Reinforcement learning-based adaptive strategies for climate change adaptation: An application for flood risk management

Kairui Feng[a], Ning Lin[a,*], Robert E. Kopp[b,c], Siyuan Xian[a], Michael Oppenheimer[d,e,f]

[a] Department of Civil and Environmental Engineering, Princeton University, Princeton, NJ, USA.

[b] Department of Earth and Planetary Sciences, Rutgers University, New Brunswick, NJ, USA.

[c] Rutgers Climate and Energy Institute, Rutgers University, New Brunswick, NJ, USA.

[d] School of Public and International Affairs, Princeton University, Princeton, NJ, USA.

[e] Department of Geosciences, Princeton University, Princeton, NJ, USA

[f] High Meadows Environmental Institute, Princeton University, Princeton, NJ, USA

*Ning Lin **Email:** nlin@princeton.edu

Author Contributions: K.F., N.L., and S.X. designed the study. R.E.K. and M.O. contributed valuable insights and expertise in the field of climate change modeling and practical policy implications. K.F. performed the research and generated the results. S.X. prepared building data and vulnerability functions. All authors analyzed the results and drafted the manuscript.

Competing Interest Statement: The authors declare no competing interests.

Keywords: Climate Change Adaptation; Reinforcement Learning; Coastal Flood Protection; Climate Uncertainty; Adaptive Policy Design

Abstract

Climate change is posing unprecedented challenges, necessitating the development of effective climate adaptation. Conventional computational models of climate adaptation frameworks inadequately account for our capacity to learn, update, and enhance decisions as exogenous information is collected. Here we investigate the potential of reinforcement learning (RL), a machine learning technique that exhibits efficacy in acquiring knowledge from the environment and systematically optimizing dynamic decisions, to model and inform adaptive climate decision-making. To illustrate, we derive adaptive stratigies for coastal flood protections for Manhattan, New York City, considering continouse observations of sea-level rise throughout the 21st century. We find that, when designing adaptive seawalls to protect Manhattan, the RL-derived strategy leads to a significant reduction in the expected cost, 6% to 36% under the moderate emissions scenario SSP2-4.5 (9% to 77% under the high emissions scenario SSP5-8.5), compared to previous methods. When considering multiple adaptive policies (buyout, accommodate, and dike), the RL approach leads to a further 5% (15%) reduction in cost, showcasing RL's flexibility in addressing complex policy design problems when multiple policies interact. RL also outperforms conventional methods in controlling tail risk (i.e., low probability, high impacts) and avoiding losses induced by misinformation (e.g., biased sea-level projections), demonstrating the importance of systematic learning and updating in addressing extremes and uncertainties related to climate adaptation. The analysis also reveals that, given the large uncertainty and potential

misjudgment about climate projection, “preparing for the worst” is economically more beneficial when adaptive strategies, such as those supported by the RL approach, are applied.

Significance Statement

Traditional risk mitigation frameworks are inadequate for the problems posed by a changing climate, given the substantial uncertainty in climate projections. This research highlights the potential of reinforcement learning (RL) as a powerful approach for modeling adaptive climate decision-making. By focusing on coastal flood protection strategies for Manhattan, New York City, the study demonstrates that the RL-based design can lead to substantial cost reductions compared to conventional methods. Furthermore, this study shows RL's high ability to handle complex policy designs, extreme loses, and potential expert misjudgment and, more generally, the critical role of systematic learning and updating in climate change adaptation. In addition, the analysis reveals that, “preparing for the worst” is economically more beneficial when adaptive strategies are applied.

Main Text

Introduction

The world's climate is changing, and depending on future emissions, it may continue to change at unprecedented rates in recent human history. Planners face the daunting task of developing policies and making investment decisions for climate change adaptation in an environment that consists of complex, interlinked systems with manifold uncertainties.

As the future unfolds, planners are expected to learn and respond to the new situation by adapting their plans to the new reality (1). Such flexible adaptations, if strategically planned, offer advantages over pre-planned strategies by addressing dynamic risks and uncertainties. Firstly, as long as flexibility is feasible, option holders can plan to invest in stages, facilitating an initial action at a relatively low cost. Secondly, flexible adaptation options enable adjusting actions or plans in response to unexpected future states, preventing catastrophic failures. Thirdly, option holders can take possible future actions into account in current decision-making to avoid overestimating the lifetime risk to be addressed.

Flexible adaptation frameworks are referred to by diverse terms such as an “adaptation pathways” (2,3), “dynamic adaptation” (4,5), or a “real options analysis” (6,7). Several analytical approaches have been developed to implement these policy frameworks in models of adaptive climate decision-making (8,9,10). These approaches can be applied to achieve cost-benefit optimal solutions, but they have not fully addressed the potential of flexible adaptation. Table 1 presents a compilation of these quantitative methods applied to environmental policy design. The methods are categorized according to their ability to (a) design dynamic policy, (b) incorporate observational data, and (c) systematically take future observations and strategy adjustments into current decision-making.

The first ability allows the decision-maker to design a dynamic path for decisions over time. For example, ref. 11 used the dynamic programming (DP) method, a classical sequential decision-making framework, to estimate the optimal path of seawall height based on the current projection of the future climate. Ref. 12 used heuristic algorithms, e.g., genetic algorithms that stochastically generate thousands of potential paths of the seawall height and strategically select better paths, for the multi-target design of coastal seawalls. These heuristic algorithms improve DP’s ability to handle the curse of dimensionality when the decision-making process involves many steps (12). These methods, however, assume a static base of information and do not directly address a key advantage of flexible policy design: the capacity to learn and thus update and improve decisions as exogenous information is collected.

The second ability allows the decision maker to design a dynamic path for decisions that could change with observations. Modeling updating processes becomes especially important as sometimes new information leads to scientific beliefs that diverge over time from the posteriori, which is known as negative learning (13). Bayesian dynamic programming (BDP) methods (14, 15), the early attempt to incorporate observations and update in flexible policy design, employ the DP model and new observations/projections when they become available to estimate the optimal path forward. Although BDP can incorporate observations and learnings into decision-making, it does not account for future learning and updating in the current decision-making and thus may overestimate the lifetime risk to be addressed. In other words, the potential for future adjustments of decisions can affect the optimality of current decisions. Decision trees or real options methods, which generate flexible plans by searching over scenario trees, can overcome this limitation (16, 17). However, the real options approach involves an event tree with scenarios exponentially increasing with the total time step of the policy pathway. Real options analysis is tractable only when the number of potential solutions and scenarios is limited. To lower the computational cost, direct policy search (DPS) approaches have been developed (18,19). These approaches model the decision at each time step as a simple function of the observation at that time step, with the parameters of the function optimized through simulations. Consequently, the intricate stochastic sequential decision problem is approximated as a parameter optimization problem. Despite their computational efficiency, these approximate approaches may still fall short of achieving true optimality in adaptive climate decision-making.

Reinforcement learning (RL) is an area of machine learning concerned with how agents ought to take action in changing environmental states to maximize their cumulative rewards (20). RL approaches systematically incorporate observations and account for future outcomes and reactions, and they support policy designs over a continuous range of future environmental states. Also, various approximations (e.g., in characterizing states and rewards) can be made in RL to achieve numerical efficiency. RL has achieved significant success in various fields, including chess playing (21,22), autonomous driving (23), and robotics control (24). RL has also been employed in addressing sequential environmental decisions with large decision spaces, e.g., power storage (25) and water management (26). However, it has not yet been used to address the large uncertainty in climate risk. Here we investigate the potential and optimality of the RL method applied to adaptive climate decision-making. More broadly, we examine the value of systematic learning and updating in climate adaptation.

As an example, we apply the RL method to model adaptive strategies to address coastal flood risk. Planned adaptation strategies to mitigate coastal flood risk include building protective structures such as seawalls, retrofitting structures (encouraged, for example, through incentives and insurance regulations), and relocation away from harm (through “retreat”, “withdrawal” or “buyout”) as discussed in ref. 27. Tropical cyclones (TCs) may lead to higher storm surge under climate change (27-31). Sea level rise (SLR) has been and will continue to be a major factor in coastal flooding. However, future SLR projections are characterized by large and deep uncertainties that currently impede the modeling of optimal risk mitigation strategies (27,32-36). Here we develop RL methods to calculate optimal coastal risk mitigation strategies (including adaptive seawall and combined strategies involving withdrawing, retrofitting, and dike) for Manhattan in New York City (NYC) that incorporate continuous SLR observations over the 21st century (Methods). The RL method efficiently handles the computational cost, which would grow exponentially as the number of SLR scenarios and temporal resolution of decision updates in traditional algorithms increase, through state and reward approximation methods (Methods; 37).

We focus on coastal flood risk management to evaluate the effectiveness of RL, in the wider realm of optimization frameworks for climate adaptation strategies. In comparison with some of the above-mentioned frameworks (DP, BDP, and DPS), our analysis shows the superior performance of the RL method in deriving flexible strategies that minimize cost and tail risk. We also found that the RL framework shows the highest ability to attain the best economic reward when the climate projection is biased. The results highlight the importance of continuous learning and systematic adaptation in combating the large uncertainty in climate projections and the potential of the RL method for modeling optimal climate adaptation strategies.

Coastal Protection for Manhattan, NYC

Superstorm Sandy in 2012 caused devastating flood damage to both New Jersey (NJ) and NYC, demonstrating the high vulnerability of the areas to storm surge flooding (38,39,40). In response, The BIG U (Fig.1a) was planned as a protective system around the low-lying topography of Manhattan. The United States Department of Housing and Urban Development (HUD) has dedicated a total of $511 million, including Rebuild by Design and National Disaster Resilience Competition funding, toward the implementation of the "BIG U" (41). A potential gateway across the New York Harbor to protect the greater New York area for an estimated cost of $119 billion and taking 25 years to build, was considered by the US Army Corps of Engineers (USACE,42). This plan, however, has been abandoned. Without the gateway, the USACE proposed two other local protection plans with multiple strategies, including building spatially varying dikes and retreat of certain communities around the NYC area. The latter USACE plans are currently under review. To develop analytical methods, we focus on the protection for lower Manhattan and make comparisons between various analytical designs with The BIG U design.

Studies seeking an analytical solution for coastal protection typically suggest that the height of levees be adapted to certain flood return levels (43-45). This approach is followed by the Netherlands, a country where a substantial portion of the land is situated below sea level (46), and also by the BIG U design for NYC (41; Fig.1a). However, this approach, in the quest for a simplified closed-form optimal solution for coastal protection, predetermines the return period of floods to design for. Also, because of the heritage of static infrastructure design, most designs or studies consider a static coastal protection strategy, such as the BIG U (New York) or Harbor Barrier (NYC-NJ area; 47). However, many coastal cities, including NYC and Shanghai (48), had to reactively elevate seawalls following significant storm surge/flood events.

Ref. 11 brought forward a viewpoint that dynamic design can improve the economic performance of the seawall and the evolving seawall height can be estimated by DP. However, ref. 11 derived a deterministic design of seawall height over the whole life cycle at the beginning of the project, and the study did not illustrate the advantage of using future observations to update the decisions. Ref. 15 quantified the value of information updating by applying DP to obtain a new seawall design given the climate projection based on new observations (BDP). Though the design in ref. 15 involved current observations in the optimization framework, it did not involve the advantage of considering potential future observations and updates. An adaptation pathway approach, partially considering this flexibility, usually enumerates decisions including seawall height over limited climate scenarios (2,3). Decision-tree-based methods or real options approaches have also been employed to reach an optimal life cycle cost in seawall design with few (<3) time steps and limited scenarios (7). On the other hand, ref. 18 employed a DPS approach to allow seawall decisions to flexibly change with sea level observations. The DPS approach is computationally efficient, but it may not achieve true optimality in the seawall design. Here we apply the RL framework to design an adaptive seawall for lower Manhattan in NYC (Fig. 1a) that will be raised over time (e.g., every ten years over the 21st century) in response to SLR observation and projection update. RL enables a systematic consideration of climate observations and decision updates to achieve the lowest lifecycle cost compared to previous methods. For benchmarks, we compare the estimated economic rewards and tail risks of the RL strategies with those of the Big U flood protection plan and other adaptation frameworks, including DP, BDP, and DPS.

Furthermore, we consider the application of RL in modeling broader strategies for coastal flood risk management. Specifically, we apply the RL framework to model an integrated strategy combining policies of regional protection (e.g., flood walls, levees, dikes), building-level retrofit/accommodation (e.g., elevating or waterproofing buildings), and coastal retreat/withdrawal (e.g., buyout by the government). These are three typical coastal flood protection measures considered in the new USACE proposal (42) and IPCC typical coastal protection types (27). In this integrated strategy, we assume the retreat zone is always lower in altitude than the accommodate zone, and the dike is built behind the accommodate zone (Fig. 1b), following ref. 49, which found that retrofitting behind seawalls is generally not cost-beneficial. Refs. 49, 50, and 51 discussed the performance (economic outcomes) of such combined defensive strategies under several potential policy pathways; here we apply RL to find the adaptive implementation time paths of these defensive strategies. The multidimensional policy suggestions given by the RL approach are compared with the one-dimensional seawall designs for lower Manhattan.

Experimental Settings

Here we employ two cases to illustrate the RL optimization framework. In Case I we design an adaptive seawall; in Case II we design three adjustable defensive strategies, including withdrawal, retrofitting, and dike. For both Case I and Case II, we minimize the net present value (NPV; 3% discount rate) of expected policy implication cost and damage. We define the decision space over a finite time horizon (starting in 2000 and ending in 2100), where choices are made sequentially in discrete time periods (every 10 years). We posit that the strategies may be adjusted in response to observed SLR every decade and updated future SLR projections. In our illustrative experiments, for every 10-year time period, we adjust the distribution of projected future SLR based on the current and past SLR observation. Specifically, for each sampled SLR observation at each time point, we assess the similarity between all SLR realizations (~80,000) and this observation using the root mean square distance. Based on the similarity, we use a discrete choice model to determine the likelihood/weight of each SLR realization to obtain an updated distribution of future SLR. The adaptation strategies at the current time point are then determined based on the adjusted/updated future SLR projection. Through performing a large number of experiments, we obtain samples of adaptation designs and access their statistics.

Coastal flooding is induced by storm tide (storm surge plus the astronomical tide) on top of the mean sea level. Storm tide risk is projected to increase for NYC and the many other areas along the U.S. East and Gulf coastline, driven by projections of increased risk associated with storm surge (27-31) and SLR (32-34) over the 21st century under climate change. In our analysis, the annual flood hazard distribution is estimated by combining the distributions of annual maximum storm tide (52) and SLR (33). Both the storm tide and SLR distributions were generated based on CMIP6. To illustrate the performance of the RL method in managing climate extremes, we consider the moderate emissions scenario, which is Shared Socioeconomic Pathway(SSP)2 4.5, which aligns with the emissions forecasts under current climate policies, as well as the very high emissions scenario, SSP5 8.5 (53). Later we design an experiment to assess the robustness of the proposed strategies when the climate projection is biased (e.g., decisions are made under the projections for SSP2 4.5, but the climate change actually experienced is closer to that projected for SSP5 8.5 due to high climate sensitivity and strong carbon cycle feedbacks). For each year over the planning horizon, flood levels are sampled from the annual flood hazard distribution. A static inundation analysis is performed to estimate the inundation and total damage for each flood level (see Methods). If the seawall is higher than the flood height, no damage is assumed for the area protected by the seawall. If a property is inundated, the damage/loss is estimated by the vulnerability function developed by the Federal Emergency Management Agency (FEMA, 54). The annual damage/loss distribution is thus obtained and used in the optimization analysis.

We first employ a one-dimensional seawall design problem to illustrate the optimization framework (Case I). We assume the presence of a seawall around the BIG U area (Fig. 1a). We employ the RL technology to obtain the optimal time path of seawall height over the 21st century in response to each realization of the SLR observation. To illustrate the advantage of the RL-based design, we compare it with the BIG U's current plan (based on the flood return level) and with static optimal (SO; based on cost-benefit analysis), DP, BDP, and DPS strategies in terms of their total expected cost, risk of extreme losses, and costs/losses when the SLR and TC climatology projections are biased (e.g., due to uncertainties in emissions and in the climate and sea-level response to the emissions).

For the same region, we then apply RL to derive the optimal multi-dimensional design including the three adjustable defensive strategies: i) withdrawal from at-risk areas, ii) improving resistance to damage, and iii) construction of a dike (Case II, Fig. 1b). All three strategies are adaptive over time. We assume the decision maker first plans the ground elevation of the dike, as once the foundation of the dike is fixed, it cannot be relocated. They then gradually buy out the properties in the retreat zone, retrofit the buildings in the accommodate zone, and building up the dike. Different from direct buyout/retrofitting all the properties at once gradually applying these plans could both save the time value of the investment and buyout/retrofit those properties at the most cost-efficient time point given evolving climate conditions. The multidimensional policy suggestions given by the RL approach are analyzed and compared with the one-dimensional seawall designs. These arrangements for the retreat and accommodation zones may not be the only possibilities. Here we focus on mathematical optimizations under the single economic target. The RL method may be extended in the future to consider other objectives such as social inequality.

Results

Case I. Dynamic Seawall Design

In Case I, we design the seawall around lower Manhattan. Different methods suggest different seawall height time series, as shown in Fig. 2. We consider SLR in NYC relative to the sea level in the year 2000 (7.0 inches above NAVD 88). The median of SLR projection under SSP2 4.5 for the end of the 21st century is lower than under SSP5 8.5 by 1.5 ft. The storm tide return level under SSP2 4.5 is also significantly lower than that under SSP5 8.5. As a result, the seawall design under SSP2 4.5 is significantly lower than that under SSP5 8.5. As a reference, the Big U original design is shown as the green line, which is static at 16 ft. This height was determined based on the 100-year flood height from the FEMA flood map and the upper 1% projected SLR in 2050 (41). Based on life-cycle cost analysis, for SSP2 4.5 (SSP5 8.5), we found the SO level (black curve) for the seawall around the Big U area to be 15.2 (19.5) ft. For DP (blue line), under SSP2 4.5, the seawall height starts at 11.3 ft, and it reaches a final level of 15.0 ft in 2070, close to the SO level. Under SSP5 8.5, the seawall height starts at 9.1 ft, which is substantially lower than SO, and it increases over time quickly to 16.5 ft by 2050 and 25.1 ft by the end of the century. RL strategy suggests a stochastic path of seawall height changing with observed SLR conditions (as well as the BDP and DPS strategies, which are not shown in the figure for clarity). The red line shows the 50% quantile seawall height suggested by RL. The red shade shows the probability density of the RL strategy, varying with the SLR observation. The median initial seawall height is 9.1 (9.1) ft, and the median final height is 15.4 (20.9) ft, both of which are similar to or lower than the DP strategy. However, under extreme SLR cases (<1%), the final seawall height determined by RL could be around 17.8 (29.1) ft. On the other hand, if the sea level remains at a lower level (<1%), the RL algorithm would suggest a final seawall height of only around 13.4 (18.3) ft.

Different seawall strategies would lead to different expected life-cycle total costs (sum of investment and damage). Under SSP2 4.5 (SSP5 8.5), the total cost for NYC would be on average 1.40 (6.45) billion dollars under the Big U design, 1.34 (2.91) billion for the SO, 1.08 (1.89) billion for the DP, 0.95 (1.59) billion for the BDP, 0.98 (1.61) billion for the DPS, and 0.89 (1.45) billion for the RL strategy. The results indicate that under the uncertain climate change projection, a slight change in the coastal protection design considering cost-benefit (from 16-ft Big U to 15.2 (19.5)-ft SO strategy) will lead to a ~5% (55%) lower total cost. Considering dynamic design can earn an additional ~20% (35%, comparing DP with SO) to ~50% (35%, comparing RL with SO) for the NYC coastal protection problem. Compared to these previous methods, the RL strategy leads to a 6% to 36% (9%-77%) reduction of the expected total costs.

RL strategies have cost advantages over other strategies because RL is designed to systematically respond to SLR observations. To illustrate the dynamic response of the RL strategy to SLR observations over time, Fig. 3 presents example trajectories of seawall height design under three SLR observations. For each case, the time history of the final-stage (2100) SLR estimation with associated uncertainty is also shown (mid-panel). Fig. 3a shows a case where the observed SLR is near the high end of the projected distribution at the beginning, and it increases with an even deeper slope towards the end of the century. The uncertainty of the final-stage SLR estimation does not narrow until late in the century. The designed path of seawall height thus increases over time (from blue to yellow). The planned final seawall height increases from 20.5 ft estimated at 2030 to the final level of 25.9 ft. Fig. 3b illustrates a highly uncertain SLR scenario, where the SLR starts on the low end of the projected distribution, becomes an extreme case by 2020, and then changes back to the median of the distribution towards the end of the century. The uncertainty of the final-stage SLR estimation remains large until the end of the century. This scenario represents a “negative learning” case (13). Accordingly, the recommended final seawall height keeps changing from 19.7 ft at the beginning, to 25.5 ft in 2050, and back to 20.5 ft in 2100. Fig. 3c shows a case in which the SLR starts at the lower end and gradually increases to the median level. The uncertainty of the final-stage SLR estimation significantly narrows from the mid-century. In this case, the planned final seawall height does not change significantly over time (from 18.2 ft to 20.7 ft).

In general, predicting SLR more accurately in its early stages, if possible, would lead to a better estimation of policy expenditures. Conversely, if SLR remains highly uncertain over a prolonged period, the implemented decisions will deviate significantly from the initial estimations, resulting in a more uncertain budget. For example, when negative learning occurs, new information leads to increasing divergence of even confident projections from the true outcome (13,55), so caution must be exercised in foreclosing options prematurely.

Case II. Multi-Dimensional Risk Management Strategy

In Case II, we discuss a combination of three adjustable defensive strategies: i) withdrawal from at-risk areas, ii) improving resistance to damage, and iii) construction of a dike. The retreat zone is always lower in altitude than the accommodate zone, and the dike is built behind the accommodate zone. Given the potential retreat zone boundary and accommodate zone boundary, the RL algorithm can be applied to search for the optimal design of these zones (i.e., the temporal evolution of the retreat and accommodation zones towards their boundaries and the temporal evolution of the dike height) and estimate the expected total cost. Thus, the RL framework is applied to all possible zone boundaries to search for the optimal design.

Searching over potential retreat and accommodation zone boundaries for SSP2 4.5 (SSP5 8.5), the analysis suggests building a dike on the 12 (17)-ft high ground and withdrawing the properties located below the 6 (8)-ft ground elevation. Fig. 4 shows how various costs change with different dike foundation elevations, given the optimal retreat zone boundary of 6 (8) ft. The total cost is separated into four parts: property damage, buyout cost (inside retreat zone), retrofit cost (inside accommodate zone), and dike construction cost. The total cost (blue curve) decreases with the dike foundation elevation until it reaches 12 (17) ft (i.e., the optimal level), and then the total cost increases with the dike foundation elevation. The dike construction cost (orange curve) decreases with the dike foundation elevation. The buyout cost (yellow curve) increases initially with the dike foundation elevation. When the dike foundation elevation is higher than 6 (8) ft, the buyout cost is constant as the buyout zone is limited to regions with ground elevation lower than 6 (8) ft. If the dike foundation height is lower than 6 (8) ft, then no accommodate zone will be designated. The retrofit cost (purple curve) increases consistently when the dike foundation elevation is larger than 6 (8) ft since all the properties between the buyout zone and the dike foundation elevation should be retrofitted. The damage (green curve) is not sensitive to the dike foundation elevation when the dike foundation elevation is lower than 6 (8) ft. However, when the dike foundation elevation is larger than 6 (8) ft, the damage decreases due to the retrofit of properties inside the accommodate zone. The damage slightly increases when the dike foundation elevation is larger than 12 (17) ft. This increase occurs because it would not be cost-beneficial for properties above 12 (17) ft to be retrofitted, and when those properties are not protected by a dike, the potential damage increases. Under the optimal strategies combination for Case II, the expected life-cycle cost for the flood management project is 0.85 (1.24) billion, which saves 5% (15%) compared to the optimal value we obtained in Case I (0.89 (1.45) billion).

Given the optimal withdrawal and accommodate zone boundaries, the RL strategy suggests gradual relocation/retrofit of properties in retreat/accommodation zones (Fig. 5). Under SSP 2 4.5, the withdrawal mainly happens between 2040-2060 (magenta curves). The development of the accommodation zone starts in the first decade; however, it could last over the century (blue curves). The dike has two significant elevation time points, one at 2030 and the other at 2070. The median final dike height is 17 ft above the mean sea level, but 5 ft above the dike foundation with a lower 5% level of 3 ft and an upper 5% level of 7 ft. Under SSP5 8.5, during the first 20 years, those properties under the 5-ft ground elevation will be relocated for most SLR scenarios, and the majority of the properties in the retreat zone will be relocated before 2090 under any SLR condition considered (magenta curves). Also, the majority of the properties inside the accommodation zone should be retrofitted before 2090 (blue curves). In most cases, the results do not suggest building the dike in the first half of the 21st century. In other words, our findings indicate that initiating dike construction before 2050 is not imperative. This result provides significant flexibility in policy implementation, especially since the dike construction represents the sole sunk cost across the three types of policies. The median final dike height is 22 ft above the mean sea level, but 5 ft above the dike foundation with a lower 5% level of 3 ft and an upper 5% level of 12 ft.

Compared to Case I, where the anticipated final dike height stands at 15.4 (20.9) ft above the foundation, the constructed dike height for Case II is substantially lower. Unlike in Case I, the dike in Case II, when combined with withdrawal and resistance strategies, is not intended to safeguard properties located in low-lying areas. Compared to the SSP5 8.5 scenario, the suggested retreat and accommodation zones are narrower, and the dike height is lower under the SSP2 4.5 scenario, given its lower storm surge and SLR projection. However, the RL result suggests starting to build the dike earlier, in 2030, as earlier protection is needed when much of the coastal areas is not retreated or retrofitted.

Similar to Fig. 3, Fig. 6 shows how the designed zones and dike heights change over time with the three illustrative SLR scenarios. In general, the results closely resemble those of Case I. Under the first case (Fig. 6a), where the sea level rises sharply with increasing uncertainty, the projected final dike level increases quickly, from 22 ft to 27 ft. The retreat and accommodation zones are projected to be fully developed earlier as time goes on (zone development curves shifting to the left). However, when there is increased uncertainty in SLR projections, the plans may exhibit more substantial changes than in Case I. Notably, in Fig. 6b, where negative learning occurs, the accommodation and retreat zone are projected to develop rapidly early on, but later the projection scales back (i.e., the zones do not need to be fully developed until later in the century). The projected final dike level in the early decades is very high (10 ft above the foundation), while the final construction is much lower (5 ft). In Fig. 6c, where the SLR gradually increases with narrowing uncertainty, the three protection measures develop slowly over time and the final results are close to the early projection.

Tail Risk

To investigate the tail risk from different design strategies, we show the tail of the distribution of the present value of the total cost (damage and construction) of the BIG U design, SO, DP, BDP, DPS, RL, and combined multi-dimensional RL strategies under SSP2 4.5 (SSP5 8.5) in Fig. 7. First, we compare the tail risk for the strategies in Case I. At a 1% exceedance level, the total cost under the Big U seawall design is 2.7 (10.1) billion. The corresponding cost for SO and DP strategies is around 4.0 (4.8) and 2.7 (2.5) billion, for DPS and BDP strategies is around 1.3 (2.7) and 1.2 (2.5) billion, and for RL strategy is around 1.2 (2.1) billion. At the 0.1% exceedance level, total cost for the Big U/SO/DP/BDP/DPS strategies is 8.7 (1200) /15.2 (522) /9.9 (215) /1.2 (3.5) /1.1 (1.1) times, respectively, that for the RL strategy. There is still a small chance (>0.01 %) for the static, DP, or BDP strategy to bear an extreme total cost exceeding 20 billion under SSP5 8.5, but this probability for the RL strategy is negligible (no realization within one million samples). These results demonstrate the outstanding risk control ability of the RL method. Usually, a strategy with a lower expected total cost will hold a higher uncertainty (risk) in the cost. However, the RL and, to a lesser extent, BDP and DPS strategies outperform other strategies at controlling both expected cost and risk, demonstrating the benefits of "observing and updating."

Secondly, we compare the RL strategy in Case I with the multiple strategies (MS) in Case II. Compared to the RL strategy, employing multiple strategies does not guarantee a lower total cost for 1% level events; in fact, it results in a 5% higher cost for SSP2 4.5 and a 5% lower cost for SSP5 8.5. This difference is induced because the RL and MS methods adhere to the same optimization framework, where a lower expected cost may be associated with a higher degree of risk. For example, comparing to build a seawall initially to protect the low-lying community, gradually retreating residents may lead to large losses if a severe storm surge occurs at the early stage. However, the tail risk for MS is still very small, compared to other methods.

Robustness of Decisions under Misinformation and Uncertainty of Climate Projection

The uncertainty surrounding future climate outcomes arises not only from climate modeling but also from factors related to emission scenarios and unknown physics. Expert opinion and model outcomes on the dynamical response of ice sheets and its impact on sea levels is characterized by increasing uncertainty after the middle of the 21st century and ambiguity due to lack of expert consensus thereafter (27,33,56). Under extreme cases, one could envision positive feedbacks that led to a SLR response closer to that in SSP5 8.5 under median climate sensitivity and carbon cycle feedback even when the emissions are moderate. This uncertainty leads to the pressing question: how robust is the decision-making framework under RL when the selected climate projection or scenario by policymakers diverges from the actual trend?

This section evaluates the robustness of the discussed analytical approaches for climate adaptation in achieving economic gains under biased climate projections through a counterfactual experiment. Two benchmark scenarios are employed: one where the coastal protection is designed based on the a version of the SLR projection for SSP5 8.5 that incorporates high-end, low-confidence estimates of ice-sheet loss based on structured expert judgement (henceforth, SSP5 8.5 LC; 62,63) and matched with reality, and another where the coastal protections are designed based on the standard, medium-confidence projections for SSP2 4.5 (SSP2 4.5 MC) that do not incorporate deeply uncertain ice-sheet processes with the potential to drive rapid ice sheet losses (See Methods on details of SLR scenarios). Additionally, two counter experiments are conducted: one where the coastal protections are planned for SSP5 8.5 LC while reality exhibits SSP2 4.5 MC , and the other where the coastal protections are planned for SSP2 4.5 MC while reality exhibits SSP5 8.5 LC. These experiments assume that policymakers do not alter their belief in future scenarios over time. In other words, policymakers who believe in the SSP5 8.5 LC scenario, even if the observation aligns with SSP2 4.5 MC and SLR appears to be low, consider the observation a lower probability case within SSP5 8.5 LC, and vice versa.

Table 2 presents the expected costs for each of the four cases. Overall, employing the RL framework results in lower expected total loss and policy expenditure (total cost of construction, residents' withdrawal, and structure retrofitting, i.e., expected total loss minus the expected damage) in every case, and the multiple strategies framework is superior to the seawall strategy. Additionally, the difference between the expected losses of cases without and with the correct climate scenario belief for the same climate scenario could be defined as the “bias loss.” It is generally observed that implementing flexible adaptation strategies, especially RL and MS (multiple strategies with RL), reduces the bias loss. For example, with the belief that SSP2 4.5 MC (SSP5 8.5 LC) will occur, the bias loss under SO is 11.58 (1.02) billion, which is 79% (90%) larger than that under MS. Moreover, bias loss tends to be lower for plans designed under the belief that SSP5 8.5 LC will occur (while in reality SSP2 4.5 MC occurs) compared to those designed under the belief that SSP2 4.5 MC will occur (while in reality SSP5 8.5 LC occurs). For example, the bias loss for SO, DP, and RL under the belief that SSP2 4.5 MC would happen is 11.5, 17.0, and 17.0 times of that under the belief that SSP5 8.5 LC would happen.

Discussion

In this study, we analyzed the RL method, among various optimization approaches, for flood adaptive design and applied the analysis to Manhattan, NYC. The methods hold potential for broader application in various climate adaptation scenarios, provided that the life-cycle benefit serves as the optimization objective. The RL framework exhibits a versatile capability in achieving optimal decisions, particularly when the temporal evolution of climate or environment can be estimated probabilities conditioned on the current state's information.

In terms of minimizing the expected life-cycle cost, there exists a performance ranking from RL, BDP/DPS, and DP to SO methods. The cost for learning-based adaptive methods is lower than non-adaptive methods because the adaptive methods can respond to the observation and adjust strategies to control risk. Also, the initial investments suggested by adaptive methods are lower than those by non-adaptive methods. Specifically, the static strategies need to be “conservative” as they are determined to cover the large uncertainty and risk for their entire lifecycle while adaptive strategies can be more “aggressive” at the start as they can adjust themselves over time according to future observations.

As both the large uncertainty and high initial investment reduce stakeholders’ willingness to implement protective strategies proactively rather than waiting for disasters to strike, the adaptive design provides a promising approach for climate adaptation. Stakeholders may be more willing to invest in a policy or project that can respond to future scenarios, especially given the lower initial cost.

The RL approach has advantages over other adaptation decision-making methods that can consider only a limited number of pre-defined climate change scenarios. Also, the RL method could flexibly be extended to coordinate multiple risk-mitigation policies. As shown in this paper, when applying multiple types of measures at the same time, the combined flood risk mitigation strategies are expected to be more effective than a single risk mitigation strategy, especially in controlling the total cost. This observation highlights the importance of coordinating multiple policies to address diverse environmental adaptation challenges. This collective approach might hold greater significance than optimizing individual policies independently, and the RL framework can facilitate such a collective approach by synergizing multiple types of stratigies and navigating intricate problem spaces.

From a project management perspective, it is well understood through the lens of the Capital Asset Pricing Model (CAPM; 59) that expected returns typically rise in tandem with risk. The trade-off between risk and return is a foundational principle in finance and investment theory; this tradeoff forms a Pareto front. For example, robust decision-making methods, such as minimax regret and information gap analysis, have been applied to formulate decisions aimed at mitigating extreme potential outcomes on one end of the Pareto front (5, 60). In this context, the distinctive performance of the RL approach is particularly noteworthy. It not only minimizes economic losses, but also effectively manages the tail risk, ensuring the avoidance of extreme impacts. The efficacy of RL comes from its systematic incorporation of new observations into dynamic decision-making, outperforming the Pareto front of frameworks that employ only current information and/or static adaptation. In addition to risk and return, the RL framework can be extended to consider other objectives such as social inequality (61) by incorporating these objectives with weighting factors into the optimization.

Additionally, our findings indicate that the bias loss tends to be lower when decisions are developed assuming high impact scenarios while low impact scenarios occur in reality, compared to plans created assuming low impact scenarios while in reality high impact scenarios take place. This result indicates the significant role of extremes in contributing to the total loss. Adapting to low or moderate impact scenarios when in reality high impact scenarios occur makes it difficult to avoid impacts from extremes. Therefore, planners are advised to adopt systematically adaptive decision-making tools, such as RL, and a "prepare-for-the-worst" approach (if multiple scenarios are equally likely) when designing adaptation strategies under uncertain climate change scenarios.

Adaptation decisions in our modeling framework are a highly simplified compared to real-world adaptation decisions (56; 57; 58). In our model, we seek simply to minimize the net financial cost of coastal damages and adaptation measures. In the real world, many values are at stake in adaptation decisions, and different players have both different values and different power to make their values influence final decisions. In addition, the idealized, continually updated adaptation decisions we find minimize cost neglect the practicalities of political economy: for example, whether funding and political will are continually available to operationalize them. Future modeling efforts could start to represent potential frictions. Nonetheless, models such as ours can be valuable guides to the players involved in such complex processes, and our results highlight the potential order-of-magnitude value that could be achieved through an iterative and flexible approach to urban coastal adaptation.

Materials and Methods

Simulated storm surge events

The current (1981-2000) and future (2081-2100) annual storm tide distributions come from ref. 52 for both SSP5 8.5 and SSP2 4.5, modeled using a coupled climatological-hydrodynamic model. The simulation methods have been used in previous coastal adaptation analyses (e.g., ref. 47). The storm tide distributions are linearly interpolated to the analysis time points over the 21st century.

Projection of sea level rise

We employed sea-level projections produced by the Intergovernmental Panel on Climate Change Sixth Assessment Report (AR6; 33, 62) using the Framework for Assessing Changes To Sea-level (FACTS; 63). AR6 produces four alternative probability distributions (‘workflows’) for trajectories of future global-mean and local relative sea-level rise for each SSP. Workflow 1f employs ice-sheets calibrated to the Ice Sheet Model Intercomparison Project (64), while workflow 2f substitutes Antarctic Ice Sheet projections based on the Linear Response Model Intercomparison Project 2 (LARMIP2; 65). AR6 judged these two workflows to represent sea-level processes in which there is at least a medium level of evidence and agreement, and thus 'medium confidence.' Workflow 3f substitutes the Marine Ice Cliff Instability-representing model of ref. 66, while workflow 4 substiutes for both Antarctic and Greenland ice sheets the structured expert judgement projections of ref. 67. AR6 judged these two workflows to represent processes for which there is limited evidence and agreement, and thus 'low confidence.'

The sea-level projections take into account ocean thermal expansion and dynamics, cryosphere and land-water storage change, vertical land motion, and spatially varying responses of the geoid and the lithosphere to shrinking land ice. This approach resulted in a dataset of 20,000 individual SLR trajectories generated for each SSP and workflow from 2000 to 2100; each trajectory contains data points at ten-year intervals, providing a comprehensive collection of SLR projections. In the main analysis, we combine the four workflow projections under each SSP using equal weighting. In the robustness analysis, SSP2 4.5 MC is based on workflow 1f, while SSP5 8.5 is based on workflow 4.

Building-level information

Data for the 43,000 buildings in Manhattan, NYC, have been processed from the MapPLUTO database of the NYC Department of City Planning. This database contains various information about each building: the number of stories, the building type, the year of construction, the year of renovation, the building's assessed value, and the square footage. We consider future property development based on statistical projection (logistic regression of total property volume on the year of construction and unified by House Price Index trends).

The LiDAR digital elevation model (DEM) data for all of NYC, with a resolution of 1 foot, has been obtained from the Department of Environmental Protection. To estimate damage from each flood level on the flood distribution, we apply DEM to estimate the inundation from each flood level. Then the property damage is estimated based on the vulnerability function, which maps the percentage of property loss to inundation depth given the building type.

The cost of mitigation measures considered in this study includes flood wall construction and building retrofit (including elevation and making lower floors waterproof). Here we consider the two main sources for flood wall construction; the unit construction cost of the flood wall used is $ 2.2 million per mile length per foot height, and each elevation has a fixed cost of $1.6 million per mile length (47). The elevation and waterproofing costs for buildings differ for structure types and are obtained from FEMA (54).

Total Loss under Case I

For Case I, we minimize the expected total cost for the Manhattan seawall project under Case I. The objective function for seawall height management, which is defined as the expected life-cycle cost for T years with a discounting rate r, can be separated into two parts: expected damage for the protected area and the construction cost. Considering that the seawall cannot be upgraded at arbitrary time points in reality, here we assume that we upgrade the seawall height (A_t) at the end of every δ years and solve this optimization problem at discrete times. Here we assume that T is divisible by δ and $\vec{A}$ (a vector of seawall height in the time sequence; regarding the initial seawall height as A_1 and $A_0 = 0$ ft) is a k dimension vector ($k = T/\delta$). Here we assume that the expected damage for the considered area for a specific time (t) is related to current seawall height $\mathrm{A}_{\left\lfloor \frac{t}{\delta} \right\rfloor + 1}$. The cost of construction is a function of both $\mathrm{A}_{\left\lfloor \frac{t}{\delta} \right\rfloor + 1}$ and $\mathrm{A}_{\left\lfloor \frac{t}{\delta} \right\rfloor}$. Under these settings, the objective function (NPV of the life-cycle cost) is:

$$L(\vec{A}_t) = \int_0^T \left[D(\vec{A}, s_t, p_t) + C(\vec{A})\right] e^{-rt} \approx \sum_{i=1}^{k} \sum_{m=1}^{\delta} [D(A_i, s_{i\delta+m}, p_{i\delta+m}) + C(A_i, A_{i-1}, m)] e^{-r(i\delta+m)} \qquad (1)$$

with the constraint that seawall height should never decrease ($A_{i+1} \geq A_i, \forall i \geq 0$).

The construction cost is calculated as a linear function of the seawall increment and happens only at the beginning of each period (increment should be larger than 1.0 ft per time under Case I, no minimum increment under Case II).

$$C(A_i, A_{i-1}, m) = \begin{cases} \int_{A_{i-1}}^{A_i} (C_b + C_i l(A_i)) l(h) dh, if\ m = 1 \\ 0, else \end{cases} \quad (2)$$

where C_b is the unit price for the seawall ($/mile/ft) and *l(h)* is the length of the seawall that needs to be constructed. *l(h)* is the length of the coastline below the elevation level *h*. The fixed cost (staging, foundation, etc.) is linear with the total length of construction with a unit price C_i ($/mile). The expected damage ($D(\vec{A}, s_t, p_t)$) could be calculated given SLR distribution (s_t) and annual storm tide distribution (p_t) for time point *t*. The *i* th increment of the seawall covers the city for δ years, and for a specific year *y* in that period, the damage function can be calculated as:

$$D(A_t, s_y, p_y) = \int_{A_t}^{+\infty} e(t) d(x) \left[\int_{-\infty}^{+\infty} P_t(x - s_y) f_t(s_y) ds_y\right] dx \qquad (3)$$

in which x is the potential flood height; $e(t)$ is the development projection (future exposure divided by the initial exposure of the city); $d(x)$ is the damage estimation for the protected urban area under a given flood height x (based on the local building distribution and fragility and digital elevation model); P_t is the CDF of projected annual storm tide distribution; and $f_t(s_y)$ is the distribution of SLR at year *y*. Here the convolution of storm tide and SLR distributions is applied to calculate the flood height distribution (68).

Under these settings, we build the framework for DP, BDP, DPS, and RL to solve the optimization problem of adaptive seawall height.

Benchmark Algorithm: Dynamic Programming

The DP method performs the optimization for the objective function (Eq. 1) with projected future SLR and storm tide distributions. Because the seawall built before one specific time point will not affect future action while the possible future action may affect the decision on current seawall height, the problem can be solved backwardly (11). This method successfully converts a multi-dimensional sequential decision-making problem into numerous decoupled one-dimensional problems, and a one-dimensional search strategy can be used to solve each of the one-dimensional problems.

Specifically, for each time period (i.e., for each one-dimensional problem), we calculate the best action (A_t; assuming $A_{t-1} = 0$) and lowest value (V_t) such that

$$A_t = argmin_{A_t} \sum_{m=1}^{\delta} \left(D(A_t, s_{t\delta+m}, p_{t\delta+m}) + C(A_t, A_{t-1})\right) e^{-r(t\delta+m)} + V_{t+1}(A_t) e^{-r(t+1)\delta} \quad (4)$$

and

$$V_t(A_{t-1}) = min_{A_t} \sum_{m=1}^{\delta} \left(D(A_t, s_{t\delta+m}, p_{t\delta+m}) + C(A_t, A_{t-1})\right) e^{-r(t\delta+m)} + V_{t+1}(A_t) e^{-r(t+1)\delta} \quad (5)$$

where V_t is the expected total damage and construction cost from time t to the end of the planning horizon under the optimal strategy for the given climate projection.

Benchmark Algorithm: Bayesian Dynamic Programming

The BDP approach follows ref. 15. Whenever we observe a sea level at a specific time y_0, we update our projection of SLR. Moreover, based on the updated information, DP analysis is applied to the time window from time y_0 to the end of the life cycle.
To illustrate the effect of information updates, we apply the conditional probability of future SLR given current conditions. Accordingly, the damage function $D(A_t, s_y, p_y)$ is changed to a

conditional form. Assuming we have observed the sea level time series $s_{1:y_0}$ before year y_0, for any year y later than y_0, we update Eq. (3) to

$$D\left(A_t, s_y | s_{1:y_0}, p_y\right) = \int_{A_t}^{+\infty} e(t) d(x) \left[\int_{-\infty}^{+\infty} P_t\left(x - s_y\right) f\left(s_y \middle| s_{1:y_0}\right) ds_y\right] dx \quad (6)$$

Here we use a neighborhood-based sampling algorithm (Ruiz and Lorenzo 2002) to estimate the conditional density function of future SLR ($f\left(s_y \middle| s_{1:y_0}\right)$) given observed sea level time series $s_{1:y_0}$. The seawall decision is updated by simple adjustment of Eq.4:

$$A_t = argmin_{A_t} \sum_{m=1}^{\delta} \left(D(A_t, s_{t\delta+m} | s_{1:t}, p_{t\delta+m}) + C(A_t, A_{t-1})\right) e^{-r(t\delta+m)} + V_{t+1}(A_t) e^{-r(t+1)\delta} \quad (7)$$

Benchmark Algorithm: Direct Policy Search

The DPS approach is adopted from ref. 18. Whenever we observe a sea level (s_t) at a specific time t, we update the seawall height with the following equation:

$$A_t = \beta_0 + \beta_1 \cdot t + \beta_2 \cdot s_t + \beta_3 \cdot t^2 + \beta_4 \cdot s_t^2 + \beta_5 \cdot s_t \cdot t \ (8)$$

where β_i are parameters that determine our decision. The optimal parameters are determined by simulations for each climate scenario. For a given set of parameters, the expected total loss is calculated by Eq. 1. Then, we enumerate over potential sets of β_i, until the expected total loss comes to a local minimum.

Proposed Algorithm: Reinforcement Learning

The RL framework considers potential future observations and updates, by changing Eq. 7 to

$$A_t(s_{1:t}) = argmin_{A_t} \sum_{m=1}^{\delta} \left(D(A_t, s_{t\delta+m} | s_{1:t}, p_{t\delta+m}) + C(A_t, A_{t-1})\right) e^{-r(t\delta+m)} + \mathbb{E}V_{t+1}\left(A_t(s_{1:t})\right) e^{-r(t+1)\delta} (9)$$

where the last term is the expectation of total damage and construction cost from time *t* to the end of the planning horizon under future strategy updates in response to possible future climate conditions, and it can be written as:

$$\mathbb{E}V_{t+1}\left(A_t(s_{1:t})\right) = \int_{-\infty}^{+\infty} V_{t+1}(A_{t+1}(s_{1:t}, s_{t+1}), A_t) f(s_{t+1} | s_{1:t}) ds_{t+1} \ (10)$$

By definition (Eq. 5), we could also rewrite Eq. 10 as:

$$\mathbb{E}V_{t+1}\left(A_t(s_{1:t}, s_{t+1})\right) = \int_{-\infty}^{+\infty} \left(\sum_{m=1}^{\delta} \left(D(A_t, s_{t+1}, p_{t\delta+m}) + C(A_t, A_{t-1})\right) e^{-r(t\delta+m)} + \mathbb{E}V_{t+2}(A_{t+1}(s_{1:t+1})) e^{-r(t+1)\delta}\right) f(s_{t+1} | s_{1:t}) ds_{t+1} \quad (11)$$

The expected reward at the current step depends only on the next step's estimation of the value function (i.e., reward approximation).

The complexity of designing RL algorithms is much larger than BDP since to obtain the current optimal seawall height, one needs to enumerate all the potential future SLR realizations. First, an unbiased design requires a large number of SLR scenarios (e.g., ~over 80,000 scenarios used in this study to sufficiently cover the large uncertainty space). (The traditional decision trees-driven framework, considering 10 time steps, would lead to a total of 80,000 to the power of 10 branches of decision trees.) The second challenge comes from estimating the reward (life-cycle benefit minus cost) of each policy decision, where the computational burden grows exponentially as the time resolution of the decision-making process becomes finer (e.g., computational time increases by 30 times if the analysis time resolution changes from every 20 years to every 10 years). In a conventional decision tree framework, it is necessary to compute all decisions and cumulative costs for every time step within each scenario. As the time resolution of the SLR process becomes finer, an exponentially growing computational burden results.

Here we apply backward approximate dynamic programming (BADP), a typical RL algorithm, to overcome the computational challenges. To implement BADP, we first simulate a large number of sea level realizations (~ 80,000). The main target of BADP is to design a look-up table, which is the $A_t(s_{1:t})$ for all the realizations simulated, to tell the policymaker what to do when they observe new information. For every possible observation of SLR, there is an optimal action given in the table produced by BADP. Note that for the final year, the seawall height could easily be determined by historical sea level records as there would be no future update. Thus, for each sea level realization group (80,000 groups; state approximation) over the time horizon, $s_{1:k}$ (where k is the final decision time point) we can determine the final stage seawall height based on the update distribution of SLR from time point k to k+1 and record that in the look-up table. Then we could use Eq. 8 to find the optimal seawall height A_{k-1} for each period, 1:k-1, backwardly, based on the updated distribution of SLR for each time period. By solving the problem backward and applying the reward approximation method under the Bellman optimality condition (allowing condensing the state tree into a single step), the optimal seawall height at any given time step could be obtained. The memory needed to store all the potential solutions under real SLR possibilities can be infinite since the SLR possibilities grow exponentially as time goes on. Here, we approximate the look-up table by restricting it to the realizations (~80,000) we simulated.

The computational cost of RL is higher than for the other methods. It takes approximately 20 minutes to run the RL algorithm for a single region under one climate scenario, whereas BDP and DPS methods take only a few minutes, and DP takes several seconds. Furthermore, the complexity of the current solutions generated by the RL framework poses challenges for interpretation by policymakers because it produces an extensive array of rule sets. Future research aims to address these issues by implementing information distillation to substantially reduce the number of rules and improve the computational efficiency of the RL framework.

Total Loss under Case II

Under Case II, the objective can be separated into four parts: expected damage, the dike construction cost, relocation cost in the retreat zone and retrofit cost in the accommodation zone. The total loss could be written as:

$$L(\vec{A}_t) = \int_0^T \left[D(\vec{A}, s_t, p_t, \vec{w}, \vec{rs}) + C(\vec{A}) + C_w(\vec{w}) + C_{rs}(\vec{rs})\right]e^{-rt} \approx \sum_{i=1}^{k}\sum_{m=1}^{\delta}[D(A_i, s_{i\delta+m}, p_{i\delta+m}, w_t, rs_t) + C(A_i, A_{i-1}, m) + C_w(w_i, w_{i-1}, m) + C_{rs}(rs_i, rs_{i-1}, m)]e^{-r(i\delta+m)} \quad \textit{(12)}$$

Here $\vec{w}$ and $\vec{rs}$ are the time-changing boundaries where the local planner is implementing buyout and retrofit policy, respectively. The cost to implement retreat (withdrawal; C_w) / accommodation (resistance; C_{rs}) zone policy is determined by the retreat/accommodation zone difference between timesteps.

$$C_w(w_i, w_{i-1}, m) = \begin{cases} \int_{w_{i-1}}^{w_i} VT(h)dh, if\ m = 1 \\ 0, else \end{cases} \quad (13)$$

where $VT(h)$ is the total value of properties at a given height *h*. The planner may buy out all the properties within the newly proposed retreat zone.

$$C_{rs}(rs_i, rs_{i-1}, m) = \begin{cases} \int_{rs_{i-1}}^{rs_i} s^* \cdot VT(h)dh, if\ m = 1 \\ 0, else \end{cases} (14)$$

where s^* is the factor of average retrofit cost for the properties inside the resistance zone compared to the total value, and the planner (or property owners) may retrofit all the properties within the newly proposed retrofit zone. For those properties that could be elevated, we assume them to be elevated following the FEMA criteria (54). For those that could not be elevated, we also follow ref. 54 to consider upgrading waterproof layers for the basement and ground floor in these properties.

To construct the RL algorithm, we could separately solve the optimization problem for each of the three zones: retreat zone, accommodation zone, and dike-protected zone. The damage inside each zone is fully decoupled; thus, the strategy implemented in a zone will not impact other zones once the height of the withdrawal boundary (*wb*) and the height of the dike foundation (*df*) are determined. As a result, we could implement a similar RL framework to the optimization problem in each zone to obtain the optimal policy implementation area $\vec{w}$ and $\overrightarrow{rs}$ as in the one-dimensional case.

For example, the governing equation for the withdrawal boundary is

$$w_t(s_{1:t}) = argmin_{w_t} \sum_{m=1}^{\delta} \left(D(w_t, s_{t\delta+m} | s_{1:t}, p_{t\delta+m}) + C_w(w_t, w_{t-1}) \right) e^{-r(t\delta+m)} + \mathbb{E}V_{t+1}\left(w_t(s_{1:t}, s_{t+1})\right) e^{-r(t+1)\delta} \quad (15)$$

and the reward function V is similar to that in Eq. (5) except being extended to include the cost for implementing withdrawal and buyout decisions.

We enumerate the potential height of the withdrawal boundary (*wb*) and dike foundation (*df*) to find the strategies that make the flood management project reach the global optimal.

Data Availability

All codes and data (excluding initial data for building properties with individual building names) will be released upon publication at Zenodo.

Acknowledgments

This work is supported by the U.S. National Science Foundation (1652448 and 2103754 as part of the Megalopolitan Coastal Transformation Hub), C3.ai Digital Transformation Institute (C3.ai DTI Research Award), and Princeton HMEI-STEP Graduate Fellowship. We thank Weiliang Jin (Princeton University) and Prabhat Hegde (Dartmouth College) for discussions on the algorithm design and Klaus Keller (Dartmouth Colleage) for valuable suggestions. We thank the IPCC Sixth Assessment Report sea-level projection authors for developing and making the sea-level rise projections available, multiple funding agencies for supporting the development of the projections, and the NASA Sea Level Change Team for developing and hosting the IPCC AR6 Sea Level Projection Tool.

References

1. Keller, Klaus, Casey Helgeson, and Vivek Srikrishnan. "Climate risk management." Annual Review of Earth and Planetary Sciences 49 (2021): 95-116.
2. Ranger N, Reeder T, Lowe J (2013) Addressing 'deep' uncertainty over long-term climate in major infrastructure projects: four innovations of the Thames Estuary 2100 Project. EURO J Decis Proces 1:233–262.
3. Nicholls RJ, Hanson SE, Lowe JA, Warrick RA, Lu X, Long AJ (2014) Sea-level scenarios for evaluating coastal impacts. Wiley Interdiscip Rev Clim Chang 5:129–150

4. Hallegatte S (2009) Strategies to adapt to an uncertain climate change. Glob Environ Chang Hum Policy Dimens 19:240–247.
5. Haasnoot M, Kwakkel JH, Walker WE, Ter Maat J (2013) Dynamic adaptive policy pathways: a method for crafting robust decisions for a deeply uncertain world. Glob Environ Chang 23:485–498.
6. Neufville R (2003) Real options: dealing with uncertainty in systems planning and design. Integr Assess 4:26–34.
7. Woodward M, Kapelan Z, Gouldby B (2014) Adaptive flood risk management under climate change uncertainty using real options and optimization. Risk Anal 34:75–92
8. Rammel, C., & van den Bergh, J. C. (2003). Evolutionary policies for sustainable development: adaptive flexibility and risk minimising. Ecological economics, 47(2-3), 121-133.
9. Benson, M. H., & Garmestani, A. S. (2011). Embracing panarchy, building resilience and integrating adaptive management through a rebirth of the National Environmental Policy Act. Journal of environmental management, 92(5), 1420-1427.
10. Huntjens, P., Pahl-Wostl, C., Rihoux, B., Schlüter, M., Flachner, Z., Neto, S., ... & Nabide Kiti, I. (2011). Adaptive water management and policy learning in a changing climate: a formal comparative analysis of eight water management regimes in Europe, Africa and Asia. Environmental Policy and Governance, 21(3), 145-163.
11. Lickley, M. J., Lin, N., and Jacoby, H. D. (2014). "Analysis of coastal protection under rising flood risk." *Climate Risk Management*, 6, 18 – 26.
12. Keller, K., Robinson, A., Bradford, D. F., & Oppenheimer, M. (2007). The regrets of procrastination in climate policy. *Environmental Research Letters*, *2*(2), 024004.
13. Oppenheimer, M., O'Neill, B. C., & Webster, M. (2008). Negative learning. *Climatic Change*, *89*, 155-172.
14. de Bruin, Kelly, R. B. Dellink, Arjan Ruijs, L. Bolwidt, Arwin van Buuren, Jaap Graveland, R. S. De Groot et al. "Adapting to climate change in The Netherlands: an inventory of climate adaptation options and ranking of alternatives." *Climatic change* 95 (2009): 23-45.
15. van der Pol, Thomas D., Ekko C. van Ierland, and Silke Gabbert. "Economic analysis of adaptive strategies for flood risk management under climate change." *Mitigation and adaptation strategies for global change* 22, no. 2 (2017): 267-285.
16. Anda, J., Golub, A., & Strukova, E. (2009). Economics of climate change under uncertainty: Benefits of flexibility. *Energy Policy*, *37*(4), 1345-1355.
17. Zhang, M., Zhou, D. and Zhou, P., 2014. A real option model for renewable energy policy evaluation with application to solar PV power generation in China. *Renewable and Sustainable Energy Reviews*, *40*, 944-955.
18. Garner, Gregory G., and Klaus Keller. "Using direct policy search to identify robust strategies in adapting to uncertain sea-level rise and storm surge." Environmental Modelling & Software 107 (2018): 96-104.
19. Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., Reed, P.M., Asce, A.M., 2016. Curses, tradeoffs, and scalable management: advancing evolutionary multiobjective direct policy search to improve water reservoir operations. J. Water Resour. Plann. Manag. 142. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000570.
20. Wiering, Marco A., and Martijn Van Otterlo. "Reinforcement learning." *Adaptation, learning, and optimization* 12, no. 3 (2012): 729.
21. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M. and Dieleman, S., 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, *529*(7587), 484-489.
22. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A. and Chen, Y., 2017. Mastering the game of go without human knowledge. *nature*, *550*(7676), 354-359.

23. Feng, S., Sun, H., Yan, X., Zhu, H., Zou, Z., Shen, S. and Liu, H.X., 2023. Dense reinforcement learning for safety validation of autonomous vehicles. *Nature*, *615*(7953), 620-627.
24. Degrave, J., Felici, F., Buchli, J., Neunert, M., Tracey, B., Carpanese, F., Ewalds, T., Hafner, R., Abdolmaleki, A., de Las Casas, D. and Donner, C., 2022. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, *602*(7897), 414-419.
25. Xiong, Rui, Jiayi Cao, and Quanqing Yu. "Reinforcement learning-based real-time power management for hybrid energy storage system in the plug-in hybrid electric vehicle." *Applied energy* 211 (2018): 538-548.
26. Herman, Jonathan D., Julianne D. Quinn, Scott Steinschneider, Matteo Giuliani, and Sarah Fletcher. "Climate adaptation as a control problem: Review and perspectives on dynamic water resources planning under uncertainty." *Water Resources Research* 56, no. 2 (2020): e24389.
27. Oppenheimer, M., Glavovic, B., Hinkel, J., van de Wal, R., Magnan, A. K., Abd-Elgawad, A., et al. (2019). Sea Level rise and implications for low lying islands, coasts and communities. In H.-O. Pörtner, D. C. Roberts, V. Masson-Delmotte, P. Zhai, M. Tignor, E. Poloczanska, et al. (Eds.) IPCC Special Report on the Ocean and Cryosphere in a Changing Climate (Vol. 355, pp. 126–129)
28. Lin, N., Emanuel, K., Oppenheimer, M., & Vanmarcke, E. (2012). Physically based assessment of hurricane surge threat under climate change. *Nature Climate Change*, *2*(6), 462-467.
29. Marsooli, R., Lin, N., Emanuel, K., & Feng, K. (2019). Climate change exacerbates hurricane flood hazards along US Atlantic and Gulf Coasts in spatially varying patterns. *Nature communications*, *10*(1), 3785.
30. Lin, N., Kopp, R. E., Horton, B. P., & Donnelly, J. P. (2016). Hurricane Sandy's flood frequency increasing from year 1800 to 2100. *Proceedings of the National Academy of Sciences*, *113*(43), 12071-12075.
31. Gori, A., Lin, N., Xi, D., & Emanuel, K. (2022). Tropical cyclone climatology change greatly exacerbates US extreme rainfall–surge hazard. Nature Climate Change, 12(2), 171-178.

32. Kopp, R. E., Horton, R. M., Little, C. M., Mitrovica, J. X., Oppenheimer, M., Rasmussen, D. J., ... & Tebaldi, C. (2014). Probabilistic 21st and 22nd century sea-level projections at a global network of tide-gauge sites. *Earth's future*, *2*(8), 383-406.
33. Fox-Kemper, B. et al. in Climate Change 2021: The Physical Science Basis (eds Masson-Delmotte, V. et al.) Ch. 9 (IPCC, Cambridge Univ. Press, 2021).
34. Kopp, R. E., Kemp, A. C., Bittermann, K., Horton, B. P., Donnelly, J. P., Gehrels, W. R., ... & Rahmstorf, S. (2016). Temperature-driven global sea-level variability in the Common Era. *Proceedings of the National Academy of Sciences*, *113*(11), E1434-E1441.
35. Kopp, R. E., Oppenheimer, M., O'Reilly, J. L., Drijfhout, S. S., Edwards, T. L., Fox-Kemper, B., ... & Xiao, C. (2023). Communicating future sea-level rise uncertainty and ambiguity to assessment users. *Nature Climate Change*, 1-13.
36. Sriver, Ryan L., Nathan M. Urban, Roman Olson, and Klaus Keller. "Toward a physically plausible upper bound of sea-level rise projections." Climatic Change 115 (2012): 893-902.
37. Powell, W. B. (2016). Perspectives of approximate dynamic programming. *Annals of Operations Research*, *241*, 319-356.
38. Xian, S., Lin, N., & Hatzikyriakou, A. (2015). Storm surge damage to residential areas: a quantitative analysis for Hurricane Sandy in comparison with FEMA flood map. Natural Hazards, 79(3), 1867-1888.
39. Hatzikyriakou, A., Lin, N., Gong, J., Xian, S., Hu, X., & Kennedy, A. (2015). Component-Based Vulnerability Analysis for Residential Structures Subjected to Storm Surge Impact from Hurricane Sandy. Natural Hazards Review, 05015005.

40. Miura, Yuki, Huda Qureshi, Chanyang Ryoo, Philip C. Dinenis, Jiao Li, Kyle T. Mandli, George Deodatis, Daniel Bienstock, Heather Lazrus, and Rebecca Morss. "A methodological framework for determining an optimal coastal protection strategy against storm surges and sea level rise." *Natural Hazards* 107 (2021): 1821-1843.
41. Big U, THE BIG U 2015, https://rebuildbydesign.org/work/funded-projects/the-big-u/
42. US Army Corps of Engineerings, NY/NJ Harbor Wide Gate/Beach Restoration NY/NJ Harbor and Tributaries Study, 2019.
43. Van Dantzig, David. "Economic decision problems for flood prevention." Econometrica: Journal of the Econometric Society (1956): 276-287.
44. Eijgenraam, Carel, Ruud Brekelmans, Dick den Hertog, and Kees Roos. "Optimal strategies for flood prevention." *Management Science* 63, no. 5 (2017): 1644-1656.
45. Zwaneveld, P., Gerard Verweij, and S. van Hoesel. "Safe dike heights at minimal costs: An integer programming approach." *European Journal of Operational Research* 270, no. 1 (2018): 294-301.
46. Vrijling, J. K. "Probabilistic design of water defense systems in The Netherlands." *Reliability engineering & system safety* 74, no. 3 (2001): 337-344.
47. Aerts, J. C., Botzen, W. W., Emanuel, K., Lin, N., de Moel, H., and Michel-Kerjan, E. O. (2014). “Evaluating flood resilience strategies for coastal megacities.” *Science*, 344(6183), 473–475.
48. Xian, S., Lin, N., & Kunreuther, H. (2017). Optimal house elevation for reducing flood-related losses. *Journal of hydrology*, *548*, 63-74.
49. Ceres, Robert L., Chris E. Forest, and Klaus Keller. "Optimization of multiple storm surge risk mitigation strategies for an island city on a wedge." *Environmental Modelling & Software* 119 (2019): 341-353.
50. van Berchum, Erik C., William Mobley, Sebastiaan N. Jonkman, Jos S. Timmermans, Jan H. Kwakkel, and Samuel D. Brody. "Evaluation of flood risk reduction strategies through combinations of interventions." *Journal of Flood Risk Management* 12, no. S2 (2019): e12506.
51. Ceres, Robert L., Chris E. Forest, and Klaus Keller. "Trade-offs and synergies in managing coastal flood risk: A case study for New York City." *Journal of Flood Risk Management* 15, no. 1 (2022): e12771.
52. Xi, Dazhi, Ning Lin, and Avantika Gori. "Increasing sequential tropical cyclone hazards along the US East and Gulf coasts." *Nature Climate Change* 13, no. 3 (2023): 258-265.
53. Pielke Jr, Roger, Matthew G. Burgess, and Justin Ritchie. "Plausible 2005-2050 emissions scenarios project between 2 and 3 degrees C of warming by 2100." *Environmental Research Letters* (2022).
54. FEMA 2009 Homeowner’s Guide to Retrofitting. Second edition. US Department of Homeland Security: Federal Insurance and Mitigation Administration (FEMA), Washington, D.C.
55. MJ Small and S Xian, A human-environmental network model for assessing coastal mitigation decisions informed by imperfect climate studies. Global Environmental Change Volume 53, November 2018, Pages 137-145
56. Kopp, Robert E., Elisabeth A. Gilmore, Christopher M. Little, Jorge Lorenzo-Trueba, Victoria C. Ramenzoni, and William V. Sweet. "Usable science for managing the risks of sea-level rise." Earth's future 7, no. 12 (2019): 1235-1269.
57. Rasmussen, D. J., Kopp, R. E., & Oppenheimer, M. (2023). Coastal defense megaprojects in an era of sea-level rise: politically feasible strategies or Army Corps fantasies?. Journal of Water Resources Planning and Management, 149(2), 04022077.
58. Sovacool, B., & Linnér, B. O. (2016). The political economy of climate change adaptation. London: Palgrave Macmillan.
59. Sharpe, William F. "Capital asset prices: A theory of market equilibrium under conditions of risk." *The journal of finance* 19.3 (1964): 425-442.
60. Hall, J. W., Lempert, R. J., Keller, K., Hackbarth, A., Mijere, C., & McInerney, D. J. (2012). Robust climate policies under uncertainty: A comparison of robust decision

making and info-gap methods. *Risk Analysis: An International Journal*, *32*(10): 1657-1672.
61. Siders, Anne R. "Social justice implications of US managed retreat buyout programs." *Climatic change* 152, no. 2 (2019): 239-257.
62. Garner, G. G., T. Hermans, R. E. Kopp, A. B. A. Slangen, T. L. Edwards, A. Levermann, S. Nowikci, M. D. Palmer, C. Smith, B. Fox-Kemper, H. T. Hewitt, C. Xiao, G. Aðalgeirsdóttir, S. S. Drijfhout, T. L. Edwards, N. R. Golledge, M. Hemer, G. Krinner, A. Mix, D. Notz, S. Nowicki, I. S. Nurhati, L. Ruiz, J-B. Sallée, Y. Yu, L. Hua, T. Palmer, B. Pearson, 2021. IPCC AR6 Sea Level Projections. Version 20210809. Dataset
63. Kopp, Robert E., Gregory G. Garner, Tim HJ Hermans, Shantenu Jha, Praveen Kumar, Aimée Slangen, Matteo Turilli et al. "The Framework for Assessing Changes To Sea-level (FACTS) v1. 0-rc: A platform for characterizing parametric and structural uncertainty in future global, relative, and extreme sea-level change." EGUsphere (2023): 1-34.
64. Nias, Isabel J., Sophie Nowicki, Denis Felikson, and Bryant Loomis. "Modeling the Greenland Ice Sheet's Committed Contribution to Sea Level During the 21st Century." *Journal of Geophysical Research: Earth Surface* 128, no. 2 (2023): e2022JF006914.
65. Levermann, A., Winkelmann, R., Albrecht, T., Goelzer, H., Golledge, N. R., Greve, R., ... & Van De Wal, R. S. (2020). Projecting Antarctica's contribution to future sea level rise from basal ice shelf melt using linear response functions of 16 ice sheet models (LARMIP-2). *Earth System Dynamics*, *11*(1), 35-76.
66. DeConto, R. M., Pollard, D., Alley, R. B., Velicogna, I., Gasson, E., Gomez, N., ... & Dutton, A. (2021). The Paris Climate Agreement and future sea-level rise from Antarctica. *Nature*, *593*(7857), 83-89.
67. Bamber, J. L., Oppenheimer, M., Kopp, R. E., Aspinall, W. P., & Cooke, R. M. (2022). Ice sheet and climate processes driving the uncertainty in projections of future sea level rise: Findings from a structured expert judgement approach. *Earth's Future*, *10*(10), e2022EF002772.
68. Lin, N. and Shullman, E. (2017). “Dealing with hurricane surge flooding in a changing environment: part I. risk assessment considering storm climatology change, sea level rise, and coastal development.” Stochastic Environmental Research and Risk Assessment, 31(9): 2379–2400.

Figures and Tables

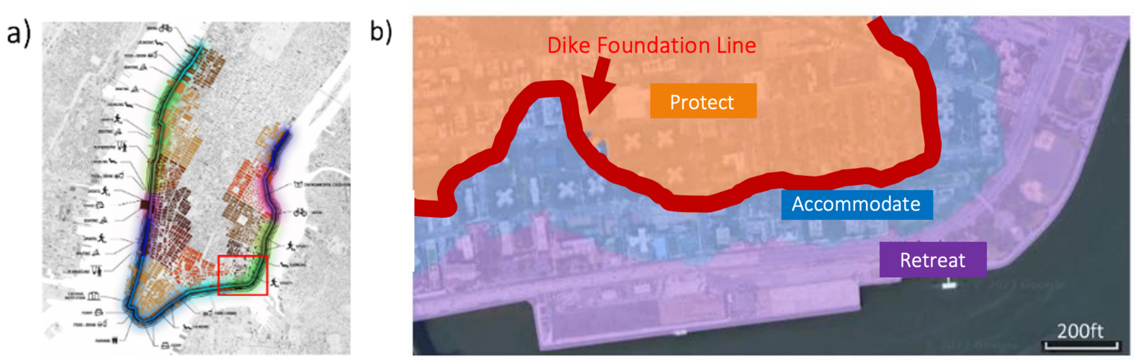

Figure 1. Study Area in NYC. a) The “BIG U” protected region in Lower Manhattan. b) Illustrative distribution on Digital Elevation Map (DEM) of the

withdrawal/accommodation/dike-protected zones, shown as an example for the area marked with a red rectangular in panel (a). The retreat zone (purple) is lower in altitude than the accommodation zone (blue). The dike protects regions beyond the accommodation zone (orange).

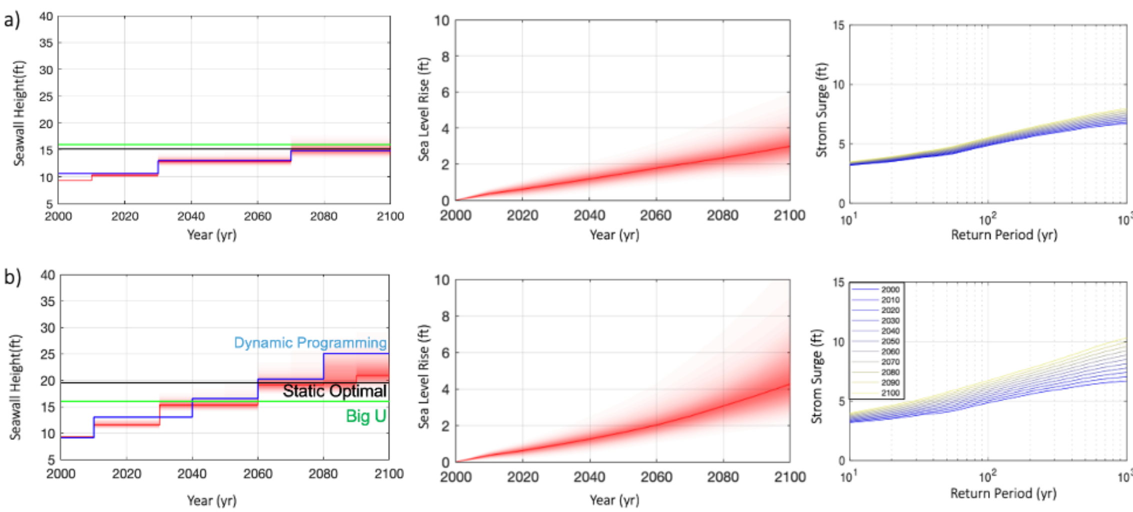

Figure 2. Analysis of seawall height strategies suggested by different models (left panel) and SLR projection (middle panel) and storm tide projection (right panel) under a) SSP2 4.5 and b) SSP5 8.5. Green line shows the “Big U” level. Black line shows the static-optimal level. In the left panel, blue curve shows the dynamic optimal level. Red curve shows the medium strategy by RL; the red shade shows the probability density function (PDF) of the seawall height by RL, with a darker color corresponding to a higher probability. In middle panel, the center curve shows the median and the shade shows the PDF of SLR projection. In right panel, the storm tide return level is shown for each decade.

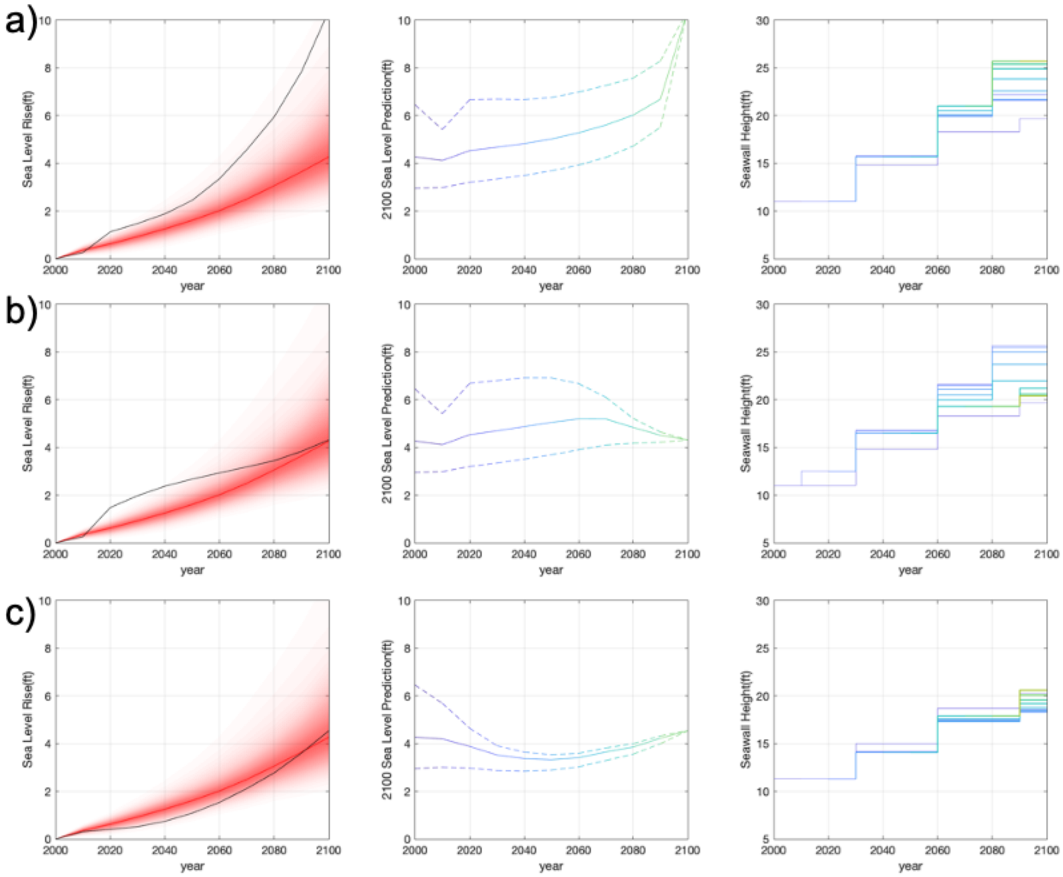

Figure 3. Sample paths of seawall height under different SLR realizations under SSP5 8.5. The first column shows selected SLR realization (black) over the background SLR distribution (red). The middle column shows the 2100 SLR prediction conditioned on the SLR realization over the 21st century (5%-95% confidence interval). The third column shows the seawall height time series planned at different time points, where different colors indicate the different time points: blue for the decision made at the beginning of the 21st century, and yellow for the decision made close to the end of the 21st century (similar to the middle column).

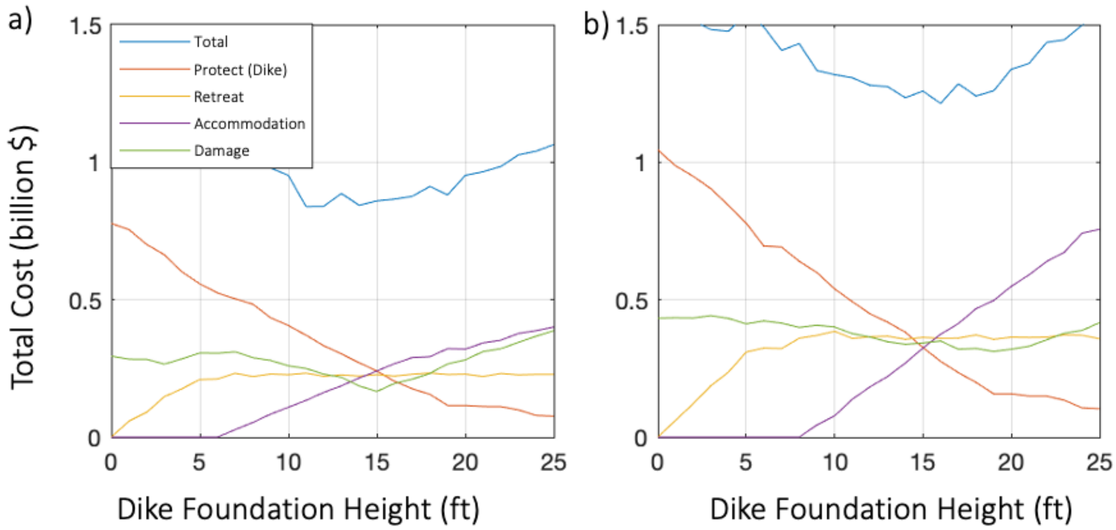

Figure 4. The composition of total cost under the multidimensional flood risk mitigation strategies given different ground elevations of dike under a) SSP2 4.5 and b) SSP5 8.5

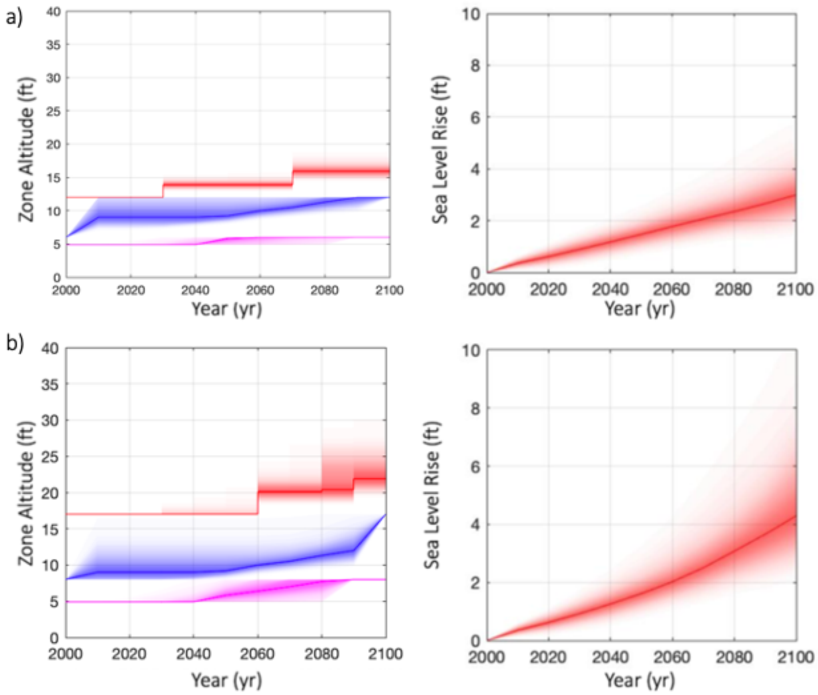

Figure 5. Analysis of retreat/accommodation zone and dike coverage by RL (left panel) for a) SSP2 4.5 and b) SSP5 8.5, under the projected SLR (right panel). Magenta curve shows the medium of the retreat zone boundary. Blue curve shows the medium of the accommodation zone boundary. Red curve shows the medium dike height. The shade in the left panel shows PDF, indicating the probability of certain retreat/accommodation zone boundary or dike height, with a darker color corresponding to a higher probability. The shade in the right panel shows PDF of SLR projection.

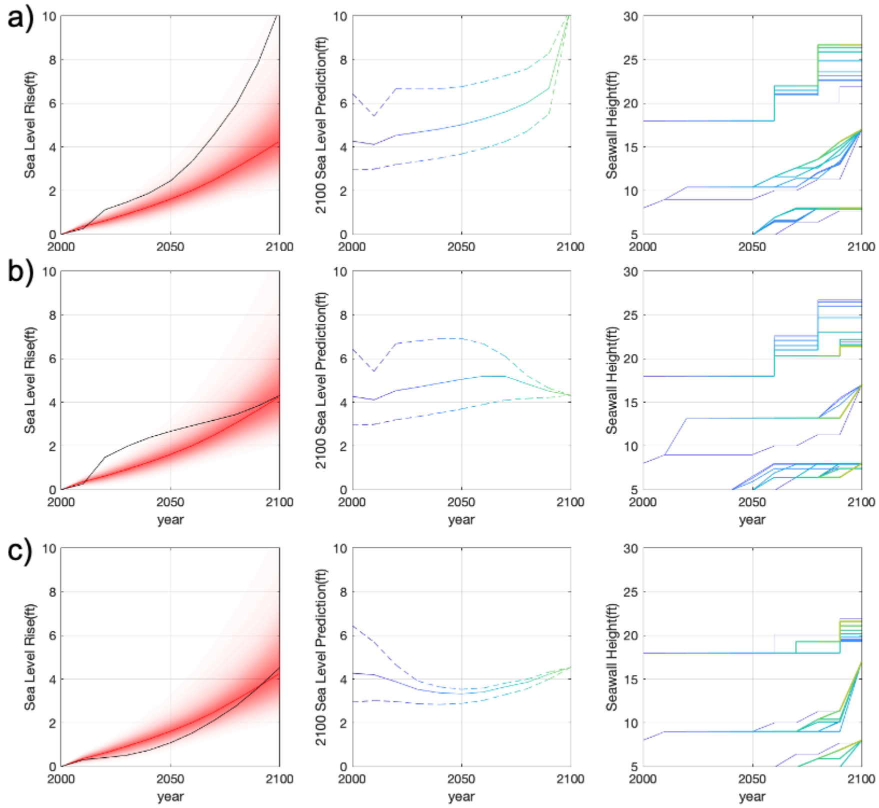

Figure 6. Sample paths of retreat zone boundary, accommodation zone boundary and dike height under different SLR realizations under SSP5 8.5 (same paths as discussed in Figure. 3). The first column shows selected SLR realization (black) over the background SLR distribution (red). The middle column shows the 2100 SLR prediction conditioned on the SLR realization over the 21st century (5%-95% confidence interval). The third column shows the seawall height time series planned at different time points, where different colors indicate the different time points: blue for the decision made at the beginning of the 21st century, and yellow for the decision made close to the end of the 21st century (similar to the middle column).

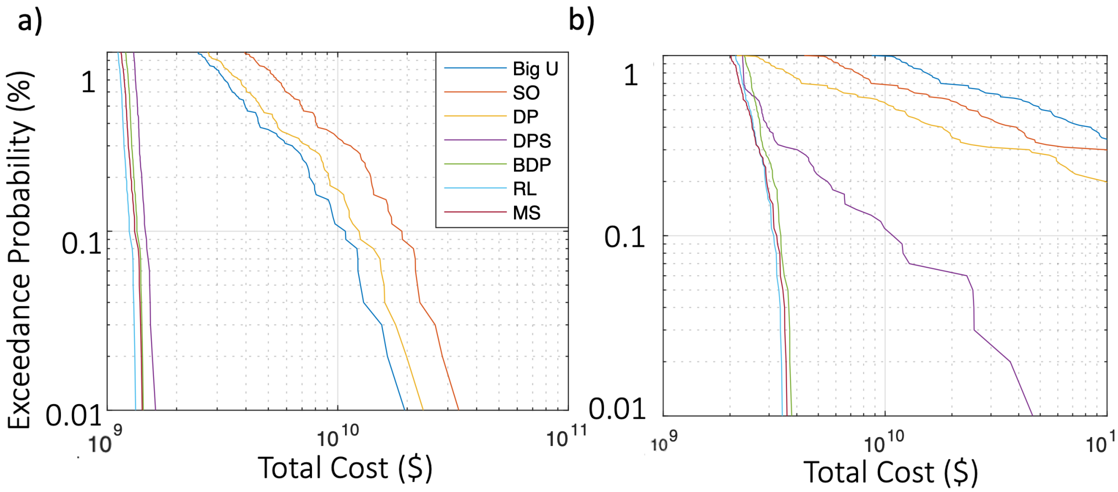

Figure 7. Quantile of total cost under different strategies under a) SSP2 4.5 and b) SSP5 8.5.

Table 1. Different Policy Design Frameworks, including design based on return period of the hazard and cost-benefit optimal statistic and dynamic strategies. The frameworks considered in this study are bolded.

Algorithm	Dynamic Design/Policy	Update with new observations	Systematic Design/Policy	Ref
Return Period (Big U)	-	-	-	Big U,2015
Static Optimal (SO)	-	-	-	Van Dantzig, 1956
Dynamic Programming (DP)	●	-	-	Lickley et al., 2014
Classic Heuristic	●	-	-	Keller et al., 2007
Adaptation Policy Pathway	●	○	-	Ranger et al., 2013; Nicholls et al., 2014
Bayesian DP (BDP)	●	●	-	Bruin, 2009; van der Pol et al., 2017
Decision Tree	●	●	○	Anda 2009
Real Options	●	●	○	Zhang et al., 2014
Direct Policy Search (DPS)	●	●	○	Garner et al., 2018; Giuliani et al., 2016
Reinforcement Learning (RL)	●	●	●	Proposed in this paper

- : no ability
○ : high computational cost
● : full ability

Table 2. Expected loss, policy expenditure, and bias loss under different scenarios and different optimization frameworks (billion USD). Four scenarios are included: projected/assumed to be SSP5 8.5 LC or SSP2 4.5 MC with reality of SSP5 8.5 LC or SSP2 4.5 MC. In comparison to the Big U strategy, six optimization frameworks are considered: static optimal (SO), dynamic programming (DP), Bayesian dynamic programming (BDP), direct policy search (DPS), reinforcement learning (RL), and RL for multiple strategies (MS).

	Projection	Reality	Big U	SO	DP	BDP	DPS	RL	MS
Expected Loss	SSP 5 8.5 LC	SSP 5 8.5 LC	10.21	3.25	2.38	1.67	1.78	1.47	1.25
Policy Expenditure			1.32	2.33	1.35	1.23	1.29	1.15	0.93
Expected Loss	SSP 2 4.5 MC		10.21	14.83	8.63	4.36	3.87	3.85	3.70
Policy Expenditure			1.32	1.21	0.77	0.84	0.96	0.90	0.98
Bias Loss			-	11.58	6.25	2.69	2.09	2.38	2.45
Expected Loss	SSP 2 4.5 MC	SSP 2 4.5 MC	1.37	1.31	1.02	0.93	0.97	0.89	0.84
Policy Expenditure			1.32	1.21	0.77	0.65	0.73	0.60	0.58
Expected Loss	SSP 5 8.5 LC		1.37	2.33	1.38	1.08	1.22	1.03	0.95
Policy Expenditure			1.32	2.33	1.35	0.95	1.15	0.88	0.69
Bias Loss			-	1.02	0.36	0.15	0.25	0.14	0.11