

Long-term adaption decisions via fully and partially observable Markov decision processes

Olga Špačková & Daniel Straub

Engineering Risk Analysis Group, Technical University Munich (olga.spackova@era.bgu.tum.de)

Abstract

Long term decisions, such as the design of infrastructure systems and buildings or the planning of risk mitigation measures, should be made in consideration of the uncertain future. The initial design of a system determines its flexibility, i.e. its ability to cope with potential future changes. Increasing flexibility is generally considered to be a good approach to dealing with future uncertainty, such as climate change uncertainty, but its effects have not been systematically investigated. We propose the use of Markov Decision Processes combined with Influence Diagrams to solve adaptation planning problems. This framework can identify the optimal system type and capacity and determine the value of flexibility. It is here applied to two numerical examples: Planning of a waste water treatment plant under uncertainty in future population growth and planning of a flood protection system under uncertain climate change scenarios. Based on these idealized examples, it is shown that for flexible systems a lower initial capacity of the system is recommendable, while for inflexible systems a conservative design (with high safety factors) should be applied. The value of flexibility is shown to be high when significant learning is expected in the future, i.e. if information gathered in the future significantly reduces uncertainty.

Keywords

Decision-making under uncertainty; infrastructure systems; Bayesian updating; Markov decision processes; climate change; flexibility; changeability; real options; cost-benefit analysis.

1 Introduction

Infrastructure systems typically have a lifetime of many decades and the future demand on their capacity is highly uncertain (Hall et al., 2014; Yzer et al., 2014). This is the case for water distribution and wastewater treatment, energy production and distribution, transportation and flood protection systems. Similarly, land use planning has long term implications, and buildings

have user requirements that can change rapidly and cannot be predicted with certainty. The design and planning of such systems should therefore address uncertainty on future climate, socioeconomic development, changing societal preferences and changing needs arising from new knowledge and information. An overview of example planning problems and related future uncertainties is given in Table 1.

Table 1. Example infrastructure planning problems and their demand uncertainties

Aim/problem	Future uncertainties	Demand parameter(s)	Observability full/partial (indicative)
River flood protection	Magnitude of flood events	Flood discharge/volume	P
	Damage potential	No. of inhabitants, monetary value	F
Storm surge protection	Magnitude of flood events	Water level, wave height	P
	Damage potential	No. of inhabitants, monetary value	F
Waste water collection and treatment	Volume of waste water to be treated	Volume/population equivalent	F
	Magnitude of extreme rainfall events, flash floods	Rainfall intensity (short term)	P
Water resources	Demand on water resources	No. of inhabitants, economic/industrial descriptors	F
	Availability of water resources (rainfall/drought characteristics)	Precipitation volume (long-term)	P
Transport /power infrastructure design	Demand on the infrastructure	No. of inhabitants/users, economic/industrial descriptors	F
	Environmental loads (temperature, precipitation, weather extremes)	Design temperature and precipitation	P
Building design	Utilization of the building, user requirements	Building use characteristics	F
	Environmental loads (temperature, weather extremes)	Design temperature, design wind/snow load	P
Land use planning, urban planning	Need of the land use categories,	No. of inhabitants, economic/industrial descriptors	F
	Environmental loads (weather extremes,	Design temperature,	P

Systematic approaches to deal with the uncertain future are often referred to as adaptation strategies, especially in the context of climate change uncertainty. Frankhauser et al. (1999) and Hallegatte (2009) summarize and qualitatively assess adaptation strategies from multiple sectors. In some problem settings, the best adaptation strategy is obvious, especially if the number of alternatives is limited. However, in most cases there exist a potentially large number of alternative solutions: the decision maker is not only choosing from different types of adaptation strategies (e.g. building a new water reservoir or reducing water losses), but can also decide about the dimensions/capacity of the proposed infrastructures (e.g. capacity of the water reservoir) and often about the flexibility of the system as well, i.e. its ability to be changed in the future without excessive costs, as defined later in Sec. 2. Decisions analysis can help identifying the best solution in these cases.

Hallegatte et al. (2012) and Cox (2012) provide an overview of decisions support methodologies that have been used in research and practice for dealing with the uncertain future. The commonly used approach to account for future uncertainties and for evaluating the flexibility of engineering systems is the so-called Real Option (RO) analysis (Martins et al., 2015). The term was adopted from the field of finance. The original formulation, with quantities such as volatility, current and exercise price or expiration date, and the related models and evaluation methods, are not directly applicable to engineering questions. Nevertheless, the term RO analysis is today used for a broad spectrum of methods, which typically involve decision trees for representing the problem and Monte Carlo (MC) simulation and different optimization algorithms to evaluate them. For example, Neufville et al. (2006) propose to use RO for evaluating different designs of infrastructure and buildings with uncertain future number of users. Buurman et al. (2009) present the application of RO to design of maritime safety systems under uncertain future terrorism threats. Gersonius et al. (2012) and Woodward et al. (2014) use RO analysis for optimizing the water systems for urban drainage and flood defense, respectively, under climate change uncertainty. Marques et al. (2014) utilize RO analysis to find optimal design of water supply system under uncertain future demand. Martínez Ceseña et al. (2013) provide an overview of RO applications in the planning of energy generation projects under different types of uncertainty, including the uncertainty on future energy prices or uncertainty on the probabilistic model describing the resource characteristics (wind, water). Fawcett et al. (2014) use RO for highway design under traffic growth uncertainty. Further studies use the same methods without explicitly calling them RO analysis, e.g. (Basupi and Kapelan, 2014). For complex problems with many decision alternatives, the RO analysis using decision trees becomes intractable, the examples presented in the literature thus focus on rather simpler problems. Additionally, they mostly do not consider the fact that also the future decisions will be made under uncertainty, even if reduced.

This paper proposes the use of (fully observable) Markov Decision Processes (MDPs) and Partially Observable Markov Decision Processes (POMDPs) for optimizing long-term adaptation strategies under uncertainty. The proposed approach is general enough to replace RO analysis with decision trees for most applications, and it offers additional possibilities and advantages: (1) Modeling of the uncertain demand as a Markov process increases the computational efficiency of the model and thus the ability to tackle complex problems. (2) The introduction of partial observability (i.e. use of POMDPs) allows explicit modeling of the fact that the demand on the

system will also be uncertain at the time of the future decisions. This is particularly relevant for problems where the system is designed for an extreme event, i.e. where the demand on the systems corresponds e.g. to a flood discharge, extreme precipitation or extreme wind speed. The estimates of frequency/magnitude of such extreme events are highly uncertain because they rely on limited historic data and imperfect models and the uncertainty thus remains also at the time of future decisions. (See Table 1 for an indication of partial/full observability and Section 3.3 for a more detailed discussion.) (3) Using influence diagrams for graphically representing the problem facilitates communication of the model and its assumptions. Influence diagrams are significantly more concise than classical decision trees, which is a major advantage when representing complex problems with many decision options and many decision points in time. A comparable approach has been proposed by (Nishijima, 2015) for the case of climate adaptation, but without considering partial observability of the processes. More recently, (Pozzi et al., under review) use hidden-model MDP for optimizing climate adaptation decisions and they investigate the effect of future observations and learning on the decisions. The flexibility of the systems is not included in any of these studies.

The decision criterion is formulated in analogy to classical static Cost Benefit Analysis (CBA), which defines the optimum as the solution that maximizes the difference between expected benefits and expected costs. Risk, which is defined as expected damage, can be considered as an additional expected cost. The objective function in static CBA is formulated as (Špačková and Straub, 2015):

$$\mathbf{v}_{opt} = \underset{\mathbf{v}}{\operatorname{argmax}}[b(\mathbf{v}) - c(\mathbf{v}) - r(\mathbf{v})] \quad (1)$$

where \mathbf{v} are parameters characterizing the system such as height of flood protection dikes, capacity of a waste water treatment plant or the size of a restricted area in a land use plan. $b(\mathbf{v})$ and $c(\mathbf{v})$ are the expected net present values of benefits and costs, and $r(\mathbf{v})$ is the net present value of risk.

Unlike classical static CBA, the proposed framework uses a dynamic formulation, which explicitly models the future uncertainties in these quantities and the decisions that will be made in the future, conditionally on future observations (the preposterior principle). In the dynamic formulation, the objective function of Eq. (1) is modified as described in Sec. 3.1 (see Eq. (4)). The future decisions are thus optimized jointly with the initial capacity. Also the optimal timing of future system adjustments is automatically identified.

In the proposed dynamic formulation, the ability of the systems to be changed in the future is explicitly modeled. A quantitative measure of flexibility adapted from (Špačková et al., 2015) is proposed, which is a function of the costs for adjusting the system in the future when new needs arise. The decision model can be applied in two ways: (1) When the flexibility of the system is fixed, the optimal initial capacity of the system is selected taking into account its flexibility. (2) When the flexibility is an optimization parameter, the model can be used to answer the question, how much should one pay initially to make the system flexible.

2 Flexibility

Flexibility describes the potential of a system to be changed in the future without excessive costs. Flexibility has been recognized as an important quality in engineering design (Fricke and Schulz, 2005; Ross et al., 2008; Saleh et al., 2009). Alternatively, the terms adaptability, reversibility or changeability are used in the same context, with slightly different meanings. Examples of flexible systems in infrastructure planning include:

- flood protection systems whose possible future upgrades are facilitated by reserving land for future construction of additional dikes or extension of existing ones (Vrijling et al., 2009);
- an urban drainage system where the backbone parts are oversized to allow extension of the system in the future or which enables cheap adjustment of the retention capacity in the future (Radhakrishnan et al., 2014);
- a motorway that is built broader so that additional lanes can be established in the future without the need for extensive changes to the road structure (Lamb et al., 2011);
- a building design enabling cheap future conversion, e.g. dividing large apartments into smaller ones, turning an office building into a block of flats or expanding a building (Allahaim and Alfaris, 2010; Boehland, 2003).

We propose a measure of flexibility following (Špačková et al., 2015), which expresses how costly it is to adjust the capacity v of a system in the future, relative to the initial investment at time t_0 . The flexibility φ is defined as

$$\varphi = \frac{c_0(v_{ad}) - c_A(v_{or}, v_{ad})}{c_0(v_{or})}, \quad (2)$$

where $c_0(v)$ is the cost of implementing a system with capacity v initially; $c_A(v_{or}, v_{ad})$ is the undiscounted cost of a future adjustment of the system capacity from v_{or} to v_{ad} .

Flexibility $\varphi = 0$ corresponds to an inflexible system and $\varphi = 1$ to a fully flexible system. For fully flexible systems with $\varphi = 1$, the cost of increasing the capacity from v_{or} to v_{ad} equals the difference between building to these capacities initially, $c_A(v_{or}, v_{ad}) = c_0(v_{ad}) - c_0(v_{or})$. For such systems, the number of steps to reach the final capacity is irrelevant; the total costs are the same if one builds to the final capacity at once or if one adjusts the capacity at every time step. For inflexible systems with $\varphi = 0$, the cost of increasing the capacity to v_{ad} is equal to cost of building to this capacity initially, irrespective of the existing capacity v_{or} , therefore $c_0(v_{ad}) = c_A(v_{or}, v_{ad})$. When the capacity of an inflexible system is changed, the existing system cannot be used, the system has to be built completely anew. Flexibility can also take negative values, which occur if the original system must be fully replaced and removal costs are invoked in addition to the cost of the new system. The flexibility of an irreversible system is $-\infty$, because its adjustment costs are infinite.

Because this measure of flexibility is defined in function of costs, it is suitable to be utilized in a quantitative optimization. This is in contrast to previously proposed measures of flexibility (e.g. Kumar, 1987; Ross et al., 2008; Sönmez and Koç, 2015), which primarily focus on quantifying the number of ways the system can be changed. In these measures, costs are included only as a constraint for accepting/rejecting the options, if at all.

For optimizing infrastructure planning, one requires the adjustment cost relative to the initial cost. It follows from Eq. 1 that the cost of changing the capacity from v_{or} to v_{ad} is

$$c_A(\varphi, v_{ad}, v_{or}) = c_0(v_{ad}, \varphi) - \varphi \cdot c_0(v_{or}, \varphi). \quad (3)$$

The difference in costs for systems with different flexibility is illustrated in Figure 1. Note that in this figure we assume that the initial costs independent of flexibility φ (unlike in Eq. (3)).

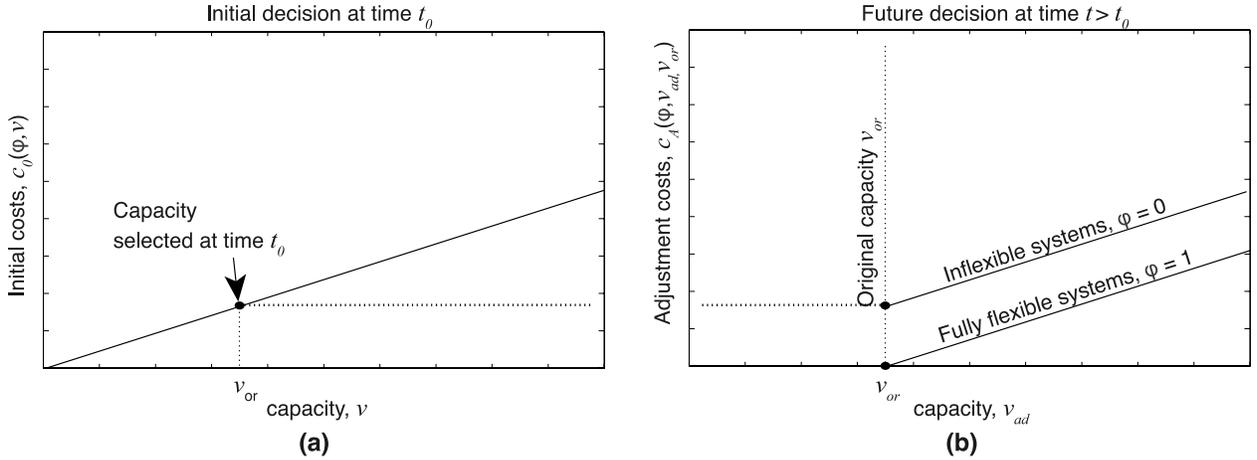


Figure 1: Illustration of the costs for a fully flexible vs. an inflexible system. (a) Initial costs when the system is first implemented; (b) adjustment costs for adapting the system to new needs in the future, when the original capacity is v_{or} . Adapted from (Špačková et al., 2015).

The proposed definition of flexibility is kept simple on purpose. In most infrastructure planning problems, the flexibility cannot be computed exactly, nor is it a constant value. It varies in function of the initial capacity v_{or} . Furthermore, typically only a discrete set of the planning alternatives are evaluated, and the flexibility can thus be quantified only for those. Nevertheless, the value of flexibility can be determined approximately based on simple considerations, and the concept then enables general statements and conclusion. If a more detailed model is desired, the flexibility can be expressed in function of v_{or} and v_{ad} .

In this paper, the capacity γ is expressed as a single scalar value defined on a continuous scale. It can be extended to the multidimensional case, i.e. a vector of capacity values, if a single value is not sufficient to express the system performance.

3 The planning problem

The aim is to optimize the capacity v of an infrastructure system with given flexibility over the system lifetime. In this section, different formulations of this decision problem are proposed, which should be selected depending on the characteristics of the problem. We first introduce the variables that are used in the model and formulate the objective of the optimization (Sec 3.1). In section 3.2, the value of flexibility is defined, which allows to quantify the gain from establishing a flexible system as compared to an inflexible one. Section 3.3 discusses the difference between

fully and partially observable random processes. Finally, solution strategies and algorithms for the proposed models are presented in Sec. 3.4.

3.1 Modelling framework

The system demand is determining the design of the infrastructure system. Examples of demand variables are given in the 3rd column of Table 1. The uncertain evolution of the demand in the future is modeled through a discrete-time Markov process $\{\theta_i, i \in \mathbf{T}\}$, where $\mathbf{T} = \{1, 2, \dots, N\}$ are the discrete time steps (e.g. years), as illustrated in Figure 2. This process is commonly called the “core process”. It is described by the transition probability $p(\theta_{i+1}|\theta_i, v_i)$, i.e. the probability of the demand being in state θ_{i+1} at time $i + 1$ given that it was in state θ_i at time i and capacity v_i was selected.

In many real applications, a direct modeling of the demand as a Markov process is not possible, because the Markov independence property, i.e. the assumption that the future of the system is independent on the past given the present, is not fulfilled. However, even non-Markovian processes can be transformed to Markovian processes through augmentation of the state space with additional variables (Rachelson et al., 2008). In such cases, the variable θ_i in each step is replaced with a set of variables Θ_i . This approach is demonstrated in the numerical example in Sec. 5 on optimizing flood protection, where uncertainty on future extreme rainfall events is modeled using 2 variables: one variable representing the climate scenario and the related trend (i.e. a “memory” in the stochastic process) and a second variable representing the annual maximum rainfall in each year.

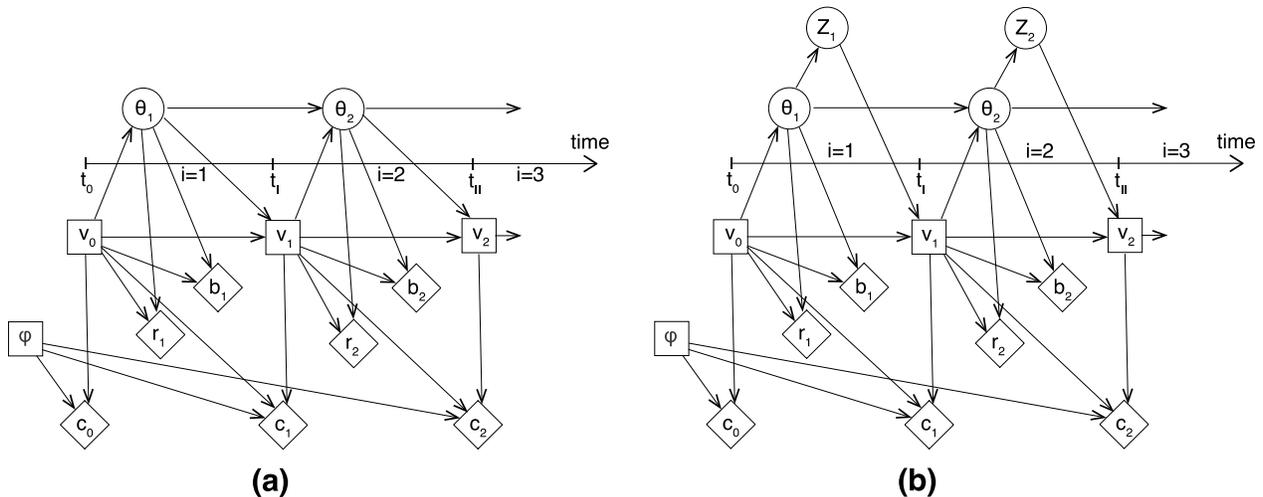


Figure 2: Influence diagrams of the planning problem: (a) Fully observable Markov Decision Process (MDP) (b) Partially Observable Markov Decision Process (POMDP). Circles denote the random variables, squares the decision nodes and diamonds the utility (benefit/costs) nodes. The arrows represent the dependence structure and the availability of information for making decisions. An arrow pointing towards a decision indicates that the corresponding parent node is known at the time of making the decision. The semantics of these influence diagrams follow those of (Jensen and Nielsen, 2007).

If the core process is not directly observable, i.e. if the demand θ_i cannot be determined with certainty even at time i , a random variable Z_i representing the observation at time step i is included in the model (Figure 2b). Observation Z_i provides indirect information about the

demand θ_i . The relationship between θ_i and Z_i is described by the conditional probability $p(z_i|\theta_i)$, which corresponds to the likelihood function. The stochastic process $\{Z_i, i \in \mathbf{T}\}$ is called the “observation process”.

Decisions on the capacity v are (potentially) made at times t_0, t_1, \dots, t_N . At t_0 , the initial capacity v_0 is selected, at future time steps t_1, \dots, t_N the capacity can be adjusted based on the current state of knowledge. In Figure 2, decisions are made at each time step (e.g. every year). In many applications, such frequent adjustment decisions are not possible or reasonable. Decisions are then considered only every Δt -th time step, as discussed in Sec. 3.4.

The system demand and capacity determine the benefits (expected gains) and risks (expected losses) throughout the lifetime of the system. Let $b(\theta_i, v_{i-1})$ denote the benefits and $r(\theta_i, v_{i-1})$ the risk in the i th time step. They are functions of the demand θ_i and the capacity v_{i-1} , which is selected at time t_{i-1} . $c_i(\varphi, v_i, v_{i-1})$ is the cost of adjusting the system capacity from v_{i-1} to v_i according to Eq. (3).

Benefits minus costs and risk is the reward W . The goal of the optimization is finding the optimal initial capacity that maximizes the expected value of the discounted life-time reward for given flexibility:

$$v_{0,opt}(\varphi) = \underset{v_0}{\operatorname{argmax}} W_0(\varphi, v_0), \quad (4)$$

where $W_0(\varphi, v_0)$ is the expected reward for given flexibility and initial capacity:

$$W_0(\varphi, v_0) = -c_0(\varphi, v_0) + \mathbb{E} \left[\sum_{i=1}^N [b_i(\theta_i, v_{i-1}) - r_i(\theta_i, v_{i-1}) - c_i(\varphi, v_i, v_{i-1})] \right]. \quad (5)$$

$c_0(\varphi, v_0)$ are the initial investment costs, and $\mathbb{E}[\cdot]$ is the expected value over all future demands, obtained with the future optimal decisions v_1, \dots, v_N . This is a sequential decision problem, in which the optimal initial capacity v_0 can only be determined by simultaneously optimizing all future capacities (Kochenderfer et al., 2015).

In some cases, the aim can additionally be to select the optimal flexibility of the system:

$$\varphi_{opt} = \underset{\varphi}{\operatorname{argmax}} W_0(\varphi, v_{0,opt}(\varphi)). \quad (6)$$

3.2 Value of flexibility

The value of flexibility is the additional cost that a decision maker should be willing to spend for a flexible system relative to an inflexible one. It can be determined as

$$V(\varphi) = W_0^*(\varphi, v_{0,opt}(\varphi)) - W_0(\varphi = 0, v_{0,opt}(\varphi = 0)), \quad (7)$$

where $W_0^*(\varphi, v_{0,opt}(\varphi))$ is computed according to Eq. (5) under the premise that flexibility is for free, i.e. that the initial cost $c_0(\varphi, v_0)$ is the same for all flexibilities φ (as in Figure 1).

3.3 Fully or partially observable demands

If the system demand is fully observable, its current value is known at the time of making the decision; it is only the uncertainty about its future development that remains. In practice, there is always some uncertainty in observations of the demand, but in many applications this uncertainty is negligible compared to other uncertainties in the model. The problem then can be modeled as a fully observable Markov Decision Process (MDP), shown in Figure 2a. If significant uncertainty about the state of demand at the times of the decisions exists, the problem must be modeled as Partially Observable Markov Decision Process (POMDP) shown in Figure 2b.

To give examples, the number of inhabitants and the built area of a city, which determine the demand on the sewerage system, can typically be observed with high accuracy. If these are the only variables representing the future demand, the process can be considered as fully observable and modeled as MDP. In contrast, values of assets at risk (buildings, infrastructure) and their vulnerability to a natural hazard are only indirectly deducible from socioeconomic data and available damage models. Significant uncertainty therefore remains in estimating the consequences of a hazard event. Similarly, discharge measurements provide limited information about the full probabilistic distribution of the discharge maxima needed for the design of flood protection. The parameters of the probabilistic distribution of the annual maximum discharge process are thus only partially observable. In these cases, POMDP should be used to model the decision problem.

Table 1 summarizes the observability of demand parameters in selected systems. This list is only indicative and should not be understood as fixed rules on the modeling approach. The ultimate selection of the model type is application-specific and depends on the data availability.

3.4 Evaluation of the decision processes

A number of algorithms for solving MDPs and POMDPs has been developed in the last decades, and this continues to be an active field of research (Monahan, 1982; Kaelbling et al., 1998; Braziunas, 2003; Poupart, 2005, Kochenderfer et al., 2015). Applications of MDPs and POMDPs can be found in different sectors, including planning of inspections in engineering systems (Ellis et al., 1995; Corotis et al., 2005; Robelin and Madanat, 2007; Luque and Straub, 2013; Memarzadeh and Pozzi, 2015), autonomous driving (Brechtel et al., 2014), medical treatment (Alagoz et al., 2009), personal assistance (Hoey et al., 2010) or natural resource management (Williams, 2011).

This section summarizes the solution of MDPs and POMDPs that should help the reader in understanding the principles of the model. The notation utilized here differs partly from conventions found in the classical literature on Markovian decision processes, which is mainly rooted in artificial intelligence applications. The notation utilized in this paper is based on the conventions used in the field of Bayesian networks and decision models; graphical models are used as part of the explanation. We focus on the solution of finite horizon models, where a bounded time interval (e.g. 100 years) is considered. For this purpose, the principle of backward induction can be applied, in which the model is evaluated backwards from the last decision to the initial one. We furthermore restrict the presentation to discrete-state processes; continuous demand parameters are discretized. The computational time needed for solving the decision models depends on the definition of the action and state spaces. It increases with the space of the variables in the core process, with the number of possible actions (possible capacity values), and

with the number of decision time instances. Unfortunately, utilization of existing MDP/POMDP softwares for solution of the proposed models is not straightforward because of time-dependent risk and cost in the proposed model. The algorithms utilized for the proposed model are described in Table 2 and Table 3, they were implemented in Matlab.

3.4.1 Markovian Decision Process (MDP)

The MDP model is depicted in Figure 2a. Its evaluation proceeds backwards starting from the last decision at time t_{N-1} and ending at t_0 . The algorithm of the evaluation is described in the following text and summarized in Table 2.

The optimal capacity at time t_i is found conditionally on flexibility φ , previous capacity v_{i-1} and demand θ_i as

$$v_{i,opt}(\varphi, v_{i-1}, \theta_i) = \underset{v_i}{\operatorname{argmax}} W_i(\varphi, v_i, v_{i-1}, \theta_i). \quad (8)$$

$W_i(\varphi, v_i, v_{i-1}, \theta_i)$ is the expected reward from all time steps following the decision v_i :

$$W_i(\varphi, v_i, v_{i-1}, \theta_i) = \Lambda_i(\varphi, v_i, v_{i-1}, \theta_i) + \Omega_i(\varphi, v_i, \theta_i), \quad (9)$$

where $\Lambda_i(\varphi, v_i, v_{i-1}, \theta_i)$ is the expected immediate reward in time step $i + 1$ and $\Omega_i(\varphi, v_i, \theta_i)$ is the expected sum of future rewards following the next decision v_{i+1} .

The expected immediate reward is calculated as¹

$$\begin{aligned} \Lambda_i(\varphi, v_i, v_{i-1}, \theta_i) &= -d(t_i) \cdot c_i(\varphi, v_i, v_{i-1}) \\ &+ \sum_{\theta_{i+1}} d(t_{i+1}) \cdot [b_{i+1}(v_i, \theta_{i+1}) - r_{i+1}(v_i, \theta_{i+1})] p(\theta_{i+1} | \theta_i, v_i) \end{aligned} \quad (10)$$

where $b_{i+1}(v_i, \theta_{i+1})$ and $r_{i+1}(v_i, \theta_{i+1})$ are the benefits and risk in time step $i + 1$. $c_i(\varphi, v_i, v_{i-1})$ is the adjustment cost associated with the decision at time step i , it is equal to $c_A(\varphi, v_i, v_{i-1})$ of Eq. (3). $p(\theta_{i+1} | \theta_i, v_i)$ is the transition probability characterizing the Markov process as introduced in Sec. 3.1. $d(\cdot)$ is a discount factor which is calculated as

$$d(t_i) = \frac{1}{(1 + \rho)^{t_i}} \quad (11)$$

¹ We use a discrete representation of the random variables (the demand θ_i). However, the framework can be extended to continuous RVs. In such case, the summation over the states of θ_{i+1} is replaced by an integral over the domain of θ_{i+1} and the probability mass functions $p(\theta_{i+1} | \theta_i, v_i)$ and $p(\theta_1)$ are replaced with probability density functions $f(\theta_{i+1} | \theta_i, v_i)$ and $f(\theta_1)$.

where ρ is the annual discount rate and t_i is expressed in years.

The expected sum of future rewards $\Omega_i(\varphi, v_i, \theta_i)$ is

$$\Omega_i(\varphi, v_i, \theta_i) = \sum_{\theta_{i+1}} W_{i+1}(\varphi, v_{i+1,opt}, v_i, \theta_{i+1})p(\theta_{i+1}|\theta_i, v_i) \quad (12)$$

where $W_{i+1}(\varphi, v_{i+1,opt}, v_i, \theta_{i+1})$ is the expected reward associated with the optimal decision at time t_{i+1} , which in the MDP terminology is called the $(N - i)$ -step value function. Its value is known from the previous evaluation step since the evaluation proceeds backwards. For the second-last time step (when the last decision is made), it is $\Omega_{N-1}(\varphi, v_{N-1}, \theta_{N-1}) = 0$, and the optimal capacity v_{N-1} is thus selected based only on the immediate reward.

The optimal initial capacity at time t_0 can be found using Eq. (4). Using the results from the previous step of the backwards induction, the expected reward is calculated as

$$W_0(\varphi, v_0) = \Lambda_0(\varphi, v_0) + \Omega_0(\varphi, v_0), \quad (13)$$

where $\Lambda_0(\varphi, v_0)$ is the expected immediate reward and $\Omega_0(\varphi, v_0)$ is the expected sum of future rewards following the decision at time t_1 . The expected immediate reward equals

$$\Lambda_0(\varphi, v_0) = -c_0(\varphi, v_0) + \sum_{\theta_1} d(t_1) \cdot [b_1(v_0, \theta_1) - r_1(v_0, \theta_1)]p(\theta_1|v_0), \quad (14)$$

$c_0(\varphi, v_0)$ is the initial cost for implementing the system with capacity v_0 and flexibility φ .

The expected sum of future rewards is

$$\Omega_0(\varphi, v_0) = \sum_{\theta_1} W_1(\varphi, v_{1,opt}, v_0, \theta_1)p(\theta_1|v_0). \quad (15)$$

The MDP evaluation procedure is summarized in the algorithm of Table 2.

Table 2: Algorithm for evaluating the MDP.

<p>1. Discount all values of benefits, costs and risk to a selected time point, typically t_0.</p> <p>Loop over all possible values of flexibility φ</p> <p>Perform the evaluation of the sequential decision process:</p> <p>2. Optimize capacity v_{N-1} at the last decision time t_{N-1} conditionally on previous capacity v_{N-2} and actual demand θ_{N-1}, i.e. find $v_{N-1,opt}(\varphi, v_{N-2}, \theta_{N-1})$ following Eqs. (8)-(10), with future rewards $\Omega_{N-1}(\varphi, v_{N-1}, \theta_{N-1})$ in Eq. (9) being equal to 0.</p> <p>3. Store the value function $W_{N-1}(\varphi, v_{N-1,opt}, v_{N-2}, \theta_{N-1})$ to be used in the next step.</p> <p style="padding-left: 40px;">Loop starting with $i = N - 2$ and continuing backwards to $i = 1$:</p> <p style="padding-left: 80px;">a. optimize capacity v_i at time t_i, , conditionally on previous capacity v_{i-1} and actual demand θ_i, i.e. find $v_{i,opt}(\varphi, v_{i-1}, \theta_i)$ following Eqs. (8)-(12).</p> <p style="padding-left: 80px;">b. Store the value function $W_i(\varphi, v_{i,opt}, v_{i-1}, \theta_i)$ to be used in the next step.</p> <p style="padding-left: 40px;">End loop</p> <p>4. Find the initial optimal capacity $v_{0,opt}(\varphi)$ following Eqs. (4) and (13)-(15).</p> <p>5. Store the value function $W_0(\varphi, v_{0,opt})$</p> <p>End loop</p> <p>6. If applicable, find the optimal flexibility φ following Eq. (6).</p>

In many applications, time intervals Δt between planning decisions are large. (The design of a sewerage system is not reassessed every year.) This is included in the analysis by considering possible changes of capacity v_i only at time steps $i = \Delta t, (2 \cdot \Delta t), \dots (m \cdot \Delta t)$. The evaluation proceeds according to Eqs. (8) and (9), wherein v_{i-1} is replaced with $v_{i-\Delta t}$. The immediate expected reward is the sum of benefits, risk and costs until the next decision, and Eq. (10) is replaced with

$$\begin{aligned}
 \Lambda_i(\varphi, v_i, v_{i-\Delta t}, \theta_i) &= -d(t_i) \cdot c_i(\varphi, v_i, v_{i-\Delta t}) \\
 &+ \sum_{j=i+\Delta t}^{j=i+\Delta t} \sum_{\theta_j} d(t_j) \cdot [b_j(v_i, \theta_j) - r_j(v_i, \theta_j)] p(\theta_j | \theta_i, v_i),
 \end{aligned} \tag{16}$$

The calculation of the future expected reward in Eq. (12) is replaced with

$$\Omega_i(\varphi, v_i, \theta_i) = \sum_{\theta_{i+\Delta t}} W_{i+\Delta t}(\varphi, v_{i+\Delta t,opt}, v_i, \theta_{i+\Delta t}) p(\theta_{i+\Delta t} | \theta_i, v_i). \tag{17}$$

The conditional probability of the future demand θ_{i+k} for $k = j$ in Eq. (16) and $k = \Delta t$ in Eq. (17) is calculated as

$$p(\theta_{i+k}|\theta_i, v_i) = \sum_{\theta_{i+k-1}} \dots \left[\sum_{\theta_{i+2}} \left[\sum_{\theta_{i+1}} p(\theta_{i+1}|\theta_i, v_i) \cdot p(\theta_{i+2}|\theta_{i+1}, v_i) \right] \cdot p(\theta_{i+3}|\theta_{i+2}, v_i) \right] \dots \cdot p(\theta_{i+k}|\theta_{i+k-1}, v_i) = \sum_{\theta_{i+1}} \dots \sum_{\theta_{i+k-1}} \prod_{j=i+1}^{i+k} p(\theta_j|\theta_{j-1}, v_i) \quad (18)$$

Finally, the optimal initial capacity is evaluated using Eqs. (4) and (13). Eqs. (14) and (15) are adjusted in the same way as Eqs. (10) and (12), see Eqs. (16) and (17).

3.4.2 Partially Observable Markovian Decision Process (POMDP)

The POMDP model is depicted in Figure 2b. The true demand θ_i is unknown when making a decision, but an observation Z_i is available, which is related to the demand through the conditional probability $p(z_i|\theta_i)$, the likelihood.

Because the true demand is not known, the past is not independent of the future given the present knowledge, and the Markov property of the process $\{\theta\}$ cannot be directly exploited as in MDPs. To overcome this, PMODPs are solved by introducing a so-called belief state β_i , which represents a probability distribution over the uncertain demand:

$$\beta_i(\theta_i) = p(\theta_i) \quad (19)$$

The belief state β_i is what is available to the planner when deciding upon v_i , it represent his knowledge of the demand. It summarizes the past history, or more technically: it is a sufficient statistic of past observations. For a given β_i , the future is independent of past observations. A POMDP can thus be transformed to a MDP, in which the model of the demand is replaced by the model of the belief state on the demand. This procedure does however lead to a significant increase of the state space and hence computational costs.

The process of transforming the POMDP to a MDP with belief states is illustrated in Figure 3 in graphical form. In Figure 3a, nodes β_i representing the belief state are added, the demands θ_i are now defined conditionally on the belief state β_i . Because the belief state is a sufficient statistic, the variables θ_i and θ_{i+1} become independent for given β_i and no arc between them is necessary. The core and observation process, i.e. nodes θ_i and z_i can now be eliminated (removed) from the chart following the procedure described in Shachter (1986), see also Straub and Der Kiureghian (2010) and Špačková (2012). The final influence diagram is shown in Figure 3b. Note that the resulting influence diagram corresponds to a fully observable MDP such as the one of Figure 2a, in which the variables representing the demand are replaced with the variables representing the belief state.

The solution of the MDP in Figure 3b therefore follows the procedure described in Section 3.4.1, Eqs. (8)-(15), where the variables θ_i are replaced with β_i for $i = 1, \dots, N$.

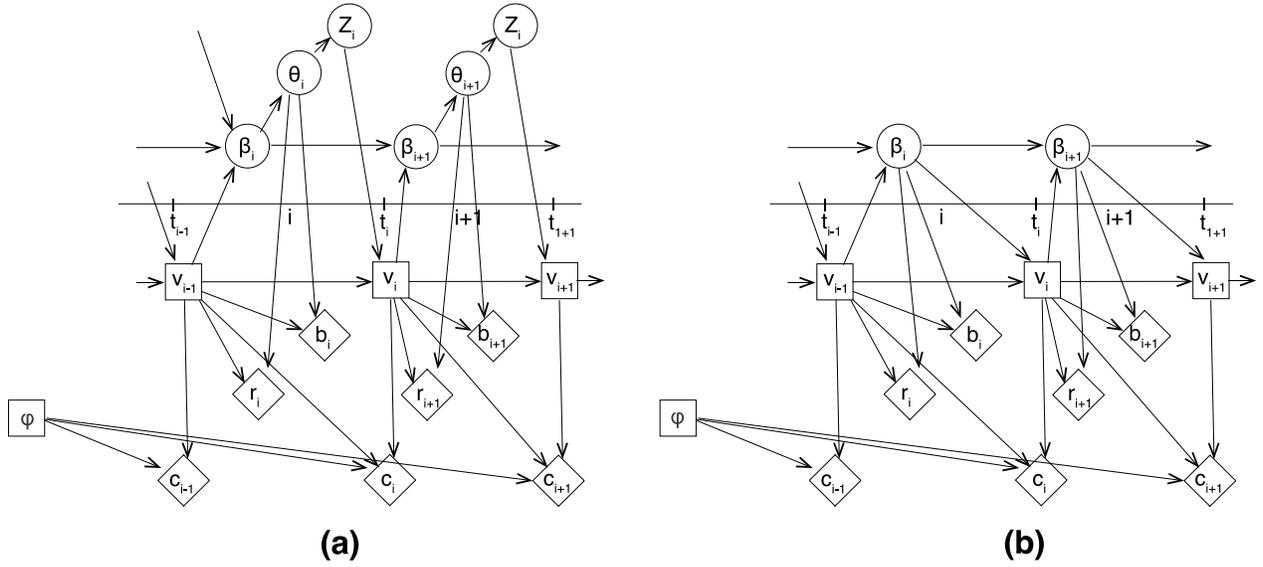


Figure 3: Solving the POMDP: (a) introducing belief states, (b) decision graph after elimination of variables representing the state of nature and observation that can be solved as a fully observable MDP.

To establish the MDP of Figure 3b, the belief state transition function $p(\beta_{i+1}|\beta_i, v_i)$ must be determined by summing $p(z_i|\beta_i)$ over all values of z_i that lead to a specific value of β_{i+1} :

$$\begin{aligned}
 p(\beta_{i+1}|\beta_i, v_i) &= \sum_{z_i} p(\beta_{i+1}|\beta_i, v_i, z_i)p(z_i|\beta_i) \\
 &= \sum_{z_i} I[\beta_{i+1} = p(\theta_{i+1}|\beta_i, v_i, z_i)] \cdot \sum_{\theta_i} p(z_i|\theta_i)\beta_i(\theta_i),
 \end{aligned} \tag{20}$$

where $I[\beta_{i+1} = p(\theta_{i+1}|\beta_i, v_i, z_i)]$ is the indicator function:

$$I[\beta_{i+1} = p(\theta_{i+1}|\beta_i, v_i, z_i)] = \begin{cases} 1, & \beta_{i+1} = p(\theta_{i+1}|\beta_i, v_i, z_i) \\ 0, & \text{otherwise} \end{cases}. \tag{21}$$

The probability distribution of θ_{i+1} conditional on β_i, v_i, z_i is computed by

$$p(\theta_{i+1}|\beta_i, v_i, z_i) \propto \sum_{\theta_i} p(\theta_{i+1}|\theta_i, v_i)p(z_i|\theta_i)\beta_i(\theta_i) \tag{22}$$

Eq. (22) describes a sequential Bayesian updating of the process, since the belief state $\beta_i(\theta_i)$ includes all information available up to z_{i-1} .

A procedure for numerically solving Eq. (20) via Monte Carlo sampling is summarized in Table 3.

Table 3: Algorithm for determining the conditional probability distributions of belief states using MC sampling.

<ol style="list-style-type: none"> 1. Generate $N_{S,1}$ samples of the demand θ_1 from the prior probability distribution $p(\theta_1) = \beta_1$, which represents the “true” demand. 2. For each sample of θ_1, sample an observation Z_1 from the distribution $p(z_1 \theta_1)$. 3. For each sample of Z_1, update the probability distribution of θ_2 with Z_1 following Eq. (22), to determine a corresponding sample of the belief state β_2. 4. Using the samples, estimate the transition probability of the belief state $p(\beta_2 \beta_1, v_1)$. <p>Loop over $i = 2 : N - 1$</p> <p style="padding-left: 2em;">Loop over $j = 1 : N_{S,i-1}$ (the samples from the previous cycle)</p> <ol style="list-style-type: none"> 5. Sample $N_{S,i} \leq N_{S,i-1}$ samples of the demand θ_i using the j-th sample of belief state β_i. 6. For each sample of θ_i, sample the observation Z_i from the distribution $p(z_i \theta_i)$. 7. Update the probability distribution of θ_{i+1} with the sampled observation Z_i following Eq. (22), to determine a sample of the belief state β_{i+1} conditional on the j-th sample of the belief state β_i. 8. Using $N_{S,i}$ samples, estimate the transition probability of belief state $p(\beta_{i+1} \beta_i, v_i)$. <p>Endloop</p> <ol style="list-style-type: none"> 9. Estimate the transition probability of belief state $p(\beta_{i+1} \beta_i, v_i)$ from the $N_{S,i} \times N_{S,i-1}$ samples of β_{i+1} conditional on β_i. <p>Endloop</p>
--

To formulate the reduced MDP of Figure 3b, it further necessary to express the benefit and risk from Eqs. (10) and (14) as a function of the belief state:

$$b_{i+1}(v_i, \beta_{i+1}) = \sum_{\theta_{i+1}} b_{i+1}(v_i, \theta_{i+1}) \beta_{i+1}(\theta_{i+1}) \quad (23)$$

$$r_{i+1}(v_i, \beta_{i+1}) = \sum_{\theta_{i+1}} r_{i+1}(v_i, \theta_{i+1}) \beta_{i+1}(\theta_{i+1}) \quad (24)$$

If decisions are made only in intervals of length Δt , i.e. at times $t = \Delta t, (2 \cdot \Delta t), \dots (m \cdot \Delta t)$, the procedure described in Eqs. (16)-(18) is used, whereby the demand θ_i is replaced with the belief state β_i . Eq. (20) for computing the transition probability of belief states is replaced with

$$p(\beta_{i+\Delta t}|\beta_i, v_i) = \sum_{z_i} \dots \sum_{z_{i+\Delta t-1}} I[\beta_{i+\Delta t} = p(\theta_{i+\Delta t}|\beta_i, v_i, z_i, \dots, z_{i+\Delta t-1})] \cdot p(z_i, \dots, z_{i+\Delta t-1}|\beta_i), \quad (25)$$

where $I[\beta_{i+\Delta t} = p(\theta_{i+\Delta t}|\beta_i, v_i, z_i, \dots, z_{i+\Delta t-1})]$ is the indicator function:

$$\begin{aligned}
I [\beta_{i+\Delta t} = p(\theta_{i+\Delta t} | \beta_i, v_i, z_i, \dots, z_{i+\Delta t-1})] \\
= \begin{cases} 1 & \text{if } \beta_{i+\Delta t} = p(\theta_{i+\Delta t} | \beta_i, v_i, z_i, \dots, z_{i+\Delta t-1}) \\ 0 & \text{otherwise} \end{cases}
\end{aligned} \tag{26}$$

and

$$\begin{aligned}
p(z_i, \dots, z_{i+\Delta t-1} | \beta_i) \propto \sum_{\theta_i} p(z_i | \theta_i) \beta_i(\theta_i) \sum_{\theta_{i+1}} p(z_{i+1} | \theta_{i+1}) p(\theta_{i+1} | \theta_i) \dots \\
\sum_{\theta_{i+\Delta t-1}} p(z_{i+\Delta t-1} | \theta_{i+\Delta t-1}) p(\theta_{i+\Delta t-1} | \theta_{i+\Delta t-2})
\end{aligned} \tag{27}$$

Eqs. (25) and (27) are evaluated using sampling methods, in analogy to the algorithm of Table 3.

4 Application to the optimization of a waste water treatment plant design

In a Waste Water Treatment Plant (WWTP) design, the goal is to select the optimal capacity to which the WWTP should be built, taking into account the uncertainty in future demand, associated with population, economic and technological developments. Both the demand and the capacity of the WWTP are expressed in the Population Equivalent (PE) units (Gillot et al., 1999; Benedetti et al., 2006).

It is assumed that decisions on possible extension of the WWTP will be carried out every 15 years; the overall planning horizon is 90 years. The annual discount rate is 0.02. At the time of the future decisions, the actual demand (in PE) at that time will be known, the process is thus fully observable.

The model is summarized in Figure 4. The uncertain demand in [PE] is modeled as a homogeneous Markov process, where the conditional probability distribution of the demand θ_i in year i , $f(\theta_i | \theta_{i-1})$, is normal distributed with time dependent mean $\mu_i = \theta_{i-1} + k$ and standard deviation σ , which is assumed to be constant over time. The initial demand is $\theta_0 = 40 \times 10^3$ [PE], the trend representing the mean annual increase of the demand is $k = 0.5 \times 10^3$ [PE] and the standard deviation describing random fluctuations is $\sigma = 2 \times 10^3$ [PE]. At the end of the planning horizon, the expected demand is 85×10^3 [PE] and its standard deviation is close to 20×10^3 [PE]. This model of population growth is hypothetical, but it is straightforward to replace it with another model for a specific area if such is available. Principles of forecasting the demographic changes can be found for example in Keyfitz and Caswell (2005).

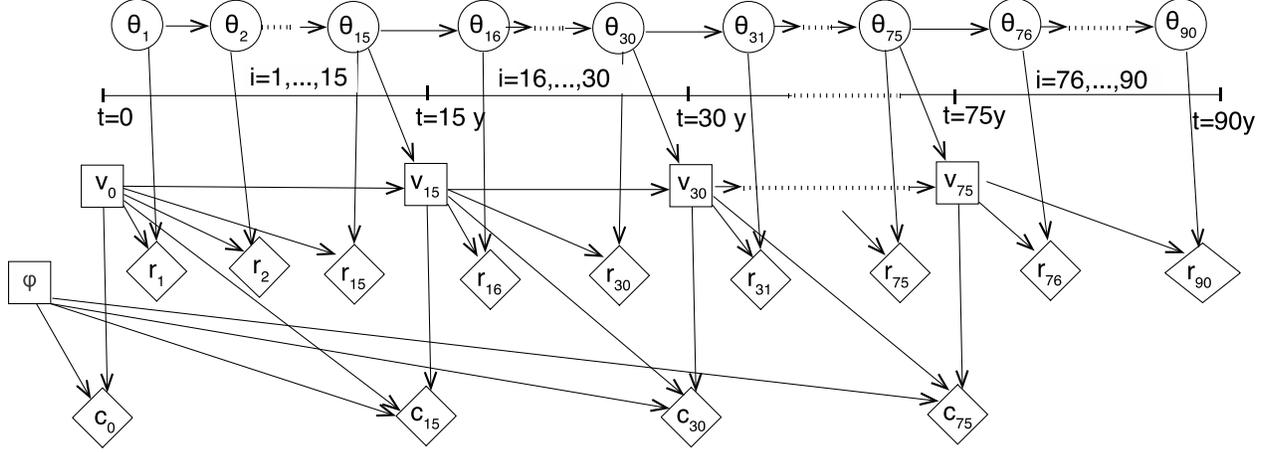


Figure 4: Influence diagram of the WWTP design problem. θ_i is the demand in year i , v_j is decision on capacity at time t_j , c_j is the cost associated with this decision, φ is flexibility and r_i is the risk (i.e. the expected costs of not having the waste water treated)

The cost of building to capacity v_0 initially is

$$c_0(\varphi, v_0) = \gamma(\varphi) \cdot \xi \cdot \sqrt{\frac{v_0}{1000}} [\times 10^6 \text{ Euro}] \quad (28)$$

where $\xi = 0.7$ is a scaling constant and $\gamma(\varphi)$ is a coefficient representing the price of the flexibility. For zero flexibility $\gamma(0) = 1$, for partial flexibility $\gamma(0.5) = 1.2$ and for full flexibility $\gamma(1) = 1.5$. Similarly shaped cost functions were found by (Gillot et al., 1999).

The adjustment cost of a future extension from v_{i-1} to v_i is

$$c_A(\varphi, v_i, v_{i-1}) = \gamma(\varphi) \cdot b \cdot \left[\sqrt{\frac{v_i}{1000}} - \varphi \sqrt{\frac{v_{i-1}}{1000}} \right] [\times 10^6 \text{ Euro}] \quad (29)$$

The cost functions are illustrated in Figure 5.

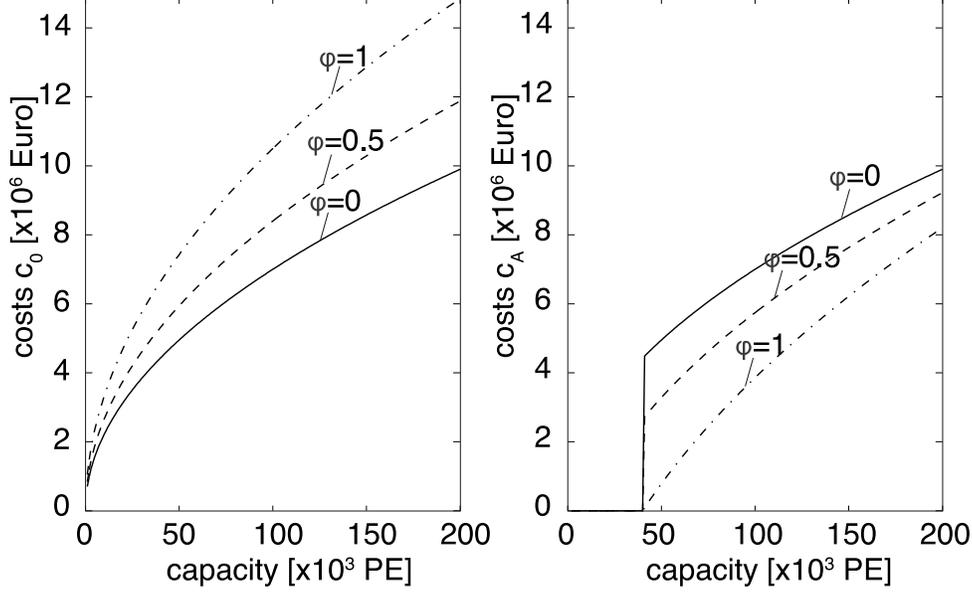


Figure 5: Investment cost for WWTP with different flexibilities φ as a function of the capacity: (a) initial costs $c_0(\varphi, v_0)$ at time $t_0 = 0$, (b) adjustment costs $c_A(\varphi, v_{adj}, v_{or})$ at future time step when the original capacity is $v_{or} = 40 \times 10^3$ PE.

Insufficient capacity of the WWTP leads to costs that correspond to fines for exceeding the effluent standards and/or to the costs of assuring alternative treatment of the waste water. The expected value of these costs (in our model denoted as risk) in year i is a function of the actual capacity v_j of the WWTP and of the demand θ_i :

$$r_i = 0 \text{ for } \theta_i \leq v_j$$

$$r_i = (\theta_i - v_j) * \beta = -\Delta_i * \beta [\times 10^6 \text{ Euro}], \text{ for } \theta_i > v_j, \quad (30)$$

where $\beta = 50$ is a scaling constant, Δ_i is the reserve $\Delta_i = v_j - \theta_i$. This simple model assumes that the risk is zero when the demand is smaller than the actual capacity and linearly increasing otherwise.

4.1 Results and discussion

The model in Figure 4 is evaluated following the procedure described in Sec. 3.4.1. In this example, the benefits are constant with respect to v and φ , hence only cost and risk are taken into account.

Figure 6 shows the main result of the decision model, i.e. the optimal initial capacity of the WWTP, which minimizes the total expected life-time reward (objective function) that is plotted with the solid line, for systems with different flexibilities. For the inflexible system (Figure 6a), the optimal initial capacity equals 69×10^3 PE, for the fully flexible system (Figure 6c) the optimal initial capacity is significantly lower at 56×10^3 PE. As expected, a more conservative design with larger reserve is recommendable for an inflexible system, because future adjustments of the system are more expansive. However, even for the inflexible system, the optimal initial

capacity is lower than the mean demand of 85×10^3 PE expected at the end of the life-time. The reason for such an under-design being optimal is the effect of discounting that favors delaying investments, in combination with the uncertainty that allows for the possibility that the mean demand at the end of service life will be lower. Figure 6 also shows the contribution of the immediate risk and costs computed with Eq. (10), which are associated with the first decision period of 15 years, and of the sum of future expected rewards computed with Eq. (12). It can be seen that due to the discounting, the effect of the future reward is relatively small compared to the immediate reward.

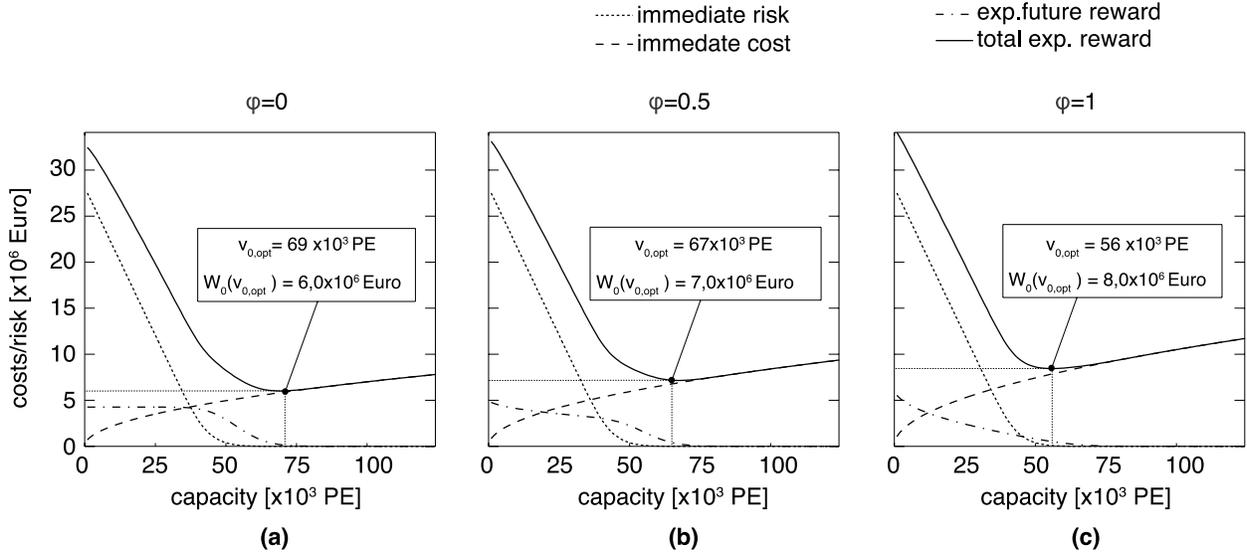


Figure 6: Optimization of the initial capacity of WWTP for different flexibilities.

For illustration purposes, Figure 7 shows three randomly selected scenarios of the demand over the lifetime of the WWTP and the corresponding optimal decisions on capacity for WWTPs with different flexibilities. The optimal initial capacity at time $t = 0$ corresponds to the one found in Figure 6. Every 15 years it is decided whether the capacity is kept or increased, depending on actual demand. To further illustrate the decision process, Figure 8 shows the conditional optimization of capacities at time $t = 75$ yr for the demand scenario from Figure 7a. Because this is the last decision before the end of service life, it is based entirely only the immediate reward, which represents the risk and costs for years 75 to 90.

The scenario in Figure 7a represents a future where the demand growth approximately follows the predicted mean growth (shown by the dotted line). For this scenario, the inflexible system with $\varphi = 0$ (dashed line with circles) is never adjusted, even if the actual demand in the second half of the planning period reaches or even exceeds the capacity of the WWTP. The exceedance of the capacity is not sufficiently large to justify the high cost associated with the adjustment of the system capacity, as can be seen from Figure 8a: the risk is not large enough to outweigh the immediate costs for increasing the capacity from 69×10^3 PE to any higher value. It can, however, also be observed that the objective function (i.e. solid line in the figure) has two local minima and the second local minimum at capacity 82×10^3 PE is not significantly higher than the global minimum at capacity 69×10^3 PE. This indicates that the optimal decision is sensitive to the model assumptions. In contrast, the flexible system with $\varphi = 1$ is adjusted at almost every decision point under the scenario of Figure 7a. The optimization of the last decision

is shown in Figure 8c. It can be observed that the objective function for the flexible system has only one global minimum. The reason is the fact that there is no jump in the adjustment cost function, as can also be observed from Figure 5b. Therefore if the observed demand is larger than the one predicted at the last decision time, it is always better to adjust the system, because the adjustment costs are significantly lower than in the case of non- or partly flexible systems.

In the scenario of Figure 7b, the demand first decreases and then slightly grows in the second half of the planning time. In such a scenario, the partly and inflexible systems (dashed line with circles and dashed line with crosses) should never be adjusted as the initially selected capacity is sufficient over the whole life time. The fully flexible system (dashed line with stars) should only be adjusted once at time $t = 60\text{yr}$, when the demand exceeds the initial demand at time $t = 0$.

In the scenario of Figure 7c, the demand increases faster than the mean prediction, especially in the second half of the lifetime. In such a scenario, the partly and non-flexible systems (dashed line with circles and dashed line with crosses) are adjusted at time $t = 60\text{yr}$ after the observed demand exceeds the WWTP capacity. The fully flexible system (dashed line with stars) is adjusted every 15 years and it thus keeps an unused capacity (reserve) during the whole lifetime.

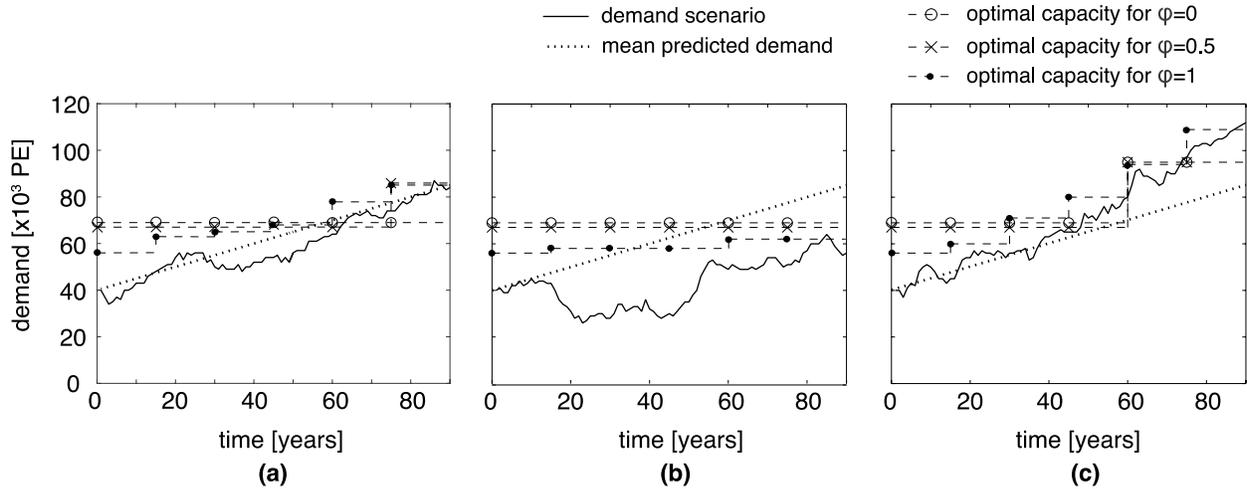


Figure 7: Alternative scenarios of future demand and the resulting optimal decisions on capacity of WWTP for different flexibilities: (a) scenario where demand oscillates around the predicted mean growth, (b) scenario where demand grows significantly slower than expected, (c) scenario where demand grows faster than expected.

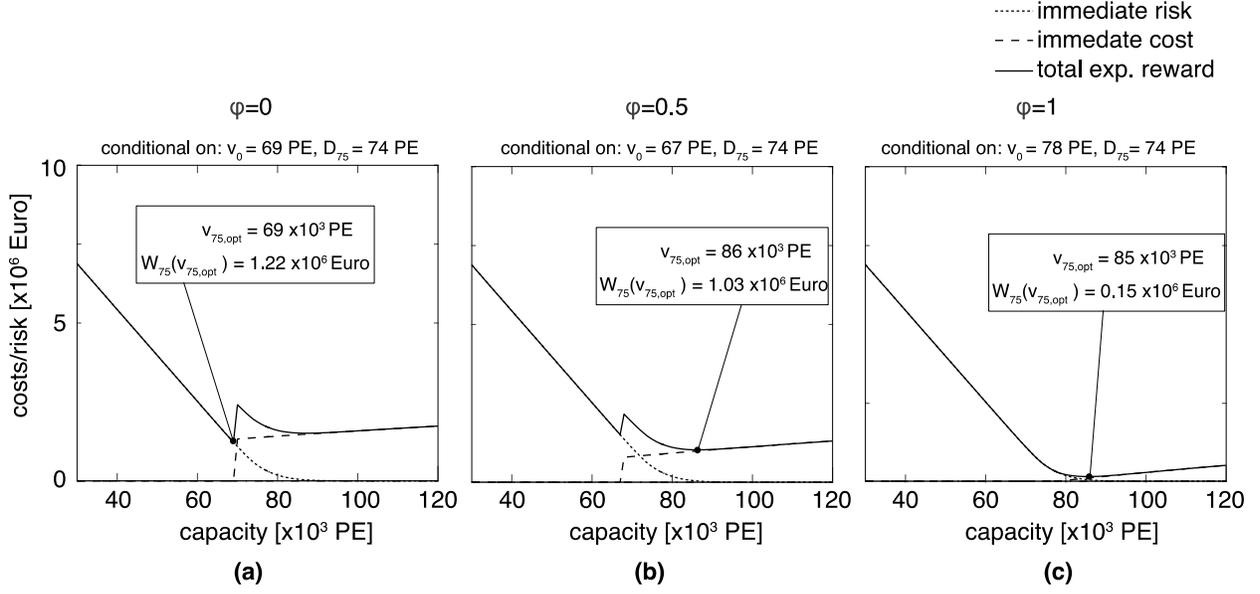


Figure 8: Optimization of future decisions at time $t = 75$ years for systems with different flexibility, conditionally on original capacity v_0 and demand D_{75} observed at the time of the decision, which corresponds to the scenario shown in Figure 7a.

4.1.1 Value of flexibility for WWTP

The value of flexibility is determined following Section 3.2. The optimization is repeated with free flexibility, i.e. the coefficients $\gamma(\varphi)$ representing the price of flexibility in Eqs. (28) and (29) are set equal to 1 for all values of φ . The corresponding optimal initial capacity for varying flexibilities is shown in Table 4. The value of flexibility is then obtained as the difference between the expected cost associated with the optimal strategy for $\varphi = 0$ and that associated with larger values of φ .

Table 4: Resulting optimal capacity of WWTP for different flexibilities

	$\varphi = 0$	$\varphi = 0.5$	$\varphi = 1$
Optimal capacity $v_{o,opt}$	69×10^3 PE	68×10^3 PE	57×10^3 PE
Expected reward $W_0(\varphi, v_{o,opt})$	6.0×10^6 Euro	6.0×10^6 Euro	5.7×10^6 Euro

The relative value of flexibility with respect to the initial investment costs is shown in Figure 9. The value of the fully flexible system is approximately 6% of the initial costs of building to the optimal protection level. Overall, the value of flexibility increases non-linearly with flexibility, and a large degree of flexibility is necessary to lead to noticeable benefits.

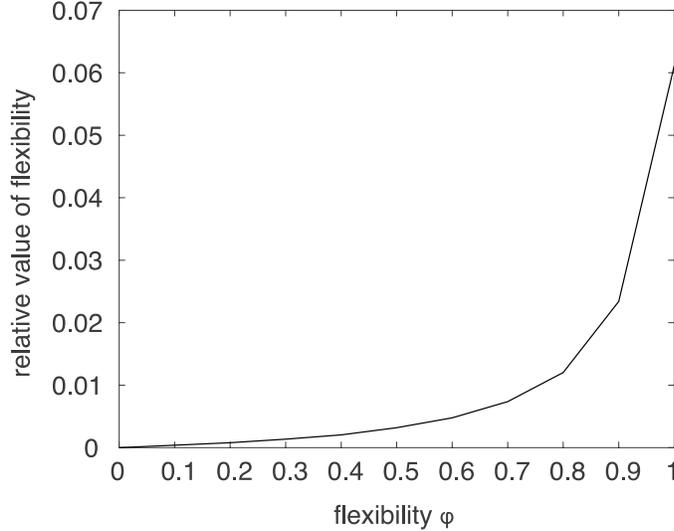


Figure 9: Relative value of flexibility of the WWTP with respect to the initial costs.

5 Application to risk-based optimization of flood protection measures

In many regions, climate change is now considered when designing flood protection systems. For example, the state of Bavaria requires a 15% increase in the design flood capacity to account for effects of climate change (LfU Bayern, 2005). Considering the significant associated uncertainties, such decisions should ideally take into account the flexibility of the protection strategies.

Here, we consider the implementation of a new river flood defense at time $t_0 = 0$, whose capacity may be adjusted in the future at times $t = 30$ and $t = 60$ years, the total planning horizon is 90 years. The annual discount rate is 0.02. The assumptions made in the example are based on realistic case studies previously analyzed at our institute (Ellinger, 2015; Perosa, 2016). However, direct transfer of the results to other catchments is not straightforward because the risk, cost as well as the impact of climate are site specific.

An influence diagram for the decision problem is shown in Figure 4. The decision problem is modeled by a POMDP. The uncertainty in the future demand is modeled using two random variables: climate scenario S_i representing the uncertainty in future climate and annual maximum discharge D_i . The probability distribution of the annual maximum discharge is defined conditional on the climate scenario.

Scenario S_i is in this example modeled as a discrete random variable. Only three scenarios are considered: (A) no change in the extreme flood discharge distribution, (B) moderate increase of extreme discharge, where the mean annual maximum discharge as well as the magnitude of a 100-year event increase by 16% in 100 years and (C) significant increase of extreme discharge, where the mean annual maximum discharge as well as the magnitude of a 100-year event increase by 32% in 100 years. This simple model is selected for illustrative purposes, as it facilitates an understanding of model principles and the interpretation of the results. For real applications, an ensemble of climate scenarios can be used or the parameters of the discharge distribution may be modeled as continuous random process variables. In the model, the random

variable scenario S_i has three states and its transition probability matrix is the 3×3 unit matrix. All scenarios are a-priori considered to be equally probable, i.e. $\Pr(S_1 = A) = \Pr(S_1 = B) = \Pr(S_1 = C) = 1/3$.

The annual maximum discharge D_i is described by a Gumbel distribution with mean $\mu(i)$ and standard deviation $\sigma(i)$ conditional on the climate scenario as shown in Table 1. Both $\mu(i)$ and $\sigma(i)$ are assumed to be linearly increasing with time; the c.o.v. is 0.8 at all times. We consider that the observed discharges can be used to update the probability of the different scenarios, as indicated by the links from the D_i to the observation variables Z_i . We neglect measurement errors, and the observed discharge Z_i equals the true discharge D_i . Nevertheless, because one only observes the annual discharge maxima and indirectly learns on the climate scenario, the Markov process is a partially observable one. The climate scenario S_i is not directly observed here.

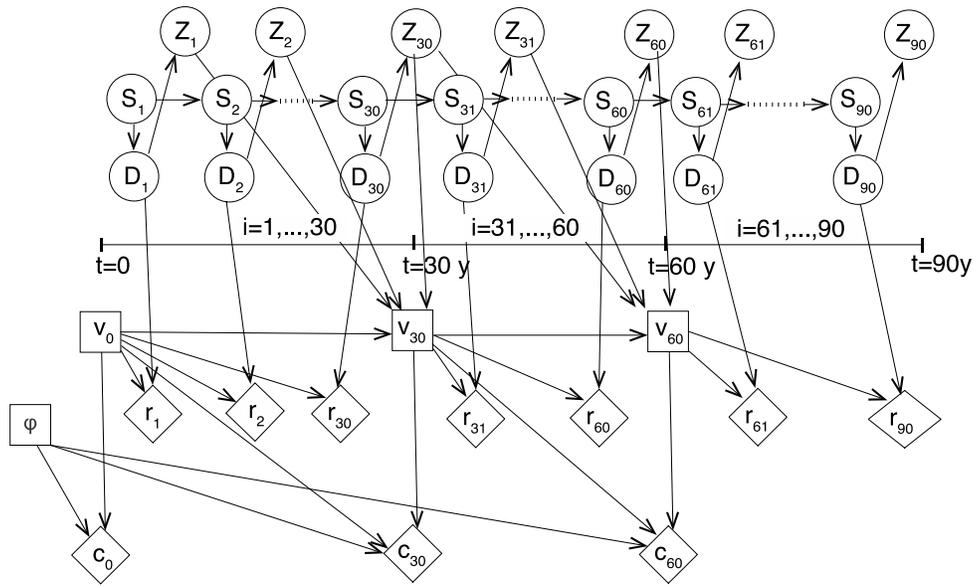


Figure 10: Influence diagram for the example decision problem for optimization of the flood protection system.

Table 5: Mean and standard deviation of the annual maximum discharge D_i [m^3/s] for different climate scenarios

Scenario S	Mean $\mu(i)$	St.dev. $\sigma(i)$
(A) no change	1200	960
(B) moderate incr.	$1200+2*i$	$960+1.6*i$
(C) significant incr.	$1200+4*i$	$960+3.2*i$

The belief state is the probability distribution over the possible demand states. Here, the demand is represented by the scenario S_i , which then determines the distribution of extreme discharges. Since scenario S_i is a discrete random variable with three states in this example, its full probability distribution $p(s_i)$ is determined by 2 parameters $p_A = \Pr(S_i = A)$ and $p_B = \Pr(S_i = B)$, which form the belief state.

Risk (expected damage) in the i th time step for given scenario S_i and for given capacity v_i equals

$$r(S_i, v_i) = \int e \cdot \kappa^{v_i}(d) \cdot f_{D_i|S_i}(d) dd \quad (31)$$

where $e = 9 \times 10^8$ Euro is the exposure, $f_{D_i|S_i}(d)$ is the conditional Gumbel PDF of annual maximum rainfall (demand) for a given scenario and $\kappa^{v_i}(d)$ is the vulnerability (relative damage) for given capacity of the protection system v_i :

$$\kappa^{v_i}(d) = \begin{cases} 0, & d \leq \max[d_0, v_i] \\ \frac{1}{d_{max} - d_0} \cdot (d - d_0), & \max[d_0, v_i] < d < d_{max} \\ 1, & d \geq d_{max} \end{cases} \quad (32)$$

The damage function (the product of vulnerability and exposure $e \cdot \kappa^{v_i}(d)$) is shown in Figure 11. The vulnerability is zero for discharges smaller than the capacity of the flood defense system, neglecting the possibility that the flood defense fails at discharges below the design discharge. Without a flood defense, the first damages occur at discharge $d_0 = 3200$ m³/s, which corresponds to a 26-year event under current climate conditions. The maximum vulnerability of 1 is reached at discharge $d_{max} = 6000$ m³/s, which corresponds to a 1100-year event under current climate conditions. For higher discharges it is assumed that all exposed objects are flooded and the maximum damage of 9×10^8 Euro is incurred. Between d_0/v_i and d_{max} , the vulnerability increases linearly.

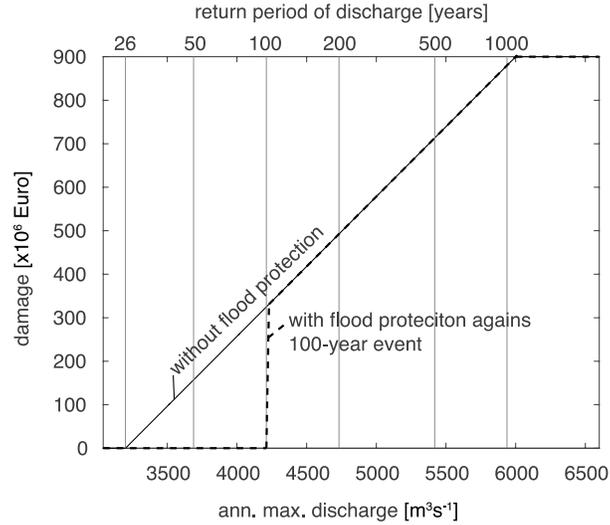


Figure 11: Idealized damage functions for the case without any protection measures (solid line) and for case where the area is fully protected against a 100-year event under current climate conditions (dashed line). Return periods correspond to current climate conditions (scenario A).

Three different levels of flexibility are considered: full flexibility $\varphi = 1$, intermediate flexibility $\varphi = 0.5$ and zero flexibility $\varphi = 0$. The initial and adjustment costs for the different capacities grow linearly with the capacity as shown in Figure 12. They are defined as

$$c_0(\varphi, v_0) = \gamma(\varphi) \cdot b \cdot (v_0 - d_0), \quad (33)$$

$$c_A(\varphi, v_i, v_{i-1}) = \gamma(\varphi) \cdot b \cdot [(v_i - d_0) - \varphi \cdot (v_{i-1} - d_0)], \quad (34)$$

where $b = 1.3 \cdot 10^5$ Euro is a scaling constant and $\gamma(\varphi)$ is a coefficient representing the price of the flexibility: $\gamma(0) = 1$, $\gamma(0.5) = 1.05$ and $\gamma(1) = 1.1$.

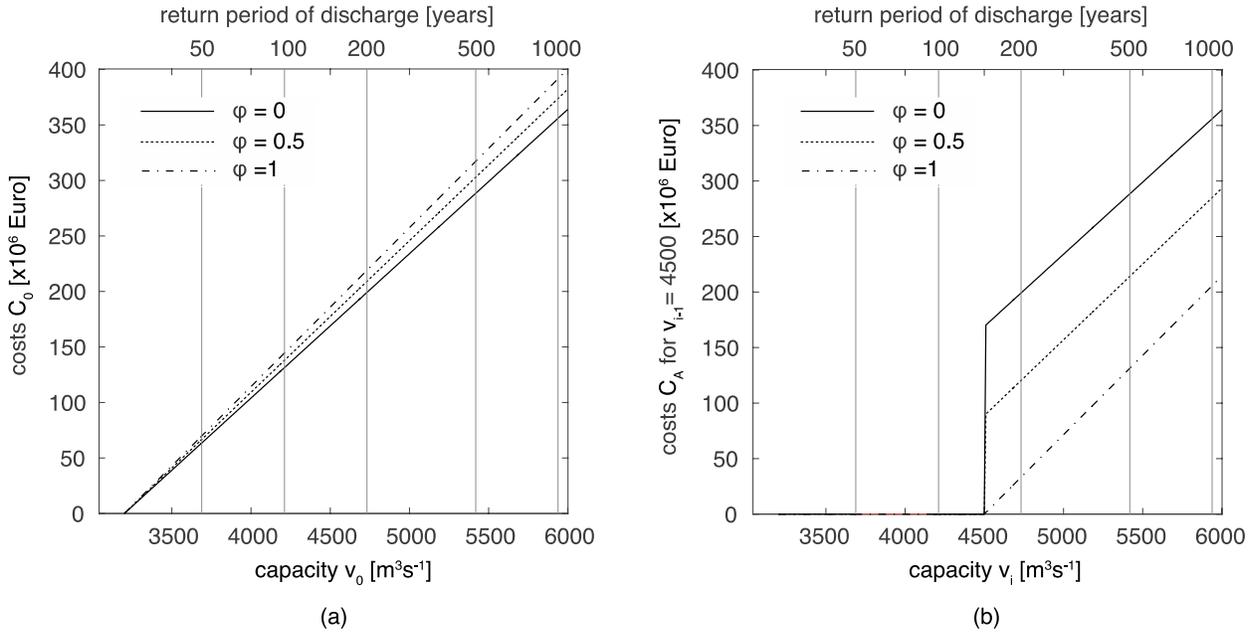


Figure 12: Costs of flood protection systems with different flexibilities φ as a function of the capacity (design discharge in m^3/s) of the system: (a) initial costs $c_0(\varphi, v_0)$ at time $t_0 = 0$, (b) adjustment costs $c_A(\varphi, v_i, v_{i-1})$ assuming that the original capacity selected at the previous time step is $4500 m^3/s$. Return periods correspond to current climate conditions (scenario A).

5.1 Classical optimization without consideration of future changes in demand

The optimal capacity of the system is first determined using the traditional static approach, where future uncertainties are not accounted for. It is assumed that there is no trend in the mean annual maximum discharge. The optimum is determined using Eq. (1). Since the benefits are disregarded in this example, Eq. (1) can be reformulated to a minimization of the sum of discounted risk and cost over the whole planning horizon. The results of the optimization are shown in Figure 13. The optimal capacity equals $4800 m^3/s$, which corresponds to a 219-year event.

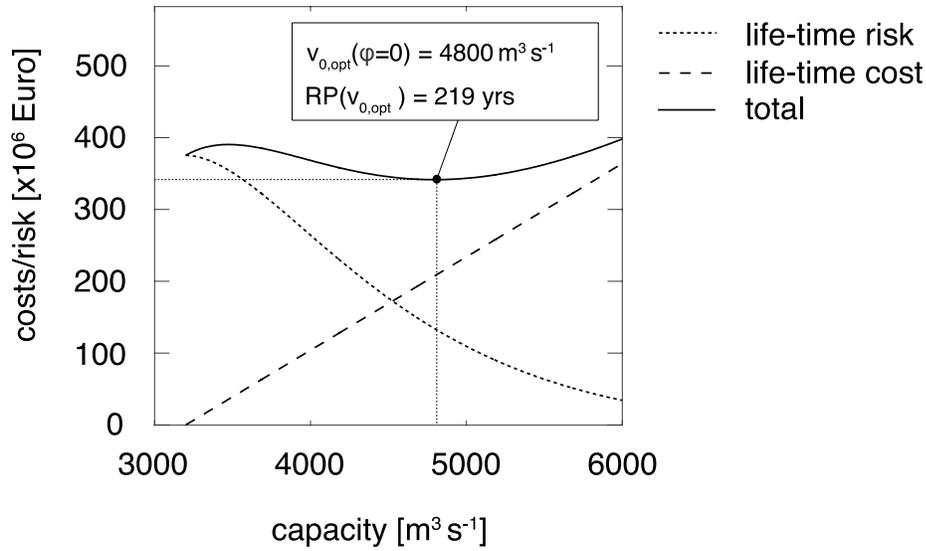


Figure 13: Static optimization of the capacity of the flood protection system, without accounting for system flexibility and assuming that there is no change in the mean annual maximum discharge (corresponding to scenario A).

5.2 Full decision model

The model in Figure 10 is evaluated using the procedure described in Sec. 3.4.2. Figure 14 shows boxplots of updated probabilities of scenarios after 30 years of rainfall observations, conditionally on the true scenario. These results show that the data is only mildly informative. For example, it can be seen from Figure 14a that if scenario A is the underlying true one, the posterior probability of A based on the rainfall observations is 0.35, for B it is 0.33 and for C it is 0.31. The spread of posterior probabilities is, however, significant, which means that the extreme rainfall observations gathered during the first 30 year can lead to erroneous conclusions. The reason for this is the fact that the differences among the scenarios are not pronounced in the first 30 years, and that trends in extremes are generally difficult to observe (Tebaldi et al., 2006).

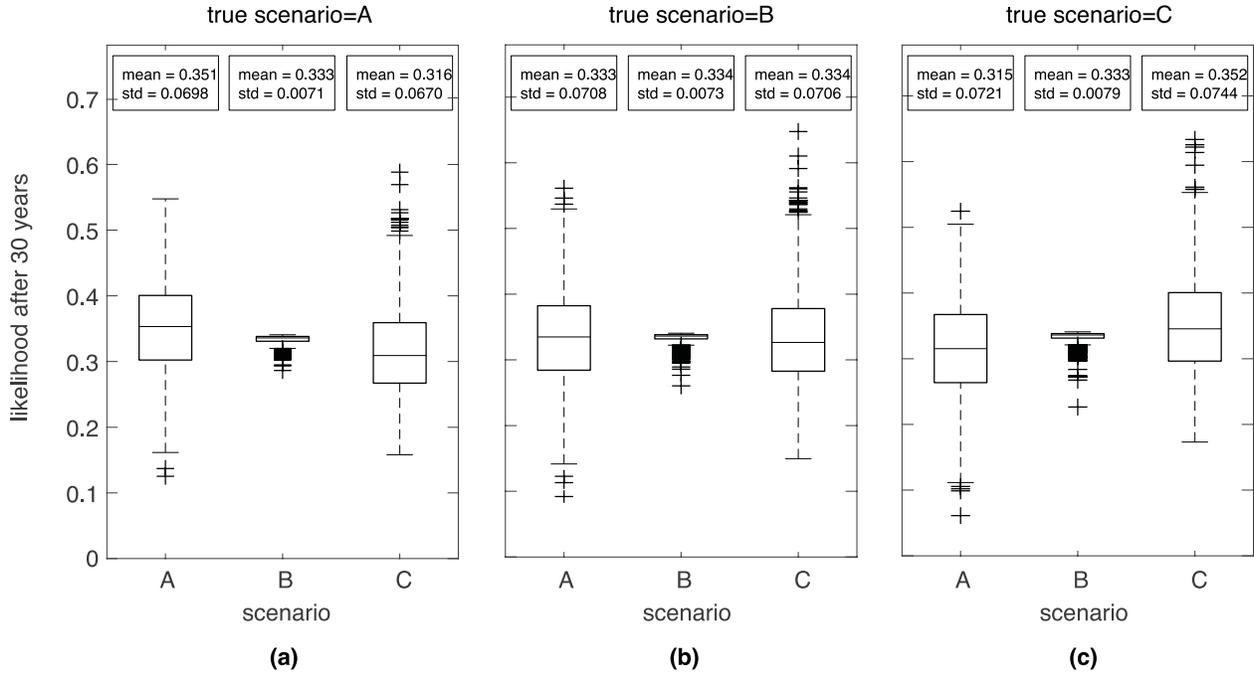


Figure 14: Boxplot of posterior probabilities of climate scenarios $\Pr(S|Z_1, \dots, Z_{30})$ updated with 30 years of annual maximum rainfall observations, conditional on a true scenario.

Figure 15 shows the probability distribution of the belief state at the decision time $t = 30$ years. A belief state here corresponds to a combination of $\Pr(S = A)$ and $\Pr(S = B)$. In agreement with the results in Figure 14, it can be observed that the spread in $\Pr(S = B)$ is significantly lower than in $\Pr(S = A)$.

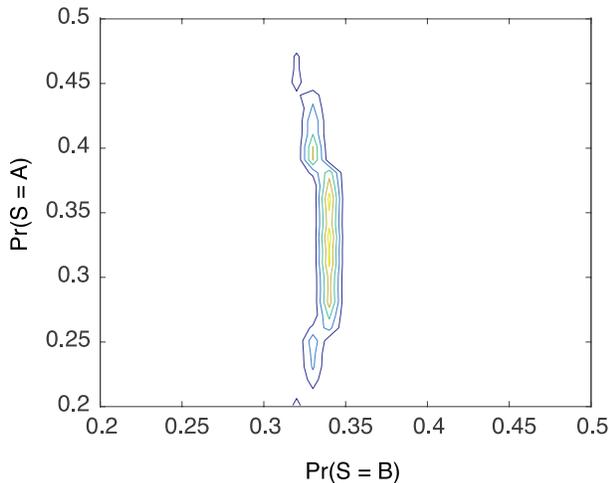


Figure 15: Probability distribution over the belief states after 30 years of annual maximum rainfall observations, $p(\beta_{30})$.

Figure 16 shows the conditional probability of belief states at time $t = 60$ yr (when the second future decision is to be made), given that the belief state at time $t = 30$ yr (when the first future

decision is to be made) is $\Pr(S = A|Z_1, \dots, Z_{30}) = \Pr(S = B|Z_1, \dots, Z_{30}) = 0.34$. A similar pattern can be observed as in Figure 15, but the variability in the belief states is higher. It is not unlikely that scenario A has a posterior probability close to zero or one after collecting data for 60 years.

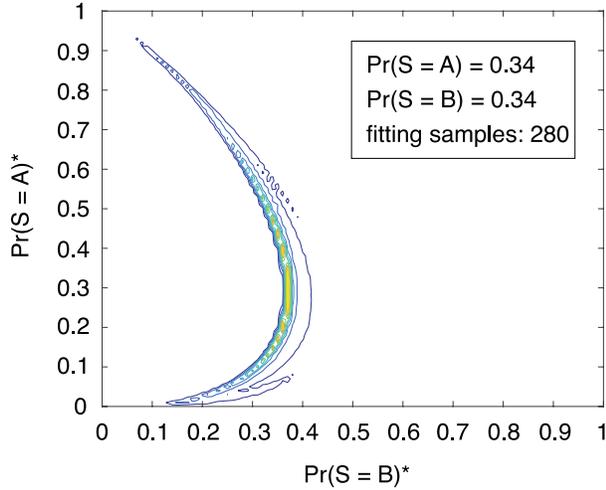


Figure 16: Conditional probability distribution over the belief states after 60 years of annual maximum rainfall observations, $p(\beta_{60}|\beta_{30})$, for a belief state β_{30} at 30yr corresponding to $\Pr(S = A|Z_1, \dots, Z_{30}) = \Pr(S = B|Z_1, \dots, Z_{30}) = 0.34$.

Figure 17 shows the results of the optimization of the initial capacity v_0 for different system flexibilities. For an inflexible flood protection system with $\varphi = 0$, the optimal capacity is $v_{0,opt} = 5240 \text{ m}^3/\text{s}$, corresponding to a 400-year rainfall event under current climate. For a fully flexible system with $\varphi = 1$, a slightly lower capacity $v_{0,opt} = 5090 \text{ m}^3/\text{s}$ is optimal (a 320-year rainfall event in current climate), because the capacity can be adjusted in the future without incurring extra costs. A comparison with the static optimization results in Table 6 shows that considering future climate change leads to higher optimal design capacities. In the case of an inflexible system, the increase is in the order of 10%, for a fully flexible system the increase is reduced to about 6%.

Table 6: Resulting optimal designs

	Static optimization	Sequential optimization		
	$\varphi = 0$	$\varphi = 0$	$\varphi = 0.5$	$\varphi = 1$
Design discharge	4800m ³ /s	5240m ³ /s	5180m ³ /s	5090m ³ /s
Design return period	220yr	400yr	360yr	320yr

A comparison of the expected rewards W_0 associated with the three values of flexibility and the corresponding optimal design shows that the optimum is achieved for $\varphi = 0$. The additional cost of flexibility is higher than its benefit, and it is more economical to select inflexible systems together with a conservative design (higher initial capacity) of the flood protection system.

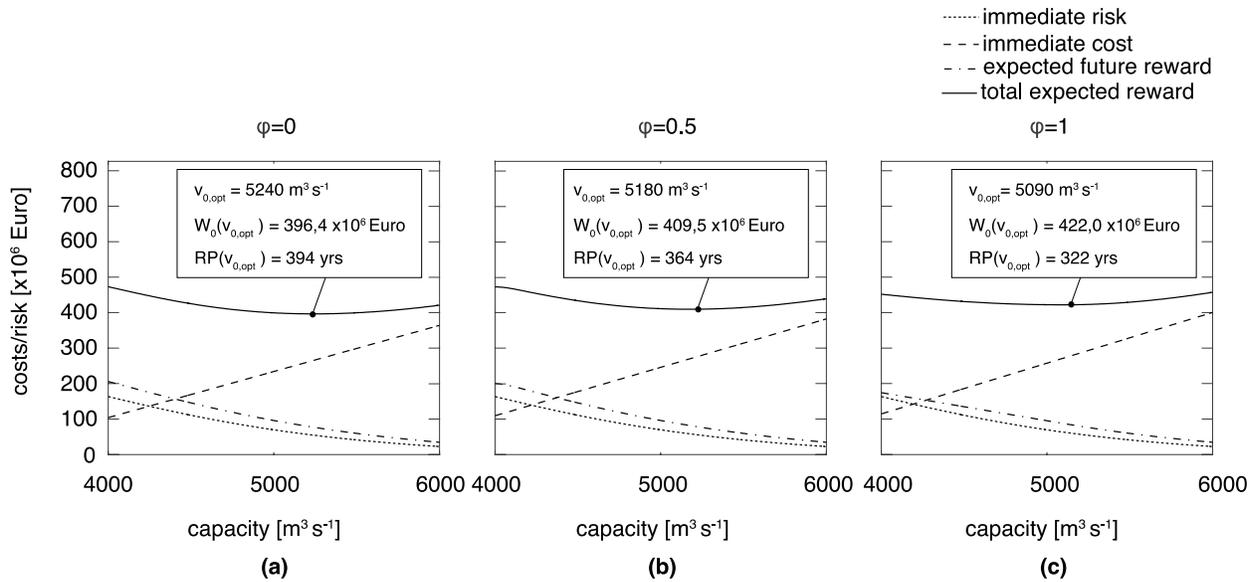


Figure 17: Results of the optimization of initial capacity for different flexibilities. $v_{0,opt}$ is the optimal initial capacity in m^3/s and p_{exc} is the exceedance probability of the capacity based on the current PDF of maximum annual rainfall.

To further investigate the characteristics of the decision problem, Figure 18 shows two randomly selected realizations of future annual maximum discharge observations and the related optimal system capacities. The optimal initial capacity at time $t = 0$ corresponds to the capacity found in Figure 17. Figure 19 shows the conditional optimization of future capacities at time $t = 30$ years for the demand realization from Figure 18a.

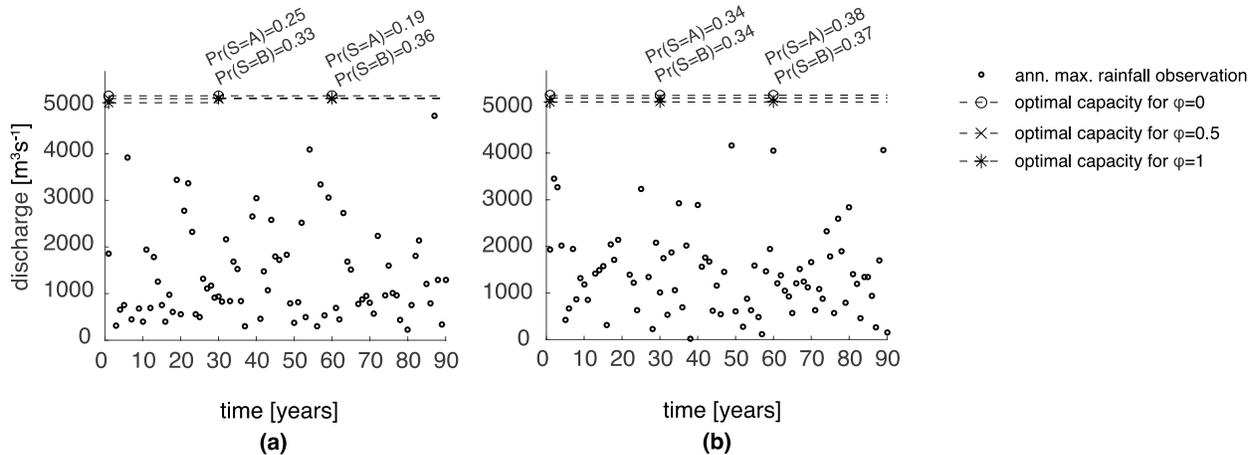


Figure 18: Alternative realizations of future discharge observations and the resulting optimal decisions on capacity of the flood defense system for different flexibilities: (a) realization where scenario C, i.e. strong increase, becomes most likely, (b) realization where all scenarios remain approximately equally likely.

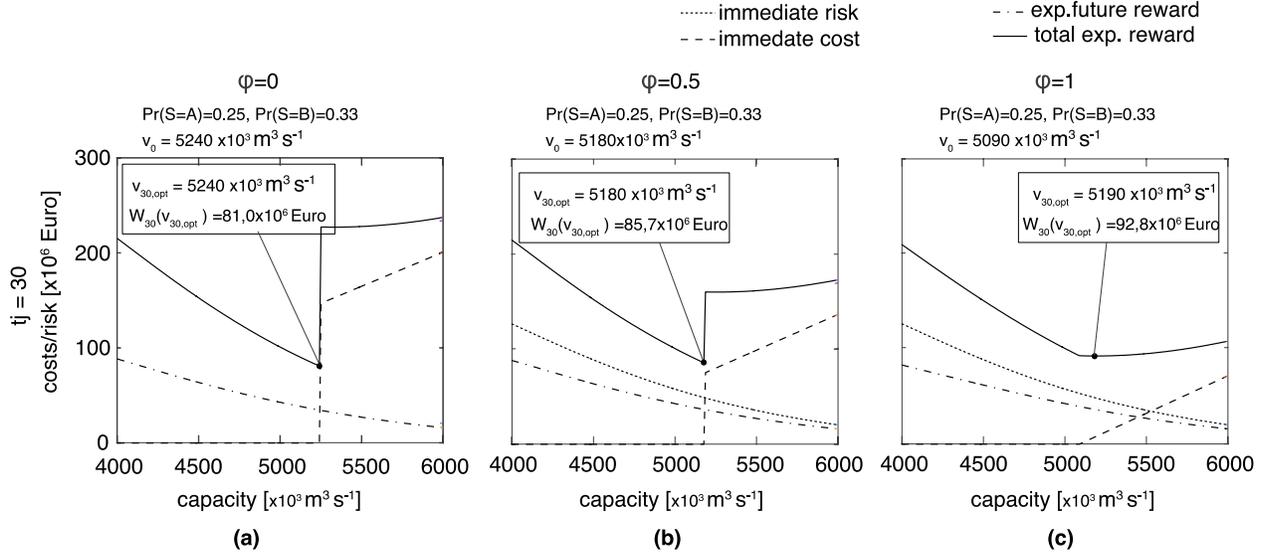


Figure 19: Optimization of future decisions at time $t = 30$ conditionally on original capacity v_0 and observed demand corresponding to the realization shown in Figure 18a.

In both realizations of Figure 18, the capacities of the inflexible and partly flexible system are kept constant. The costs of adjusting the capacity of these systems is so high that it does not outweigh the possible increase in the predicted risk due to a higher frequency of flood events. This can be observed in Figure 19, which shows the sudden jump in the cost function.

The capacity of the fully flexible system is changed at time t_{30} in Figure 18a, when the observations indicate an increase in the likelihood of scenario C. Figure 19c shows the corresponding optimization, where the objective function is minimized for capacity $v_{30,opt} = 5190 \text{ m}^3/\text{s}$. For the realization in Figure 18b, the fully flexible system is not changed throughout the lifetime.

5.3 Value of flexibility for flood risk mitigation measures

The value of flexibility is determined following the procedure described in Section 3.2. The optimization is repeated with free flexibility, i.e. the coefficients $\gamma(\varphi)$ representing the price of flexibility in Eqs. (28) and (29) are set equal to 1 for all values of φ . The corresponding optimal initial capacity for varying flexibilities is shown in Table 7.

Table 7: Optimal designs assuming that the flexibility is for free.

	Static optimization		Sequential optimization	
	$\varphi = 0$		$\varphi = 0$	$\varphi = 1$
Design discharge	4800m ³ /s		5240m ³ /s	5220m ³ /s
Design return period	220yr		400yr	380yr

The value of flexibility is presented in Figure 20. The value of partial flexibility ($0.1 \leq \varphi \leq 0.9$) is equal to zero, i.e. it is not worth to implement partly flexible systems if they are more costly than the inflexible ones. The fully flexible system is beneficial as compared to the non-flexible but the benefit is less than 0.1% of the initial costs.

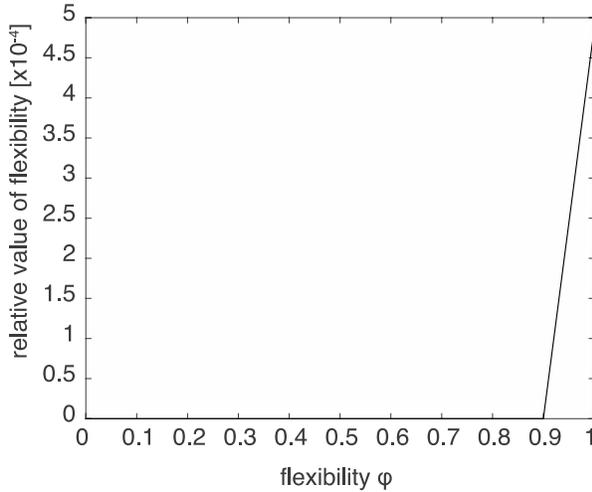


Figure 20: Relative value of flexibility of the flood protection system.

6 Discussion

We present a quantitative decision model, which can improve the design and planning of long-life systems that face future uncertainties. The two numerical examples of Sections 4 and 5 demonstrate the applicability of the proposed decision models to the optimization of the capacities of a Waste Water Treatment Plant (WWTP) and a flood protection system. The results show that the optimal design is influenced by the flexibility of the system, i.e. its ability to be modified in the future. By neglecting the flexibility, the planning can lead to suboptimal decisions. In both investigated examples, the associated objective functions are quite flat and sometimes have multiple local minima. As a result, the optimum is rather robust with respect to model assumptions. This facilitates a generalization of the results obtained from the case studies.

In both examples, the value of flexibility is found to be limited. In the WWTP example, it is up to 6% of the initial design costs, in the flood protection example, it is less than 0.1% of the initial design costs. The latter value is arguably an underestimation, because the amount of information used for learning the likelihood of climate scenarios was very limited. In practice, it can be expected that one will learn significantly more through improved climate models and additional data. Nevertheless, it can be concluded that the value of flexibility is limited because of (a) the effect of discounting and (b) the fact that increasing capacity is a no-regret strategy, in particular in the second example. In both examples, the real effect of the flexibility would come at later times in the service life, but because of discounting, the present value of the benefit associated with late changes is small. Increasing capacity is a no-regret strategy, because of the inherent uncertainty associated with extreme events. Adding additional capacity will always lead to a further reduction of risk, therefore the net cost of an overdesign is limited. However, this effect

may not be present in other applications. For example, an overdesign of a transportation system, of which parts remain unused later, has no benefits and is thus not a no-regret strategy.

A comparison of the results from the WWTP example, a fully observable Markov Decision Process (MDP), with the flood risk example, a Partially Observable Markov Decision Process (POMDP), provides insights of the effect of the learning process on the optimal decisions. The more informative future observations are, the higher is the value of flexibility and the more significant adjustments are likely to be made to the system. In the WWTP example, the demand parameter is fully observable and future observations thus significantly reduce uncertainty. As a consequence, the capacity of the WWTP is likely to be changed several times during its lifetime and the value of flexibility is higher. On the contrary, the flood risk example represents a case where the future observations bring only a very limited knowledge on the true state of nature (the parameters of the extreme discharge distribution) and the learning effect is very low. As a consequence, no significant changes of the system are likely in the future and the value of flexibility is low.

For the purpose of illustration, some of the assumptions made in the examples (e.g. the low number of scenarios utilized in the modeling of the climate change uncertainty in the flood risk example) are simplifying reality. For real applications, the uncertainty in the distribution of the extreme events (due to climate or due to limited historic records) should ideally be modeled with a continuous model of the parameters of probabilistic distribution as shown for example in (Dittes et al., 2016).

In this paper, the capacity of the system was modeled as a continuous variable. In practical applications, determining the costs and benefits (risks) associated with each possible capacity is often infeasible. For example, in flood mitigation planning, a large variety of different measures, and combinations thereof, can be implemented (e.g., dykes combined with retention areas, mobile flood barriers, warning systems) and each measure has one or more parameters to be optimized (height of a dyke, volume of the retention, type of the mobile barriers, etc.). In such cases it is only realistic to evaluate a limited number of combinations of these measures. The continuous optimization presented in this paper would thus be replaced with a discrete optimization, similarly to the approach presented in (Špačková and Straub, 2015). The discrete formulation would reduce computational cost of the proposed decision model.

7 Conclusions

We propose the use of (Partially Observable) Markov Decision processes for planning of long-term adaptation strategies for civil infrastructure and structures. The approach is applicable to a wide range of systems that have relevant uncertainty in the demand on the system, both at present and in the future. The approach allows identifying the optimal system design taking into account the flexibility of the system, i.e. its ability to be changed in the future without excessive costs. For this purpose, we define a quantitative measure of flexibility, based on costs for the initial construction and adjustments. The approach can also help identifying the optimal timing for changing the system as well as the value of flexibility, i.e. the additional costs that a decision maker should be willing to pay for increasing system flexibility. It explicitly includes the fact that also future decisions will be made under uncertainty, even if reduced through additional

information. The use of an influence diagrams for graphical representation of the decision process helps communicating the model assumptions.

The approach is illustrated with two examples considering an optimization of the design of a Waste Water Treatment Plant (WWTP) under uncertain future population growth and the planning a of a flood protection system under climate change uncertainty. In both examples, the value of system flexibility is shown to be limited, because of the fact that changes are likely to be made in the farther future and the associated net present value is small because of discounting. Additionally, increasing the system capacity is a no-regret strategy in both applications, hence an overdesign of the systems still leads to an additional reduction in system risk, even if it is sub-optimal. These conditions are, however, not found in all application cases. At present, we therefore have only limited knowledge on the value of flexibility in different adaptation planning problems. Future work should include a systematic study of the value of flexibility under different system characteristics and under different types of uncertainty. The models proposed in this paper can be used for such investigations to derive generalized planning recommendations.

References

- Alagoz, O., Hsu, H., Schaefer, A.J., Roberts, M.S., 2009. Markov Decision Processes: A Tool for Sequential Decision Making under Uncertainty. *Med. Decis. Making*. doi:10.1177/0272989X09353194
- Allahaim, F., Alfaris, A., 2010. Towards Changeability. The Adaptable Building Design (ABD) Framework. Presented at the 5th International Conference of the Arab Society for Computer Aided Architectural Design (ASCAAD 2010), Morocco, pp. 189–200.
- Basupi, I., Kapelan, Z., 2014. Evaluating flexibility in water distribution system design under future demand uncertainty. *J. Infrastruct. Syst.* 21, 04014034.
- Benedetti, L., Bixio, D., Vanrolleghem, P.A., 2006. Assessment of WWTP design and upgrade options: balancing costs and risks of standards' exceedance. *Water Sci. Technol.* 54, 371–378. doi:10.2166/wst.2006.571
- Boehland, J., 2003. Future-proofing your building: Designing for flexibility and adaptive reuse. *Environ. Build. News* 28, 310–314.
- Braziunas, D., 2003. Pomdp solution methods. University of Toronto, Toronto.
- Brechtel, S., Gindele, T., Dillmann, R., 2014. Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs, in: 17th International IEEE Conference on Intelligent Transportation Systems (ITSC). pp. 392–399. doi:10.1109/ITSC.2014.6957722
- Buurman, J., Zhang, S., Babovic, V., 2009. Reducing Risk Through Real Options in Systems Design: The Case of Architecting a Maritime Domain Protection System. *Risk Anal.* 29, 366–379. doi:10.1111/j.1539-6924.2008.01160.x
- Corotis, R.B., Hugh Ellis, J., Jiang, M., 2005. Modeling of risk-based inspection, maintenance and life-cycle cost with partially observable Markov decision processes. *Struct. Infrastruct. Eng.* 1, 75–84. doi:10.1080/15732470412331289305
- Cox, L.A., 2012. Confronting Deep Uncertainties in Risk Analysis. *Risk Anal.* 32, 1607–1629. doi:10.1111/j.1539-6924.2012.01792.x
- Dittes, B., Špačková, O., Straub, D., 2016. Managing uncertainty in flood predictions: Flexible protection strategies vs. safety factors. *Water Resour. Res.* under review.
- Ellinger, M., 2015. Categorization and modelling of flood protection measures for risk-based optimization (MSc. Thesis). Technische Universität München, Munich, Germany.
- Ellis, H., Jiang, M., Corotis, R.B., 1995. Inspection, Maintenance, and Repair with Partial Observability. *J. Infrastruct. Syst.* 1, 92–99. doi:10.1061/(ASCE)1076-0342(1995)1:2(92)
- Fankhauser, S., Smith, J.B., Tol, R.S.J., 1999. Weathering climate change: some simple rules to guide adaptation decisions. *Ecol. Econ.* 30, 67–78. doi:10.1016/S0921-8009(98)00117-7
- Fawcett, W., Urquijo, I.R., Krieg, H., Hughes, M., Mikalsen, L., Gutiérrez, Ó.R.R., 2014. Cost and Environmental Evaluation of Flexible Strategies for a Highway Construction Project under Traffic Growth Uncertainty. *J. Infrastruct. Syst.*

- Fricke, E., Schulz, A.P., 2005. Design for changeability (DfC): Principles to enable changes in systems throughout their entire lifecycle. *Syst. Eng.* 8, no–no. doi:10.1002/sys.20039
- Gersonius, B., Ashley, R., Pathirana, A., Zevenbergen, C., 2012. Climate change uncertainty: building flexibility into water and flood risk infrastructure. *Clim. Change* 116, 411–423. doi:10.1007/s10584-012-0494-5
- Gillot, S., De Clercq, B., Defour, D., Simoens, F., Gernaey, K., Vanrolleghem, P., 1999. Optimization of wastewater treatment plant design and operation using simulation and cost analysis, in: *Proceedings of 72nd Annual WEF Conference and Exposition*. New Orleans, USA, pp. 9–13.
- Hall, J.W., Otto, A., Tran, M., Barr, S., Alderson, D., 2014. A National Model for Strategic Planning of Infrastructure Systems. *American Society of Civil Engineers*, pp. 2821–2829. doi:10.1061/9780784413609.284
- Hallegatte, S., 2009. Strategies to adapt to an uncertain climate change. *Glob. Environ. Change, Traditional Peoples and Climate Change* 19, 240–247. doi:10.1016/j.gloenvcha.2008.12.003
- Hallegatte, S., Shah, A., Brown, C., Lempert, R., Gill, S., 2012. Investment Decision Making Under Deep Uncertainty -- Application to Climate Change (SSRN Scholarly Paper No. ID 2143067). *Social Science Research Network*, Rochester, NY.
- Hoey, J., Poupart, P., Bertoldi, A. von, Craig, T., Boutilier, C., Mihailidis, A., 2010. Automated handwashing assistance for persons with dementia using video and a partially observable Markov decision process. *Comput. Vis. Image Underst., Special issue on Intelligent Vision Systems* 114, 503–519. doi:10.1016/j.cviu.2009.06.008
- Jensen, F.V., Nielsen, T.D., 2007. *Bayesian networks and decision graphs*. Springer.
- Kaelbling, L.P., Littman, M.L., Cassandra, A.R., 1998. Planning and acting in partially observable stochastic domains. *Artif. Intell.* 101, 99–134. doi:10.1016/S0004-3702(98)00023-X
- Keyfitz, N., Caswell, H., 2005. *Applied mathematical demography*. Springer.
- Kochenderfer, M.J., Amato, C., Chowdhary, G., How, J.P., Reynolds, H.J.D., Thornton, J.R., Torres-Carrasquillo, P.A., Üre, N.K., Vian, J., 2015. *Decision Making Under Uncertainty: Theory and Application*. MIT Press.
- Kumar, V., 1987. Entropic measures of manufacturing flexibility. *Int. J. Prod. Res.* 25, 957–966. doi:10.1080/00207548708919888
- Lamb, M., Collis, R., Deix, S., Krieger, B., Hautiere, N., IFSTTAR, F., 2011. The forever open road. Defining the next generation road, in: *AIPCR World Congress*, Mexico.
- LfU Bayern, 2005. Sonderbericht: Klimaänderungsfaktoren bei Planungen für den Hochwasserschutz, Bayerisches Landesamt für Umwelt, in: *Gewässerkundlicher Jahresbericht Für Bayern*. München.
- Luque, J., Straub, D., 2013. Algorithms for optimal risk-based planning of inspections using influence diagrams, in: *Proc. 11th International Probabilistic Workshop IPW11*.

- Marques, J., Cunha, M., Savić, D., 2014. Using real options in the optimal design of water distribution networks. *J. Water Resour. Plan. Manag.*
- Martínez Ceseña, E.A., Mutale, J., Rivas-Dávalos, F., 2013. Real options theory applied to electricity generation projects: A review. *Renew. Sustain. Energy Rev.* 19, 573–581. doi:10.1016/j.rser.2012.11.059
- Martins, J., Marques, R.C., Cruz, C.O., 2015. Real Options in Infrastructure: Revisiting the Literature. *J. Infrastruct. Syst.* 21, 04014026. doi:10.1061/(ASCE)IS.1943-555X.0000188
- Memarzadeh, M., Pozzi, M., 2015. Integrated Inspection Scheduling and Maintenance Planning for Infrastructure Systems. *Comput.-Aided Civ. Infrastruct. Eng.* n/a–n/a. doi:10.1111/mice.12178
- Monahan, G.E., 1982. State of the Art—A Survey of Partially Observable Markov Decision Processes: Theory, Models, and Algorithms. *Manag. Sci.* 28, 1–16. doi:10.1287/mnsc.28.1.1
- Neufville, R. de, Scholtes, S., Wang, T., 2006. Real Options by Spreadsheet: Parking Garage Case Example. *J. Infrastruct. Syst.* 12, 107–111. doi:10.1061/(ASCE)1076-0342(2006)12:2(107)
- Nishijima, K., 2015. Concept of decision graphical framework for optimising adaptation of civil infrastructure to a changing climate. *Struct. Infrastruct. Eng.* 0, 1–7. doi:10.1080/15732479.2015.1020496
- Perosa, F., 2016. Cost functions of river flood protection measures (Technical report). TU München, Munich, Germany.
- Poupart, P., 2005. Exploiting structure to efficiently solve large scale partially observable Markov decision processes (PhD thesis). University of Toronto.
- Pozzi, M., Memarzadeh, M., Klima, K., 2016. Hidden-model processes for adaptive management under uncertain climate change. *J. Infrastruct. Syst.* under review.
- Rachelson, E., Quesnel, G., Garcia, F., Fabiani, P., 2008. A Simulation-based Approach for Solving Generalized Semi-Markov Decision Processes, in: *Proc. of the 2008 Conference on ECAI 2008: 18th European Conference on Artificial Intelligence*. IOS Press, Amsterdam, The Netherlands, The Netherlands, pp. 583–587.
- Radhakrishnan, M., Quang, C.N.X., Pathirana, A., Phi, H.L., Ashley, R., 2014. Evaluation Of Retrofitting Options In Urban Drainage Systems Based On Flexibility: A Case Study For Nhieu Loc - Thi Nghe Basin In Ho Chi Minh City. *Int. Conf. Hydroinformatics*.
- Robelin, C.-A., Madanat, S.M., 2007. History-Dependent Bridge Deck Maintenance and Replacement Optimization with Markov Decision Processes. *J. Infrastruct. Syst.* 13, 195–201. doi:10.1061/(ASCE)1076-0342(2007)13:3(195)
- Ross, A.M., Rhodes, D.H., Hastings, D.E., 2008. Defining changeability: Reconciling flexibility, adaptability, scalability, modifiability, and robustness for maintaining system lifecycle value. *Syst. Eng.* 11, 246–262. doi:10.1002/sys.20098

- Saleh, J.H., Mark, G., Jordan, N.C., 2009. Flexibility: a multi-disciplinary literature review and a research agenda for designing flexible engineering systems. *J. Eng. Des.* 20, 307–323. doi:10.1080/09544820701870813
- Shachter, R.D., 1986. Evaluating Influence Diagrams. *Oper. Res.* 34, 871–882. doi:10.1287/opre.34.6.871
- Sönmez, Ö.E., Koç, V.T., 2015. On Quantifying Manufacturing Flexibility: An Entropy Based Approach, in: *Proceedings of the World Congress on Engineering*.
- Špačková, O., 2012. Risk management of tunnel construction projects: Modelling uncertainty of construction time (cost) estimates for risk assessment and decision-making. (Doctoral thesis, Czech Technical University, Prague). Czech Technical University, Prague.
- Špačková, O., Dittes, B., Straub, D., 2015. Risk-based optimization of adaptable protection measures against natural hazards. Presented at the 12th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP12, Vancouver, Canada.
- Špačková, O., Straub, D., 2015. Cost-Benefit Analysis for Optimization of Risk Protection Under Budget Constraints. *Risk Anal.* 35, 941–959. doi:10.1111/risa.12310
- Straub, D., Der Kiureghian, A., 2010. Bayesian Network Enhanced with Structural Reliability Methods: Methodology. *J. Eng. Mech.* 136, 1248–1258. doi:10.1061/(ASCE)EM.1943-7889.0000173
- Tebaldi, C., Hayhoe, K., Arblaster, J.M., Meehl, G.A., 2006. Going to the Extremes. *Clim. Change* 79, 185–211. doi:10.1007/s10584-006-9051-4
- Vrijling, J.K., Kanning, W., Kok, M., Jonkman, S.N., 2009. Designing robust coastal structures, in: *Proc. of the 5th Coastal Structures International Conference*. Presented at the 5th Coastal Structures International Conference, Venice, Italy.
- Williams, B.K., 2011. Resolving structural uncertainty in natural resources management using POMDP approaches. *Ecol. Model.* 222, 1092–1102. doi:10.1016/j.ecolmodel.2010.12.015
- Woodward, M., Kapelan, Z., Gouldby, B., 2014. Adaptive Flood Risk Management Under Climate Change Uncertainty Using Real Options and Optimization. *Risk Anal.* 34, 75–92. doi:10.1111/risa.12088
- Yzer, J.R., Walker, W.E., Marchau, V., Kwakkel, J.H., 2014. Dynamic adaptive policies: A way to improve the cost-benefit performance of megaprojects. *Environ. Plan. B Plan. Des.* 41.