

Prins, Taco; van der Ploeg, Frederick; van den Bremer, Ton S.

Working Paper

Sea level rise and optimal flood protection under uncertainty

Tinbergen Institute Discussion Paper, No. TI 2025-030/VI

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Prins, Taco; van der Ploeg, Frederick; van den Bremer, Ton S. (2025) : Sea level rise and optimal flood protection under uncertainty, Tinbergen Institute Discussion Paper, No. TI 2025-030/VI, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/316217>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2025-030/VI
Tinbergen Institute Discussion Paper

Sea Level Rise and Optimal Flood Protection under Uncertainty

Taco Prins¹

Frederick van der Ploeg²

Ton S. van den Bremer³

¹ University of Amsterdam

² University of Amsterdam, University of Oxford, Tinbergen Institute

³ University of Amsterdam, Delft University of Technology, Tinbergen Institute

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <https://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

Sea Level Rise and Optimal Flood Protection under Uncertainty

Taco Prins*, Frederick van der Ploeg[†], and Ton S. van den Bremer[‡]

Revised April 2025

Abstract

We analyse optimal investment in one of the most important forms of climate adaptation: flood protection. Investments to build and heighten dykes and surge barriers involve considerable adjustment costs, so that their construction locks in the level of flood protection for some time. Investment decisions must take into account both economic and sea level rise uncertainty over a horizon of several decades, where the latter is to a large extent driven by global warming. We put forward a tractable macro-finance DSGE model that includes flood risk. We obtain solutions for optimal flood protection as a function of these uncertainties, costs, and preferences regarding impatience, risk aversion and intertemporal substitution. Sea level rise uncertainty always leads to more flood protection. Economic uncertainty leads to less (more) protection if the elasticity of substitution is greater (less) than one. We illustrate our results with a calibrated case study for the Netherlands.

Keywords: Sea level rise, flood risk, macroeconomic risk, climate adaptation, discounting, risk aversion, intertemporal substitution

JEL subject codes: F64, Q51, Q54

*University of Amsterdam, Faculty of Economics and Business, 1018 WB, Amsterdam, The Netherlands. Email: t.a.prins@uva.nl.

[†]University of Oxford, Department of Economics, OXCARRE, Manor Road Building, Oxford, OX1 3UQ, UK, and University of Amsterdam, Faculty of Economics and Business, 1018 WB, Amsterdam, The Netherlands. Email: rick.vanderploeg@economics.ox.ac.uk. Also affiliated with CEPR and CESifo.

[‡]University of Amsterdam, Faculty of Economics and Business, 1018 WB, Amsterdam, The Netherlands; and Delft University of Technology, Faculty of Civil Engineering and Geosciences, 2628 CN, Delft, The Netherlands. Email: t.s.vandenbremer@uva.nl.

1 Introduction

Climate change is widely expected to increase the global costs of weather-related disasters (e.g. IPCC 2022). There is evidence that, already in the first two decades of this century, more than half of the recorded economic damage due to such disasters was attributable to climate change (Newman and Noy 2023). Because the success of efforts to mitigate climate change is uncertain and will not be felt for decades, there is a need to adapt to a world in which natural disasters occur more frequently. Barrage and Furst (2019) and Bakkensen and Barrage (2022) show that beliefs about flood risks affect housing construction and housing prices. Fried (2022) and Hsiao (2024) discuss optimal flood control by investing in seawalls and stilts in, respectively, a heterogeneous-agent macro model and a dynamic spatial model of urban development and flooding.

The prospect of increasing flood risk due to climate change calls for a review of current flood protection policies in many countries. There is a great uncertainty about the trajectory of global temperatures and even greater uncertainty about sea level rise and the frequency of extreme weather events, because the underlying geophysical processes are poorly understood (see, e.g. Kopp et al. 2017, Haasnoot et al. 2020). Meanwhile, investment in flood protection must be forward-looking. A sensible flood protection strategy takes into account both the uncertainty of economic and sea level rise, as well as the inevitable trend increase in sea level as global warming progresses.

We study optimal investment in flood protection within the context of a macro-finance growth model with both sea level rise uncertainty and economic uncertainty, where the risk of flooding can be reduced by investment in dykes. Our framework takes into account regular macroeconomic uncertainty and the risk of rare macroeconomic disasters. These do not depend on the risk of flooding or (uncertainty in) sea level rise but do affect the social discount rate, which in turn affects optimal flood protection investment. We allow the aversion to risk to differ from the aversion to intertemporal fluctuations by using Epstein–Zin preferences, which imply that prudence decreases in the ease of intertemporal substitution. Capital in final goods production and flood protection is hard to adjust, which we capture with intertemporal adjustment costs in investments in the economy-wide capital stock and in the stock of flood protection (dykes).

We show that the effect of macroeconomic uncertainty on optimal flood protection depends on the size of the elasticity of intertemporal substitution, while sea level rise uncertainty always increases optimal investment in flood protection. To illustrate our proposed framework, we calibrate it to the Netherlands, where the challenge of increasing flood risk

is felt particularly keenly. Much of the land mass in the Netherlands, including almost the whole economic core in the western part of the country (the ‘Randstad’), lies below sea level. Unsurprisingly, throughout Dutch history, land reclamation and flood protection have been major components of public investment.

At the heart of the Dutch flood protection strategy has been a system of safety norms for individual polders or ‘dyke ring areas’, based on cost-benefit analysis following the contribution of Van Dantzig (1956), which appeared just after the large North Sea flood of 1953. This seminal study compares the net present value of investment costs and expected damages due to flooding, where economic growth and the evolution of flood risk are deterministic processes, in contrast to our framework. Our contribution is to show how flood protection standards should be adjusted to take economic and climate uncertainty into account, paying due attention to aversion to risk and aversion to intertemporal fluctuations.

Our analytical framework builds on Pindyck and Wang (2013), who study the consequences of catastrophes for optimal consumption, wealth, and welfare in an AK-growth model. They find that their representative agent has a very high willingness to pay (about half the capital stock) to eliminate disaster risk and fluctuations of economic growth. Their findings contrast starkly with a simple net-present-value calculation of the expected cost of disasters. Their basic framework of disasters has also been used in the context of flood protection by two earlier studies.

Douenne (2020) differs from Pindyck and Wang (2013) by assuming that the arrival rate of disasters decreases in current adaptation spending. He calibrates his model to the different regions of the United States, and investigates the effect of disasters on economic growth — a priori ambiguous — and welfare. Hong et al. (2023) place their work much more explicitly in the context of increasing disaster risk due to uncertain climate change. They allow firms to engage in private adaptation spending to curb the proportion of their capital that is exposed to disaster risk. They let aggregate adaptation spending reduce the tail risk of the damage distribution from disasters for all firms, which requires public intervention. They also let Bayesian learning about the disaster arrival rate proxy for uncertainty about the severity of climate change. They thus derive partial analytical solutions, and give insights into optimal spending to prevent flood damage from hurricanes¹.

Relative to these earlier contributions, one of our innovations is that we consider the level of flood protection as a stock that cannot be adjusted instantaneously. Canonical

¹Hong et al. (2023) also provide theoretical and empirical evidence that the landfall of tropical cyclones in a country tends to lower Tobin’s average q and raise the equity premium. In their model, a learning effect causes the perceived risk of further disasters to rise, which hurts investment.

contributions on investment under uncertainty, such as Abel et al. (1996), note that uncertainty creates an option value of ‘waiting and seeing’ if investment is irreversible. This causes such investment to be deferred when learning takes place in the future, a point made in the context of flood protection investment by Van der Pol et al. (2014). However, if investment serves to mitigate an unknown disaster risk, the effect of uncertainty on the timing of investment is less obvious. In our model, we provide closed-form solutions for the level of flood protection that is optimal in the long run, and solve for the full path of flood protection investments numerically.

Section 2 presents our stochastic dynamic programming problem for optimal investment in flood protection with macroeconomic and flood risks. Section 3 presents the solution. Section 4 calibrates our framework to the Netherlands. Section 5 offers a quantitative assessment of optimal investment in flood protection, and presents a sensitivity analysis with respect to preferences, macroeconomic and flood risk parameters, and dyke heightening costs. Section 6 concludes and offers suggestions for further research.

2 DSGE Framework for Optimal Flood Protection

Here, we present our dynamic stochastic general equilibrium (DSGE) model of disaster risk and optimal flood protection. For the macro-finance building blocks, we draw on Pindyck and Wang (2013) and Hambel et al. (2024), who consider the consequences of catastrophes for optimal consumption, wealth, and welfare in an AK-growth model. They show that disaster risk itself, even in the absence of a disaster strike, affects the consumption-saving decision. Tsai and Wachter (2018) clarify the asset-pricing implications of this type of model. An innovation relative to Barro (2009) and Barro (2015) is the allowance for adjustment costs in capital investment, which make disasters more costly.

2.1 Preferences

The utility of the representative agent has three separate preference parameters: the (utility) rate of time preference ρ , the coefficient of relative risk aversion (CRRA) γ , and the elasticity of intertemporal substitution (EIS) ψ . The future is discounted more if ρ is higher. If the CRRA is higher, society is willing to pay more to avoid the risk of low consumption. The elasticity of intertemporal substitution (EIS) is ψ . Note that $1/\text{EIS}$ determines how averse society is to fluctuations in utility over time (more averse if the EIS is lower). The CRRA and the EIS are conceptually distinct, and Epstein–Zin preferences based on Kreps and Porteus (1978) and adapted for continuous time by Duffie and Epstein (1992) allow them

to be analysed separately. The value function V_t is thus implicitly defined by

$$V_t = E_t \left[\int_t^\infty f(C_s, V_s) ds \right], \quad (1)$$

where the aggregator function is given by

$$f(C_s, V_s) = \frac{\rho}{1 - \psi^{-1}} \frac{C_s^{1-\psi^{-1}} - ((1-\gamma)V_s)^\omega}{((1-\gamma)V_s)^{\omega-1}}, \quad (2)$$

and $\omega \equiv \frac{1-\psi^{-1}}{1-\gamma}$. In the macro-finance literature, the CRRA and the EIS have been helpful to better explain the equity premium puzzle (e.g. Bansal and Yaron, 2004)². We show later how important it is to distinguish CRRA and the inverse of EIS to gain a better understanding of the effects of various kinds of uncertainty on optimal flood protection.

2.2 Production and capital dynamics

We assume a ‘AK’ production technology for output

$$Y_t = AK_t, \quad (3)$$

where K_t is the stock of capital and is a broad measure of capital, including intangible capital, and A is the exogenous and constant productivity of the capital stock (or the output-capital ratio). This ‘AK’ structure leads to endogenous growth. Consumption C_t is output net of investment in production capital $I_{K,t}$ and in flood protection $I_{H,t}$, so that

$$C_t = Y_t - I_{K,t} - I_{H,t}. \quad (4)$$

Productive capital evolves according to

$$dK_t = \left(I_{K,t} - \frac{1}{2}\theta_K \frac{I_{K,t}^2}{K_t} - \delta_K K_t \right) dt + \sigma_K K_t dB_{K,t} - (1 - Z_e)K_t dJ_{e,t} - (1 - Z_f(H_t, h_t))K_t dJ_{f,t}, \quad (5)$$

where $\theta_K > 0$ denotes the adjustment cost parameter and $\delta_K \geq 0$ the depreciation rate. Investment is thus subject to quadratic adjustment costs. Capital growth is subject to regular

²If the CRRA is constrained to equal the inverse of the EIS, as in a standard additively separable aggregator, a high aversion to consumption risk implies a high aversion to fluctuations in aggregate consumption over time. This is restrictive, since a high equity premium then implies a high risk-free interest rate.

fluctuations captured by geometric Brownian motion $B_{K,t}$ with σ_K denoting the volatility and to discrete downwards jumps from economic disasters $J_{e,t}$ and flooding disasters $J_{f,t}$.³ The jump processes are Poisson processes. The first jump process $J_{e,t}$ has a fixed arrival rate $\lambda_e \geq 0$, and represents economic disasters that cannot be influenced by policy makers (cf. Barro, 2009). When an economic disaster strikes, a proportion $1 - Z_e$ of the capital stock is permanently wiped out, where the survival fraction Z_e is a random variable whose probability density function follows a power law, $f_{Z_e}(Z_e) = \chi_e Z_e^{\chi_e - 1}$, where $\chi_e > 0$ denotes the distribution parameter (cf. Pindyck and Wang, 2013). The expected loss from an economic disaster equals $\mathbb{E}[1 - Z_e] = 1/(\chi_e + 1)$. The second jump process $J_{f,t}$ captures flooding disasters and has an endogenous arrival rate depending on the stock of flood protection H_t and sea level rise h_t , i.e. $\lambda_f(H_t, h_t)$ (see Section 2.4). The (random) survival fraction Z_f is also described by a power-law probability density function, $f_{Z_f}(Z_f) = \chi_f Z_f^{\chi_f - 1}$, where $\chi_f(H_t, h_t) > 0$ now depends on the stock of flood protection and sea level rise (see Section 2.5).

2.3 Investment in flood protection capital

The variable H_t captures accumulated investment in flood protection, net of depreciation. Since the flood arrival rate is modelled as a function of the height of dykes relative to the sea level, it can be thought of as flood protection capital or dyke height. Henceforth, we will refer to it as ‘flood protection’ or ‘dyke height’, always denominated in metres. Its evolution is deterministic and given by

$$\frac{dH_t}{dt} = \frac{1}{\nu_t} \left(I_{H,t} - \frac{1}{2} \theta_{H,t} \frac{I_{H,t}^2}{H_t} \right) - \delta_H H_t, \quad (6)$$

where $\theta_{H,t} \equiv \theta_H K_0 / K_t$ with θ_H (units m·yr/Euros) the initial adjustment cost parameter for flood protection investment, $\nu_t \equiv \nu K_t / K_0$ with ν a parameter describing the initial cost of flood protection investment relative to productive capital investment (units Euros·m⁻¹), and $\delta_H \geq 0$ the depreciation rate of dykes. Investment in flood protection is also subject to quadratic adjustment costs. To keep the ‘AK’ structure, we let the cost of flood protection relative to that of economy-wide investment increase in the economy-wide capital stock and the adjustment cost parameter decrease in the economy-wide capital stock, both due to economy-wide technical progress. Adjustment costs depend on the flood protection

³The effects of regular fluctuations and the risk of rare macroeconomic disasters on asset prices are analysed extensively in the macro-finance literature (e.g. Pindyck and Wang, 2013; Tsai and Wachter, 2018). We focus here on their effect on optimal investment in climate adaptation.

investment rate relative to the existing flood protection stock, while total investment costs increase one-to-one with the productive capital stock (and thus output). A country with a more developed flood protection stock thus has a greater capacity for further investment.⁴

2.4 Flood risk

The single flood risk disaster process $J_{f,t}$ in our framework represents flooding in different dyke ring areas, which by definition are separate events. The arrival rate of each of these events is driven by local (i.e. dyke-ring-specific) sea level rise and local flood protection. In addition, flood events in different areas can be positively correlated when several dykes fail in a storm, or negatively correlated when the weakest dyke fails along a river (reducing the probability of flooding downstream). Accounting for such interactions complicates matters considerably.⁵ We defer an extended treatment of these issues to Section 4, where we calibrate the model to the Netherlands taking account of the various dyke ring areas. We can include many separate flooding disaster processes in our framework provided that the local flood protection level and sea level rise only depend on aggregate variables, and spatial heterogeneity in economic growth is ruled out. To keep notation simple, we do not include these here.

We specify the flood arrival rate as a function of dyke height and sea level rise,

$$\lambda_f(H_t, h_t) = P_0 e^{-\alpha(H_t - H_0 - h_t)}, \quad (7)$$

where P_0 is the initial flood arrival rate, $H_t - H_0$ is the amount of dyke heightening that has taken place between times 0 and t , and h_t is cumulative sea level rise over the same period. Dyke heightening relative to cumulative sea level rise decreases flood risk through the parameter $\alpha > 0$. Conversely, flood risk increases if the mean sea level increases relative to the height of the dykes. Cumulative sea level rise requires dykes to be raised.

The functional form in (7) stems from Van Dantzig (1956), who observed that the probability of a tide exceeding a certain height is roughly a decreasing exponential function

⁴We abstract from the possibility that flood protection becomes relatively cheaper (more expensive) over time as output rises, which would motivate deferring (bringing forward) investment. Empirically, Jonkman et al. (2013) report that the unit cost of flood defenses increase with GDP in a set of (cross-sectional) case studies including the Netherlands, the United States (New Orleans), and Vietnam. This is confirmed by Nicholls et al. (2019), who use a large sample. The cost of labour and the price of land increase with GDP, but the cost of raw materials does not necessarily. Technological innovation can depress absolute costs.

⁵See e.g. Zwaneveld and Verweij (2018), who solve such a problem to find the appropriate flood protection strategy in the IJssel lake area. In the Netherlands, to estimate the arrival rate of flooding in each dyke ring area, relative to the legal norms, events are treated as independent except for a major positive interaction (‘cascade’) into South Holland province. See Ministry of Infrastructure and Water Management (2014).

of that height relative to the mean sea level. Hence, if one thinks of flooding as overtopping, then sea level rise causes flood risk to increase exponentially in the absence of investment in flood protection. The functional form for flood risk has been widely used in applied work in the Netherlands, where before 2017 flood protection standards required structures to resist overtopping from storm surges and river discharges with a certain return period. It was a key input for the ‘Waterwet’ legislation adopted by the Dutch parliament in 2017, which established the flood protection standards that must be achieved by 2050.⁶

2.5 Flood damages and flood protection

Expected damages from a flood increase with dyke height H_t and sea level rise h_t , mainly through the expected inundation depth. Dyke height alters the damage distribution (conditional on flooding taking place). In our model, a higher stock of flood protection and a higher sea level rise should thus lower the flood damage distribution parameter χ_f (i.e. $\partial\chi_f/\partial H_t < 0, \partial\chi_f/\partial h_t < 0$) and increase expected damages (since $\mathbb{E}[1 - Z_f] = 1/(\chi_f + 1)$). As described in Section 4, our empirical data contains parameters for the effect on expected damages of a one metre increase in the dyke height and sea level. These are used to calibrate a flood damage distribution parameter of the exponential form,⁷

$$1 + \chi_f(H_t, h_t) = (1 + \chi_f(H_0, h_0))e^{-\zeta_H(H_t - H_0) - \zeta_h h_t}, \quad (8)$$

where the parameters $\zeta_H > 0$ and $\zeta_h > 0$ describe the percentage increase in expected monetary damages from a one metre increase in the dyke height and the sea level, respectively. To have meaningful solutions, we require $\zeta_H < \alpha$ (see our calibration in Section 4).

We let climate change and sea level rise affect both the frequency of flooding disasters, via (7), and the severity of disasters, via (8). We thus follow the best available evidence on the effects of climate change.⁸ In the context of flooding, policymakers may respond to increased disaster frequency by heightening flood protection structures. The combined

⁶The formula can be found as equation (2.2) on page 20 of an overview provided by Kind (2011), in Dutch, of the official ‘WV21’ research project. This explains that local correction factors take into account sources of dyke failure other than overtopping risk, see *ibid* p.26. Eijgenraam et al. (2017) provide a summary in English of the computational approach underpinning WV21.

⁷This function is obtained by differentiating expected damages from flooding $\mathbb{E}[1 - Z_f] = 1/(1 + \chi_f(H, h))$ with respect to the flood protection stock H_t , setting $\partial\mathbb{E}[1 - Z_f]/\partial H = -\zeta_H\mathbb{E}[1 - Z_f]$, and solving the resulting differential equation to obtain (8), and similarly for sea level rise h_t .

⁸According to the IPCC’s Sixth Assessment Report (Working Group I, 2021), there is for example ‘high confidence that sea level rise will lead to a higher possibility of extreme coastal water levels’ (*ibid.* p. 1592), and ‘high confidence that the occurrence and magnitude of compound flooding in coastal regions will increase in the future due to both sea level rise and increases in heavy precipitation’ (*ibid* p. 1600).

effect of climate change and adaptation on disaster frequency is ambiguous, but its effect on expected capital destruction is positive. Our results in Section 5 will quantify this insight.

A further implication of our process for capital accumulation (5) is that damages sustained from a flooding disaster are permanent. When a flooding disaster strikes, a fraction of the capital stock is wiped out, and subsequent expected economic growth is constant albeit from a lower (post-disaster) capital stock. The presence of a ‘rebound’ or ‘mean-reversal’ effect would introduce a term structure for discount rates especially for long-lasting risky assets such as real estate (see Giglio et al. 2021). However, there is currently no evidence that macroeconomic disasters, or flooding disasters in particular, are followed by higher economic growth⁹ and we thus abstract from the resulting term structure.

2.6 Sea level rise

Global mean sea level rise has been accelerating, reaching a rate of 3.2 to 4.2 mm/year between 2006 and 2018 (IPCC 2022). The implications for future sea level rise are highly uncertain. Given a path of the global surface temperature, the physical processes of ice sheet melting and collapse, as well as thermal expansion and reduced land water storage, will each contribute to total sea level rise. However, they will do so over a long timescale. Because (especially) the behaviour of the Antarctic Ice Sheet is poorly understood, current observations are consistent with vastly different potential outcomes from 2050 onwards. ‘Finding the planet on a “moderate” sea-level rise pathway over the first half of the 21st century thus cannot exclude “extreme” outcomes subsequently’ (Kopp et al., 2017, p.1228). To capture the generally right-skewed prediction of future sea level rise projections, we specify a stochastic process, which approximates sea level rise as a power-law (skewed) transformation of an Ornstein–Uhlenbeck process,

$$h_t = X_t^{1+\theta_X} - X_0^{1+\theta_X}, \quad \text{with} \quad dX_t = \eta_X(\bar{X} - X_t)dt + \sigma_X dB_{X,t}, \quad (9)$$

where $B_{X,t}$ is the Brownian motion describing this stochastic process, X_t is the Ornstein–Uhlenbeck process with steady-state mean \bar{X} , volatility σ_X and mean reversion coefficient η_X , and $\theta_X > 0$ allows for (positive) skewness in the resulting distribution.¹⁰

⁹See Cerra and Saxena (2008) for a broad set of macroeconomic disasters and Hsiang and Jina (2014) for a large multi-country study where losses from hurricanes turn out to deepen over time (i.e. no recovery). Batten (2018), in a review article on climate change as a macroeconomic risk, concludes that ‘the ‘no recovery’ hypothesis seems to be supported by the largest number of empirical studies’ (ibid. p. 25).

¹⁰The constant X_0 ensures that, at $t = 0$, the process is well-defined and the expected rate of sea level rise is strictly positive. The process becomes ill-defined when $X_t \leq 0$. The probability that this occurs is very small, and our calibration demonstrates this. Formally, equation (9) has the form $h_t = X_t^{1+\theta_X} - X_0^{1+\theta_X}$ for

We assume that the Brownian motion $B_{X,t}$ is uncorrelated with the Brownian motion in the evolution for capital $B_{K,t}$ and, in doing so, ignore the effect of climate change on the economy other than through flooding. The skewness parameter $\theta_X > 0$ delivers both an initially increasing pace of sea level rise and a right skew in its distribution at any point in time. The steady-state mean \bar{X} in the process for X_t ensures that the pace of sea level rise eventually slows.¹¹ Our stochastic process for sea level rise in (9) can be thought of as a ‘model emulator’ of highly complex climate and sea level rise models (see Section 4).

3 Characterisation of Optimal Flood Protection

The social planner maximises the value from current and future consumption subject to the production function (3), the resource constraint (4), the evolution of the stock of productive capital (5)-(7) and that of flood protection (8)-(10). Policymakers can instantaneously adjust the consumption rate and the investment rates in productive capital and flood protection, which are the control variables. Capital, flood protection, and the level of cumulative sea level rise are the state variables, so that the value function is $V(K, H, h)$. This corresponds to a dynamic programming problem with Hamilton–Jacobi–Bellman (HJB) equation

$$\begin{aligned}
0 = \max_{I_K, I_H} \{ & f(C, V) + \left(I_{K,t} - \frac{1}{2} \theta_K \frac{I_{K,t}^2}{K_t} - \delta_K K \right) V_K + \frac{1}{2} \sigma_K^2 K^2 V_{KK} + \\
& \left(\frac{1}{\nu_t} \left(I_{H,t} - \frac{1}{2} \theta_H \frac{I_{H,t}^2}{H_t} \right) - \delta_H H \right) V_H + \eta_h(h) V_h + \frac{1}{2} \sigma_h(h)^2 V_{hh} + \\
& \lambda_e \mathbb{E}[V(Z_e K, H, h) - V(K, H, h)] + \lambda_f(H, h) \mathbb{E}[V(Z_f K, H, h) - V(K, H, h)] \},
\end{aligned} \tag{10}$$

where $\eta_h = (1/dt)\mathbb{E}_t[dh]$ and $\sigma_h(h)^2 = (1/dt)\mathbb{E}_t[(dh)^2]$ can be obtained from (9).

The first line in the HJB equation shows the instantaneous utility of consumption plus the expected change in the value function due to changes in the capital stock plus the Itô term due to the continuous variations in K from the Brownian motion. The second line shows the increase in the value function from increases in flood protection and decrease from increases in sea level rise plus an Itô term due to the stochasticity of sea level rise. The third line captures the expected decrease in the value function from economic and flooding disasters, which is more negative if the arrival rate or severity of these disasters is greater.

$X_t > 0$ and $h_t = 0$ for $X_t \leq 0$, which is also implemented in the numerical method. The effect of the atom at $h_t = 0$ in the distribution for sea level rise is negligibly small.

¹¹Such an inflection point is projected to occur within a few decades (SSP2-4.5) to about a century (SSP5-8.5). Regardless of the pathway, sea level continues to rise for several centuries.

Using (4) and taking derivatives of the HJB equation gives the optimality conditions¹²

$$\left(1 - \theta_K \frac{I_{K,t}}{K_t}\right) V_K = \frac{1}{\nu_t} \left(1 - \theta_{H,t} \frac{I_{H,t}}{H_t}\right) V_H = f_C(C, V). \quad (11)$$

The marginal benefit of each type of investment, multiplied by the rate at which the consumption good can be turned into capital, decreases in the investment rate, and must be set to the marginal cost of investing (i.e. the marginal utility of consumption).

For a given value function $V(K, H, h)$, (11) and (4) allow one to express the controls (I_K, I_H, C) in terms of (K, H, h) . Since the cost of investment in flood protection and the damage from flooding rise proportionally with the capital stock, the problem is homogeneous in K , and the value function has the form $V(K, H, h) = \frac{1}{1-\gamma} (g(H, h)K)^{1-\gamma}$, so that we only need to solve the (reduced-form) HJB equation for the function $g(H, h)$. In doing so, the number of states is reduced from three to two, and computational costs become less. The homogeneity of the value function results in $i_K \equiv \frac{I_K}{K}$ and $i_H \equiv \frac{I_H}{K}$ being functions of H and h only. We define $\phi_K(i_K) \equiv i_K - \frac{1}{2}\theta_K i_K^2$ and $\phi_H(i_H, H) \equiv \frac{K_0}{\nu} \left(i_H - \frac{1}{2}\theta_H K_0 \frac{i_H^2}{H}\right)$.

3.1 Optimal flood protection without sea level rise

We first obtain some analytical results without sea level rise ($\eta_h = \sigma_h = 0$) as then the optimal flood protection policy gives rise to convergence to a balanced growth path with constant flood risk. The resulting long-run flood arrival rate gives insight into the numerical solution of our fully calibrated model (see Section 5).

3.1.1 Long-run flood arrival rate

Along the balanced growth path, the stock of flood protection H is constant and from (6) we have that maintenance costs equal depreciation of the stock of flood protection. The resulting flood protection investment, $i_H(H)$, increases in H . Increasing maintenance costs of flood protection, together with a diminishing marginal effect on flood risk, leads to a steady-state value of the flood protection level H^* . Exploiting the homogeneity of the problem in K , we can characterise H^* analytically (see Appendix A). We introduce the shorthands $\Gamma(H, h) \equiv \mathbb{E}[1 - Z_f^{1-\gamma}]/(1 - \gamma)$ for the expected risk-adjusted damages from a flood event and $\Theta_H(i_H, H) \equiv \frac{\theta_H}{2\nu} \frac{(i_H K_0)^2}{H}$ for the adjustment costs in flood protection investment.

¹²Because utility is concave in consumption and the marginal product of investment is decreasing in the investment rates i_K and i_H , the second-order conditions for maximisation are met.

Result 1: *The optimal long-run level of flood protection H^* satisfies*

$$-\left(\frac{\partial \lambda_f(H^*, h)}{\partial H} \Gamma(H^*, h) + \lambda_f(H^*, h) \frac{\partial \Gamma(H^*, h)}{\partial H}\right) \frac{\phi'_H(i_H, H^*)}{\phi'_K(i_K)} = r^* + \delta_H + \frac{\partial \Theta_H(i_H, H^*)}{\partial H}. \quad (12)$$

The risk- and growth-adjusted discount rate is

$$r^* = \rho + (\psi^{-1} - 1)(g_e - \frac{1}{2}\gamma\sigma_{\text{total}}^2), \quad (13)$$

where expected growth net of the expected impact of economic and flooding disasters is

$$g_e \equiv \phi_K(i_K) - \delta_K - \lambda_e \mathbb{E}[1 - Z_e] - \lambda_f(H^*, h) \mathbb{E}[1 - Z_f]$$

and macroeconomic volatility including the effect of economic and flooding disasters is¹³

$$\sigma_{\text{total}}^2 \equiv \sigma_K^2 + 2\lambda_e \mathbb{E}[1 - Z_e] \mathbb{E}[1 - Z_e^{1-\gamma}] + 2\lambda_f(H^*, h) \mathbb{E}[1 - Z_f] \mathbb{E}[1 - Z_f^{1-\gamma}].$$

Equation (12) states that the reduction in expected damages from one unit of investment in dykes must equal the user cost of flood protection. The reduction in expected damages from flooding resulting from investing in dykes (the left-hand side of (12)) is the resulting marginal reduction in the probability of flooding *times* the expected risk-adjusted damages if a flood occurs (the first term in the big brackets). From this, the existing probability of flooding *times* the marginal increase in expected risk-weighted damages conditional on flooding is subtracted (the second term), which corrects for the phenomenon that if flooding occurs with higher dykes, the damage is greater (due to the greater inundation depth of flooding). The sum of these two terms must be positive for investment in flood protection to be optimal. Together, they are multiplied by ϕ'_H/ϕ'_K to convert from units of flood protection to units of physical capital.

The user cost of flood protection equals the risk- and growth-adjusted discount rate, r^* , plus maintenance costs, where the latter equals the depreciation rate of flood protection, $\delta_H \geq 0$, minus the marginal reduction in adjustment costs, $\partial \Theta_H / \partial H < 0$ (the right-hand side of (12)). This user cost potentially depends on the flood arrival rate through the project-specific discount rate r^* as well as the flood protection stock H .

Overall, it is more helpful to think in terms of an optimal flood arrival rate than an optimal flood protection stock.

¹³This term corresponds to volatility, where the effect of economic and climate disasters is corrected for risk aversion γ .

Result 2: *If the flood arrival rate λ_f and the flood damage distribution parameter χ_f depend exponentially on dyke height H (as in (7)-(8)), the damage ratio Z_f follows a power function distribution,¹⁴ and $\alpha > \zeta_H$, we obtain*

$$\lambda_f(H^*, h) \frac{\alpha - \zeta_H (1 + \chi_f(H^*, h)) / (1 + \chi_f(H^*, h) - \gamma)}{1 + \chi_f(H^*, h) - \gamma} \frac{\phi'_H(i_H, H^*)}{\phi'_K(i_K)} = r^* + \delta_H + \frac{\partial \Theta_H(i_H, H^*)}{\partial H}. \quad (14)$$

With zero adjustment costs and risk aversion ($\theta_K = \theta_H = \gamma = 0$), (14) implies

$$\lambda_f(H^*, h) = \frac{\nu}{K_0} (r^* + \delta_H) \left(\frac{1 + \chi_f(H^*, h)}{\alpha - \zeta_H} \right). \quad (15)$$

Hence, according to (15), the probability of flooding is proportional to the initial relative cost of flood protection investment relative to productive capital investment (ν) and the user cost of flood protection ($r^* + \delta_H$). Furthermore, (optimal) flood risk is inversely proportional to expected damages from a flood (i.e. $\mathbb{E}[1 - Z_f] = 1/(1 + \chi_f)$). Finally, flood risk decreases (i.e. more protection) in the marginal effect of dyke heightening on flood risk α , and increases in the percentage increase in expected damages from a one metre increase in the dyke height ζ_H (less protection).

3.1.2 Risk- and growth-adjusted discount rate

With a unit elasticity of intertemporal substitution ($\psi = 1$), the discount rate (13) needed to evaluate optimal flood protection investment from (12) collapses to $r^* = \rho$. The discount rate, with its effect on optimal flood protection, is then unaffected by uncertainty whether regular or coming from economic or flooding disasters. (Uncertainty about the size of flooding damages does affect optimal flood protection directly.) The discount rate is also unaffected by uncertainty if society has no risk aversion ($\gamma = 0$).

To gain better insight into the discount rate (13), we unpack it as

$$r^* = \rho + \psi^{-1} g_e - g_e - \frac{1}{2} (1 + \psi^{-1}) \gamma \sigma_{\text{total}}^2 + \gamma \sigma_{\text{total}}^2 \quad (16)$$

(cf. (13) in Barro (2009), for the price-dividend ratio, and extended by Pindyck and Wang (2013) and Van den Bremer and Van der Ploeg (2021)). The first two terms form the Keynes–Ramsey rule, which give the discount rate as the pure rate of time preference plus income growth multiplied by intergenerational inequality aversion (the inverse of the

¹⁴In which case, we can obtain $\Gamma(H, h) = 1/(\chi_c(H, h) + 1 - \gamma)$ and $\mathbb{E}[1 - Z_c] = 1/(\chi_c(H, h) + 1)$. We require $\chi_c(H, h) > \gamma$.

elasticity of intertemporal substitution). The stochasticity in the model means that we must account for the expected growth rate instead of a deterministic growth rate in the first two terms. The third term corrects for expected flood damages growing at same rate as the economy ($-g_e$), since they destroy a fraction of the capital stock.

The fourth term captures the increase in saving, and reduction in the discount rate, due to prudence (cf. Kimball, 1990)¹⁵. The term $\frac{1}{2}\gamma\sigma_{\text{total}}^2$ is the reduction in economic growth that would offset exactly the effect of uncertainty on expected welfare growth and the discount rate. The total effect of this precautionary premium on the discount rate is multiplied by the coefficient of relative prudence, which for our preferences equals $1 + \psi^{-1}$. The fifth term indicates that damages are pro-cyclical (due to their proportionality to the capital stock) and thus offer insurance value, which increases the discount rate and depresses flood protection investment.

From (13), we see that if $\psi > 1$, macroeconomic uncertainty (via σ_K and λ_e) increases the discount rate and depresses flood protection investment (see also Section 5). But, if $\psi < 1$, macroeconomic uncertainty curbs the discount rate and increases flood protection.

3.2 Optimal flood protection with sea level rise uncertainty

In our simulations we consider trends in the sea level rise and investigate what the effects of sea level rise uncertainty are on flood protection. We find that flooding (as opposed to economic) uncertainty always increases flood protection investment.

4 Application: Flood Management in The Netherlands

About a quarter of the Netherlands lies below sea level and more than half is vulnerable to flooding, including the entire economic core in the west of the country. This makes the Netherlands especially vulnerable to sea level rise. But, the Netherlands is a wealthy, highly densely populated country with the resources to defend the coastline against sea level rise even under more extreme scenarios, at least for the coming decades.¹⁶ Abandonment or inland development, the subject of recent research by Desmet et al. (2021), Balboni (2021), and Hsiao (2023), is not currently considered by policymakers.

¹⁵Kimball (1990) defines the ‘equivalent precautionary premium’ for a source of uncertainty for a two-period consumption-savings model as the certain reduction in future wealth that will increase the marginal value of future consumption by the same amount as the presence of uncertainty.

¹⁶See the six studies forming the ‘Knowledge Programme Sea Level Rise’ (‘Kennisprogramma Zeespiegelstijging’) or KP ZSS, in Dutch, discussed below. The literature list cites a summary available in English.

The hydrological map shown in Figure 1 illustrates the exceedance norms that govern flood protection policy in the Netherlands — for each flood protection segment, the maximum annual probability of failure is given. The areas shaded in very light grey are subject to flooding, and are surrounded by flood protection segments indicated with different colours. Those marked in red have the lowest (i.e. strictest) exceedance norms; those marked in blue have the highest exceedance norms. In general, segments have tighter norms if they protect dyke ring areas that are densely populated, are subject to flooding from the sea, or lie low. Since some dyke ring areas are not completely flat, the location of a breach in the dyke ring can matter greatly for damages as well. This is why the norms apply at the disaggregated dyke ring segment level and not the higher dyke ring level.

Figure 1 can also be used to distinguish between several flood-prone parts of the Netherlands, each with a different adaptation strategy. The southwestern part of the country, south of and including Rotterdam, is mostly protected by storm surge barriers that were expensive to construct. Given ongoing sea level rise, costly fixed investment will have to be undertaken anew in the future. By contrast, the preferred adaptation strategy for the western and northern (or ‘Wadden’) coastlines is sand suppletion. Sand is deposited along the shoreline to protect it as the sea level rises. The parts that are not along the coast are also affected by sea level rise. For the IJssel lake area (centre of the map), the most cost-effective strategy is to strengthen only the outer Afsluitdijk that dams it off from the North Sea. The capacity of pumps moving water from the IJssel lake into the sea must also be increased. For dyke ring areas along rivers, dyke heightening will be necessary, especially nearer to the coast, because sea level rise hinders the outflow of water. Climate scenarios with sea level rise also feature higher peak discharge rates from the rivers.¹⁷

4.1 Current flood policy framework

The most recent revision of flood protection standards in the Netherlands occurred in 2017, when a new set of exceedance norms was enacted in law. This resulted in tighter norms especially around the major rivers. The required engineering work is currently underway, with a target date for completion of 2050.¹⁸ Underpinning the latest set of flood protection standards is an official research project called WV21 (‘Waterveiligheid 21-ste eeuw’), which combines engineering and economic research. The economic component is a cost-benefit analysis in the spirit of van Dantzig (1956), comparing the fixed and variable costs of

¹⁷The information in the above two paragraphs is taken from KP ZSS.

¹⁸The work is necessary because of the tighter norms, but also because research found that existing norms were often not met around rivers due to a failure to anticipate structural failure.



Figure 1: Map of the Netherlands with the annual exceedance norms that must be met for each dyke ring segment by 2050, ranging from 1 in 300 (dark blue) to 1 in 100,000 (red-brown) years. Source: Ministry of Infrastructure and Water Management 2016, p.10.

strengthening flood defences with the discounted flow of benefits in the form of avoided flood damages. The discount rate is not endogenous in the analysis, but is taken from guidelines drawn up to harmonise public investment spending across government departments. To show what happens when it depends on economic and flood risk, we first calibrate the preference parameters to match the risk-free rate and risk premium used by the Dutch government. We then vary different sources of economic uncertainty, as well as sea level

rise uncertainty, thus changing these rates and the optimal exceedance norms.

4.2 Flooding in six dyke ring zones

We use data of different degrees of granularity corresponding to segment level data at the finest resolution, area level data at intermediate degree of resolution, and dyke ring zone at the coarsest resolution to obtain our aggregated data for flooding, sea level rise, maintenance and investment and (maximum) damages from flooding (see Appendix B for a discussion of how we do this and what assumptions are needed). We use flood arrival rates at the zone level, derived from annual failure probabilities of 198 primary ring segments including more than 5 million unique flood events or breaches from 1,846 simulated scenarios with assigned correlated probabilities from Kolen and Nicolai (forthcoming). Aggregating our remaining data to the highest dyke ring zone level, we add six independent flooding disaster processes to the evolution of capital, i.e.

$$dK_t = \left(I_{K,t} - \frac{1}{2}\theta_K \frac{I_{K,t}^2}{K_t} - \delta_K K_t \right) dt + \sigma_K K_t dB_{K,t} - (1 - Z_e) K_t dJ_{e,t} - \sum_{i=1}^{i=6} (1 - Z_{f,i}(H_t, h_t)) K_t dJ_{f,i,t}, \quad (17)$$

where the arrival of flooding in each zone evolves over time according to

$$\lambda_{f,i,t}(H_t, h_t) = P_{0,i} e^{-\alpha(H_t - H_0 - \zeta h_t)} \quad (18)$$

with the aggregate variables defined above. The initial arrival rate of flooding in each zone, $P_{0,i}$, is taken directly from Kolen and Nicolai (forthcoming)¹⁹. For each zone, we take the damages estimated by Kolen and Wouters (2007), corrected for economic growth, to be the 90th percentile of the distribution²⁰. The distribution of damages becomes more unfavourable over time due to sea level rise. We approximate this with the variables in the WV21 data set describing, at the dyke ring area level, the percentage increase in expected damages from a one metre increase in dyke height and sea level rise.

¹⁹The initial flood arrival rate for each zone follows from the standards set for 2050 in the Water Law ('Waterwet') legislation of 2017. Engineering work is currently underway to bring down flood risk along the major rivers, and to counter the increase in risk from climate change.

²⁰I.e., as the 10th percentile of the survival fraction conditional on a disaster strike, which is assumed to follow a power distribution. We do this because Kolen and Wouters indicate that the arrival rate of the events they study is one tenth of the prevailing flood protection norm in the affected areas. The scenario with flooding across zones is assigned to the western coast zone, which has the lowest arrival rate of flooding. Damages include casualties multiplied by a value of 6.7 million euros per death, as used by WV21.

4.3 Sea level rise

Our model of sea level rise along the North Sea Coast (9) can be calibrated to different emission scenarios. Figure 2 displays the projections of the Dutch meteorological agency (KNMI) for sea level rise, under the SSP2-4.5 and SSP5-8.5 scenarios for greenhouse gas emissions, along with our approximation (black lines). The path of sea level rise is similar for the next two decades under a low- and high-emissions scenario, but the trends diverge strongly after about 2060. Figure 2 indicates that our (transformed) Ornstein–Uhlenbeck process with skewness (9), calibrated to the pathways, closely approximates the projected trends of sea level rise and its distribution. For each SSP, we obtain the parameters in (9) by matching to the median of the projection in 2005, 2050, 2100 and 2150, as well as the 5-95% confidence interval in 2100.




Figure 2: Comparison between our model (9) and KNMI projections for cumulative sea level rise and its distribution along the Dutch North Sea coast for the SSP2-4.5 and SSP5-8.5 scenarios. The shaded red area shows the 5-95% confidence band from Van Dorland et al. (2023). The lines in the figure are the fitted 5th, 50th, and 95th percentiles of our model. Red dots indicate the median of the KNMI projection for sea level rise, which is given for 2050 and 2100.

The KNMI projections also include several more extreme ‘low-likelihood high-impact’ scenarios that can occur when warming exceeds 2 to 3°C, when ice cliffs around Antarctica may collapse under their own weight and suddenly speed up ice loss (Van Dorland et al. 2023; see also DeConto et al. 2021). We abstract from these scenarios. More broadly, we reduce the uncertainty facing policymakers by calibrating the stochastic process for sea level

rise separately to the KNMI projections under different SSP scenarios, and assuming that the parameters of the stochastic process are known. Given the specification of the process, and currently observed sea level rise, this translates into known probability densities for total sea level rise at any point in the future. In reality, policymakers do not know what the future paths of emissions and global temperatures will be. Further, given the uncertainties in the physical processes of sea level rise, there ‘will not be a single, uniquely valid approach for estimating the probability of different levels of future change’ (Kopp et al. 2017, p. 1230). We do not address this more fundamental uncertainty here.

The WV21 data contains estimates at the dyke ring segment level of the expected rate of water level increase under a baseline climate scenario and the associated increase in the flood arrival rate. We aggregate these to obtain values for α and ζ in equation (18).

4.4 Calibration: Floods and flood protection

Tables 1 and 2 summarise the calibrated parameters for the flooding and flood protection part of our model (with further details in Appendix B).

Area	Initial flood arrival rate $P_{0,i}$ [year ⁻¹]	Disaster size parameter $\chi_{f,i}$	Dyke height dependence of damages $\zeta_{H,i}$ [m ⁻¹]	Sea level dependence of damages $\zeta_{h,i}$ [m ⁻¹]
Western coast and SWD	1/4,000	60	0.12	0.11
Wadden coast	1/300	179	0.10	0.09
IJssel lake	1/800	461	0.04	0.06
Rhine-Meuse: I	1/400	277	0.11	0
Rhine-Meuse: II	1/300	1093	0.13	0
Rhine-Meuse estuary	1/400	158	0.11	0.10

Table 1: Parameters for the six dyke ring ones. Source: Kolen and Wouters (2007) and Kolen and Nicolai (forthcoming). Note that a smaller flooding size parameter $\chi_{f,i}$ implies a larger expected disaster size. The parameters $\zeta_{H,i}$ and $\zeta_{h,i}$ indicate the percentage increase in expected damages per metre increase in the dyke height and sea level, respectively.

4.5 Calibration: Economy

We set the adjustment cost parameter in the capital investment function to $\theta_K = 39.08$, while the annual volatility and annual arrival rate for economic disasters are set to $\sigma_K = 0.02$ and $\lambda_e = 0.088$, respectively. This combination of parameters reproduces the empirical fact that volatility of aggregate consumption is much less than of dividends (cf. Hambel

Parameter name	Unit	Symbol	Value	Source
Flood risk curvature	metre^{-1}	α	3.73	WV21
Mean water level increase relative to SLR		ζ	1.02	WV21
Flood investment cost	Euros/metre	ν	$42 \cdot 10^9$	See Appendix B
Flood adjustment cost	$\text{metre} \cdot \text{year} / \text{Euros}$	θ_H	$4.57 \cdot 10^{-10}$	ibid.
Depreciation rate	year^{-1}	δ_H	0.01	ibid.
Initial protection stock	metre	H_0	0.52	ibid.
SLR constant	$\text{metre}^{\frac{1}{1+\theta_x}}$	X_0	0.61	SSP2-4.5
			0.42	SSP5-8.5
SLR rate	$(1/\text{year})^{\frac{1}{1+\theta_x}}$	η_X	0.014	SSP2-4.5
			0.0087	SSP5-8.5
SLR skewness		θ_X	5.99	SSP2-4.5
			3.27	SSP5-8.5
SLR volatility	$\text{metre}/\text{year}^{1/2}$	σ_X	0.0056	SSP2-4.5
			0.0070	SSP5-8.5
SLR steady-state mean	$\text{metre}^{\frac{1}{1+\theta_x}}$	\bar{X}	1.04	SSP2-4.5
			1.34	SSP5-8.5

Table 2: Parameters governing the rate of sea level rise, the convexity of flood risk in sea level rise, and investment and maintenance costs. For sea level rise parameters, we give in turn the values for the SSP2-4.5 and SSP5-8.5 scenario; see Van Dorland et al. (2023). The steady-state means of the Ornstein–Uhlenbeck process X and the transformed process for sea level rise h differ. The former are in the last two rows of the table, while the latter are 1.27 m (SSP2-4.5) and 3.49 m (SSP5-8.5).

et al., 2024). The survival fraction of economic disasters follows a power distribution with distribution parameter $\chi_e = 8$, which means an average disaster destroys 11% of capital. The productivity of capital, including human capital, is set to $A = 0.113$, matching Pindyck and Wang (2013) and close to the value in Hambel et al. (2024). We match the annual risk-free discount rate and the macroeconomic risk premium used by the Dutch government in cost-benefit analysis of -1% and 3.25% , respectively, using the expressions in Section 3.1, which hold without sea level rise.²¹ The coefficient of relative risk aversion that matches the risk premium is $\gamma = 4.30$. Given this value, we match the risk-free rate by setting $\rho = 0.049/\text{year}$ if $\psi = 0.5$ and $\rho = 0.025/\text{year}$ if $\psi = 1.5$.

²¹The inaccuracy from the assumption of constant flood protection and flood risk is marginal, because flooding is only a very small part of overall disaster risk for the Netherlands. Given optimal policy, sea level rise does not change this except in the more extreme climate scenarios.

5 Quantitative Assessment

We first show the optimal investment in flood protection over time for two pathways for sea level rise (SSP2-4.5 and SSP5-8.5) under our baseline calibration.²² We then examine the effects of economic and sea level rise uncertainty on optimal investment. Both sources of uncertainty turn out to be important. Our simulations indicate that the effect of economic uncertainty on investment depends on the elasticity of intertemporal substitution, while sea level rise uncertainty unambiguously increases flood protection investment.

The figures below include two panels. The left panel shows optimal investment in flood protection over time. The right panel shows the return period of flooding disasters, given sea level rise and optimal investment in flood protection. From the disaster arrival rates in the dyke ring zones, we calculate the annual probability of a flood occurring anywhere.²³ The return period is defined as the inverse of that number and corresponds to the exceedance norms used in flood protection policy in practice.

Since the parameter values in our calibration are not precisely known, we offer results with alternative calibrations of flood risk and flood protection investment costs. Finally, we provide an overview of the welfare cost of flood risk and sea level rise for all our results.

5.1 Baseline results: no sea level rise uncertainty

To demonstrate the basic mechanisms at work, we first show results without sea level rise uncertainty.²⁴ Thus, we replace each stochastic process for sea level rise by a deterministic process that follows has the same mean.²⁵ Flood protection investment is proportional to the aggregate capital stock or GDP (provided there is no sea level rise uncertainty), and is thus subject to regular macroeconomic shocks and disasters. This means that flood protection investment as percentage of output is deterministic. The left panel of Figure 3 indicates that optimal flood protection investment is significantly higher under the more

²²We solve the HJB equation of our model using finite differences. Specifically, we use the implementation in Achdou et al. (2022), which allows use of MATLAB's sparse matrix tools. Since the flood arrival rate is exponential, it becomes very large in the low- H , high- h corner of the grid. We therefore set the maximum annual flood arrival rate to 1 to ensure convergence. Doubling or halving this maximum arrival rate does not affect the solutions.

²³The combined annual probability of flooding is $1 - \prod_{i=1}^{i=6} e^{-\lambda_{f,i}} \approx \sum_{i=1}^{i=6} \lambda_{f,i}$ if arrival rates $\lambda_{f,i}$ are small. The vertical axis on the right panels is calculated as $(1 - \prod_{i=1}^{i=6} e^{-\lambda_{f,i}})^{-1}$.

²⁴The optimal path of the flood protection stock depends on the optimal investment paths. For subsequent results, we will use the year 2100 as an appropriate investment horizon, but we show a longer timescale here to illuminate the logic underlying the results.

²⁵The volatility parameter σ_X is set to zero, and the mean-reversion parameter η_X is adjusted to ensure the deterministic process equals the mean of the skewed stochastic process with negligible error.

pessimistic global warming scenario SSP5-8.5, and more so as the years progress and the differences in global warming and sea level rise between these two scenarios increases. Sea level rise begins to flatten several decades earlier under the SSP2-2.5 scenario, which causes flood protection investment to fall. In the SSP5-8.5 scenario adjustment costs are high, so investment has to be high too in the first decades to get the same level of protection.²⁶

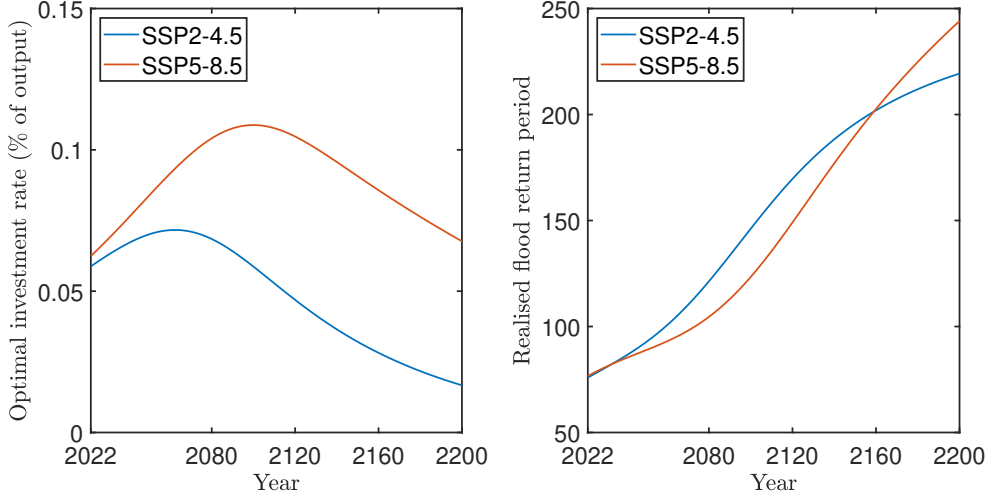


Figure 3: Optimal flood protection investment and the return period of flooding for the global warming scenarios SSP2-4.5 and SSP5-8.5 without sea level rise uncertainty (baseline results).

The right panel of Figure 3 shows that optimal exceedance norms are significantly tighter than the current norms. They increase over time, because investment in flood protection becomes cheaper as the existing flood protection stock increases, while expected damages from flooding increase with sea level rise. Furthermore, adjustment costs of flood protection fall over time as the economy grows.²⁷ For these reasons, flood protection policy eventually becomes more stringent in the SSP5-8.5 than in the SSP2-4.5 scenario, despite the higher investment costs. Once sea level rise stops, flood protection converges to its long-run level (see Section 3). Although the the optimal long-run level of flood protection is considerably higher than its present level in the Netherlands, initial investment is close to actual investment – we provide a fuller comparison with current policy in section 5.8.²⁸

²⁶The shape of the flood protection investment paths is clearly influenced by the specification of adjustment costs. In the flood protection investment function, the adjustment cost parameter is divided by the flood protection stock, so that investment capacity rises linearly with it.

²⁷Adjustment costs fall as $\theta_{H,t}$ falls and ν_t rises in (6).

²⁸Around 360 million Euros per year has been allocated by the Dutch government to flood protection

Sea level rise does not affect the risk-free rate and the macroeconomic risk premium under the SSP5-8.5 scenario much compared with the baseline without sea level rise uncertainty (see Figure 9 in Appendix C). The risk-free rate is only marginally lower and more so for lower ψ , as then the intertemporal substitution effect resulting from lower consumption growth is larger (see the second term on the right-hand side of (16)).²⁹ The risk premium is only marginally higher.

5.2 Effects of higher risk of rare macroeconomic disasters

Figure 4 shows what happens if the risk of macroeconomic disasters is increased by around 44%. This lowers the discount rate and thus increases optimal flood protection investment if the elasticity of intertemporal substitution is below unity (i.e. $\psi = 0.5$). As a result, the right panel indicates that the time for a flood to occur increases significantly. Conversely, if the elasticity of intertemporal substitution is above unity (i.e. $\psi = 1.5$), the higher macroeconomic disaster risks increases the discount rate and thus diminishes optimal flood protection investment. Now it becomes more likely that a flood will occur (lower flood return period). The reason for these opposite effects of macroeconomic disaster risk on the discount rate (13) is that the discount rate contains a prudence and an insurance term as can be seen from (16). For $\psi < 1$, the former dominates, so the discount rate goes down and flood protection investment goes up. For $\psi > 1$, the reverse holds.

Macroeconomic disaster risk has a considerable effect on flood protection investment. The effect is asymmetric, since it diminishes over time if $\psi = 1.5$ but not if $\psi = 0.5$.³⁰ Given sea level rise and the exponential shape of the flood arrival rate, investment cannot remain below its baseline level for very long if $\psi = 1.5$. The higher investment rate if $\psi = 0.5$ is up to a point self-reinforcing, because a larger accumulated flood protection stock makes further investment more efficient through lower adjustment costs. A higher degree of regular macroeconomic uncertainty (higher volatility σ_K) yields similar results, though quantitatively smaller (see Appendix D, along with the corresponding figures for the SSP5-8.5 scenario).

reinforcement until 2050, which is less than in our optimal scenarios.

²⁹Another reason may be that flood protection spending must go up in the near future, increasing the marginal utility of consumption relative to the present. This effect outweighs higher economic (flood) risk pushing up the discount rate for $EIS > 1$.

³⁰For $\psi = 1.5$, the difference relative to baseline becomes negative at some point, because investment costs are high when the accumulated flood protection stock is low.




Figure 4: Effects of a higher risk of macroeconomic disasters on optimal flood protection investment and the return period under the SSP2-4.5 scenario without sea level rise uncertainty. The disaster arrival rate is increased by 44%, so the discount rate falls to 1% for $\psi = 0.5$.

5.3 Effects of sea level rise uncertainty

Figure 5 shows the optimal path for stochastic flood protection investment and the return period with stochastic sea level rise as in (9) and compares them to the baseline where sea level rise follows the mean of the stochastic process.³¹ Optimal flood protection investment paths differ, because the optimal policy response differs when sea level rise is stochastic even when the mean is kept constant. Furthermore, the average of any nonlinear response of optimal investment to variations in sea level rise will differ. With quadratic adjustment costs, the former dominates the latter effect.

5.4 Effects of doubling flood risk

The true flood arrival rate in the Netherlands is only known approximately. Because the level of flood protection is high, major disasters presuppose either meteorological conditions that have not yet been observed or unforeseen failures in flood defences. Such probabilities

³¹The figures display the average of 50,000 simulated paths of flood protection investment and the return period with stochastic sea level rise. Because the process for sea level rise is skewed, which is a convex transformation, the mean of the stochastic process lies above its median. We ensure that the deterministic process coincides with the mean of the stochastic process by slightly increasing the rate of sea level rise parameter η_X when we shut down the volatility parameter σ_X . See Appendix F for details.

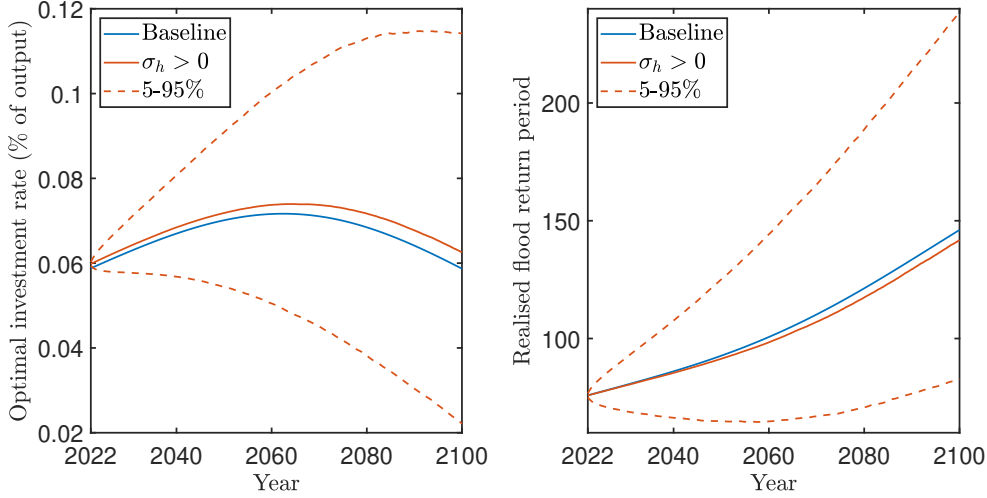


Figure 5: Effects of stochastic sea level rise on optimal flood protection investment and the return period under the SSP2-4.5 scenario. The blue line shows the deterministic path; the red line shows the mean of the stochastic path; and the dashed lines show the 5 and 95 percentiles of the stochastic paths. In the right panel, the red line shows the inverse of the mean flood arrival rate.

are inherently difficult to estimate.³² While the WV21 studies led to an upwards revision in estimated flood risk along rivers, some argue that the probability of many scenarios is overestimated.³³ The return on flood protection investment is determined by flood risk, i.e. the flood arrival rate times expected risk-adjusted flood damages as in (12). Changing either will thus have a similar effect on optimal investment. Figure 6 shows the effects of doubling the initial flood arrival rate in all zones relative to the baseline results.

It turns out that this has little influence on optimal exceedance norms in the long run. To understand this, recall that the long-run optimality condition for the flood arrival rate equates the marginal reduction in risk-adjusted damages from flood protection to its user cost. If initial flood risk is greater, more needs to be invested to bring it down, but that

³²The exponential functional form we use to approximate the flood arrival rate stems from the observed distribution of high water levels along the coast. The initial arrival rates in each zone come from Kolen and Nicolai (forthcoming) and rely in part on subjective expert judgment. The results for damages from Kolen and Wouters (2007) are based on computer-simulated inundation patterns.

³³See, ABN AMRO, ‘Climate change and the Dutch housing market: Insights and policy guidance based on a comprehensive literature review’, 2024, and references therein. The criticism is essentially that the correlation between different extreme events is neglected. Kolen and Nicolai (forthcoming), on which we base our calibration of flood risk, explicitly take this correlation into account.

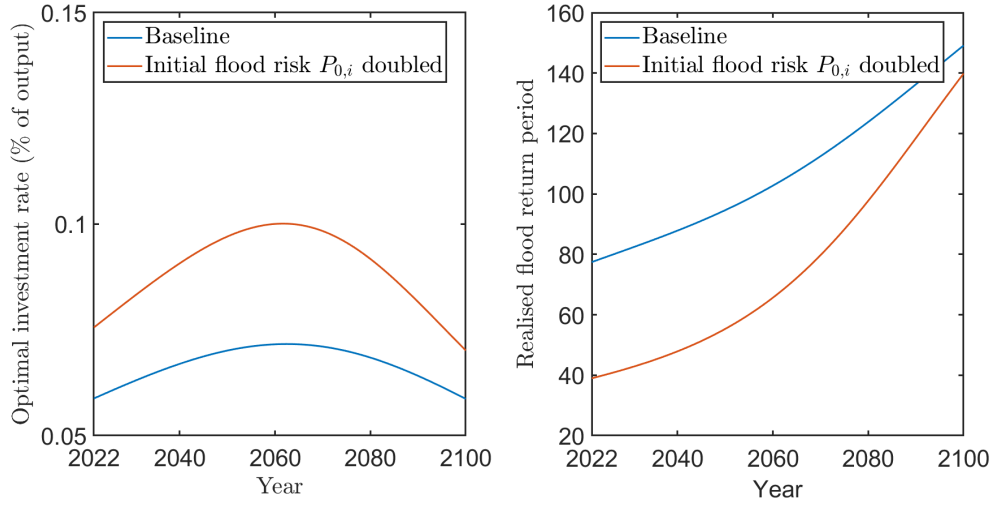


Figure 6: Effects of doubling the initial flood arrival rate, $P_{0,i}$ in each zone on optimal flood protection investment and the return period under the SSP2-4.5 scenario with no sea level rise uncertainty.

matters only if the user cost changes with the level of flood protection.³⁴ The flood arrival rate does take many decades to converge to its baseline level. Early on, there is not enough spare investment capacity to reduce flood risk at an acceptable cost. The capacity only becomes available once sea level rise has flattened.

5.5 Effects of doubling dyke heightening costs

The cost projections in the KP ZSS studies, which we use for calibration, are not very precise. Within the confidence intervals provided, the cost of dyke heightening can be half or double the central estimate. The budget allocated by the Dutch government to achieve the updated flood protection standards by 2050 has overrun by a factor two.³⁵ Figure 7 therefore shows the effects of doubling the cost of dyke heightening.

Doubling investment costs leads to higher, more costly investment, but flood protection falls and in the long run the optimal exceedance norms are roughly halved. Holding the flood protection stock H constant, the marginal damage reduction from investment is in the long run proportional to the flood arrival rate (see the discussion of Results 1 and 2 in

³⁴For example, if total maintenance costs are very nonlinear in the level of flood protection.

³⁵This refers to the HWBP programme (Hoogwaterbeschermingsprogramma). Revised cost projections from 2023 to 2050 are roughly twice the original budget, see for example, in Dutch, Ministry of Infrastructure and Water Management 2024, p.2.

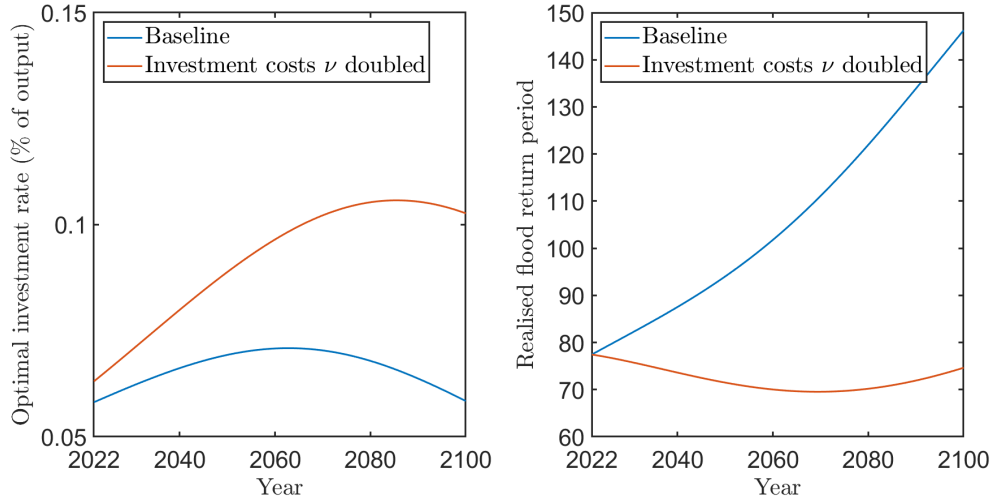


Figure 7: Effects of doubling investment costs on optimal flood protection investment and return period under the SSP2-4.5 scenario with no sea level rise uncertainty. Relative to the baseline, the cost parameter ν is doubled while the adjustment cost parameter θ_H is halved. This way, investment capacity is unaffected (adjustment costs stay constant given dyke heightening).

Section 3). If the user cost changes, the flood arrival rate must change by the same amount to equate the marginal benefit and cost of flood protection investment. In the WV21 data the expected damages from a flood event are increasing in the stock H , which in turn is tied to the flood arrival rate. However, the above reasoning still holds approximately.

5.6 Cumulative dyke heightening

Table 3 displays cumulative dyke heightening, in metres between the years 2022 and 2100, along the optimal path of each scenario. This dyke heightening is averaged over dyke ring areas, weighted by investment cost (see Appendix B for details). Expected sea level rise by 2100 is 0.64 m under the SSP2-4.5 scenario and 0.22 m higher by the end of the century under the SSP5-8.5 scenario with very little mitigation and much more global warming and sea level rise. Doubling flood risks increases dyke heightening, while doubling investment costs curbs dyke heightening. If both initial flood risk and investment costs are doubled, the effect on optimal policy nets approximately to zero³⁶. Higher economic disaster risk leads to more (less) dyke heightening if the elasticity of intertemporal substitution is below

³⁶The net effect is not exactly zero because the optimal discount rate moves down a little.

(above) one. The total amount in Euros of flood protection investment between 2022 and 2100, along the optimal path and discounted at the risk-adjusted discount rate is given in parentheses.³⁷

	SSP2-4.5	SSP5-8.5
Baseline	0.72 (15)	0.91 (18)
With sea level rise uncertainty	0.73 (16)	0.92 (18)
Doubling of initial flood risk	0.89 (20)	1.06 (23)
Doubling of investment costs	0.54 (19)	0.73 (25)
Doubling of flood risk and costs	0.72 (30)	0.91 (36)
Higher economic disaster risk ($\psi = 0.5$)	0.83 (19)	1.01 (21)
Higher economic disaster risk ($\psi = 1.5$)	0.72 (14)	0.89 (17)

Table 3: Cumulative dyke heightening between 2022 and 2100 expressed in metres, and in parentheses the associated discounted investment cost in billion Euros (2022 price level).

5.7 Welfare costs of disaster risk

Pindyck and Wang (2013) show that the welfare cost of disaster risk in their setting can be expressed as an equivalent permanent consumption tax. If there is homogeneity in the capital stock, such a tax is non-distorting and boils down to a loss of wealth. Since this also holds in our model, we use the same welfare measure (see Appendix G). For each of the simulations discussed above, we solve the model with zero sea level rise. We then derive the permanent consumption tax that equates welfare to its value with sea level rise for the different simulations. Table 4 thus presents the willingness to pay (WTP), i.e. the percentage of consumption one is prepared to sacrifice for ever into the future, to avoid sea level rise for each of the simulations, both for the SSP2-4.5 and the SSP5-8.5 scenario.

The baseline results shows that society is prepared to sacrifice permanently 0.085% of consumption to avoid sea level rise in the SSP2-4.5 scenario. However, in the SSP5-8.5 scenario with more emissions and global warming and sea level rise, society is willing to sacrifice almost one-and-half times as much. If sea level rise is stochastic, society is willing to sacrifice a little more in each scenario. If the flood arrival rate were to double, society is willing to increase the sacrifice to avoid sea level rise to a permanent cut in aggregate consumption of 0.12% in the SSP2-4.5 and to 0.16% in the SSP5-8.5 scenario. Finally, if

³⁷Since investment costs scale with output in our model, they are stochastic even if sea level rise is not. We use the expected growth rate to calculate expected costs. Although the optimal discount rates and expected economic growth vary between scenarios, we use their baseline values everywhere for comparability.

	SSP2-4.5	SSP5-8.5
Baseline	0.085	0.12
With sea level rise uncertainty	0.090	0.13
Doubling of initial flood risk	0.12	0.16
Doubling of investment costs	0.14	0.19
Doubling of flood risk and costs	0.16	0.23
Higher economic disaster risk ($\psi = 0.5$)	0.10	0.15
Higher economic disaster risk ($\psi = 1.5$)	0.080	0.11

Table 4: WTP to avoid sea level rise, expressed as a permanent consumption tax (%).

dyke heightening costs double, society is willing to pay roughly one-and-half times as much in each of the two global warming scenarios.

In line with previous contributions in the macro finance literature, we find that society’s willingness to pay to eliminate heightened disaster risk exceeds its expected financial cost. Comparing our baseline model solutions with and without sea level rise, we can calculate the present discounted value (PDV) of increased adaptation investment and flood damages due to sea level rise. This sum is 12 billion Euros under SSP2-4.5, and 16 billion Euros under SSP5-8.5. Meanwhile, the present discounted value of the equivalent (in welfare terms) consumption tax is 17 billion Euros under SSP2-4.5 and 24 billion Euros under SSP5-8.5³⁸.

Since optimal exceedance norms weakly depend on the sea level, the financial cost of sea level rise is mostly that of additional investment. In our baseline calibration, the cost of increased adaptation investment due to sea level rise is roughly 5 times that of increased flood damages under both the SSP2-4.5 and the SSP5-8.5 scenario. Somewhat surprisingly, within each pathway, the effect on welfare of sea level rise uncertainty is muted. Although investment paths are strongly affected by sea level rise uncertainty (see Figure 5), the flood arrival rate can be kept under control without excessive adjustment costs in our calibrated model.³⁹ Most of the variation in investment costs and flood damages then averages out. In limiting total investment capacity, adjustment costs do generate a more sizeable interaction effect between the cost of sea level rise and initial flood risk.

Table 4 also shows results for higher macroeconomic disaster risk when the rate of time preference is adjusted to ensure that the risk-free rate remains matched to the data.

³⁸The PDV of adaptation investment, expected flood damages, and the consumption tax is calculated by simulating the model forward for 1000 years, correcting for the correlation of investment costs, flood damages, and consumption with output, and discounting by the optimal risk-adjusted rate of 2,25%.

³⁹This finding aligns with the results of the KP ZSS studies to which we calibrate the model. There, it is argued that due to fixed costs required investment is in fact weakly concave in sea level rise between 2050 and 2200, provided it follows a predictable path. See, in Dutch, *ibid* p.29.

With $\psi < 1$, this higher risk depresses the discount rate and increases flood protection investment and thus increases the return period. As a result, society is willing to forsake more consumption to avoid sea level rise. With $\psi > 1$, we obtain the opposite effects.⁴⁰

5.8 Comparison with current policies

We find that optimal policy requires a continuous tightening of flood protection standards. The flood protection norms established by legislation in 2017, which we use to calculate the initial flood return period, are thus not sufficiently stringent. In the simulation where we shut down sea level rise it is optimal to more than double the return period of flooding in the long run, from 76 to 194 years. Once we allow for sea level rise, which makes flood events more destructive, the optimal return period in the long run increases further to 242 years (SSP2-4.5) or 378 years (SSP5-8.5).

The Dutch government has initiated a research programme examining the cost and feasibility of maintaining the current flood protection norms in the face of sea level rise. First findings of this programme, called KP ZSS (‘Kennisprogramma Zeespiegelstijging’), were released in 2023. It found that necessary dyke heightening in most dyke ring zones, averaged by the length of dykes, are a little less than sea level rise itself.⁴¹ Our analysis weights the required heightening of segments by the investment cost and finds that it slightly exceeds sea level rise.⁴² KP ZSS also estimated the overall cost of maintaining flood protection standards to be roughly 300 million Euros per year given sea level rise of 0.75 metres by 2200 (ibid p. 32). This estimate is identical to our own because we use it for calibration.

We can also compare our optimal flood protection investment rates with actual expenditures. There is no unified budget for flood protection spending in the Netherlands, but most strengthening of primary flood defences is currently funded under a programme called HWBP (‘Hoogwaterbeschermingsprogramma’). The most recent cost estimate, revised upwards from 360 million Euros, is 857 million Euros per year.⁴³ At 0.076% of output in 2024, this is close to optimal investment in most of our scenarios. Since the optimal flood protection level in our results rises faster than what is planned under HWBP, our baseline calibration of investment costs may thus be somewhat optimistic.

⁴⁰Note that the baseline for the two different values of ψ is the same as ρ is adjusted to match the risk-free interest rate. The welfare cost in the baseline is thus the same for both values of ψ .

⁴¹For example, given sea level rise of 0.75 metres by 2200, required heightening along the coastline is 0.70 metres, and less along the rivers. See KP ZSS 2023 (in Dutch) p.24.

⁴²This can be seen from the fact that $\zeta > 1$. Another difference with the KP ZSS methodology is that we assume the norms are exactly satisfied, while KP ZSS allows for initial dyke height to exceed the norms.

⁴³See, in Dutch, Ministry of Infrastructure and Water Management 2024, p.2. The lower bound for costs is put at 561 million euros per year and the upper bound at twice that level, 1.2 billion euros per year.

6 Conclusion

We have put forward a tractable DSGE framework to assess optimal flood protection investments in the face of gradual sea level rise as well as macroeconomic and flood risks. Policy makers can curtail the flooding risk by investing in flood protection. Our framework allows policymakers to differ in their aversion to intertemporal fluctuations and their aversion to risk, where prudence increases in the former. We have obtained two sets of results.

First, we have derived a rule for the optimal long-run level of flood protection as function of relative risk aversion, the elasticity of intertemporal substitution, and uncertainty of economic growth. We have shown that regular economic uncertainty and especially the risk of rare macroeconomic disasters leads to less (more) flood protection if the elasticity of substitution is greater (less) than one. Second, we have calibrated our model to the Netherlands and shown that sea level rise uncertainty leads to more flood protection. Third, we have shown how changes in flood risk and in dyke heightening costs affect flood protection investment and the flood return period.

Within each climate scenario, our quantitative analysis revealed that doubling initial flood risk or investment costs changes optimal dyke heightening by up to a quarter (by less when sea level rise is greater). Increasing economic disaster risk affects optimal dyke heightening through the optimal discount rate. If the elasticity of intertemporal substitution is less (greater) than 1, total dyke heightening increases (reduces). Sea level rise uncertainty by comparison has a limited effect on optimal policy despite the fact that it is modelled as a skewed distribution.

Our results suggest that the welfare cost of sea level rise under an optimal flood protection policy correspond to up to a 0.23% permanent drop in consumption. In comparison, Desmet et al. (2021) use a model of coastal retreat, abstracting from flood protection and (within-model) sea level rise uncertainty, and find a figure of 1.05% for the welfare loss in PDV terms. This is higher, since it does not allow flood protection to react to flood risk. Hsiao (2023) also shows that adaptation lowers the social cost of sea level rise. For the Netherlands flood protection investment is very cheap compared to the value of land.

This brings us to avenues for future research. First, our analysis deals with *risk*, i.e. uncertain outcomes with known probabilities. But as Barnett et al. (2020) point out, given the wide variety of economic, climate and flood research, one might want to allow for *ambiguity* with unknown weights for alternative possible models and for *misspecification* to capture unknown flaws in the approximation of models. Since many parts of our model are approximate (e.g. the exponential functional form for flood risk), and parameter values are

unknown, it would be of interest to use the robust control techniques developed by Hansen and Sargent (2007) to deal with this type of uncertainty, ambiguity and mis-specification. Second, it is of interest to allow policymakers to learn over time about the parameters driving the stochastic process for sea level rise rather than conditioning all information on past data. Third, policymakers may recognise that fixed costs and irreversibilities in flood protection investments in the presence of uncertainty can lead to hold-up problems. This could be analysed using the real option theory outlined in Dixit and Pindyck (1994). With data on fixed costs, one could examine and test how the ‘saw-tooth’ pattern of investment is affected by flood uncertainty. Finally, our application has been to flood control in a flood-prone rich country. It would be of interest to extend our analysis to allow for heterogeneity among households (rich versus poor, borrowing-constrained or not) and to develop applications to flood-prone developing regions where adaptation is more challenging due to political, borrowing, and other constraints.

Acknowledgement We are grateful for funding from the insurance company a.s.r. We wish to thank Matthijs Kok and Bas Kolen for sharing their knowledge of Dutch flood protection policy with us, and Jarl Kind furthermore for sharing important data sets from the official WV21 research project. We also thank seminar and conference participants at the 2024 Conference on Energy and the Environment in Mannheim, SURED 2024 in Ascona, EAERE 2024 in Leuven, EEA 2024 in Rotterdam, and the 2024 LSE Environment Week in London. Taco Prins thanks colleagues at the University of Amsterdam for helpful advice, in particular Marcelo Pedroni and Andrea Titton.

References

- IPCC. Special report on the ocean and cryosphere in a changing climate. *IPCC Publication*, 2022.
- Rebecca Newman and Ilan Noy. The Global Costs of Extreme Weather that are Attributable to Climate Change. *Nature Communications*, 14, 2023.
- Lint Barrage and Jacob Furst. Housing investment, sea level rise, and climate change beliefs. *Economics Letters*, 177(C):105–108, 2019.
- Laura A. Bakkensen and Lint Barrage. Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics. *The Review of Financial Studies*, 35(8):3666–3709, 2022.

- Stephie Fried. Seawalls and Stilts: A Quantitative Macro Study of Climate Adaptation. *The Review of Financial Studies*, 89(6):3303–3344, 2022.
- Robert E. Kopp, Robert M. DeConto, Daniel A. Bader, Carling C. Hay, Radley M. Horton, Scott Kulp, Michael Oppenheimer, David Pollard, and Benjamin H. Strauss. Evolving Understanding of Antarctic Ice-Sheet Physics and Ambiguity in Probabilistic Sea-Level Projections. *Earth’s Future*, 5(12):1181–1277, 2017.
- M. Haasnoot, J. Kwadijk, J. van Alphen, D. Le Bars, B. van den Hurk, F. Diermanse, A. van der Spek, G. Oude Essink, J. Delsman, and M. Mens. Adaptation to uncertain sea-level rise; how uncertainty in antarctic mass-loss impacts the coastal adaptation strategy of the Netherlands. *Environmental Research Letters*, 15(3):034007, 2020. doi: 10.1088/1748-9326/ab666c.
- David van Dantzig. Economic decision problems for flood prevention. *Econometrica*, 24(3): 276–287, 1956.
- Robert S. Pindyck and Neng Wang. The economic and policy consequences of catastrophes. *American Economic Journal: Economic Policy*, 5(4):306–339, 2013. doi: 10.1257/pol.5.4.306.
- Thomas Douenne. Disaster risks, disaster strikes, and economic growth: the role of preferences. *Review of Economic Dynamics*, 38:251–272, 2020.
- Harrison Hong, Neng Wang, and Jinqiang Yang. Mitigating disaster risks in the age of climate change. *Econometrica*, 91(5):1763–1802, 2023.
- Andrew B. Abel, Avinash Dixit, Janice C. Eberly, and Robert S. Pindyck. Options, the value of capital, and investment. *Quarterly Journal of Economics*, 111(3):753–777, 1996.
- T.D. van der Pol, E.C. van Ierland, and H.-P. Weikard. Optimal dike investments under uncertainty and learning about increasing water levels. *Journal of Flood Risk Management*, 7:308–318, 2014.
- Christoph Hambel, Holger Kraft, and Frederick van der Ploeg. Asset diversification versus climate action. *International Economic Review*, (65):1323–1355, 2024.
- Jerry Tsai and Jessica A. Wachter. Pricing long-lived securities in dynamic endowment economies. *Journal of Economic Theory*, (177):848–878, 2018.

- Robert J. Barro. Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–264, 2009.
- Robert Barro. Environmental protection, rare disasters and discount rates. *Economica*, (82):1–23, 2015.
- David M. Kreps and Evan L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, 1978.
- Darrell Duffie and Larry G. Epstein. Stochastic differential utility. *Econometrica*, 60(2): 353–394, 1992. doi: 10.2307/2951600.
- Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004. doi: 10.1111/j.1540-6261.2004.00670.x.
- Sebastiaan N. Jonkman, Marten M. Hillen, Robert J. Nicholls, Wim Kanning, and Mathijs van Ledden. Costs of adapting coastal defences to sea-level rise— new estimates and their implications. *Journal of Coastal Research*, 290:1212–1226, 2013. doi: 10.2112/jcoastres-d-12-00230.1.
- Robert J. Nicholls, Jochen Hinkel, Daniel Lincke, and Thomas van der Pol. Global investment costs for coastal defense through the 21st century. *World Bank Policy Research Paper 8745*, 2019.
- Peter Zwaneveld and Gerard Verweij. Economic decision problems in multi-level flood prevention. *Bureau for Economic Policy Analysis Discussion Paper*, 2018.
- Ministry of Infrastructure and Water Management. De veiligheid van Nederland in kaart: Eindrapportage (Dutch) [The safety of the Netherlands mapped: Final report]. 2014.
- Jarl Kind. Maatschappelijke kosten-batenanalyse Waterveiligheid 21e eeuw (Dutch) [Social cost-benefit analysis flood protection 21st century]. *Deltares Policy Paper*, 2011.
- Carel Eijgenraam, Ruud Brekelmans, Dick den Hertog, and Kees Roos. Optimal strategies for flood prevention. *Management Science*, 63(5):1644–1656, 2017.
- IPCC. Climate Change 2021: The Physical Science Basis. *Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, 2021.

- Stefano Giglio, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber. Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies*, (34):3527–3571, 2021.
- Valerie Cerra and Sweta Chaman Saxena. Growth Dynamics: The Myth of Economic Recovery. *American Economic Review*, 98(1):439–57, 2008.
- Solomon Hsiang and Amir Jina. The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones. *NBER Working Paper 20352*, 2014. doi: 10.3386/w20352.
- Sandra Batten. Climate change and the macro-economy: a critical review. *Bank of England Staff Working Paper No. 706*, 2018.
- Ton S. van den Bremer and Frederick van der Ploeg. The risk-adjusted carbon price. *American Economic Review*, 111(9):2782–2810, 2021.
- Miles S. Kimball. Precautionary saving in the small and in the large. *Econometrica*, 58(1): 53–73, 1990.
- Klaus Desmet, Robert E. Kopp, Scott A. Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H. Strauss. Evaluating the economic cost of coastal flooding. *American Economic Journal: Macroeconomics*, 13(2):444–486, 2021. doi: 10.1257/mac.20180366.
- Clare Balboni. In harm’s way? Infrastructure investments and the persistence of coastal cities. *Working Paper*, 2021.
- Allan Hsiao. Sea level rise and urban adaptation in Jakarta. *Working Paper*, 2023.
- Ministry of Infrastructure and Water Management. Achtergronden bij de normering van de primaire waterkeringen in Nederland (Dutch) [Background information to the norms of primary flood defences in the Netherlands]. *Report*, 2016.
- Bas Kolen and Robin Nicolai. New framework to generate event set for risk based spatial planning and financial decisions regarding Dutch flood management. *Forthcoming*.
- B. Kolen and C.A.H. Wouters. Als het tóch misgaat: Overstromingsscenario’s voor rampenplannen (Dutch) [If it nonetheless goes wrong: Inundation scenarios for disaster planning]. 2007.

- R. Van Dorland, J. Beersma, J. Bessembinder, N. Bloemendaal, H. van den Brink, M. Brotons Blanes, S. Drijfhout, R. Haarsma, I. Keizer, F. Krikken, D. Le Bars, G. Lenderink, E. van Meijgaard, J.F. Meirink, T. Reerink, F. Selten, C. Severijns, P. Siegmund, A. Sterl, B. Overbeek, H. de Vries, B. Wichers Schreur, and K. van der Wiel. KNMI national climate scenarios 2023 for the Netherlands. *KNMI Scientific Report*, WR-23-02, 2023.
- Robert M. DeConto, David Pollard, Richard B. Alley, Isabella Velicogna, Edward Gasson, Natalya Gomez, Shaina Sadai, Alan Condron, Daniel M. Gilford, Erica L. Ashe, Robert E. Kopp, Dawei Li, and Andrea Dutton. The Paris climate agreement and future sea-level rise from Antarctica. *Nature*, (593):83–89, 2021.
- Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *Review of Economic Studies*, 89:45–86, 2022.
- Mirjam Bani, Ester Barendregt, Marieke Blom, Sander Burgers, Carola de Groot, Rianne Hordijk, Anne Nobel, Sandra Phlippen, and Bram Vendel. Climate change and the Dutch housing market: Insights and policy guidance based on a comprehensive literature review. *ABN Amro Publication*, 2024.
- Ministry of Infrastructure and Water Management and Water Boards. Eindadvies lange termijn programmering (Dutch) [Final recommendations long-term planning]. *Report*, 2024.
- Ministry of Infrastructure and Water Management. System-level analysis of flood protection: Supraregional report (Dutch) [Systeemanalyse waterveiligheid: Bovenregionale rapportage]. 2023.
- Michael Barnett, William Brock, and Lars Peter Hansen. Pricing Uncertainty Induced by Climate Change. *The Review of Financial Studies*, 33(3):1024–1066, 2020.
- Lars Peter Hansen and Thomas J. Sargent. Recursive robust estimation and control without commitment. *Journal of Economic Theory*, 136(1):1–27, 2007.
- Avinash Dixit and Robert S. Pindyck. Investment Under Uncertainty. *Princeton University Press*, 1994.
- KP ZSS. How can the Netherlands cope with sea level rise? *Interim Report for the Sea Level Rise Knowledge Programme*, 2023.

Tweede Kamer der Staten-Generaal. Vaststelling van de begrotingsstaat van het Deltafonds voor het jaar 2023 (Dutch) [Determination of the budgetary state of the Delta Fund for the year 2023]. *Stuk 36 200 J-2*, 2022.

Adrien Bilal and Esteban Rossi-Hansberg. Anticipating climate change across the United States. *Working paper*, 2023.

Lint Barrage. The fiscal costs of climate change in the United States. *Working Paper*, 2024.

Working Group Discount Rate. Rapport werkgroep discontovoet (Dutch) [Report of the working group discount rate]. *Report to the Ministry of Finance*, 2020.

Appendix

A: Derivations

Derivation of Result 1

We have the aggregator function

$$f(C_s, V_s) = \frac{\rho}{1 - \psi^{-1}} \frac{C^{1-\psi^{-1}} - ((1 - \gamma)V)^\omega}{((1 - \gamma)V)^{\omega-1}} \quad (19)$$

with $\omega \equiv \frac{1-\psi^{-1}}{1-\gamma}$. The investment function net of adjustment costs is $\Phi_K(I_{K,t}, K_t)$ and gives the rate at which the consumption good can be converted into productive capital. It is homogeneous of degree one in investment and capital with quadratic adjustment costs, so

$$\Phi_K(I_{K,t}, K_t) \equiv I_{K,t} - \frac{1}{2}\theta_K \frac{I_{K,t}^2}{K_t}. \quad (20)$$

Similarly, the investment function for flood protection net of adjustment costs is defined by

$$\Phi_H(I_{H,t}, K_t, H_t) \equiv \frac{1}{\nu_t} \left(I_{H,t} - \frac{1}{2}\theta_{H,t} \frac{I_{H,t}^2}{H_t} \right). \quad (21)$$

The HJB equation is

$$\begin{aligned}
0 = \max_{I_K, I_H} \{ & f(C, V) + (\Phi_K(I_K, K) - \delta_K K) V_K(K, H, h) + (\Phi_H(I_H, K, H) - \delta_H H) V_H(K, H, h) + \\
& \frac{1}{2} \sigma_K^2 K^2 V_{KK}(K, H, h) + \eta_h(h) V_h(K, H, h) + \frac{1}{2} \sigma_h(h)^2 V_{hh}(K, H, h) + \\
& \lambda_e \mathbb{E}[V(Z_e K, H, h) - V(K, H, h)] + \sum_{i=1}^{i=n_c} \lambda_{f,i}(H, h) \mathbb{E}[V(Z_{f,i} K, H, h) - V(K, H, h)] \}.
\end{aligned} \tag{22}$$

Using $V(K, H, h) = \frac{1}{1-\gamma} (g(H, h) K)^{1-\gamma}$ and (6), and dividing by $(g(H, h) K)^{1-\gamma}$, we obtain

$$\begin{aligned}
0 = \max_{i_K, i_H} \{ & \left(\frac{\rho}{1-\psi^{-1}} \left(\left(\frac{A - i_K - i_H}{g(H, h)} \right)^{1-\psi^{-1}} - 1 \right) + (\phi_K(i_K) - \delta_K) - \frac{1}{2} \gamma \sigma_K^2 + \right. \\
& (\phi_H(i_H, H) - \delta_H H) \frac{g_H(H, h)}{g(H, h)} + \eta_h(h) \frac{g_h(H, h)}{g(H, h)} + \frac{1}{2} \sigma_h(h)^2 \left(\frac{g_{hh}(H, h)}{g(H, h)} - \gamma \left(\frac{g_h(H, h)}{g(H, h)} \right)^2 \right) \\
& \left. - \lambda_e \frac{1}{\chi_e + 1 - \gamma} - \sum_{i=1}^{i=n_c} \lambda_{f,i}(H, h) \frac{1}{\chi_{f,i}(H, h) + 1 - \gamma} \right\}.
\end{aligned} \tag{23}$$

Using the investment adjustment cost functions

$$\phi_K(i_{K,t}) = i_{K,t} - \frac{1}{2} \theta_K i_{K,t}^2 \tag{24}$$

and

$$\phi_H(i_{H,t}) = \frac{1}{\nu} (i_{H,t} - \frac{1}{2} \theta_H \frac{i_{H,t}^2}{H_t}), \tag{25}$$

we obtain the first-order optimality conditions for i_K ,

$$-\rho \left(\frac{A - i_K - i_H}{g(H, h)} \right)^{-\psi^{-1}} \frac{1}{g(H, h)} + (1 - \theta_K i_K) = 0, \tag{26}$$

and for i_H ,

$$-\rho \left(\frac{A - i_K - i_H}{g(H, h)} \right)^{-\psi^{-1}} \frac{1}{g(H, h)} + \frac{1}{\nu} \left(1 - \frac{\theta_H}{H} i_H \right) \frac{g_H(H, h)}{g(H, h)} = 0. \tag{27}$$

Plugging (26) into the HJB equation gives

$$\begin{aligned}
A - i_K - i_H = & -\frac{1 - \psi^{-1}}{1 - \theta_K i_K} \left(-\frac{\rho}{1 - \psi^{-1}} + (\phi_K(i_K) - \delta_K) - \frac{1}{2} \gamma \sigma_K^2 + \right. \\
& (\phi_H(i_H, H) - \delta_H H) \frac{g_H(H, h)}{g(H, h)} + \eta_h(h) \frac{g_h(H, h)}{g(H, h)} + \frac{1}{2} \sigma_h(h)^2 \left(\frac{g_{hh}(H, h)}{g(H, h)} - \gamma \left(\frac{g_h(H, h)}{g(H, h)} \right)^2 \right) \\
& \left. - \lambda_e \frac{1}{\chi_e + 1 - \gamma} - \sum_{i=1}^{i=n_c} \lambda_{f,i}(H, h) \frac{1}{\chi_{f,i}(H, h) + 1 - \gamma} \right). \tag{28}
\end{aligned}$$

Without sea level rise, the terms involving η_h and σ_h drop out of (28). There are three unknowns: i_H , i_K , and H . Given that there is a steady state for the flood protection stock, let i_H be the (lowest) level of flood protection investment such that $\phi_H(i_H, H) = \delta_H H$. Then, (28) becomes

$$\begin{aligned}
A - i_K - i_H + \theta_K(-A i_K + i_K^2 + i_K i_H) + (1 - \psi^{-1})(i_K - \theta_K \frac{1}{2} i_K^2) = \\
-(1 - \psi^{-1}) \left(-\frac{\rho}{1 - \psi^{-1}} - \delta_K - \frac{1}{2} \gamma \sigma_K^2 - \lambda_e \frac{1}{\chi_e + 1 - \gamma} - \sum_{i=1}^{i=n_c} \lambda_{f,i}(H, h) \frac{1}{\chi_{f,i}(H, h) + 1 - \gamma} \right). \tag{29}
\end{aligned}$$

Rearranging terms gives

$$\begin{aligned}
& \frac{1}{2} \theta_K (1 + \psi^{-1}) i_K^2 - \left(\frac{1}{\psi} + (A - i_H) \theta_K \right) i_K = \\
-A - i_H - (1 - \psi^{-1}) & \left(-\frac{\rho}{1 - \psi^{-1}} - \delta_K - \frac{1}{2} \gamma \sigma_K^2 - \lambda_e \frac{1}{\chi_e + 1 - \gamma} - \sum_{i=1}^{i=n_c} \lambda_{f,i}(H, h) \frac{1}{\chi_{f,i}(H, h) + 1 - \gamma} \right). \tag{30}
\end{aligned}$$

The third condition comes from taking the derivative of (23) with respect to H (the envelope condition). The first-order optimality conditions give $g(H, h)^{(\psi^{-1}-1)} = (A - i_K - i_H)^{\frac{\psi^{-1}-1}{1-\psi}} \left(\frac{\rho}{1 - \theta_K i_K} \right)^{-1}$. Using this and $\frac{g_H(H)}{g(H)} = \nu \frac{(1 - \theta_K i_K)}{(1 - \frac{\theta_H}{H} i_H)}$ gives

$$\begin{aligned}
& -\nu(A - i_K - i_H(H)) \frac{(1 - \theta_K i_K)^2}{(1 - \frac{\theta_H}{H} i_H(H))} + \frac{1}{2} \frac{\theta_H}{H^2} i_H(H)^2 \frac{(1 - \theta_K i_K)}{(1 - \frac{\theta_H}{H} i_H(H))} - \nu \delta_H \frac{(1 - \theta_K i_K)}{(1 - \frac{\theta_H}{H} i_H(H))} + \\
& \alpha \lambda_f(H, h) \frac{1}{\chi_f(H, h) + 1 - \gamma} + \lambda_f(H, h) \left(\frac{1}{\chi_f(H, h) + 1 - \gamma} \right)^2 \frac{\partial \chi_f(H, h)}{\partial H} = 0.
\end{aligned} \tag{31}$$

Finally, the risk-adjusted discount rate is the expected return on equity, $\frac{dQ_t + D_t - dt}{Q_t}$, where Q_t denotes the value of equity and D_t denotes dividends. In equilibrium, $C_t = D_t$ and $Q_t = \frac{K_t}{1 - \theta_K i_K}$. In steady state, $C_t = cK_t$ and $Q_t = qK_t$, so that $\frac{dQ_t + D_t - dt}{Q_t} = (A - i_K - i_H)(1 - \theta_K i_K) + \frac{dK_t}{K_t}$. The risk- and growth-adjusted discount rate is then $(A - i_K - i_H)(1 - \theta_K i_K)$. To obtain Result 2, replace this with r^* in the first term of (31). The formulae for r^* in the text can be recovered straightforwardly using equation (28), taking $1 - \theta_K i_K$ to the left-hand side of the equation.

Derivation of Equation (40)

Given Assumption 3 in Section 6, if there are dyke ring areas i and j with arrival rates of flooding given by equation (39), then

$$\frac{\lambda_{f,i}(H_{i,t+dt}, h_{i,t+dt})}{\lambda_{f,i}(H_{i,t}, h_{i,t})} = \frac{\lambda_{f,j}(H_{j,t+dt}, h_{j,t+dt})}{\lambda_{f,j}(H_{j,t}, h_{j,t})}, \forall t, dt \geq 0. \tag{32}$$

Writing $dH_{i,t}$ for $H_{i,t+dt} - H_{i,t}$, and similarly for other variables, it follows that

$$\alpha_i(dH_{i,t} - dh_{i,t}) = \alpha_j(dH_{j,t} - dh_{j,t}), \forall i, j, t. \tag{33}$$

According to Assumption 1, adjustment costs apply to investment summed over all areas. Since adjustment costs are the only source of nonlinearity in investment costs, at the dyke ring level we can define constant local cost factors ν_i that describe the relative cost of dyke heightening. In particular, we can write

$$dH_{i,t} = \frac{\nu}{\nu_i} \frac{i_{H,i}}{i_H} dH_t. \tag{34}$$

The rate of dyke heightening in an area is proportional to the amount of investment allocated to it, and inversely proportional to the local cost; i_H and ν are aggregates which sum over areas; and $dH_t = \phi(i_{H,t}, H_t)$. Finally, Assumption 2 specifies that the rate of

increase of the mean water level in an area is proportional to sea level rise, so that

$$dh_{i,t} = \zeta_i dh_t. \quad (35)$$

Then, some rearranging shows that

$$\alpha_i \left(\frac{\nu}{\nu_i} \frac{i_{H,i}}{i_H} \phi(i_H, H_t) - dh_{i,t} \right) = \frac{1}{\sum_i \left(\frac{\nu_i/\nu}{\alpha_i} \right)} \left(\phi(i_H, H_t) - \sum_i \frac{\nu_i \zeta_i}{\nu} dh_t \right). \quad (36)$$

Define $\alpha \equiv \frac{1}{\sum_i \left(\frac{\nu_i/\nu}{\alpha_i} \right)}$, $H_t \equiv \sum_i \frac{\nu_i}{\nu} H_{i,t}$ (directly from (34)) and $\zeta \equiv \sum_i \frac{\nu_i}{\nu} \zeta_i$. The formula for local flood risk then gives

$$d\lambda_{f,i,t} = -\alpha(dH_t - \zeta dh_t)\lambda_{f,i,t}, \quad (37)$$

which given the initial flood arrival rate $P_{0,i}$ integrates to

$$\lambda_{f,i}(H_t, h_t) = P_{0,i} e^{-\alpha(H_t - H_0 - \zeta h_t)}. \quad (38)$$

Compared to the dynamic programming problem in which local investment rates are allowed to vary independently, Assumption 1 involves a simplification. Given Assumption 2, each $h_{i,t}$ is stochastic, containing the Brownian motion from the process for aggregate sea level rise. Therefore, the policy functions $i_{H,i,t}$ that would keep the relative arrival rates of flooding constant everywhere would not be \mathcal{F}_t -measurable. The simplified problem assumes away a source of uncertainty, i.e. the relative increase in flood risk between locations.

B: Data Sources, Aggregation, and Calibration

The aforementioned WV21, which was carried out between 2005 and 2011, forms the backbone of our data relating to flood risk and flood protection. It contains, at the dyke ring segment level, estimates of the fixed and variable (per metre) costs of heightening the dykes, the local structural rate of sea level rise, and the local convexity of flood risk in the sea level⁴⁴. The second major data source is KP ZSS ('Kennisprogramma Zeespiegelstijging'), an official research programme that explores the required policy response to sea level rise in the Netherlands from 2050 onwards. The first round of findings, released in 2023, contain estimates for the amount and costs of dyke heightening, sand suppletion, and increased pump capacity necessary under different scenarios for cumulative sea level rise between 2050 and 2200. This data is provided at a more aggregated level of six zones: the south-western

⁴⁴We thank Dr. Jarl Kind for sharing the relevant WV21 data with us.

Delta, the western coast, the (northern) Wadden coast, the IJssel lake, the Rhine-Meuse area, and the Rhine-Meuse estuary (abbreviated RMM). As briefly sketched above, these zones face distinct challenges under scenarios with sea level rise.

We use two further data sources. First, to obtain current spending on maintenance and investment, we refer to the budget enacted by the Dutch parliament in 2022, covering all flood protection spending.⁴⁵ Second, as explained below, we use a third government-commissioned study from 2007 for a plausible estimate of maximum damages from flooding.

Data aggregation

The single nationwide flood protection and climate variables in our model should be viewed as aggregates that combine existing flood protection, sea level rise, and flood risk variables at a more local level. In the Netherlands, one such level is the dyke ring area. By definition, flooding in each dyke ring area is a separate event, because flooding in multiple areas cannot result from a single dyke failure. In flood protection policy, exceedance norms were traditionally set at the dyke ring area level, before being further disaggregated.⁴⁶ We make several assumptions to obtain a mapping from the disaggregated WV21 data to the aggregate variables in the model. First, the flood arrival rate in each dyke ring area depends on local sea level rise and local flood protection so that flooding in one area does not cause or prevent flooding elsewhere. Second, when flood protection investment is divided over dyke ring areas, adjustment costs apply to aggregate investment as constraints on investment capacity exist at the national level (e.g. due to a limited number of engineers). Third, sea level rise in each dyke ring area depends linearly on sea level rise along the North Sea coast. This allows for different areas to be more or less exposed to sea level rise, but does not allow for unrelated land subsidence. Finally, flood protection investment does not affect the spatial distribution of flood arrival rates.⁴⁷

Implicitly, we also assume that the spatial distribution of vulnerable productive capital remains constant over time. Since flooding in a dyke ring area destroys a fraction of the local capital stock, we would otherwise need to keep track of this distribution to calculate expected damages. There is a literature that examines how economic activity in flood-prone areas could respond to sea level rise, which contains opposing views.⁴⁸

⁴⁵See, in Dutch, Tweede Kamer der Staten-Generaal, parliamentary paper 36 200 J, 2022.

⁴⁶Since the Waterwet of 2017, they are set at the dyke ring segment level.

⁴⁷For example, if the arrival rate of flooding is halved in one dyke ring area, it must be halved in all areas. Policymakers thus cannot allocate relatively more investment to some areas than to other areas.

⁴⁸For the United States, more theoretical work (e.g. Bilal and Rossi-Hansberg 2023) indicates that migration away from flood-prone areas will reduce the amount capital at risk, with ambiguous welfare implications,

Suppose that there are $1, 2, \dots, n_f$ dyke ring areas with n_f flooding processes $J_{f,i,t}$, with the arrival rate and expected damages depending on local flood protection and sea level rise. Then, the arrival rate of flooding in each dyke ring area i is

$$\lambda_{c,i,t}(H_{i,t}, h_{i,t}) = P_{0,i} e^{-\alpha_i(H_{i,t} - H_{i,0} - h_{i,t})}. \quad (39)$$

Appendix A shows that under the above assumptions the local flood protection stock $H_{i,t}$ can be related to an aggregate H_t defined as its cost-weighted average, and local sea level rise $h_{i,t}$ to sea level rise along the North Sea coast h_t , i.e.

$$\lambda_{c,i,t}(H_t, h_t) = P_{0,i} e^{-\alpha(H_t - H_0 - \zeta h_t)}, \quad (40)$$

where α and H_0 are cost-weighted averages of variables at the dyke ring area level, and ζ weights sea level rise h_t by its average effect on the water level in dyke ring areas. All variables can be calculated from the WV21 data. The policy maker changes the arrival rate of flooding, relative to its initial level, by the same amount in all zones – this is the key simplifying assumption. The model thus remains computationally tractable, but requires an assumption on how flood events across dyke ring areas are related. To simplify these correlations, it is convenient to define an aggregated dyke ring zone level. Instead of including a disaster process for each dyke ring area, we group dyke ring areas by dyke ring zone.

Flood risk

The WV21 data contains the estimated annual failure probability for all 198 primary dyke ring segments of the Netherlands, but not probabilities for combinations of segments. To approximate these, we refer to Kolen and Wouters (2007) and Kolen and Nicolai (forthcoming). The first study was tasked with constructing ‘worst credible’ flood scenarios that could occur in the Netherlands, given flood protection levels. It found that even under extreme scenarios flooding would be mostly contained in one of six zones.⁴⁹ This was due to the distinct nature of flood risk in these zones (e.g. river-based versus sea-based flooding), and storms being at peak strength for a limited amount of time. The main exception to this

while more empirical projections are that migration will increase the amount of capital at risk. See Barrage (2024) on the fiscal cost of climate change, which makes the same assumption as we do.

⁴⁹They are roughly the same zones mentioned above as used by the KP ZSS studies, i.e. the south-western Delta, the western coast, the (northern) Wadden coast, the IJssel lake, the Rhine-Meuse area, and the Rhine-Meuse estuary (RMM). Kolen and Wouters used slightly larger and overlapping zones, however.

rule was a scenario featuring a hurricane over the North Sea battering the western coastline, which we picture in Figure 8 below.

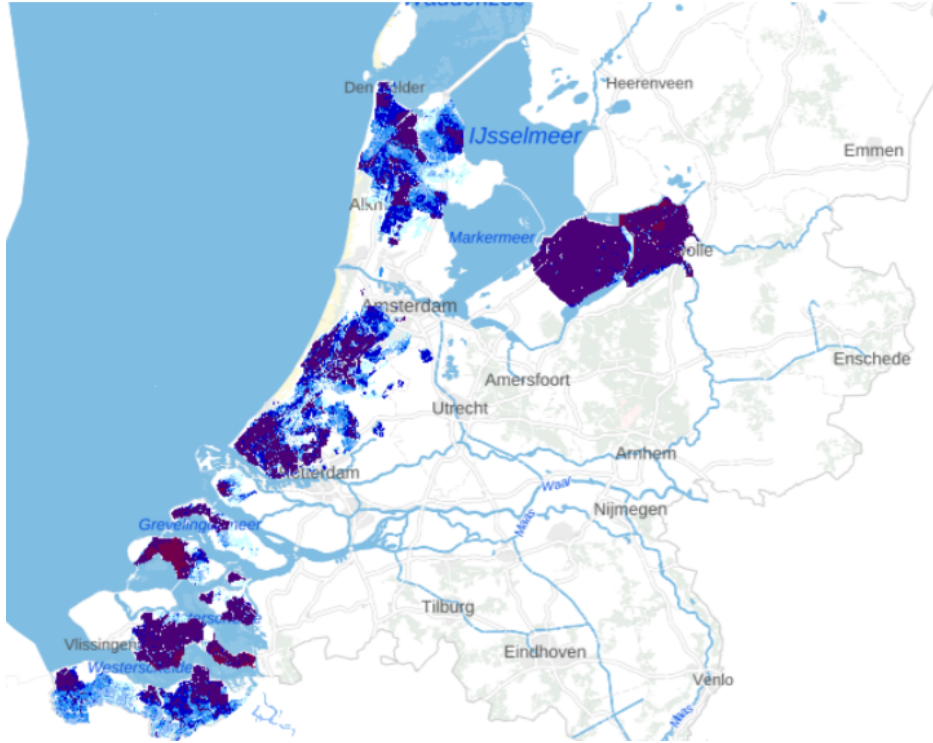


Figure 8: A map showing the affected dyke ring areas in one of the scenarios of Kolen and Wouters (2007), causing widespread damage in both the western coast and south-western Delta areas. The return period of each scenario is related to the prevalent norm in the affected zones - in this case, more than 10,000 years. Source: Kolen and Wouters (2007), accessed online through LIWO ('Landelijk Informatiesysteem Water en Overstromingen').

Moving beyond these worst cases, Kolen and Nicolai (forthcoming) construct more than 5 million unique flood events, defined as combinations of breaches, from 1,846 simulated scenarios. They then assign probabilities to each event using correlation matrices. For our purposes, accounting for all these events in the model would be technically feasible but cumbersome. Instead, we take summary statistics from the paper that give the annual probability of flooding for each of the six zones mentioned above, given the exceedance norms at the dyke ring segment level and the correlation structure of flooding within each zone. We take flooding across zones to be independent, except for the extreme scenario shown above.

Calibration of flood protection investment

Four parameters in the model relate to the cost of flood protection investment: the relative (to productive capital) cost parameter ν , the adjustment cost parameter θ_H , the average initial dyke height H_0 , and the depreciation rate of flood protection δ_H .

We obtain three calibration targets directly from empirical data:

1. From the latest budget adopted by the Dutch parliament, the projected maintenance costs for flood protection between 2022 and 2036 is 246 million Euros per year⁵⁰.
2. The KP ZSS studies suggest that the cost of additional flood protection investment necessary to maintain flood protection standards, assuming sea level rise proceeds at 0,5 centimeters per year from 2050 onwards, is 311 million Euros per year.⁵¹
3. Jonkman et al. (2013), suggest that the depreciation rate of several types of flood defenses is 1%-2% per year.

We first solve numerically the optimal control problem of minimising total investment costs subject to achieving some increase in flood protection, and then find combinations of parameters that satisfy the above requirements on the cost-minimising path.

The adjustment costs in investment cannot be observed directly, but imply a maximum level of productive investment. In the KP ZSS studies, the maximum sea level rise that can be offset in its effect on the flood arrival rate by investment in flood protection between 2050 and 2200 is put around 5,15 metres⁵². We use this number to fix adjustment costs for most results, but also provide extensive comparative statics for different parameter values.

Calibration of preference parameters

We calibrate the preference parameters so that, under the baseline values for economic volatility and economic and climate disaster risk, the optimal risk-free rate and equity

⁵⁰This figure does not account for expenditure from local Water Boards and so is a lower bound for maintenance costs.

⁵¹Two studies within the KP ZSS programme report the cost of the required strengthening of hard structures given sea level rise (295 million Euros per year), and the required volume in cubic metres of sand suppletions. The cost of sand suppletions in the Netherlands per m³ reported in Jonkman et al. (2013) is used to convert the latter figure into Euros.

⁵²The results of KP ZSS indicate that sea level rise of 3 metres can probably be met by current technology and methods, 'though only with great and far-reaching efforts' (see, in Dutch, 'Tussenbalans van het Kennisprogramma Zeespiegelstijging', p.41). A preliminary cost calculation for a potential sea level rise of 5,15 metres is reported in the same document.

returns match the rates used by the Dutch government. Currently, the discount rate applicable to most government investment projects is 2.25%. This comprises a risk-free rate of -1% and a risk premium of 3.25%. The guidelines allow different rates to be used in exceptional cases, for example a lower rate of 1.60% for projects with mostly fixed costs and risk-free returns, but these are not used in the cost-benefit analysis for flood protection investment. The commission establishing the discount rates to be used is the ‘Werkgroep discontovoet’. For its latest report, see, in Dutch, ‘Rapport Werkgroep discontovoet 2020’. (Given that the era of negative interest rates has gone, a new commission is currently re-considering the discount rate to be used.) The official 2020 report on the discount rate refers mostly to market returns enjoyed by private investors as a basis for these estimates. As set out in the report, the commission uses this indirect approach to welfare analysis because it sees the values of preference parameters as too uncertain for practical purposes. A calibrated Ramsey rule is used only as a robustness check on the discount rate obtained from market returns. In contrast, we show how economic and climate uncertainty affect optimal exceedance norms once they are incorporated explicitly in the policy framework.

C: Optimal risk-free rate and risk premium




Figure 9: Optimal risk-free rate and risk premium under the SSP2-4.5 and SSP5-8.5 scenarios, compared to a scenario with no sea level rise.

D: Effects of Higher Regular Macroeconomic Uncertainty

The two figures below show what happens if regular macroeconomic volatility σ_K is increased by around 44% in the SSP2-4.5 and SSP5-8.5 scenarios. They complement Figure (4) in the main text, which considers economic disaster risk. The effects are similar but quantitatively much smaller, reflecting the lesser importance (relative to disaster risk) of regular volatility in the optimal discount rate.

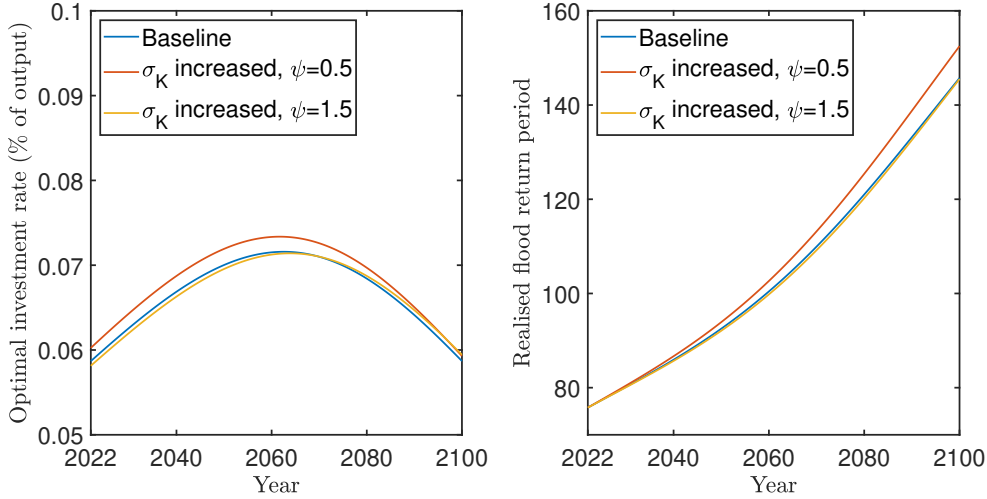


Figure 10: Effects of higher macroeconomic volatility σ_K on optimal flood protection investment and the return period under the SSP2-4.5 scenario without sea level rise uncertainty. The volatility is increased by 44%.

The third figure shows what happens if the risk of macroeconomic disasters is increased by around 44% in the SSP5-8.5 scenario, instead of SSP2-4.5. It is otherwise the exact analogue of Figure (4) in the main text. The effects are similar: higher economic uncertainty increases (decreases) investment if $\psi = 0.5$ ($\psi = 1.5$). The difference is asymmetric since investment does not fall far below its baseline level.

E: Effects of Sea Level Rise with Arctic Ice Sheet Collapse

Figure 13 shows the projection by Van Dorland et al. (2023) of sea level rise given abrupt collapse of the Arctic Ice Sheet. No confidence interval is given for that projection. The calibration of the skewed Ornstein–Uhlenbeck process seeks to match projected sea level rise in the extreme scenario in 2040, 2050, and 2060 as the 99th percentile of the process. The median of the process is fixed to match observed sea level rise in 2005. The data provided by




Figure 11: Effects of higher macroeconomic volatility σ_K on optimal flood protection investment and the return period under the SSP5-8.5 scenario without sea level rise uncertainty. The volatility is increased by 44%.




Figure 12: Effects of a higher risk of macroeconomic disasters on optimal flood protection investment and the return period under the SSP5-8.5 scenario without sea level rise uncertainty. The disaster arrival rate is increased by 44%, so that the discount rate falls to 1%.

Van Dorland et al. (KNMI) end in 2060. The parameter values of the calibrated Ornstein–Uhlenbeck process with skew are $X_0 = 0.75$, $\eta_X = 0.0023$, $\theta_X = 9.06$, $\sigma_X = 0.0035$,

$\bar{X} = 1.43$. The calibration is not satisfactory because the very high skewness parameter translates to very high median sea level rise after 2100. We therefore do not calculate optimal policy for this scenario. To allow for the possibility of Arctic Ice Sheet collapse, a separate ‘tipping point’ process would need to be constructed.

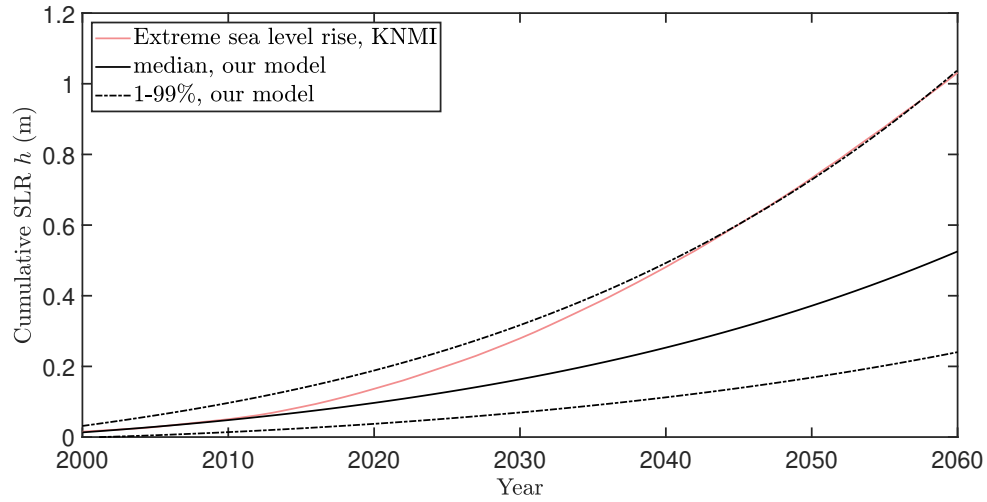


Figure 13: The red line is projected sea level rise given abrupt collapse of the Arctic Ice Sheet provided by KNMI. The solid black line is the fitted median, and the dotted lines are the 1st and 99th percentiles of the fitted process.

F: Modelling Sea level rise uncertainty

We now discuss how we compare optimal policy with and without sea level rise uncertainty. Since the Ornstein–Uhlenbeck process is skewed, setting the volatility parameter σ_X to zero reduces expected sea level rise. To correct for this, we adjust the parameter η_X .

As discussed in subsection 2.6, we calibrate a process for sea level rise that is written as $h_t = X_t^{1+\theta_X} - X_0^{1+\theta_X}$, with $dX_t = \eta_X(\bar{X} - X_t)dt + \sigma_X dB_{X,t}$. The probability density function of X_t is $\frac{e^{-\frac{(X - X_0 e^{-\eta_X t} - \bar{X}(1 - e^{-\eta_X t}))^2}{2\sigma_X^2(1 - \exp(-2\eta_X t))/(2\eta_X)}}}{\sigma_X \sqrt{2\pi(1 - \exp(-2\eta_X t))/(2\eta_X)}}$. We also have $\frac{dh_t}{dX_t} = (1 + \theta_X)X_t^{\theta_X}$, which we can rewrite as $\frac{dh_t}{dX_t} = (1 + \theta_X)(h_t + X_0^{1+\theta_X})^{\frac{\theta_X}{1+\theta_X}}$. By change of variables, the PDF of h_t can be written as

$$f_t(h) = \frac{e^{-\frac{((h + X_0^{1+\theta_X})^{\frac{1}{1+\theta_X}} - X_0 e^{-\eta_X t} - \bar{X}(1 - e^{-\eta_X t}))^2}{2\sigma_X^2(1 - \exp(-2\eta_X t))/(2\eta_X)}}}{\sigma_X \sqrt{2\pi(1 - \exp(-2\eta_X t))/(2\eta_X)}(1 + \theta_X)(h + X_0^{1+\theta_X})^{\frac{\theta_X}{1+\theta_X}}}. \quad (41)$$

The mean of h_t at time t is $\int_{-\infty}^{\infty} f_t(h)h dh$, which can be found by numerical integration.

If we compare optimal policy with and without sea level rise uncertainty, we adjust the rate of sea level rise in the latter case to ensure that the mean of both processes are the same in 2100. The alternative process is $dX'_t = \eta'_X(\bar{X} - X'_t)dt$, with $X'_0 = X_0$, so that $h'_t = (X_0 e^{-\eta'_X t} + \bar{X}(1 - e^{-\eta'_X t}))^{1+\theta_X} - X_0^{1+\theta_X}$, and we set η'_X so that $h'_t = \mathbb{E}[h_t]$ in 2100 (in practice, the means do not visibly differ anywhere). We get $\eta'_X = 0.0143$ ($\eta_X = 0.0140$) for SSP2-4.5, $\eta'_X = 0.00876$ ($\eta_X = 0.00866$) for SSP5-8.5, and $\eta'_X = 0.00479$ ($\eta_X = 0.00472$) for the scenario of extreme sea level rise.

G: Welfare Measure - Permanent Consumption Tax

Here we show that a permanent consumption tax does not change optimal investment rates in capital and flood protection, so that it can be used to compare welfare with and without sea level rise. It extends the derivations of Appendix C in Pindyck and Wang (2013) to our model. Let $V(K, H, \tau)$ denote the value of consumption with sea level rise h fixed at its 2022 level, but with a permanent tax removing a part of consumption. Write $V(K, H, \tau) = \frac{1}{1-\gamma}(g(H, \tau)K)^{1-\gamma}$. With the tax and without sea level rise, the reduced-form HJB equation becomes

$$0 = \max_{i_K, i_H} \left\{ \left(\frac{\rho}{1 - \psi^{-1}} \left(\left(\frac{(1 - \tau)(A - i_K - i_H)}{g(H, \tau)} \right)^{1 - \psi^{-1}} - 1 \right) + (\phi_K(i_K) - \delta_K) - \frac{1}{2} \gamma \sigma_K^2 + \right. \right. \\ \left. \left. (\phi_H(i_H, H) - \delta_H H) \frac{g_H(H, \tau)}{g(H, \tau)} - \lambda_e \frac{1}{\chi_e + 1 - \gamma} - \sum_{i=1}^{i=n_c} \lambda_{f,i}(H) \frac{1}{\chi_{f,i}(H) + 1 - \gamma} \right\}. \quad (42)$$

The first-order optimality condition for i_K becomes

$$-(1 - \tau) \rho \left(\frac{(1 - \tau)(A - i_K - i_H)}{g(H, \tau)} \right)^{-\psi^{-1}} \frac{1}{g(H, \tau)} + (1 - \theta_K i_K) = 0, \quad (43)$$

which gives

$$c = \left(\frac{(1 - \tau) \rho}{1 - \theta_K i_K} \right)^\psi g(H, \tau)^{1 - \psi}. \quad (44)$$

The first-order optimality condition for i_H becomes

$$-(1 - \tau) \rho \left(\frac{(1 - \tau)(A - i_K - i_H)}{g(H, \tau)} \right)^{-\psi^{-1}} \frac{1}{g(H, \tau)} + \frac{1}{\nu} \left(1 - \frac{\theta_H}{H} i_H \right) \frac{g_H(H, \tau)}{g(H, \tau)} = 0, \quad (45)$$

which gives

$$i_H = \frac{H}{\theta_H} (1 - \nu (1 - \theta_K i_K) \frac{g(H, \tau)}{g_H(H, \tau)}). \quad (46)$$

Finally, plugging (43) into (42) gives

$$A - i_K - i_H = -\frac{1 - \psi^{-1}}{1 - \theta_K i_K} \left(-\frac{\rho}{1 - \psi^{-1}} + (\phi_K(i_K) - \delta_K) - \frac{1}{2} \gamma \sigma_K^2 + \right. \\ \left. (\phi_H(i_H, H) - \delta_H H) \frac{g_H(H, \tau)}{g(H, \tau)} - \lambda_e \frac{1}{\chi_e + 1 - \gamma} - \sum_{i=1}^{i=n_c} \lambda_{f,i}(H) \frac{1}{\chi_{f,i}(H) + 1 - \gamma} \right). \quad (47)$$

Now suppose that $g(H, \tau) = (1 - \tau)g(H, 0)$ for $0 \leq \tau < 1$. Then, $\frac{g_H(H, \tau)}{g(H, \tau)} = \frac{g_H(H, 0)}{g(H, 0)}$, so any i_K and i_H that solve (46) and (47) for $\tau = 0$ also solve these equations for $0 \leq \tau < 1$. Plugging these solutions in (44) gives:

$$(1 - \tau)(A - i_K(\tau = 0) - i_H(\tau = 0)) = \left(\frac{(1 - \tau) \rho}{1 - \theta_K i_K(\tau = 0)} \right)^\psi g(H, \tau)^{1 - \psi}, \quad (48)$$

$$(1 - \tau)(A - i_K(\tau = 0) - i_H(\tau = 0))^{\frac{1}{1-\psi}} \left(\frac{\rho}{1 - \theta_K i_K(\tau = 0)} \right)^{\frac{-\psi}{1-\psi}} = g(H, \tau), \quad (49)$$

$$(1 - \tau)g(H, 0) = g(H, \tau). \quad (50)$$

This proves the supposition. Pindyck and Wang argue that the consumption tax is non-distortionary, because it lowers consumption by the same fraction in all periods and so does not affect households' intertemporal marginal rate of substitution. This reasoning continues to hold. The economy is not necessarily on a balanced growth path, due to the presence of the flood protection stock H , but its costs are scaled so that its optimal level is independent of K and hence of the consumption tax.