Safe Dike Heights At Minimal Costs: The Nonhomogeneous Case

Ruud Brekelmans
Department of Econometrics & Operations Research,
CentER, Tilburg University, the Netherlands

Dick den Hertog
Department of Econometrics & Operations Research,
CentER, Tilburg University, the Netherlands

Kees Roos
Delft University of Technology, the Netherlands

Carel Eijgenraam
CPB Netherlands Bureau for Economic Policy Analysis

Working paper, 2012

Abstract

Dike height optimization is of major importance to the Netherlands as a large part of the country lies below sea level and high water levels in rivers can cause floods. A cost-benefit analysis is discussed in Eijgenraam et al. (2010), which is an improvement of the model Van Dantzig (1956) introduced after a devastating flood in the Netherlands in 1953. We consider the extension of this model to nonhomogeneous dike rings, which may also be applicable to other deltas in the world. A nonhomogeneous dike ring consists of different segments with different characteristics with respect to flooding and investment costs. The individual segments can be heightened independently at different moments in time and by different amounts, making the problem considerably more complex than the homogeneous case. We show how the problem can be modeled as a MINLP problem and present an iterative algorithm that can be used to solve the problem. Moreover, we consider a robust optimization approach to deal with uncertainty in the model parameters. The method has been implemented and integrated in software, which is used by the government to determine how the safety standards in the Dutch Water Act should be changed.

Keywords: flood prevention, MINLP, cost-benefit analysis, robust optimization

1 Introduction

In the Netherlands, dike rings, consisting of dunes, dikes and structures, protect a large part of the country against flooding. After the serious flood in 1953, a cost-benefit model was developed by Van Dantzig (1956) to determine optimal dike heights. In Eijgenraam et al. (2010) we improve and extend Van Dantzig's model. In that paper we show how to properly include economic growth in the cost-benefit model, and how to address the question when to invest in dikes. All these models consider dike rings that consist of a homogeneous dike. This means that all parts in the dike ring have the same characteristics with respect to investment costs, flood probabilities, water level rise, etc. The objective of the cost-benefit analysis is to find an optimal balance between investment costs and the benefit of reducing flood damages, both as a result of heightening dikes. The question then becomes when and how much to invest in the homogeneous dike ring.

Most dike rings in the Netherlands, however, are nonhomogeneous, consisting of different segments that each have different characteristics. Differences occur, for instance, if a dike ring protects against more than one river, each with different characteristics, or if a dike ring contains a sluice. Currently, there are dike rings with up to ten segments in the Netherlands. In this nonhomogeneous case, it is not necessary and not desirable to enforce that all these segments are heightened simultaneously and by exactly the same amount. Hence, the decision problem for the nonhomogeneous case concerns when and how much to invest in each individual dike segment.

In the current paper, which is based on Eijgenraam (2007b), Den Hertog and Roos (2009) and Brekelmans et al. (2009), we consider the extension of the homogeneous case in Eijgenraam et al. (2010) to the nonhomogeneous case. The research has been carried out as part of a project initiated by the government. The project's main goal is to support decision-making with respect to setting new safety standards for the dike rings in the Netherlands. These safety standards can be derived from an optimal investment strategy and the resulting flood probabilities. How this can be done is explained in Eijgenraam (2006) and Eijgenraam (2007a). Here we confine ourselves to a description of the first stage: finding the optimal investment strategy. In order to lay a firm base for the new standards, all 53 dike rings in the Netherlands need to be analyzed thoroughly. This requires that particular scenarios can be analyzed within a reasonable amount of time, where each scenario represents a certain instance of the model parameters such as economic growth, interest rate, water level rise, flood characteristics, investment costs and so on.

It is shown in Eijgenraam et al. (2010) that the homogeneous case can be solved explicitly for a specific choice of the investment cost function. For other investment cost functions the homogeneous case can be conveniently solved using a dynamic programming method. Unfortunately, this is not possible for the nonhomogeneous case, since the state space explodes if a dike ring consists of multiple dike segments. We show how the nonhomogeneous dike height optimization problem can be modeled as a Mixed Integer Nonlinear Programming (MINLP) problem. The model building blocks have many useful convexity properties and, depending on the choice of investment cost function, the problem can be written as a convex MINLP problem. To take into account the uncertainties with respect to many of the model parameters, we also investigate how robust solutions can be obtained. Due to the complexity of the model it is not possible to consider the robust optimization method of Ben-Tal et al. (2009). Therefore, we consider the so-called regret criterion in combination with a finite set of scenarios.

In addition to the MINLP formulation of the decision problem, we construct an iterative optimization algorithm that speeds up the solution time considerably. The algorithm has been implemented in AIMMS, which has subsequently been integrated in user-friendly software to perform the dike ring analysis. The government will update the safety standards in the Dutch Water Act based on, among other things, a new Cost-Benefit Analysis carried out by the water-consultancy company Deltares.

This paper is organized as follows. In Section 2 we define the nonhomogeneous dike height optimization problem and transform it into a MINLP problem. Additionally, we consider a robust optimization version of this problem. Section 3 discusses how the problem can be solved in practice and introduces an optimization algorithm that can be used for this purpose. Numerical results are presented in Section 4 and concluding remarks are given in Section 5.

2 Nonhomogeneous Dike Height Optimization Problem

2.1 Problem Formulation

In this section we present the nonhomogeneous dike height optimization problem. The problem is an extension of the homogeneous problem introduced by Eijgenraam et al. (2010). We shall elaborate on the differences with the homogeneous case, and the reader is referred to Eijgenraam et al. (2010) for the foundation of the common model parts. A dike ring, which protects a certain area of land against water floods, is said to be nonhomogeneous if it consists of, say L $(L > 1)$ different segments. All segments can be heightened independently of each other. Moreover, each segment has its own properties with respect to investment costs and flood probabilities. To indicate the dependence of a model parameter on a particular dike segment, a subscript ℓ $(\ell = 1, \ldots, L)$ will be added to this parameter. The set of all segments is denoted by $\mathcal{L}$.

The objective is to find an investment plan that minimizes the expected total costs. Only investments in the finite planning horizon $[0, T)$ are considered. An investment plan is represented by a tuple $(\mathbf{U}, \mathbf{t})$, with $\mathbf{U} \in \mathbb{R}_+^{L\times(K+1)}$ and $\mathbf{t} = (t_0, t_1, \ldots, t_K)^T$. The vector $\mathbf{t}$ represents the possible timings of dike segment heightenings, where $t_0 = 0 < t_1 < \cdots < t_K < T$. Hence, $K + 1$ is an upper bound on the number of segment heightenings in the planning horizon. For notational convenience, we denote $t_{K+1} = T$. The matrix $\mathbf{U}$ represents the segment heightenings, where the element $\mathbf{U}_{\ell k} = u_{\ell k}$ is the heightening (cm) of segment ℓ at time t_k $(\ell = 1, \ldots, L,\ k = 0, \ldots, K)$. Of course, heightenings are assumed to be nonnegative. If $u_{\ell k} = 0$, then this means that segment ℓ is not heightened at time t_k.

The ℓ-th row of $\mathbf{U}$, with the $K+1$ heightenings of dike segment ℓ, is denoted by $\mathbf{u}^{(\ell)}$.

Throughout the remainder of this paper we use the following notation for the cumulative segment heightening and the absolute segment height at time t ($t \geq 0$):

$$h_{\ell t} = \sum_{k: t_k \leq t} u_{\ell k}, \quad \text{and} \quad H_{\ell t} = H^{-}_{\ell 0} + h_{\ell t}.$$

where $H^{-}_{\ell 0}$ is the absolute height of segment ℓ immediately prior to a possible heightening at time $t = 0$. For notational convenience, we also use $h_{\ell k} = h_{\ell t_k}$ and $H_{\ell k} = H_{\ell t_k}$. Note that it follows from this definition that the segment height is a nondecreasing step function. Moreover, this implicitly means that heightenings are measured at the moment that the investment actions are completed. A lead time is not modeled. As a consequence, the existence of a sometimes considerable lead time between the moment that it becomes clear that action is desirable and the final completion of the investment has to be taken care of in the definition of an appropriate safety standard and the accompanying official test procedures (see Eijgenraam 2006).

The flood probability of segment ℓ at time t is given by

$$P_{\ell t} = P^{-}_{\ell 0} \exp\left(\alpha_\ell(\eta_\ell t - h_{\ell t})\right), \tag{1}$$

with $P^{-}_{\ell 0}$ (1/year) the initial flood probability, α_ℓ (1/cm) the parameter of the exponential distribution for extreme water levels and η_ℓ (cm/year) the structural increase of the water level. Both the hydraulic conditions and the quality of the dike segment are summarized by one indicator: height above the level that corresponds to the flood probability $P_{\ell 0}$. This presupposes that actual problems with piping and the quality of some of the structures are solved before further improvements in the safety level are considered. Studies by Rijkswaterstaat (the implementing agency of the Ministry of Infrastructure and Environment) (Silva and Stijnen (2005) and Stijnen et al. (2006)) confirm that under these conditions overtopping and overflow will indeed be the determining failure mechanism and that the weakest segment fully determines the flood probability of the entire dike ring. Hence, we define the flood probability of the entire dike ring at time t by $P_t = \max_{\ell \in \mathcal{L}} P_{\ell t}$.

A property that all segments have in common is that they protect the same area of land. Hence, if there is a flood, the damage does not depend on the segment in which a breach occurs. Furthermore, the potential damage costs increase in time with the economic growth rate γ. The damage costs do, however, also depend on the resulting height of the water level within a dike ring after a flood. In particular, along rivers the damage costs increase by the rise in the height of the lowest segment (in absolute height). Putting all this together yields the following damage costs, at time t, in the case of a nonhomogeneous dike ring:

$$V_t = V^{-}_0 \exp\big(\gamma t + \zeta(\min_{\ell \in \mathcal{L}} H_{\ell t} - \min_{\ell \in \mathcal{L}} H_{\ell 0})\big),$$

with V^{-}_0 the initial damage costs and ζ (1/cm) the parameter that represents the increase in damage costs depending on the height of the lowest dike segment.

The expected damage costs at time t is given by the product of the flood probability and the damage costs:

$$S_t = P_t V_t = \max_{\ell \in \mathcal{L}} S^{-}_{\ell 0} \exp\big(\beta_\ell t - \alpha_\ell h_{\ell t} + \zeta(\min_{\ell \in \mathcal{L}} H_{\ell t} - H^{-}_{\ell_0 0})\big), \tag{2}$$

where $S^{-}_{\ell 0} = P^{-}_{\ell 0} V^{-}_0$, $\beta_\ell = \alpha_\ell \eta_\ell + \gamma$ and $\ell_0 = \arg\min_\ell H^{-}_{\ell 0}$. By using the fact that the segment heights remain unchanged in the interval $[t_k, t_{k+1})$, the total expected damage in this interval can be written as

$$\int_{t_k}^{t_{k+1}} S_t \exp(-\delta t) dt = \exp(-\zeta H_{\ell_0 0}) \int_{t_k}^{t_{k+1}} \exp(-\delta t + \zeta \min_{\ell \in \mathcal{L}} H_{\ell t}) \max_{\ell \in \mathcal{L}} \big(S^{-}_{\ell 0} \exp(\beta_\ell t - \alpha_\ell h_{\ell k})\big) dt, \tag{3}$$

where δ is the discount rate.

From an optimization point of view there are two problems with the integral in (3):

(i) The minimum absolute segment height $\min_\ell H_{\ell t}$ cannot be incorporated in an optimization model as a convex constraint.

(ii) Even though the segment heights do not change during the interval $[t_k, t_{k+1})$, the segment flood probabilities $P_{\ell t}$ as defined by (1) increase monotonically in time. Hence, the segment ℓ for which the maximum flood probability is obtained may change during the interval $[t_k, t_{k+1})$.

If we want to use (3) in a MINLP model, then we have to make some assumptions about these two issues. The minimum operator in (3) refers to the fact that the size of the damage depends on the segment that is lowest in absolute height. Since in practice it is usually clear which of the segments along rivers is the lowest in absolute height, it is assumed that this segment is known in advance. Let this dike segment be denoted by ℓ^*. It turns out that, for the dike rings in the Netherlands, this assumption is always satisfied.

An obvious approach to dealing with the maximum operator in (3) is to interchange the integral and the maximum operator. Note that this yields a lower bound for (3), which introduces an error only if the segment for which the maximum is obtained changes within the interval $[t_k, t_{k+1})$. Clearly, the effect of the error will be more serious if the length of the interval is longer, and consequently this should be taken into account when defining the intervals. In the implementation of the MINLP model to be introduced in Section 2.2, we shall make sure that these intervals are small enough to guarantee a sufficiently accurate approximation.

Using the two assumptions from above, (3) can be approximated by

$$\mathcal{E}_k(\mathbf{U}, \mathbf{t}) = \max_{\ell \in \mathcal{L}} \frac{S^-_{\ell 0}}{\beta_{1\ell}} \exp\bigl(\zeta(H_{\ell^* t_k} - H^-_{\ell_0 0}) - \alpha_\ell h_{\ell k}\bigr) \Bigl[\exp(\beta_{1\ell} t_{k+1}) - \exp(\beta_{1\ell} t_k)\Bigr], \tag{4}$$

with $\beta_{1\ell} = \beta_\ell - \delta$. The total expected damage in the planning horizon $[0, T)$ is then approximated by

$$\mathcal{E}(\mathbf{U}, \mathbf{t}) = \sum_{k=0}^{K} \mathcal{E}_k(\mathbf{U}, \mathbf{t}).$$

Note that for a fixed investment plan, it is possible to evaluate the size of the approximation error, since we can accurately evaluate the minimum and maximum operators in (3). This evaluation can be used to obtain a true comparison between investment plans with different discretization schemes.

To take into account the period after the planning horizon, it is assumed that there are no changes to the expected damage after T, and hence no more investments are required. Thus, the discounted expected damage after the planning horizon is $S_T \int_T^\infty \exp(-\delta t) dt$, which can be approximated analogously to (4), i.e.,

$$\mathcal{R}(\mathbf{U}, \mathbf{t}) = \max_{\ell \in \mathcal{L}} \frac{S^-_{\ell 0}}{\delta} \exp\bigl(\beta_{1\ell} T - \alpha_\ell h_{\ell K} + \zeta(H_{\ell^* t_K} - H^-_{\ell_0 0})\bigr). \tag{5}$$

The investment costs associated with the heightening of segment ℓ at time t_k depend, of course, on the actual amount of the heightening. The costs, however, are assumed to be independent of the heightening of other segments, regardless of the moments of these heightenings. We use the same investment cost function as introduced by Eijgenraam et al. (2010), and refer to it as *exponential* investment costs. For any *positive* heightening $u_{\ell k}$, the exponential investment costs are given by

$$\mathcal{I}^e_{\ell k}(\mathbf{u}^{(\ell)}) = (c_\ell + b_\ell u_{\ell k}) \exp\bigl(-\lambda_\ell \textstyle\sum_{i=0}^{k} u_{\ell i}\bigr), \qquad \mathbf{u}^{(\ell)} \in \mathbb{R}^{K+1}_+. \tag{6}$$

Hence, the investment costs depend on the amount of the heightening and the amount of the total heightening up to time t_k. Since there are no investment costs when there is no heightening, the investment cost function is discontinuous at zero, i.e.,

$$\mathcal{I}_{\ell k}(\mathbf{u}^{(\ell)}) = \begin{cases} \mathcal{I}^e_{\ell k}(\mathbf{u}^{(\ell)}) & \text{if } u_{\ell k} > 0, \\ 0 & \text{if } u_{\ell k} = 0. \end{cases}$$

As in Eijgenraam et al. (2010), we also consider the *quadratic* investment costs

$$\mathcal{I}^q_{\ell k}(\mathbf{u}^{(\ell)}) = \phi_{\ell 0} + \phi_{\ell 1} u_{\ell k} + \phi_{\ell 2}\bigl(\textstyle\sum_{i=0}^{k} u_{\ell i}\bigr)^2, \qquad \mathbf{u}^{(\ell)} \in \mathbb{R}^{K+1}_+. \tag{7}$$

The total discounted investment costs in the planning horizon, in the case of the exponential investment cost function, are then given by

$$\mathcal{I}(\mathbf{U}, \mathbf{t}) = \sum_{\ell=1}^{L} \sum_{k=0}^{K} \mathcal{I}_{\ell k}(\mathbf{u}^{(\ell)}) \exp(-\delta t_k).$$

Since the objective is to minimize the sum of the investment costs and expected damage costs, the resulting optimization model can now be formulated as

$$\begin{aligned} &\min \mathcal{I}(\mathbf{U},\mathbf{t}) + \mathcal{E}(\mathbf{U},\mathbf{t}) + \mathcal{R}(\mathbf{U},\mathbf{t}) \\ &\text{s.t. } \mathbf{U} \in \mathbb{R}_+^{L\times(K+1)}, \quad t_0 = 0 < t_1 < \cdots < t_K < T. \end{aligned} \tag{8}$$

2.2 MINLP Model

This section discusses how the general dike height optimization problem (8) can be transformed into a mathematical optimization model that can be solved using optimization solvers. The problem as stated by (8) can be considered as a Non-Linear Programming (NLP) model since the decision variables $\mathbf{U}$ and $\mathbf{t}$ are continuous and the objective function's components are clearly nonlinear. From an optimization point of view, however, there are some issues that prevent us from actually solving the problem as stated by (8): the discontinuity of the investment cost functions at zero, and the approximation error of the expected damage in (4). The latter issue forces us to discretize the planning horizon, since continuous time variables could result in large intervals and consequently serious approximation errors. The discontinuity of the investment cost function can be resolved by discretization of the heightenings as well, or by adding binary decision variables that indicate whether a heightening is actually greater than zero or not. If both the moments and the amounts of the heightenings are discretized, then, theoretically, the problem can be solved using the dynamic programming approach as in the homogeneous case in Eijgenraam et al. (2010). Unfortunately, the state space grows too large if multiple segments are considered, which implies that the dynamic programming approach is not applicable. Therefore, we consider a MINLP approach with discretization of the planning horizon.

Next, the reformulation of problem (8) into a MINLP model is discussed. We assume that a *discretization scheme* $\mathbf{t} = (t_0, \ldots, t_{K+1})$ with $t_0 = 0 < t_1 < \cdots < t_K < t_{K+1} = T$ has been prefixed. The MINLP model then becomes:

$$\min \sum_{\ell=1}^{L}\sum_{k=0}^{K} \exp(-\delta t_k)(c_\ell y_{\ell k} + b_\ell u_{\ell k}) \exp\big(-\lambda_\ell \textstyle\sum_{i=0}^{k} u_{\ell i}\big) + \sum_{k=0}^{K} E_k + R \tag{9a}$$

$$\text{s.t. } E_k \geq \frac{S_{\ell 0}^-}{\beta_{1\ell}} \exp\big(\zeta(H_{\ell^* k} - H_{\ell_0 0}^-) - \alpha_\ell h_{\ell k}\big)\Big[\exp(\beta_{1\ell} t_{k+1}) - \exp(\beta_{1\ell} t_k)\Big], \quad \ell = 1,\ldots,L,\ k = 0,\ldots,K, \tag{9b}$$

$$R \geq \frac{S_{\ell 0}^-}{\delta} \exp\big(\beta_{1\ell} T - \alpha_\ell h_{\ell K} + \zeta(H_{\ell^* K} - H_{\ell_0 0}^-)\big), \quad \ell = 1,\ldots,L, \tag{9c}$$

$$h_{\ell k} = \sum_{i=0}^{k} u_{\ell i}, \quad \ell = 1,\ldots,L,\ k = 0,\ldots,K, \tag{9d}$$

$$H_{\ell k} = H_{\ell 0}^- + h_{\ell k}, \quad \ell = 1,\ldots,L,\ k = 0,\ldots,K, \tag{9e}$$

$$0 \leq u_{\ell k} \leq y_{\ell k} M, \quad y_{\ell k} \in \{0,1\}, \quad \ell = 1,\ldots,L,\ k = 0,\ldots,K, \tag{9f}$$

$$h_{\ell k}, H_{\ell k}, E_k, R \in \mathbb{R}, \quad \ell = 1,\ldots,L,\ k = 0,\ldots,K. \tag{9g}$$

The objective function (9a) includes the exponential investment costs with the fixed cost component c_ℓ multiplied by $y_{\ell k}$. The binary variables $y_{\ell k}$ combined with (9f) are required to ensure that either $u_{\ell k} = 0$ and the investment costs in the objective function are zero, or $u_{\ell k} > 0$ and the investment costs are equal to $\mathcal{I}_{\ell k}^e(\mathbf{u}^{(\ell)})$. Of course, it is also possible to use the quadratic investment cost function instead of the exponential. In that case both the fixed cost component and the quadratic component have to be multiplied by the binary variable $y_{\ell k}$, which yields the objective function

$$\sum_{\ell=1}^{L}\sum_{k=0}^{K} \exp(-\delta t_k)\big[\phi_{\ell 0} y_{\ell k} + \phi_{\ell 1} u_{\ell k} + \phi_{\ell 2}\big(\textstyle\sum_{i=0}^{k} u_{\ell i}\big)^2 y_{\ell k}\big] + \sum_{k=0}^{K} E_k + R \tag{10}$$

that can be used instead of (9a). In (9f), M denotes an upper bound of the highest possible dike heightening. The auxiliary variables E_k and R represent the expected damage costs in $[t_k, t_{k+1})$ and $[T, \infty)$ respectively. Constraints (9b) and (9c) are used to model the damage costs as convex constraints without using the maximum operator, as does occur in (4).

It is clear that the optimal solution to problem (9) is fully determined by the decision variables $u_{\ell k}$ ($\ell = 1, \ldots, L,\ k = 0, \ldots, K$). These decision variables could be considered the "pure" decision variables of problem (9), which, together with the discretization scheme $\mathbf{t}$, represent the investment plan $(\mathbf{U}, \mathbf{t})$ that answers the fundamental questions of *when* and *how much* should be invested in dike heightening. Moreover, for a fixed investment plan $(\mathbf{U}, \mathbf{t})$, the objective function (9) is equal to

$$Z(\mathbf{U}, \mathbf{t}) = \mathcal{I}(\mathbf{U}, \mathbf{t}) + \mathcal{E}(\mathbf{U}, \mathbf{t}) + \mathcal{R}(\mathbf{U}, \mathbf{t}).$$

2.3 Robust Optimization

It is clear that model (9) requires the input of several parameters, which in practice are often uncertain. In this section these uncertainties are taken into account. We assume, however, that all parameters are fixed, i.e. that they do not change over time. After all, we are interested in finding a robust solution, that is, a solution that performs well for a broad range of realistic instances of model parameters.

In the literature, several approaches have been proposed to deal with uncertain parameters. One class of methods assumes that a certain probability distribution can be projected on the uncertain parameter's values. For the application of dike height optimization in the Netherlands, the availability of this information is very unrealistic. The robust optimization approach by Ben-Tal et al. (2009) assumes that the uncertain parameters are contained in a so-called uncertainty set. However, depending on the shape of this uncertainty set, this would produce either a trivial or an untractable model.

For the reasons mentioned above, we consider the regret approach combined with a finite set of scenarios. A scenario represents an instance of all (uncertain) model parameters that the decision-maker sees as a possible outcome of these parameters. Of course, it is the joint responsibility of the institution that actually performs the cost-benefit analysis and the decision-maker to create a set of scenarios that is a representative reflection of the space of uncertain parameter values. Let $\mathcal{S} = \{1, \ldots, S\}$ denote this finite set of S scenarios. We add a superscript s to the model parameters to indicate the reference to the parameter values of scenario s.

The fundamental problem of parameter uncertainty is that a decision on the investment plan has to be made before the uncertain parameter values become known. Once the actual scenario is known, a measure of the quality of the decision can therefore be obtained by the difference between the costs of the chosen investment plan within the framework of the actual scenario and the costs of the optimal investment plan belonging to the actual scenario. This difference is called the regret. Formally: the regret of a given investment plan $(\mathbf{U}, \mathbf{t})$ for scenario s is defined by

$$\text{regret}(\mathbf{U}, \mathbf{t}, s) = Z^s(\mathbf{U}, \mathbf{t}) - Z^{s*}, \tag{11}$$

where Z^{s*} is the optimal expected total costs for scenario s, i.e., $Z^{s*} = \min_{\mathbf{U}}\{Z^s(\mathbf{U}, \mathbf{t}) : \mathbf{U} \in \mathbb{R}_+^{L\times(K+1)}\}$.

We use the regret criterion to produce two robust optimization approaches. The first approach is to find an investment plan that minimizes the *average regret* over all scenarios. Mathematically, this is equivalent to

$$\min_{\mathbf{U}}\left\{ \tfrac{1}{S}\sum_{s=1}^{S}\big(Z^s(\mathbf{U}, \mathbf{t}) - Z^{s*}\big) : \mathbf{U} \in \mathbb{R}_+^{L\times(K+1)} \right\}. \tag{12}$$

Note that it can be shown that this is equivalent to minimizing the total costs over all scenarios, because the Z^{s*} only affect the objective value, and not the optimal robust solution.

The second robust optimization approach is to minimize the *maximum regret* over all scenarios, which can be written as

$$\min_{z,\mathbf{U}}\left\{ z \in \mathbb{R} : z \geq \text{regret}(\mathbf{U}, \mathbf{t}, s)\, \forall s \in \mathcal{S},\ \mathbf{U} \in \mathbb{R}_+^{L\times(K+1)} \right\}. \tag{13}$$

Both these robust optimization approaches can be formulated as a MINLP model very similar to (9). The robust models have to include different variables for the auxiliary variables E_k and R for each scenario. However, the decision variables $u_{\ell k}$ and $y_{\ell k}$ are common for all scenarios, since we are looking for only one investment plan.

3 Implementation Issues

One of the project goals, set by Deltares, is that the problem can be solved for all major dike rings in a reasonable amount of time without the necessity to tune the algorithm's settings for specific dike rings. Hence, we are not designing a solution algorithm for one particular instance of the problem. This section discusses the implementation of a solution algorithm for the nonhomogeneous dike height optimization problem. A heuristic algorithm is needed because MINLP (9) cannot be solved in reasonable time for dike rings with more than 6 segments. The algorithm presented here has been implemented in AIMMS and the software company HKV has integrated this model in the software package OptimaliseRing (Duits 2009a,b), used by the actual performers of the cost-benefit analysis.

Section 3.1 discusses the selection of the MINLP solver used for the dike height optimization problem. Sections 3.2 and 3.3 discuss the two main ideas of the iterative optimization algorithm, which is presented in Section 3.4. An illustrative example of the algorithm is given in Section 3.5.

3.1 Solver Selection

To be able to solve a MINLP one needs a MINLP solver. MINLPs can be highly complex and therefore the performance of MINLP solvers can vary vastly for different problems. The performance differences may pertain to the quality of the final solution as well as to the solution time. An extensive overview of MINLP solvers is given by Bussieck and Vigerske (2011). The most notable distinction one can make between MINLP solvers is whether the solver aims to solve convex or nonconvex MINLPs. The first class of solvers is usually based on the outer approximation algorithm by Duran and Grossmann (1986) using either NLP or MIP relaxations. These solvers ensure global optimality only for convex problems, but they can serve as a heuristic for nonconvex problems. The second class of solvers ensures global optimality, and to reach this goal a convexification step (Tawarmalani and Sahinidis 2002) has to be added to the algorithm.

The nonhomogeneous dike height optimization problem formulated by MINLP (9) is not a convex MINLP. However, in Appendix A it is shown that the problem possesses many useful convexity properties. For example, if the quadratic investment cost function is used, the MINLP model (9) can be reformulated into a convex MINLP model, see Appendix A.4. This is not the case for the exponential investment cost function. Hence, it is not immediately obvious whether a convex or nonconvex solver is most suitable for solving the dike height optimization problem. We have tested the solver performance for two typical solvers: AOA from the class of convex solvers, and BARON from the class of nonconvex solvers. Both solvers can be called from the modeling tool AIMMS, which we used to implement the dike height optimization problem. AOA is an outer approximation algorithm inside AIMMS and uses general purpose NLP and MIP solvers to solve subproblems. In our experiments we used CONOPT for the NLP subproblems, and CPLEX for the MIP subproblems.

After a limited number of experiments it turned out that BARON's solution time is too long for our practical purposes. Apparently, the effort that BARON has to make to guarantee a global solution is too high. Even to obtain a reasonable suboptimal solution requires an unacceptable amount of time. The results of AOA, on the other hand, did yield promising results, and as a consequence we selected AOA as our preferred solver for the nonhomogeneous dike height optimization problem.

3.2 Discretization Scheme

The MINLP (9) requires a discretization scheme $t_0 = 0 < t_1 < \cdots < t_K < t_{K+1} = T$ to be defined. Several things should be kept in mind when choosing a discretization scheme. First of all, the discretization scheme determines the problem size. For any solution method that is selected to solve MINLP (9), the problem size will be an important factor with respect to the solution time. Besides the number of segments L, which is exogenously determined, the problem size is determined by the discretization of the planning horizon $[0, T)$ into $K + 1$ decision moments: $t_0 = 0 < t_1 < \cdots < t_K < T$. It is important to note that the problem formulation (9) does not require equidistant interval sizes, even though the most natural choice of a discretization scheme would satisfy this property. In Appendix B.2.1 it is shown how this flexibility can be exploited. Moreover, it is interesting to compare solutions obtained by applying different discretization schemes to MINLP (9) for the same underlying dike height optimization problem.

Secondly, the interval sizes $t_k - t_{k-1}$ are important for the accuracy of the expected damage determined by equation (4). The reason for this is that the expected damage is approximated by interchanging the integral and maximum operators in the original definition of the expected damage. The approximation will be more accurate if small interval sizes are used. Hence, from this perspective it is desirable to have many small intervals, and consequently a large K. Computational results have shown, however, that for almost all realistic problems the approximation is very acceptable for equidistant intervals of ten years. Note that this boils down to "rounding" the decision moment up or down with a maximum of five years, which is not that much if you consider the intervals between segment heightenings.

Finally, another aspect to keep in mind is that the moments at which a dike heightening can occur are restricted by the discretization scheme. Therefore, having smaller intervals offers more flexibility to determine the exact moments of dike heightenings, which can improve the objective value of the problem.

A good discretization scheme should balance the diversity of the investment plans it allows and the accuracy of the approximations in the model against the solution time of the resulting problem. We try to find this optimal balance in two stages. In the first stage, a rough grid of possible decision moments t_k $(k = 0, 1, \ldots, K)$ is selected, which offers less flexibility in selecting an investment plan and could result in less accurate approximations of the expected damage. The advantage, however, is that less decision variables are required and that the solution time decreases significantly. A rough discretization scheme could therefore be used to obtain a good initial solution to the problem, which can be exploited in the second stage of the algorithm. The problem can now be solved with a finer discretization scheme, and by considering only investment plans in the neighborhood of the initial solution it is not necessary to increase the number of decision variables. Numerical evidence shows that this is a good approach to obtain solutions on a finer grid of decision moments.

Appendix B.2 describes in greater detail how the discretization scheme can be initialized and how it can be refined in the second stage by using the results from the first stage.

3.3 Nearly-Redundant Constraints

With respect to the dike height optimization problem, it is obvious that the optimal solution will satisfy some general properties that can be deduced using common sense or by mathematical analysis of the problem. If a certain property of the MINLP's optimal solution is known, then the feasible region of the model can be reduced by adding this property to the model as a constraint. This constraint is not supposed to change the MINLP's optimal solution, and it is therefore called redundant. However, by including a redundant constraint the MINLP solver may be able to find the optimal solution a lot quicker.

If a redundant constraint is added to an optimization problem, then it is important to determine whether this constraint is redundant for *all* possible instances of the problem, or just for a particular instance of interest. For example, in practice it is unlikely that a dike heightening smaller than 1cm can ever be optimal, yet it is fairly obvious that instances can be invented for which this is in fact optimal. For the dike height optimization problem, many such "rules of thumb" can be formulated that cannot be proven to apply to all possible instances of the general problem. Moreover, many rules of thumb include some kind of wild guess, such as the 1cm boundary in the previous example. This results in two contradictory objectives. On the one hand, it is desirable to set the boundary very tight so that a large part of the original feasible region can be neglected. On the other hand, the boundary has to be set such that the resulting constraint is indeed redundant for all instances of interest.

Due to the dike height optimization problem size, which of course depends on L and K, it is very interesting to consider redundant constraints. Moreover, it can prove useful to use "nearly" redundant constraints, which for our purposes are defined as constraints that are satisfied by optimal solutions or almost optimal solutions to most practical instances of MINLP (9). Nearly-redundant constraints can be used in an iterative setting, for example to find a good solution very quickly in the initial stage of an algorithm. In the subsequent stages of the algorithm, the actual redundancy of the constraints can be tested by using the solution obtained in the initial stage.

The remainder of this section discusses two nearly-redundant constraints that turned out to be particularly useful for the dike height optimization problem in the Netherlands.

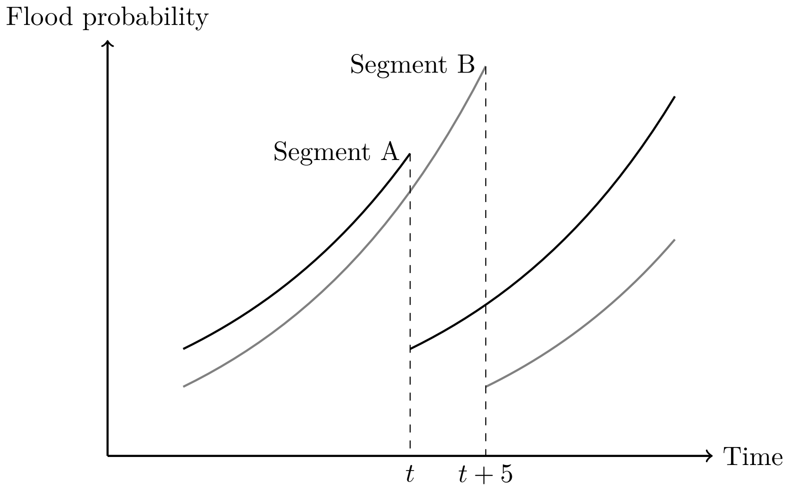

Figure 1: Flood probabilities of two dike segments having heightenings in quick succession.

3.3.1 Period Between Heightenings of the Same Segment.

Consider two successive heightenings of the same dike segment ℓ. Since fixed costs are incurred for each segment heightening, it makes little sense to schedule the two heightenings shortly after each other. In such a case, a single dike heightening carried out at the same time as the first heightening with an amount equal to the sum of the two heightenings will be cheaper, because it saves on fixed costs and the expected damage in the intermediate period is smaller. Of course, the definition of "shortly after each other" is not self-evident, and will depend on the parameter settings. However, if we consider the practical application in the Netherlands, then a minimum period of 40 years between dike heightenings of the same segment seems very reasonable, except for sand suppletion along the coast. A possible exception to this rule applies to the start of the planning horizon. In many cases, dike heightenings of different segments turn out to occur at the same point in time. Hence, at the start of the planning horizon a dike segment has to be brought into the same "rhythm" as the other segments. In such a case the optimal period between the first two dike heightenings of a segment can be smaller than defined by this rule of thumb, because the benefit of the first dike heightening is very high.

For the property described above, the following constraint can be added to the MINLP formulation:

$$\sum_{i:\, t_k < t_i \leq t_k + Y_\mathrm{p}} y_{\ell i} \leq 1, \qquad \forall \ell, k : Y_\mathrm{s} \leq t_k \leq T - Y_\mathrm{p}, \tag{14}$$

where Y_p is the minimum period between dike heightenings, and Y_s is the number of years at the start of the planning horizon that is not taken into account.

3.3.2 Groups of Segments Heightened at the Same Time.

A common property of the optimal solution to many dike height optimization problems is that the heightening of different segments takes place at the same point in time. This does not imply that all segments should always be heightened at the same time. The point is that it is very unlikely that two segments are heightened in quick succession. How to explain this? Recall that the driving motivation to heightening dikes is to decrease the expected damage by decreasing the flood probability. Since the flood probability is equal to the maximum flood probability, measured across all dike segments, the overall benefit of updating segment A five years prior to segment B is usually very low. This is illustrated by Figure 1, where the heightening of segment A at time t decreases the total flood probability by just a small amount, since the flood probability of segment B is just slightly smaller than segment A's flood probability immediately prior to segment A's heightening.

It is very easy to enforce that all segments are heightened simultaneously, with the possible exception of a short period at the start of the planning horizon. This is achieved by the constraints

$$y_{1k} = y_{\ell k}, \qquad \ell = 2, \ldots, L,\ k : t_k \geq Y_\mathrm{f}, \tag{15}$$

with Y_f the interval at the start of the planning horizon that is excluded. This constraint performs remarkably well when it is added to MINLP (9): the solution time decreases drastically for problems with many segments, and it does not increase the optimal objective value for many practical problems, so the constraint is actually redundant for these problems. Unfortunately, for some problems constraint (15) will increase the optimal objective value by an unacceptable amount. Therefore, it cannot be permanently added to the problem. However, constraint (15) can still be used in an iterative algorithm to quickly obtain a reasonable solution to the original problem.

The reason that constraint (15) sometimes affects the optimal solution of the problem is that some dike rings have one or more segments with very different characteristics than the other segments. Therefore, we would like to relax constraint (15) such that it only applies to segments that have more or less the same characteristics. Let $\mathcal{G}$ be such a subset of $\mathcal{L}$. The relaxed constraints then become

$$y_{\ell' k} = y_{\ell k}, \qquad \forall \ell, k : \ell \in \mathcal{G},\ \ell \neq \ell',\ t_k \geq Y_\mathrm{f}, \tag{16}$$

with ℓ' an arbitrary segment in $\mathcal{G}$. Note that (15) is a special case of (16) with $\mathcal{G} = \mathcal{L}$.

In Appendix B.3 two methods are proposed that can be used to partition $\mathcal{L}$ into two disjoint subsets $\mathcal{G}_1$ and $\mathcal{G}_2$ for which constraint (16) will be imposed.

3.4 Iterative Algorithm

This section presents the iterative heuristic algorithm based on the ideas of the previous sections. The iterative algorithm is based on the following generic MINLP, which consists of the basic model (9), the constraint (14), two instances of constraint (16) and constraints related to refining the discretization scheme:

$$\min \sum_{\ell \in \mathcal{L}} \sum_{k=0}^{K} \exp(-\delta t_k)(c_\ell y_{\ell k} + b_\ell u_{\ell k}) \exp\big(-\lambda_\ell \textstyle\sum_{i=0}^{k} u_{\ell i}\big) + \sum_{k=0}^{K} E_k + R \tag{17a}$$

$$\text{s.t. } E_k \geq \frac{S_{\ell 0}}{\beta_{1\ell}} \exp\big(\zeta(H_{\ell^* k} - H_{\ell_0 0}) - \alpha_\ell h_{\ell k}\big)\Big[\exp(\beta_{1\ell} t_{k+1}) - \exp(\beta_{1\ell} t_k)\Big], \quad \ell = 1, \ldots, L,\ k = 0, \ldots, K, \tag{17b}$$

$$R \geq \frac{S_{\ell 0}}{\delta} \exp\big(\beta_{1\ell} T - \alpha_\ell h_{\ell K} + \zeta(H_{\ell^* K} - H_{\ell_0 0})\big), \qquad \ell = 1, \ldots, L, \tag{17c}$$

$$\sum_{i:\, t_k < t_i \leq t_k + Y_\mathrm{p}} y_{\ell i} \leq 1, \qquad \forall \ell, k : Y_\mathrm{s} \leq t_k \leq T - Y_\mathrm{p}, \tag{17d}$$

$$y_{\ell_1 k} = y_{\ell k}, \qquad \forall \ell, k : \ell \in \mathcal{G}_1,\ \ell \neq \ell_1,\ t_k \geq Y_\mathrm{f}, \tag{17e}$$

$$y_{\ell_2 k} = y_{\ell k}, \qquad \forall \ell, k : \ell \in \mathcal{G}_2,\ \ell \neq \ell_2,\ t_k \geq Y_\mathrm{f}, \tag{17f}$$

$$h_{\ell k} = \sum_{i=0}^{k} u_{\ell i}, \qquad k = 0, \ldots, K,\ \ell \in \mathcal{L}, \tag{17g}$$

$$y_{\ell k} = 0, \quad u_{\ell k} = 0 \qquad k \in F,\ \ell \in \mathcal{L}, \tag{17h}$$

$$H_{\ell k} = H_{\ell 0} + h_{\ell k}, \qquad k = 0, \ldots, K,\ \ell \in \mathcal{L}, \tag{17i}$$

$$0 \leq u_{\ell k} \leq y_{\ell k} M, \quad y_{\ell k} \in \{0, 1\}, \qquad k = 0, \ldots, K,\ \ell \in \mathcal{L}, \tag{17j}$$

$$h_{\ell k}, H_{\ell k}, E_k, R \in \mathbb{R}, \qquad k = 0, \ldots, K,\ \ell \in \mathcal{L}, \tag{17k}$$

where ℓ_1 and ℓ_2 are arbitrary elements from the sets $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively.

The basic idea of the iterative algorithm is as follows. First, a discretization scheme is generated that does not result in too many (binary) decision variables for the generic problem, such that it can be solved in a reasonable amount of time. Second, the generic MINLP (17) is solved for different instances of the nearly-redundant constraints (17e) and (17f). This should give at least one fairly good solution. Finally, the algorithm zooms in on the best solution found so far by refining the discretization scheme in the neighborhood of segment heightenings and defining the set of uninteresting decision moments F

(cf. Appendix B.2.2). Constraints (17h) guarantee that no solutions in the uninteresting region of the time horizon are considered. Additionally, the tolerance for the stopping criterion could be tightened to obtain the best possible result out of this final solve.

Besides a maximum solution time per solve of the generic model, a relative tolerance stop criterion is used in the AIMMS implementation. This relative tolerance stop criterion is satisfied if the objective value's lower bound, obtained by AOA's MIP subproblem, is within a certain percentage of a solution's objective value.

The iterative optimization algorithm is formally defined by:

1. Define a discretization scheme $\mathbf{t}$ according to the method in Section B.2.1.
2. Solve the generic MINLP (17) for all promising partitions $\mathcal{G}_1$ and $\mathcal{G}_2$ according to Appendix B.3 and compute the true objective values by getting rid of the approximation errors in the expected damage.
3. Select the solution with the lowest true objective value and select corresponding segment subsets $\mathcal{G}_1$ and $\mathcal{G}_2$.
4. Refine the discretization scheme according to the method in Appendix B.2.2 and determine which decision moments that are outside the range of interest, i.e., define the set $F \subset \{0, 1, \dots, K\}$.
5. Tighten the relative stop criterion and resolve MINLP (17) with the selected subsets $\mathcal{G}_1$ and $\mathcal{G}_2$ from step 3 and the updated discretization scheme and associated set F from step 4.

In the next section, an illustrative example of the iterative algorithm is presented. In Appendix B.4 the performance of the algorithm is analyzed to show that the reduction in solution time is obtained at no or a very small deterioration of the objective value.

3.5 Example Iterative Algorithm

This section discusses some results of the iterative algorithm from Section 3.4. Recall that the algorithm iteratively solves the generic MINLP (17) for different partitions $\{\mathcal{G}_1, \mathcal{G}_2\}$ of $\mathcal{L}$, and different discretization schemes. In Appendix B.3 it is proposed to use partitions by selecting individual segments or by ranking the segments based on flood probability related parameters. In addition, partitions based upon differences in investment costs are proposed. In the current database used by Deltares such instances are not present, hence these partitions are not applied in practice. However, these cost-based partitions are tested on artificial test problems and appear to be valuable if significant differences in investment costs are found.

The iterative algorithm has been applied to dike ring 17, *IJsselmonde*, which has 6 segments, using the exponential investments cost function. Table 1 gives a summary of the algorithm's iterations. Note that the iteration numbers correspond to the subsequent solves of the generic MINLP and not the steps of the algorithm. The third and fourth column show the generic MINLP's objective value and the true objective value of this solution, respectively. The last column shows the solution time of the solve in minutes, where the decimal fraction refers to the fraction of one full minute.

The first iteration corresponds to the case $\mathcal{G}_1 = \mathcal{L}$, which forces simultaneous heightenings of all dike segments. In the subsequent solves one segment is allowed to have heightenings at different times than all other segments (cf. (18) in Appendix B.3). There are only marginal improvements in objective value for most segments, except for segment *5-Oude Maas*. In the next three iterations of the algorithm the segments are partitioned into two subsets based on the segments' flood probability parameters (cf. (20) in Appendix B.3). In iteration 8 the segments are partitioned into a subset with segments *1-Nieuwe Maas* and *5-Oude Maas*, and a subset with the other four segments. It can be seen that this solution yields an additional benefit compared to iteration 6, where only dike segment *5-Oude Maas* is released from the other segments. The other two partitions tried in iteration 9 and 10 do not yield any further improvement, hence iteration 8 gives the best objective value of all MINLP variations. Therefore, the solution from iteration 8 is used to refine the discretization scheme and the generic MINLP is solved using this refined discretization scheme, with the relative tolerance tightened from 1% to 0.25% in iteration 11. Again the objective value improves, which indicates that the initial discretization scheme was too rough and that the timing of the dike heightenings can be improved.

Another observation to be drawn from Table 1 is that there are indeed differences between the MINLP objective and the true objective. The difference represents the approximation error of the expected

Iteration	Description	MINLP objective (M€)	True objective (M€)	Solution time (min)
1	All at the same time	388.9774	389.4514	0.12
2	Free segment 1-Nieuwe Maas	388.4750	389.0692	0.16
3	Free segment 2-Nieuwe Maas	388.8734	389.3527	0.31
4	Free segment 3-Noord	388.8734	389.3527	0.16
5	Free segment 4-Noord	388.9774	389.4514	0.26
6	Free segment 5-Oude Maas	379.3647	380.4492	0.34
7	Free segment 6-Oude Maas	388.8734	389.3527	0.37
8	Sorted groups no. 1	377.7331	378.7116	0.34
9	Sorted groups no. 2	382.7661	383.6943	0.55
10	Sorted groups no. 3	388.8734	389.3527	0.61
11	Refine best solution	377.0474	377.3725	0.10

Table 1: Algorithm iterations for dike ring 17.

damage, caused by interchanging the integral and maximum operator in (3). This approximation error is actually quite moderate, especially if one considers the impact that many uncertain model parameters could have on this objective value.

4 Numerical Results

As discussed in Section 3, the optimization algorithm has been implemented in AIMMS using the AOA solver. All numerical results in this section were obtained using AIMMS 3.8.5 with CPLEX 11.2 and CONOPT 3.14G on a PC with an Intel Core 2 CPU processor.

A database with data about the dike rings in the Netherlands was provided by Deltares. This database contains all relevant parameters for the nonhomogeneous dike height optimization problem, the only exception being the parameters for the quadratic investment cost function. These missing parameters were estimated by approximating the exponential investment cost function (see Brekelmans et al. 2009, Appendix D).

The main results for the nonhomogeneous dike height optimization problem are presented in Section 4.1. In Section 4.2 the results of the homogeneous model are compared to the results of the nonhomogeneous model. Robust optimization results are discussed in Section 4.3.

4.1 Results Nonhomogeneous Problem

4.1.1 Overview Dike Rings.

A selection of the dike rings from Deltares' database were optimized by our optimization algorithm. For all experiments we used common values for the discount rate per year ($\delta = 0.0247$) and the economic growth rate per year ($\gamma = 0.019$). A summary of the results for the exponential investment costs is shown in Table 2. The first two columns give the dike ring number along with the number of segments in the dike ring. The third column gives the MINLP model's objective value of the algorithms final iteration. The fourth column gives a true evaluation of this objective value that does not suffer from an approximation error in the expected damage. It can be seen that the MINLP's objective is indeed a lower bound and that the approximation error is very modest, which indicates that the approximation of the expected damage is suitable for our MINLP model.

The fifth column in Table 2 gives the solution time in minutes. There does not appear to be a clear relationship between the number of segments and the solution time. This is mainly due to the fact that the discretization scheme is created in such a way that the number of resulting decision variables does not depend on the number of segments. In other words, a dike ring with more segments has a rougher discretization scheme than a dike ring with less segments, as explained in Appendix B.2.1.

For the same set of experiments, Table 3 shows the moments of the first three updates of the dike rings, which could correspond to one or more segment heightenings taking place at the same point in

Dike ring	Segments	MINLP objective (M€)	True objective (M€)	Solution time (min)
10	4	107.51	107.51	0.52
13	4	10.38	10.38	0.07
14	2	94.04	94.04	0.54
16	8	1044.45	1046.08	6.24
17	6	377.05	377.37	3.33
21	10	217.40	217.71	2.23
22	5	373.98	374.08	7.62
36	6	395.65	395.65	60.19
38	3	136.26	136.29	59.33
43	8	486.72	488.10	1.65
47	2	16.57	16.57	8.54
48	3	42.92	42.92	2.77

Table 2: Results optimization algorithm for a selection of dike rings.

time. In addition, the table shows the effect the heightenings have on the dike ring's flood probability, i.e., the flood probabilities just before and just after the update are listed. It turns out that for these parameter settings the safety standards for many dike rings have to be tightened in comparison with the current safety standards in the Dutch Water Act. The Dutch government has anticipated these results by reserving additional funds for future dike heightening. For the new safety standards in this example, there are five out of twelve dike rings that require immediate segment heightenings at $t = 0$. The results also clearly indicate that the flood probabilities just prior to a heightening decrease over time. This is a result of the economic growth, which increases the damage costs if a flood occurs, and therefore it is beneficial to let the flood probabilities decrease over time.

	First heightening			Second heightening			Third heightening		
Dike ring	t	$P_-(t)$	$P_+(t)$	t	$P_-(t)$	$P_+(t)$	t	$P_-(t)$	$P_+(t)$
10	68	6.6×10^{-4}	6.7×10^{-5}	156	1.2×10^{-4}	1.4×10^{-5}	244	2.5×10^{-5}	2.9×10^{-6}
13	140	1.8×10^{-4}	1.6×10^{-5}	244	3.7×10^{-5}	2.5×10^{-6}	-	-	-
14	36	1.5×10^{-4}	2.3×10^{-5}	104	4.6×10^{-5}	6.5×10^{-6}	168	1.3×10^{-5}	1.8×10^{-6}
16	0	5.0×10^{-4}	2.8×10^{-4}	40	3.7×10^{-4}	7.7×10^{-5}	105	1.2×10^{-4}	2.5×10^{-5}
17	20	3.8×10^{-4}	9.1×10^{-5}	81	1.9×10^{-4}	1.3×10^{-5}	165	4.3×10^{-5}	2.9×10^{-6}
21	0	5.0×10^{-4}	2.5×10^{-4}	45	5.2×10^{-4}	5.3×10^{-5}	120	1.5×10^{-4}	1.4×10^{-5}
22	7	5.2×10^{-4}	4.5×10^{-5}	100	1.1×10^{-4}	8.5×10^{-6}	200	2.3×10^{-5}	1.2×10^{-6}
36	36	1.1×10^{-3}	1.7×10^{-4}	102	4.1×10^{-4}	6.3×10^{-5}	165	1.5×10^{-4}	2.4×10^{-5}
38	0	6.7×10^{-4}	2.7×10^{-4}	28	4.6×10^{-4}	1.9×10^{-5}	126	8.6×10^{-5}	3.2×10^{-6}
43	0	2.7×10^{-4}	2.7×10^{-4}	30	4.6×10^{-4}	3.9×10^{-5}	120	9.7×10^{-5}	7.3×10^{-6}
47	30	2.5×10^{-4}	1.2×10^{-5}	120	4.0×10^{-5}	1.2×10^{-5}	200	1.6×10^{-5}	5.8×10^{-7}
48	0	2.8×10^{-4}	1.2×10^{-5}	77	3.0×10^{-5}	2.9×10^{-6}	154	7.1×10^{-6}	6.6×10^{-7}

Table 3: Moments (in years measured from the start of the planning horizon) of the first three dike ring updates and the flood probabilities just before ($P_-(t)$) and after ($P_+(t)$) the updates.

4.1.2 Dike Ring 17.

In Section 3.5 we discussed the iterations of our algorithm applied to dike ring 17. Here we take a closer look at the resulting solution. Figures 2 and 3 give a graphical overview of the final solution obtained with the iterative algorithm. Figure 2 shows the cumulative heightenings of the six segments during the 300-year planning horizon. Figure 3 shows the resulting segment flood probabilities. It can be seen that the two segments *1-Nieuwe Maas* and *5-Oude Maas* are not heightened together with the other segments at $t = 20$. Figure 3 also shows why it is not necessary to heighten these two segments: their flood probabilities are still very low compared to the other segments. Although in this particular example there is a moment at which not all segments are heightened simultaneously, the figure clearly demonstrates why simultaneity very frequently leads to very good, or even optimal, results. Recall that a dike ring's flood probability is determined by the maximum segment flood probability. Hence, if a single segment is not heightened simultaneously with the other segments, then it is likely that this segment's flood probability will become, or even remain, the dike ring's maximum flood probability. The benefit of heightening the other segments, in terms of decreasing the expected damage, is therefore usually smaller than the incurred investment costs.

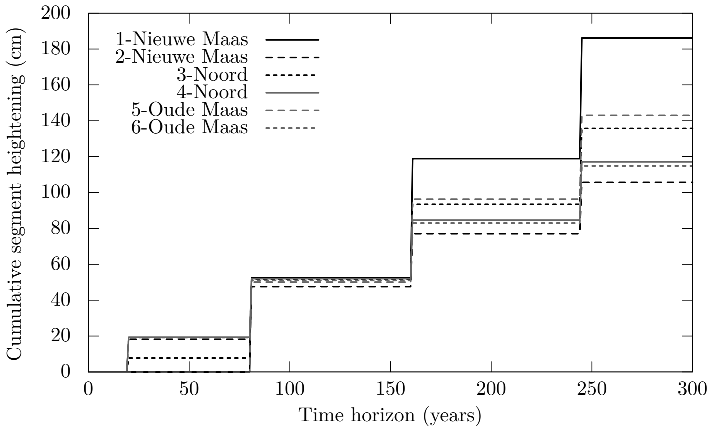

Figure 2: Cumulative segment heightening dike ring 17.

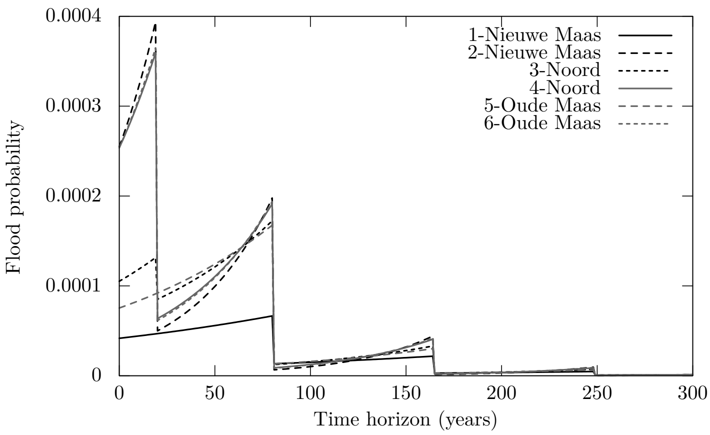

Figure 3: Segment flood probabilities dike ring 17.

4.1.3 Dike Ring 16.

Dike ring 17, discussed in Section 4.1.2, is an example of a dike ring where, in the optimal solution, not all segments are always heightened simultaneously. Very often this is the case, however. Dike ring 16, *Alblasserwaard en Vijfheerenlanden*, is a dike ring with eight segments. Figure 4 shows the flood probabilities of the dike ring's eight segments for the solution obtained with the iterative algorithm and the quadratic investment cost function. It can be seen that all segments are heightened simultaneously after the start of the planning horizon.

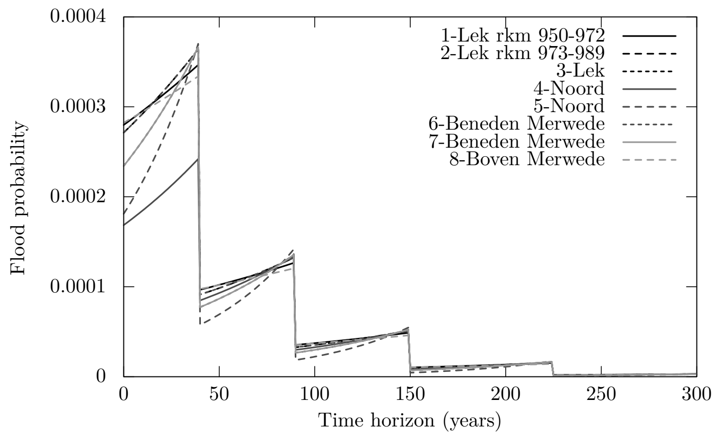

Figure 4: Segment flood probabilities dike ring 16.

4.1.4 Exponential vs. Quadratic Investment Costs.

The most important criterion for selecting an investment cost function is of course that it corresponds as closely as possible to the actual investment costs. On the other hand, the effect of a particular investment cost function on the optimization process has to be taken into account. This effect can be on the solution time as well as on the quality of the returned solution. With this in mind, convex investment cost functions are preferred for both types of effects. Hence, with regard to the two investment cost functions introduced in Section 2.2, the exponential and quadratic, the quadratic is preferred. With respect to the solution time, it is easy to verify whether this is indeed correct. Numerical experiments (see Brekelmans et al. 2009, Section 4.4) for thirteen dike rings have shown that the solution time for the exponential investment cost function is, on average, a factor of three higher than for the quadratic investment cost function.

As it is impossible to approximate an exponential investment cost function with a quadratic function without approximation errors, it may be expected that, with respect to both the location and the objective value of the optimal solution, different results are obtained for the two investment cost functions. This conjecture is confirmed by the numerical results (see Brekelmans et al. 2009, Section 4.4). Another interesting observation is that if the optimal solutions for both investment cost functions represent different investment plans, then the objective values for both solutions, evaluated for both investment cost functions, usually differ by no more than one percent. This indicates that there can be several structurally different solutions that have nearly optimal objective values. This is an interesting finding with respect to robust optimization: there might be reasons to prefer a slightly sub-optimal solution to the globally optimal solution of one set of parameters, for instance in the event of parameter uncertainty.

4.2 The Homogeneous vs. Nonhomogeneous Problem

If the nonhomogeneous dike height optimization problem is compared to the homogeneous case, then it is clear that the nonhomogeneous case is more complex and, certainly for dike rings with many segments, computationally more demanding. In this section we investigate whether this additional effort is justified, that is, does it yield an additional benefit to model a dike ring using multiple nonhomogeneous segments instead of a single homogeneous segment?

Consider a dike ring with more than one segment and suppose we want to obtain an investment plan using the homogeneous case. The main problem for this approach is to convert the nonhomogeneous segments' characteristics into the characteristics of a single homogeneous segment. For investment costs parameters such as c_ℓ and b_ℓ it seems obvious to aggregate them over all segments. For other parameters such as λ_ℓ, α_ℓ and η_ℓ it is not obvious at all. A reasonable approach could be to take a (weighted) average of these parameters over all segments. Clearly, this method cannot always result in a satisfactory conversion of the nonhomogeneous segments into a homogeneous segment, and this is of course exactly the reason why we are considering nonhomogeneous dike rings in the first place. However, it is interesting to analyze the quality of the resulting investment plans. Once an investment plan for the homogeneous case has been obtained, then it can be transformed back into an investment plan for the original dike ring with multiple segments, and it can be evaluated using the nonhomogeneous model.

Recall the dike rings 16 and 17 from the previous sections with eight and six dike segments, respectively. For each dike ring, the segments have been reduced to a single segment according to the method described above. Subsequently, the homogeneous cases have been solved and transformed back into the original nonhomogeneous problems. Table 4 shows the results of this method as well as the previous results obtained with the nonhomogeneous model. Clearly, the solutions obtained with the homogeneous model are inferior to the solutions from the nonhomogeneous case.

Dike ring	Obj. homogeneous solution (M€)	Obj. nonhomogeneous solution (M€)
16	3410	1029
17	396	377

Table 4: Comparison of results of homogeneous and nonhomogeneous cases. Table gives the objective values of the solutions obtained with (1) the reduction to the homogeneous case and (2) the regular nonhomogeneous case. Both solutions are evaluated using the nonhomogeneous model's true evaluation of the objective.

The results for the homogeneous case are probably very sensitive to the method used to reduce the segment characteristics. In Eijgenraam (2005) a qualitative reasoning is given why it is likely that the timing of the heightenings for the different segments is the same. Further, it would be the segment with the biggest value for β_ℓ and—since γ is the same—the biggest value for $\alpha_\ell \eta_\ell$ that is critical for the length of the time interval between two subsequent heightenings. (Both suppositions are now confirmed e.g., in Figures 3 and 4.) Therefore the parameters of this critical segment were used in the calculation of the damage costs. The investment costs are calculated as a weighted average. The amount of the heightening differs between the segments, but has a fixed relation to the calculated time interval, say τ. Using the exponential cost function the heightening per segment is calculated as

$$u_{\ell k} = \frac{\beta_\ell \cdot \tau}{\alpha_\ell - \zeta + \lambda_\ell}.$$

So, the differences between a homogeneous and a nonhomogeneous calculation can be smaller for other methods than used for Table 4, but there is no obvious method that performs well under all circumstances. Therefore, these results demonstrate that solving the dike height optimization problem using the homogeneous model is not sensible if the dike ring is in fact nonhomogeneous.

4.3 Robust Optimization Results

In this section we discuss the solutions obtained by the two robust optimization approaches proposed in Section 2.3, which is desirable if model parameters are uncertain. We consider a particular example with five scenarios, i.e., instances of model parameters, for dike ring 29, *Walcheren*. The scenarios have been created by selecting a default scenario, which would have been used if the standard approach had been applied, and subsequently we consider a low and high variation of the climate and economic model parameters. Table 5 gives an overview of how three different levels of a model parameter (low, average, high) are combined in the five scenarios. The climate scenarios are based on an analysis by the KNMI (Royal Netherlands Meteorological Institute). The exact values for these parameters are taken from the database provided by Deltares.

Let us first consider each of the five scenarios individually, i.e., the iterative algorithm from Section 3.4 is applied to each of the five instances of model parameters separately. This yields five different investment plans, say *sol-1* to *sol-5*, that are optimal for the five scenarios, respectively. The corresponding objective values are shown in the column labeled *Optimal objective* in Table 6. It can be seen that scenario 5 is the worst-case scenario since it has the highest optimal objective of all scenarios.

Now consider the situation where the parameter values are unknown and can take on the values defined in the five scenarios, and that only a single investment plan can be selected. Suppose that the default scenario's optimal investment plan, that is *sol-3*, is selected. The optimal objective value for scenario 3 is 79.50, but what happens if scenario 5 occurs? The resulting objective value has to be worse

No.	Climate	η	α	P_0	Economic	γ
1	low	low	high	low	low	low
2	low	low	high	low	high	high
3	avg.	avg.	avg.	avg.	avg.	avg.
4	high	high	low	high	low	low
5	high	high	low	high	high	high

Table 5: Scenario definitions based upon climate and economic model parameters.

Scenario				Regret for solution						
No.	Climate	Econ.	Optimal objective	sol-1	sol-2	sol-3	sol-4	sol-5	avg-regret	max-regret
1	low	low	60.93	0.00	12.44	5.89	7.70	26.15	13.26	12.13
2	low	high	90.48	69.70	0.00	3.78	8.97	6.89	2.22	3.14
3	avg.	avg.	79.50	15.00	2.43	0.00	2.49	12.97	4.01	4.00
4	high	low	79.77	25.36	3.36	1.94	0.00	11.65	2.77	2.84
5	high	high	107.74	1690.81	27.30	101.51	64.22	0.00	10.11	12.13
		avg regret	N/A	360.17	9.106	22.624	16.676	11.532	6.474	6.848
		max regret	N/A	1690.81	27.30	101.51	64.22	26.15	13.26	12.13

Table 6: Results robust optimization dike ring 29.

than the optimal objective value for scenario 5, which is 107.74, and in this case it is equal to 209.24. This means that the regret of *sol-3* for scenario 5 equals $209.24 - 107.74 = 101.51$, or $\text{regret}(U_{\text{sol}-3}, 5) = 101.51$, where $U_{\text{sol}-3}$ denotes the investment plan *sol-3*. The regret of the solutions *sol-1* to *sol-5* for all five scenarios are listed in Table 6.

Since it is not known in advance which scenario will occur, it is interesting to look at the aggregate measures shown in the last two rows of Table 6: the average regret and the maximal regret. These measures relate to the robust optimization approaches introduced in Section 2.3. The last two columns in Table 6 show the regrets and aggregate measures for the two robust approaches. Note that these results were obtained by applying the iterative algorithm to a modified version of the generic MINLP (17), similar to the robust optimization models in Section 2.3. It thus appears that investment plans exist that, over the set of all possible scenarios, perform much better than the single-scenario solutions *sol-1* to *sol-5*. Hence, if there is uncertainty about the model parameters, then it is important not to focus on a single scenario but to consider robust solutions instead.

5 Concluding Remarks

In this paper we considered the dike height optimization problem: what is the optimal dike investment strategy to protect against floods? This is a very important problem in the Netherlands and the government needs to deal with this problem by setting dike ring safety standards. The homogeneous case of this problem is addressed in Eijgenraam et al. (2010). In this paper we consider the extension to the nonhomogeneous case, which is very relevant in practice. This entails that dike rings consist of multiple dike segments, which can be heightened independently of each other. We modeled the nonhomogeneous case as a MINLP model. We also developed an optimization algorithm that reduces the solution time of the problem for large instances. The model and algorithm have been implemented in a software package that is used by the government to determine revised safety standards to be incorporated in the Dutch Water Act.

For the numerical example carried out in Section 4.1, it appears that the safety standards for many dike rings have to be tightened. According to these results, some of the dike rings require immediate heightening. The Dutch government has already reserved additional funds for the required future dike heightenings. Moreover, the numerical experiments show that the nonhomogeneous model yields results that are superior to the results obtained by considering a dike ring to be homogeneous.

It is important to obtain solutions that are robust against uncertainties in model parameters. The model has been further extended to incorporate a robust optimization approach based on the regret criterion. We show that the average and maximum regret can be reduced significantly by applying this robust optimization approach.

The nonhomogeneous dike height optimization problem has been developed for the situation in the Netherlands. It would be interesting to see if the model can be applied to other deltas in the world in which dike rings are nonhomogeneous.

A Convexity Analysis

In this appendix we consider the convexity properties of the exponential investment cost function, the quadratic investment cost function and the expected damage as they have been defined in the MINLP model in Section 2. Additionally, we consider alternative formulations of the MINLP model (9) in Section A.4.

A.1 Exponential Investment Costs

In this section the convexity properties of the exponential investment cost function are considered. Recall the the exponential investment costs of for a positive heightening of segment ℓ at time t_k:

$$\mathcal{I}^e_{\ell k}(\mathbf{u}) = (c_\ell + b_\ell u_k)\exp\bigl(-\lambda_\ell \textstyle\sum_{i=0}^k u_i\bigr), \qquad \mathbf{u} \in \mathbb{R}^{K+1}_+,\ \ell = 1,\dots,L,\ k = 0,\dots,K.$$

Let $c_\ell, b_\ell, \lambda_\ell \geq 0$. It is easy to verify that if $\lambda_\ell = 0$ or $b_\ell = 0$ or $k = 0$, then $\mathcal{I}^e_{\ell k}(\mathbf{u})$ is convex on $\mathbb{R}^{K+1}_+$, otherwise $\mathcal{I}^e_{\ell k}(\mathbf{u})$ is not convex.

A.2 Quadratic Investment Costs

In this section the convexity properties of the quadratic investment cost function are considered. Recall the the quadratic investment costs of for a positive heightening of segment ℓ at time t_k:

$$\mathcal{I}^q_{\ell k}(\mathbf{u}) = \phi_{\ell 0} + \phi_{\ell 1} u_k + \phi_{\ell 2}\bigl(\textstyle\sum_{i=0}^k u_i\bigr)^2, \qquad \mathbf{u} \in \mathbb{R}^{K+1}_+,\ \ell = 1,\dots,L,\ k = 0,\dots,K.$$

Let $\phi_{\ell 2} \geq 0$. It is easy to verify that $\mathcal{I}^q_{\ell k}(\mathbf{u})$ is convex on $\mathbb{R}^{K+1}_+$.

A.3 Expected Damage

Let $\mathbf{t} = (t_0,\dots,t_K)$ denote a discretization scheme. Let $\mathcal{E}_{\ell k}(\mathbf{u},\mathbf{t})$ be the approximation of $\int_{t_k}^{t_{k+1}} S_{\ell t}\exp(-\delta t)dt$, similarly to in Section 2.1, i.e.

$$\mathcal{E}_{\ell k}(\mathbf{u},\mathbf{t}) = \frac{S_{\ell 0}}{\beta_{1\ell}}\exp\bigl(\zeta(H_{\ell^* t_k} - H_{\ell_0 0}) - \alpha_\ell h_{\ell k}\bigr)\Bigl[\exp(\beta_{1\ell} t_{k+1}) - \exp(\beta_{1\ell} t_k)\Bigr],$$

for $\ell = 1,\dots,L,\ k = 0,\dots,K$ with $\mathbf{u} \in \mathbb{R}^{K+1}_+$. Further, let $\mathcal{E}_k(\mathbf{U},\mathbf{t})$ be the approximation of the expected damage in $[t_k, t_{k+1})$ given by

$$\mathcal{E}_k(\mathbf{U},\mathbf{t}) = \max_{\ell\in\mathcal{L}} \mathcal{E}_{\ell k}(\mathbf{u}^{(\ell)},\mathbf{t}),$$

for $k = 0,\dots,K$ with $\mathbf{U} = [\mathbf{u}^{(1)},\dots,\mathbf{u}^{(L)}]^T \in \mathbb{R}^{L\times(K+1)}_+$. Finally, let $\mathcal{E}(\mathbf{U},\mathbf{t})$ be the approximation of the expected damage in $[0,T)$:

$$\mathcal{E}(\mathbf{U},\mathbf{t}) = \sum_{k=0}^K \mathcal{E}_k(\mathbf{U},\mathbf{t}).$$

It is easy to verify that $\mathcal{E}_{\ell k}(\mathbf{u},\mathbf{t})$, $\mathcal{E}_k(\mathbf{U},\mathbf{t})$ and $\mathcal{E}(\mathbf{U},\mathbf{t})$ are all convex in $\mathbf{u}$ and $\mathbf{U}$, respectively.

A.4 Alternative MINLP Formulations

The MINLP model (9) is one way of modeling the nonhomogeneous dike height optimization problem. However, there are alternatives as well. From a modeling point of view, these alternatives are not interesting, though from a numerical point of view, alternative representations of the model could result in a more efficient solution process.

An obvious alternative formulation can be obtained by removing the variables $h_{\ell k}$ and $H_{\ell k}$ and corresponding equality constraints (9d) and (9e) from the model and making appropriate substitutions elsewhere in the model.

A more interesting alternative formulation of MINLP model (9) is obtained by introducing auxiliary variables $I_{\ell k}$, which represent the investment costs for segment ℓ at time t_k, and substitute these in the place of the investment costs in the objective function (9a). Additionally, the following constraints have to be added to the MINLP:

$$
\begin{aligned}
&I_{\ell k} \geq \mathcal{I}^e_{\ell k}(\mathbf{u}^{(\ell)}) - (1 - y_{\ell k}) I^{\max}_{\ell} && \ell = 1, \ldots, L,\ k = 0, \ldots, K, \\
&I_{\ell k} \geq 0 && \ell = 1, \ldots, L,\ k = 0, \ldots, K.
\end{aligned}
$$

If the constant $I^{\max}_{\ell}$ is larger than the highest reasonable investment costs for segment ℓ, then these two constraints ensure that $I_{\ell k} = 0$ if $y_{\ell k} = 0$, and $I_{\ell k} = \mathcal{I}^e_{\ell k}(\mathbf{u}^{(\ell)})$ if $y_{\ell k} = 1$. Of course, this formulation can also be applied to the quadratic investment costs. The advantage of this formulation is that, if the quadratic investment cost function $\mathcal{I}^q_{\ell k}(\mathbf{u}^{(\ell)})$ is used, this incorporates the investment costs as a convex constraint in the model, which renders the entire MINLP model as a convex optimization problem. Note that this is not true if the exponential investment cost function is used, or if the original formulation in the objective value (9a) is used. Convexity of an optimization problem is of great importance for optimization purposes, since it guarantees that a local optimum is a globally optimal solution. The question remains of course whether it is numerically efficient to use this alternative formulation. This question is addressed in Appendix B.1.

B Implementation Details

This appendix discusses some implementation details of the optimization algorithm for the nonhomogeneous dike height optimization problem. In Appendix B.1 we consider the numerical performance of different variations of the MINLP model. In Appendix B.2 we discuss in detail how the discretization scheme is initialized and refined in the iterative optimization algorithm presented in Section 3.4. In Appendix B.3 we discuss how to partition the set of segments for the use of nearly redundant constraints that enforce heightenings of groups of segments at the same time. In Appendix B.4 we demonstrate the performance of the algorithm compared to the original MINLP model (9).

B.1 Comparison of MINLP Formulations

In this section the numerical performances of two different MINLP formulations of the dike height optimization problem are compared. The first MINLP formulation is (9) applied with quadratic investment costs, i.e., with objective function (10) instead of (9a). The second MINLP formulation is the convex alternative formulation of the first formulation as explained in Appendix A.4.

A numerical test is carried out for three dike rings by solving both MINLP formulations for each dike ring. No time limit or relative stop criterion is used, so the problems are solved until, with certainty, AOA cannot find an improvement of the best solution found. For convex problems this implies that a global optimum has been found.

The solutions for all three dike rings, in particular the objective values and corresponding investment plans, obtained by both MINLP formulations were found to be completely identical. Hence, the solutions obtained by the first MINLP are identical to the solutions obtained by the second MINLP, which we know to be the global optima of the underlying problems. It cannot be proven that this will always be the case, and moreover, in practice another stop criterion will be used. However, the solution times of the first MINLP are much better than the solution times of the second MINLP. The solution times of the second MINLP were 2 to 80 times longer than of the first MINLP. Based on this result the first MINLP formulation is preferred, i.e., the investment costs are directly included into the objective function.

B.2 Implementation Discretization Scheme

In Appendix B.2.1 a method is proposed that can be used to select an initial discretization scheme. Subsequently, Appendix B.2.2 considers how an existing discretization scheme can be refined, for the use in an iterative solution algorithm.

B.2.1 Initialization of Discretization Scheme.

The decision for a discretization scheme can be divided into three different steps that, to some extent, can be considered independently of each other:

(i) selecting the length of the planning horizon T,

(ii) selecting the number of time intervals K, and

(iii) selecting the interval sizes $t_k - t_{k-1}$ for $k = 1, 2, \ldots, K + 1$.

The first logical step in determining the discretization scheme is to select the length of the planning horizon T. If we consider K fixed, then increasing T implies that the average interval size increases, and consequently approximation accuracy of the expected damage, given by (4), diminishes. Another reason for choosing a value of T that is not too large is that, due to discounting, the effect of a decision becomes smaller as the time span of the decision becomes larger. On the other hand, if we choose T small, then the approximation of expected damage after T, given by (5), becomes inaccurate or simply too large. Numerical experiments have shown that, for the application to dike rings in the Netherlands, $T = 300$ yields an approximation of the expected damage after T that is satisfactory in the sense that it does not affect the first part of the investment plan. This is important, because in the end, this first part is used to determine the new safety standards. Using a smaller planning horizon can result in significant approximation errors, which can result in a different investment plan. On the other hand, a larger planning horizon does not offer many benefits.

The second step is to choose the number of decision moments. The value of K is an important factor regarding the number of decision variables. The investment plan is characterized by $L(K + 1)$ continuous decision variables, but additionally, we also need $L(K + 1)$ binary decision variables and several auxiliary decision variables in the MINLP. Especially the binary variables have a strong impact on the solution time of a MINLP solver. The MINLP's solution time explodes if the number of discrete decision variables increases, making it important to choose the value of K cautiously, especially for problems with multiple segments. The fact that the combination of the number of dike segments L and the number of decision moments $K + 1$ determines the number of decision variables of the MINLP forces us, for reasons of computation time, to use different values of K for dike rings that have a different number of dike segments. If, as a rule of thumb, we take the number of integer decision variables as a starting point and try to fix this at B, then the value of K can be determined by

$$K = \left\lceil \frac{B}{L} \right\rceil.$$

The third and final step in selecting a discretization scheme is choosing the interval sizes $t_k - t_{k-1}$. The most obvious choice is to choose equidistant intervals, i.e.,

$$t_k = \frac{kT}{K + 1}, \qquad k = 0, 1, \ldots, K + 1.$$

The MINLP formulation does not require the possible dike heightening moments t_k to be equidistant. Since costs are discounted, the impact of decisions at the start of the planning horizon is higher than decisions at the end of the planning horizon. This property can be exploited by choosing smaller intervals at the start of the planning horizon and larger intervals at the end, thereby saving decision variables towards the end of the planning horizon. A possible approach that uses this idea is to divide the planning horizon into two parts: $[0, T_1)$ and $[T_1, T)$, and subsequently to assign a fraction φ of the $K + 1$ integer decision variables per dike segment to the first part, and a fraction $1 - \varphi$ to the second part.

There are various alternative methods to define a good discretization scheme, especially regarding the last two steps of the process described here. For example, the values of t_k could be restricted to integer values for reasons of convenience. Considering the time span of the dike height optimization problem, this is not a serious limitation.

B.2.2 Refining an Existing Discretization Scheme.

The main reason for not choosing small interval sizes is that it requires a large value of K and consequently many decision variables and a large solution time. The consequence is that approximation (4) can be less accurate and that the flexibility in selecting the timing of segment heightenings is limited. One option to deal with this issue is to generalize the concept of a discretization scheme in relation to MINLP (9) and to use multiple discretization schemes for the same dike height optimization problem.

In MINLP (9), K determines the number of decision variables $u_{\ell k}$ and $y_{\ell k}$, but on the other hand K also determines the number of auxiliary variables $I_{\ell k}$ and E_k. This relationship does not necessarily have to be like this. Recall that especially the number of binary variables $y_{\ell k}$ is an important factor for the MINLP's solution time. If K is kept fixed, then the number of binary variables can be reduced by enforcing $y_{\ell k} = 0$ for an appropriately chosen subset F of $\{0, 1, \ldots, K\}$, which can be realized by simply substituting $y_{\ell k} = 0$ for $k \in F$ in MINLP (9). Another way of looking at this is that the $y_{\ell k}$, $k \in F$, are demoted from decision variables to parameters. Notice that, due to constraint (9f), also the substitution $u_{\ell k} = 0$, $k \in F$ can be carried out, thereby reducing the number of decision variables even further. The advantage of this approach is that the approximation of the expected damage does not suffer from the reduction in the number of decision variables.

This generalization of MINLP (9) can be exploited in an iterative approach. Suppose a solution has been obtained by using an equidistant discretization scheme with ten year intervals. Assume that in the resulting investment plan, for a certain segment the first heightening takes place at $t = 50$ and the second at $t = 150$. It is, intuitively, very unlikely that exactly the same timing of segment heightenings would have been found for the discretization scheme with one year intervals, since this discretization scheme offers much more flexibility in investment plans. However, the obtained solution gives a good idea of the timing of the dike heightenings that might have been obtained by using this "ideal" discretization scheme. It is, for instance, very unlikely that a dike heightening at $t = 95$ would have been obtained. More specifically, we expect to obtain dike heightenings in the neighborhood of the dike heightenings found in the selected discretization scheme. This idea can be used to refine the initial discretization scheme. Refining an existing discretization scheme is done by (1) adding decision moments to the discretization scheme, and (2) fixing the decision variables for the decision moments in a subset of the resulting discretization scheme.

First, we discuss adding decision moments to a discretization scheme. Suppose that a dike heightening takes place at decision moment t_k. Now consider the interval (t_{k-1}, t_{k+1}). We want to add decision moments in this interval to be able to explore other solutions in the neighborhood of the first solution. Since we are not necessarily interested in adding decision moments close to t_{k-1} or t_{k+1}, two intervals are defined using a fraction ξ as shown in Figure 5. Note that the parameter ξ is not necessarily ≤ 1. If $\xi > 1$, then the new intervals stretch beyond the existing decision moments t_{k-1} and t_{k+1}. Moreover, contrary to the example, the interval sizes $t_k - t_{k-1}$ and $t_{k+1} - t_k$ are not necessarily equal depending on the initial discretization scheme.

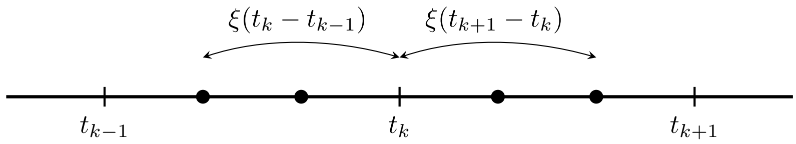

Figure 5: Refinement of decision moments.

Now it has to be decided how many decision moments will be added to both intervals. This is determined by the number Q, and the additional decision moments will be added using equal interval sizes.

For example, if $\xi = 2/3$ and $Q = 2$, then this yields the additional decision moments marked by the dots in Figure 5. Another sensible parameter choice is $\xi = 1/2$ and $Q = 1$, which yields exactly one additional decision moment half-way in between the existing decision moments.

Finally, once a new discretization scheme has been defined, then it makes sense to use the generalization of MINLP (9) and fix the decision variables outside the intervals of interest equal to zero. This has a positive effect on the solution time since the number of decision variables decreases.

B.3 Selecting Segment Partitions

In Section 3.3.2 we propose to use nearly-redundant constraints that forces heightenings of a certain group of segments to occur simultaneously. Here we discuss how these groups can be defined.

Even though it is easy to formulate constraint (16), it is less obvious how to choose the subset $\mathcal{G}$. Moreover, if one group of segments with similar characteristics is selected, then the remaining segments often have something in common as well. We would therefore like to impose constraint (16) on this group as well, especially if this group is large. It is of course possible to partition $\mathcal{L}$ into more than two subsets, but unless there is a compelling reason to do so, based upon specific expert dike ring knowledge, it is almost impossible to provide general guidelines for selecting these partitions.

If we restrict ourselves to partitioning $\mathcal{L}$ into two subsets, $\mathcal{G}_1$ and $\mathcal{G}_2$, for which constraint (16) will be imposed, then there still are

$$\sum_{n=0}^{\lfloor L/2 \rfloor} \binom{L}{n}$$

ways of doing so. For dike rings with more than four dike segments, this number is simply too large to try all possibilities. The number of partitions might be reduced, however, if we are able to select promising partitions from the total set.

One particular useful partition is $\mathcal{G}_1 = \mathcal{L}$ and $\mathcal{G}_2 = \emptyset$, which represents the special case (15). Another set of interesting partitions are the partitions that allow an investigation of whether a single segment ℓ should be heightened simultaneously with the others or not. These are the partitions

$$\mathcal{G}_1 = \mathcal{L} \setminus \{\ell\} \quad \text{and} \quad \mathcal{G}_2 = \{\ell\}, \quad \ell = 1, \ldots, L. \tag{18}$$

Note that, in this case, $\mathcal{G}_2$ does not yield any additional constraint according to (16), since $\mathcal{G}_2$ has cardinality 1.

If L is large, then there are many partitions left with more than one segment in both subsets. A good selection method for these partitions may be derived by important segment properties that can be ordered. For example, consider the flood probability. One of the most important reasons for dike segments to have non-simultaneous dike heightenings is the varying development of flood probability over time. The increase in flood probability for a certain dike segment can be much slower than the increase for other segments, as a result of which the dike segment has to be heightened less often than the other segments. The development of a segment's flood probability is partly dependent on $\alpha_\ell \eta_\ell$, and hence it might be worthwhile to select the partition of $\mathcal{L}$ on this basis. Let $\sigma : \mathcal{L} \to \mathcal{L}$ be a permutation of $\mathcal{L}$ such that

$$\alpha_{\sigma(\ell)}\eta_{\sigma(\ell)} \leq \alpha_{\sigma(\ell+1)}\eta_{\sigma(\ell+1)}, \qquad \ell = 1, 2, \ldots, L-1. \tag{19}$$

The subsets of $\mathcal{L}$ for which (16) will be included in the MINLP based upon this idea are then given by

$$\mathcal{G}_\ell^1 = \{\sigma(1), \ldots, \sigma(\ell)\} \text{ and } \mathcal{G}_\ell^2 = \{\sigma(\ell+1), \ldots, \sigma(L)\}, \quad \ell = 2, \ldots, L-2. \tag{20}$$

Another interesting permutation of $\mathcal{L}$ can be derived by considering the investment costs. For example, if two dike segments have equal properties except that one has much higher fixed investment costs, then it might be optimal for this segment to be heightened less frequently but more comprehensively than the other segment. A typical example is the situation where a flood gate in a dike ring is modeled as a separate segment. Heightening a flood gate usually involves extremely high fixed costs in comparison with the variable costs. If only a single flood gate, or another extreme cost-characteristic segment, is present, then this situation is already tackled by the individual partition of single segments such as in (18). However, if more than one of these segment types are present, then a good approach is to partition the segments based upon the quotient of the investment costs and its derivative. Let

$$F_\ell(u) = \left. \frac{\mathcal{I}_{\ell 1}(\mathbf{u})}{\frac{\partial}{\partial u_1} I(\mathbf{u})} \right|_{\mathbf{u}=(u,0,\ldots,0)}$$

be this quotient, and sort the segments according to $F_\ell(u)$ for, say $u = 100$cm. Partitions of $\mathcal{L}$, based upon this ordering, can now be derived similar to (20).

B.4 Optimality Iterative Algorithm

Recall that, ideally, we would like to solve the MINLP (9). Unfortunately, this problem's solution time can explode very suddenly as K increases for dike rings with many segments, making the use of MINLP (9) not very suitable in practice. For selected problems it is of course possible to compare the performance of the iterative algorithm defined in Section 3.4 to MINLP (9). Table 7 shows the results for dike ring 16 obtained with three different methods. The first method is the iterative algorithm, which has been applied with two initial discretization schemes (stage 1): one with $K = 14$ and one with $K = 25$ as indicated by the second column of Table 7. In the last step of the iterative algorithm, a refined discretization scheme (stage 2) is used and the corresponding value of K and the cardinality of the associated set F are shown in the third and fourth column respectively. The fifth column shows the solution time and the last two columns give the MINLP objective as well as the true objective. As might be expected, the solution time increases if a finer discretization scheme is used. The objective values for the initial discretization schemes with $K = 14$ and $K = 25$ are a good example why investment plans should be compared using the true evaluation. The true objective for $K = 25$ is lower than the objective for $K = 14$, albeit less than 0.1%, and since the true objective is based on a fixed and unique discretization scheme we can truly say that this investment plan is the best of the two. The two MINLP objectives give an opposite perception. However, since these two objectives result from different discretization schemes, they cannot be adequately compared.

Method	K (init)	K (refine)	$\#F$ (refine)	Solution time (min)	MINLP objective (M€)	True objective (M€)
Iterative algorithm	14	25	8	2.9	1027.7735	1029.3511
Iterative algorithm	25	34	20	21.6	1027.8737	1028.5856
Single & refine	14	25	8	434.1	1027.7735	1029.3511
Single & refine (forced)	14	23	9	0.2	1027.7455	1029.3472
Single & refine (forced)	25	34	20	1.5	1027.8341	1028.5409

Table 7: Results for dike ring 16 obtained with different methods.

The second method shown in Table 7, indicated by "single & refine", is the use of MINLP (9) with an initial and a refined discretization scheme similar to the discretization schemes in our iterative algorithm. In particular, constraints (17d), (17e) and (17f) are not used in this method. For the initial discretization scheme with $K = 14$, this gives the same final investment plan as the iterative algorithm but with a solution time of 434.1 min. instead of 2.9 min. This method's solution time clearly indicates that the nearly-redundant constraints used in the iterative algorithm have a significant impact on the model's solution time.

As a third method we have also included the extension of the previous method that forces segments to be heightened simultaneously, i.e. with constraint (15) added to the optimization model. This basically boils down to using only the first and last iteration of the iterative algorithm and ignoring all the intermediate iterations. By now it will have become clear that for dike ring 16, constraint (15) really is a redundant constraint and can safely be included without affecting the optimal solution. In general, this is not known beforehand and therefore is not applicable in practice. However, if this information is available, then it may deliver quite substantial savings in solution time.

It cannot be guaranteed that, in general, the iterative algorithm yields the global optimal solution. However, the results from Table 7 indicate that it is more beneficial to apply the iterative algorithm with a finer discretization scheme than to ensure that a global solution has been found for a rougher discretization scheme.

Acknowledgment: We wish to thank all members of the project team "Optimal safety standards" for offering us their very useful knowledge and expertise. We are particularly grateful to Jarl Kind, Carlijn Bak (both Rijkswaterstaat and Deltares) and Matthijs Duits (HKV) for their extensive support. We would like to emphasize that this project was a multi-disciplinary project and that our model relies heavily on many research results carried out by the following organizations: CPB Netherlands Bureau for Economic Policy Analysis, Rijkswaterstaat, Deltares, HKV, KNMI, and Ministry of Infrastructure and Environment. Without them, this project would not have been possible.

References

Ben-Tal, A., L. E. Ghaoui, A. Nemirovski. 2009. *Robust Optimization*. Princeton Series in Applied Mathematics, Princeton University Press.

Brekelmans, R., D. den Hertog, C. Roos. 2009. Computing safe dike heights: A prototype optimization algorithm. CAR report, CentER Applied Research, Tilburg University, The Netherlands.

Bussieck, M. R., S. Vigerske. 2011. MINLP solver software. *Wiley Encyclopedia of Operations Research and Management Science* .

Dantzig, D. van. 1956. Economic decision problems for flood prevention. *Econometrica* **24**(3) 276–287.

Duits, M. T. 2009a. OptimaliseRing. Technische documentatie van een numeriek rekenmodel voor de economische optimalisatie van veiligheidsniveaus van dijkringen. Versie 2.0. HKV rapport PR1377, HKV LIJN IN WATER. In Dutch.

Duits, M. T. 2009b. OptimaliseRing. Testrapport. Versie 2.0. HKV rapport PR1377, HKV LIJN IN WATER. In Dutch.

Duran, M., I. Grossmann. 1986. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Mathematical Programming* **36**(3) 307–339.

Eijgenraam, C. J. J. 2005. Veiligheid tegen overstromen. kosten-baten analyse voor ruimte voor de rivier, deel I. CPB Document 82, CPB Netherlands Bureau for Economic Policy Analysis, The Hague. In Dutch.

Eijgenraam, C. J. J. 2006. Optimal safety standards for dike-ring areas. CPB Discussion Paper 62, Netherlands Bureau for Economic Policy Analysis, The Hague.

Eijgenraam, C. J. J. 2007a. From optimal to practical safety standards for dike-ring areas. *Water Science & Technology* **56**(4) 113–124.

Eijgenraam, C. J. J. 2007b. Oplossing van het veiligheidsprobleem bij meer dan één dijkvak per dijkring. Tech. rep., CPB Netherlands Bureau for Economic Policy Analysis, The Hague. In Dutch.

Eijgenraam, C. J. J., R. Brekelmans, D. den Hertog, C. Roos. 2010. Safe Dike Heights At Minimal Costs: The Homogeneous Case. Working paper, Tilburg University.

Hertog, D. den, C. Roos. 2009. Computing safe dike heights at minimal costs, Second version. CAR report, CentER Applied Research, Tilburg University, The Netherlands.

Silva, W., J. W. Stijnen. 2005. *Rampenbeheersingsstrategie overstromingen Rijn en Maas: Achtergrondrapportage veiligheid en rivierkunde*. No. 2005.024 in RIZA rapport, Rijkswaterstaat, RIZA. In Dutch.

Stijnen, J. W., C. A. H. Wouters, E. O. F. Calle, H. Knoeff, A. C. W. M. Vrouwenvelder. 2006. Vergelijking leidraad waterkeringen. HKV rapport PR 1176.10, HKV lijn in water. In Dutch.

Tawarmalani, M., N. V. Sahinidis. 2002. *Convexification and Global Optimization in Continuous and Mixed-Integer Nonlinear Programming: Theory, Algorithms, Software, and Applications*. Kluwer Academic Publishers.