## Appendix 1

The appendix includes:

1.1 A detailed description of the multi-objective optimization procedure.
1.2 The impact of random seeds on the results.
1.3 Runtime dynamics of the multi-objective evolutionary algorithm.

### 1.1 Evolutionary Multiobjective Optimization

We use MOEAs to optimize the three problem formulations (P1, P2, and P3). The optimization finds the optimal strategies for single objective formulations (P1 and P2) and the Pareto approximate front that captures the trade-offs between all objectives for P3. Here, we use the recently developed and benchmarked MOEA, the BORG MOEA (Hadka and Reed 2013). While the advantage of using MOEAs lies in their ability to identify Pareto approximate fronts for multi-objective problems, the BORG MOEA can also be used to solve the single objective formulations such as P1 and P2. Solving the single objective formulations (P1 and P2) yields a single approximately optimal pollution strategy per formulation while solving the multi-objective formulation (P3) yields a wide range of strategies that span the trade-off between various objectives. If $n_{p3}$ solutions form the Pareto approximate set for P3, optimizing P1, P2, and P3 together present the stakeholders with a total of $n_{p3}$+2 potential management strategies to choose from. This process of generation of alternative strategies for stakeholders to compare and contrast forms the first part of MORDM.

MOEAs are being used to solve many-objective problems in a rapidly growing body of literature as their population-based search enables the direct approximation of problems' Pareto frontiers in a single optimization run (Purshouse and Fleming 2003, Fleming et al. 2005, Aguirre et al. 2013, Lygoe et al. 2013, Morino and Obayashi 2013, Reed et al. 2013, Woodruff et al. 2013). For example, while solving a five-objective problem, the MOEAs simultaneously solve the five single-objective problems, ten two-objective problems, ten three-objective problems, and five four-objective problems. This has significant computational advantages over weight-based approaches to solve multi-objective problems that assign different weights to each objective and solve a single objective problem by maximizing the aggregated weighted objective. In order to capture the entire trade-off space, such approaches need to perform the optimization many times by varying the weights, which becomes intractable as the number of objectives goes beyond three (Teytaud 2007).

The BORG MOEA is able to solve high dimensional problems for non-linear threshold based models by adaptively using multiple (in this case six) search strategies. Search strategies refer to the operators that the algorithm uses to search the space of feasible solutions. Most multi-objective optimization algorithms employ a single search operator. But the BORG MOEA simultaneously uses six search strategies and assigns higher probability of use to a search strategy based on its ability to identify solutions on or close to the Pareto approximate front. Other features include random re-start of the search when no significant progress is observed (or the algorithm is trapped in a local optima), epsilon-dominance archiving (Laumanns et al. 2002), adaptive population sizing (Kollat and Reed 2007), and a steady state algorithm structure (Deb et al. 2005). Comparative analysis has demonstrated the efficacy of the BORG MOEA on a range of test problems as well as real world problems many of which require optimization under uncertainty (Hadka and Reed 2012, Reed et al. 2013). The BORG MOEA is also relatively easy-to-use, has an underlying theoretical proof of convergence, and is highly scalable on parallel computing systems, thus increasing its potential usage across a wide variety of disciplines.

The framework used for performing multi-objective optimization of the lake problem is shown in Figure A1.1. In order to optimize under stochastic uncertainty, we randomly sample 100 SOWs out of 10000 SOWs for each evaluation of the lake problem in a function call of the BORG MOEA. In any given evaluation of a lake's pollution strategy (parameter 'a' in Equation 1), the relatively small number of SOWs (100) sampled is likely to yield 'noisy' objective function values. However, the small sample size drastically reduces computational demands. Evolutionary heuristics of the Borg MOEA underlying the Darwinian selection reward those solutions whose performance minimally varies across these small number of samples and have been demonstrated to be capable of maintaining high quality search in uncertain spaces (Miller and Goldberg 1996, Smalley et al. 2000, Reed et al. 2013). In other words, even if the BORG MOEA optimizes the objective function values across only 100 randomly sampled SOWs from the total space of 10000 SOWs at every function call, its underlying structure enables it to identify strategies that perform well across the entire ensemble of uncertain futures.

This framework therefore allows for multi-objective optimization under uncertainty (i.e., identifies a strategy that performs well over many SOWs drawn from the baseline well-characterized uncertainty distribution). The robustness analysis under well-characterized uncertainty tests each solution against all 10000 SOWs, serving as a validation of this computational savings using small number of samples. Moreover, all results presented in the study are objective performance re-evaluated against all 10000 SOWs. We employ a parallelized multi-master version of the BORG MOEA that runs on a cluster with 8 islands (different evolving populations that search the space) across 64 nodes to speed up the optimization procedure. The algorithm's search operator parameters were set at default values based on recommendations discussed in (Hadka and Reed 2013).

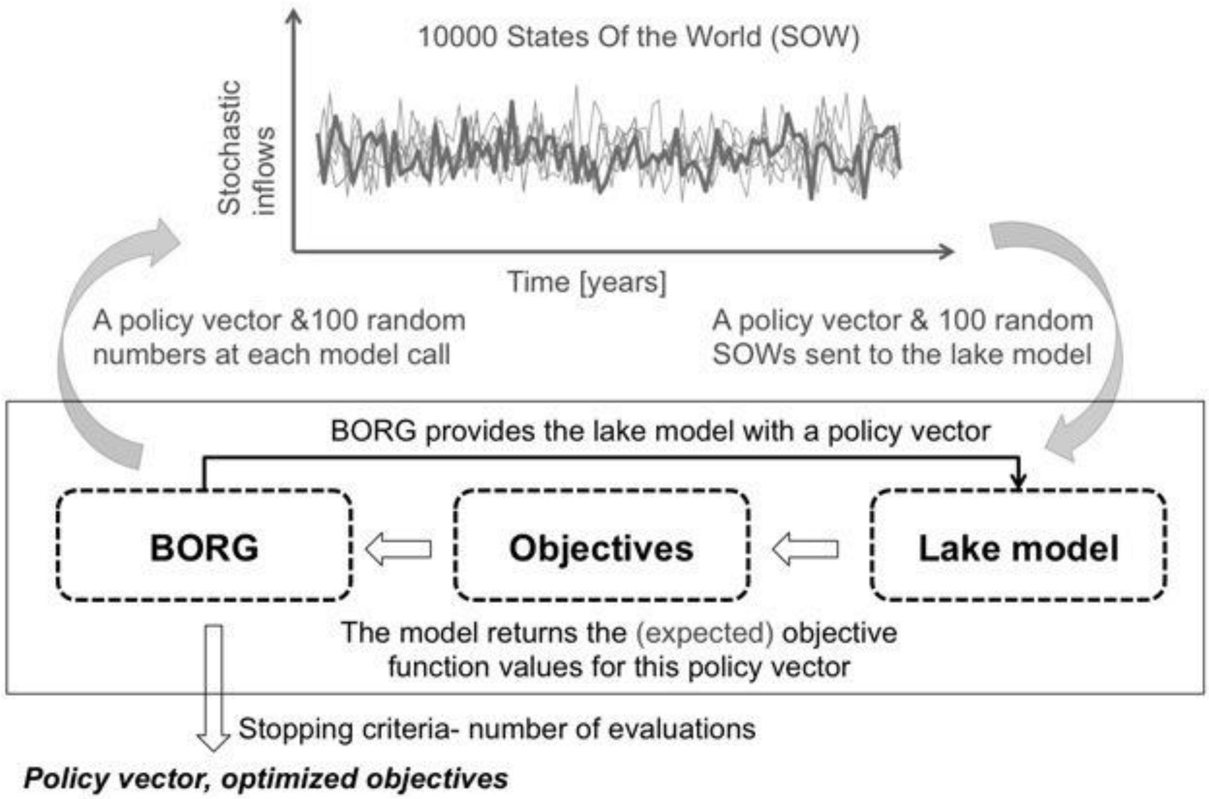


Figure A1.1 The robust optimization framework implemented using the BORG MOEA. This framework is used to optimize the various formulations of the lake model. The procedure within the solid black rectangle is repeated until a stopping criterion is met. Gray components show how the BORG MOEA implements optimization in the presence of stochastic uncertainty.

1.2 Impact of Random Seeds on Resulting Compromise Strategy

This analysis was carried out in order to establish the reliability of the results, i.e., to ensure that the conclusions of the study are independent of a random start of the algorithm.

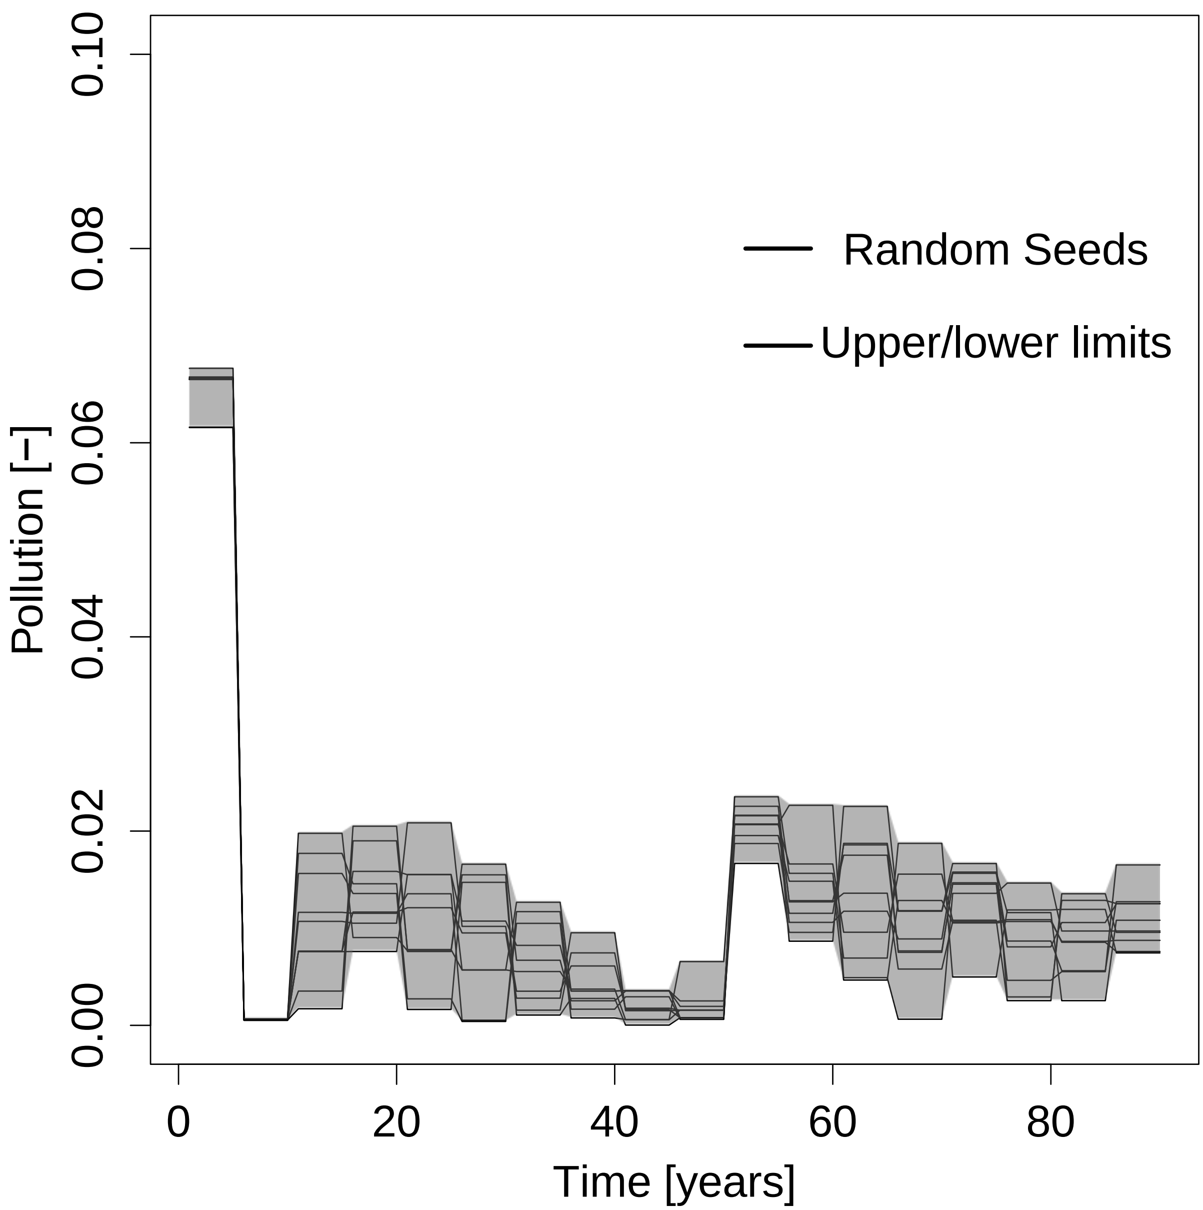


Figure A1.2 Analyzing the impact of random seeds on the results. Using ten random seeds to start the BORG MOEA, ten different Pareto approximate fronts are obtained for the stochastic multi-objective formulation (P3). Each Pareto approximate front is evaluated to arrive at the *compromise* strategy using the minimum tolerable windows approach. Each strategy is associated with five objectives - expected utility, utility of current generation, utility of future generation, phosphorus in the lake, reliability. The strategy that maximizes the minimum objective among the across all strategies is identified as the *compromise* strategy. Figure shows the identified *compromise* strategies across ten random seeds, and the envelope bounding their lower and upper limits.

1.3 Runtime Dynamics of the BORG MOEA

The appendix presents the runtime dynamics of the algorithm. Runtime dynamics show that the algorithm converged to the Pareto approximate set within one million function evaluations.

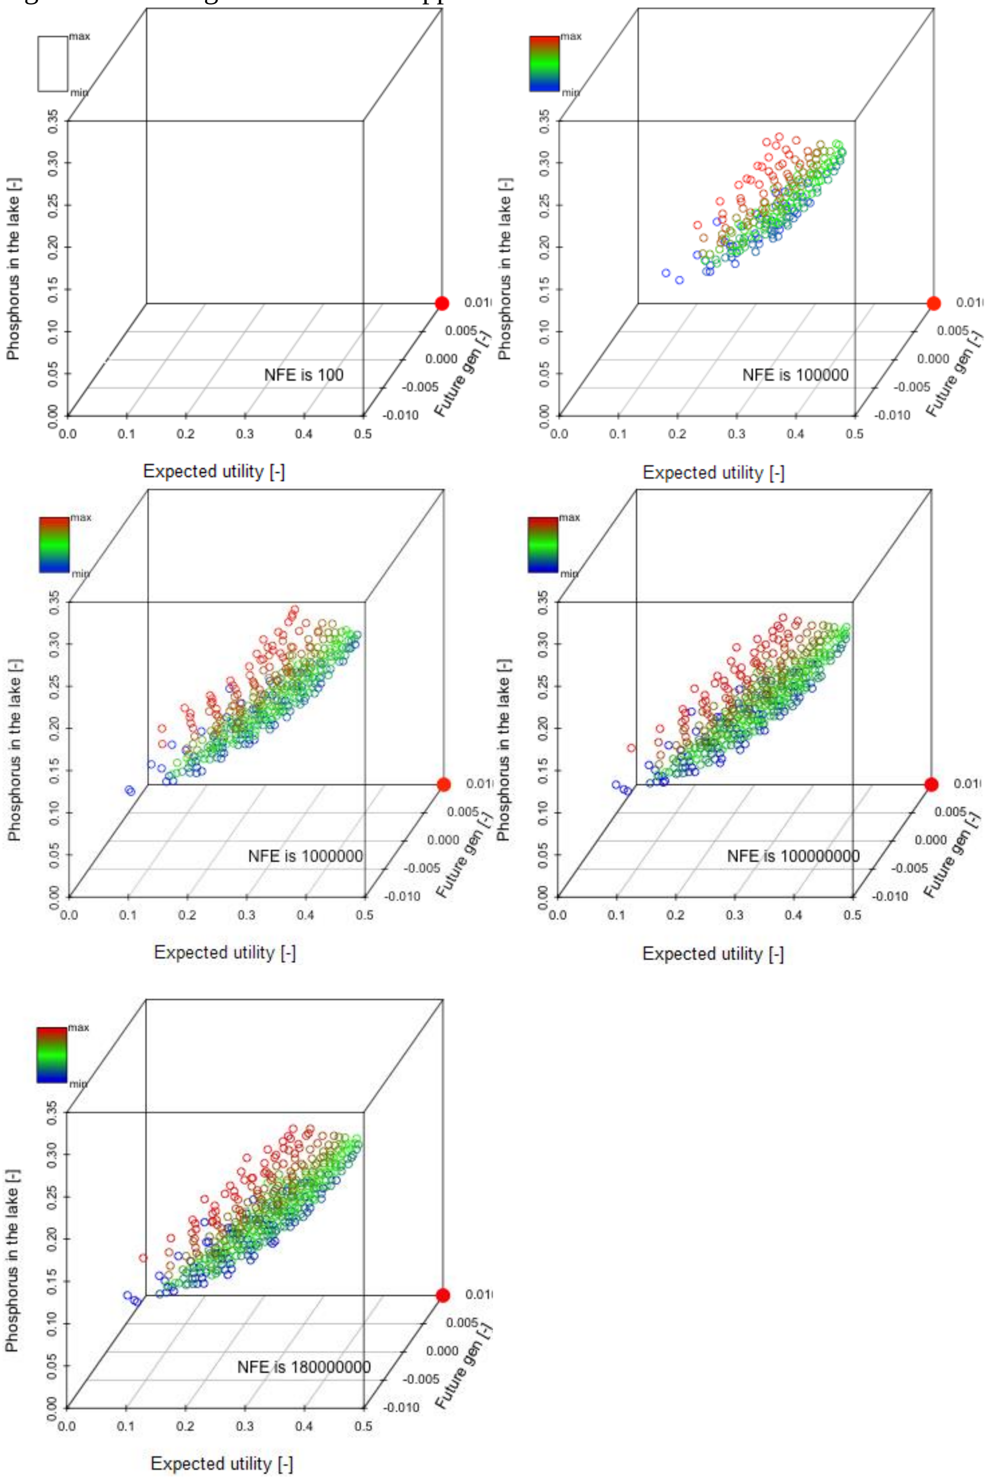


Figure A1.3 Runtime dynamics for the BORG MOEA - The figure shows the evolution of the Pareto approximate front with number of function evaluations (NFE) as the BORG MOEA explores the objective space. Each plot shows expected utility on the x-axis, utility of future generations on the y-axis, amount of phosphorus in the lake on the z-axis. The color represents the utility of the first generation. Reliability objective is not shown here. The ideal point is shown as the red point on the bottom right corner. Snapshots of the algorithm's search are plotted at the following NFEs – one hundred, hundred thousand, one million, hundred million, and 180 million. The figure shows that the algorithm converges to the Pareto approximate front within 1 million function evaluations. For the results presented in this study, we ran the parallel version of the algorithm on eight nodes, each node running approximately 180 million evaluations.

## References


Aguirre, H., A. Oyana, and K. Tanaka. 2013. Adaptive ε-Sampling and ε-Hood for Evolutionary Many-Objective Optimization.*in* R. C. Purshouse, P. J. Fleming, C. M. Fonseca, S. Greco, and J. Shaw, editors. Evolutionary Multi-Criterion Optimization. Springer Berlin Heidelberg, Sheffield.

Deb, K., M. Mohan, and S. Mishra. 2005. Evaluating the epsilon-Domination Based Multi-Objective Evolutionary Algorithm for a Quick Computation of Pareto-Optimal Solutions. Evolutionary Computation 13:501-525.

Fleming, P., R. Purshouse, and R. Lygoe. 2005. - Many-Objective Optimization: An Engineering Design Perspective. - 3410.

Hadka, D., and P. Reed. 2012. Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. Evol. Comput. 20:423-452.

Hadka, D., and P. Reed. 2013. Borg: An auto-adaptive many-objective evolutionary computing framework. Evol. Comput. 21:231-259.

Kollat, J. B., and P. Reed. 2007. A framework for Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization (VIDEO). Environmental Modelling & Software 22:1691-1704.

Laumanns, M., L. Thiele, K. Deb, and E. Zitzler. 2002. Combining convergence and diversity in evolutionary multiobjective optimization. Evol. Comput. 10:263-282.

Lygoe, R., M. Cary, and P. Fleming. 2013. A Real-World Application of a Many-Objective Optimisation Complexity Reduction Process. 7811.

Miller, B. L., and D. E. Goldberg. 1996. Optimal Sampling for Genetic Algorithms. ASME Press.

Morino, H., and S. Obayashi. 2013. Knowledge Extraction for Structural Design of Regional Jet Horizontal Tail Using Multi-Objective Design Exploration (MODE). 7811:656-668.

Purshouse, R. C., and P. J. Fleming. 2003. Evolutionary many-objective optimisation: an exploratory analysis. Pages 2066-2073 Vol.2063 *in* Evolutionary Computation, 2003. CEC '03. The 2003 Congress on.

Reed, P. M., D. Hadka, J. D. Herman, J. R. Kasprzyk, and J. B. Kollat. 2013. Evolutionary multiobjective optimization in water resources: The past, present, and future. Advances in Water Resources 51:438-456.

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

Teytaud, O. 2007. On the hardness of offline multi-objective optimization. Evol. Comput. 15:475-491.

Woodruff, M., P. Reed, and T. Simpson. 2013. Many objective visual analytics: rethinking the design of complex engineered systems. Structural and Multidisciplinary Optimization 48:201-219.