Environmental Modelling and Software 113 (2019) 59–72

Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

How can learning-by-doing improve decisions in stormwater management? A Bayesian-based optimization model for planning urban green infrastructure investments

Fengwei Hung*, Benjamin F. Hobbs

Department of Environmental Health & Engineering, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD, USA

ARTICLE INFO

Keywords:
Adaptive management
Green infrastructure
Learning
Stormwater management
Stochastic programming

ABSTRACT

Urban stormwater management is shifting its attention from traditional centralized engineering solutions to a distributed and greener approach, namely Green Infrastructure (GI). However, uncertainties concerning GI's efficacy for reducing runoff and pollutants are a barrier to the adoption of GI. One strategy to deal with the uncertainty is to implement GI adaptively, in which stormwater managers can learn and adjust their plans over time to avoid undesired outcomes. We propose a new class of GI planning methods based on two-stage stochastic programming and Bayesian learning, which accounts for projected information gains and decision makers' objectives and willingness to accept risk. In the hypothetical example, the model identifies four categories of investment strategies and quantifies their benefits and costs: all-in, greedy investment plus deferral, mixed investments plus deferral, and learn-and-adjust. Which strategy is optimal depends on the user's risk attitudes, and the alternatives' costs and risks.

1. Introduction

Green infrastructure (GI), sometimes referred to as "Low Impact Development" or "Best Management Practices," is a distributed approach that reduces urban stormwater runoff through on-site infiltration, storage, and evaporation to improve water quality in downstream water courses. Examples of GI practices are rain gardens, rain barrels, tree trenches, permeable pavement, and green roofs. GI is gaining popularity due to its potential social and economic benefits (Environmental Protection Agency, 2010). Among major US cities, Philadelphia is the first that has committed to GI as the primary solution for stormwater problems (PWD, 2011). Starting in 2011, Philadelphia plans to invest $2.4 billion on GI over 25 years. Other cities, such as Washington DC, Chicago, Portland, and Seattle, are also implementing GI at site or neighborhood scales (Environmental Protection Agency, 2010).

However, the runoff and pollution reductions provided by GIs depend on, for example, their specific designs, local climate characteristics, underlying soil properties, and vegetation (Montalto et al., 2011). Also, GI performance may deteriorate with time. Maintenance costs and the upstream and downstream relationships of multiple GIs add further to the uncertainty of the efficacy and expense of GI, which complicates stormwater management planning, and can result in a failure to meet management goals (Chocat et al., 2007; Roy et al., 2008).

Researchers have developed tools for designing, sizing and selection of GI for stormwater management at the site scale (Lee et al., 2005; Loáiciga et al., 2013; Massoudieh et al., 2017; Morales-Torres et al., 2016; Zhen et al., 2004), as well as watershed-scale investment planning models for managing water quality and runoff volume (McGarity, 2012; Montaseri et al., 2015; Sebti et al., 2016). Yet uncertainty has not been recognized by these optimization frameworks. Because of the uncertain interactions of GI, climate, and human activities, it is crucial to recognize risks and the opportunity to learn when planning (Barton et al., 2012; Williams and Brown, 2014); i.e., to manage adaptively.

An experienced planner or stormwater engineer may have a good sense of the extent of potential learning from experimentation, such as field trials of GI. However, we argue that formalizing the learning process and estimating the value of learning can be worthwhile even when quantification is difficult. This has not been done in previous GI optimization models. In this paper, we propose a new method for GI investment planning based on two-stage stochastic programming (TSP), in which we incorporate projected information gains and decision makers' risk attitudes and objectives. The projected knowledge gains are assumed to be a function of the amount of investments of in GI. These gains are, called "learning" throughout this paper, and are used to update our knowledge beliefs concerning of the GI cost-effectiveness

* Corresponding author. Ames Hall 313, 3400 N Charles St, Baltimore, MD, 21218, USA.
E-mail address: hfengwe1@jhu.edu (F. Hung).

https://doi.org/10.1016/j.envsoft.2018.12.005
Received 29 December 2017; Received in revised form 1 November 2018; Accepted 18 December 2018
Available online 21 December 2018

in the later stages based on the idea of Bayesian Inference.

To our knowledge, this is the first time that risk aversion and learning have been included in dynamic optimization models for adaptive stormwater management. This method aims at providing guidance regarding the choice of GI practices, investment timing and amount, risks of failing to meet management goals, and maximization of benefits (including social and economic benefits).

In the hypothetical example in Section 4, the goal of the two-stage decision problem is to maximize cumulative stormwater reductions given a fixed budget for the 40-year planning horizon, where early investments (at year 1) may result in learning that updates our understanding of the GI cost-effectiveness in later decision stage (at year 11). In the example, we identify four broad classes of strategies, which we call ‘all-in’, ‘greedy investment plus deferral’, ‘mixed investment plus wait-and-see’, and ‘learn-and-adjust.’ This range of possibilities involves tradeoffs between cost, expected stormwater reduction and risks. The best choice in a given situation depends both on case study-specific parameters and the decision maker's willingness to accept risks.

The remainder of this paper is organized as follows. In the next section, we summarize the literature on adaptive management, risks associated with GI, and the use of stochastic programming to model learning. There are indeed previous models for water resource management that include uncertainty and risk aversion and models for managing ecosystems and climate change that represent learning; yet models that simultaneously consider both risk-aversion and learning for adaptive management have not been proposed or implemented in stormwater management or other related fields. Section 3 explains the building blocks of the proposed method and illustrates the ideas with a decision tree example. Section 4 shows the results from a hypothetical example to illustrate the method which is formulated as Mixed Integer Linear Programming with risk constraints. Sections 5 discusses the sensitivity of the solutions and presents conclusions. Finally, the mathematical formulation is presented in the Appendix.

2. Literature review

2.1. Adaptive stormwater management

Adaptive management is a framework for long-term planning under uncertainty. It was first proposed by Holling (1978) for environmental assessment and management, and by Walters and Hilborn (1976) in ecological management. It is now widely recognized and applied in those contexts (Lovell and Taylor, 2013; McCarthy and Possingham, 2007; Southwell et al., 2016; Williams, 2011).

Adaptive management can be characterized as a process of “learning to manage by managing to learn” (Bormann et al., 1994), in which decisions at each planning stage must balance the cost and benefits to future decisions of learning to reduce uncertainty, versus investing to yield immediate benefits (Williams, 2011). This paper presents a general methodology for quantitative implementation of adaptive management that is designed to provide useful information in the form of quantified risks and the costs and benefits of learning.

2.2. Risks in stormwater management with GI

Researchers have recognized that there exist risks in managing stormwater and have tried to address them. Chocat et al. (2007) pointed out three particular types of risks involved with GI:

1. transferring maintenance responsibility to end-users may result in degraded performance,
2. uncertainties in the transition path from the centralized system to decentralized systems, and
3. residual contamination in stormwater retained and treated by GI, which could cause human health concerns.

In addition, Roy et al. (2008) flagged the risks of underestimating costs of GI installation and maintenance, missing management goals, loss of functionality, and community resistance to installations, particularly in municipalities of lower socio-economic status with relatively few natural environmental assets.

However, quantifying the risks associated with stormwater management is challenging. A few studies have tried to estimate the risks of stormwater pollution, for example, stormwater pollution risks to ecosystems (Novotny and Witte, 1997), overflow risks from detention basins (Guo, 2002), and sediment pollution risks from combined sewer overflows (CSOs) in receiving waters (Rossi et al., 2005). Yet neither optimization models nor decision support tools, in general, have been developed to manage risks or evaluate adaptive strategies in stormwater management.

2.3. Stochastic programming applications and learning

Two-stage (or, more generally, multi-stage) stochastic programming (TSP) is a logical way to model the implementation of adaptive management, as it can analyze “here-and-now” investments while considering opportunities to change course in future stages, depending on what is learned and their probabilities (Shapiro and Philpott, 2007). Generalizations of TSP, including stochastic dynamic programming and partially observable Markov decision processes, describe sequential decision processes where decisions are made in each stage based on the decision maker's beliefs concerning the underlying state of nature and the expected value of some objective (Monahan, 1982; White, 1991). Chance-constrained and risk-averse programming are other optimization approaches that use probability distributions and can be compatible with TSP (Housh et al., 2013; Krokhmal et al., 2001; Wang and Huang, 2014).

Other approaches to decision making under uncertainty attempt to avoid the use of probabilities. One is robust decision making (RDM), which was developed for the problems with deep uncertainty (i.e., uncertainty is highly subjective such that different experts’ estimates of probability distributions may be very different from each other) (Lempert and Collins, 2007). RDM generates plausible scenarios and attempts to identify strategies that work well across all scenarios, instead of using a probability distribution and an objective based on (probability-weighted) expected value. RDM applications can be found in the climate change and water supply literature (Beh et al., 2017; Kasprzyk et al., 2013; Lan et al., 2015; Mortazavi-Naeini et al., 2015). However, RDM does not naturally lend itself to quantifying the value of learning in adaptive management, so our approach instead relies on TSP.

Risk aversion, in which decision makers place greater weight on poor outcomes compared to use of expected (probability-weighted) values, has been considered in many environmental management problems (e.g., Baker, 2009; Chao and Hobbs, 1997; Piantadosi et al., 2008). The particular implementation of risk aversion that we use, conditional value at risk (CVaR, explicitly defined below), has been applied in water supply allocation and storage problems to explore risk-expected value tradeoffs (Paydar and Qureshi, 2012; Piantadosi et al., 2008; Webby et al., 2007), but not in water quality management. Ours is the first use of CVaR to address risk in water pollution investments.

Learning has been considered in a wide range of environmental applications, usually in the form of Bayesian analysis. Kim et al. (2003) and Jacobi et al. (2013) applied stochastic dynamic programming to fishery and water quality management, respectively, where research and monitoring (learning) are optimized in the first stage followed by decisions concerning land and ecosystem management actions in the second stage. Harrison (2007) developed a Bayesian programming method for water quality control problems. Varouchakis et al. (2016) applied Bayesian decision analysis in reservoir construction planning problem. A number of analyses of uncertainty and learning can be found in the climate policy literature (see Golub et al., 2014). For

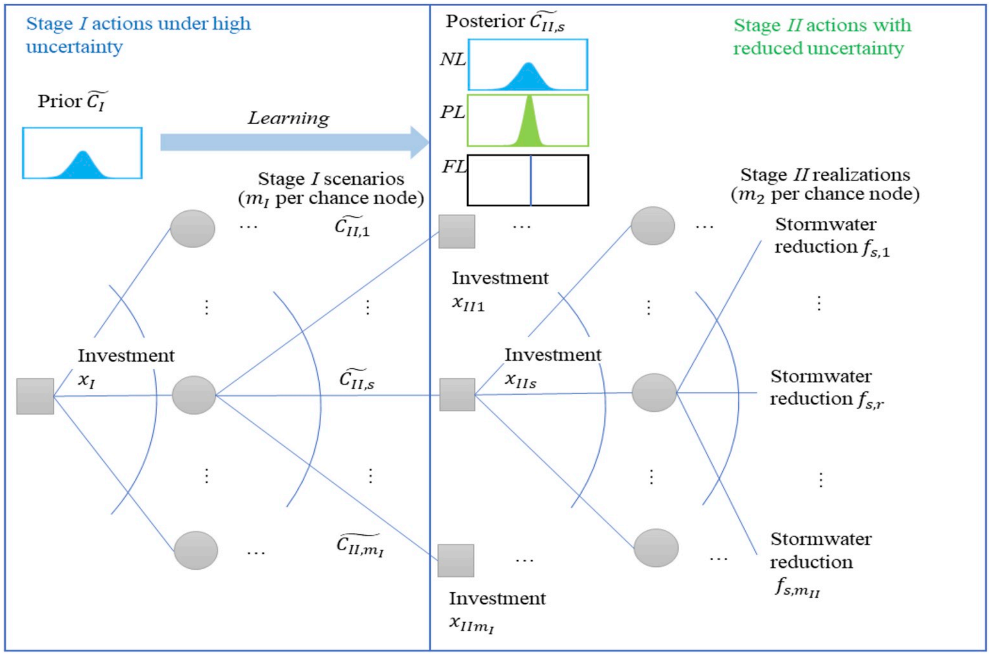

Fig. 1. Schematic decision tree example of two-stage decision making with uncertainty and learning to reduce uncertainty.

instance, Baker (2009) presented a TSP model for optimizing carbon emissions abatement with learning about uncertain climate change damages, while Webster et al. (2017) applied TSP to consider optimal research into climate mitigation technologies whose performance is uncertain. In their stochastic dynamic programming models, the state variables (e.g., the belief of climate change and mean water level) are updated by Bayes' law with new observations, which is the type of "learning" we implement in our models below.

In summary, we have found that stochastic programming has been used to implement adaptive management in related contexts, but not in GI planning itself. Learning has been addressed, but the construction of the likelihood function is often challenging, which limits the applicability of Bayesian approaches. The method presented in the following section uses simple "learning curve functions" to depict the relationship of the "here-and-now" GI investment and learning, which provides an intuitive way to express the amount of learning while avoiding the need to explicitly construct likelihood functions.

3. Methodology

The primary objective of urban stormwater management is to reduce runoff in order to prevent sewer overflows and improve water quality in streams. While the problem can be formulated as a multi-objective optimization, for simplicity, our formulation considers only one objective to be optimized (i.e., runoff reduction), subject to constraints on the other objectives (e.g., a fiscal budget and the risk of undesired outcomes). In this section, we only explain the general scheme of the method, and present the full mixed integer linear program implementation of the TSP in the Appendix.

Section 3.1 presents a decision tree example that we use to illustrate the underlying logic of the method for representing learning and adaptive GI investment planning. Section 3.2 presents the basic stochastic programming framework, and then explains the modeling of risk-aversion and learning in that framework.

3.1. Decision tree representation of the adaptive GI investment planning model

In the case of GI planning with two investment decision stages, denoted by Stages *I* and *II*, the manager needs to make investment planning with several distinct GIs whose performance (stormwater reductions) are uncertain and independent. We summarize the manager's problem as follows:

- *Decisions:* choose the GI investment portfolio in Stage *I* given a budget shared with both stages while considering the opportunity to change course in Stage *II* in order to:
  - o maximize the *objective* of expected reductions in total stormwater runoff over a multidecadal time horizon, while
  - o satisfying a *risk constraint* concerning the annual runoff reduction in Stage *II* and a fixed budget shared by both stages.
- Assumptions:
  - o The manager is risk-averse regarding the possibility of failing to meet the State *II* reduction target in annual stormwater reduction.
  - o Investment in a particular type of GI in Stage *I* will result in learning if the investment exceeds some threshold. This learning can be taken advantage of in Stage *II* by
    - investing more in that type of GI, if it turns out to be highly effective, or
    - investing in other types, if the first type is instead found to be relatively ineffective.
  - o Learning about one type of GI provides no information about the performance of others.
  - o GIs are installed in parallel (i.e., no interactions among GI installations) so that the reduction in stormwater can be summed

We next use a decision tree schematic to illustrate the underlying logic of the method. Interested readers can refer to the Appendix for details of the mathematical formulation.

3.1.1. Decision tree example

A decision tree involves decision nodes (squares), chance nodes (circles) and outcomes (at the end of the tree branches). The chronologic sequence of decisions and chance events proceeds from left to right. A decision node represents a point in time when decisions must be made among the alternatives represented as arcs exiting the right side of the node; a chance node stands for random events, each represented by arcs exiting the right whose probabilities sum to 1; and the outcomes show the objective function values realized for a particular sequence of decisions and random events represented by the nodes and arcs leading from the left-most node to the terminal node on the right. Fig. 1 is a schematic in that only some of the nodes and branches are shown for clarity.

In this example, the manager can make investments in either or both Stages *I* and *II*, which can consist of investments in any or all of *n* types of GI. In general, each type can have different levels of efficiency and uncertainty in stormwater reduction per $ investment.

Following the terminology of Bayesian statistics, the distribution of the GI performance in Stage *I* (representing our current understanding about the technology), is called the prior distribution ($\widetilde{C_I}$), and the updated distribution in Stage *II* (representing our new knowledge given the experience learned from the Stage *I* investment) is called the posterior distribution ($\widetilde{C_{IIs}}$). In probability, the posterior distribution is a conditional probability of the outcome given what is learned.

We assumed three types of learning as follows:

- No learning (NL): the posterior distribution is the same as the prior distribution
- Partial learning (PL): the variance of the posterior distribution is less than the prior distribution but not 0, and
- Full learning (FL): the posterior distribution has zero variance (certainty)

PL and FL are achieved for a particular GI type if investment x_I in that type exceeds thresholds, respectively (denoted as Th^{PL} and Th^{FL}, $Th^{PL} < Th^{FL}$). Which of the types of learning that happens depends on which of the thresholds are met. That is, if the investment in Stage *I* (denoted x_I) is:

- less than Th^{PL}, nothing is learned and the posterior distribution is the prior distribution;
- between Th^{PL} and Th^{FL}, then PL takes place; and
- otherwise, FL takes place.

We defer the discussion of the learning and investment relationship to the next subsection and focus our discussion here on the decision problem. To account for uncertainty, m_I scenarios (indexed by *s*) are generated after the Stage *I* decision is made, and m_{II} scenarios (indexed by *r*) are defined to follow each Stage *II* decision. The total number of scenarios is $m_I * m_{II}$. Each Stage *I* scenario contains a realization of stormwater reduction and a learning outcome (the posterior distribution) from the Stage *I* investment, while each State *II* scenario describes a realized reduction for investments in the second stage. To help the reader understand the process, we present a numerical example of a PL branch of the decision tree (Fig. 2).

In this example, the manager has two investment options: GI_1, whose performance is normally distributed (mean and variance equal to 0.5 gallon per $ per year (gal/$/yr) and 0.04 (gal/$/yr)^2, respectively); and GI_2, which provides 0.48 gal/$/yr stormwater reduction with certainty. The total budget for the two-stage planning problem is $100M, and the investment decisions are made at year 1 and 11. We assume that, once the first stage's GI is installed, it will continue providing stormwater reduction until the end of the planning horizon and, therefore, the installations in Stage *I* can provide stormwater reduction for 40 years while, if installed in Stage *II,* a GI reduces stormwater only for 30 years. Costs are not discounted in this example. The manager's objective is to maximize the total stormwater reduction over the 40-year planning horizon. However, the manager may also be concerned with the risk of very low reductions occurring in Stage *II*, as we explain below.

Moving forward in time (from left to right in Fig. 2), first the manager decided to invest $40M in GI_1 and none in GI_2 because she knows that the former investment would result in PL and reduce the performance uncertainty in GI_1 to an assumed value of 0.01 (gal/$/yr)^2. However, she does not know which scenario *s* she would be in prior to entering Stage *II*. Either one of the two Stage *I* scenarios could happen. Scenario $s = 1$ consists of a realization of a performance of 0.7 gal/$/yr reduction for the investment made in GI_1 in Stage *I*, and a posterior distribution with mean and variance equal to 0.675 gal/$/yr and 0.01 (gal/$/yr)^2, respectively. Meanwhile scenario $s = 2$ could consist of a realization of 0.6 gal/$/yr reduction for that GI, and a posterior distribution with mean and variance equal to 0.325 gal/$/yr and 0.01 (gal/$/yr)^2, respectively.

Note that those two realized performance values are consistent with the assumed prior mean and variance (i.e., mean: (0.675 + 0.325)/2 = 0.5; variance: $((0.675–0.5)^2 + (0.325–0.5)^2)/2 + 0.01 = 0.04$; more details of the preservation of the prior mean and variance can be found in the Appendix). Note also that, although we only have two scenarios for each chance node for brevity, a large number of scenarios is recommended to better represent the random distributions. In Stage *II*, the manager then makes another investment decision, which depends on which scenario she is in. If she is in scenario $s = 1$, the Stage *II* investment would be to allocate the remaining part of the budget ($60M) in GI_1 in order to take advantage of its high expected performance in the posterior distribution. Whereas, if she is in scenario $s = 2$, she would rather invest the remaining budget in GI_2 since the expected value of the posterior distribution for GI_1's performance is lower than 0.48 gal/$/yr. The final outcomes at the end of the tree are the total stormwater reduction for the 40-year planning horizon.

We assume equal probability for each scenario (50%) so that we can calculate the stormwater reduction under each combination of first- and second-stage scenarios (*s,r*), as shown in Table 1, and the expected reduction of the first-stage decision (x_I = [$40M, $0]), 1840 MG.

In contrast, if learning is not considered, a risk-neutral manager would invest all her budget in GI_1 in Stage *I* ("*all-in*") and receive an expected value of 2000 MG stormwater reduction (= 0.5*40 yr*0.7 gal/$/yr + 0.5*40 yr*0.3 gal/$/yr). However, if she is risk-averse and would like to make sure the *annual reduction* she received at the end of the planning horizon would be higher than some management goal, she may want to instead consider the investment option in Fig. 2. Although the investment presented in Fig. 2 provides less total stormwater reduction (1840 MG = 0.25*2515 + 0.25*2155 + 0.5*1,344, next to last column of Table 1; this is 160 MG less than the *"all-in"* case), the worst annual reduction (40.8 MG/yr, under $s = 2$) is higher than the worst reduction of the "*all-in*" strategy ($100M*0.3 gal/$ = 30 MG/yr). This example is just one branch of the decision tree in which the PL takes place; there could be other investment strategies that show different tradeoffs between expected reduction and risk in the NL, FL, and other PL branches, which are not shown.

It is worth noting that the mean of particular posteriors for particular scenarios *s* could be higher or lower than the prior's mean; however, the overall mean is preserved when averaged across all scenarios, while the posterior variance is always reduced because of our learning assumption. Moreover, the realizations of the prior may not correspond to the means of the posteriors. For example, the first-stage realization for a particular GI might be a disappointingly low stormwater reduction per $ of investment. Yet as a result of this experience, engineers may also have learned how to improve that GI's design, which improves the mean of the posterior distribution of gal/$ for that GI. As another example, the reduction realized in the first stage could be high, however, the stormwater managers may find that the best locations have been taken, so that remaining sites would cost more for installing GI, which results in a lower mean of the posterior distribution compared to the prior distribution.

Rather than defining a decision tree by discretizing values of Stage *I* and *II* investments *x* in this manner and then solving it by standard backward induction methods, we instead solve a stochastic optimization problem with the same logical structure using mixed integer linear programming, as described in the Appendix. This allows us to treat *x* as

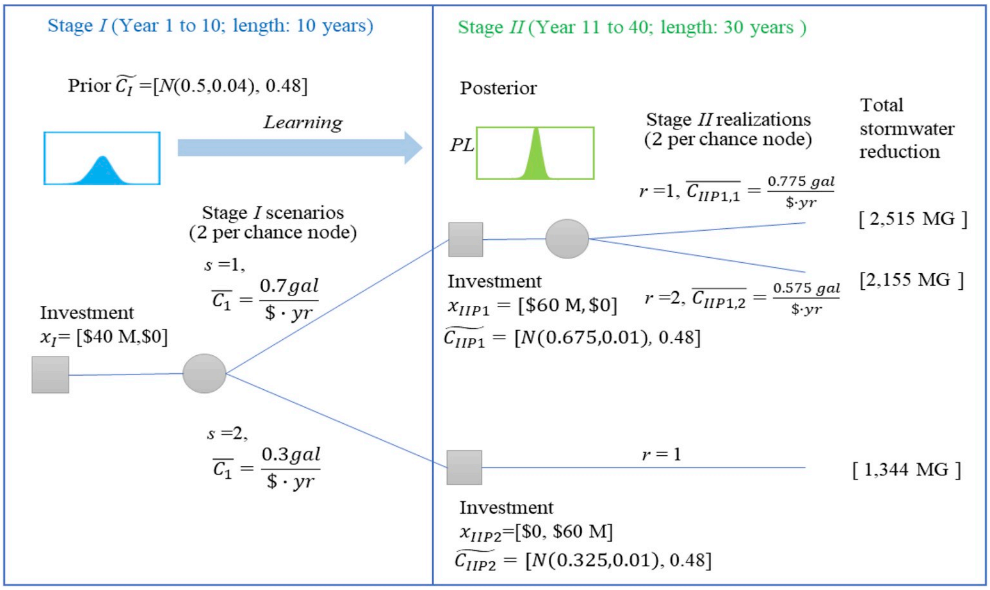

Fig. 2. Numerical example: portion of a GI planning decision tree in which the total budget for the 40-year planning horizon is $100M; the Stage *I* investment results in PL for the uncertain GI_1; and the Stage *II* decision can be either to invest further in GI_1 or instead make all subsequent investments in GI_2.

continuous variables, and risk constraints can also be imposed (which is done only with difficulty in a decision tree). The proposed method can handle a large number of scenarios and GI types and is, in general, easy to solve. Furthermore, our formulation allows the user to define learnings (discussed in the next subsection). By changing the fiscal budget and target CVaR (our risk metric in the proposed method), tradeoffs between expenditures, risk, and expected stormwater reductions can be explored, as we illustrate with the example described in Section 4.

3.1.2. Learning and learning curve functions

Learning is the process of understanding the system of interest. Specifically, in statistics, learning is the process of updating one's beliefs about the system from new observations or new findings, where a belief is expressed as the distribution of a random variable. The traditional Bayesian approach assumes a prior distribution for a random variable based on current understanding, while representing learning by the adjustments in the distribution in the later stages (Kelly and Kolstad, 1999). For example, such adjustments might depend on what is learned and include: *a.* a reduction in the variance of the probability distribution, and/or *b.* a change in the mean of the distribution, as shown in the decision tree example, above.

In practice, the prior distribution (of some uncertain variable $\widetilde{C_I}$) can be updated using Bayes' Law with actual observations after research or monitoring take place. Although the actual observation is not available at the first planning stage, we will usually expect that more GI investment will result in a more significant reduction in uncertainty (of GI performance) in later stages. The underlying assumption is that implementation experience can be transformed into the knowledge of the systems through a process which is described as a learning curve function (Berglund and Söderholm, 2006; Ferioli et al., 2009; Söderholm and Sundqvist, 2007).

We adopt learning curve functions in our model and use them to describe how the Bayesian updating process will be affected by the level of investment. For example, learning curves, denoted *G(x)*, might show how the variance changes using 1-step, 2-step, or linear functions, as shown in Fig. 3.

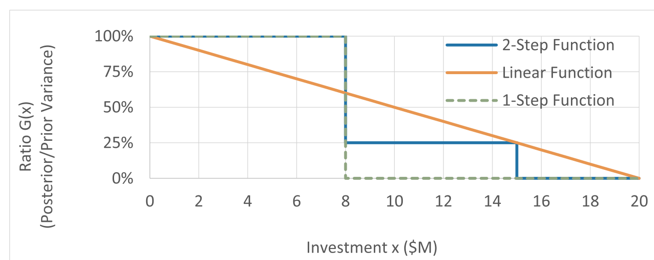

Fig. 3. Examples of learning curves that predict the fraction of reduction in variance as a function of the investment.

In the figure, the 1-step function shows that if the investment is larger than or equal to $8M, learning occurs and the variance is reduced to zero. Otherwise, the variance will remain unchanged. The 2-step function has three states: no learning (NL), partial learning (PL) and full learning (FL) with 100%, 25%, and 0% variance, respectively, corresponding to investment under $8M, between $8M and $15M, and over $15M. On the other hand, the linear function shows that the variance reduction is proportional to any investment less than $20M.

Learning curves can be very general, for instance representing shifts in the mean of the uncertain variable such that the average over the posteriors is better than the prior because learning has improved the technology. The users can choose the form and parameterization of the learning curves to reflect expert judgment, statistical analyses, or both.

Table 1
The total stormwater reduction of each of the three scenarios in the PL branch example (Fig. 2).

Scenario *s,r*	Time Horizon (year)		Realization (gal/$)		Investment ($M)		Total Reduction (MG)	Annual Reduction in Stage II (MG/yr)
	Stage *I*	Stage *II*	Stage *I*	Stage *II*	Stage *I*	Stage *II*	Objective	Risk Metric
1,1	40	30	0.7	0.775	40	60	2515	74.5
1,2	40	30	0.7	0.575	40	60	2155	62.5
2,1	40	30	0.3	0.48	40	60	1344	40.8

3.2. Two-stage stochastic programming with CVaR constraints

This section introduces a general mathematical statement of the proposed method, which can help readers to understand the full formulation in the Appendix.

3.2.1. Basic two-stage stochastic programming (TSP)

A generic formulation of a TSP model is as follows. Risk aversion is ignored for the moment, so the objective (Eq. (1)) is to maximize the expected reductions in runoff subject to a resource/budget constraint and the non-negativity constraints (Eqs. (2) and (3)):

Choose $\{x_I, x_{II}\}$ in order to Maximize

$$f_I(x_I, x_{II}) = E[\widetilde{C_I}]x_I + E_{s\in S}[\widetilde{C_{IIs}}x_{IIs}] \quad (1)$$

Subject to:

$$A(x_I + x_{IIs}) - B \le 0, \quad \forall\, s \in S \quad (2)$$

$$x_I \ge 0, \;\; x_{II} \ge 0 \quad (3)$$

Notation:

- S is the set of future scenarios and $s\in S$ is the index of the scenarios
- x_I and $x_{IIs} \in R^{1\times n}$ are the decision vectors representing the \$ investments in n types of GI in Stage I and in scenario s, Stage II, respectively
- $x_{II} = \{x_{IIs},\;$ for all $s \in S\}$
- $f_I(x_I, x_{II})$ is the overall expected benefit function (e.g., expected reductions in stormwater runoff)
- $\widetilde{C_I} \in R^{n\times 1}$ is a random vector representing the GI's performance coefficients (e.g., gal/\$/yr) in Stage I
- $\widetilde{C_{IIs}} \in R^{n\times 1}$ is a random vectors representing the GI's performance coefficients in scenario s, Stage II
- $A \in R^{m\times n}$ and $B \in R^{m\times 1}$ are constant matrices for linear resource constraints

The probability distribution of $\widetilde{C_I}$ is the prior distribution while $\widetilde{C_{IIs}}$ is the posterior distribution, given that scenario s has occurred. Updating the prior of $\widetilde{C_I}$ to the posterior of $\widetilde{C_{IIs}}$, given s, is the process of "learning." Matrices A and B are deterministic in the problem formulations above, but they could also be stochastic.

A special case of the TSP model is one with a deterministic second-stage problem. However, in the simplest case in which the decision maker is risk-neutral, it makes no difference whether the second-stage problem has residual uncertainty in the objective function or not, since the objective is to maximize expected performance − we can just use the expected value of $\widetilde{C_{II}}_s$ in the second stage objective. This is however not true if the decision maker is risk-averse, in which case, a formulation that includes the decision maker's risk attitude is needed. This generalization is the subject of the next subsection.

3.2.2. Risk-averse optimization

A decision maker's risk-averse preferences can be modeled with constraints on some risk measures (see Artzner et al., 1999). Let $\widetilde{f(x)}$ be the uncertain value of the objective function (assuming that $\widetilde{f(x)}$ is to be maximized), given decisions x. Conditional Value at Risk (CVaR) is the risk measure used in this paper, and is a popular measure because of certain mathematical properties it possesses. To explain CVaR, we first introduce another popular risk metric, Value at Risk (VaR). VaR is the value of the α-quantile of the uncertain objective $\widetilde{f(x)}$, given that the decisions are x. α is a specified number ranging from 0 to 1 indicating the probability of the least acceptable outcome to the decision maker. CVaR is the conditional mean of the outcomes worse than $VaR_\alpha(x)$, given α and the decisions x. Therefore, $CVaR_\alpha \le VaR_\alpha$. The mathematical expressions of the two risk measures are as follows:

- VaR:

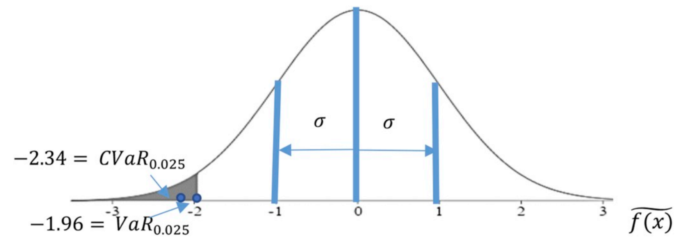

Fig. 4. Standard deviation (σ), VaR_α, and $CVaR_\alpha$ on Normal (0,1) (shaded area α≤ 0.025) (Note: $\widetilde{f(x)}$ is the uncertain value of the objective function, not the probability density.).

$$VaR_\alpha(x) = Min_y\{Prob(\widetilde{f(x)} \le y) \ge \alpha\} \quad (4)$$

- CVaR:

$$CVaR_\alpha(x) = \frac{1}{\alpha}\int_0^\alpha VaR_y(x)dy \quad (5)$$

Fig. 4 illustrates the meaning of standard deviation (σ), VaR, and CVaR for a standard normal distribution (mean = 0, variance = 1) and α= 0.025.

The VaR and CVaR measures are indices of "bad outcomes", and so more is desired if $\widetilde{f(x)}$ is to be maximized (as is the case here with stormwater reductions). Thus, to represent a limit on risk, they would be constrained by imposing lower bounds. We chose to constrain CVaR, because it has two appealing properties: (a) ease of computation (a linear program can be used fx1 when there are discrete scenarios) (Rockafellar and Uryasev, 2000); and (b) CVaR reflects the distribution of outcomes below VaR, penalizing distributions whose tails encompass more extreme values.

By adding CVaR constraints, the optimization model can represent risk-averse preferences. If the constraint is binding, the result is likely to be an investment that increases the expected performance in the worst-case scenarios (those with a cumulative probability less than α). However, this will likely be at the expense of a deterioration in the expected value of the objective (1). That is, there will be tradeoffs between risk and expected performance.

Adding CVaR constraints in TSP directly makes it a non-linear problem because of the nonlinearity of the probability distribution function. However, Rockafellar and Uryasev (2000) proposed a linear programming formulation with Monte Carlo sampling that approximates the CVaR problem. We adopt a variant of that model (Krokhmal et al., 2001) that maximizes the expectation of $f(x)$ subject to CVaR constraints (6):

$$\begin{cases} z_k \ge \tau - f_k(x), \;\; for\ all\ k \in \{1,2,\ldots,m\}. \\ \quad \tau - \frac{1}{1-\alpha}\sum_{s\in S} P_k z_k \ge CVaR_\alpha, \end{cases} \quad (6)$$

where k is the index of m realizations of the random variables; τ is an auxiliary variable for VaR_α calculation; z_k is an auxiliary variable for the calculation of the loss exceeding VaR_α under realization k, which has probability P_k; $f_k(x)$ is the function of some metric that we would like to constrain the risk on under realization k; and $CVaR_\alpha$ is a specified number representing the least acceptable outcome (e.g., the smallest runoff reduction that is acceptable). Our implementation is somewhat more complex to account for the two-stage nature of the problem in which overall performance depends on realizations in both Stage I (earlier years) and Stage II (later years). More details are provided in the Appendix.

4. A hypothetical example

4.1. Overview of assumptions

The modeling framework summarized above is general in that it gives us the flexibility to build versions based on a range of assumptions about learning to assess whether considering learning could significantly change the optimal Stage *I* investment decisions. In the illustrative example of this section, we implement two versions: the Multi-level Learning (ML) model and the Technology Improvement (TI) model, the details for both of which are presented in the Appendix. The ML model assumes that the first-stage investment for a given GI type could yield either NL, PL or FL (as in the two-step learning function shown in Fig. 3). Meanwhile, the TI model assumes that learning not only reduces the variance but also increases the expected value of the posterior distributions of stormwater reductions. The TI model can be viewed as having two learning curves for each GI type showing the impact of learning upon the mean and variance, respectively, of the posterior distributions, whereas the ML model only has a single learning curve function corresponding to just the posterior variance.

In this two-stage planning example, we make the following assumptions.

1. The stormwater manager can invest in three GI types, called GI_1, GI_2, and GI_3;
2. The types have random first-stage performance (stormwater reduction in gal/\$/yr) levels, $\widetilde{C_I^1}$, $\widetilde{C_I^2}$, and $\widetilde{C_I^3}$, respectively;
3. A budget of \$100M is available, which can be optimally split between Stages *I* and *II*. The planning horizon is 40 years: Stage *I* begins at year 1 and Stage *II* will begin at year 11;
4. The decision variables are how to distribute the \$100M budget between the stages, and among the GI types. The distribution of the budget allocated to Stage *II* among GI types depends on the scenario *s*;
5. The objective function is to maximize the total probability-weighted annualized runoff reduction (in million gallons/year, MG/yr) resulting from the investments, with the reductions summed over a 40-year planning horizon. Once a GI type is installed, it will continue to provide stormwater benefits until the end of the planning horizon (year *40*);
6. Constraints on investment include the overall budget constraint, nonnegativity for all investments in each stage, and definitions of CVaR and learning relationships (defined above; also see the Appendix). A prespecified lower bound is placed on $\text{CVaR}_{\alpha=0.1}$. That is, the manager specifies a minimum acceptable stormwater reduction corresponding to the conditional expectation of the worst 10% of the outcomes. In the examples, we vary $\text{CVaR}_{\alpha=0.1}$ (MG/yr) from 0 to 60 to represent increasing degrees of risk aversion. The relationships between first stage investments and learning are summarized in Sections 4.2 and 4.3.

The final set of assumptions concern the three GI types. They represent three distinct technologies: GI_1 is a low-cost and well-understood technology (a low variance for $\widetilde{C_I^1}$) with the lowest learning thresholds for both PL and FL; GI_2 is more cost-efficient than GI_1 (in terms of expected gal/\$/yr performance) but the uncertainty associated with its performance and its learning thresholds are also higher; and GI_3 is a new technology with little empirical data for its cost-effectiveness but is believed to have the highest expected cost-efficiency as well as the most uncertainty about its gal/\$/yr performance, and also the highest learning thresholds for both PL and FL. Table 2 shows the assumptions made about their first-stage performance in this example.

The learning curves (variance reduction for each GI type) used in the ML model are shown in Fig. 5. There, $G_i(x) = 1, i \in \{1,2,3\}$ signifies that the investment is not enough to trigger learning (NL); $G_i(x) = 0.25$ means the investment results in partial learning (PL) and the variance is reduced to one-fourth of the original value; and $G_i(x) = 0$ means that perfect information is provided in Stage *II* (FL), so that variance is reduced to zero. The PL (75% reduction) and FL (100% reduction) thresholds are \$5M and \$10M respectively for GI_1; \$8M and \$15M respectively for GI_2, and \$15M and \$30M respectively for GI_3.

Table 2
GI performance assumptions for the hypothetical example (All distributions assumed to be independent).

gal/\$/yr	$\widetilde{C_I^1} \sim$ Normal (μ_I^1=0.4, σ_I^1=0.05)
	$\widetilde{C_I^2} \sim$ Normal (μ_I^2=0.5, σ_I^2=0.20)
	$\widetilde{C_I^3} \sim$ Normal (μ_I^3=0.6, σ_I^3=0.25)

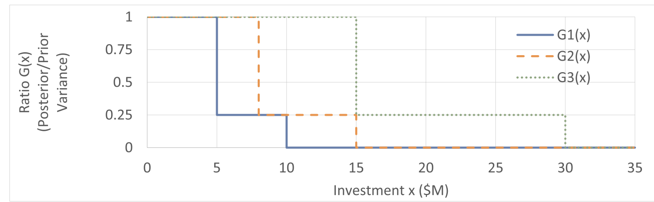

Fig. 5. Variance learning curve functions used in the ML and the TI models.

In the rest of Section 4, we show the Stage *I* decisions resulting from the ML and TI models (Sections 4.2 and 4.3, respectively) for a range of CVaR values, which represent different levels of risk aversion. Also, we show the tradeoff between expected stormwater reduction and risk, indicating that a substantial improvement in risk can be purchased with a slight deterioration in expected reduction. The specific learning assumptions used in the ML and TI models are described in their respective sections.

4.2. Results from the ML model (assumes no technological improvement)

To represent uncertainty in GI performance, we generated 1000 Stage *I* scenarios *s*, as well as 5 Stage *II* scenarios *r* for each Stage *I* scenario. We generated more State *I* scenarios because computational experience shows that a larger sample size in Stage *I* can help solutions to converge, but that increasing sample sizes in the later stage has less of an effect (Ji et al., 2005). The issue of solution convergence is discussed further in Section 5.1.

The ML model is solved for values of CVaR ranging from 16 MG/yr to 42 MG/yr under the \$100M budget, and the resulting first stage investments in the three GI types are shown in Fig. 6 (left axis). The figure also shows the probability-weighted average stormwater reduction, averaged over the 40-year time horizon (right axis). At a CVaR of 16 MG/yr, the CVaR constraint no longer binds and the risk-neutral solution results, i.e., this is the solution that maximizes the expected annual stormwater reduction, achieving 60 MG/yr. Higher values of the

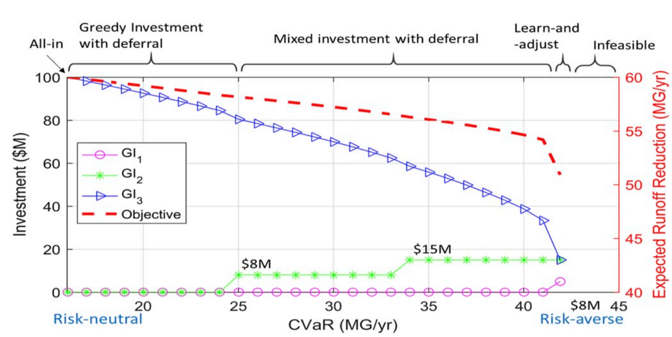

Fig. 6. Stage *I* investment solutions (left axis) and the objective function values (right axis) from ML model for $\textbf{CVaR}_{0.1}$ in the range [16,42].

CVaR lower the risk of “poor” stormwater results, but at the expense of lower expected performance. CVaR can be feasibly increased as far as 42 MG/yr, which represents a 163% increase; at that level, expected stormwater reduction, however, falls about 13%, from 60 to 52 MG/yr. Whether that deterioration in expected performance is justified by the reduced risk of abysmal performance is a value judgment that the decision makers must make, informed by the results of the model.

Depending on the degree of risk aversion, the optimal strategy falls into one of four groups. The results suggest that a risk-neutral decision maker should invest all \$100M in GI_3 in Stage *I*, because GI_3 has the largest expected benefit, which is shown as the “All-in” solutions (when CVaR ≤ 16 MG/yr). There appear to be insufficient incentives to invest in other technologies or to wait to obtain better estimates of GI performance. However, if the manager is risk-averse, she may prefer to save some budget for future investment or mix her investment with more than one GI type or to take advantage of learning. For example, when 16< CVaR ≤ 24 MG/yr, the model suggests investments in GI_3 but also to save some budget for the later stage in case GI_3turns out to be worse than expected (called “greedy investment with deferral” strategy). When 24< CVaR ≤ 41 MG/yr, the model suggests investments in both $GI_2$and GI_3 while saving some budget to take advantage of learning from FL or PL (called “mixed investment with deferral” strategy) in GI_2 and GI_3. Finally, when 41< CVaR ≤ 42 MG/yr, the model suggests investing only for learning (PL for GI_1 and GI_3; FL for GI_2) and saving the rest of the budget to invest in the second stage in order to minimize the risk (i.e., maximize CVaR) (which we call the “learn-and-adjust” strategy). Not surprisingly, the most investment strategies suggest investing early because “to wait” means that there is no reduction in the first 10 years, which lowers the total runoff reduction over the 40-year time horizon.

4.3. Results from the TI model (assumes technology improvement)

The TI model adds another set of learning curve functions (Fig. 7, denoted $H_i(x)$ for $i \in \{1,2,3\}$) to the above ML model such that, if the learning criteria are met, the expected values of GI performance $\widetilde{C_{II}^1}$, $\widetilde{C_{II}^2}$ and $\widetilde{C_{II}^3}$ will improve by 50%, 20% and 0%, respectively. As a result, compared to the ML model of Section 4.2, the model has a higher incentive to make “here-and-now” (Stage *I*) investments for reducing uncertainty and increasing the expected values (Fig. 8).

Moreover, under the technology improvement assumption, the CVaR value can be as high as 51 MG/yr and still yield a feasible solution because of the increase in performance resulting from Stage *I* investment, while, in the ML model, the CVaR cannot exceed 42 MG/yr. Finally, when CVaR is set to 39 MG/yr or higher, the TI model also suggests an investment of \$5M in GI_1 to obtain PL, whereas this technology is not invested in by the ML model except the most risk-averse case. This is because GI_1, being more uncertain, has a higher potential to increase its performance so that it may become a preferable option in Stage *II*, if its uncertain second stage performance turns out to be relatively high. But if we do not invest in GI_1 in Stage *I*, we may be disappointed by the achieved runoff reduction in Stage *II* if GI_2 and GI_3 both turn out to be less efficient than expected.

Fig. 8 shows the model's objective (expected stormwater reduction

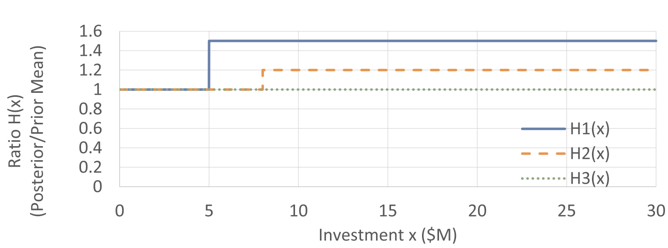

Fig. 7. Learning curve functions for expected value improvement in the TI model.

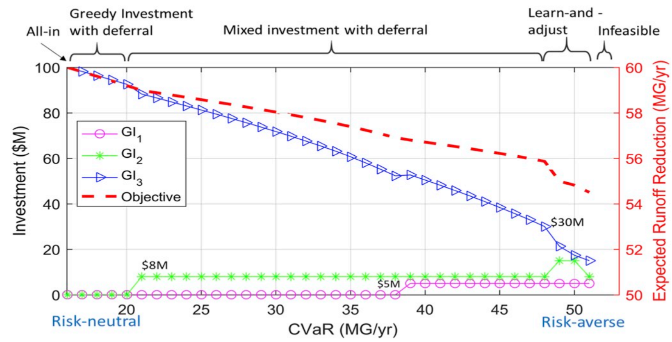

Fig. 8. Stage *I* investments and objective function values from the TI model, $\mathbf{CVaR}_{0.1}$ in the range [16, 51].

over 40 years) as a function of CVaR. It reveals that, without considering risk aversion, the objective value is 60 MG/yr (the same as in Fig. 6), which like the ML model is an outcome of the “all-in” strategy. It turns out that under risk neutrality, the expected gains in technology effectiveness in Stage *II* are not justified by the loss of efficiency resulting from the Stage *I* investments necessary to achieve those gains. However, in the presence of risk-aversion, CVaR can be increased to 51 MG/yr, if the decision maker is willing to accept a loss of 5.6 MG/yr (from 60 to 54.4) in expected value. This occurs from investing the minimal amount needed in Stage *I* to achieve the technology improvements and posterior variance reductions that are possible in all three technologies. Also, when applying a “mixed investment plus wait-and-see” strategy (CVaR in the range of [21,47]), the objective curve is higher than the results in Fig. 5, which is a result of technology improvements. The technology improvement assumption can be viewed as indicating that a drop in cost or an increase in efficiency for a particular GI is forecasted as possible in the near future. Then the TI model can help answer the question of how Stage *I* investments can help expedite these improvements by providing several distinct solutions (mixes of GIs) for the decision makers to ponder, with some solutions emphasizing reducing uncertainty and others focused on improving expected stormwater reductions.

5. Discussion

Section 4 has shown that, in this hypothetical example, the ML and TI models both suggest four distinct investment strategies, depending on the learning assumptions, decision maker's risk attitudes, and the particular parameter values assumed for the model. In a more realistic case study, the choices of the model input assumptions need to reflect the conditions and the opportunities in the study area; these will include the probability distributions and learning, costs, resource, and other physical and regulatory constraints. However, the user should keep in mind that, as with the results of any model, the precise numerical results should be viewed skeptically and that the most valuable outputs are insights as to which investment alternatives appear most economic and why. In this paper, we focus on the definition and illustration of the method; future work will apply the method to a realistic case study.

In this section, we consider two crucial implementation issues. The first issue is solution instability due to scenario sampling errors, in which the optimal investments may be a function of the particular sample that is taken. Section 5.1 presents a solution stability analysis to examine the variability of the solutions and the resulting risk levels. The second issue, which is discussed in Section 5.2, is the impact of alternative learning assumptions since they are based on expert judgment and are likely to be imprecise and significantly impact the solution.

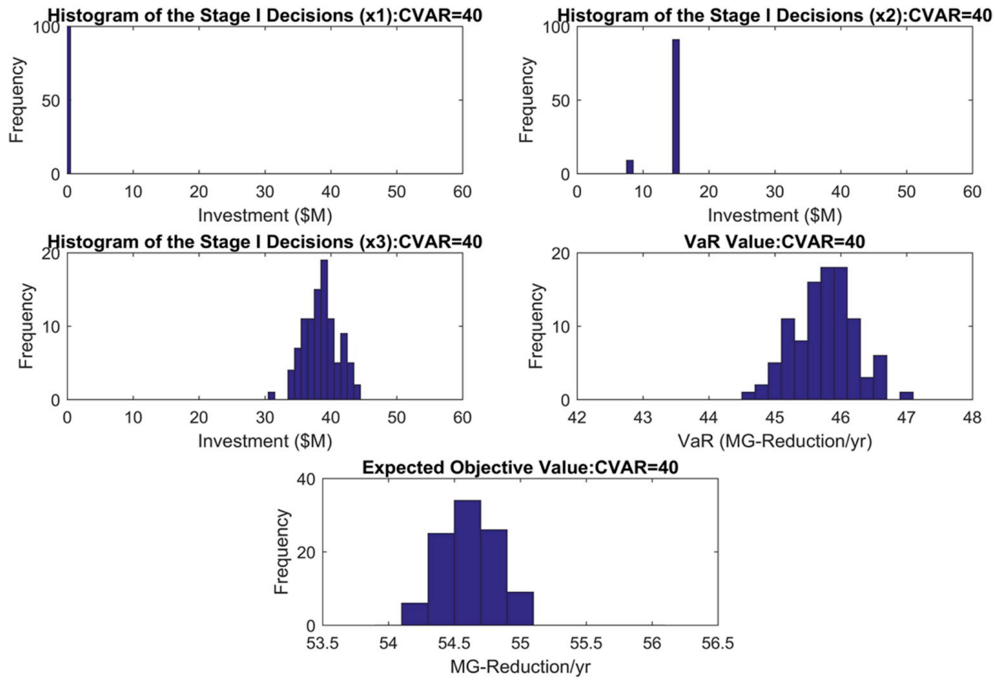

Fig. 9. Results of 100 Monte Carlo simulation runs of the ML model for solution stability analysis.

5.1. Are solutions robust to sample error?

The technique we used to solve the ML and TI models in Section 4.1 and 4.2 relies on generating scenarios to approximate the random distributions, which are subject to sample error. A large sample size can reduce this error but will increase computational costs. Our numerical example presented in Section 4.1 has 14,013 decision variables and 18,013 constraints, respectively, which takes about 10–30 min to solve on a conventional laptop (Intel Core i7, 4 GB RAM). Therefore, it is not practical to include much larger samples. A solution stability analysis can provide information about solution variability due to sampling errors, which also provides alternative solutions for the decision maker to consider.

To evaluate whether the sample size is adequate and to assess the resulting variability of the solutions, we perform a solution stability analysis by running a Monte Carlo simulation of the ML model (Section 4.1) 100 times, with each of the 100 solutions obtained setting $CVaR_{0.1}$ to 40 MG/yr. Recall that the number of scenarios is 1000 and 5 for Stage *I* and *II*, respectively.

Fig. 9 shows the results from the Monte Carlo simulation runs. The histograms of the first stage solutions indicate that it is optimal to invest $x_1 =$ \$0 in GI_1 in all solutions, and $x_2 =$ \$8M (PL) in GI_2 for 9% of the solutions and $x_2 =$ \$15M (PL) for 91% of the solutions. For GI_3, it suggests a range of $x_3 =$ \$31Mto \$43M (FL) with an average investment of \$38.5M. Meanwhile, the $CVaR_{0.1}$ value is set to 40 but the 90% exceedance value $VaR_{0.1}$can vary from 45 to 47 MG/yr and the expected annualized runoff reduction can vary from 54.2 to 55 MG/yr. Thus, the pattern of investment is somewhat stable in the face of sample error, but the precise level of risk and expected performance of the system can vary. Larger sample sizes would result in smaller variations. The stormwater manager can choose an investment decision and run the Monte Carlo simulation to evaluate the variability of the objective and CVaR values. For example, Fig. 10 shows the histograms of the objective and CVaR values of the initial GI investment plan of {\$0M, \$15M, \$30M} from 100 simulation runs, which are in the ranges of (53.7, 54.1) and (40.3,43), respectively. This analysis provides additional information regarding the stability of the resulting expected outcomes and risk levels.

5.2. How do first-stage recommendations depend on learning opportunities and risk aversion?

Decision makers can adjust learning assumptions to explore whether the possible benefits of learning can justify significantly different portfolios of the first stage investments. A trivial example is that a risk-neutral decision maker would apply the same “all-in” strategy in the previous ML and TI examples (Figs. 6 and 8), even though the latter has an additional technology improvement assumption. On the contrary, a risk-averse decision maker (with the CVaR set to 36 MG/yr), would likely adopt a “mixed investment with deferral” strategy, but the investment portfolio in the ML example (Fig. 6) consists of only GI_2 and GI_3 whereas the investment portfolio in the TI example (Fig. 8) also includes an investment in GI_1.

Besides, our method can be used to evaluate active learning opportunities by applying a weight of zero to the first stage reduction in the objective, so that the objective of the planning problem only focuses on the second stage reduction and considered the first stage decision as representing possibilities for active experiments. As an example, we modify the TI model, changing the technology improvement assumption to increase the effectiveness of GI_1, GI_2 and GI_3 by 60%, 20%, and 0%, respectively. Then the model would suggest a Stage *I* investment portfolio of \$5M, \$8M, and \$15M for GI_1, GI_2 and GI_3, respectively, for a risk-neutral decision maker, but in the risk-averse case would recommend decreasing the investment in GI_3 because of its high uncertainty (Case 1 in Fig. 11). If the learning assumption is changed to increase the effectiveness of GI_1, GI_2 and GI_3 by 60%, 30%, and 10%, respectively, then the model would suggest investing \$8M and \$15M in GI_2 and GI_3, respectively, in Stage *I* for a risk-neutral decision maker, but would recommend adding an investment of \$5M in GI_1 for a risk-averse decision maker (Case 2 in Fig. 11). GI_1 is not included in the risk-neutral case because of its low efficacy without learning but is added in the risk-averse case because of its low uncertainty and improved efficacy in the second stage. Thus, whether risk aversion results in adding or subtracting types of investments from the first stage decisions depends on the magnitude of technology improvement in the learning assumptions.

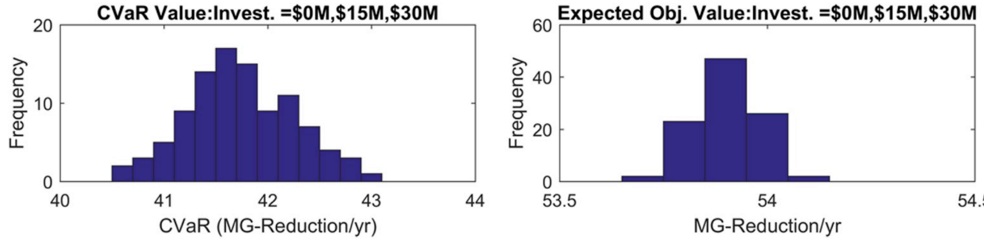

Fig. 10. Results of 100 Monte Carlo simulation runs of the ML model for the first-stage investment of {$0M, $15M, $30M}.

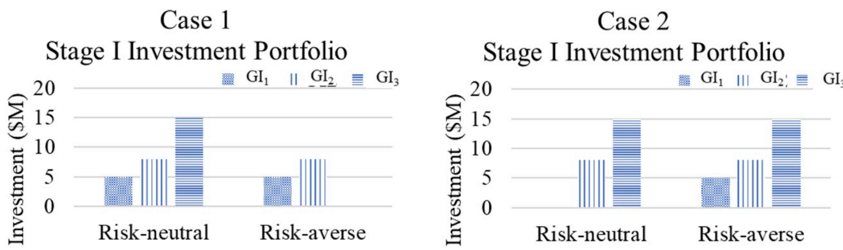

Fig. 11. Stage I investment portfolios for two alternative learning assumption cases.

6. Conclusion and final remarks

We have presented a new class of methods for adaptive stormwater management incorporating decision maker's risk attitudes and learning followed by a hypothetical numerical example. Although the stochastic programming-based method is developed for stormwater management, it can also be applied to other adaptive management problems involving uncertainty and learning, e.g., natural resources management and climate change.

We have described the general model set-up above, with detailed formulations in the Appendix. Two different models are presented, which reflect two distinct assumptions about the effect of learning: the multi-level learning (ML) model which assumes learning can only reduce variance, and technology improvement (TI) model which assumes learning would also improve the overall performance of GI. Then the models are tested in the numerical example where the numerical results show that four basic types of investment strategies could be optimal (Figs. 6 and 8): *all-in* (invest all the budget in the first planning stage in one technology); *greedy investment with deferral* (invest most budget in one technology in the first stage and save a small amount of budget for the later uses); *mixed investment with deferral* (having a mix of investments in the first stage to achieve partial or full learning (PL and FL) and saving a portion of the budget for the later stage to take advantage of the learning outcomes); and *learn-and-adjust* (early investment mainly for learning PL and/or FL to defer the major decisions to the second stage when the learning outcomes are revealed).

In the results from the TI model (Fig. 8), we see that possible technology improvement can provide additional incentive to invest in learning early (Stage *I*) and invest in technologies that have potential to improve their effectiveness. However, the optimal investment strategy would depend on the decision maker's risk attitudes, the learning assumptions, and the particular parameter values. Finally, we show that the solutions of the first-stage investments and the resulting CVaR estimates can be sensitive to sampling error in choosing scenarios. Solution stability analyses (5.1) can provide the decision makers more information concerning the risks (VaR and CVaR) associated with the recommended GI portfolios.

Future work includes a more realistic case study where in which the learning assumptions and the prior distributions of the GI performance are derived from discussions with the experts and stakeholders. Another future direction could be to improve the computation efficiency by applying heuristic search or scenario s reduction methods.

Acknowledgment

This research was funded by United States Environmental Protection Agency (USEPA) grant number 83555501. Its contents are solely the responsibility of the grantee and do not necessarily represent the official views of the USEPA. Further, USEPA does not endorse the purchase of any commercial products or services mentioned in the publication. We thank C. Harman, A. McGarity, and D. Sheer for their guidance.

Appendix A. Model Formulations

We now present the technical details of the model. Following Rockafellar and Uryasev (2000) CVaR modeling method, we discretize continuous random distributions by taking a sample of realizations and assigning a probability to each realization. Unfortunately, to represent many possible degrees of learning under a continuous learning curve function, an impractically large number of samples is required. Therefore, to avoid the resulting computational problems, we use step functions to approximate the learning curves so that we only need to generate a reasonably number of realizations of future scenarios, one set of realizations per step. For clarity, our model denotes random variables with a tilde (~).

A.1. Overall model formulation

The proposed method allows us to build models under various assumptions about learning to assess whether considering learning could significantly change optimal “here-and-now” decisions. However, the exact formulation depends on the learning assumptions. As examples, Sections A.2 shows the formulation of a Multi-level Learning (ML) model which assumes that learning can only reduce the uncertainty of GI performance, whereas the Technology Improvement (TI) model in Section A.3 assumes that learning can reduce uncertainty and, meanwhile, stimulate improvement in the expected performance. The TI model can be viewed as an extension of the ML model since it has the assumption that learning would improve overall GI performance in addition to the uncertainty reduction. Both models are applied to solve the hypothetical problem in Section 4.

A.1.1. Prior and posterior distributions

Following the terminology of Bayesian statistics, the distribution of the GI performance in Stage *I* (representing our current understanding about the technology) is called the prior distribution, and the updated distribution in Stage *II* (representing our new knowledge given the experience learned from the Stage *I* investment) is called the posterior distribution. In probability theory, the posterior distribution is defined as the conditional probability given what is learned. The proposed method assumed that distribution parameters (e.g., mean and variance) of the posterior are functions of the Stage *I* investment, which can be derived from the data and experts' judgments. The assumptions come from the idea that more experience with GI would reduce the uncertainty of its performance (we are interested in the overall performance of many GI practices but not the performance of individual GIs) and could result in better designs and installations and therefore more reduction of stormwater. These functions are named learning curve functions.

A.1.2. Learning curve functions

The prior distributions, denoted $\widetilde{C_I}$, are assumed normal distributions with a mean μ and a variance σ^2 (μ and σ^2 are vectors whose elements, μ_i and σ_i^2, $i \in \{1,2,3\}$, are the mean and variance of the three GI practices). The variance is the measure of uncertainty applied in this example. The learning curve function for uncertainty (variance) reduction is in the form of a two-step function (see Fig. 3) indicating which of the three learning cases (no learning (NL), partial learning (PL) and full learning (FL)) would happen, given an investment. The learning curve function for uncertainty reduction of a GI type (denoted $G(x_I)$, x_I being the Stage *I* investment in that GI) is defined as follows for both ML and TI models.

$$G(x_I) = \begin{cases} \sigma^2, \quad if\ Th^{PL} > x_I & (NL\ takes\ place) \\ \beta\sigma^2,\ \beta \in (0,1),\ if\ Th^{PL} \le x_I < Th^{FL} & (PL\ takes\ place) \\ 0, \quad if\ Th^{FL} \le x_I & (FL\ takes\ place) \end{cases} \tag{A.1}$$

where β is a user specified number to represent the variance reduction under PL, and Th^{PL} and Th^{FL} are the thresholds for investments needed for the PL and FL cases, respectively. For example, the two-step function in Fig. 3 has $\beta = 0.25$, $Th^{PL} = 8$, and $Th^{FL} = 15$. The function says that if the investment is higher than the learning threshold of FL (Th^{FL}), the variance of the posterior distribution is reduced to 0 (certainty). If the investment is between the investment thresholds of PL and FL, then the variance is reduced to $\beta\sigma^2$. Otherwise, the variance of the posterior distribution remains the same as the prior.

Meanwhile, the learning curve function for mean improvement used in the TI model, denoted $H(x_I)$, is assumed to have only one level in the hypothetical example (Section 4) with a threshold equal to the threshold for PL, Th^{PL}, as shown as below.

$$H(x_I) = \begin{cases} \gamma\mu,\ \gamma > 1,\ if\ Th^{PL} \le x_I \\ \mu, \quad if\ x_I < Th^{PL} \end{cases} \tag{A.2}$$

where γ is scaling constant that adjusts the posterior mean. The function says that if an investment exceeds the threshold for PL (Th^{PL}), the average performance would increase to $\gamma\mu$. In a more general case, $G(x_I)$ could have more than one PL option representing different levels of variance reduction, and $H(x_I)$ could have multiple levels as well.

Table A.1
Priors and posteriors of ML and TI models in NL, FL and PL cases assuming that priors and posteriors are normal distributions

Learnings	Prior/NL (Posterior)	FL (Posterior)	PL (Posterior)
Multi-level Learning (ML) Model	$\widetilde{C_I}(\mu, \sigma^2)$	C_{IIFs} (certainty)	$\widetilde{C_{IIPs}}(C_{IPs}, \beta\sigma^2)$
Technology Improvement (TI) Model	$\widetilde{C_I}(\mu, \sigma^2)$	$C_{IIFs,\ TI}$ (certainty)	$\widetilde{C_{IIPs,TI}}(C_{IPs,TI}, \beta\sigma^2)$

Table A1 illustrates how the posteriors change with the learning assumptions. $\widetilde{C_{IIPs}}$ and $\widetilde{C_{IIPs,TI}}$ denote the posterior (Stage *II*) distributions learned in the PL case for the ML and the TI models, respectively. C_{IIFs} and $C_{IIFs,\ TI}$ denote the Stage *II* realized performance value learned in the FL case for the ML and the TI models, respectively, and C_{IPs} and $C_{IPs,TI}$ denote the expected values of the posterior distributions in the PL case for the ML and the TI models, respectively. The values of C_{IIFs}, $C_{IIFs,\ TI}$, C_{IPs}, and $C_{IPs,TI}$ are random samples generated consistent with the assumptions concerning the prior and posterior distributions.

The notation for the realized values and distributions are scalars for one GI and are vectors for multiple GI types, as in the formulations of the ML and TI models. More details of the sampling processes are presented with the model formulations in Section A.2 and A.3.

A.2. ML model

The ML model maximizes the probability-weighted annualized stormwater reductions over the entire time horizon (million gallon/year, MG/yr), subject to resource constraints (e.g., budget), logical constraints ("learning" and "auxiliary" constraints) that relate the amount of learning (NL, PL, or FL) for each GI type to the amount of the first stage investment x_I in that GI, and the CVaR constraints that place a lower bound upon the annual stormwater reduction in the second stage resulting from all GI investments that have been made, representing the risk of failing to meet the management goal. Note that the mathematical notation of the model formulations is in vector form instead of scalars.

A.2.1. Scenarios and random sample sets

Instead of directly working with the probability density functions and solving a non-linear programming problem, the proposed method uses scenarios to approximate the random distributions so that the problem could be solved by linear programming. The random variables in the hypothetical example are the GI performance in stormwater reduction per $M investment per year, and the scenarios are the possible outcomes of the random GI performance in the learning cases in Stage *I* and *II*. A scenario consists of a set of values representing a possible outcome, which are the realized GI performance in Stage *I* and *II*, for each GI technology and each of the learnings (i.e., NL, PL and FL cases). The performances of the GI practices are assumed to be statistically independent. That is, the performance of one GI gives us no information about the performance of the other GI. The sampling procedure is described as follows:

1. Stage *I*: m_I samples are drawn from $\widetilde{C_I}$ (the prior distribution vector of the GI performance, in which the elements are the prior distributions of each GI), denoted $\overline{C_{Is}}$, s∈ $S_I = \{1,2,...,m_I\}$, for the realizations of performance for each GI resulting from Stage *I*'s investment.
2. Stage *II*:
   - The NL case: The sampling process for each scenario s is the same as described in Stage *I* (since the posterior and prior distributions are the same), and the samples are denoted by $C_{IINs,r}$, $s \in S_I = \{1,2,...,m_I\}$, $r \in S_{II} = \{1,2,...,m_{II}\}$.
   - The FL case: The assumption of no uncertainty in Stage *II* means one realization per scenario (denoted C_{IIFs}), unlike the NL case where m_{II} realizations are needed to represent the posterior uncertainty.
   - The PL case: The generation of posterior realizations is more complex in PL case. The posterior distributions in the PL case, denoted $\widetilde{C_{IIPs}}$, is assumed to have a reduced variance ($\beta\sigma^2$, $\beta \in (0,1)$) according to the learning curve, $G(x_I)$ and a mean of μ (the prior mean). From the Law of Total Variance, the prior variance must equal the preposterior variance: $Var(\widetilde{C_I}) = E[Var(\widetilde{C_{IIPs}})] + Var(E[\widetilde{C_{IIPs}}])$. From this we can derive that $Var(E[\widetilde{C_{IIPs}}]) = (1-\beta)\sigma^2$, $\beta \in (0,1)$. Similarly, from the Law of Total Expectation, we can derive $E[\widetilde{C_I}] = E[E[\widetilde{C_{IIPs}}]] = \mu$, as it should be if there is no technology improvement considered. Therefore, the posterior mean $E[\widetilde{C_{IIPs}}]$ of Stage *II* realizations $C_{IIPs,r}$ for the PL case given scenario s is itself a random variable with mean equal to μ and variance equal to $(1-\beta)\sigma^2$, denoted $\widetilde{C_{IP}}$. This allows us to generate realizations of $\widetilde{C_{IIPs}}$ for each scenario s by first drawing a posterior mean, denoted C_{IPs}, for that scenario s from $\widetilde{C_{IP}}$. Then a sample set of realizations of individual $\widetilde{C_{IIPs,r}}$, $r \in S_{II}$ can be drawn from a distribution with mean C_{IPs} and variance $\beta\sigma^2$, denoted $C_{IIPs,r}$, $r \in S_{II}$. This process is repeated for each scenario s.

It is worth noting that the above process implies that mean performance of the posterior $\widetilde{C_{IP}}$is statistically independent of the observed Stage *I* GI performance, $\overline{C_{Is}}$, in that scenario. This represents the situation in which experience with implementation results in learning not only in the form of observations of the performance of Stage *I* investments, but also concerning the hydrological characteristics and installation and maintenance costs for the remaining candidate sites. For example, the managers and engineers may develop new GI siting strategies and designs based on the new knowledge of the watershed which improves the performance of GIs, or they may realize that the remaining GI candidate sites have poor soil properties which would lower the overall GI performance. Alternatively, it is possible to assume nonzero or even perfect correlation of Stage *II* performance with what is observed in Stage *I,* in which case the main information obtained in Stage *I* is the observed performance of that stage's investments.

A.2.2. Mathematical formulation

The ML model formulation, which is a specific implementation of the general model in Section 4.2, is as follows.Decision variables. x_I*:* the Stage *I* investment vector whose elements $x_{I,i}$, $i \in N = \{1,2,...,n\}$ represent the investment in n GI types, $x_I \in R^{1\times n} x_{IINs}$: the Stage *II* investment decision vector for the NL case in scenario s whose elements $x_{IINs,i}$ represent the investment in n GI types, $s \in S_I$, $x_{IINs} \in R^{1\times n} x_{IIPs}$: the Stage *II* investment decision vector for the PL case in scenario s whose elements $x_{IIPs,i}$ represent the investment in n GI types, $s \in S_I$, $x_{IIPs} \in R^{1\times n} x_{IIFs}$: the Stage *II* investment decision vector for the FL case in scenario s whose elements $x_{IIFs,i}$ represent the investment in n GI types, $s \in S_I$, $x_{IIFs} \in R^{1\times n} L_F$: a binary vector whose elements $L_{F,i}$ indicate whether (=1) or not (=0) FL occurs for each of the n GI types, $L_F \in R^{n\times 1} L_P$: a binary vector whose elements $L_{P,i}$ indicate whether or not PL occurs for each of the n GI types, $L_P \in R^{n\times 1} L_N$: a binary vector whose elements $L_{N,i}$ indicate whether or not NL occurs for each of the n GI types, $L_N \in R^{n\times 1}\tau$: an auxiliary variable used to calculate VaR_α, $\tau \in R z_{s,r}$: the stormwater reduction below τ in scenario *s,* $s \in S_I$, $r \in S_{II} = \{1,2,...,m_{II}\}$, $z_{s,r} \in R$Constants. A: a matrix of resource (land, budget, etc.) consumption rates per unit of investment in each GI type for each of k resource constraints, $A \in R^{k\times n} B$: resource upper bounds. $B \in R^k \mu$: the mean of the prior distribution of stormwater reduction rates for the n GI types at Stage *I*. $\mu = E[\widetilde{C_I}]$, $\mu \in R^{1\times n} C_{Is}$: the realization of the stormwater reduction rate in Stage *I* in scenario s for $s \in S_I$, $C_{Is} \in R^{1\times n} C_{IPs}$: the expected value of the posterior of the stormwater reduction rate for the PL case in scenario s for all $s \in S_I$, $C_{IPs} \in R^{1\times n} C_{IIFs}$: the second stage realization of the stormwater reduction rate for the FL case in scenario s for $s \in S_I$, $C_{IIFs} \in R^{1\times n} C_{IINs,r}$: the second stage realization of the stormwater reduction rate for the NL case in scenario (*s, r*), $s \in S_I$ and $r \in S_{II}$, $C_{IINs,r} \in R^{1\times n} C_{IIPs,r}$: the second stage realization of the stormwater reduction rate for the PL case in scenario (*s, r*), $s \in S_I$ and $r \in S_{II}$, $C_{IIPs,r} \in R^{1\times n}$ $CVaR_\alpha$: the user-specified tolerable risk level for the second stage annual stormwater reduction, $\alpha \in (0,1)$. $CVaR_\alpha \in R$

M: a large number (e.g., $M = 10^6$).

T^I, T^{II}: the number of years in the planning horizon of Stages *I* and *II*, respectively, T^I, $T^{II} \in R Th^{FL}$, Th^{PL}: the investment thresholds for FL and PL, respectively, Th^{FL}, $Th^{PL} \in \mathrm{R}^{1\times n}$. $Th_i^{FL}$and Th_i^{PL}, $i \in N = \{1,2,...,n\}$ are the elements of their respective vectors.Objective and constraints.

$$Maximize\ f_I(x_I, x_{IIN1}, ..., x_{IINm_I}, x_{IIF1}, ..., x_{IIFm_I}, x_{IIP1}, ..., x_{IIPm_I}) = \mu x_I + \frac{T^{II}}{T^I + T^{II}}\left(\frac{1}{m_I}\sum_{s=1}^{m_I}(\mu x_{IINs} + C_{IIFs}x_{IIFs} + C_{IPs}x_{IIPs})\right) \tag{A.3}$$

Subject to:

$$A(x_I + x_{IINs} + x_{IIFs} + x_{IIPs}) - B \le 0,\ \forall\ s \in S_I, \text{(Resource Constraints)} \tag{A.3.1}$$

$$\begin{cases} -x_{I,i} + Th_i^{PL} L_{p,i} \le 0 \\ -x_{I,i} + Th_u^{FL} L_{F,i} \le 0 \\ x_{I,i} - ML_{F,i} \le Th_i^{FL} \\ x_{I,i} - M(L_{P,i} + L_{F,i}) \le Th_i^{PL} \\ L_{N,i} + L_{P,i} + L_{F,i} = 1 \end{cases},\ \forall\ u \in \{1,2, .., n\},\ \text{(Learning Constraints)} \tag{A.3.2}$$

$$\begin{cases} x_{IINs,i} - ML_{N,i} \leq 0 \\ x_{IIPs,i} - ML_{P,i} \leq 0 \\ x_{IIFs,i} - ML_{F,i} \leq 0 \end{cases}, \quad \forall\ u \in \{1,2,..,n\} \text{and} \quad \forall\ \ s \in S_I \quad \text{(Auxiliary Constraints)} \tag{A.3.3}$$

$$\begin{cases} z_{s,r} \geq \tau - f_{IIs,r}(x_I, x_{IINs}, x_{IIFs},\ x_{IIPs}),\ \forall\ s \in S_I,\ \ r \in S_{II} \\ \tau - \frac{1}{(1-\alpha)m_I m_{II}} \sum_{s=1}^{m_I} \sum_{r=1}^{m_{II}} z_{s,r} \geq CVaR_\alpha \end{cases}, \text{ (CVaR Constraints)} \tag{A.3.4}$$

$$f_{IIs,r}(x_I, x_{IINs}, x_{IIFs},\ x_{IIPs}) = C_{Is}x_I + C_{IINs,r}x_{IINs} + C_{IIFs}x_{IIFs} + C_{IIPs,r}x_{IIPs},\ \forall\ s \in S_I,\ \ r \in S_{II} \tag{A.3.5}$$

$$x_I,\ x_{IINs}, x_{IIPs},\ x_{IIFs} \geq 0,\ \forall\ s \in S_I \ \text{ (Non − negativity)} \tag{A.3.6}$$

The objective A.3 maximizes the expected annual average stormwater reduction over the entire time horizon, with the second stage reduction being discounted by the ratio of the second stage planning horizon to the total planning years. Constraint A.3.1 sets a lower bound for the resources, e.g., budget and suitable sites. Constraints A.3.2 and A.3.3 are logical constraints that identify which learning type happens in the second stage by setting the upper bound of the corresponding second stage decision variable to M while restricting the upper bounds of the other second stage decision variables to zero. Constraints A.3.4 enforce a lower bound on CVaR, which requires the evaluation of the stormwater reduction in each scenario (s,r) by Eq. (A.3.5). Finally, constraints A.3.6 force the investment decisions to be non-negative.

A.3. The TI model

We assume that the investment could result in a technology improvement which increases the mean of the posteriors. Therefore, the realized reduction in Stage II in the FL case (denoted $C_{IIFs,TI}$) should be sampled from a distribution with mean $\gamma\mu$, γ>1 to account for "improvement by learning." In the PL case, the posterior mean $C_{IPs,TI}$ for each s is generated from $\widetilde{C_{IP,TI}}$, where $\text{Var}(\widetilde{C_{IP,TI}}) = (1-\beta)\sigma^2$, $\beta \in (0,1)$ and $E[\widetilde{C_{IP,TI}}] = \gamma\mu$, denoted $C_{IPs,TI}$. The random second stage samples in the PL case are generated from $\widetilde{C_{IIs,TI}}$ ($C_{IPs,TI}$, $\beta\sigma^2$). The objective f_I(A.3) and $f_{IIs,r}$ (A.3.5) in the ML model are revised as follows while the form of the constraints remains unchanged.

Objective:

$$Maximize\ f_I\left(x_I, x_{IIN1}, ..., x_{IINm_I}, x_{IIF1},...,x_{IIFm_I},\ x_{IIP1},...,x_{IIPm_I}\right) Maximize\ f_I\left(x_I, x_{IIN1}, ..., x_{IINm_I}, x_{IIF1},...,x_{IIFm_I},\ x_{IIP1},...,x_{IIPm_I}\right) = \mu x_I + \frac{T^{II}}{T^I + T^{II}}\left(\frac{1}{m_I}\sum_{s=1}^{m_I}(\mu x_{IINs} + C_{IIFs,TI}x_{IIFs} + C_{IPs,TI}x_{IIPs})\right) \tag{A.4}$$

$$f_{IIs,r}(x_I, x_{IINs}, x_{IIFs},\ x_{IIPs}) = C_{Is}x_I + C_{IINs,r}x_{IINs} + C_{IIFs,\ TI}x_{IIFs} + C_{IIPs,r,\ TI}x_{IIPs},\ \text{ for all } s \in S_I,\ \ r \in S_{II} \tag{A.4.1}$$

References

Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. Math. Finance 9 (3), 203–228.

Baker, E., 2009. Optimal policy under uncertainty and learning about climate change: a stochastic dominance approach. J. Publ. Econ. Theor. 11 (5), 721–747.

Barton, D.N., Kuikka, S., Varis, O., Uusitalo, L., Henriksen, H.J., Borsuk, M., La Hera, A.D., Farmani, R., Johnson, S., Linnell, J.D.C., 2012. Bayesian networks in environmental and resource management. Integrated Environ. Assess. Manag. 8 (3), 418–429.

Beh, E.H.Y., Zheng, F., Dandy, G.C., Maier, H.R., Kapelan, Z., 2017. Robust optimization of water infrastructure planning under deep uncertainty using metamodels. Environ. Model. Software 93, 92–105.

Berglund, C., Söderholm, P., 2006. Modeling technical change in energy system analysis: analyzing the introduction of learning-by-doing in bottom-up energy models. Energy Pol. 34, 1344–1356.

Bormann, B.T., Cunningham, P.G., Brooks, M.H., Manning, V.W., Callopy, M.W., 1994. Adaptive ecosystem management in the Pacific Northwest. General Technical Report. .

Chao, P.T., Hobbs, B.F., 1997. Decision analysis of shoreline protection under climate change uncertainty. Water Resour. Res. 33 (4), 817.

Chocat, B., Ashley, R., Marsalek, J., Matos, M.R., Rauch, W., Schilling, W., Urbonas, B., 2007. Toward the sustainable management of urban stormwater. Indoor Built Environ. 16 (3), 273–285.

Environmental Protection Agency, 2010. Green Infrastructure Case Studies: Municipal Policies for Managing Stormwater with Green Infrastructure.

Ferioli, F., Schoots, K., van der Zwaan, B.C.C., 2009. Use and limitations of learning curves for energy technology policy: a component-learning hypothesis. Energy Pol. 37, 2525–2535.

Golub, A., Narita, D., Schmidt, M.G.W., 2014. Uncertainty in integrated assessment models of climate change: alternative analytical approaches. Environ. Model. Assess. 19 (2), 99–109.

Guo, J.C.Y., 2002. Overflow risk analysis for stormwater quality control basins. J. Hydrol. Eng. 7 (6), 428–434.

Harrison, K.W., 2007. Two-stage decision-making under uncertainty and stochasticity: Bayesian programming. Adv. Water Resour. 30 (3), 641–664.

Holling, C.S., 1978. Adaptive Environmental Assessment and Management. John Wiley & Sons.

Housh, M., Ostfeld, A., Shamir, U., 2013. Limited multi-stage stochastic programming for managing water supply systems. Environ. Model. Software 41, 53–64.

Jacobi, S., Hobbs, B., Wilcock, P., 2013. Bayesian optimization framework for cost-effective control and research of non-point-source sediment. J. Water Resour. Plann. Manag. 139 (5), 534–543.

Ji, X., Zhu, S., Wang, S., Zhang, S., 2005. A stochastic linear goal programming approach to multistage portfolio management based on scenario generation via linear programming. IIE Trans. 37, 957–969.

Kasprzyk, J.R., Nataraj, S., Reed, P.M., Lempert, R.J., 2013. Many objective robust decision making for complex environmental systems undergoing change. Environ. Model. Software 42, 55–71.

Kelly, D.L., Kolstad, C.D., 1999. Bayesian learning, growth, and pollution. J. Econ. Dynam. Contr. 23, 491–518.

Kim, B.J., Hobbs, B.F., Koonce, J.F., 2003. Multicriteria Bayesian analysis of lower trophic level uncertainties and value of research in Lake Erie. Hum. Ecol. Risk Assess. 9 (4), 1023–1057.

Krokhmal, P., Palmquist, J., Uryasev, S., 2001. Portfolio optimization with conditional value-at-risk objective. J. Risk 4 (2), 43–68.

Lan, F., Lin, W.H., Lansey, K., 2015. Scenario-based robust optimization of a water supply system under risk of facility failure. Environ. Model. Software 67, 160–172.

Lee, J.G., Heaney, J.P., Lai, F., 2005. Optimization of integrated urban wet-weather control strategies. J. Water Resour. Plann. Manag. 131 (4), 307–315.

Lempert, R.J., Collins, M.T., 2007. Managing the risk of uncertain threshold responses: comparison of robust, optimum, and precautionary approaches. Risk Anal. 27 (4), 1009–1026.

Loáiciga, H.A., Asce, F., Sadeghi, K.M., Shivers, S., Kharaghani, S., 2013. Stormwater control measures: optimization methods for sizing and selection. J. Water Resour. Plann. Manag. 139, 554–564.

Lovell, S.T., Taylor, J.R., 2013. Supplying urban ecosystem services through multifunctional green infrastructure in the United States. Landsc. Ecol. 28 (8), 1447–1463.

Massoudieh, A., Maghrebi, M., Kamrani, B., Nietch, C., Tryby, M., Sassan, A., Panguluri, S., 2017. A flexible modeling framework for hydraulic and water quality performance assessment of stormwater green infrastructure. Environ. Model. Software 92, 57–73.

McCarthy, M.A., Possingham, H.P., 2007. Active adaptive management for conservation. Conserv. Biol. 21 (4), 956–963.

McGarity, A.E., 2012. Storm-water investment strategy evaluation model for impaired urban watersheds. J. Water Resour. Plann. Manag. 138 (2), 111–124.

Monahan, G.E., 1982. State of the art—a survey of partially observable Markov decision processes: theory, models, and algorithms. Manag. Sci. 28 (1), 1–16.

Montalto, F., Behr, C., Yu, Z., 2011. Accounting for uncertainty in determining green infrastructure cost-effectiveness. In: Economic Incentives for Stormwater Control. CRC Press, pp. 71–100.

Montaseri, M., Hesami Afshar, M., Bozorg-Haddad, O., 2015. Development of simulation-optimization model (MUSIC-GA) for urban stormwater management. Water Resour. Manag. 29 (13), 4649–4665.
Morales-Torres, A., Escuder-bueno, I., Andr, I., Perales-momparler, S., 2016. Decision Support Tool for energy-efficient, sustainable and integrated urban stormwater management. Environ. Model. Software 84, 518–528.
Mortazavi-Naeini, M., Kuczera, G., Kiem, A.S., Cui, L., Henley, B., Berghout, B., Turner, E., 2015. Robust optimization to secure urban bulk water supply against extreme drought and uncertain climate change. Environ. Model. Software 69, 437–451.
Novotny, V., Witte, J.W., 1997. Ascertaining aquatic ecological risks of urban stormwater discharges. Water Res. 31 (10), 2573–2585.
Paydar, Z., Qureshi, M.E., 2012. Irrigation water management in uncertain conditions-Application of modern portfolio theory. Agric. Water Manag. 115, 47–54.
Philadelphia Water Department, 2011. Green City Clean Waters.
Piantadosi, J., Metcalfe, A.V., Howlett, P.G., 2008. Stochastic dynamic programming (SDP) with a conditional value-at-risk (CVaR) criterion for management of storm-water. J. Hydrol. 348, 320–329.
Rockafellar, R., Uryasev, S., 2000. Optimization of conditional value at risk. J. Risk 2 (3), 21–41.
Rossi, L., Krejci, V., Rauch, W., Kreikenbaum, S., Fankhauser, R., Gujer, W., 2005. Stochastic modeling of total suspended solids (TSS) in urban areas during rain events. Water Res. 39, 4188–4196.
Roy, A.H., Wenger, S.J., Fletcher, T.D., Walsh, C.J., Ladson, A.R., Shuster, W.D., Thurston, H.W., Brown, R.R., 2008. Impediments and solutions to sustainable, watershed-scale urban stormwater management: lessons from Australia and the United States. Environ. Manag. 42, 344–359.
Sebti, A., Fuamba, M., Bennis, S., 2016. Optimization model for BMP selection and placement in a combined sewer. J. Water Resour. Plann. Manag. 142, 04015068.
Shapiro, A., Philpott, A., 2007. A tutorial on stochastic programming. Retrieved 2018. Available. https://www2.isye.gatech.edu/people/faculty/Alex_Shapiro/TutorialSP.pdf.
Söderholm, P., Sundqvist, T., 2007. Empirical challenges in the use of learning curves for assessing the economic prospects of renewable energy technologies. Renew. Energy 32 (15), 2559–2578.
Southwell, D.M., Hauser, C.E., McCarthy, M.A., 2016. Learning about colonization when man-aging metapopulations under an adaptive management framework. Ecol. Appl. 26 (1), 279–294.
Varouchakis, E.A., Palogos, I., Karatzas, G.P., 2016. Application of Bayesian and cost benefit risk analysis in water resources management. J. Hydrol. 534, 390–396.
Walters, C.J., Hilborn, R., 1976. Adaptive control of fishing systems. J. Fish. Res. Board Can. 33, 145–159.
Wang, S., Huang, G.H., 2014. An integrated approach for water resources decision making under interactive and compound uncertainties. Omega 44, 32–40.
Webby, R.B., Adamson, P.T., Boland, J., Howlett, P.G., Metcalfe, A.V., Piantadosi, J., 2007. The Mekong-applications of value at risk (VaR) and conditional value at risk (CVaR) simulation to the benefits, costs and consequences of water resources development in a large river basin. Ecol. Model. 201 (1), 89–96.
Webster, M., Fisher-Vanden, K., Popp, D., Santen, N., 2017. Should we give up after solyndra? Optimal technology R&D portfolios under uncertainty. J. Assoc. Environ. Resour. 4, 123–151.
White, C.C., 1991. A survey of solution techniques for the partially observed Markov decision process. Ann. Oper. Res. 32, 215–230.
Williams, B.K., 2011. Adaptive management of natural resources-framework and issues. J. Environ. Manag. 92, 1346–1353.
Williams, B.K., Brown, E.D., 2014. Adaptive management: from more talk to real action. Environ. Manag. 53, 465–479.
Zhen, X.-Y., Yu, S.L., Lin, J.-Y., 2004. Optimal location and sizing of stormwater basins at watershed scale. J. Water Resour. Plann. Manag. 130 (4), 339–347.