Kirsch, B. R., Characklis, G. W., and Zeff, H. B. (2013). Evaluating the Impact of Alternative Hydro-Climate Scenarios on Transfer Agreements: A Practical Improvement for Generating Synthetic Streamflows. *Journal of Water Resources Planning and Management*, 139(4):396–406.

KITA, H., ONO, I., and KOBAYASHI, S. (2000). Multi-parental extension of the unimodal normal distribution crossover for real-coded genetic algorithms. *Transactions of the Society of Instrument and Control Engineers*, 36(10):875–883.

Klemm, M., Supinski, B., and Board, O. (2019). *OpenMP Application Programming Interface Specification Version 5. 0*. Independently Published.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Hart, Schaffner and Marx, New York. reprint 1, 1933.

Kollat, J. and Reed, P. (2007). A computational scaling analysis of multiobjective evolutionary algorithms in long-term groundwater monitoring applications. *Advances in Water Resources*, 30(3):408–419.

Korteling, B., Dessai, S., and Kapelan, Z. (2013). Using information-gap decision theory for water resources planning under severe uncertainty. *Water resources management*, 27(4):1149–1172.

Kuczera, G. (1992). Water supply headworks simulation using network linear programming. *Advances in Engineering Software*, 14(1):55–60.

Kwadijk, J. C. J., Haasnoot, M., Mulder, J. P. M., Hoogvliet, M. M. C., Jeuken, A. B. M., Krogt, R. A. A. v. d., Oostrom, N. G. C. v., Schelfhout, H. A., Velzen, E. H. v., Waveren, H. v., and Wit, M. J. M. d. (2010). Using adaptation tipping points to prepare for climate change and sea level rise: a case study in the Netherlands. *Wiley Interdisciplinary Reviews: Climate Change*, 1(5):729–740.

Kwakkel, J. and Haasnoot, M. (2015). How robust is a robust policy? A comparative analysis of alternative robustness metrics for supporting robust decision analysis. In *EGU General Assembly Conference Abstracts*, volume 17 of *EGU General Assembly Conference Abstracts*, page 8937.

Kwakkel, J. H. (2017). The exploratory modeling workbench: An open source toolkit for exploratory modeling, scenario discovery, and (multi-objective) robust decision making. *Environmental Modelling & Software*, 96:239–250.

Kwakkel, J. H. and Haasnoot, M. (2019). *Supporting DMDU: A Taxonomy of Approaches and Tools*, pages 355–374. Springer International Publishing, Cham.

Kwakkel, J. H., Haasnoot, M., and Walker, W. E. (2014). Developing dynamic adaptive policy pathways: a computer-assisted approach for developing adaptive strategies for a deeply uncertain world. *Climatic Change*, 132(3):373–386.

Kwakkel, J. H., Haasnoot, M., and Walker, W. E. (2016a). Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty. *Environmental Modelling & Software*, 86:168–183.

Kwakkel, J. H. and Jaxa-Rozen, M. (2016). Improving scenario discovery for handling heterogeneous uncertainties and multinomial classified outcomes. *Environmental Modelling and Software*, 79:311–321.

Kwakkel, J. H., Walker, W. E., and Haasnoot, M. (2016b). Coping with the Wickedness of Public Policy Problems: Approaches for Decision Making under Deep Uncertainty. *Journal of Water Resources Planning and Management*, 142(3):01816001.

Labadie, J. (2011). Modsim-dss water rights planning water resources management & river operations decisión support system. *Documentation and User Manual. Colorado State University*.

Lall, U. (1995). Recent advances in nonparametric function estimation: Hydrologic applications. *Reviews of Geophysics*, 33(S2):1093–1102.

Laumanns, M., Thiele, L., Deb, K., and Zitzler, E. (2002). Combining convergence and diversity in evolutionary multiobjective optimization. *Evolutionary computation*, 10(3):263–282.

Lempert, R. J. (2002). A new decision sciences for complex systems. *Proceedings of the National Academy of Sciences*, 99(Supplement 3):7309–7313.

Lempert, R. J., Bryant, B. P., and Bankes, S. C. (2008). Comparing Algorithms for Scenario Discovery. Technical Report WR-557-NSF, RAND.

Lempert, R. J. and Collins, M. (2007). Managing the risk of an uncertain threshold response: Comparison of robust, optimimum, and precautionary approaches. *Risk Analysis*, 27(4):1009–1026.

Lempert, R. J., Groves, D. G., Popper, S. W., and Bankes, S. C. (2006). A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios. *Management Science*, 52(4):514–528.

Lettenmaier, D. P. and Burges, S. J. (1978). Climate change: Detection and its impact on hydrologic design. *Water Resources Research*, 14(4):679–687.

Liu, Y., Gupta, H., Springer, E., and Wagener, T. (2008). Linking science with environmental decision making: Experiences from an integrated modeling approach to supporting sustainable water resources management. *Environmental Modelling and Software*, 23(7):846–858.

Lopez-Nicolas, A., Pulido-Velazquez, M., Rougé, C., Harou, J., and Escriva-Bou, A. (2018). Design and assessment of an efficient and equitable dynamic urban water tariff. application to the city of valencia, spain. *Environmental Modelling & Software*, 101:137–145.

Loucks, D. P. and Da Costa, J. R. (2013). *Decision support systems: Water resources planning*, volume 26. Springer Science & Business Media.

Loucks, D. P. and van Beek, E. (2017). *Water Resource Systems Planning and Management*. Springer.

Lownsbery, K. E. (2014). Quantifying the Impacts of Future Uncertainties on the Apalachicola-Chattahoochee-Flint Basin. Master's thesis, University of Massachusetts Amherst.

Lund, J. (2013). Some curious things about water management. *Journal of Water Resources Planning and Management*, 139(1):1–2.

Lund, J. R. and Israel, M. (1995). Water transfers in water resource systems. *Journal of Water Resources Planning and Management*, 121(2):193–204.

Maass, A., Hufschmidt, M., Dorfman, R., Thomas Jr, H., Marglin, S., and Fair, G. (1962). Design of water resource system. *Soil Science*, 94(2):135.

Maier, H., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L., Cunha, M., Dandy, G., Gibbs, M., Keedwell, E., Marchi, A., Ostfeld, A., Savic, D., Solomatine, D., Vrugt, J., Zecchin, A., Minsker, B., Barbour, E., Kuczera, G., Pasha, F., Castelletti, A., Giuliani, M., and Reed, P. (2014). Evolutionary algorithms and other metaheuristics in water resources: Current status, research challenges and future directions. *Environmental Modelling & Software*, 62:271–299.

Martins, E. S. and Stedinger, J. R. (2000). Generalized maximum-likelihood generalized extreme-value quantile estimators for hydrologic data. *Water Resources Research*, 36(3):737–744.

Matrosov, E. S., Harou, J. J., and Loucks, D. P. (2011). A computationally efficient open-source water resource system simulator–application to london and the thames basin. *Environmental Modelling & Software*, 26(12):1599–1610.

Matrosov, E. S., Huskova, I., Kasprzyk, J. R., Harou, J. J., Lambert, C., and Reed, P. M. (2015). Many-objective optimization and visual analytics reveal key trade-offs for londons water supply. *Journal of Hydrology*, 531:1040–1053.

McNabb, D. E. (2019). The population growth barrier. In *Global Pathways to Water Sustainability*, pages 67–81. Springer International Publishing.

McPhail, C., Maier, H. R., Kwakkel, J. H., Giuliani, M., Castelletti, A., and Westra, S. (2018). Robustness metrics: How are they calculated, when should they be used and why do they give different results? *Earthś Future*, 6(2):169–191.

Miller, B. and Goldberg, D. (1996). ptimal sampling for genetic algorithms. Intelligent Engineering Systems Through Artificial Neural Networks. *ANNIE conference*.

Milly, P. C. D., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., and Stouffer, R. J. (2008). Stationarity is dead: Whither water management? *Science*, 319(5863):573–574.

Moallemi, E. A., Zare, F., Reed, P. M., Elsawah, S., Ryan, M. J., and Bryan, B. A. (2019). Structuring and evaluating decision support processes to enhance the robustness of complex human–natural systems. *Environmental Modelling & Software*, page 104551.

Moody, P. and Brown, C. (2013). Robustness indicators for evaluation under climate change: Application to the upper Great Lakes. *Water Resources Research*, 49(6):3576–3588.

Moody's (2017). Rating methodology: Us municipal utility revenue debt. techreport, Moody's.

Moody's (2019). Research announcement: Moody's: Us regulated utilities sector outlook for 2019 remains negative. techreport, Moody's.

Morris, M. D. (1991). Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2):161–174.

Mortazavi-Naeini, M., Kuczera, G., and Cui, L. (2014). Application of multiobjective optimization to scheduling capacity expansion of urban water resource systems. *Water Resources Research*, 50(6):4624–4642.

Mozenter, Z. D., Yates, A. J., Schnier, K. E., Hughes, J. A., and Characklis, G. W. (2018). Understanding water utility attitudes toward water transfers and risk: Pretest results. *Journal of Water Resources Planning and Management*, 144(4):06018002.

Murphy, K. P. (2012). *Machine Learning*. MIT Press Ltd.

Nair, S. and Howlett, M. (2017). Policy myopia as a source of policy failure: adaptation and policy learning under deep uncertainty. *Policy & Politics*, 45(1):103–118.

NCDENR (2002). Jordan Lake Water Supply Storage Allocation Round Three Hearing Officers' Report. Technical report, North Carolina Department of Environment and Natural Resources, Division of Water Resources.

Nicklow, J., Reed, P., Savic, D., Dessalegne, T., Harrell, L., Chan-Hilton, A., Karamouz, M., Minsker, B., Ostfeld, A., Singh, A., and Zechman, E. (2010). State of the Art for Genetic Algorithms and Beyond in Water Resources Planning and Management. *Journal of Water Resources Planning and Management*, 136(4):412–432.

Olmstead, S. M. (2010). The economics of managing scarce water resources. *Review of Environmental Economics and Policy*, 4(2):179–198.

Orange Water and Sewer Authority (2010). Water Shortage Response Plan.

Padula, S., Harou, J. J., Papageorgiou, L. G., Ji, Y., Ahmad, M., and Hepworth, N. (2013). Least economic cost regional water supply planning – optimising infrastructure investments and demand management for south east england's 17.6 million people. *Water Resources Management*.

Palmer, R. N. and Characklis, G. W. (2009). Reducing the costs of meeting regional water demand through risk-based transfer agreements. *Journal of Environmental Management*, 90.

Pareto, V. (1896). *Cours D'Economie Politique*. Rouge, Lausanne.

Paulson, E., Badruzzaman, M., Triana, E., Cherchi, C., Stewart, N., Hsin Sun, Y., and Jacangelo, J. (2018). Framework for evaluating alternative water supplies:

Balancing cost with reliability, resilience and sustainability. Technical report, Water Research Foundation.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Potter, K. W. (1976). Evidence for nonstationarity as a physical explanation of the Hurst Phenomenon. *Water Resources Research*, 12(5):1047–1052.

Quinn, J. D., Reed, P. M., Giuliani, M., and Castelletti, A. (2017). Rival framings: A framework for discovering how problem formulation uncertainties shape risk management trade-offs in water resources systems. *Water Resources Research*, 53(8):7208–7233.

Quinn, J. D., Reed, P. M., Giuliani, M., Castelletti, A., Oyler, J. W., and Nicholas, R. E. (2018). Exploring how changing monsoonal dynamics and human pressures challenge multireservoir management for flood protection, hydropower production, and agricultural water supply. *Water Resources Research*.

Reed, P. M. and Hadka, D. (2014). Evolving Many-Objective Water Management to Exploit Exascale Computing. *Water Resources Research*, 50:8367–8373.

Reed, P. M., Hadka, D., Herman, J. D., Kasprzyk, J. R., and Kollat, J. B. (2013). Evolutionary Multiobjective Optimization in Water Resources: The Past, Present and Future. *Advances in Water Resources*, 51:438–456.

Reis, J. and Shortridge, J. (2019). Impact of uncertainty parameter distribution on robust decision making outcomes for climate change adaptation under deep uncertainty. *Risk Analysis*.

Saltelli, A. (2002). Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145(2):280–297.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Galtelli, D., Saisana, M., and Tarantola, S. (2008). *Global Sensitivity Analysis: The Primer*. Wiley, West Sussex, England.

Saltelli, A., Ratto, M., Tarantola, S., and Campolongo, F. (2006). Sensitivity analysis practices: Strategies for model-based inference. *Reliability Engineering and System Safety*, 91(10):1109–1125.

Schapire, R. (1999). A short introduction to boosting. *Japanese Society For Artificial Intelligence*.

Schapire, R. E. (1990). The strength of weak learnability. *Mach. Learn.*, 5(2):197–227.

Schneller, G. O. and Sphicas, G. P. (1983). Decision making under uncertainty: Starr's domain criterion. *Theory and Decision*, 15(4):321–336.

Scruton, R. (2013). *The Uses of Pessimism: And the Danger of False Hope*. OXFORD UNIV PR.

Shafer, K. and Fox, R. (2017). An equitable water future, a national briefing paper. techreport, US Water Alliance.

Shah, A. and Ghahramani, Z. (2016). Pareto frontier learning with expensive correlated objectives. In *International Conference on Machine Learning*, pages 1919–1927.

Sheer, D. P. (2010). Dysfunctional Water Management: Causes and Solutions. *Journal of Water Resources Planning and Management*, 136(1):1–4.

Sieber, J. (2006). Weap water evaluation and planning system.

Sieber, J. and Purkey, D. (2015). *WEAP: Water Evaluation And Planning System. User Guide for WEAP.* Stockholm Environment Institute, U.S. Center, 11 Curtis Avenue, Somerville, MA 02144 USA.

Simon, H. A. (1959). Theories of decision-making in economics and behavioral science. *The American Economic Review*, 49(3):253–283.

Singh, A. and Minsker, B. S. (2008). Uncertainty-based multiobjective optimization of groundwater remediation design. *Water Resources Research*, 44(2):n/a–n/a. W02404.

Singh, R., Reed, P. M., and Keller, K. (2015). Many-objective robust decision making for managing an ecosystem with a deeply uncertain threshold response. *Ecology and Society*, 20(3).

Smalley, J. B., Minsker, B. S., and Goldberg, D. E. (2000). Risk-based in situ bioremediation design using a noisy genetic algorithm. *Water Resources Research*, 36(10):3043–3052.

Sobol, I. (2001). Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation*, 55(1-3):271–280.

Stallman, R. M. and Community, G. D. (2017). *Using the GNU Compiler Collection For GCC Version 7.2.0*. GNU Press a division of the Free Software Foundation, 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Starr, M. K. (1962). *Product Design and Decision Theory*. Prentice-Hall, Englewood Cliffs, N.J.

Stedinger, J. R. (1993). Frequency analysis of extreme events. *in Handbook of Hydrology.*

Stedinger, J. R. and Taylor, M. R. (1982). Synthetic streamflow generation: 1. Model verification and validation. *Water Resources Research*, 18(4):909–918.

Storn, R. and Price, K. (1997). Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4):341–359.

Tang, Y., Reed, P., and Kollat, J. (2007). Parallelization strategies for rapid and robust evolutionary multiobjective optimization in water resources applications. *Advances in Water Resources*, 30(3):335–353.

The City of Raleigh Public Utilities Department (2011). City of Raleighś Water Shortage Response Plan.

Tian, X., Galelli, S., and de Neufville, R. (2018). Impact of operating rules on planning capacity expansion of urban water supply systems. *Urban Water Journal*, 15(7):654–661.

Triangle J, C. o. G. (2014). Triangle Regional Water Supply Plan. Technical report, Jordan Lake Partnership.

Trindade, B., Reed, P., and Characklis, G. (2019). Deeply uncertain pathways: Integrated multi-city regional water supply infrastructure investment and portfolio management. *Advances in Water Resources*, page 103442.

Trindade, B., Reed, P., Herman, J., Zeff, H., and Characklis, G. (2017). Reducing regional drought vulnerabilities and multi-city robustness conflicts using many-objective optimization under deep uncertainty. *Advances in Water Resources*, 104:195–209.

Tsoukias, A. (2008). From decision theory to decision aiding methodology. *European Journal of Operational Research*, 187(1):138–161.

Tsutsui, S., Yamamura, M., and Higuchi, T. (1999). Multi-parent recombination with simplex crossover in real coded genetic algorithms. In *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-Volume 1*, pages 657–664. Morgan Kaufmann Publishers Inc.

USWA, U. S. W. A. (2019a). Fourth annual value of water index, value of water campaign.

USWA, U. S. W. A. (2019b). Strengthening utilities through consolidation: The financial impact.

Vaux Jr, H. J. and Howitt, R. E. (1984). Managing water scarcity: An evaluation of interregional transfers. *Water resources research*, 20(7):785–792.

Vrugt, J. A. and Robinson, B. A. (2007). Improved evolutionary optimization from genetically adaptive multimethod search. *Proceedings of the National Academy of Sciences*, 104(3):708–711.

Walker, W., Harremoës, P., Rotmans, J., Van der Sluijs, J., Van Asselt, M., Janssen, P., and Krayer von Krauss, M. (2003). Defining Uncertainty: A Conceptual Basis for Uncertainty Management in Model-Based Decision Support. *Integrated Assessment*, 4(1).

Walker, W. E., Haasnoot, M., and Kwakkel, J. H. (2013). Adapt or perish: a review of planning approaches for adaptation under deep uncertainty. *Sustainability*, 5(3):955–979.

Walker, W. E., Rahman, S., and Cave, J. (2001). Adaptive policies, policy analysis, and policy-making. *European Journal of Operational Research*, 128(2):282–289.

Wang, T. and de Neufville, R. (2005). Real options "in" projects. *Real Options Annual International Conference*.

Ward, V. L., Singh, R., Reed, P. M., and Keller, K. (2015). Confronting tipping points: Can multi-objective evolutionary algorithms discover pollution control tradeoffs given environmental thresholds? *Environmental Modelling and Software*, 73:27–43.

Watkins, D. W. and McKinney, D. C. (1997). Finding robust solutions to water resources problems. *Journal of Water Resources Planning and Management*, 123(1):49–58.

Watson, A. A. and Kasprzyk, J. R. (2016). Incorporating Deeply Uncertain Factors into the Many Objective Search Process. *Environmental Modelling and Software*. (In-Review).

Welsh, W. D., Vaze, J., Dutta, D., Rassam, D., Rahman, J. M., Jolly, I. D., Wallbrink, P., Podger, G. M., Bethune, M., Hardy, M. J., et al. (2013). An integrated modelling framework for regulated river systems. *Environmental Modelling & Software*, 39:81–102.

Wilchfort, O. and Lund, J. R. (1997). Shortage management modeling for urban water supply systems. *Journal of Water Resources Planning and Management*, 123(4):250–258.

Woodruff, M., Reed, P., and Simpson, T. (2013). Many-Objective Visual Analytics: Rethinking the Design of Complex Engineered Systems. *Structural and Multidisciplinary Optimization*, 48:201–219.

Wu, W., Dandy, G. C., Maier, H. R., Maheepala, S., Marchi, A., and Mirza, F. (2017). Identification of optimal water supply portfolios for a major city. *Journal of Water Resources Planning and Management*, 143(9):05017007.

Yates, D., Sieber, J., Purkey, D., and Huber-Lee, A. (2005). Weap21a demand-, priority-, and preference-driven water planning model: part 1: model characteristics. *Water International*, 30(4):487–500.

Zandvoort, M., Campos, I. S., Vizinho, A., Penha-Lopes, G., Lorencová, E. K., van der Brugge, R., van der Vlist, M. J., van den Brink, A., and Jeuken, A. B. (2017). Adaptation pathways in planning for uncertain climate change: Applications in portugal, the czech republic and the netherlands. *Environmental Science & Policy*, 78:18–26.

Zatarain Salazar, J., Reed, P. M., Herman, J. D., Giuliani, M., and Castelletti, A. (2016). A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. *Advances in Water Resources*, 92:172–185.

Zeff, H., Kasprzyk, J., Herman, J., Reed, P., and Characklis, G. (2014). Navigating Financial and Supply Reliability Tradeoffs in Regional Drought Portfolios. *Water Resources Research*, 50(6):4906–4923.

Zeff, H. B. and Characklis, G. W. (2013). Managing water utility financial risks through third-party index insurance contracts. *Water Resources Research*, 49(8):4939–4951.

Zeff, H. B., Herman, J. D., Reed, P. M., and Characklis, G. W. (2016). Cooperative drought adaptation: Integrating infrastructure development, conservation, and water transfers into adaptive policy pathways. *Water Resources Research*, 52(9):7327–7346.

Zhu, J. and Hastie, T. (2005). Kernel logistic regression and the import vector machine. *Journal of Computational and Graphical Statistics*, 14(1):185–205.