Optimal Adaptation to Uncertain Climate Change

Graeme Guthrie
graeme.guthrie@vuw.ac.nz

September 10, 2022

Abstract

Enormous public investment will occur as communities adapt to climate change. Much of this investment will be irreversible and the future benefits are currently uncertain. The real options embedded in adaptation projects are therefore potentially important and their existence needs to be incorporated into investment decision-making. Standard real options analysis is inadequate for this purpose because the future arrival of information about climate change is unlikely to conform to the stochastic processes typically used in real options analysis. This paper presents a new framework that reflects current uncertainty about climate change and how that uncertainty might change over time. Optimal investment depends on current beliefs regarding the severity of future climate change, how quickly these beliefs are likely to change in the future, and current economic conditions. Most of the net benefits of optimal investment can be captured if investment timing is decided using a simple alternative decision-making rule.

Keywords: climate change; adaptation; cost–benefit analysis; real options analysis

This research is funded by the Deep South Challenge as part of its Living With Uncertainty programme.

Find more at https://sites.google.com/view/graemeguthrie

Optimal Adaptation to Uncertain Climate Change*

Graeme Guthrie†

September 10, 2022

Abstract

Enormous public investment will occur as communities adapt to climate change. Much of this investment will be irreversible and the future benefits are currently uncertain. The real options embedded in adaptation projects are therefore potentially important and their existence needs to be incorporated into investment decision-making. Standard real options analysis is inadequate for this purpose because the future arrival of information about climate change is unlikely to conform to the stochastic processes typically used in real options analysis. This paper presents a new framework that reflects current uncertainty about climate change and how that uncertainty might change over time. Optimal investment depends on current beliefs regarding the severity of future climate change, how quickly these beliefs are likely to change in the future, and current economic conditions. Most of the net benefits of optimal investment can be captured if investment timing is decided using a simple alternative decision-making rule.

Keywords: climate change; adaptation; cost–benefit analysis; real options analysis

JEL classification numbers: G31, Q54, R42

1 Introduction

Adapting to climate change will require substantial investment of public funds over the coming decades. Investment of this scale creates enormous potential for inefficient decisions involving the timing, scale, and location of investment. Fortunately, modern decision-support tools can potentially improve the efficiency of this investment. In particular, typical climate-change adaptation projects have the characteristics needed for real options analysis to be important: future benefits are uncertain; reversing investment is costly, if it is possible at all; and decisions can be spread over time as uncertainty is resolved. This paper presents a new real-options framework for analysing investment decisions related to climate-change adaptation and demonstrates its application using the example of a decision-maker tasked with upgrading an urban stormwater system in response to climate change.

The paper makes three main contributions. First, it introduces a theoretical framework that incorporates climatic and economic volatility in a single model. The main features of this framework are the way that it models uncertainty about future climate change, how this uncertainty evolves over time, and how this injects volatility into investment payoffs. Second, the paper constructs an optimal investment policy for an adaptation problem in which the capacity of infrastructure to cope with extreme weather events is a continuous variable and the decision-maker decides how quickly to increase capacity. The value of investment delay options is economically significant, but the relative importance of climatic and economic volatility varies over time. Third, the paper shows that almost all of the net benefits of investment can be captured if

*This research is funded by the Deep South Challenge as part of its Living With Uncertainty programme.
†Email: graeme.guthrie@vuw.ac.nz.

investment timing is decided using a simple alternative to "full" real options analysis. If the decision-maker uses this alternative rule, then she invests if and only if the present value of the project's benefits is greater than or equal to the sum of the investment expenditure and a specific approximation of the value of delaying investment. If investment decisions are made using this alternative rule, then—for the parameterisation used here—almost all of the net benefits of investment are received.

Turning to the first contribution, there is currently considerable uncertainty surrounding future climatic conditions (Pindyck, 2021). Climate change happens slowly, we do not know how much CO_2 and other greenhouse gases will be emitted over the coming decades, and we are uncertain how the climate will change as a result. However, this uncertainty will fall over time due to a combination of scientific progress, statistical evidence, and observing weather patterns (van der Pol et al., 2017b). The arrival of new climate information will inject volatility into forecasts, but recent research suggests that meaningful learning about some important aspects of climate change will take 20–50 years to occur (Urban et al., 2014; Lee et al., 2017).[1] Intra-regime variability means that learning about some aspects of climate change will be even slower (Kopp et al., 2014). The theoretical framework introduced in this paper models this uncertainty by assuming that the outcome of climate change will match one of two possible scenarios. The decision-maker observes a stream of information that she uses to continuously update her subjective beliefs regarding the probability of each scenario holding. This injects volatility into the payoff from delaying investment in an adaptation project. The arrival of new climate information is slow at first, so this source of volatility is initially small, but accelerates as the two climate scenarios diverge.

The use of Bayesian decision theory has a long history in environmental decision-making. For example, in a paper that preceded the literature on real options analysis, Davis et al. (1972) explicitly consider the option to delay choosing the capacity of a system of dikes in order to gather additional data on actual river flows. Delaying investment allows the decision-maker to update her beliefs regarding the parameters that affect the probability distribution of peak annual river flows used to calculate the expected costs of flood damage. A small but growing literature adopts a similar approach to issues surrounding climate change. For example, van der Pol et al. (2014) analyse investment in flood-protection infrastructure using a model in which uncertainty about the rate of water-level increase is completely resolved at a single future date. In contrast, van der Pol et al. (2017a) assume that uncertainty is reduced, but not eliminated, at a single future date. De Bruin and Ansink (2011) go one step further and use a model in which uncertainty about the impact of climate change is eliminated in two steps, whereas Guthrie (2019) assumes uncertainty falls gradually over time. One feature common to all of these papers is that climate-induced volatility in investment payoffs comes from changes in *beliefs* about some unobservable aspect of climate change. That is, volatility comes from the future arrival of new information about the extent of climate change. This is a key feature of the approach adopted here.

The paper's second contribution is the construction of an optimal policy for increasing an asset's capacity to cope with extreme weather events. The benefits of investment depend on the distribution of the severity of these weather events, which vary as more climate information becomes available. Initially, climatic uncertainty falls slowly, so that most of the volatility in the investment benefits results from fluctuations in an economic factor. However, climatic uncertainty eventually begins to fall, which increases the volatility of the investment benefits. Thus, the value of the option to delay investment varies over time, depending on the rate with which new climate information arrives. The value of the delay option contains a second time-dependent component, reflecting the trajectories of key climate variables under climate change.

[1] Woolley et al. (2020, p. 173) observe that projected sea-level rise in the near term for the four RCP scenarios is so close that it may require another two decades of monitoring to conclusively determine which trajectory applies.

This induces relatively fast growth in benefits in the short term, which slows down when the climate approaches its new equilibrium. The optimal investment policy can be expressed in several forms, including a threshold for the benefit–cost ratio. This threshold exceeds one by a substantial margin, reflecting the value of the delay option. The threshold is initially high, especially when the decision-maker believes climate change is going to be relatively severe, but falls over time due to the declining growth rate in benefits. This is offset by the increased rate with which new climate information arrives, which increases the investment threshold.

These results originate from the stochastic properties of the investment payoff, which ultimately stem from the model of climate uncertainty. The paper thus contributes to the part of the real options literature that investigates the implications of non-standard stochastic structures. Most real options models assume that one or more exogenous state variables evolve according to geometric Brownian motion, but some authors have explored different structures, especially when considering investment in climate-change adaptation. For example, Truong et al. (2018) assume bushfires occur with a Poisson intensity that is itself stochastic. In contrast, Guthrie (2019) assumes the decision-maker learns about the underlying climate scenario by observing the frequency of extreme weather events, which injects volatility into the payoff from delaying investment. Most recently, Lee and Zhao (2021) build a real options model of climate-change adaptation that combines a Brownian motion process with extreme events, represented by Poisson jumps with a hyper-exponential jump size distribution.

Finally, the paper's third contribution is to show that most of the benefits of investment can be captured using a simple alternative to "full" real options analysis. This is important because, although the benefits from using real options analysis can be significant, it is regarded by many practitioners as complex and resource-intensive (Dawson et al., 2018; Wreford et al., 2020). It is not widely used, particularly when evaluating relatively small projects. This is serious because much adaptation spending will fund relatively small projects under the jurisdiction of authorities with limited analytical resources. The importance of the flexibility in these projects is too costly to ignore. Decision-makers need alternative approaches that are simple enough to be useful for evaluating small- and medium-scale adaptation decisions, yet retain a degree of economic rigour, as recommended in recent reviews of decision-support tools for adaptation assessment (Watkiss et al., 2015; Wreford et al., 2020). This paper provides one such approach.

Optimal investment involves comparing the present value of the benefits from investing and the total cost of investing, which equals the sum of the investment expenditure and the present value of the real option to wait and invest some time in the future. The alternative rule involves replacing the "fully optimal" value of the delay option with its value assuming investment is delayed until the best *fixed* future date. This is much simpler to calculate than the full option value, because all that is required is a series of standard static cost–benefit calculations, one for each possible future investment date. The approximate option value is the maximum of these values. If the decision-maker uses this alternative rule, then she invests if and only if the present value of the project's benefits is greater than or equal to the sum of the investment expenditure and the approximate option value. Equivalently, she invests if no fixed future investment date implies a greater net present value than investing immediately.

Under the alternative policy, investment is delayed significantly past the date when the present value of the benefits from investment equals the required expenditure. Nevertheless, it still occurs earlier than under the welfare-maximising investment policy. However, for the next few decades at least, the acceleration in investment is moderate. The calculation underlying the alternative investment policy assumes the decision-maker will wait and invest at a fixed future date, rather than exploiting future climatic and economic information as it arrives. That is, the calculation underlying the alternative investment policy assumes the decision-maker delays the investment, but not the decision when to invest—that is made now. This only introduces substantial errors if there is a significant risk of new information arriving that would cause the decision-maker to regret the investment decision that is implicit in the approximate option value. For the situation we are examining, the prospect of such bad news is low, at least for the foreseeable future. Due to the relatively low volatility of the economic factor and—for the foreseeable future—the low volatility of the decision-maker's beliefs regarding the magnitude of climate change, there is very little chance that any subsequent information will be significant enough to make the net present value of investing negative. In other words, if the decision-maker uses the alternative rule, then she delays investment so long that there is little chance that she will subsequently regret investing. It is not surprising that the added value from delaying the *decision* is so small.

For at least the first 40 years, the welfare losses from adopting the alternative rule are small. The welfare performance of the alternative investment rule is worse if the economic factor has a lower expected growth rate or higher volatility. These situations increase the likelihood that past investments will subsequently be stranded, which increases the importance of delaying the investment decision and reduces the performance of the alternative investment rule. If there is more noise in the climatic signal, and therefore less volatility in the decision-maker's beliefs, the probability of post-investment stranding is smaller, which improves the performance of the alternative investment rule. The effect is insignificant in the short term, because in this case the climate scenarios are so similar. However, for longer horizons, there is a noticeable improvement in the alternative rule's performance when the climatic signal is noisier. Although the alternative investment rule performs reasonably well over a wide range of parameter values, there will likely still be situations where it results in a significant reduction in welfare. In particular, if the economic factor is volatile enough and grows slowly enough, and if there is enough noise in the climatic signal, then there is still a place for "full" real options analysis.

This paper's final contribution is therefore to the subset of the real options literature that tries to link the technique to real world decision-making. Some papers in this literature attempt to derive forms of real options analysis that are simple enough to be useful in practice. For example, Copeland and Antikarov (2003, 2005) have developed a "marketed asset disclaimer" approach that is popular among some corporate-finance practitioners of real options analysis. The remaining papers in this literature attempt to interpret pragmatic decision-making techniques as tractable approaches that (deliberately or otherwise) capture the value of the real options embedded in projects. For example, McDonald (2000) and Malchow-Møller and Thorsen (2005) evaluate various rules of thumb from a real options perspective, whereas Boyle and Guthrie (2006) and Wambach (2000) concentrate on the payback method.

The rest of the paper is organised as follows. Section 2 presents the model set-up. The theoretical framework has two main components: a description of how beliefs about future climate change evolve over time, which is developed in Section 3, and a description of how the benefits of investment are calculated, which is developed in Section 4. The analysis of the model begins in Section 5, which derives the conditions that an optimal investment policy must satisfy. It continues in Section 6, which investigates these optimal policies, and Section 7, which compares the performance of these policies with a simpler alternative. Finally, Section 8 offers some concluding remarks. All proofs are contained in an appendix.

2 Model set-up

A decision-maker is responsible for upgrading an urban stormwater system. The system's capacity at any given date t equals q_t, which the decision-maker is able to increase over time by investing in system upgrades. At any point in time, she can increase capacity by any positive amount Δq by spending $c \cdot \Delta q$, for some constant $c > 0$. There is no limit to the number of times that investment can occur, but increases in capacity are irreversible. The decision-maker's objective is to minimise the present value of all relevant costs, including flooding costs and the investment expenditure incurred when upgrading the system. The social discount rate is the constant r.

There are two possible climate scenarios, "good" and "bad." At any date t, the expected cost associated with flooding over the next increment of time equals $\Gamma_g(q_t, t)x_t\,dt$ if the good scenario holds and $\Gamma_b(q_t, t)x_t\,dt$ if the bad scenario holds, where Γ_g and Γ_b are decreasing, convex functions of q. That is, extra capacity reduces expected flooding costs, but the amount of the reduction becomes smaller as q becomes larger. The economic factor x_t, which is observable, evolves according to the geometric Brownian motion

$$dx_t = \mu x_t\,dt + \sigma x_t\,d\xi_t,$$

where μ and σ are known constants and ξ_t is a Wiener process. This variable determines the expected economic consequences of flooding for any given capacity of flood defences. It reflects the depreciated replacement cost of the vulnerable assets, clean-up costs, disruption costs, environmental damage, and health effects.

The decision-maker does not know which climate scenario holds. Instead, she assigns a probability to each scenario being the true one and then updates these probabilities as more information about the state of the climate arrives. The expected flooding cost between dates t and $t + dt$, as measured by the decision-maker, therefore equals

$$\big(p_t\Gamma_b(q_t, t) + (1 - p_t)\Gamma_g(q_t, t)\big)\,x_t\,dt,$$

where p_t denotes the date t probability that the bad climate scenario holds. The new climate information takes the form of a noisy signal of the average annual maximum one-day rainfall λ_t.[2] At any given date t, λ_t takes the values $\lambda_g(t)$ and $\lambda_b(t)$ in the good and bad scenarios, respectively, for some functions λ_g and λ_b. The decision-maker observes the path of a variable, z_t, which evolves according to

$$dz_t = \lambda_t\,dt + \theta\,d\zeta_t, \tag{1}$$

where θ is a constant and ζ_t is a Wiener process that is uncorrelated with ξ_t.[3]

3 Processing future climate information

In equation (1), λ_t equals $\lambda_b(t)$ if the bad scenario holds and $\lambda_g(t)$ if the good scenario holds. The information contained in the observed path of z_t allows the decision-maker to update her beliefs regarding the climate scenario. Over the next short period of time lasting Δt years, the decision-maker observes

$$\frac{\Delta z_t}{\Delta t} \approx \lambda_t + \frac{\theta\varepsilon}{\sqrt{\Delta t}},$$

where $\varepsilon \sim N(0, 1)$. That is, she observes the sum of the unobservable value of the climate factor, λ_t, and a noise term that is (approximately) normally distributed with mean zero and variance $\theta^2/\Delta t$. If θ is relatively small, then the noisy signal is relatively precise and the observed value of $\Delta z_t/\Delta t$ gives a relatively accurate estimate of λ_t. In this case, provided that $\lambda_b(t)$ and $\lambda_g(t)$ are not too similar, the noisy signal will give the decision-maker a relatively clear indication of which climate scenario actually holds.

The next lemma describes the evolution of the decision-maker's beliefs.[4]

[2]In principle, the actual flooding costs reveal additional information about the climate scenario. We ignore this information here, mainly in the interests of tractability but also because the actual flooding cost is potentially a very noisy indicator of the expected flooding costs. We are implicitly assuming that the noise here is so great that the information content is negligible.

[3]This significantly simplifies the analysis. For example, if the noise term $d\zeta_t$ were correlated with the shock that influences the path of the economic factor, $d\xi_t$, then the decision-maker would need to incorporate the information content of growth in the economic factor when updating her subjective probability that the bad scenario holds.

[4]Appendix A.1 contains a heuristic proof of this result. See Harrison (2013, pp. 138–141) for more on Brownian models of dynamic inference.

Lemma 1 *The subjective probability that the bad climate holds evolves according to*

$$dp_t = \frac{p_t(1-p_t)(\lambda_b(t)-\lambda_g(t))}{\theta} d\zeta_t. \tag{2}$$

■

As Lemma 1 makes clear, the decision-maker's beliefs regarding which climate scenario holds evolve stochastically. Uncertainty surrounding future values of the subjective probability that the bad scenario holds will turn out to be an important source of volatility and a crucial determinant of the option value embedded in typical adaptation projects. Equation (2) shows that p_t will be most volatile when the identity of the actual climate scenario is most uncertain, because the coefficient $p_t(1-p_t)$ is maximised when $p_t = 1/2$. In this situation, the new information contained in the path of z_t has the most effect on p_t. In contrast, when the decision-maker is almost certain the good scenario holds ($p_t \approx 0$) or she is almost certain the bad scenario holds ($p_t \approx 1$), new information will have little effect on her beliefs, so that the volatility of p_t will be small. Equation (2) also shows that p_t will be more volatile when θ is smaller; that is, when the signal contains little noise, so that the decision-maker can rely on the information contained in the observed path of z_t to update her beliefs. Finally, if the difference between $\lambda_b(t)$ and $\lambda_g(t)$ is smaller then p_t will be less volatile. In this case, even a relatively precise signal may be unable to distinguish between the two possible values of the climate factor, so the decision-maker will be unable to infer much about the identity of the underlying climate scenario.

Consider the situation facing Dunedin City Council, New Zealand, which is planning to spend $35 million in the next ten years upgrading its stormwater system in order to reduce flooding during times of heavy rain (Dunedin City Council, 2019, 2022).[5] Suppose the annual maximum one-day rainfall for the area is drawn from a generalised extreme value distribution and that the good and bad scenarios correspond to RCP2.6 and RCP8.5, respectively. In order to focus attention on the key issues, further suppose that only the location parameter of this distribution varies over time and across scenarios; the scale and shape parameters are constant and take the same values in the two scenarios. The initial distribution matches the estimated distribution generated by the NIWA High Intensity Rainfall Design System (HIRDS).[6] Its location, scale, and shape parameters equal 75.71, 22.54, and 0.1908, respectively. If we allow the location parameter to increase over time to match the means of the distributions generated by HIRDS under RCP2.6 and RCP8.5. then the paths for the average annual maximum one-day rainfall are

$$\lambda_b(t) = 132.95 - 39.03e^{-0.0083t}$$

and[7]

$$\lambda_g(t) = 98.03 - 4.12e^{-0.0788t}.$$

The blue curve in Figure 1 plots the current density function of the annual maximum one-day rainfall. The yellow and green curves plot the density functions in 2095 in the good and bad scenarios, respectively. Climate change shifts the distribution rightwards, with the distribution shifting further in the bad scenario than in the good scenario. For example, Dunedin's average annual maximum one-day rainfall is currently 94mm. Under RCP2.6 and RCP8.5, this increases to 98mm and 112mm, respectively, in 2095. Figure 2 plots the two possible paths for the average annual maximum one-day rainfall $\lambda(t)$. The two trajectories are reasonably similar in the short

[5]Hughes et al. (2019) evaluate the impact of climate change on New Zealand stormwater systems.

[6]Extreme rainfall projections for any New Zealand location can be viewed at https://hirds.niwa.co.nz/. Data from the "Leith at Pine Hill" site (ID 508510), which we use here, is also used by the Otago Regional Council to monitor Dunedin's rainfall.

[7]HIRDS generates separate distributions for the periods 2031–2050 and 2081–2100. The bad- and good-scenario paths used here for the location parameter have the functional form $\alpha_0 - \alpha_{1b}e^{-\alpha_{2b}t}$ and $\alpha_0 - \alpha_{1g}e^{-\alpha_{2g}t}$, respectively, with the parameters chosen so that: $\lambda_b(t)$ matches the mean of the RCP8.5 distribution at $t \in \{0, 20, 75\}$; $\lambda'_g(0) = \lambda'_b(0)$; and the limiting value of $\lambda_g(t)$ matches the mean of the RCP2.6 distribution at $t = 75$.

Figure 1: Distribution of the annual maximum one-day rainfall in Dunedin

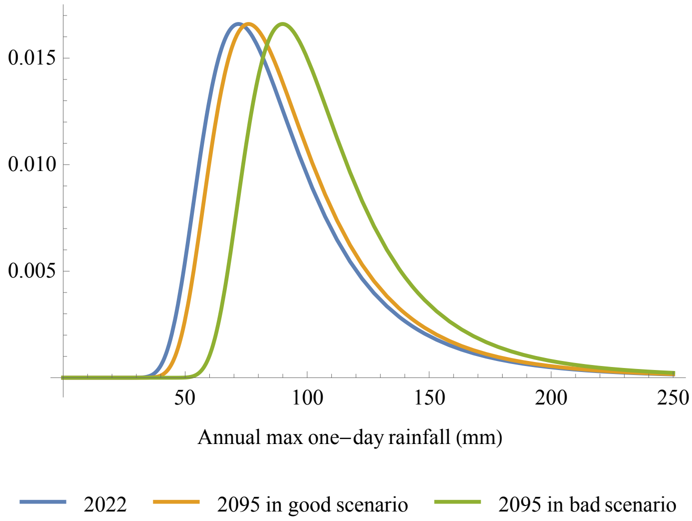

Notes. The blue curve plots the estimated density function for the annual maximum one-day rainfall in Dunedin, New Zealand, given current climatic conditions. The estimated density functions for the same variable in 2095 are shown by the yellow (good scenario) and green (bad scenarios) curves.

term, but then they diverge. This means there is not much potential for observation-based learning in the short run, but there will be in the long-run. This will have an impact on the value of a project's embedded flexibility.

Suppose the decision-maker initially believes the two climate scenarios are equally likely (that is, $p = 1/2$) and she observes the path of the variable z_t described in equation (1), with $\theta = 15$. That is, each year the decision-maker observes a signal of the average annual maximum one-day rainfall with a standard deviation of 15mm. The left-hand graph in Figure 3 plots two simulated paths for the subjective probability that the bad scenario holds, both starting at $p_0 = 1/2$. The right-hand graph plots the corresponding paths for the average annual maximum one-day rainfall, $p_t\lambda_b(t) + (1 - p_t)\lambda_g(t)$. Very little happens to the subjective probabilities during the first ten years, due to the similarities between the two scenarios during this period. However, the subjective probabilities are highly volatile for the next 50 years. The true scenario is revealed by the time 80 years have passed, but even as late as 50 years from now, there is uncertainty about which scenario holds. The behaviour of the average annual maximum one-day rainfall in the right-hand graph illustrates how changes in beliefs drive the volatility of the expected rainfall intensity (and, ultimately, the volatility of the project's expected benefits). The expected value of the average annual maximum one-day rainfall evolves smoothly for the first ten years, but is volatile for the next 50 years. Eventually, the average annual maximum one-day rainfall moves smoothly along one of the two possible paths—but which path it will be is not known for several decades. Figure 3 highlights the three distinct phases of the decision-maker's information about climate change. Initially, there is substantial uncertainty about climate change, but too little information arriving for the uncertainty to fall significantly. This is followed by a period of substantial volatility, during which the decision-maker's uncertainty falls. The simulated paths for p_t start to diverge during this period because the possible paths for the average annual maximum one-day rainfall start to diverge, so that the noisy signal contains more information. Finally, the uncertainty is largely resolved and the decision-maker knows which path the climate is evolving along.

This behaviour is confirmed by the graphs in Figure 4, which plot the distributions of p

Figure 2: Evolution of the average annual maximum one-day rainfall in Dunedin

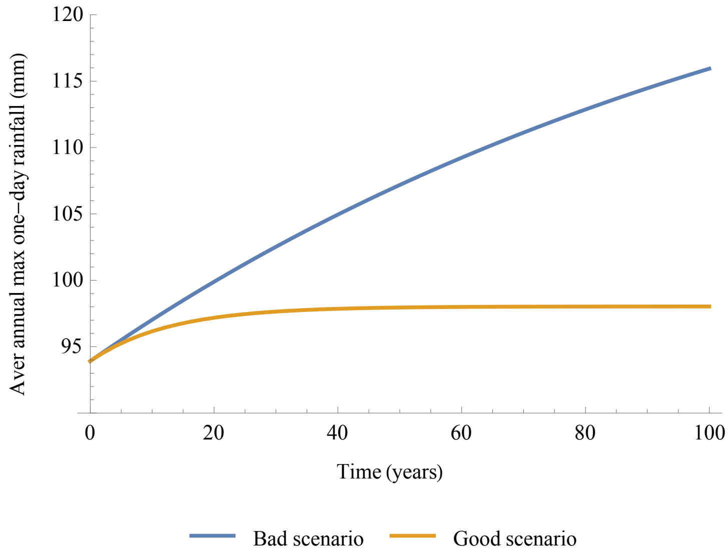

Notes. The graph plots the evolution of the average annual maximum one-day rainfall in Dunedin under the two climate scenarios.

for four horizons.[8] The top left-hand graph shows that ten years from now, the decision-maker has learned little about the true climate scenario: the subjective probability is close to its initial value of 1/2 for all 10,000 simulations. The top right-hand graph shows there will be considerable variation in this probability 20 years from now; the bottom left-hand graph shows the very beginning of a bimodal distribution 40 years from now. The bottom right-hand graph shows that by the time 80 years have passed, the decision-maker will be almost certain about the true climate scenario: the subjective probability will be extremely close to zero or one.[9]

4 Measuring expected flooding costs

4.1 Flooding-cost function

The functions $\Gamma_b(q,t)$ and $\Gamma_g(q,t)$ determine the expected flooding cost, given the stormwater system's capacity q and the current date t. In practical applications, this function is often calculated using detailed hydrological modelling in a two-step process. The first step is to calculate the macroscopic damage function, which gives the flooding cost as a function of the scale of any breach of the stormwater system. The second step averages over all possible breach scales.

When calculating the cost associated with a breach, we need to know how the event affects the surrounding land. Specifically, for each potential flooding event, we need to know the depth of the floodwaters on each unit of land in the locality. Analysts often use either resource-intensive hydrodynamic modelling or simpler flood-fill algorithms to determine the distribution of the depth of floodwaters in a locality as a function of the amount by which the peak rainfall exceeds the capacity of the stormwater system.[10] This information is combined with a depth–damage

[8] The graphs are constructed by simulating 10,000 paths for p each starting with $p = 1/2$ and spanning 80 years, using the same method to simulate each path as is used to construct the paths in Figure 3. We extract the values of p at each $t \in \{10, 20, 40, 80\}$ from each simulated path. The graphs in Figure 4 plot the histograms for each of these slices.

[9] If the decision-maker initially believes that the bad scenario is twice as likely as the good one, then the long-run distribution of p will be a discrete distribution with $p = 1$ being twice as likely to occur as $p = 0$.

[10] Flood-fill algorithms identify all land that is below the assumed water level and is not cut off from the water

Figure 3: Two simulated paths for the decision-maker's beliefs regarding the climate scenario

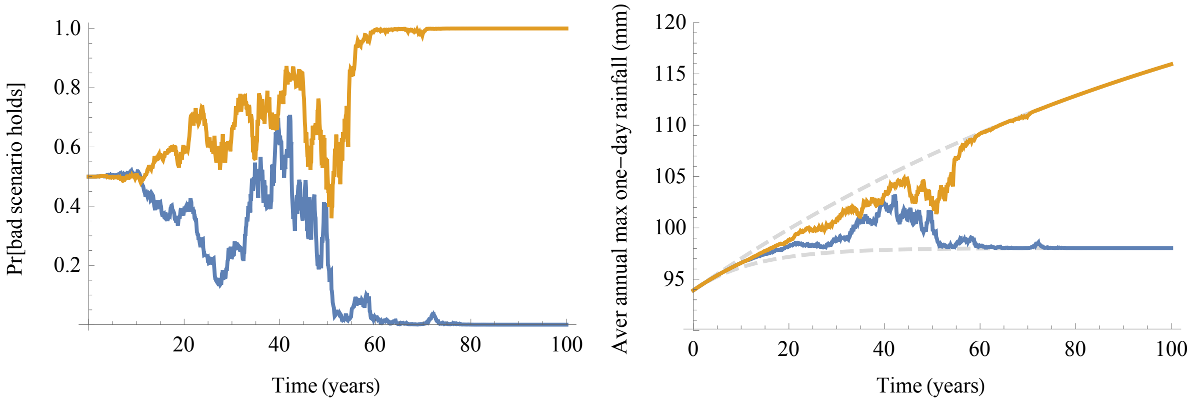

Notes. The left-hand graph plots two simulated paths for the subjective probability that the bad scenario holds, both starting at $p_0 = 1/2$. The right-hand graph plots the corresponding paths for the expected value of the climate factor, $p_t\lambda_b(t) + (1 - p_t)\lambda_g(t)$.

function, which estimates the cost of damage to a flooded property as a function of the depth of the floodwaters on that property.[11] Summing these costs over the entire locality yields the macroscopic damage function, which gives the total flooding cost as a function of the scale of the breach. That is, the macroscopic damage function equals

$$M(b) \equiv \int_0^\infty m(b,l)n(l)\,dl,$$

where $m(b,l)$ is the number of properties that experience a flood with level l following a breach with scale b, $n(l) \cdot x$ is the economic cost if a property experiences a flood with depth l, and we integrate over all possible flood levels. Taking the scenario-specific expected values of $M(b)$ over all possible breach scales b gives the functions Γ_b and Γ_g.

In order to focus on the role of climate uncertainty, and not the hydrological characteristics of a particular locality, we assume the macroscopic damage function takes the form $M(b) = \max\{0, b\}$, where the breach b equals the annual maximum one-day rainfall R minus the capacity of the local stormwater system q. Suppose the distribution functions of the annual maximum one-day rainfall are $F_b(R;t)$ and $F_g(R;t)$ in the bad and good scenarios, respectively, at date t. The expected flooding cost at date t therefore equals $x_t\Gamma_b(q_t,t)$ in the bad scenario, where

$$\Gamma_b(q_t,t) = E_t[\max\{0, R_t - q_t\}] = \int_{q_t}^\infty (R - q_t)F_b'(R;t)\,dR.$$

The definition of $\Gamma_g(q_t,t)$ is similar. As required, these are decreasing and convex functions of the system's capacity, q_t. That is, investing in greater capacity reduces expected flooding costs, but with diminishing returns. We will see in Section 5 that the optimal investment policy depends on the marginal avoided flooding cost. The next lemma shows that this quantity has a particularly simple form.

source by land that lies above the assumed water level. That is, low lying areas that are not connected to the flood source are assumed to be unaffected by the flood. Boettle et al. (2011) use a flood-fill algorithm to construct a flooding cost function for Copenhagen; Prahl et al. (2018) use one to estimate flooding cost functions for the 600 largest European coastal cities

[11]Damage to buildings is just one component of flooding costs. Depth–damage functions can also be constructed for damage to infrastructure. Similarly, depth–disruption functions can be constructed to estimate the costs of the disruption to the transport sector caused by flooding events (Pregnolato et al., 2017).

Figure 4: Distribution of the future probability that the bad scenario holds, when $p_0 = 1/2$

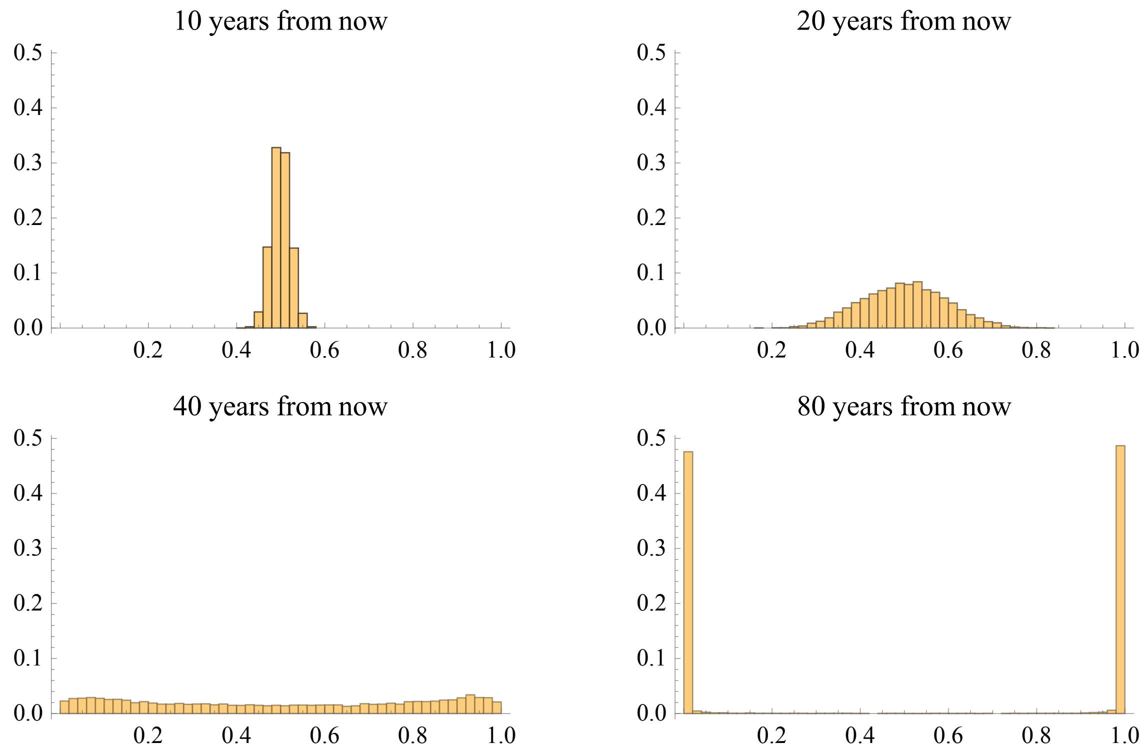

Notes. The graphs plot the distributions of p for four different horizons. Each histogram is drawn using 10,000 paths for p, each starting with $p = 1/2$ and each calculated with the method used to construct Figure 3.

Lemma 2 *Suppose the stormwater system has capacity q and the annual maximum one-day rainfall has distribution function $F(R, t)$ at date t. If the macroscopic damage function takes the form $\max\{0, R_t - q_t\}$, then increasing the stormwater system's capacity by the infinitesimal amount dq reduces the date t expected flooding cost by $x_t(1 - F(q, t))dq$.* ■

Climate change shifts the distribution of R rightwards, which reduces $F(q; t)$ and therefore increases the marginal benefit of investment.

The cost functions Γ_b and Γ_g will not give perfect results for any particular situation, but they will be a good approximation in many situations. For example, the functional form of Γ_b and Γ_g generates expected flooding costs that are a decreasing function of the capacity of the current flood-protection infrastructure, and which are larger when the annual maximum one-day rainfall is more variable. Furthermore, the benefits of investment are negligible when the system's capacity is high, they are substantial when it is low, and investing in additional capacity has diminishing returns. These are all realistic properties and play important roles in determining optimal investment policies.

4.2 The economic factor x_t

Recall that the expected flooding cost over the next increment of time equals $\Gamma_b(q_t, t)x_t\,dt$ if the bad scenario holds and $\Gamma_g(q_t, t)x_t\,dt$ if the good scenario holds. The role of the exogenous state variable x_t is to capture the effects of intertemporal changes in the depreciated replacement cost of the vulnerable assets, as well as changes in clean-up costs, disruption costs, environmental damage, and health-related costs. The behaviour of x_t determines the answer to the question: how would the costs associated with a flooding event change if an identical event occurred one year later? In order to keep the model tractable, we assume that x_t evolves according to geometric Brownian motion. All we need to specify regarding this variable are the drift μ and volatility σ of this stochastic process.

We select values for the model's parameters to approximate the situation facing Dunedin City Council's upgrade of its stormwater system. The investment programme will involve improving maintenance of stormwater grates, upgrading pumping stations, installing backflow prevention valves, and replacing old wastewater pipes. The slow rate of climate change means this ten-year programme will be just one part of the ongoing work needed to maintain and upgrade the stormwater system in response to climate change. The benefits generated by system upgrades comprise avoided flooding costs. The main sources of uncertainty surrounding the magnitude of these benefits are climatic (the frequency and severity of extreme rainfall events) and economic. Although the overall budget is significant, much of this investment programme can be scaled up or down as new information about future flood risk arrives.

Many factors potentially influence how x_t evolves over time. Perhaps the most important influence is the depreciated replacement cost of vulnerable assets (Merz et al., 2010; Penning-Rowsell et al., 2013). Over time, this quantity fluctuates because of a combination of depreciation of the existing assets, spending on maintenance and improvement of existing assets as the economy grows, investment in additional assets, and changes in the prices of inputs into the repair and replacement process. Let C_n denote the level of an index of real construction costs, N_n the number of dwellings, and DRC_n a depreciated replacement cost index, all at the end of year n, and let the constant δ denote the rate of depreciation. If M_{t+1} dwellings are added during the next year, then one estimate of the depreciated replacement cost of assets at the end of year $n+1$ equals[12]

$$DRC_{n+1} = DRC_n \left(\frac{C_{n+1}}{C_n}\right)\left(1 - \delta + \frac{M_{t+1}}{N_t}\right).$$

The term inside the first pair of large brackets on the right-hand side captures the effect of changes in construction costs; the term inside the second pair of large brackets captures the effects of depreciation and additions to the housing stock. We assume a depreciation rate of $\delta = 0.015$ and combine a national construction cost index with building consent data for Dunedin City to construct DRC_n for the period 1996–2022.[13] This will be our main proxy for the economic factor. However, changes in local GDP, or some other measure of economic activity, are also potentially a useful proxy for changes in the disruption costs associated with flooding.[14] Changes in local population might give useful information about changes in the health effects and number of fatalities associated with flooding. Changes in traffic volumes might give useful information about disruption costs when transport infrastructure is affected by flooding.

We estimate the drift and volatility for a collection of proxies for the economic factor and then make an overall assessment of plausible values of μ and σ. One issue that arises with most of these proxies is that they are unlikely to evolve according to geometric Brownian motion. In particular, shocks to these variables are typically serially correlated, whereas geometric Brownian motion assumes that changes in any two non-overlapping time periods are independent. Positive serial correlation in annual growth rates means that a relatively small shock to the current depreciated replacement cost of vulnerable assets can have a larger impact on the long-run depreciated replacement cost because it leads to a persistent sequence of changes in subsequent

[12]Statistics New Zealand reports building consents per capita. We assume a constant household size of 2.45 individuals and calculate the implied value of M_{t+1}/N_t. Comparing census data from 2006 and 2018 shows that the average household size of a Dunedin household fell from 2.49 to 2.40 over this period.

[13]This approach will overestimate the average growth rate in the vulnerable housing stock if Dunedin's growth has occurred in areas not prone to flooding. However, most of the growth in the proxy variable comes via the construction cost index, so the effects of intra-city variations on the parameter estimates are likely to be minor.

[14]To be clear, we would not be using the change in GDP in the aftermath of a flood to measure the cost of flooding. Rather, we would be supposing that if GDP grows by ten percent over time, then—all else equal—expected disruption costs due to flooding grow by ten percent as well.

Table 1: Proxies for the drift and volatility of the economic factor x_t

Variable	Long-run mean of $\Delta \log x$	Effect of a one std dev shock on $\Delta \log x$					
		1	2	3	5	10	∞
Depreciated replacement cost	0.015	0.023	0.039	0.049	0.061	0.069	0.070
Regional GDP (total)	0.029	0.026	0.035	0.038	0.040	0.040	0.040
Regional per capita GDP (total)	0.017	0.024	0.030	0.032	0.032	0.032	0.032
Regional per capita GDP (total) × local pop'n	0.022	0.026	0.034	0.036	0.037	0.037	0.037
Regional GDP (owner-occupied property)	0.047	0.038	0.054	0.061	0.065	0.066	0.066

Notes. We fit an AR(1) model to annual changes in the natural logarithm of each variable and then calculate the effect of a one-standard deviation annual shock over the subsequent ten years.

years. Therefore, the short-run volatility of these quantities will likely underestimate the variability of their long-run values—and the latter is more important for the types of investment projects we consider. In particular, the value of the real option to delay investment ultimately depends on the volatility of the present value of the flow of a project's benefits, rather than the short-term volatility of the flow itself, and shocks to long-run benefit flows give a good indication of the former.

In order to address this issue, we focus on the long-term effects of shocks to the variables above. Specifically, for each variable, we fit an AR(1) model to annual changes in the natural logarithm of the variable and then calculate the effect of a one-standard deviation annual shock on the variable over the subsequent ten years. Table 1 reports the results of this procedure. The first row shows the results for the depreciated replacement cost index, which suggest that $\mu = 0.015$ and $\sigma = 0.070$ are reasonable estimates of the drift and volatility of the economic factor. The remaining rows of Table 1 show results for various measures of local economic activity. GDP is available only for the wider Otago region (rather than just Dunedin City), so we report results for regional GDP, per capita regional GDP, and the product of per capital regional GDP and the population of Dunedin city. We also report results for the component of regional GDP corresponding to the services generated by owner-occupied property. The values of the drift parameter are higher than the drift of depreciated replacement cost, which is not surprising because none of these measures of economic activity allow for depreciation. In contrast, the values of the volatility parameter are lower than the volatility of depreciated replacement cost, which reflects the relatively high volatility of real construction costs. Bearing in mind the importance of depreciation and construction (repair) costs for the calculation of flooding costs, we adopt $\mu = 0.015$ and $\sigma = 0.070$ as our baseline parameter values. The three curves in Figure 5 plot the density functions of the distribution of x_t for $t \in \{10, 20, 40\}$, assuming that $x_0 = 1$. Even though the volatility might seem low, the long time frames involved in the decision-maker's investment problem mean that there will be considerable variation in the economic factor over relevant time frames.

We set the discount rate equal to $r = 0.05$, consistent with the (real) social discount rate recommended by the New Zealand Treasury (Treasury, 2022). Without any loss of generality, we choose the units of expenditure so that the cost of building new capacity is $c = 1$.

5 Theory

This section describes the approach the decision-maker uses to evaluate investment policies. She chooses the capacity of the stormwater system at each date, subject only to the constraint that, due to the irreversibility of investment, this quantity can never decrease. There are no fixed upgrade costs, the marginal cost of upgrading the system is constant, and the expected flow of flooding-related costs,

$$(p_t \Gamma_b(q_t, t) + (1 - p_t)\Gamma_g(q_t, t))\, x_t,$$

Figure 5: Distribution of future values of the economic factor

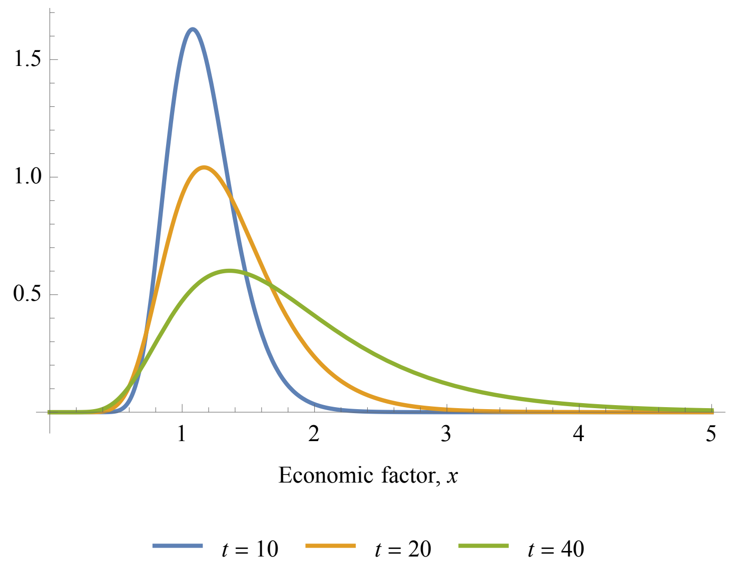

Notes. The graph plots the density functions for the distribution of the economic factor at dates 10, 20, and 40, assuming its current value equals 1. The drift and volatility parameters equal $\mu = 0.015$ and $\sigma = 0.070$.

is a decreasing, convex function of the system's capacity. The decision-maker can therefore restrict attention to policies of instantaneous control, as the best policy of this type achieves an overall level of costs that is less than or equal to that resulting from any other investment policy.[15]

When evaluating investment policies, the decision-maker seeks to minimise the present value of all future flooding-related costs, including expenditure on upgrading the stormwater system. We denote this quantity by $W(t, p, x; q)$. This section presents the conditions that determine optimal investment policies using a sequence of three lemmas to develop the formulation of the decision-maker's problem that we will subsequently solve. Many papers use smooth-pasting conditions to describe these conditions, but we use variational inequalities instead as these are more closely related to the numerical approach used to solve the decision-maker's investment-timing problem.[16] The first of these lemmas presents the variational inequalities that the value function W satisfies if the decision-maker adopts an optimal upgrade policy.

Lemma 3 *If the decision-maker adopts an optimal upgrade policy, then the value function* $W(t, p, x; q)$ *satisfies*

$$c + W_q(t, p, x; q) \geq 0$$

and

$$\mathcal{D}W(t, p, x; q) - rW(t, p, x; q) + (p\Gamma_b(q, t) + (1 - p)\Gamma_g(q, t))x \geq 0$$

[15] See Dixit (1993) for an introduction to instantaneous control policies and Stokey (2008) for a more detailed discussion. The optimality of an instantaneous control policy for a similar problem is proven on pp. 360–367 of Dixit and Pindyck (1994). If the investment cost features a fixed cost, so that increasing system capacity by Δq requires expenditure of $f + c \cdot \Delta q$ for some positive constant f, then the decision-maker should adopt a policy of impulse control, which involves lumpy (rather than continuous) investment.

[16] Appendix B outlines the numerical method used here, which is based on the operator-splitting approach of Ikonen and Toivanen (2004, 2009).

everywhere, where

$$\mathcal{D} = \frac{\partial}{\partial t} + \frac{1}{2}\left(\frac{p(1-p)(\lambda_g(t)-\lambda_b(t))}{\theta}\right)^2 \frac{\partial^2}{\partial p^2} + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2}{\partial x^2} + \mu x \frac{\partial}{\partial x}.$$

The first condition holds with equality in the investment region and the second condition holds with equality in the waiting region. ■

The left-hand side of the first inequality in Lemma 3 equals the marginal effect of investment on the present value of all future costs. If the decision-maker increases capacity slightly, she incurs immediate expenditure of c and the present value of all subsequent costs changes by W_q (that is, it falls by $-W_q$). Provided the decision-maker has adopted an optimal policy, the sum of these two terms is either positive (when it is optimal not to invest) or zero (when the decision-maker is investing). In contrast, the second inequality in Lemma 3 captures the effect of delaying investment. If the decision-maker holds capacity constant for the next dt units of time, she incurs expected flooding costs of $(p\Gamma_b(q,t) + (1-p)\Gamma_g(q,t))x\,dt$ and the present value of future costs increases by $\mathcal{D}W\,dt$. The amount by which the sum of these two terms exceeds $rW\,dt$ equals the cost of waiting rather than following an optimal upgrade policy. Therefore, the left-hand side of the second inequality is positive when investing is optimal and is zero when waiting is optimal.

It is convenient to focus on the marginal value of capacity. Rearranging the first-order Taylor series expansion of $W(t,p,x;q+dq)$ shows that

$$W(t,p,x;q) = (-W_q(t,p,x;q)\,dq) + W(t,p,x;q+dq) + o(dq)$$

as $dq \to 0$. That is, the present value of the future costs associated with a stormwater system with capacity q, $W(t,p,x;q)$, equals the sum of the present value of the costs associated with the marginal unit of capacity, $-W_q(t,p,x;q)\,dq$, and the present value of the future costs associated with an upgraded system with capacity $q+dq$, $W(t,p,x;q+dq)$. The next lemma presents the variational inequalities satisfied by the marginal value, W_q.

Lemma 4 *If the decision-maker adopts an optimal upgrade policy, then $W_q(t,p,x;q)$ satisfies*

$$W_q(t,p,x;q) + c \geq 0$$

and

$$\mathcal{D}W_q(t,p,x;q) - rW_q(t,p,x;q) + \left(p\frac{\partial\Gamma_b(q,t)}{\partial q} + (1-p)\frac{\partial\Gamma_g(q,t)}{\partial q}\right)x \leq 0$$

everywhere. The first condition holds with equality in the investment region and the second condition holds with equality in the waiting region. ■

The most useful feature of the variational inequalities in Lemma 4 is that q only appears explicitly via the nonhomogeneous term in the second inequality.[17] That is, q assumes the same status as parameters such as the scale and shape of the rainfall distribution—it can be treated as a constant that affects the level of the future flooding costs. This means that, rather than solving for an optimal upgrade policy for the stormwater system as a whole, the decision-maker can solve for an optimal policy for investing in each additional unit of capacity individually. In other words, the problem can be broken into a collection of simpler problems, each one involving deciding when to install a particular unit of capacity.

The decision-maker's problem in Lemma 4 is equivalent to the problem of when to exercise a standard investment-timing option of the type introduced by McDonald and Siegel (1986). In this equivalent problem, there is no flow of benefits (or costs) until the incremental upgrade option is exercised, but exercising the option initiates a flow of avoided flooding costs. Specifically,

[17] This is most obvious if we use new notation for the dependent variable W_q. Writing $W_q = U$, for some function U, makes it clear that q appears only via the nonhomogeneous term.

exercising the option incurs lump-sum expenditure of $c\,dq$ and initiates a perpetual benefit flow of

$$-\left(p\frac{\partial\Gamma_b(q,t)}{\partial q}+(1-p)\frac{\partial\Gamma_g(q,t)}{\partial q}\right)x\,dq.$$

If this option is exercised at date t, the present value of this benefit flow equals

$$(p_t v_b(t;q)+(1-p_t)v_g(t;q))\,x_t\,dq,$$

where

$$v_b(t;q)=-\int_t^\infty e^{-(r-\mu)(s-t)}\frac{\partial\Gamma_b(q,s)}{\partial q}\,ds \tag{3}$$

and

$$v_g(t;q)=-\int_t^\infty e^{-(r-\mu)(s-t)}\frac{\partial\Gamma_g(q,s)}{\partial q}\,ds. \tag{4}$$

We can isolate the present value of the incremental upgrade option by comparing the present value of all future costs with and without this option. The present value of all future costs equals $W(t,p,x;q)$ if the decision-maker has the option to install all units of capacity greater than q. In contrast, it equals

$$(pv_b(t;q)+(1-p)v_g(t;q))\,x\,dq+W(t,p,x;q+dq)$$

if she only has the option to install all units of capacity greater than $q+dq$, where the latter is the sum of the present value of the flooding costs that the upgrade option would avoid and the present value of all other costs. The present value of the incremental upgrade option is therefore equal to

$$\begin{aligned}V(t,p,x;q)\,dq &\equiv \big((pv_b(t;q)+(1-p)v_g(t;q))\,x\,dq+W(t,p,x;q+dq)\big)-W(t,p,x;q)\\ &= \big(\big(pv_b(t;q)+(1-p)v_g(t;q)\big)\,x+W_q(t,p,x;q)\big)\,dq.\end{aligned}$$

The next lemma describes the variational inequalities satisfied by this present value.

Lemma 5 *If the decision-maker adopts an optimal upgrade policy, then the present value of the incremental upgrade option satisfies*

$$V(t,p,x;q)\geq\big(pv_b(t;q)+(1-p)v_g(t;q)\big)\,x-c$$

and

$$\mathcal{D}V(t,p,x;q)-rV(t,p,x;q)\leq 0$$

everywhere. The first condition holds with equality in the investment region and the second condition holds with equality in the waiting region. ■

Lemma 5 describes the investment-timing problem that we will solve and is the basis of the analysis in the rest of the paper. That is, rather than solving the overall system upgrade problem described in Lemma 3, in principle we solve it one "incremental upgrade project" at a time. The variational inequalities in Lemma 5 are typical for optimal stopping problems. The first one reflects the fact that the investment payoff equals the present value of the flooding costs avoided due to the increased capacity, minus the investment expenditure. The second inequality reflects that fact that—from the perspective of the incremental upgrade option—there is a zero net-benefit flow during the period prior to investment.

To conclude this section, consider the special case where the model's parameters take the values specified in Sections 3 and 4. The left-hand graph in Figure 6 plots the benefit flow multipliers $-\frac{\partial\Gamma_b(q,t)}{\partial q}$ and $-\frac{\partial\Gamma_g(q,t)}{\partial q}$, whereas the right-hand graph plots the investment payoff multipliers $v_b(t;q)$ and $v_g(t;q)$ defined in equations (3) and (4), all as functions of t. That is,

Figure 6: Investment benefit functions

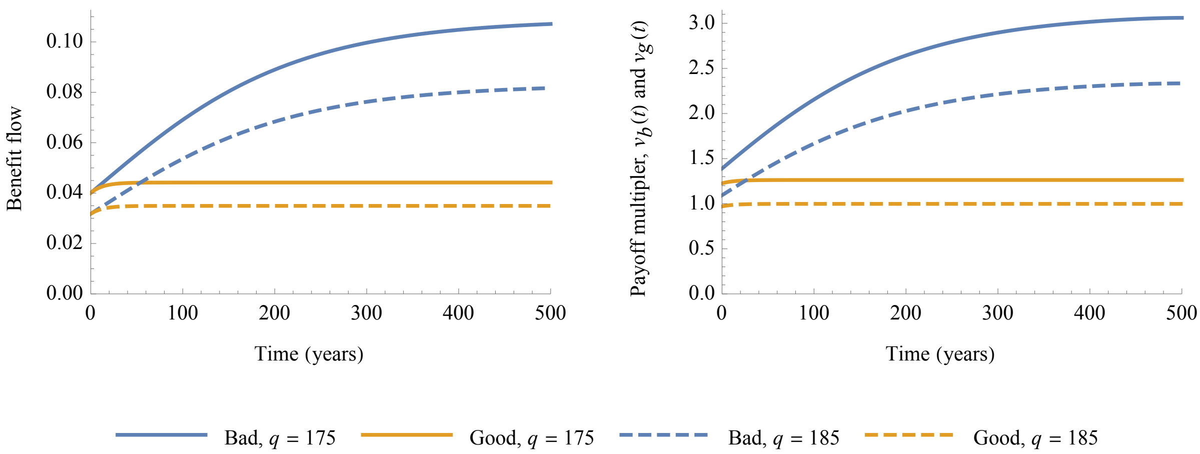

Notes. The left-hand graph plots the benefit flow multipliers $-\frac{\partial \Gamma_b(q,t)}{\partial q}$ and $-\frac{\partial \Gamma_g(q,t)}{\partial q}$, whereas the right-hand graph plots the investment payoff multipliers $v_b(t;q)$ and $v_g(t;q)$ defined in equations (3) and (4), all as functions of t. The solid and dashed curves correspond to a system capacity of 175mm/day and 185mm/day, respectively.

if the bad scenario holds, the benefit flow generated by an incremental system upgrade equals the product of the height of the blue curve in the left-hand graph and the current level of the economic factor x. The present value of the benefits generated by an incremental system upgrade equals the product of the height of the blue curve in the right-hand graph and the current level of the economic factor x. If the good scenario holds, the multipliers are the heights of the corresponding yellow curves. The solid and dashed curves correspond to a system capacity of 175mm/day and 185mm/day, respectively. The benefits are more valuable when the bad scenario is more likely to hold and when the system's current capacity is relatively low. Delaying investment does not significantly increase the multiplier if the good scenario holds, but can lead to a very significant increase if the bad scenario holds.

6 Optimal investment

Each incremental upgrade option is a perpetual American call option that gives its holder (the decision-maker) the right to pay c in exchange for a perpetual flow of benefits. The most important properties of such options are traditionally illustrated using graphs that plot the present value of the option as a function of the present value of the underlying asset. Figure 7 adopts a similar approach. Each graph plots the present value of an incremental upgrade option as a function of the economic factor x for a different combination of the stormwater system's current capacity (the columns) and the evaluation date (the rows). The blue curves show results when the decision-maker is certain the good scenario holds, the green curves show results when she is certain the bad scenario holds, and the yellow curves show results when she believes the two scenarios are equally likely. If x is to the left of the black dot on the curve corresponding to the current level of p, then the decision-maker should delay investment; if it is to the right of this point, then she should invest immediately.

Figure 7 reveals considerable information about the optimal investment policy. Consider the top pair of graphs, corresponding to the situation at date 0. In both cases, the dots are clustered together, showing that the optimal investment threshold for the economic factor is reasonably insensitive to p. However, the three points shift rightwards when the stormwater system's capac-

Figure 7: Present value of incremental upgrade options

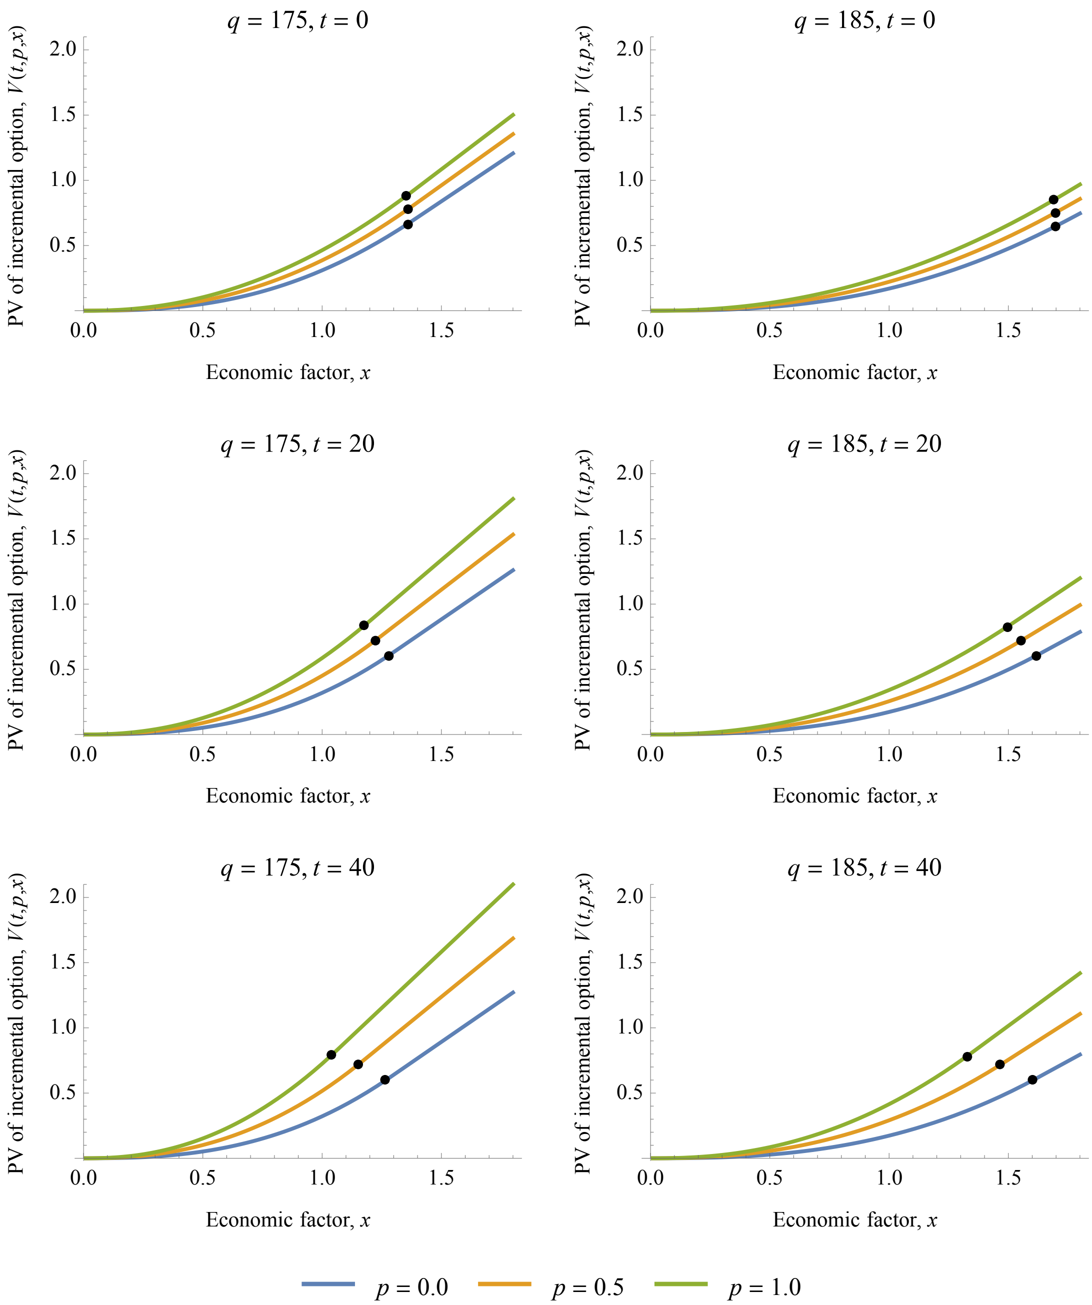

Notes. Each graph plots the present value of an incremental upgrade option as a function of the economic factor for a different combination of the stormwater system's current capacity (the columns) and the evaluation date (the rows). The blue curves show results when the decision-maker is certain the good scenario holds, the green curves show results when she is certain the bad scenario holds, and the yellow curves show results when she believes the two scenarios are equally likely to hold. The black dots split each curve into waiting and investing regions.

ity increases. That is, the investment test becomes more demanding as q increases.[18] In contrast, the dots are spread out in the middle pair of graphs, which show the situation at date 20, indicating that the optimal investment threshold becomes more sensitive to p over time. If the bad climate scenario is more likely, then the investment threshold is less demanding. The dots are spread out even further in the bottom pair of graphs, indicating that the optimal investment threshold's sensitivity to p grows as we look further into the future. The dots shift left as we move down each column; that is, the investment test becomes less demanding as we move further into the future.

The reason why the optimal thresholds in Figure 7 are initially almost independent of p can be traced back to the properties of climate change. We might expect the optimal investment threshold to be sensitive to p, as this parameter determines the expected future path of rainfall intensity, which affects the benefit flow initiated by investment. However, the project's benefit flow initially has a high expected growth rate and low volatility, as suggested by Figures 2 and 6. This means that the only reason to invest immediately is that the short-term benefits are large enough to compensate for incurring expenditure now rather than at some future date.[19] These short-term benefits are independent of p because of the short-term similarity of the possible paths for rainfall intensity, as can be seen in Figure 2.[20] As we move further into the future, the expected growth rate of the benefit flow falls and climate volatility rises, so that long-term considerations become increasingly important to the decision-maker. This is reflected in the optimal investment threshold becoming more sensitive to p as t increases.

Figure 7 also gives an indication of just how valuable the decision-maker's investment-timing option is. In particular, the height of each dot equals the present value of the option at the time it is exercised. That is, the height shows the amount by which the present value of the project's benefits must exceed the required expenditure in order for investment to be optimal. Recall that this expenditure is $c = 1$. Thus, the graphs in Figure 7 show it can be optimal to delay upgrading the system until the present value of the benefits from investing exceed the required expenditure by 50% or more. The graphs show that—when the investment test is expressed in these terms—the threshold for investment can be much more demanding when the climate scenario is more likely to be bad; that is, the dots on the green curves are significantly higher than their counterparts on the yellow and blue curves.

The form of the model means it is natural to express the investment policy as a threshold $\hat{x}(t, p; q)$ for the economic factor. This formulation leads to a particularly simple rule for exercising an incremental upgrade option: if the system currently has capacity q then it is optimal to exercise the incremental upgrade option the first time that $x_t \geq \hat{x}(t, p_t, q)$. If the decision-maker adopts this policy then she will invest in just enough new capacity at each date to prevent x_t from exceeding $\hat{x}(t, p_t, q_t)$. That is, she follows a policy of barrier control. This means the capacity of the stormwater system may remain constant for long periods of time, but its capacity will start to increase again once flooding-related costs rise to a sufficiently high level. Investment can be triggered by increases in x (that is, the costs avoided due to investment rise), p (that is, floods become more likely), and t (that is, the progression of climate change).

Investment policies are often implemented using a threshold for the benefit–cost ratio (BCR), which equals the present value of a project's future benefits divided by the present value of the investment expenditure. Recall that the present value of the incremental upgrade option's future benefits equals

$$\big(p_t v_b(t; q) + (1 - p_t) v_g(t; q)\big)\, x_t$$

[18]This is essential for us to be able to treat each incremental upgrade as an independent project. It means that the decision-maker will exercise the option to upgrade a system with capacity of 175mm/day before she will upgrade an otherwise identical system with capacity of 185mm/day.

[19]That is, the long-term benefits generated by an incremental system upgrade appear in both the payoff from investing immediately and the value of waiting. They therefore have little effect on the current investment decision.

[20]If the paths started at different levels or diverged more rapidly, then the present value of the short-term benefits would be more sensitive to p, which would lead the x threshold to be more sensitive to p.

Figure 8: Investment threshold in terms of the benefit–cost ratio

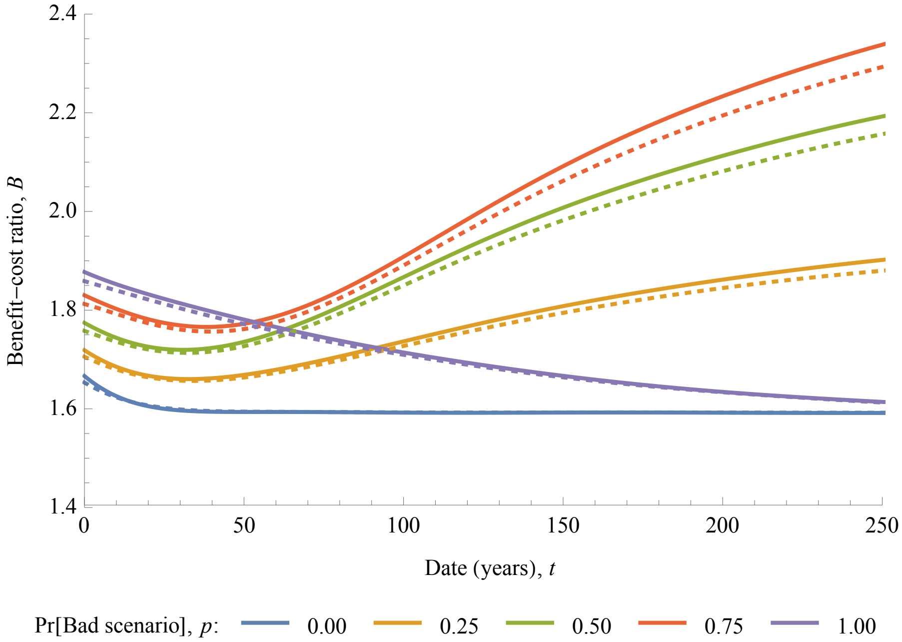

Notes. The graph plots the investment threshold for the benefit–cost ratio, $\hat{B}(t, p; q)$, as a function of t. The solid curves show the optimal threshold when $q = 175$ for the indicated levels of p. The dotted curves show the optimal threshold when $q = 185$, for the same levels of p.

at date t, if the system has capacity q at that date. Dividing this expression by the lump-sum investment expenditure shows that the benefit–cost ratio from undertaking "project q" at date t equals

$$BCR(t, p_t, x_t; q) = \frac{x_t}{c} \left(p_t v_b(t; q) + (1 - p_t) v_g(t; q) \right). \tag{5}$$

The decision-maker will therefore upgrade a system with capacity q the first time the corresponding benefit–cost ratio is greater than or equal to

$$\hat{B}(t, p_t, q) \equiv \frac{\hat{x}(t, p_t, q)}{c} \left(p_t v_b(t; q) + (1 - p_t) v_g(t; q) \right), \tag{6}$$

where $\hat{x}(t, p_t, q)$ denotes the optimal investment threshold for x.

The threshold for the benefit–cost ratio has a useful economic interpretation. The amount by which it exceeds one equals the value of the real option to delay investment, expressed as a proportion of the required investment expenditure. For example, if the delay option had no value then the optimal investment threshold would be $\hat{B} = 1$. Figure 8 plots the investment threshold for the benefit–cost ratio, $\hat{B}(t, p; q)$ in equation (6), as a function of the date t, for five different probabilities that the bad scenario holds. The solid curves show the threshold when $q = 175$; the dotted curves show the threshold when $q = 185$.[21] The value of the real option to delay investment is economically significant. That is, across a range of dates and climate probabilities, investment is only socially optimal if the benefits are very significantly greater than the required expenditure (at least 60% higher in all cases). As we will see shortly, the causes of this high option value vary. In the near term, the high threshold for the benefit–cost ratio is due to the

[21] When expressed in terms of the BCR, the investment test is slightly less demanding when the system's capacity is higher. However, this is not inconsistent with the approach of treating each incremental upgrade as an independent project, because Figure 7 shows the decision-maker will exercise the $q = 175$ option before she exercises the $q = 185$ option.

possibility that the climate scenario is bad (so that the expected growth rate in benefits is high); in the longer-term, the high threshold for the benefit–cost ratio is due to uncertainty about the climate scenario (and the resulting high level of p-volatility).

In the short-term, the BCR threshold is a strongly increasing function of p. That is, the investment test is more demanding when the bad scenario is believed to be more likely. Equivalently, the option to delay investment is initially more valuable when the bad scenario is believed to be more likely. Consider the situation when $p = 1$ (so that the bad climate scenario holds) and the level of the economic factor is such that the present value of the project's benefits is slightly greater than the investment expenditure. If the decision-maker had a now-or-never investment option then it would clearly be optimal to invest. However, suppose investment can be delayed. The high initial growth rate in the project's benefit flow means that delaying investment by one year reduces the present value of the benefit flow, measured at date 0, by just 2.8%, whereas the present value of the investment expenditure falls by 4.9%. Delay will therefore be optimal. Indeed, the present value of the benefits from immediate investment will need to exceed the investment expenditure by a considerable margin before immediate investment is optimal. In contrast, if $p = 0$ (so that the good climate scenario holds), delaying investment by one year reduces the present value of the benefit flow, measured at date 0, by 3.2%. The decision-maker's incentive to delay investment is therefore not as strong in this case, the delay option has less value, and the optimal BCR threshold for optimal investment is lower.

Now consider what happens as t increases, starting with the extreme cases where $p = 0$ and $p = 1$ (that is, when the true climate scenario is known with certainty and p is therefore constant). In both cases, the optimal BCR threshold falls steadily to a common limiting value. Volatility in the economic factor induces volatility in the investment payoff, which contributes to the value of the real option to delay investment. However, economic-factor volatility is constant over time and the investment option is perpetual, so the time dependence of the optimal BCR threshold must come from another source. That source is the declining growth rate of the project's benefit flow. As time passes and rainfall intensity converges to its new long-run level, the investment payoff multiplier ($v_b(t)$ or $v_g(t)$) increases, but at a declining rate. All else equal, this means that the incentive to delay investment falls as t increases. For example, from the point of view of date 100, delaying investment by one year in the bad scenario reduces the present value of the benefit flow, measured at date 100, by 3.2%, so that delay is more costly at date 100 than it is at date 0. Once rainfall intensity has converged to its long-run value (so that the investment payoff multiplier has converged to its long-run value), then the value of the delay option derives solely from the behaviour of the economic factor. The investment payoff is proportional to a stochastic variable that evolves according to geometric Brownian motion. As in the standard investment-timing model of McDonald and Siegel (1986), the optimal threshold for the BCR equals $\beta/(\beta - 1)$, where β is a simple function of the parameters μ, σ, and r. For this example, $\beta = 1.61$.[22]

The behaviour of the optimal BCR threshold is more complicated in the intermediate cases where $0 < p < 1$ (that is, when there is uncertainty about the true climate scenario). The effect of the high growth rate in rainfall intensity (in this case, expected rainfall intensity, given that there is climate uncertainty) remains, which contributes to the optimal BCR threshold falling over time. Now, however, there is an offsetting effect, which appears during the period when much is learnt about the state of the climate. During this period, the possible paths for rainfall intensity $\lambda(t)$ start to diverge, which means that the noisy signal of the true level of rainfall intensity is more informative. That is, climate uncertainty starts to fall and p volatility starts to rise, which leads to an increase in the volatility of the investment payoff. As is the case with other investment-timing problems, when the investment payoff becomes more volatile, the value

[22]The BCR threshold converges to its limiting value much more quickly when $p = 0$ because the trajectory of rainfall intensity converges to its long-run value much more quickly in the good climate scenario (when $p = 0$) than in the bad scenario. That is, the high-growth period lasts for a much shorter period of time.

of waiting increases, which explains the increase in the BCR threshold as time increases. This increase in BCR threshold works in the opposite direction to the one discussed above (for the cases with $p = 0$ and $p = 1$). At first it is small, so the BCR threshold is decreasing in t, but eventually it dominates the other effect and the BCR threshold is increasing in t.[23]

Each entry in Table 2 reports the optimal BCR threshold for a different combination of parameters. The first row of the table shows results for the baseline parameter values and the large panels vary one parameter at a time, holding all other parameters at their baseline values. The table shows that the BCR threshold exceeds one by an economically significant margin across a large selection of parameter values. Unless the drift in the economic factor is very small, the optimal investment policy requires that the present value of the project's benefits exceed investment expenditure by at least 50%. Premia of 80% and above are common. It is optimal to wait for a larger BCR when the economic factor has higher volatility and a higher expected growth rate. In the short–medium term, the amount of noise in the climate signal has a negligible impact on the optimal BCR threshold. If investment does not occur for many decades, then the optimal BCR threshold is slightly decreasing in the amount of noise in the signal (and therefore slightly increasing in the volatility of the probability that the bad scenario holds).

7 Approximating the optimal investment policy

The focus so far has been on the decision-maker's optimal investment policy, which is derived using real options analysis. The benefits from using real options analysis can be significant, but it is not widely used, particularly when evaluating relatively small projects. This is especially serious because much adaptation spending will fund relatively small projects under the jurisdiction of authorities with limited analytical resources. The importance of the flexibility in these projects is too costly to ignore, so we need approaches that are simple enough to be useful for evaluating small- and medium-scale adaptation decisions, yet retain a degree of economic rigour. This section evaluates the potential of one such approach.

7.1 Approximating the value of the delay option

Recall that it is optimal to invest as soon as the net present value from investing is greater than or equal to the value of the option to delay investment, provided that the option value is consistent with the decision-maker adopting an optimal investment policy in the future. Equivalently, it is optimal to invest as soon as

$$\big(p_t v_b(t; q_t) + (1 - p_t) v_g(t; q_t)\big)\, x_t \geq c + V(t, p_t, x_t, q_t), \qquad (7)$$

where the option value $V(t, p_t, x_t, q_t)$ satisfies the conditions in Lemma 5. That is, the present value of the project's benefits (the left-hand side of equation (7)), must be greater than or equal to the opportunity cost of investment. This opportunity cost equals the sum of the investment expenditure and the present value of the delay option that is destroyed by the act of investment. It is this option value that is difficult to calculate in practical applications. Therefore, this subsection describes a relatively straightforward approach to estimating $V(t, p_t, x_t, q_t)$ that, in many situations, results in a close approximation to the correct value.

The alternative approach to valuing the delay option considered here assumes that if the decision-maker delays investment then she will invest at the best *fixed* future date. This is obviously not an optimal policy, but we are only using this assumption for the purposes of

[23] The BCR threshold reaches very high levels for large t when $0 < p < 1$. In this situation, the true climate scenario will be revealed very quickly, so that the investment payoff is very volatile and the option value of waiting is very high. However, this combination of circumstances has little practical relevance as the climate scenario will almost certainly be known by then (or soon will be), in which case p will be very close to either 0 or 1. However, the results are included here to explain the non-monotonic nature of the BCR thresholds over horizons that *are* of practical relevance.

Table 2: Optimal investment using a threshold for the benefit–cost ratio

θ	σ	μ	$p = 0.25$				$p = 0.50$				$p = 0.75$			
			$t = 0$	$t = 40$	$t = 80$	$t = 120$	$t = 0$	$t = 40$	$t = 80$	$t = 120$	$t = 0$	$t = 40$	$t = 80$	$t = 120$
Baseline case														
15	0.070	0.015	1.72	1.66	1.70	1.76	1.77	1.72	1.81	1.92	1.83	1.76	1.84	1.98
Volatility of economic factor, σ														
15	0.050	0.015	1.65	1.58	1.63	1.69	1.71	1.65	1.73	1.84	1.76	1.69	1.78	1.92
15	0.060	0.015	1.68	1.62	1.67	1.73	1.73	1.68	1.77	1.90	1.78	1.73	1.80	1.96
15	0.070	0.015	1.72	1.66	1.70	1.76	1.77	1.72	1.81	1.92	1.83	1.76	1.84	1.98
15	0.080	0.015	1.76	1.70	1.74	1.81	1.82	1.77	1.84	1.97	1.87	1.81	1.88	2.03
15	0.090	0.015	1.81	1.75	1.79	1.87	1.86	1.81	1.89	2.03	1.92	1.86	1.91	2.06
Drift of economic factor, μ														
15	0.070	0.005	1.38	1.36	1.40	1.46	1.41	1.39	1.48	1.59	1.44	1.41	1.50	1.63
15	0.070	0.010	1.53	1.48	1.53	1.58	1.56	1.53	1.61	1.72	1.60	1.56	1.65	1.78
15	0.070	0.015	1.72	1.66	1.70	1.76	1.77	1.72	1.81	1.92	1.83	1.76	1.84	1.98
15	0.070	0.020	2.00	1.91	1.96	2.02	2.07	1.99	2.08	2.20	2.15	2.06	2.12	2.29
15	0.070	0.025	2.39	2.28	2.31	2.38	2.51	2.39	2.47	2.61	2.62	2.49	2.55	2.70
Noise in climate information, θ														
5	0.070	0.015	1.72	1.68	1.75	1.81	1.77	1.77	1.95	2.10	1.83	1.81	2.05	2.32
10	0.070	0.015	1.72	1.67	1.73	1.80	1.77	1.73	1.85	1.98	1.83	1.78	1.91	2.10
15	0.070	0.015	1.72	1.66	1.70	1.76	1.77	1.72	1.81	1.92	1.83	1.76	1.84	1.98
20	0.070	0.015	1.72	1.66	1.69	1.74	1.77	1.71	1.77	1.87	1.83	1.76	1.79	1.90
25	0.070	0.015	1.72	1.66	1.68	1.72	1.77	1.71	1.75	1.83	1.83	1.76	1.77	1.84

Notes. Each entry in the table reports the optimal investment threshold for the benefit–cost ratio for the indicated parameter combination. All parameters other than those listed take their baseline values.

estimating the option value. If the decision-maker does delay investment, then she will continue to evaluate the project, gather relevant climatic and economic information, and use it to recalculate the option value—but each time she calculates the option value, she will assume that any delayed investment will occur at the best fixed future date. That is, the scope of this assumption is strictly limited to calculating an approximate value for the delay option.

Suppose it is currently date t and the decision-maker decides to exercise the incremental upgrade option at some fixed future date $t+s$. The present value of the investment payoff at the time investment occurs equals

$$\left(p_{t+s}v_b(t+s;q) + (1-p_{t+s})v_g(t+s;q)\right) x_{t+s} - c.$$

As shocks to p and x are independent, the implied date-t present value equals

$$Z(s) \equiv e^{-(r-\mu)s}\left(p_t v_b(t+s;q) + (1-p_t)v_g(t+s;q)\right) x_t - e^{-rs}c.$$

Therefore, if the decision-maker were to adopt a policy of investing at the best fixed future date, then the present value of the delay option would equal $\sup_{s>0} Z(s)$. We use this quantity as our approximation of the true option value, $V(t, p_t, x_t, q_t)$, in equation (7).

The decision-maker invests as soon as the present value of the benefits from investment is greater than or equal to the sum of the investment expenditure and $\sup_{s>0} Z(s)$. This rule can be expressed in terms of a threshold for x, as shown in the next lemma.

Lemma 6 *Suppose that $\mu \geq 0$. The present value of the benefits from exercising the incremental upgrade option at date t is greater than or equal to $c + \sup_{s>0} Z(s)$ if and only if x_t is greater than or equal to the threshold $\hat{x}^*(t, p_t)$, where*

$$\hat{x}^*(t, p_t) = \frac{rc}{1 - p_t F_b(q_t, t) - (1-p_t)F_g(q_t, t)},$$

where F_b and F_g are the cumulative distribution functions of rainfall intensity in the bad and good climate scenarios, respectively. ■

Note that the threshold in Lemma 6 depends on p_t. Thus, even though the underlying calculation of the value of the delay option assumes that investment occurs at a fixed future date, when implemented this rule results in an investment date that depends on the probability that the bad scenario holds. In particular, if the decision-maker delays investment, then she will take advantage of new information that arrives after this decision. The fixed-date assumption applies *only* to the calculation of the option value.

The investment threshold in Lemma 6 has a very natural interpretation. Lemma 2 implies that if the incremental upgrade option is exercised at date t, then the immediate benefit flow equals

$$\left(p_t(1 - F_b(q_t, t)) + (1-p_t)(1 - F_g(q_t, t))\right) x_t = \left(1 - p_t F_b(q_t, t) - (1-p_t)F_g(q_t, t)\right) x_t.$$

The investment rule in Lemma 6 is therefore equivalent to exercising the incremental upgrade option if and only if the immediate benefit flow is greater than or equal to rc. That is, for the situation considered here, the decision-maker just has to compare the initial benefit flow and the opportunity cost of the funds needed to invest, both measured over the next increment of time.

Each graph in Figure 9 plots the investment threshold for the economic factor as a function of the probability p for a different date t. The three curves correspond to three potential investment rules. The bottom (blue) curve shows the threshold when the decision-maker adopts a policy of investing as soon as the present value of the project's benefits is greater than or equal to the required investment expenditure. This "zero NPV" policy has the same form as the policy described by equation (7) if the value of the delay option is set equal to zero. The middle (yellow)

Figure 9: Three different investment thresholds in terms of x

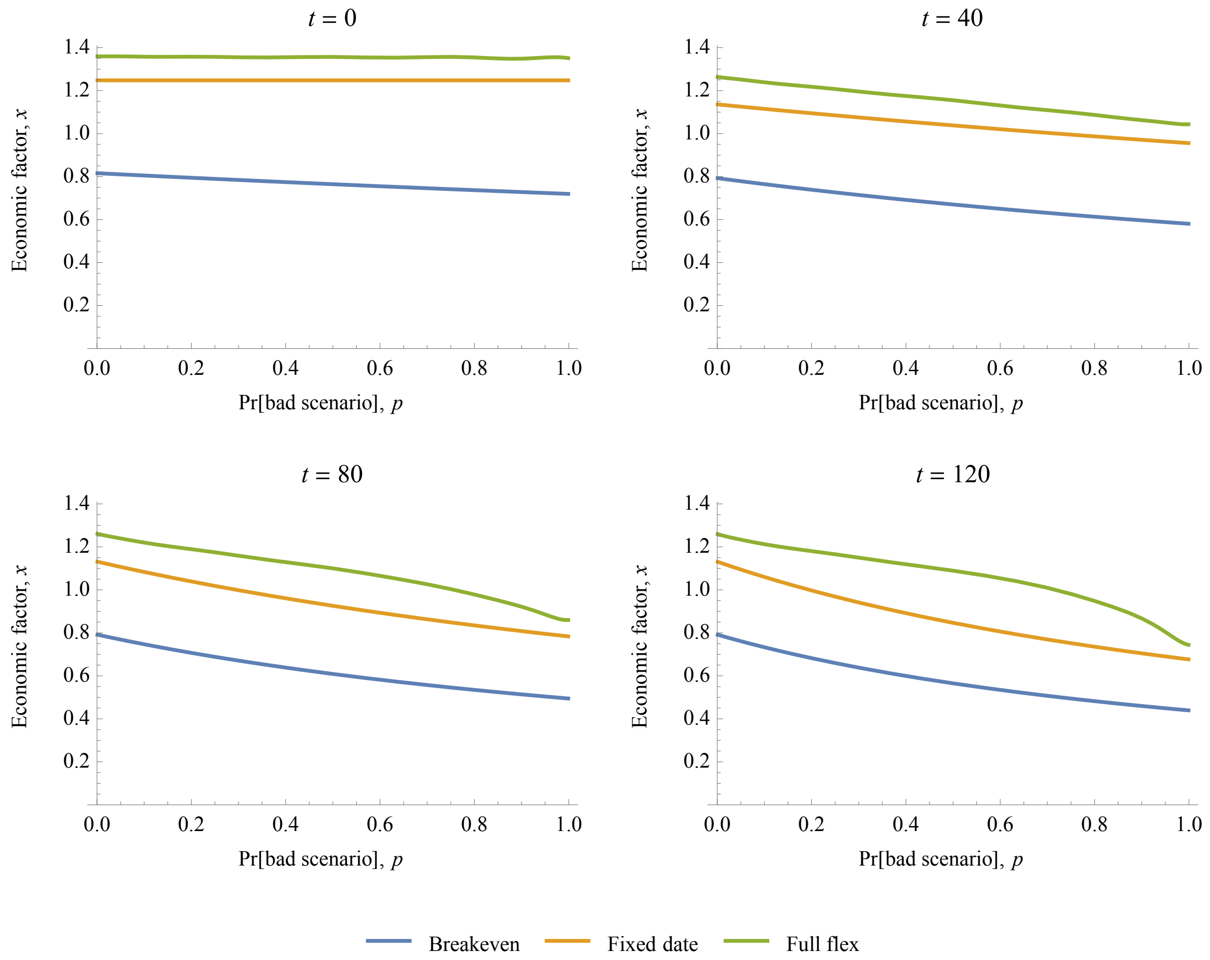

Notes. Each graph plots the investment threshold for the economic factor as a function of the probability p for a different date t. The blue curve shows the threshold when the decision-maker adopts a policy of investing as soon as the present value of the project's benefits is greater than or equal to the required investment expenditure. The yellow curve corresponds to the approximate policy in Lemma 6, in which the value of the delay option is calculated assuming investment occurs at the best fixed future date. Finally, the green curve corresponds to the optimal investment policy.

curve corresponds to the approximate policy in Lemma 6. It also has the same form as the policy described by equation (7), but in this case the value of the delay option is calculated assuming that investment occurs at the best fixed future date. Finally, the top (green) curve corresponds to the optimal investment policy, in which the option value in equation (7) is calculated assuming the decision-maker adopts an optimal investment policy in the future. Figure 10 shows the same information, but in terms of the benefit–cost ratio.

The three investment policies represented in Figures 9 and 10 have progressively more demanding investment thresholds. To see why, consider equation (7). The zero-NPV policy sets $V = 0$. In contrast, the approximate policy sets it equal to the present value of the investment payoff assuming that investment occurs at the best fixed future date. This quantity is positive, explaining why the second threshold is greater than the first. Similarly, the optimal policy sets V equal to the present value of the investment payoff assuming investment at the optimal future date. This quantity must be greater than or equal to the present value assuming investment occurs at the best *fixed* date, which is why the third threshold is at least as great as the second

Figure 10: Three different investment thresholds in terms of the benefit–cost ratio

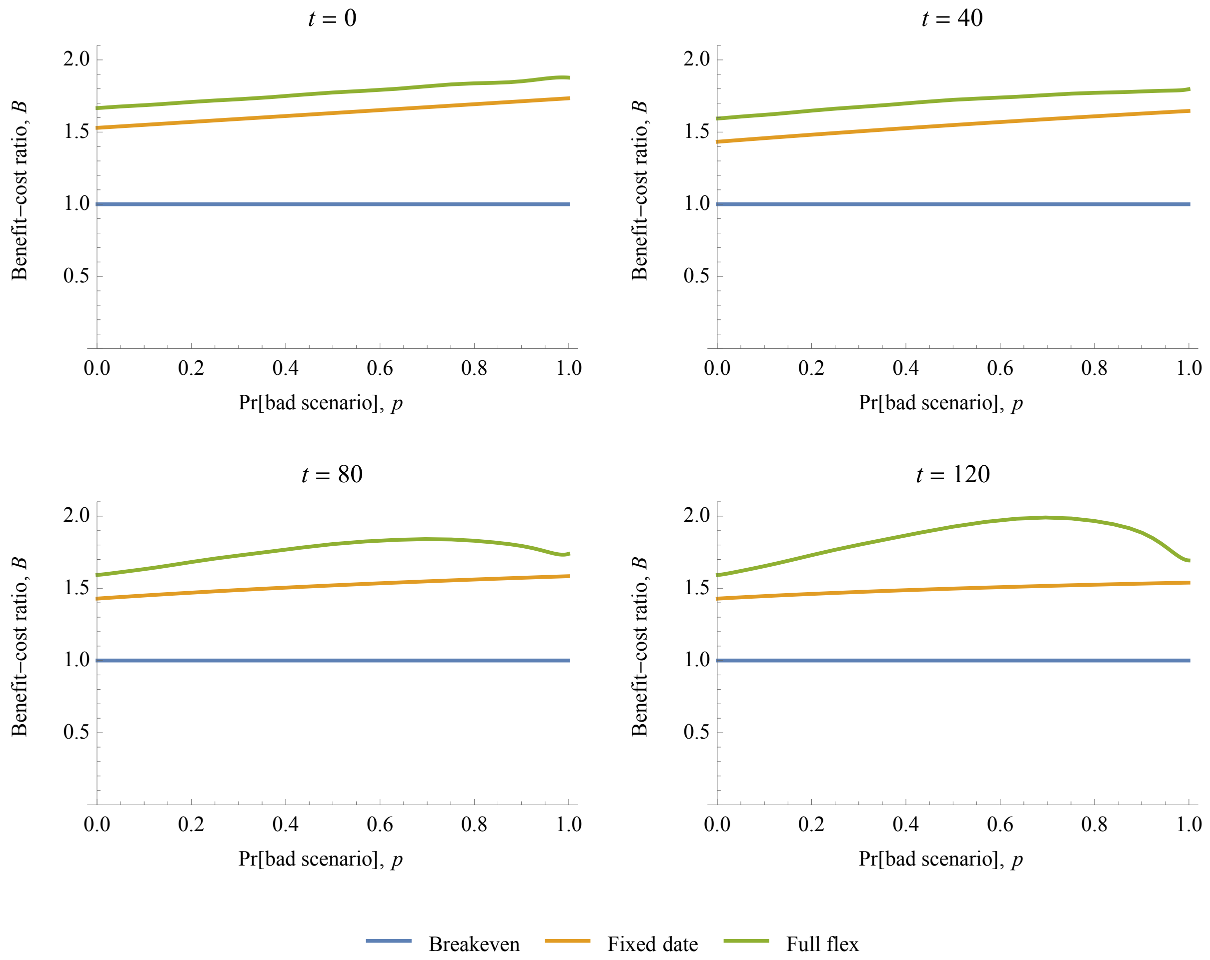

Notes. Each graph plots the investment threshold for the benefit–cost ratio as a function of the probability p for a different date t. The blue curve shows the threshold when the decision-maker adopts a policy of investing as soon as the present value of the project's benefits is greater than or equal to the required investment expenditure. The yellow curve corresponds to the approximate policy in Lemma 6, in which the value of the delay option is calculated assuming investment occurs at the best fixed future date. Finally, the green curve corresponds to the optimal investment policy.

threshold. The approximate policy underestimates the value of the delay option and therefore accelerates investment. The question we will investigate shortly is whether this acceleration leads to an economically significant reduction in overall welfare.

Even under the alternative policy, investment can be delayed significantly past the date when the present value of the benefits from investment equals the required investment expenditure. Consider the situation at date $t_0 = 0$ if the decision-maker believes the two climate scenarios are equally likely ($p_0 = 1/2$). This is the situation shown in the top left-hand graphs in Figures 9 and 10. The breakeven level of the economic factor is $x_0 = 0.76$. The very steep time profile of avoided flooding costs (evident in the left-hand graph in Figure 6), means that the short-term benefits from investing are low for a breakeven project. For the example here, if investment is delayed for exactly five years, the expected flooding costs during the period leading up to the delayed investment are just 0.1430, whereas the delay reduces the present value of the investment expenditure by 0.2212. For a ten-year delay, the corresponding figures are 0.2670 and 0.3935. Thus, significant delay is appropriate even for a project currently in a breakeven position. In this situation, if x equals the breakeven level, then the best fixed investment date is actually $t^* = 24.14$. The present value of investing at this date—no matter what happens between now and then—is 0.1691. In summary, the high growth rate in the initial benefit flow means that it can be better to delay investment a considerable length of time beyond the date when the present value of the project's benefits cover the required investment expenditure even under the suboptimal policy in which option valuations assume a fixed investment date. This explains the large gaps between the blue (breakeven policy) and yellow (best fixed date policy) curves in Figures 9 and 10.

An important question to consider is whether incorporating the value of the arrival of new information in the value of the delay option leads to a large additional delay. The gaps between the yellow and green investment thresholds in Figures 9 and 10 suggest that, for the next few decades, the resulting delays are moderate. When the gaps subsequently increase, it is for levels of p that indicate considerable remaining uncertainty about the true climate scenario. The distributions of p in Figure 4 suggest that these situations are unlikely to occur.[24] In summary, incorporating the value of the arrival of new information in the value of the delay option leads to moderate delays in investment, but is unlikely to lead to large additional delays.

The intuition for this behaviour can be found in the two values of the delay option. The approximate value assumes the decision-maker will wait and invest at a fixed future date, whereas the correct value assumes she will follow an optimal future investment policy that exploits future climatic and economic information as it arrives. That is, the first option value assumes the decision-maker delays the investment, but not the decision when to invest—that is made now. In contrast, the second option value assumes the decision-maker delays the decision as well. The second option value will only be significantly higher than the first one if there is a significant risk of new information arriving that would cause the decision-maker to regret the investment decision that is implicit in the approximate option value. For the situation we are examining, the prospect for such bad news is low, at least for the next few decades. Recall from the discussion following Lemma 6 that under the approximate approach, the decision-maker invests when the flow of benefits equals rc. By this stage, the net present value of investing is much greater than zero. Due to the low volatility of the economic factor and—for the next few decades—the low volatility of the probability p, there is very little chance that any subsequent bad news will be bad enough to make this net present value negative. In other words, if the decision-maker uses the approximate rule, then she delays investment so long that there is little chance that she will subsequently regret investing. It is not surprising that the added value from delaying the decision is so small.

This explanation is closely related to the "bad news principle" of irreversible investment (Bernanke, 1983). When investment can be delayed, the decision-maker's willingness to invest immediately depends on the potential for future events to strand the asset. That is, of all possible future outcomes, it is only the unfavourable ones (the "bad news") that affect the current propensity to invest. This happens because a decision-maker who does not invest sacrifices a short-term benefit flow in return for the option *not to invest* in the future. This option will have no value if future information favours the project. Its value comes from the prospect that future information will lead the decision-maker to regret having invested previously. For the situation we are examining, there is very little chance that any subsequent bad news will be bad enough to make this net present value negative. In other words, if the decision-maker uses the approximate rule, then she delays investment so long that there is little chance that she will subsequently

[24] The assumption that x evolves according to geometric Brownian motion implies that economic volatility is constant over time. In contrast, the volatility of the probability p varies over time. The pattern of widening gaps between yellow and green curves in Figures 9 and 10 suggests that the results are being driven by climate volatility. Further evidence can be found along the $p = 0$ and $p = 1$ boundaries, which show what happens when there is no climate volatility. Consistent with economic volatility being constant, the gap between the yellow and green thresholds is almost constant over time along these boundaries.

regret investing.[25]

7.2 Welfare implications of adopting an approximate policy

Our aim is to identify a simple decision-making rule that keeps the welfare costs of deviating from optimal decision-making low. The relevant issue is whether welfare falls significantly if a simple rule is used, not whether the resulting policy is greatly different from the optimal one. Even if use of a simple tool results in a large error in investment timing, if the loss of welfare is small then the decision-making rule may be suitable for practical use. In contrast, for some projects a simple rule may potentially lead to a small error in investment timing but a large loss of welfare. For these projects, the simple rule is not appropriate. This subsection therefore uses the present value of future welfare to evaluate the performance of the policy described in Section 7.1.

For this comparison, for each policy we need to calculate welfare, as measured by the present value of the investment option. Welfare under the optimal policy is already calculated as part of the construction of the optimal policy, when we solve the variational inequalities in Lemma 5. Construction of welfare under the approximate policy is more subtle. The policy is derived by assuming that each time the decision-maker calculates the value of the investment option, she assumes that investment occurs at the best fixed future date. However, this value is *not* the value of the option if she actually adopts this policy. Consider the numerical example discussed earlier when $t = 0$, the two scenarios are equally likely ($p_0 = 1/2$), and the economic factor is equal to the breakeven level ($x_0 = 0.7647$). In this situation, the present value of the investment option implied by the best fixed future investment date was 0.1691. However, if the decision-maker actually follows the approximate policy described in Lemma 6, then the present value of the investment option will equal 0.2048. The difference between the two present values reflects the value of the decision-maker's ability to respond to future information. This flexibility has value, even though her response will not be fully optimal. If she followed the optimal investment policy instead, the value of the investment option would equal 0.2076 in this case.

Fortunately, calculating the value of welfare under the approximate policy is straightforward. We simply calculate the present value of the investment option by solving the standard differential equation using the investment threshold in Lemma 6 in place of the optimal threshold. The starting point is Lemma 5, which describes the variational inequalities we need to solve to find the optimal policy. Now, however, we are simply calculating the present value of welfare associated with a given policy. Therefore, the present value function satisfies the differential equation

$$\mathcal{D}V(t,p,x;q) - rV(t,p,x;q) = 0$$

in the waiting region (that is, where $x < \hat{x}^*(t,p)$) and it satisfies

$$V(t,p,x;q) = \big(pv_b(t;q) + (1-p)v_g(t;q)\big)\,x - c$$

in the investment region (that is, where $x \geq \hat{x}^*(t,p)$).

Figure 11 plots the value of the investment option at date $t = 80$ as a function of the level of the economic factor for three different levels of the probability that the bad climate scenario holds. The solid curves show the value of the investment option if the decision-maker adopts the optimal investment policy and the dashed curves show its value if she adopts the alternative policy in Lemma 6 instead. The dashed grey lines show the investment payoffs. The lower and upper points on each curve show the alternative and optimal investment thresholds, respectively. In all cases, the height of the solid is curve is greater than or equal to the height of the corresponding dotted curve. Investment occurs immediately when x is above the investment threshold, so in this region the value of the option equals the investment payoff. As investment

[25] In this context, bad news is that climate change is less severe than expected or that the economic factor grows more slowly than expected. In both cases, the project's benefit flow is smaller than expected.

Figure 11: Value of the investment option under two investment policies

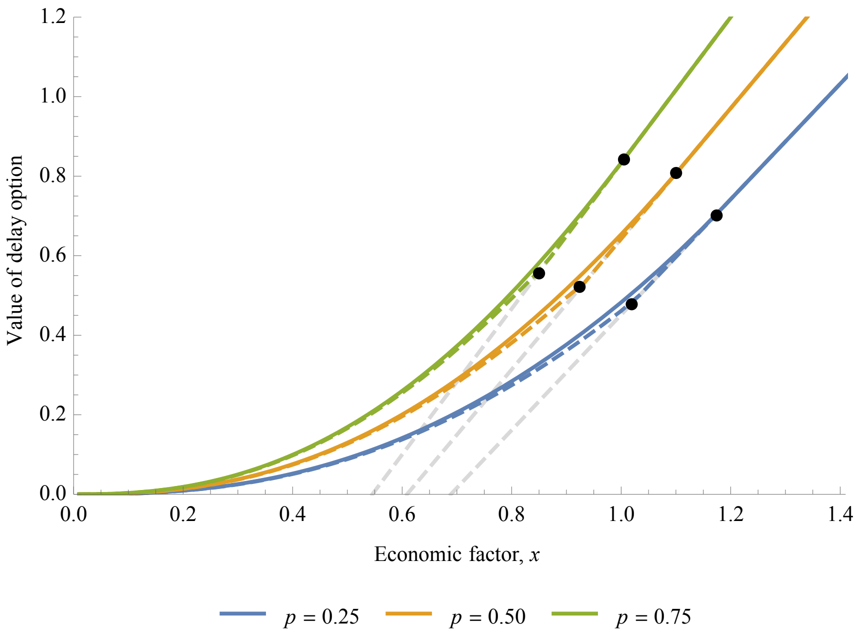

Notes. The graph plots the value of the investment option at date $t = 80$ as a function of the level of the economic factor for three different levels of the probability that the bad climate scenario holds. The solid curves show the option's value if the decision-maker adopts the optimal investment policy and the dashed curves show its value if she adopts the alternative policy in Lemma 6. The dashed grey lines show the investment payoffs. The lower and upper points on each curve show the alternative and optimal investment thresholds, respectively.

occurs too soon under the alternative rule, the dashed curves converge to the investment payoff more quickly than their solid counterparts. This explains why the solid and dotted curves start to diverge as x increases. The distance between the two curves at the point where the dashed curve joins the investment payoff is one measure of the welfare that is lost if the decision-maker adopts the alternative rule in Lemma 6: it is the present value of the investment option that is lost, measured at the (premature) investment date.

Table 3 summarises the results of sensitivity analysis into the welfare losses that result from adopting the simple investment rule in Lemma 6. Each entry reports the present value of welfare that is lost, measured at the (premature) investment date, when following this investment rule. That is, it reports

$$100\left(1-\frac{\big(p_t v_b(t;q)+(1-p_t)v_g(t;q)\big)\,\hat{x}^*-c}{V(t,p,\hat{x}^*)}\right),$$

where $\hat{x}^*$ is the alternative investment threshold in Lemma 6 and V is the function in Lemma 5. In terms of Figure 11, it is the distance between the two curves at the point where the dashed curve joins the investment payoff, expressed as a percentage of the height of the solid curve. The first row of the table shows results for the baseline parameter values. For at least the first 40 years, the welfare losses are small and largely insensitive to the probability that the bad climate scenario holds. Even after 120 years, following the alternative investment rule results in welfare losses of less than ten percent. This is consistent with the role of the bad news principle. For larger values of t, the spread between $\lambda_b(t)$ and $\lambda_g(t)$ is greater, which (from Lemma 1) means the noisy signal of rainfall intensity is more informative and the volatility of p is greater. This increases the probability that incremental system upgrades will subsequently be stranded once new climatic information becomes available, which makes delaying the investment *decision* more

Table 3: Welfare lost when using the approximate investment policy

θ	σ	μ	$p = 0.25$				$p = 0.50$				$p = 0.75$			
			$t = 0$	$t = 40$	$t = 80$	$t = 120$	$t = 0$	$t = 40$	$t = 80$	$t = 120$	$t = 0$	$t = 40$	$t = 80$	$t = 120$
Baseline case														
15	0.070	0.015	1.6	3.1	5.1	7.7	1.4	2.6	5.4	9.1	1.3	2.0	4.2	7.0
Volatility of economic factor, σ														
15	0.050	0.015	0.5	1.2	2.9	5.3	0.4	1.1	3.7	7.5	0.4	0.8	2.8	5.8
15	0.060	0.015	1.0	2.1	3.9	6.5	0.8	1.8	4.5	8.3	0.7	1.3	3.4	6.3
15	0.070	0.015	1.6	3.1	5.1	7.7	1.4	2.6	5.4	9.1	1.3	2.0	4.2	7.0
15	0.080	0.015	2.4	4.3	6.3	9.0	2.1	3.6	6.4	10.1	1.9	2.9	5.1	7.8
15	0.090	0.015	3.4	5.6	7.7	10.3	3.0	4.8	7.5	11.0	2.6	3.9	6.1	8.7
Drift of economic factor, μ														
15	0.070	0.005	12.0	22.9	30.3	37.7	10.0	18.1	27.5	36.7	8.4	13.8	21.5	29.0
15	0.070	0.010	4.2	8.1	12.1	16.8	3.6	6.6	11.9	18.1	3.1	5.1	9.2	13.8
15	0.070	0.015	1.6	3.1	5.1	7.7	1.4	2.6	5.4	9.1	1.3	2.0	4.2	7.0
15	0.070	0.020	0.7	1.3	2.2	3.6	0.6	1.1	2.6	4.7	0.5	0.9	2.0	3.7
15	0.070	0.025	0.3	0.5	1.0	1.8	0.3	0.5	1.2	2.5	0.2	0.4	1.0	2.0
Noise in climate information, θ														
5	0.070	0.015	1.6	3.6	6.7	9.7	1.4	3.4	7.5	11.6	1.3	2.6	5.7	8.6
10	0.070	0.015	1.6	3.3	5.8	8.7	1.4	2.9	6.4	10.4	1.3	2.2	4.9	7.9
15	0.070	0.015	1.6	3.1	5.1	7.7	1.4	2.6	5.4	9.1	1.3	2.0	4.2	7.0
20	0.070	0.015	1.6	3.0	4.6	6.8	1.4	2.5	4.7	8.0	1.3	2.0	3.6	6.2
25	0.070	0.015	1.6	2.9	4.2	6.2	1.4	2.4	4.2	7.1	1.2	1.9	3.2	5.4
Extreme case														
5	0.090	0.005	18.7	31.7	39.0	45.4	16.0	26.3	35.2	42.8	13.7	21.1	28.5	34.7

Notes. Each entry reports the present value of welfare that is lost, measured at the (premature) investment date, when following the simple investment rule in Lemma 6. In terms of Figure 11, it is the distance between the two curves at the point where the dashed curve joins the investment payoff, expressed as a percentage of the height of the solid curve.

valuable. In other words, the value of the delay option is more sensitive to the arrival of new climatic information when t is larger, so making no allowance for that information leads to a greater error when approximating the value of the delay option.

Each row in the first main panel of Table 3 shows results for a different level of volatility of the economic factor.[26] The welfare performance of the alternative investment rule deteriorates as the economic factor becomes more volatile. As greater economic volatility increases the likelihood that the decision-maker receives "bad news" after investing, the reduced performance of the alternative rule for large σ is consistent with the bad news principle. The second main panel of Table 3 paints a similar picture. Each row shows results for a different expected growth rate in the economic factor. Faster expected growth decreases the probability that incremental system upgrades will subsequently be stranded, which decreases the importance of delaying the investment decision and improves the performance of the alternative investment rule.

The third panel examines the effect of the noise contained in the climatic signal. A larger value of θ implies more noise in the climatic signal, and therefore less volatility in the probability that the bad scenario holds. This reduces the probability of post-investment stranding, and improves the performance of the alternative investment rule. The effect is insignificant for small values of t, because in this case the possible paths for rainfall intensity $\lambda(t)$ are so similar that there is almost no volatility in p. However, for large values of t, there is a noticeable improvement in the alternative rule's performance when the climatic signal is noisier.

The three panels show that changing individual parameters do not lead to significant welfare losses from use of the alternative investment rule. As a final check, the bottom row of Table 3 sets each of the three key parameters equal to the value for which the welfare losses are greatest. That is, the economic factor's volatility is high, its expected growth rate is low, and the climate signal has relatively little noise. The bottom row shows that in this case, the welfare losses resulting from using the alternative rule are almost 20 percent at date 0; they can reach 30 percent by date 40. Thus, although the alternative investment rule performs reasonably well over a wide range of parameter values, there will likely still be situations where it results in a significant reduction in welfare. Thus, if the economic factor is volatile enough and grows slowly enough, and if there is enough noise in the climatic signal, then there is still a place for "full" real options analysis.

7.3 Policy implications

We have constructed an optimal investment policy for upgrading an asset to cope with climate change, and also constructed a simpler alternative policy for deciding when to invest. Both of these contributions have implications for how climate-change adaptation policy should be implemented. The optimal policy shows that the value of delaying investment can be economically significant and that the optimal delays can be long. In many situations, it can be optimal to delay investment until the present value of the investment benefits exceeds the required investment expenditure by 75%; in some cases, the premium is even larger. The wedge between the minimum acceptable present value of the investment benefits and the investment expenditure equals the value of the real option to delay investment. Standard cost–benefit analysis (which ignores the value of the delay option and therefore leads to investment as soon as the present value of a project's benefits exceed its cost) will induce poor investment decisions. Unfortunately, optimal decision-making requires estimating the value of the delay option, which requires a "deep look into the future" (Bernanke, 1983). Fortunately, Section 7.2 suggests that there is a simple alternative approach available to decision-makers that in many situations results in investment timing that is close to optimal.

As it is described in Section 7.1, this method involves estimating the value of the delay option by assuming that delayed investment will occur at the best fixed future date. This option value is

[26] Based on the evidence in Section 4.2, the range of values of μ and σ in Table 3 covers plausible values of the two parameters.

repeatedly updated, so that the eventual investment date will be determined by future climatic and economic information. That is, the fixed-date assumption is used *only* for calculating the present value of the delay option. There is a much simpler way to implement this approach, which involves only a minor modification to the way that cost–benefit analysis is typically carried out. Standard cost–benefit analysis involves estimating the present value of the benefits of immediate investment and comparing them with the present value of the costs. If the former is greater than the latter, then immediate investment is recommended. That is, the implicit alternative to investing immediately is to not invest at all. The only change that is required to implement the method introduced in Section 7.1 is to carry out this analysis for *all* possible (fixed) investment dates. That is, also calculate the present value of benefits minus costs assuming that investment is carried out n years from now, for $n = 1, 2, 3, \ldots$. If the net present value from investing immediately is greater than or equal to the largest of these delayed net present values, then invest immediately, otherwise wait and reevaluate the project in the future (using all new climatic and economic information that has arrived in the meantime). This process is repeated at regular intervals until immediate investment is preferred to delaying until some fixed future date.

The sensitivity analysis in Table 3 shows this approach will give near-optimal behaviour in many situations. However, if the economic factor has high volatility and a small expected growth rate, and the climate signal does not contain too much noise, then using the alternative policy to make adaptation decisions could lead to reasonably significant welfare losses. Therefore, when these conditions are met, analysts should either use "full" real options analysis or find an approximation with better performance than the one considered here. The situations in which this will be necessary are characterised by a relatively high probability of future "bad news" resulting in the investment being stranded. In this context, "bad news" means that the scale of climate change might not be as severe as was expected, or the value of the assets vulnerable to flooding might be lower than was expected, at the time the investment occurred.

8 Conclusion

This paper analyses climate-change adaptation investment using a new theoretical framework that incorporates climatic and economic volatility in a single model. The central feature of this framework is the way that it models uncertainty about future climate change and how this uncertainty will evolve over time. The paper derives an optimal investment policy for a particular class of adaptation investments and finds that the value of the investment-timing options embedded in projects of this type is economically significant. Finally, the paper shows that most of the net benefits of investment can be captured if investment timing is decided using a simple alternative to "full" real options analysis. If the decision-maker uses this alternative rule, then she invests as soon as no fixed future investment date implies a greater net present value than investing immediately.

Some components of the model in this paper are highly stylised, but with modifications the framework could be suitable for practitioner use—at least until richer models of climate-change uncertainty become available. For example, when the model is used to analyse investment in flood risk management projects, the project benefits take the form of avoided flooding costs. In this paper, these are calculated using a stylised macroscopic damage function, but in practical applications this could be replaced by a damage function that is constructed using detailed hydrological analysis. The framework presented in this paper is sufficiently flexible to accommodate such enhancements without greatly increasing the complexity of the required calculations.

There are two areas in which the model could be enhanced. The simplest extension would be to apply the existing framework to problems involving alternative option structures. In this paper, the capacity of the underlying project is a continuous variable and the decision-maker decides how quickly to increase capacity. Many adaptation projects will have more complicated structures than this—such as the option to build in flexibility to make future investment more efficient—and we cannot assume that the alternative rule will still be effective in such situations. A more substantive extension would involve increasing the number of climate scenarios. Using two scenarios, as in the current version of the model, probably gives an adequate model of climate uncertainty over the short-to-medium term, but for very long-term projects it would be better to allow a greater variety of potential outcomes.

A Proofs

A.1 Heuristic proof of Lemma 1

If $p = \Pr[\lambda = \lambda_b]$ is the decision-maker's current subjective probability that the bad scenario holds, then Bayes' Theorem implies that her updated subjective probability is

$$\begin{aligned} p' = \Pr\left[\lambda = \lambda_b \middle| \frac{\Delta z}{\Delta t}\right] &= \frac{f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_b\right) \cdot \Pr[\lambda = \lambda_b]}{f\left(\frac{\Delta z}{\Delta t}\right)} \\ &= \frac{f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_b\right) \cdot p}{f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_b\right) \cdot p + f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_g\right) \cdot (1-p)}, \end{aligned}$$

where

$$f\left(\frac{\Delta z}{\Delta t}\middle|\lambda\right) = \frac{1}{\theta}\sqrt{\frac{\Delta t}{2\pi}} e^{\frac{-\Delta t}{2}\left(\frac{\frac{\Delta z}{\Delta t} - \lambda}{\theta}\right)^2}$$

is the density function of $\Delta z/\Delta t$, conditional on λ. The change in the probability that the bad scenario holds therefore equals

$$\Delta p \equiv p' - p = p(1-p)\left(\frac{f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_b\right) - f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_g\right)}{f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_b\right) \cdot p + f\left(\frac{\Delta z}{\Delta t}|\lambda = \lambda_g\right) \cdot (1-p)}\right).$$

Substituting in $\Delta z/\Delta t = \lambda + \theta\varepsilon/\sqrt{\Delta t}$ implies that

$$\begin{aligned} \Delta p =& p(1-p)\left(\frac{e^{\frac{-1}{2}\left(\left(\frac{\lambda-\lambda_b}{\theta}\right)\sqrt{\Delta t}+\varepsilon\right)^2} - e^{\frac{-1}{2}\left(\left(\frac{\lambda-\lambda_g}{\theta}\right)\sqrt{\Delta t}+\varepsilon\right)^2}}{e^{\frac{-1}{2}\left(\left(\frac{\lambda-\lambda_b}{\theta}\right)\sqrt{\Delta t}+\varepsilon\right)^2} \cdot p + e^{\frac{-1}{2}\left(\left(\frac{\lambda-\lambda_g}{\theta}\right)\sqrt{\Delta t}+\varepsilon\right)^2} \cdot (1-p)}\right) \\ =& \left(\frac{p(1-p)(\lambda_b - \lambda_g)\varepsilon}{\theta}\right)\sqrt{\Delta t} \\ &+ \left(\frac{p(1-p)(\lambda_b - \lambda_g)\left(2(\lambda - (p\lambda_b + (1-p)\lambda_g)) - (1-2p)(1-\varepsilon^2)(\lambda_b - \lambda_g)\right)}{\theta^2}\right)\Delta t \\ &+ o(\Delta t) \end{aligned}$$

as $\Delta t \to 0$. Immediately before we observe z, we have $E[\varepsilon^2] = 1$ and $E[\lambda] = p\lambda_b + (1-p)\lambda_g$, so that the expected value of the coefficient of Δt is zero. Therefore

$$\Delta p = \frac{p(1-p)(\lambda_b - \lambda_g)}{\theta}\left(\sqrt{\Delta t}\,\varepsilon\right) + o(\Delta t),$$

which motivates the stochastic process for p in equation (2).

A.2 Proof of Lemma 2

Suppose the annual maximum one-day rainfall has density function $f(R,t)$ at date t. If the stormwater system has capacity q, then the date-t expected flooding cost equals $x_t C(q)$, where

$$C(q) \equiv E_t[\max\{0, R-q\}] = \int_q^\infty (R-q) f(R,t)\, dR.$$

This implies that

$$C'(q) = -\int_q^\infty f(R,t)\,dR,$$

so that the avoided flooding cost for an incremental system upgrade equals

$$-x_t C'(q) = x_t \int_q^\infty f(R,t)\,dR = x_t \Pr_t[\tilde{R} > q] = x_t(1 - F(q,t)).$$

A.3 Proof of Lemma 3

Suppose the state of the system is currently $(t, p, x; q)$. If the decision-maker increases capacity to $q' \geq q$ and then waits dt units of time before reevaluating the investment programme, then the present value of all relevant costs equals

$$c(q'-q) + (p\Gamma_b(q',t) + (1-p)\Gamma_g(q',t))x\,dt + e^{-r\,dt}E[W(t+dt, p+dp, x+dx; q')].$$

She chooses q' to minimise this quantity, which we can rewrite as

$$\begin{aligned} U(t,p,x;q,q') \equiv& W(t,p,x;q') + c(q'-q) \\ &+ \Big(\mathcal{D}W(t,p,x;q') - rW(t,p,x;q') + (p\Gamma_b(q',t) + (1-p)\Gamma_g(q',t))x\Big)\,dt + o(dt) \end{aligned}$$

as $dt \to 0$. In the waiting region, this quantity is maximised by setting $q' = 0$; that is, $\frac{\partial U}{\partial q'}$ is positive at $q' = q$ and $W(t,p,x;q) = U(t,p,x;q,q)$. These conditions imply that

$$W_q(t,p,x;q) + c > 0$$

and

$$\mathcal{D}W(t,p,x;q) - rW(t,p,x;q) + (p\Gamma_b(q,t) + (1-p)\Gamma_g(q,t))x = 0,$$

respectively. On the other hand, in the investment region it is optimal to increase capacity to q^* satisfying

$$W_q(t,p,x;q^*) + c = 0,$$

so that

$$W(t,p,x;q) = U(t,p,x;q,q^*) = W(t,p,x;q^*) + c(q^*-q).$$

As q^* is independent of q, this implies that

$$W_q(t,p,x;q) = -c.$$

In addition, $W(t,p,x;q) < U(t,p,x;q,q)$, which implies that

$$\mathcal{D}W(t,p,x;q) - rW(t,p,x;q) + (p\Gamma_b(q,t) + (1-p)\Gamma_g(q,t))x > 0.$$

Thus, W satisfies the variational inequalities in Lemma 3.

A.4 Proof of Lemma 4

From Lemma 3, inside the waiting region,

$$W_q(t,p,x;q) + c > 0$$

and

$$\mathcal{D}W(t,p,x;q) - rW(t,p,x;q) + (p\Gamma_b(q,t) + (1-p)\Gamma_g(q,t))x = 0.$$

Differentiating the differential equation with respect to capacity q shows that, inside the waiting region, W_q satisfies

$$\mathcal{D}W_q(t,p,x;q) - rW_q(t,p,x;q) + \left(p\frac{\partial \Gamma_b(q,t)}{\partial q} + (1-p)\frac{\partial \Gamma_g(q,t)}{\partial q}\right)x = 0.$$

Now consider the situation inside the investing region. If the decision-maker does not invest, then the present value of all relevant costs will equal

$$(p\Gamma_b(q,t) + (1-p)\Gamma_g(q,t))x\,dt + e^{-r\,dt}E[W(t+dt,p+dp,x+dx;q)].$$

In contrast, if she increases capacity to some level q' (which need not be the optimal level), then the present value of all relevant costs will equal

$$c(q'-q) + (p\Gamma_b(q',t) + (1-p)\Gamma_g(q',t))x\,dt + e^{-r\,dt}E[W(t+dt,p+dp,x+dx;q')].$$

As some investment is optimal, the second present value will be less than the first one provided $q'-q$ is sufficiently small. That is,

$$\begin{aligned} &c(q'-q) + (p\Gamma_b(q',t) + (1-p)\Gamma_g(q',t))x\,dt + e^{-r\,dt}E[W(t+dt,p+dp,x+dx;q')] \\ &< (p\Gamma_b(q,t) + (1-p)\Gamma_g(q,t))x\,dt + e^{-r\,dt}E[W(t+dt,p+dp,x+dx;q)], \end{aligned}$$

which implies that

$$\begin{aligned} 0 >& c + \left(p\left(\frac{\Gamma_b(q',t)-\Gamma_b(q,t)}{q'-q}\right) + (1-p)\left(\frac{\Gamma_g(q',t)-\Gamma_g(q,t)}{q'-q}\right)\right)x\,dt \\ &+ e^{-r\,dt}E\left[\frac{W(t+dt,p+dp,x+dx;q')-W(t+dt,p+dp,x+dx;q)}{q'-q}\right]. \end{aligned}$$

Taking the limit as $q' \to q$ shows that

$$0 > c + (p\Gamma_{b,q}(q,t) + (1-p)\Gamma_{g,q}(q,t))\,x\,dt + e^{-r\,dt}E\left[W_q(t+dt,p+dp,x+dx;q)\right],$$

which in turn implies that

$$0 > c + W_q(t,p,x;q) + \big((p\Gamma_{b,q}(q,t) + (1-p)\Gamma_{g,q}(q,t))\,x + \mathcal{D}W_q(t,p,x;q) - rW_q(t,p,x;q)\big)\,dt.$$

The sum of the first two terms on the right-hand side equals zero, which implies that

$$(p\Gamma_{b,q}(q,t) + (1-p)\Gamma_{g,q}(q,t))\,x - rW_q(t,p,x;q) + \mathcal{D}W_q(t,p,x;q) < 0$$

in the investing region. This completes the proof that W_q satisfies the variational inequalities in Lemma 4.

A.5 Proof of Lemma 5

The first result in Lemma 4 shows that

$$V(t,p,x;q) = \big(pv_b(t;q) + (1-p)v_g(t;q)\big)\,x + W_q(t,p,x;q) \geq \big(pv_b(t;q) + (1-p)v_g(t;q)\big)\,x - c.$$

The second result shows that

$$\begin{aligned} \mathcal{D}V(t,p,x;q) - rV(t,p,x;q) &= \mathcal{D}W_q(t,p,x;q) - rW_q(t,p,x;q) \\ &\quad + \mathcal{D}\left(\big(pv_b(t;q) + (1-p)v_g(t;q)\big)\,x\right) \\ &\quad - r\left(pv_b(t;q) + (1-p)v_g(t;q)\right)x \\ &\leq -(p\Gamma_{b,q}(q,t) + (1-p)\Gamma_{g,q}(q,t))\,x \\ &\quad + \mathcal{D}\left(\big(pv_b(t;q) + (1-p)v_g(t;q)\big)\,x\right) \\ &\quad - r\left(pv_b(t;q) + (1-p)v_g(t;q)\right)x \\ &= 0, \end{aligned}$$

where the final step uses the result that

$$\mathcal{D}v_i(t;q) = \Gamma_{i,q}(q,t) + (r-\mu)v_i(t;q)$$

for $i = b, g$.

A.6 Proof of Lemma 6

Note that

$$\begin{aligned} Z'(s) &= re^{-rs}c + e^{-(r-\mu)s}\left(p_t v_b'(t+s;q) + (1-p_t)v_g'(t+s;q)\right)x_t \\ &\quad - (r-\mu)e^{-(r-\mu)s}\left(p_t v_b(t+s;q) + (1-p_t)v_g(t+s;q)\right)x_t. \end{aligned}$$

The definitions of $v_b(t;q)$ and $v_g(t;q)$ imply that

$$v_b'(t;q) = (r-\mu)v_b(t;q) + F_b(q,t) - 1$$

and

$$v_g'(t;q) = (r-\mu)v_g(t;q) + F_g(q,t) - 1.$$

It follows that

$$e^{rs}Z'(s) = rc - e^{\mu s}\left(1 - p_t F_b(q,t+s) - (1-p_t)F_g(q,t+s)\right)x_t.$$

As $F_b(R,t)$ and $F_g(R,t)$ are decreasing functions of t, it follows that $e^{rs}Z'(s)$ is a decreasing function of s (provided that $\mu \geq 0$). Therefore, if $Z'(0) \leq 0$ then $Z'(s) \leq 0$ for all $s \geq 0$, which implies that

$$\sup_{s\geq 0} Z(s) = Z(0) = \left(p_t v_b(t;q) + (1-p_t)v_g(t;q)\right)x_t - c.$$

In this case, the decision-maker will invest immediately. In contrast, if $Z'(0) > 0$, then the first-order condition $Z'(s) = 0$ has a unique positive solution, yielding the best fixed future investment date, $s^* > 0$. This implies that

$$\sup_{s\geq 0} Z(s) = Z(s^*) > Z(0) = \left(p_t v_b(t;q) + (1-p_t)v_g(t;q)\right)x_t - c,$$

in which case the decision-maker will delay investment.

These results demonstrate that the decision-maker will invest immediately if and only if $Z'(0) \leq 0$, which holds if and only if

$$\left(1 - p_t F_b(q,t) - (1-p_t)F_g(q,t)\right)x_t \geq rc.$$

This policy can be written in terms of a threshold for x:

$$x_t \geq \frac{rc}{1 - p_t F_b(q,t) - (1-p_t)F_g(q,t)}.$$

B Numerical methods

The model is formulated in (t,p,x)-space, but we change coordinates to (t,y,z) when we solve it numerically, where $y = \log(p/(1-p))$ and $z = \log x$. That is, we replace the probability that the bad scenario holds with the corresponding log-odds ratio, and the economic factor with its natural logarithm. In the waiting region, the value of the incremental upgrade option satisfies the partial differential equation

$$0 = V_t + \frac{1}{2}\left(\frac{\lambda_b(t) - \lambda_g(t)}{\theta}\right)^2 V_{yy} + \frac{1}{2}\sigma^2 V_{zz} + \frac{1}{2}\left(\frac{\lambda_b(t) - \lambda_g(t)}{\theta}\right)^2\left(\frac{e^y - 1}{e^y + 1}\right)V_y + \left(\mu - \frac{\sigma^2}{2}\right)V_z - rV.$$

Unlike the partial differential equation in the original coordinates, in these coordinates the coefficients of the second-order derivatives, V_{yy} and V_{zz}, are functions of t only.

We solve the investment-timing problem on a discrete grid. This involves converting the variational inequalities in Lemma 5 into a linear complementarity problem, which we solve using a variant of the operator-splitting approach developed by Ikonen and Toivanen (2004, 2009). The transformed coordinate system clusters grid points near the $p = 0$, $p = 1$, and $x = 0$ boundaries, but cannot actually place grid points on these boundaries. Therefore, we set up the grid so the endpoints in the y direction are close to the $p = 0$ and $p = 1$ boundaries, and the lower z grid point is close to the $x = 0$ boundary.

We solve the model out to some arbitrary large date T, which we set equal to 500 years, and assume that nothing changes past this point. In particular, we assume the investment payoff equals

$$\big(pv_b(T;q) + (1-p)v_g(T;q)\big)\, x - c$$

on or after date T. The effects of this assumption are negligible, given that rainfall intensity has almost completed the convergence process by this date. Moreover, the effects of any approximation errors will be trivial by the time everything is discounted back to date 0. We solve this problem using the projected successive over-relaxation approach, set $V(T, y, z)$ equal to the present value of the perpetual option, and then use this as the terminal condition before iterating back to date 0 in the usual way.[27]

We calculate the optimal investment threshold as follows. For each (t_i, y_i), calculate the largest grid point z_i inside the waiting region. Let g_i equal the stopping payoff at (t_i, y_i, z_i) minus the waiting payoff at the same grid point. As the grid point is in the waiting region, this number will be negative. Let g_{i+1} denote the corresponding quantity at the adjacent grid point (t_i, y_i, z_{i+1}). As this is inside the stopping region, g_{i+1} will be positive. Note that in order to calculate g_{i+1}, we need to calculate the waiting payoff. This will be readily available from the intermediate calculations used in the finite difference method. Assuming the difference between the stopping and waiting payoffs is linear in z between grid points, the payoffs are equal at

$$z^* = z_i + \left(\frac{-g_i}{g_{i+1} - g_i}\right)\Delta z,$$

where Δz is the distance between grid points in the z direction. The corresponding threshold for the economic factor is

$$\hat{x}^*(t_i, y_i) = \exp\left(z_i + \left(\frac{-g_i}{g_{i+1} - g_i}\right)\Delta z\right).$$

References

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics*, 98(1):85–106.

Boettle, M., Kropp, J. P., Reiber, L., Roithmeier, O., Rybski, D., and Walther, C. (2011). About the influence of elevation model quality and small-scale damage functions on flood damage estimation. *Natural Hazards and Earth System Sciences*, 11(12):3327–3334.

Boyle, G. and Guthrie, G. (2006). Payback without apology. *Accounting and Finance*, 1(46):1–10.

Copeland, T. and Antikarov, V. (2003). *Real Options: A Practitioner's Guide*. Texere, New York.

[27] We assume the decision-maker ignores new climate information along the upper and lower y boundaries, which is what would happen as $y \to \pm\infty$. Thus, along these boundaries, V is a function of (t, z) only. We impose numerical boundary conditions along the upper and lower z-boundaries.

Copeland, T. E. and Antikarov, V. (2005). Real options: Meeting the Georgetown challenge. *Journal of Applied Corporate Finance*, 17(2):32–51.

Davis, D. R., Kisiel, C. C., and Duckstein, L. (1972). Bayesian decision theory applied to design in hydrology. *Water Resources Research*, 8(1):33–41.

Dawson, D. A., Hunt, A., Shaw, J., and Gehrels, W. R. (2018). The economic value of climate information in adaptation decisions: Learning in the sea-level rise and coastal infrastructure context. *Ecological Economics*, 150:1–10.

De Bruin, K. and Ansink, E. (2011). Investment in flood protection measures under climate change uncertainty. *Climate Change Economics*, 2(4):321–339.

Dixit, A. K. (1993). *The Art of Smooth Pasting*. Harwood Academic Publishers, Amsterdam.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.

Dunedin City Council (2019). Talking futures for South Dunedin: Water, challenges and opportunities. Hui Presentation, 12 August.

Dunedin City Council (2022). The future. `https://www.dunedin.govt.nz/council/council-projects/south-dunedin-future/south-dunedin-future`, last accessed on 2022-02-15.

Guthrie, G. (2019). Real options analysis of climate-change adaptation: Investment flexibility and extreme weather events. *Climatic Change*, 156(1–2):231–253.

Harrison, J. M. (2013). *Brownian Models of Performance and Control*. Cambridge University Press, New York, NY.

Hughes, J., Cowper-Heays, K., Olesson, E., Bell, R., and Stroombergen, A. (2019). Stormwater, wastewater and climate change: Impacts on our economy, environment, culture and society. *Tonkin+Taylor*, Job number 1004190.

Ikonen, S. and Toivanen, J. (2004). Operator splitting methods for American option pricing. *Applied Mathematics Letters*, 17(7):809–814.

Ikonen, S. and Toivanen, J. (2009). Operator splitting methods for pricing American options under stochastic volatility. *Numerische Mathematik*, 113:299–324.

Kopp, R. E., Horton, R. M., Little, C. M., Mitrovica, J. X., Oppenheimer, M., Rasmussen, D. J., Strauss, B. H., and Tebaldi, C. (2014). Probabilistic 21st and 22nd century sea-level projections at a global network of tide-gauge sites. *Earth's Future*, 2(8):383–406.

Lee, B. S., Haran, M., and Keller, K. (2017). Multidecadal scale detection time for potentially increasing Atlantic storm surges in a warming climate. *Geophysical Research Letters*, 44(10):617–623.

Lee, S. and Zhao, J. (2021). Adaptation to climate change: Extreme events versus gradual changes. *Journal of Economic Dynamics and Control*, 133:104262.

Malchow-Møller, N. and Thorsen, J. (2005). Repeated real options: Optimal investment behaviour and a good rule of thumb. *Journal of Economic Dynamics and Control*, 29:1025–1041.

McDonald, R. and Siegel, D. (1986). The value of waiting to invest. *Quarterly Journal of Economics*, 101(4):707–728.

McDonald, R. L. (2000). Real options and rules of thumb in capital budgeting. In Brennan, M. J. and Trigeorgis, L., editors, *Project Flexibility, Agency, and Competition*, pages 13–33, Oxford. Oxford University Press.

Merz, B., Kreibich, H., Schwarze, R., and Thieken, A. H. (2010). Assessment of economic flood damage. *Natural Hazards and Earth System Sciences*, 10(8):1697–1724.

Penning-Rowsell, E., Priest, S., Parker, D., Morris, J., Tunstall, S., Viavattene, C., Chatterton, J., and Owen, D. (2013). *Flood and Coastal Erosion Risk Management: A Manual for Economic Appraisal.* Routledge, Abingdon, Oxon.

Pindyck, R. S. (2021). What we know and don't know about climate change, and implications for policy. *Environmental and Energy Policy and the Economy*, 2:4–43.

Prahl, B. F., Boettle, M., Costa, L., Kropp, J. P., and Rybski, D. (2018). Damage and protection cost curves for coastal floods within the 600 largest European cities. *Scientific Data*, 5:180034.

Pregnolato, M., Ford, A., Wilkinson, S. M., and Dawson, R. J. (2017). The impact of flooding on road transport: A depth-disruption function. *Transportation Research Part D: Transport and Environment*, 55:67–81.

Stokey, N. L. (2008). *The Economics of Inaction: Stochastic Control Models with Fixed Costs.* Princeton University Press, Princeton.

Treasury (2022). Discount rates. `https://www.treasury.govt.nz/information-and-services/state-sector-leadership/guidance/financial-reporting-policies-and-guidance/discount-rates`, last accessed on 2022-03-04.

Truong, C., Trück, S., and Mathew, S. (2018). Managing risks from climate impacted hazards—The value of investment flexibility under uncertainty. *European Journal of Operational Research*, 269(1):132–145.

Urban, N. M., Holden, P. B., Edwards, N. R., Sriver, R. L., and Keller, K. (2014). Historical and future learning about climate sensitivity. *Geophysical Research Letters*, 41:2543–2552.

van der Pol, T. D., Gabbert, S., Weikard, H., van Ierland, E. C., and Hendrix, E. M. T. (2017a). A minimax regret analysis of flood risk management strategies under climate change uncertainty and emerging information. *Environmental and Resource Economics*, 68(4):1087–1109.

van der Pol, T. D., van Ierland, E. C., and Gabbert, S. (2017b). Economic analysis of adaptive strategies for flood risk management under climate change. *Mitigation and Adaptation Strategies for Global Change*, 22(2):267–285.

van der Pol, T. D., van Ierland, E. C., and Weikard, H. (2014). Optimal dike investments under uncertainty and learning about increasing water levels. *Journal of Flood Risk Management*, 7(4):308–318.

Wambach, A. (2000). Payback criterion, hurdle rates and the gain of waiting. *International Review of Financial Analysis*, 9(3):247–258.

Watkiss, P., Hunt, A., Blyth, W., and Dyszynski, J. (2015). The use of new economic decision support tools for adaptation assessment: A review of methods and applications, towards guidance on applicability. *Climatic Change*, 132(3):401–416.

Woolley, J.-M., Eager, C., Jozaei, J., Paul, V., Paulik, R., Pearce, P., Sood, A., Stuart, S., Vincent, A., Wadhwa, S., and Zammit, C. (2020). Climate change projections and impacts for Tairāwhiti and Hawke's Bay. *NIWA Client Report*, 2020298AK. Report prepared for Envirolink, Gisborne District Council and Hawke's Bay Regional Council.

Wreford, A., Dittrich, R., and van der Pol, T. D. (2020). The added value of real options analysis for climate change adaptation. *WIREs Climate Change*, 11(3):e642.