ADVANCES IN COASTAL FLOOD RISK ANALYSIS

by

Nathan B. Geldner

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy

School of Industrial Engineering

West Lafayette, Indiana

May 2023

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL

Dr. David R. Johnson, Chair

School of Industrial Engineering, Department of Political Science

Dr. Andrew Liu

School of Industrial Engineering

Dr. Venkatesh Merwade

Lyles School of Civil Engineering

Dr. Scott Hemmerling

The Water Institute of the Gulf

Approved by:

Dr. Young-Jun Son

This dissertation is dedicated to my wife Rimi, who inspired in me the courage, determination, and self-compassion I needed to see this through.

ACKNOWLEDGMENTS

I would first like to thank my graduate advisor Professor David R. Johnson, whose mentorship and guidance made me the researcher I am today, and whose empathy and patience have exceeded my expectations at every turn, challenging me to be a better scientist and a better person. I'd like to thank Susan R. Hunter for her mentorship as well in helping me to become a more quantitatively rigorous and competent researcher. I'd like to thank Professors Venkatesh Merwade, David Yu, Thomas Hertel, and Juan Sesmero for broadening my horizons outside of industrial engineering. A special additional thanks to Professor Hertel; without his AGEC 528 course I wouldn't have met my wonderful wife. I would also like to thank Professors Klaus Keller and James Doss-Gollin for their input in structuring, framing, and providing feedback as coauthors of the article (pending submission for publication) which acts as chapter 4 of this dissertation. I would like to thank Jingya Wang, Zachariah T. Richardson, and Andrew B. Zehr who worked with me in developing the SWaMPS model used in chapter 3 of this dissertation under Professor Johnson. I would also like to thank Dr. Jordan Fischbach, Hugh Roberts, and Dr. Shubhra Misra for their leadership and everyone I've worked with from Rand Corporation, the Water Institute of the Gulf, and Louisiana's Coastal Protection and Restoration Authority for working with me on Louisiana's 2023 Coastal Master Plan and the Louisiana Watershed Initiative, which together have made up some of the most fulfilling and enjoyable work of my graduate school career. I also acknowledge the Coastal Protection and Restoration Authority and the Louisiana Office of Community Development for funding my work on those projects respectively. I also acknowledge the Fredricks N. Andrews Fellowship and the Lee A. Chaden fellowship which also funded a portion of my doctoral research. I would also like to thank my committee members for their feedback and their participation in my preliminary and final doctoral exams. I thank the Purdue University School of Industrial Engineering for giving me this opportunity. Finally, I'd like to thank my parents for their constant love and support, and my brother Aaron for repeatedly, proverbially, pulling my head out of a place where it shouldn't be.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABBREVIATIONS AND NOMENCLATURE

ADCIRC-SWAN – a numerical simulation model which estimates both surge and wave behavior associated with a tropical storm.

CBA – Cost-benefit analysis, a decision-making framework in which strategies are selected principally based on economic measures of cost effectiveness or net present value.

CLARA – the Coastal Louisiana Risk Assessment Model used in Louisiana's Coastal Master Plan

COMPOUND FLOODING – flooding driven by combined surge, wave, rainfall, and riverine dynamics.

CORRECTIVE JUSTICE – the idea that individuals who are unjustly harmed are entitled to restitution from those who benefit from said harm.

DEEP UNCERTAINTY – uncertainty which cannot be modelled probabilistically due to lack of knowledge or disagreement.

(DEONTOLOGICAL) EGALITARIANISM – The belief that every individual should receive equal consideration in governmental processes irrespective of status quo resources allocations.

DIRECT ECONOMIC DAMAGE – a measure of flood risk representing the dollar value of expected flood damage over time.

DISTRIBUTIVE JUSTICE – The idea of justice as it relates to the distribution of resources and the perception of resource distributions as fair or just.

EJPM-OS – the extended joint probability framework with optimal sampling, a probabilistic modeling framework for compound flooding.

HEC-HMS – a simulation model of the hydrologic processes of a watershed.

HEC-RAS – a numerical simulation model which estimates the flow of water over time through river channels and floodplains.

JPM-OS – the joint probability method with optimal sampling, a probabilistic modeling framework for surge- and wave-driven flooding.

K-MEANS CLUSTERING – an unsupervised machine learning method which groups together similar observations.

MORO-MS – Multi-scenario multi-objective robust optimization, a framework for generating strategies in decision-making under deep uncertainty which treats each objective in each modelled plausible future as a separate objective in a multiobjective context.

NONSTRUCTURAL PROTECTION – flood protection measures which do not affect the distribution of flood depths.

OPTIMAL SAMPLING DISCRETIZATION – a proposed method which produces a discrete approximation of a continuous random variable or a coarser discrete approximation of a discrete random variable so as to minimize the square error induced by approximation.

PRIORITARIANISM – The belief that resources should be preferentially allocated to those individuals and groups with the fewest resources or greatest social vulnerability in order to produce more equal outcomes.

PROCEDURAL JUSTICE – the idea of justice as it relates to governmental processes and the perception that those processes are fair or just.

REAL OPTIONS ANALYSIS – a framework for making decisions under uncertainty which prices in the value of flexibility and adaptivity, loosely informed by methods used in the analysis of financial options.

RESIDENCE LOSS EQUIVALENTS – a measure of flood risk representing the aggregate expected proportional damage to a home over time, such that 1 residence loss equivalent implies an expected damage equal to the replacement cost of the home.

STRUCTURAL PROTECTION – flood protection measures which affect the distribution of flood depths.

SWaMPS – the Surge and Wave Model for Protection Systems, a spatially narrower, simplified implementation of CLARA for single-polder systems (i.e., systems which can be treated as a stylized bowl with a single representative flood depth for a given storm event) based on the Larose to Golden Meadow Ring Levee system.

SYNTHETIC STORM – a discrete realization of an idealized continuous parameterization of tropical storm attributes; an idealized tropical storm characterized by a fixed set of parameters.

ABSTRACT

Much work has been done to advance the state of risk-informed decision-making to protect against coastal flooding. The state of the art as practiced in the state of Louisiana, the Coastal Louisiana Risk Assessment Model (CLARA) characterizes flood risk driven by surge and waves from tropical storms as a random process and estimates the distribution of flood depths and resulting economic damage. This dissertation identifies three key limitations of coastal flood risk assessment as applied to the state of Louisiana, proposes methods to address them, and demonstrates each method in a case study. The first limitation identified is that the CLARA model addresses surge- and wave-driven flood hazard but does not account for rainfall-driven and riverine hazard. To address this limitation, chapter 2 presents an extension of the methods used in CLARA developed as part of the Louisiana Watershed Initiative which permits characterization of compound hazard consisting of surge, rainfall, and riverine hazard from tropical cyclones. The second limitation identified is that due to its computational cost CLARA is used in Louisiana's 2023 Coastal Master Plan to evaluate pre-specified and static flood risk mitigation projects over a small set of possible future landscapes; it would be preferable to use an optimization-driven approach to generate efficient and adaptive combinations of projects which balance performance across a diverse set of possible future landscapes. To address this limitation, chapter 3 presents an analysis applying a spatially narrower but computationally inexpensive model based on CLARA called the SWaMPS model to a case study to determine the extent to which optimization-based adaptive mitigation strategies which respond to observed climate change trajectories can outperform similar static strategies. The third and final limitation identified is that while the most recent iteration of the CLARA model can support algorithmic optimization of cost-effectiveness of building-level mitigation strategies such as retrofits to increase the first floor elevation of single-family residences, the principal metric historically used in flood risk mitigation is reduction in economic damage measured in dollars, the optimization of which on the level of individual buildings would implicitly prioritize expensive structures and therefore may neglect impoverished neighborhoods. Chapter 4 addresses this limitation with a proposed alternative efficiency metric which treats individual homes as equally valuable, the optimization of which results in greater investment in impoverished neighborhoods without explicitly targeting final expected damage distributions or individual groups.

1. INTRODUCTION

Tropical cyclone-induced flooding presents a hazard to coastal regions across the globe. The Congressional Budget Office estimates an annual damage associated with tropical cyclones of $54 billion annually in the United States alone [1]. Various interventions are available to mitigate the damages associated with hurricane-induced flooding including the construction of levees (physical barriers to prevent floodwaters from entering a system) and elevation-in-place of individual structures (raising the structure on an elevated pier foundation such that it will not be impacted by a given level of flooding). As the occurrence of tropical cyclones and therefore flooding induced by tropical cyclones is inherently stochastic, effective mitigation of the associated damages requires effective characterization of coastal flood risk and identification of efficient, risk-informed and socially acceptable designs for coastal protection projects.

Historically, coastal protection projects such as levee construction have used a standards-based design process [2], [3]. A standard design load is identified, such as the 100-year surge event i.e., the magnitude of surge which will be exceeded with a 1% annual probability under stationary environmental conditions. A levee is then designed to ensure that overtopping remains below some threshold when subject to the standard design load. This approach has the benefit of being tractable and straightforward but is ultimately arbitrary as a standard for environmental decision-making. Beyond the 100-year surge event, protection afforded by a levee is characterized both by its effect on the expected annual damage caused by tropical cyclones and by its performance at a range of return periods. In order then to design protection projects which minimize damage in expectation or target multiple return periods to reflect diverse risk attitudes among decision-makers and stakeholders, we require models which characterize flood risk as a random process driving a probability distribution of flood depths and corresponding damages.

This dissertation poses three major questions which arise throughout the process of modelling coastal flood risk in support of policy decision-making. It then identifies limitations in how these questions are addressed in the current state of practice and proposes and evaluates methods which can help address those limitations. The questions are: 1) How do we characterize coastal flood hazard? 2) How do we identify high-performing strategies for mitigating the consequences of that hazard given specified performance metrics? 3) What performance metrics should we consider when evaluating mitigation strategies? In addressing these questions

sequentially, this dissertation is intended to provide novel contributions throughout the process of modelling coastal flood risk in support of policy decision-making.

The current state of practice of coastal flood risk modeling as practiced in the state of Louisiana is the Coastal Louisiana Risk Assessment model (CLARA). Using a wide range of inputs, CLARA identifies a probability distribution of flood depths for each of 126,174 grid cells along the coast and utilizes a comprehensive structure inventory with known depth-damage relationships to characterize the expected damage at a range of return periods [4]–[6]. Louisiana's Comprehensive Master Plan for a Sustainable Coast uses the CLARA model to compare the performance of competing standards-based structural and non-structural project designs. In this context a structural project denotes the creation or improvement of a protection structure such as a levee or floodwall, and a nonstructural project denotes a subsidy or cost-sharing program for individual privately owned structures to be floodproofed (up to several feet), elevated in place, or purchased outright.

While CLARA is effective in producing coastwide risk estimates with and without projects in place, it has several notable limitations corresponding to questions posed above. Firstly, in areas outside of structural protection systems it considers only surge- and wave-driven flood hazard and not flood hazard driven by rainfall or riverine dynamics. This is not an issue within a certain distance of the coast as the influence of rainfall is strictly dominated by that of surge and waves, but consideration of rainfall and riverine dynamics is critical to characterizing flood risk further inland. Secondly, CLARA's estimation of flood depths within structural protection systems is too slow to permit optimization of system designs. Instead, of protection measures evaluated by CLARA typically consist of pre-prescribed standards-based designs which are compared on a cost effectiveness basis, and non-structural protection designs are evaluated similarly conditionally upon selected protection projects. The CLARA model does not permit for the identification of cost-effective combinations of structural and nonstructural protection under a range of uncertain climate futures. Thirdly, the 2023 revision of CLARA is (for the first time) technically capable of identifying maximally cost-effective allocations of nonstructural protection resources among individual buildings across the coast (conditioned upon a fixed set of structural projects). However, such a project design has not been implemented. This is in large part because it has the potential to worsen historical social inequities, as minimizing economic damage as measured in dollars implicitly prioritizes more expensive structures. Such concerns warrant investigation into

alternative decision metrics. The projects described in this dissertation are meant to respectively address these three limitations.

Chapter 2 of this dissertation presents a novel statistical framework for characterizing compound coastal flood hazard, defined as flood hazard driven by combined surge, rainfall-driven, and riverine dynamics, which extends the probabilistic framework underlying the CLARA model. Chapter 3 of this dissertation presents an analysis using the SWaMPS model, a scaled down computationally efficient model based on CLARA, and a multiobjective evolutionary algorithm to investigate the tradeoffs between optimized flood risk mitigation interventions targeting different severities of climate future, and the extent to which these tradeoffs are mitigated by an explicitly adaptive decision process inspired by real options analysis. Chapter 4 of this dissertation investigates the extent to which targeting individual household-level interventions to maximize damage reduction as measured in dollars may neglect disadvantaged neighborhoods, and the extent to which this effect may be mitigated by using a more egalitarian measure of risk. Taken together, these three projects represent a set of advances not only in one sub-problem of coastal flood risk modeling for policy decision support, but rather reflect a holistic set of advances throughout the modeling process.

2. APPLIED JOINT PROBABILISTIC MODELING OF COMPOUND COASTAL FLOOD HAZARD: AN EXTENSION OF THE JOINT PROBABILITY METHOD WITH OPTIMAL SAMPLING

Compound coastal flooding, i.e., flooding driven by interacting pluvial, riverine, and coastal dynamics, poses a significant hazard which in some areas is much greater than can be attributed to inland or coastal dynamics separately [7]–[10]. Characterizing this hazard requires two major model components: a physically driven simulation model (or metamodel thereof) which estimates flood depths resulting from a given storm event, and a statistical model which estimates the probability distribution of the number and characteristics of storm events in a given year. A number of physically driven simulation models exist for estimating compound coastal flood hazard, though their development and improvement remain an active area of research [11]. As illustrated below, a meaningful research gap exists in the statistical modeling of compound coastal flood hazard, particularly in discretizing hazard distributions to a set of events which can be tractably simulated by computationally expensive simulation models. This article describes a framework for the statistical modeling of compound coastal flood hazard developed for the Louisiana Watershed Initiative entitled the extended joint probability method with optimal sampling (EPJM-OS) and its initial implementation in a preliminary pilot analysis for the Amite River Basin in coastal Louisiana.

2.1 Methodological Context

Statistical models of compound flooding from tropical cyclones consist of three major components. The first component, the recurrence rate, estimates the rate at which tropical cyclones occur; this is needed to convert statistical estimates of flood hazard made on a per-storm basis to an annualized basis for intuitive use in policy decision-making. The second component describes the continuous joint distribution of tropical cyclone features which drive hazard when tropical cyclones occur. The third component samples events to simulate from the joint distribution and assigns probability masses or relative likelihoods (i.e., represents the continuous joint distribution as a discrete approximation) to each event for the purpose of producing a cumulative distribution function for the simulations' outcome variables of interest (e.g., peak surge elevations, inundation depths).

The occurrence of tropical cyclones is typically modelled as a Poisson process with a rate parameter λ estimated with a capture zone or kernel function weighting approach [12], [13]. The capture zone approach simply counts and averages the number of storms passing through a specified area per year. The kernel function weighting approach applies a smoothing kernel to the travel paths or "tracks" of historical storms and integrates the resulting kernel frequency density over a length of idealized coastline or region of interest.

The continuous joint distribution of tropical cyclone features is typically captured using copulas, physically driven Monte Carlo ensembles, or joint probability methods. Copulas are common in the literature and easy to use, as they require only specified marginal distributions and simplified dependence structures between hazard drivers such as peak surge and total rainfall [14]–[16]. In practice can only produce dependency structures which match one or at most two dependency measures of the true joint distribution [17], e.g. the meta-Gaussian copula which captures rank correlation only. Physically driven Monte Carlo ensembles are less common and are generated by randomly seeding tropical cyclone vortices and evolving them with deterministic meteorological simulations [18]. This approach carries advantages for the physical realism of individual cyclones but may or may not reflect the true joint variance structure of storm features. Additionally, this approach requires a large ensemble size to fully characterize the variance of possible tropical cyclones, which may be prohibitive when using computationally expensive simulation models to evaluate resulting flood depths. Joint probability methods are uncommon in compound flood hazard analysis and more commonly used for purely coastal i.e., surge- and wave-driven, flood hazard characterization. Joint probability methods leverage empirically derived statistical relationships and conditional independence structures permitting analysts to flexibly express the joint distribution of tropical cyclone features as a series of conditional distributions or Bayesian factorization [5], [19], [20]. Most applications of these methods specifically model specifically model the characteristics of storms at a representative point in time such as the time of landfall in order to keep the dimensionality of the probabilistic models at a reasonable size. While realistically two tropical cyclones can have similar attributes at a single point in time (e.g., landfall) and diverge in their characteristics later, parametric approaches such as copulas and joint probability methods idealize tropical storm behavior as being uniquely determined by their characteristics at the representative point in time.

Continuous joint distributions of tropical cyclone features are typically discretized in one of three ways: naïve Monte Carlo sampling [21], structured samples [5], and optimization-driven subsampling of larger Monte Carlo or structured samples [4], [19]. The idealized discrete storm events in the resulting distribution are referred to as synthetic storms. Naïve Monte Carlo sampling directly samples from the continuous joint distribution but requires a large sample size. A structured sample can more efficiently span tropical cyclone parameter space but relies on a heuristic integration scheme to assign probability masses and may also require a large sample size. Optimization-driven subsampling is often used in coastal flood risk analysis to reduce the set of synthetic storms derived from either method to a size for which flood depths can be more tractably simulated [22], [23], but requires initial simulation of the original set [4] or the use of Bayesian quadrature which assumes both a closed form representation of the continuous hazard distribution and unrealistic assumptions about the variance structure of conditional flood depth exceedance probabilities [22].

2.2 Simulation Methods

While statistical modeling i.e., the parameterization of a set of storm events with corresponding probability masses and a rate of occurrence to represent the distribution of tropical cyclones which may occur and their characteristics, physically driven simulation is necessary to characterize the flood depths which arise from those characteristics. External collaborators on the preliminary pilot analysis of the Louisiana Watershed Initiative for the Amite River Basin used HEC-HMS to model upstream hydrologic behavior [24]. HEC-HMS models the complete hydrologic process of a watershed including infiltration, routing, evapotranspiration, and soil moisture accounting. A lumped model was used, modeling each sub-basin as a single unit. External collaborators used the ADCIRC-SWAN model developed for the 2023 Louisiana Coastal Master Plan to model downstream forcing from surge and waves [25]. ADCIRC numerically solves circulation and transport problems on an unstructured grid and is used to model storm surge behavior under meteorological forcing while the SWAN component computes random wind-generated waves in coastal and inland waters. External collaborators simulated flood depths in the study domain driven by upstream hydrological forcing and downstream surge and wave forcing using a two-dimensional HEC-RAS model [26] with unsteady flow routing calculated over an irregular mesh via the shallow water equations with the Eulerian-Lagrangian method. HEC-RAS

with unsteady flow routing numerically solves a simplified form of the Navier-Stokes equations in order to estimate the flow of water over time through river channels and floodplains.

2.3 Project Context for Statistical Methods

The methods used to model the continuous joint distribution of tropical cyclone features and the rate of occurrence of tropical cyclones for the Louisiana Watershed Initiative's preliminary pilot analysis of the Amite River Basin (as described below in section 2.2.1) were decided from the outset of the planning phase by external collaborators, but discretization required novel methods development. External collaborators responsible for HEC-RAS modeling reported that no more than 200 storm events could be simulated due to limited local computing resources as well as the high physical resolution of the available HEC-RAS model. This prohibited the use of naïve Monte Carlo sampling as well as any optimization-driven subsampling approach involving initial simulation of a large sample. Structured sampling and Bayesian Quadrature were prohibited by a lack of a closed-form expression for rainfall behavior as described below in section 2.2.1. This left no established method in the coastal flood risk analysis literature appropriate to meet the project's needs and prompted an urgent research question: How to select a small number of storm events for hydrodynamic simulation which accurately capture the joint distribution of surge and hydrologic behavior? While the broader outcome of the preliminary pilot analysis represents a step towards eventual statewide characterization of compound flood hazard to inform policy decision-making for flood risk mitigation, my principal contribution and the focus of this article is the proposal and implementation of a discretization approach which addresses this research question in a way which is approximately optimal in the sense of minimizing an intuitive and commonly used loss function.

2.4 Statistical Methods

The methods presented here are an extension of the joint probability method with optimal sampling (JPM-OS), which was developed for purely coastal i.e., surge- and wave-driven hazard. As the name suggests, JPM-OS uses a joint probability method to model the continuous distribution of tropical cyclone hazard and an optimization-driven subsampling scheme for discretization [27]. The joint probability method is extended to incorporate compound hazard with

the introduction of a stochastic rainfall field generator [28] which characterizes the spatiotemporal distribution of rainfall for a given synthetic storm via Monte Carlo sampling. This requires a novel approach to discretization, as the number of events required for convergence with naïve Monte Carlo sampling is in many cases too large to simulate with hydrodynamic models such as HEC-RAS, and the lack of a closed-form conditional distribution function prevents us from using Bayesian quadrature. The extended methodology presented here therefore incorporates a clustering-based optimal sampling approach which subsamples from a large probability-weighted set of events based on the boundary conditions of each event or based on a reduced-fidelity, high-speed simulation of each event, so as to minimize the integrated square error induced in subsampling.

The "extended JPM-OS" (EJPM-OS) implementation described here was developed during a preliminary pilot study for the Louisiana Watershed Initiative and applied in an illustrative case study to the Amite River Basin. This preliminary pilot analysis was performed in the interest of methods development, and a revised pilot analysis with revised methods applied to the Amite River Basin is ongoing as of March 2023. The methods of the revised pilot analysis will later be adapted for coastwide implementation. Due to ongoing methodological development and varying operational constraints arising within the Louisiana Watershed Initiative, there are notable differences, particularly in optimal sampling implementation, between the methods of the preliminary pilot analysis, the ongoing revised pilot analysis, and the planned coastwide analysis. The remainder of this section first describes the methods used in the preliminary pilot analysis, followed by a discussion of how the clustering-based optimal subsampling approach can be applied in a more general context both within and outside of the EJPM-OS framework, under varying operational constraints relating to the availability and runtimes of hydrologic and hydrodynamic models.

2.4.1 Continuous Model of Hazard Drivers

It was decided at the outset of the project that the statistical model of joint flood hazard would extend the JPM-OS method as implemented in the CLARA model used in Louisiana's 2023 Coastal Master Plan [4]. This version of CLARA used a one-dimensional capture zone (i.e., line-crossing) approach for recurrence analysis. CLARA uses a joint probability method to characterize the continuous joint distribution of five tropical storm parameters at landfall: landfall location,

central pressure, radius of maximum windspeed, heading angle, and forward velocity. The joint density function is expressed below.

$$\Lambda(c_p, r, v_f, \theta_l, x) = \Lambda_1 \cdot \Lambda_2 \cdot \Lambda_3 \cdot \Lambda_4 \cdot \Lambda_5$$

$$\Lambda_1 = f(c_p|x) = \frac{\delta}{\delta c_p}\left\{\exp\left\{-exp\left[-\frac{c_p - (a_0(x) + a_1(x)t)}{a_2(x)}\right]\right\}\right\}$$

$$\Lambda_2 = f(r|c_p) = \frac{1}{\sigma(c_p)\sqrt{2\,\pi}} e^{\frac{\left(\ln r - \bar{r}(c_p)\right)^2}{2\sigma^2(c_p)}}$$

$$\Lambda_3 = f(v_f|\theta_l) = \frac{1}{\sigma\sqrt{2\,\pi}} e^{\frac{\left(v_f - \overline{v_f}(\theta_l)\right)^2}{2\sigma^2}}$$

$$\Lambda_4 = f(\theta_l|x) = \frac{1}{\sigma(x)\sqrt{2\,\pi}} e^{\frac{\left(\theta_l - \overline{\theta_l}(x)\right)^2}{2\sigma^2(x)}}$$

$$\Lambda_5 = f(x) = \Phi(x)$$

This joint distribution was identified empirically by Resio et. al. [27] leveraging observed conditional independence relationships.

The distribution of rainfall, conditional on the five tropical cyclone parameters used in CLARA, was modelled using the stochastic rainfall generator for tropical cyclone produced by Villarini et al. [28]. This generator estimates the expected rainfall associated with a synthetic storm and samples from a parameterized model of the residual variance. In doing so it captures and samples from the aleatory uncertainty in rainfall associated with each synthetic storm. Discrepancies in water mass balance in hydrological model calibration revealed bias in the Stage IV precipitation data which had been used to calibrate the generator, taken from the National Centers for Environmental Prediction. This resulted in five equiprobable bias correction factors being applied to rainfall fields produced by the generator. Additionally, it was found by the hydrological modeling team that the distribution of antecedent conditions (e.g., soil moisture and riverine base flows) could be reasonably represented using three equiprobable cases.

2.4.2 Optimal Sampling Discretization as Implemented

The CLARA model discretizes the distribution of the 5 storm parameters around a structured sample of 645 probability-weighted synthetic tropical cyclones [29]. In the 2023 Coastal Master Plan, this structured sample was used to characterize the hazard distribution under existing

conditions and an optimization-driven subsampling approach was used to generate a reduced synthetic storm set for characterizing coastal hazard over multiple parameterized climate trajectories with various flood risk mitigation interventions [4]. In the preliminary pilot analysis of the Amite River Basin, it was decided due to computational constraints to use the CLARA subsampling approach for coastal hazard to reduce the set of 645 synthetic storms to a set of 50 storms. Stochastic rainfall fields were then generated for each of these 50 events, under each of the three antecedent conditions cases, with each of the five bias correction factors, for a total of 37,500 events.

Flood depths for discrete events were simulated via HEC-RAS. Upstream boundary conditions were modelled using HEC-HMS and downstream boundary conditions (i.e., storm surge and waves) were modelled using ADCIRC + SWAN. ADCIRC + SWAN did not pose a computational bottleneck both because only one simulation was necessary per synthetic storm (downstream hydrodynamic forcing was modelled as independent of antecedent conditions and rainfall) and because results were already available as intermediate data products from the 2023 Coastal Master Plan. Similarly, as HEC-HMS models hydrological systems as idealized networks with dynamics governed by empirically derived heuristics rather than with spatially explicit physical modeling, it was considered feasible to characterize hydrological dynamics for all 37,500 events. However, the HEC-RAS model available was computationally expensive, and with available computational resources only 200 HEC-RAS simulations were feasible. This raised the question of how to subsample 200 events from the probability-weighted set of 37,500 events while introducing as little error as possible into our representation of the distribution of compound flood hazard drivers.

Optimal Sampling Discretization for Compound Coastal Flood Risk

The goal of optimal sampling discretization for compound coastal flood risk is to approximate a continuous random variable as a discrete random variable or to approximate a discrete random variable with many possible outcomes as a discrete random variable with fewer possible outcomes. There are two loss functions commonly used to scalarize error in the approximation of continuous functions: mean square error and mean absolute error. The choice of which loss function to use depends on the desired attributes of the approximation with respect to its sensitivity to outliers; as a first order approximation, the mean minimizes mean square error

whereas the median minimizes mean absolute error. Machine learning practitioners often select mean absolute error because it is less sensitive to outliers [30]. Conversely, in approximating the distribution of compound coastal flood hazard drivers, sensitivity to outliers in the tails of the hazard distribution is desirable. Therefore, optimal sampling discretization minimizes the integrated square error of approximation (equivalent to the mean square error as the normalization constant integrates to 1) expressed as a loss function in Equation 1.

$$L_X(X') = \int_{\Omega} \left|\left|X(\omega) - X'(\omega)\right|\right|^2 dp(\omega) \tag{1}$$

Here X is the original (multivariate) random variable, X' is the discretized random variable used to approximate X, and $L_X(X')$ is the loss function or error induced by approximating X as X'. The right-hand side of the equation invokes the measure-theoretic definition of a random variable. A random variable is defined as a function $X: \Omega \to \mathbb{R}$ where Ω is a sample space consisting of possible events. Our multivariate random variables are vectors of univariate random variables $X = (X_1, X_2, \dots)$, $X' = (X_1', X_2', \dots)$. The loss function or approximation error can be interpreted as the squared Euclidean distance between the true representation of an event $\omega \in \Omega$, $X(\omega) \in \mathbb{R}^n$, and its approximated representation after discretization $X'(\omega) \in \mathbb{R}^n$, integrated over the space of events Ω with probability measure p.

In the case that X is a continuous original random variable, it can be approximated with one constructed from an arbitrarily large random sample, resulting in Equation 2. Note that in the case that $X$is already discrete Equation 2 holds with equality.

$$L_X(X') \approx \sum_{\Omega} ||X(\omega) - X'(\omega)||^2 p(\omega) \tag{2}$$

The objective of optimal sampling discretization is to select X' so as to minimize the approximated loss function. This is achieved by performing weighted k-means clustering of $X(\omega)$ and setting $X'(\omega)$ equal to the centroid of the cluster containing $X(\omega)$. This follows from equation 2. Equation 3 expresses Equation 2 in terms of outcomes $x = X(\omega)$ and set $X'(\omega) = \mu_i$ where i is selected such that $X(\omega) \in S_i$ where S_i is the cluster containing $X(\omega)$. It shows that the (approximated) value of our loss function from equation 2 is exactly equal to the within-cluster variance which is minimized by observation-weighted k-means clustering as we see in equation 3.

$$L_X(X') \approx \sum_{i=1}^{k} \sum_{x \in S_i} \left|\left|x - \mu_i\right|\right|^2 p(x) \tag{3}$$

Note that k is the number of clusters or discrete values of X', which is set a-priori based on computational constraints. While k-means clustering algorithms guarantee convergence only to locally optimal clusterings [31], repeated optimization with randomly initialized centroids ensures results which are close to globally optimal.

Weaknesses of Optimal Sampling on Boundary Conditions

The most significant weakness of the proposed optimal sampling discretization as implemented in the preliminary pilot study of the Amite River is that optimal sampling on boundary condition features i.e., surge and discharge information, is not necessarily the same as optimal sampling on peak water surface elevation which is ultimately the hazard of interest. This issue is described in more detail and greater generality below in section 2.2.3 subheading "What to Sample on". Put briefly, HEC-RAS is a deterministic simulation, so a discretization of the hazard distribution which induces no error in the distribution of boundary conditions would similarly induce no error in the distribution peak water surface elevations. While a discretization which induces no error is of course impossible both for fundamental reasons and because the features of surge and discharge behavior used for discretization are of much lower-dimension than the full spatially explicit time series used as boundary conditions, this leads us to believe that a discretization which performs well in minimizing error in the distribution of boundary conditions will similarly perform well in minimizing the error in peak water surface elevation. An additional limitation of this method in practice is that there is no observation at the exact centroid of each cluster, so the observation nearest each cluster centroid is used instead.

Pre-Processing and Implementation

Pre-processing HEC-HMS output for optimal sampling discretization begins with extracting peak discharge, runup time, and drawdown time from each of the four major inlets to the HEC-RAS domain as shown in Figure 1, lag time between time of peak surge and peak discharge, as well as peak average surge depth, runup time, and drawdown time among representative points within the study domain. Discharge runup times are calculated by treating the discharge from the time at which discharge first exceeds its mean value over the hydrograph up to the time of peak discharge as the left half of a Gaussian density function and calculating the

corresponding standard deviation. More specifically, runup time σ_l is calculated as shown below, as a function of discharge $Q(t)$, average discharge of the hydrograph $\bar{Q}$, time t, and time of peak discharge t_{peak}.

$$\sigma_l = \frac{\sqrt{\int_{\overline{Q<Q(t),t<t_{peak}}} t^2 Q(t) dt}}{\sqrt{\int_{\overline{Q<Q(t),t<t_{peak}}} Q(t) dt}}$$

Drawdown times were similarly calculated from the time of peak discharge up to the point at which discharge receded below its mean value over the hydrograph. Surge runup and drawdown times were calculated in the same way using a surge hydrograph with values averaged over a random sample of locations within the study domain. A log transformation was applied to peak, runup, and drawdown of discharge and surge due to pronounced skewness, although this is not recommended in future analyses as it reduced the relative weight of extreme events in clustering. All features were then standardized to have mean zero and standard deviation equal to 1. Events were heuristically observation-weighted according to the CLARA-derived probability mass of their corresponding synthetic storms by use of repeated observations, permitting us to treat the set of events as a random sample of equiprobable events. More specifically, each event indexed by i was repeated in the dataset n_i times as shown below, where p_i is the probability of the event and p_{min} is the smallest probability of any event.

$$n_i = \text{round}(\frac{p_i}{p_{min}})$$

Following extraction and standardization, principal component analysis was applied to the resulting dataset. The original intention of using principal component analysis was to permit small components to be dropped from analysis in case clustering runtimes or memory requirements were overly large, but this proved unnecessary. It was however helpful in holistically visualizing the performance of the sampling approach.

From this point the dataset was clustered and discretized such that the observation nearest the centroid of each cluster was assigned the summed probability mass of observations in the respective cluster. Several synthetic storms were unrepresented in the resulting discretization. In the preliminary pilot analysis this led to an adjustment referred to as an "orphan storm correction" which replaced certain cluster centroids with nearby events from unrepresented synthetic storms

so as to minimize additional error induced by the adjustment. This assured that each synthetic storm in the ADCIRC+SWAN ensemble was represented at least once in the subset of events run through HEC-RAS. However, the use of this orphan storm correction is not recommended for future applications.

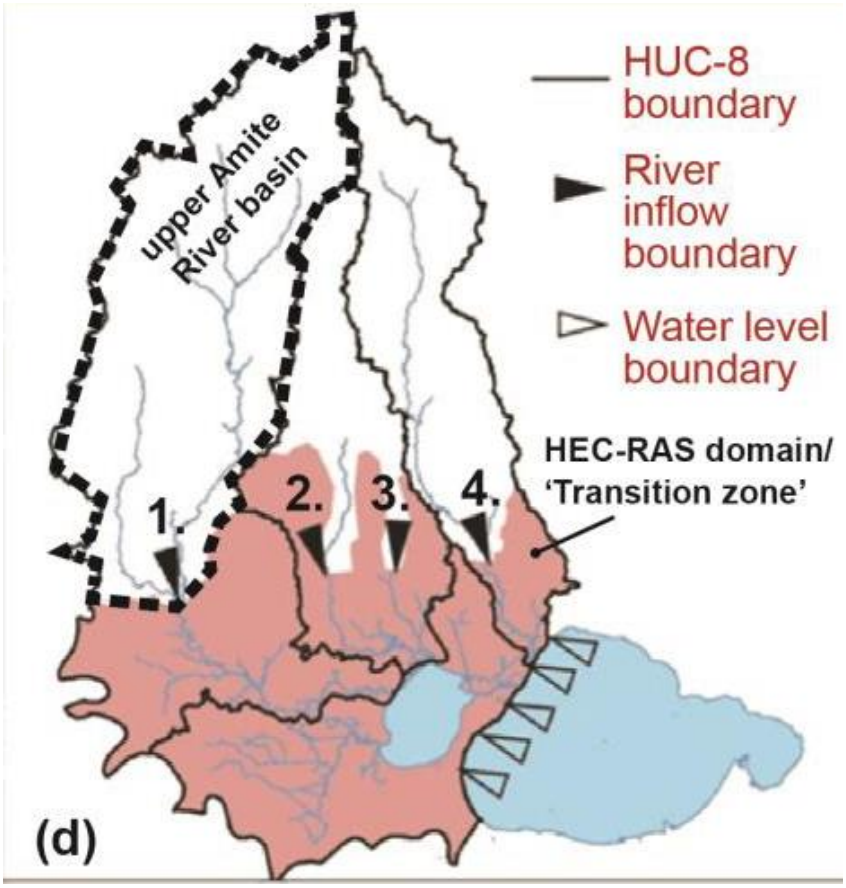

Figure 1: An illustration of the HEC-RAS model domain and its four major inlets [32].

2.4.3 Alternative Implementations of Optimal Sampling Discretization

Many aspects of the EJPM-OS implementation used in the preliminary pilot analysis of the Amite River Basin were the result of operational constraints which may or may not exist in other applications. In the interest of supporting future implementations of EJPM-OS and future uses of optimal sampling discretization in other frameworks, the section below describes below various ways in which optimal sampling discretization can be refined and used in other contexts.

Sample Sizes

Decisions about sample sizes in the preliminary pilot analysis of the Amite River Basin were largely driven by computational constraints. The preliminary pilot analysis used an optimal subsample size of 200 events because that was the number of HEC-RAS simulations that external partners could run with available resources. Computational constraints similarly governed the number of stochastic rainfall fields used. In project contexts where computational constraints are less binding, it is preferred to set these values with some empirical justification.

Computational constraints were substantially relaxed entering the (ongoing) revised pilot analysis for the Amite River Basin. Analysis was performed to identify 1) the number of stochastic rainfall fields needed for each synthetic storm in order to approximate the full variability of hydrologic responses at the level of individual storms, and 2) the number of clusters needed in optimal sampling discretization to would approximate the full variability of hydrologic response if sampling were performed at the level of individual storms. To accomplish this, external collaborators generated 1500 stochastic rainfall fields for each of three historical hurricanes—Hurricane Matthew, Hurricane Rita, and Hurricane Isaac—and distributed the rainfall fields for each historical hurricane equally among three antecedent conditions cases representing the 25^{th}, 50^{th}, and 75^{th} percentile of soil moisture levels. For each historical hurricane, I evaluated the rate at which random and optimal subsamples converged to the empirical distribution of the full sample. Results are shown in section 2.3.3.

One-Stage vs two-stage sampling

The optimal sampling discretization approach as implemented in the preliminary pilot study was effectively a two-stage sampling approach. Reducing CLARA's 645 storm set down to 50 based only on surge-wave hazard acted as the first stage, and optimal sampling discretization for compound hazard was used as a second stage. Where possible, it is recommended to use a single-stage sampling approach; if applied in the EJPM-OS implementation of the preliminary pilot, that would mean skipping the first sampling stage entirely and generating rainfall fields for all 645 storms prior to optimal sampling. In cases where single-stage sampling is infeasible e.g., if a single-stage sampling design would be prohibitively computationally expensive to run through HEC-HMS, a two-stage sampling approach is possible where both stages utilize optimal sampling

discretization for compound hazard. If applied to the EJPM-OS implementation in the preliminary pilot, that would mean a first stage of generating and sampling on a small number of rainfall fields for each of the 645 synthetic storms in the full storm set, followed by a second stage of sampling with a large number of rainfall fields for each synthetic storm selected in the previous stage.

Feature-Weighting in Pre-Processing

The outcome of optimal sampling discretization is sensitive to the choice of features to cluster on and the relative weights applied to them in pre-processing. The pre-processing step implemented for optimal sampling discretization in the preliminary pilot analysis did not explicitly consider the relative weights of the features used for clustering. 16 features were used: peak discharge and rates of runup and drawdown for four locations, time between peak discharge and peak surge depth, and runup and drawdown rates of surge depth. Because all features were standardized to have standard deviation equal to 1, and the scalar variance of a multivariate random variable is the sum of the variances of its component variables, the variance attributable to discharge behavior was four times greater than the variance attributable to surge and carried proportionally greater weight in clustering. It is likely that this is partially responsible for the absence of several synthetic storms in the optimal sampling discretization prior to orphan storm correction in the preliminary pilot analysis. Analysists can directly assign a weight w_i to each feature x_i used in clustering by multiplying each feature x_i by $\sqrt{w_i}$ after standardizing.

What to Sample on

As mentioned in section 2.2.2, subheading "Weaknesses of Optimal Sampling on Boundary Conditions", the most significant weakness of optimal sampling discretization as implemented in the preliminary pilot study is that optimal sampling on boundary condition features i.e., surge and discharge information is not necessarily the same as optimal sampling on peak water surface elevation, which is ultimately the hazard of interest. More generally, there is a multivariate random variable $X(\omega)$ representing inputs and/or forcings of the hazard outcome of interest $Y(\omega)$ associated with the random event ω, where in this application $X(\omega)$ is upstream hydrologic forcing and downstream surge forcing, and $Y(\omega)$ is the spatially varying peak water surface elevation. There is a function f such that $f\big(X(\omega)\big) = Y(\omega)$, in this case a HEC-RAS model, which

is computationally expensive to evaluate, hence our approximation $X(\omega) \approx X'(\omega)$ where $X'(\omega)$ is the discretized approximation of $X(\omega)$ via optimal sampling. Optimal sampling discretization minimizes the integrated square difference between X and X', but in an ideal world it would be preferable to minimize the integrated square difference between Y and $Y' = f(X')$. Because f is a deterministic function (here specifically HEC-RAS is a deterministic simulation), a discretization which induced no error at all such that $X' = X$ would similarly induce no error in approximating peak water surface elevations such that $f(X') = Y$. This provides reason to believe heuristically that minimizing error in approximating X should perform well in approximating Y, but due to nonlinearity in f this does not guarantee that error is minimized in approximating Y. To guarantee that error is minimized in approximating Y would require direct evaluation of $f(X)$.

Generally speaking, if it is possible to evaluate $f(X)$ then there's no need for optimal sampling in the first place, but applications may arise in which optimal sampling discretization on X is not the best approach. As with the optimal subsampling approach used in CLARA, there may be cases, particularly involving decision-making for flood risk mitigation, that it is feasible to evaluate $f(X)$ directly under existing climatological and landscape conditions, but not under the full set of possible future conditions and possible flood risk interventions under consideration. In this case, it may be preferable to directly evaluate $Y = f(X)$ under existing conditions, perform optimal sampling discretization directly on Y to arrive at Y' such that the integrated square difference between Y and Y' is minimized, and use Y' to characterize hazard under the full set of future conditions and flood risk interventions. This is most appropriate under slowly changing landscape conditions, and may suffer in performance under large structural changes that drastically affect the landscape's hydrodynamics There may also be cases where a metamodel or coarser-resolution hydrodynamic model $\hat{f}$ is available, which can inexpensively estimate $\hat{Y} = \hat{f}(X)$ such that $\hat{Y} \approx Y$ with sufficient fidelity to be useful in sampling but insufficient fidelity to be used directly for hazard estimation. In this case, analysts may wish to perform optimal sampling discretization to identify $\hat{Y}' \approx \hat{Y}$, and then simulate the selected events with the more expensive full-fidelity hydrodynamic model f such that the final hazard estimate is $Y' = f\left(\hat{f}^{-1}\left(\hat{Y}'\right)\right)$. This latter approach is particularly salient in cases where f is a hydrodynamic model incorporating rain-on-grid and does not rely on hydrologic model output for upstream boundary conditions, as sampling directly on rainfall, which is mechanically further removed from final water surface

elevations than is hydrologic discharge, may perform worse than sampling on hydrologic discharge.

2.5 Results

The preliminary pilot analysis of the Amite River Basin conducted for the Louisiana Watershed Initiative was, as said above, a preliminary analysis for the sake of methods development. Several methodological details of the preliminary pilot, some described above and others relating to the rainfall generator and hydrodynamic modeling, were later deemed to warrant iterative improvement prior to the publication of authoritative flood hazard estimates in the revised pilot analysis. At the time of writing, the revised pilot analysis is not yet complete, and its methodological improvements are out of scope for this document. For these reasons, nothing resembling a flood map will be shown here. In lieu of flood hazard estimates, presented below are indicators of the performance of optimal sampling discretization in approximating the previous distribution of hydrological and surge behavior, as well as the impact of optimal sampling discretization on the distribution of surge hazard. Also shown below are results of the analysis described in section 2.2.3 subheading "Sample Sizes", which investigates the relationship between sample size and performance of random and optimal sampling of stochastic rainfall fields at the level of individual tropical cyclones.

2.5.1 Fidelity of Optimal Sampling Discretization to the Original Distribution in the Preliminary Pilot Analysis of the Amite River Basin

The implementation of optimal sampling discretization applied in the preliminary pilot analysis of compound flood hazard in the Amite River Basin involved several implementation details which warrant refinement in future analyses, including the relative under-weighting of surge behavior and the use of a log transformation which reduced the effective weight of extreme events. In spite of this, the optimal sampling discretization performed surprisingly well in approximating a distribution characterized by 37,500 events (50 synthetic storms, 50 rainfall fields, five bias correction factors, and three antecedent conditions cases) in 16 dimensions (peak surge, surge runup and drawdown, lag time between peak surge and peak discharge, and peak discharge and discharge runup and drawdown for four inlets). Figure 1 shows the cumulative distribution functions of the first principal component of the sample and of peak discharge at the largest inlet

to the HEC-RAS domain. The optimal subsample matches the distribution of the first principal component of the original sample almost exactly. The optimal subsample appears to underestimate peak discharge for extreme events. This is expected to improve in future analyses when the log transformation is no longer applied in pre-processing.

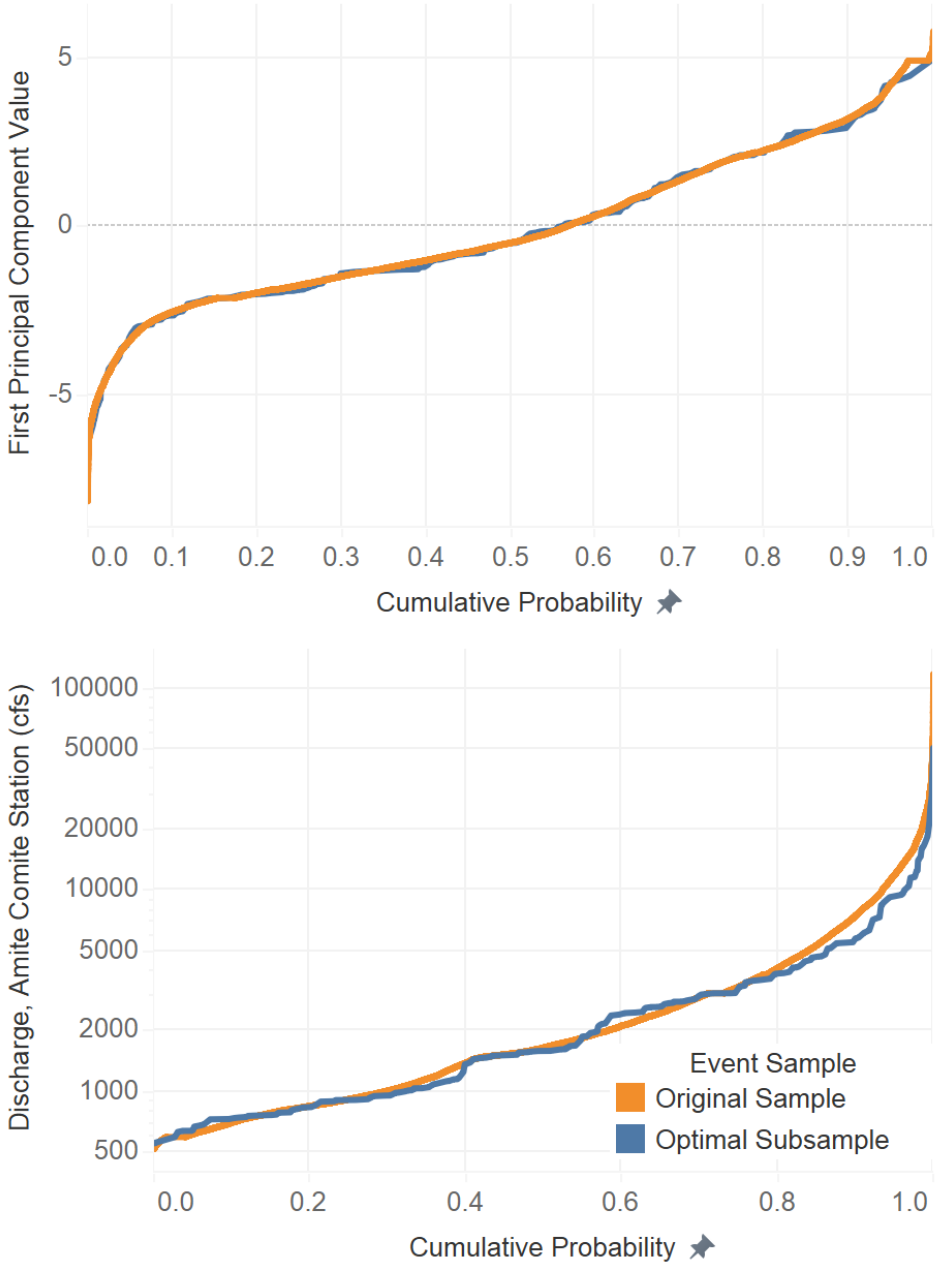

Figure 2: Cumulative distribution functions of first principal component value and peak discharge at the largest inflow to the HEC-RAS domain along the Amite River, characterized by the original sample and the optimal subsample [32].

2.5.2 Agreement in Coastal Hazard Between Optimal Sampling Discretization and Original Distribution

Optimal sampling discretization will typically result in clusters containing storm events generated from more than one synthetic storm. This results in changes to the total probability masses assigned to each synthetic storm in the optimally sampled distribution compared to the original distribution. To investigate the magnitude of this effect, comparisons were made between the hazard distribution estimated by the EJPM-OS statistical model with surge-only hazard simulated by ADCIRC + SWAN and the hazard distribution estimated by the original CLARA model with the same set of synthetic tropical storms and the same ADCIRC + SWAN simulations. The results are shown in Figure 3. They show broad agreement between the methods at the 10- and 100-year return periods, although the EJPM-OS-derived probability masses does underestimate the 50-year flood depth by about a foot in a section of the model domain. It is likely that this performance would improve with a greater relative weight placed on surge features in optimal sampling discretization and would also improve with a one-stage sampling design with a more diverse set of synthetic tropical cyclones.

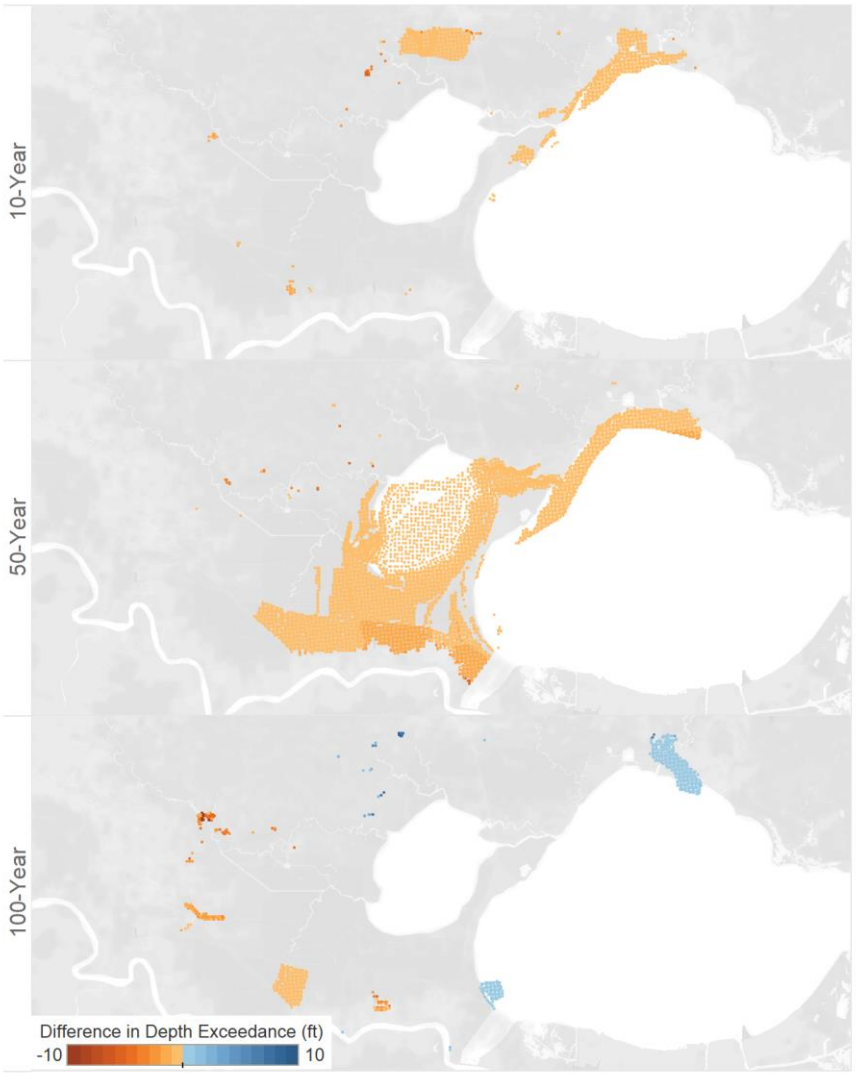

Figure 3: Difference between surge-only flood hazard estimated by EJPM-OS as implemented in the preliminary pilot analysis and previous CLARA methods, expressed in feet at the 10-, 50-, and 100-year return periods. Only pixels with a difference of 6 inches or greater are shaded [32].

2.5.3 Sample Sizes and Performance

As described above in sample 2.2.2 subheading “Sample Sizes”, at the outset of the revised pilot analysis there was a need to characterize the rate at which empirical distribution of a random sample of stochastic rainfall fields of a given hurricane converges to the true distribution, which can be reasonably approximated with a very large sample. There was a similar to know the rate at which the optimal sample discretization of the empirical distribution of a large sample converges

to the empirical distribution of said large sample. To help us accomplish this, external partners generated 1500 stochastic rainfall fields for each of three historical hurricanes—Hurricane Matthew, Hurricane Rita, and Hurricane Isaac—and distributed the rainfall fields for each historical hurricane equally among three probability-weighted antecedent conditions cases representing the 25^{th}, 50^{th}, and 75^{th} percentile of soil moisture levels. These were then run through HEC-HMS. In order to examine the sensitivity of the convergence rates of interest to the dimensionality of data extracted from HEC-HMS simulations, all analysis was performed on four sets of extracted data. One contained only peak discharge information from virtual gages at the four major inlets to the HEC-RAS domain. One contained peak discharge information from an additional six upstream virtual gage locations for a total of 10 virtual gage locations. The remaining two used the same sets of gage locations, but also contained discharge runup rate, discharge drawdown rate, and time of peak discharge, with peak discharge—deemed by external partners to be the most salient feature—scaled (after standardization) by a factor of $\sqrt{3}$ so as to account for as much variance as runup, drawdown, and peak time combined. For the sake of brevity, only results using peak discharge from four gage locations and results including runup, drawdown, and peak time from ten gage locations are shown. The convergence rates of the intermediate datasets are consistent with those shown.

It was straightforward to calculate and express the convergence rates of optimal sampling, as the total within-cluster sum of squares of a clustering divided by the total sum of squares reflects the portion of variance captured by the clustering on the interval [0, 1]. Figures 3 and 4 show the captured variance as a function of number of clusters for the extraction of peak discharge at four gage locations and the extraction of peak discharge, runup, drawdown, and peak time of 10 gage locations respectively. Non-monotonicity in these plots is due to random initialization of the k-means clustering algorithm. The figures show the best clusterings out of 100 randomly initialized executions of the k-means clustering algorithm in line with the use of random restarts in practical applications of the method. It is clear from these figures that as the dimensionality of data to be clustered increases the number of clusters necessary to reach a given level of performance increases. Tens of clusters are easily enough to capture 90% of the variance of the distribution of peak discharges among four gage locations, but hundreds are required to capture a similar portion of variance when including temporal dynamics among ten gage locations. This performance must be weighed carefully against computing budgets on the basis of expert judgement. It bears noting

here that this performance is measured at the level of individual storms. An optimal sampling discretization carried out over a set of diverse synthetic tropical storms with a fixed number of clusters per synthetic storm will necessarily yield equal or greater performance when compared to an optimal sampling discretization with the same number of clusters per synthetic storm which only permits clustering at the level of individual synthetic storms, as the latter represents a constrained special case of the former. For this reason, the performance shown below represents a lower bound on the performance of optimal sampling discretization as applied in EJPM-OS rather than a direct measure. The difference between the performance of optimal sampling discretization at the level of individual storms versus over a diverse set of synthetic storms will be investigated in future work.

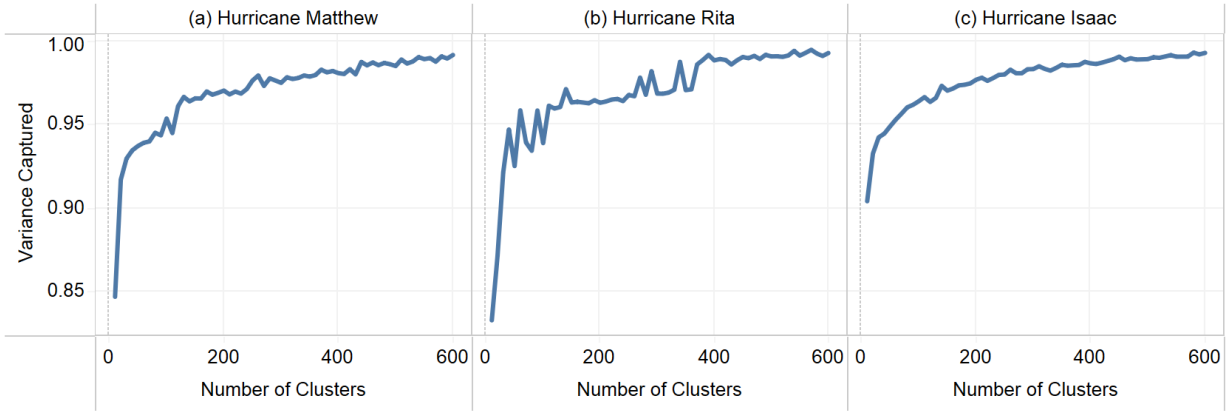

Figure 4: Variance captured in optimal sampling as a function of number of clusters for three historical hurricanes, using peak discharge from four model domain inlet gage locations.

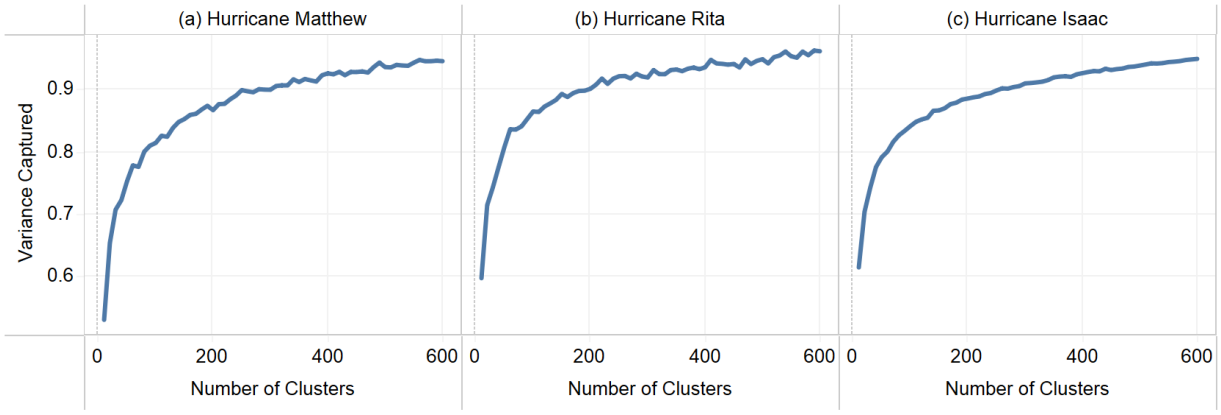

Figure 5: Variance captured in optimal sampling as a function of number of clusters for three historical hurricanes, using peak discharge, runup rate, drawdown rate, and time of peak surge from ten modelled gage locations, with peak discharge scaled by $\sqrt{3}$ after standardization.

In order to quantify the convergence of random sampling in terms that are directly analogous to the results shown for optimal sampling, this analysis used a heuristic measure which expresses the portion of variance of extracted HMS data which is captured by a random sample. This heuristic treats each random sample as the final set of centroids of a k-means clustering algorithm, assign each observation in the full dataset to the "cluster" of each element of the random sample, and calculate the portion of variance captured in line with what is shown above for optimal sampling. Results shown are the average performance over 100 replicates for each sample size. Because in EJPM-OS rainfall fields are typically generated randomly and assigned in equal number across antecedent conditions cases even where antecedent conditions cases are probability-weighted, random samples were selected without replacement equiprobably from the three antecedent conditions cases. The pseudo-clustering performance shown reflects the variance captured relative to the probability-weighted empirical distribution. Figures 5 and 6 show the performance of random sampling over peak discharge at four inlet gage locations and over peak discharge, runup rate, drawdown rate, and time of peak discharge at ten gage locations with peak discharge scaled by $\sqrt{3}$ after standardization. Consistent with our results for optimal sampling, the performance of random sampling degrades substantially as more gage locations as well as temporal dynamics are considered. These results also show greater heterogeneity in the performance of random sampling between historical hurricanes, indicating that these results are less consistently generalizable. Additionally, it is unclear whether the large amount of variance captured for very

large sample sizes actually indicates convergence to the true distribution, or whether it is simply a result of approaching the sample size of the full dataset containing 1500 randomly generated rainfall fields. Again, great care must be taken to balance the benefits of larger numbers of rainfall fields against computational constraints. These results may be utilized for future EJPM-OS implementations to inform decisions between one-stage and two-stage optimal sampling designs and, in the case of two-stage sampling designs, how to balance the number of synthetic storms to select in the first stage against the number of stochastic rainfall fields to generate per selected storm in the second stage.

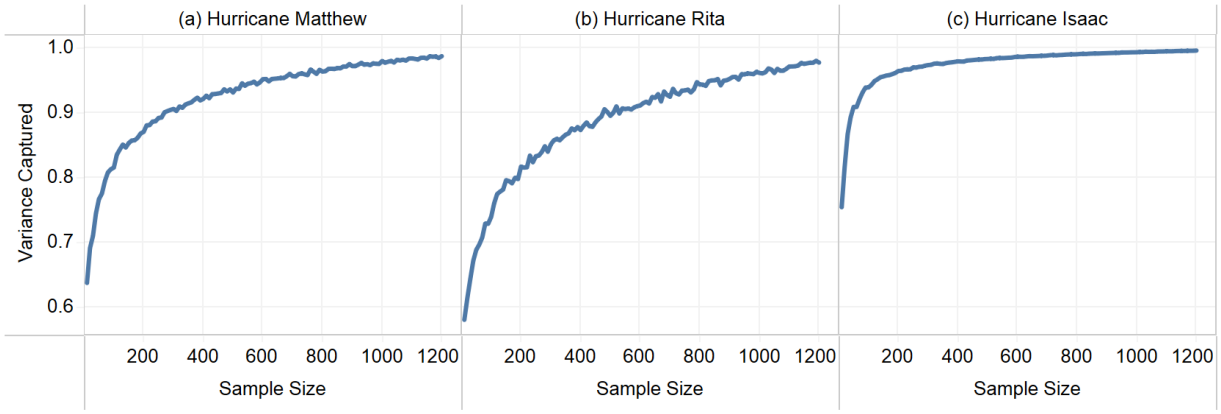

Figure 6: Variance captured in random sampling as a function of sample size for three historical hurricanes, using peak discharge from four model domain inlet gage locations.

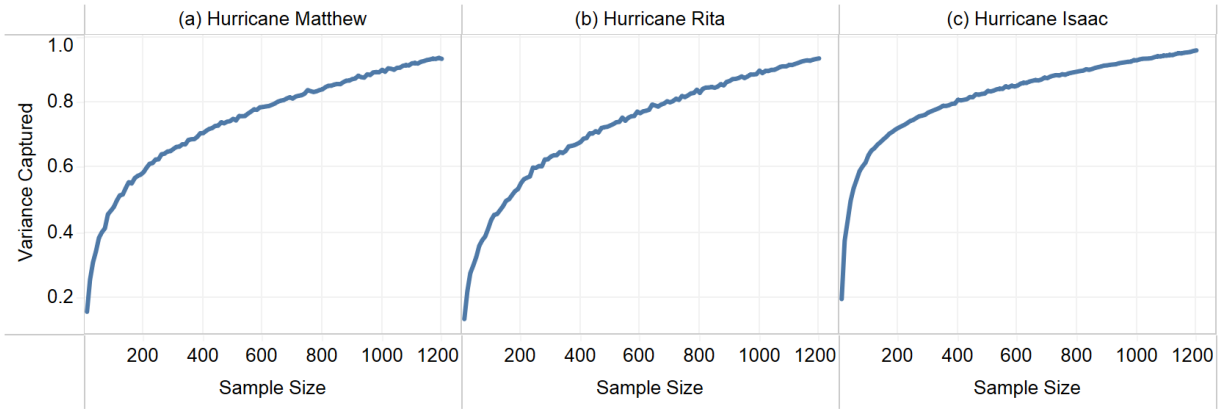

Figure 7: Variance captured in random sampling as a function of sample size for three historical hurricanes, using peak discharge, runup rate, drawdown rate, and time of peak surge from ten gage locations, with peak discharge scaled by $\sqrt{3}$ after standardization.

2.6 Conclusions

The extended joint probability method with optimal sampling (EJPM-OS) represents a novel approach for statistical modeling of compound coastal flood hazard. Preliminary implementation has shown good performance over several measures. Methodological changes for ongoing efforts of the Louisiana Watershed initiative will continue to iterate on and improve upon this approach. Future research will seek to further refine these methods, better quantify the performance of the method, and study how said performance varies as a function of various implementation decisions.

The optimal sampling discretization procedure used in EJPM-OS is highly generalizable to statistical characterization of natural hazards where the outcome of interest of a random event is calculated with a computationally expensive model, and for which boundary conditions or lower-fidelity estimates (via coarser model structure or metamodeling) can be produced more efficiently. It can be applied directly in cases where the distribution of events is represented as an empirical distribution or random sample, and it can be applied to a large Monte Carlo sample of an arbitrary joint distribution without requiring any assumptions about the variance structure of the hazard. The optimality guarantee associated with this approach, that of minimizing integrated square error in discretization, is highly appropriate from the perspective of viewing natural hazards through the lens of multivariate random processes.

2.7 Chapter Acknowledgements

I acknowledge the support and guidance of Louisiana state agency personnel from the Council of Watershed Management supporting the Louisiana Watershed Initiative (LWI), as well as the LWI Data and Modeling Transition Zone Technical Advisory Group

3. REAL OPTIONS ANALYSIS OF ADAPTIVE PLANNING IN STRUCTURAL AND NONSTRUCTURAL FLOOD RISK MITIGATION

Coastal flooding associated with tropical cyclones presents a severe and worsening hazard [1]. A variety of interventions are available to combat this hazard, which can be roughly categorized as structural and nonstructural protection. Structural protection is defined as the set of interventions which influence the probability distribution of flood depths in a given location (e.g., levee construction). Nonstructural protection is defined as the set of interventions which influence the damage associated with a given flood depth e.g. elevation-in-place retrofits to individual buildings [6]. Structural and nonstructural protection therefore interact nonlinearly and require a certain minimum model complexity to account for both, which makes generation of efficient solutions a challenging problem in its own right. The problems of generating and comparing strategies combining structural and nonstructural are made more complicated by "deep" uncertainty. Deep uncertainties are those which cannot be confidently ascribed probability distributions due lack of knowledge or disagreement between stakeholders e.g. the severity and timing of specific impacts of climate change [33]. Failure to account for deep uncertainties in decision making leaves decision makers vulnerable to regret. A flood risk mitigation strategy designed for a moderate climate future is subject to regret in the form of wasted investment in a mild climate future and regret in the form of failure to prevent flood damages in a severe climate future. However, explicitly adaptive and flexible strategies have the potential to reduce regret in decision-making under deep uncertainty in a variety of management contexts [34]–[36]. A solution which performs well under a wide range of deeply uncertain conditions is considered "robust. This chapter utilizes an optimization-driven framework for decision-making under deep uncertainty which uses multi-objective optimization to generate robust solutions which balance system performance over multiple climate futures. It leverages framework to characterize the extent to which adaptive flood risk management strategies combining structural and nonstructural protection measures are less vulnerable to regret than non-adaptive strategies. The intention of this work is to draw actionable conclusions to inform applied flood risk modeling efforts for applied policy decision-making in how they generate and compare efficient flood risk mitigation strategies under uncertainty.

A wide range of flood risk models are available in the literature, but most have limited applicability to investigating adaptive combinations of structural and nonstructural protection. Such models are typically either too computationally burdensome for optimization or too simplified to be useful in an applied setting. They often assume a single elevation for a levee or ring levee system, ignore levee fragility and wave overtopping, assume a single value for flood depths or damage when flooding occurs, fail to explicitly model surge and wave behavior as a function of sea level rise, or assume stationarity in economic and climate conditions [37]–[46]. In the case of Louisiana's 2017 Coastal Master Plan, the Coastal Louisiana Risk Assessment (CLARA) Model was sufficiently detailed to characterize the impacts of structural and nonstructural protection but was too computationally burdensome for an optimization-driven experimental design. Instead, a relatively small set of structural protection projects were evaluated separately, with nonstructural projects evaluated conditionally upon selected structural projects [47]. In order to better explore the interactions between structural and nonstructural protection in an adaptive planning context, this analysis utilizes the recently developed Surge and Wave Model for Protection Systems (SWaMPS)—a faster, lower resolution, and spatially limited version of CLARA [48] that is configured to represent the Larose to Golden Meadow Hurricane Protection Project, a ring levee structural protection system in Lafourche Parish, Louisiana.

Generation, evaluation, and comparison of flood risk mitigation strategies under conditions of deep uncertainty is challenging because the rank-ordering of mitigation strategy performance varies between deeply uncertain future conditions and because analyst knowledge of those future conditions is non-probabilistic. One common approach to decision-making under deep uncertainty is Robust Decision-Making (RDM), which stress-tests strategies under a wide range of future conditions in order to determine which are robust [33]. The choice of how to measure robustness, however, is not trivial. A satisficing robustness metric would measure the fraction of modelled futures in which strategy performance exceeds a given threshold, a regret-based metric would measure the maximum or average regret over modelled futures, and a statistically based metric would use some descriptive statistic of strategy performance over modelled futures such as a coefficient of variation [49]. All such parameterized robustness metrics however carry the notable weakness that they make implicit probabilistic assumptions about deeply uncertain future conditions. To conclude that a strategy which performs adequately in 99% of modelled deep0ly uncertain futures is more robust than one which performs adequately in 1% of modelled futures is

epistemically unsound, as without probabilistic information it is not guaranteed that the modelled futures in the former 99% have a greater summed probability of occurrence than those in in the latter 1%. For this reason the RDM framework uses robustness analysis largely for the purpose of qualitatively understanding a strategy's potential shortcomings and augmenting it, or otherwise characterizing regions of uncertain parameter space in which strategies succeed or fail and evaluating them probabilistically in a post-hoc fashion [33]. Other approaches such as Adaptation Pathways, rather than simply stress testing a prespecified strategy or set of strategies, simulates deeply uncertain futures in order to identify branching pathways of successive interventions which maintain pre-specified minimum performance levels under divergent future conditions [50], [51]. Despite the efficacy of RDM in exploring the performance of policy actions under diverse deeply uncertain future conditions and the efficacy of Adaptation Pathways in generating robust and adaptive strategies, both frameworks applied directly require a structured, human-in-the-loop decision processes which are effective in exploring combinatorial strategies but likely struggle when seeking to optimize continuous decision variables such as the height of a levee or the amount of resources to devote to nonstructural protection.

In cases where optimization-driven decision processes are required, such as in cases where there is a need to optimize where cost-effectiveness or the net present value of policy decisions with continuous decision levers, other frameworks are available. Of particular note is Real Options Analysis, which structures adaptive strategies as decision trees similarly to Adaptation Pathways, optimizing average net present value of decisions over multiple decision points simultaneously [34], [35], [52]. This is particularly suited to the present decision context in which the generation and comparison of robust and adaptive combinations of structural and nonstructural flood risk mitigation is desired. In simultaneously optimizing short-term and long-term actions, Real Options Analysis compares the near-term benefits of immediate action which carry potential regret with the flexibility of delaying action and acting adaptively when more information is available which carries the opportunity cost associated with delayed action. It should be noted that while Real Options Analysis invokes the notion of and derives inspiration from the analysis of options in the sense of the financial derivate contracts, the "real" part of Real Options Analysis as applied to decision-making under deep uncertainty refers to the fact that the "options" analyzed are not in fact financial derivatives but rather optionality associated with a choice to delay action such as the construction of a reservoir in the case of water resource management or a levee in the case of flood

risk management. For this reason the analysis structure differs from that of financial options analysis and involves optimization of arbitrary system models via metaheuristic optimization routines such as genetic algorithms and does not typically involve e.g. the Black-Scholes equation [34], [35], [52]. Real Options Analysis does however carry the same weakness as univariate robustness measures described above in that it relies on probabilistic assumptions; even where lip service is paid to the non-probabilistic nature of deeply uncertain futures, the optimization of an average net present value of decisions amounts to an expected net present value with an implicit assumption of equiprobably futures. This approach applied directly, or any straightforward mathematical programming approach which optimizes the average of objective values over modelled futures is problematic. A strategy which dominates another even in a multiobjective sense when averaged over possible futures does not necessarily do so from the perspective of diverse groups of decision-makers or stakeholders who may hold differing subjective beliefs about the relative likelihood of possible futures and does not necessarily do so in expectation under a reasonable probabilistic model which might be formed if more information were available. This research therefore seeks an approach by which to generate and compare adaptive strategies which directly considers the value of flexibility and adaptivity as does Real Options Analysis, but which does not impose epistemically unsupported rank-orderings as does typical Real Options Analysis or the use of scalarized robustness metrics.

To arrive at an implementation of Real Options Analysis that is fully appropriate in any rank-ordering of strategies under deep uncertainty requires integration with another optimization-driven approach for decision-making under deep uncertainty. One such method is Many-Objective Robust Decision-Making (MORDM). MORDM involves selecting a baseline future, optimizing strategy performance in that baseline future in a multi-objective sense, and stress-testing the resulting Pareto frontier across a range of deeply uncertain futures as in RDM [53]. Applied as originally proposed, this framework does not lend itself well to adaptive strategies as the initial optimization is over a single modeled future. An extension of this approach referred to as multi-scenario MORDM repeats the optimization phase over several parameterized futures which are discovered in an initial MORDM analysis, but this approach is similarly inappropriate here as each optimization is performed on a single parameterized future [54]. A competing framework for optimization-driven generation of robust policy decisions is Robust Optimization (RO). Robust optimization typically optimizes the average value of objectives over a set of parameterized futures

while either setting constraints on a scalarized robustness measure or simultaneously optimizing average objective values and a scalarized robustness measure in a multiobjective sense [55], [56]. As discussed above, the use of scalarized robustness measures may lead to epistemically unfounded rank-orderings of solution robustness and is therefore problematic. However, recent work by Shavazipour et al. has proposed an alternative approach for the generation of robust solutions which does not make epistemically problematic comparisons between solutions under deep uncertainty, entitled Multi-Scenario Multi-Objective Robust Optimization (MORO-MS). This approach, by simply treating each objective of interest in each modelled future state of the world as a distinct objective, avoids epistemically problematic rank orderings by only considering a solution to be dominated in the case where it is dominated in every modelled future. It has been applied in a didactic case study of the shallow lake problem [57] and in a case study of forest landscape management [58]. No study has yet applied the principles of MORO-MS to an adaptive planning context or to flood risk mitigation. The application of Multi-Scenario Multi-Objective Robust Optimization to maximize the net present value of adaptive strategies over every modelled future in a multiobjective sense represents a natural extension of Real Options Analysis in that the resulting Pareto frontier trivially contains solutions which maximize net present value of solutions in every future as well as the solution which maximizes average net present value. To our knowledge, this is the first implementation of a method which uses an optimization-driven approach to generate adaptive strategies across a range of deeply uncertain future conditions while handling comparisons between strategies under deep uncertainty with full epistemic rigor.

3.1 Methods: Initial Experiment

This analysis uses the SWaMPS model and the NSGA-II multiobjective evolutionary algorithm [59] to identify adaptive flood risk mitigation strategies consisting of structural and nonstructural mitigation which simultaneously minimize the combined net present value of mitigation costs and expected damage across three sets of sea level rise futures—low, medium, and high—within each of which individual trajectories are treated as equiprobable. Expected damage is calculated over an 80-year period, and each set of sea level rise futures contains five distinct trajectories. This optimization is performed in an adaptive context and in a static context. The adaptive context considers two decision stages – one in 2015 at the start of the planning period and one in 2055 halfway through the planning period. At the time of the second decision stage in

the adaptive context, decision-makers have observed the degree of sea level rise which has occurred thus far and take it into account in their decision-making. This is represented in the optimization problem by three distinct sets of protection measures to take at the second decision stage: one which is applied in the "low" set of sea level rise futures, one which is applied in the "medium" set of futures, and one which is applied in the "high" set of futures. The static optimization context also simultaneously optimizes performance over the low, medium, and high set of future conditions, but makes the same decision in the second stage irrespective of observed sea level rise conditions. This allows us to characterize the extent to which the adaptive decision-making of the adaptive context can better balance performance over the sets of climate futures than the one-off decision-making of the static context.

Table 1 shows an XLRM chart [60] describing the overall problem structure. Our sea level rise trajectories are generated by Brick V0.3 over multiple Representative Concentration Pathways and calibration ensemble members [61]. These trajectories are selected by first generating sea level rise trajectories for each member of Brick V0.3's calibration ensemble for each of three available representative concentration pathway, adjusting them to reflect zero sea level rise off the coast of Louisiana at the start year for this study of 2015. The sea level rise trajectories are then ordered by their total sea level rise from 2015 to 2055 and select trajectories with evenly spaced indices spanning the entire set. For the sake of computational tractability, 15 trajectories are selected, such that each set of futures contains five. It is important to note that, as shown in Figure 8, due to nonlinearity in sea level rise trajectories (resulting from e.g., sudden ice sheet collapse), the lowest "high" trajectory has a 2095 sea level substantially lower than the highest "medium" trajectory and the lowest "medium" trajectory has a 2095 sea level substantially lower than the highest "low" trajectory. This captures a more realistic degree of uncertainty following observation than similar previous work using Real Options Analysis, wherein a less extreme near-term observation is assumed to always lead to a less extreme long-term climate trajectory [34], [36], [62]. While changes in storm frequency and intensity as well as uncertain population growth, economic growth, and discount rates are also salient uncertainties in an applied context, the initial experiment excludes frequency and intensity changes from our uncertainty ensemble because they are more difficult to observe directly with a high degree of confidence, and population change, economic growth, and discount rates are excluded in order to focus narrowly on a straightforwardly orderable continuum of hazard intensities.

Table 1: XLRM Chart of present study – adaptive context

eXogenous uncertainties: - Sea level rise trajectory, binned into 3 sets: high, medium, low	Levers: (Each lever listed has one value for an upgrade in 2015, and one in each of three sets of sea level rise trajectories in 2055) - Binary indicator for levee upgrade - Continuous investment intensity for levee upgrade - Nonstructural elevation standard (NAVD88 feet)
Relations: - SWaMPS - Mixed-lever NSGAII	Measures: - Sum of net present value of mitigation and expected damages, evaluated separately over high, medium, and low sea level rise trajectories.

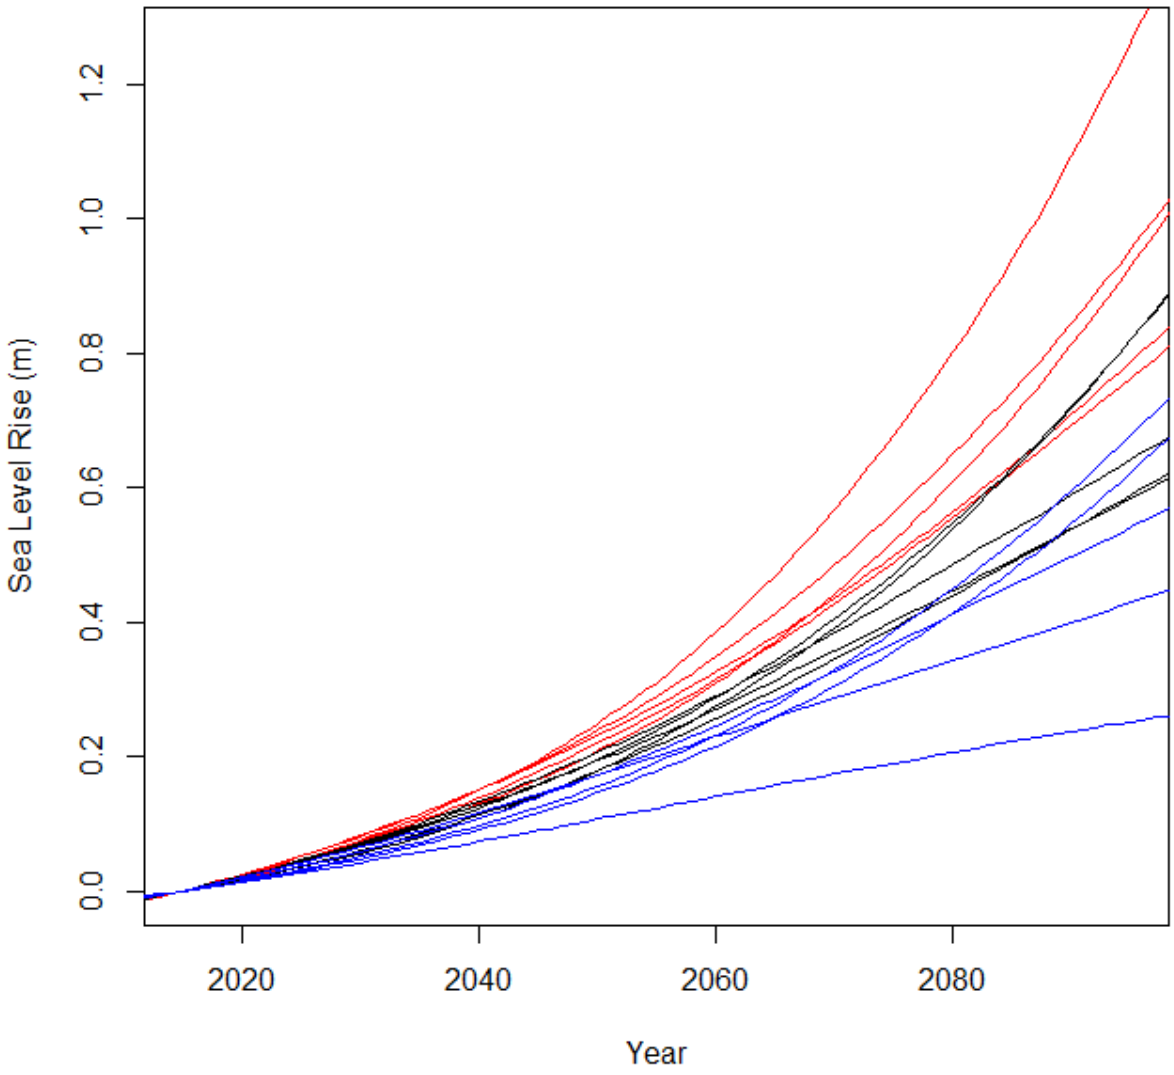

Figure 8: Sea level rise trajectories – high in red, medium in black, and low in blue.

The use of system performance across three sets of climate futures as objectives and a single second-stage system upgrade within each climate future is meant to be a tractable approximation of a more complex but intractable experimental design. Ideally, the number of sea level rise trajectories would be substantially larger. Due to the limited number of sea level rise trajectories used, the experiment as implemented groups sea level rise trajectories into three groups based on the ordering of sea level rise at the year 2055; if it were tractable to run many more sea level rise trajectories these would instead be grouped into bins based on sea level rise at 2055 with widths corresponding to the variance in statistical estimators of near-term mean sea level trends based on on-line observations. Each of these bins would have a set of decision variables corresponding to a second-stage system upgrade. This preferred experimental design would have one objective for each sea level rise trajectory. This more complex design, however, would be computationally intractable. The model currently takes approximately thirty seconds to generate a net present value of costs for a single strategy for a single sea level rise trajectory and as a result requires require three days to optimize over 15 sea level rise trajectories with 30,000 function evaluations when running in parallel on 50 cores available at the time the experiment was conducted. This limits the number of sea level rise trajectories, which limits the number of sea level rise bins which can be considered and assigned distinct decision levers as any given bin with a corresponding set of decision levers must reflect a realistic degree of uncertainty in future sea level rise following the second decision stage. Net present values of costs are averaged over sea level rise trajectory bins rather than treated as individual objectives for similar reasons; research has shown that the convergence of multi-objective optimization via multi-objective evolutionary algorithms slows as the number of objectives increases [63]. While this does mean that sea level rise trajectories within the low, medium, and high bins are treated as equiprobable, this reflects a decision to model deep uncertainty in sea level rise with a deeply uncertain ordinal hyperparameter, wherein the deeply uncertain state of the world characterized by a given value of that hyperparameter is itself described probabilistically; while this simplification reduces the realism of the experiment, it does not therefore reflect a failure to treat deeply uncertain system parameters non-probabilistically. Additionally, it would be preferred to permit second-stage protection decision to be made at an arbitrary time governed by an additional decision lever for each set of sea level rise trajectories, leading the structure of the "option" present in the experimental design to more closely resemble American options with respect to the financial derivative contracts. This

however would add dimensionality to the lever space and therefore slow convergence substantially, necessitating that the implemented experimental design limit second-stage mitigation decisions to a fixed time and more closely resemble European options with respect to the financial derivative contracts. The temporal resolution used in modelling reflects another simplification made for computational tractability: expected annual damages due to flooding are estimated via the SWaMPS model at years 0, 20, 40, 60, and 80 with discontinuities in expected damages over time at the midpoints between modelled years i.e. at years 10, 30, 50, and 70, with an additional discontinuity at year 40 where expected annual damage is calculated before and after the second-stage mitigation decision. Discounting is applied at 10-year intervals rather than annually or continuously to simplify the modelling codebase.

Optimization is handled via mixed-lever NSGAII implementation from the PyMoo package for Python [64]. While the SWaMPS model natively parameterizes structural mitigation as an increase in crest elevation to each of 12 levee reaches, permitting each reach crest height to vary independently in each upgrade stage would result in an unreasonably high-dimensional search space for optimization and would be unlikely to converge to the Pareto frontier in a reasonable number of function evaluations. Investment in structural protection upgrades is expressed as a dimensionless scalar, the parameterization of which is described below. Levee upgrades at each decision stage are controlled by a continuous, dimensionless levee height upgrade intensity and a Boolean indicator for whether an upgrade occurs at that decision stage. Using a continuous levee height upgrade without a Boolean indicator was considered, permitting a continuous levee height upgrade of zero to indicate the lack of an upgrade. However, extremely small levee height upgrades would have a proportionally small impact on flood damages but would still incur large overhead costs. The use of a Boolean indicator for levee upgrades was therefore deemed necessary to ensure that within the optimization, solutions with substantial levee upgrades would be discoverable from solutions without levee upgrades and vice versa. This was not considered necessary for nonstructural mitigation, as nonstructural mitigation does not have the same magnitude of overhead cost. Additional nonstructural mitigation applied at the second decision stage would carry additional marginal costs, as structures which had previously received improvements and had not been bought out would have to be upgraded again. This does not, however, represent the same degree of discontinuity in cost as is present in upgrades to structural protection The SWaMPS model natively supports modeling of nonstructural mitigation through a nonstructural protection

standard elevation measured in NAVD88 feet. Applicable buildings lying below this standard by zero to three feet are offered exterior floodproofing treatments. Applicable buildings lying below the standard by three to twelve feet are offered elevation-in-place retrofits. Those lying more than twelve feet below the standard are offered voluntary buyouts. The SWaMPS implementation of nonstructural protection follows that of the Coastal Louisiana Risk Assessment Model (CLARA) utilized in Louisiana's 2017 Coastal Master Plan [65]. The effectiveness of nonstructural mitigation in SWaMPS is modulated by a user-determined participation rate parameter. Participation rate in nonstructural mitigation is assumed to be 100% for the purposes of this study to best highlight the impacts of nonstructural mitigation and the interaction thereof with structural protection.

As described above, it was desirable to parameterize levee height upgrades with a single continuous scalar value rather than with a separate value controlling the upgrade magnitude for each reach. Otherwise, optimization would be unlikely to converge on an approximate Pareto frontier within a reasonable number of function evaluations. For this reason, the PyMoo implementation of NSGAII was used to compute a Pareto frontier of expected annual damage versus structural upgrade cost with one meter of sea level rise, while permitting the increase in levee crest height of each reach to vary continuously. The population size was left at the default value of 100. The resulting solution set along this approximate Pareto frontier was then sorted by cost. The least costly levee upgrade was given the index 0, and the most costly levee upgrade was given the index 99. In our static and adaptive optimization contexts described above, the dimensionless value of structural upgrade intensity is used to index the structural configurations along this Pareto frontier, with non-integer values resulting in a linear interpolation between two contiguous configurations. This initial optimization can be expressed:

$$Minimize < f_{EAD}(\boldsymbol{x}, \boldsymbol{c}, n), f_{cost}(\boldsymbol{x}, n) >, \boldsymbol{x} = < x_1, x_2, \dots, x_{12} >$$

$$s.t.$$

$$x_i \in [0,3.5]\ i \in \{1, \dots, 12\}$$

Here x_i refer to continuous levee height upgrade magnitude in meters for each of 12 levee reaches, f_{EAD} refers to the expected annual damage with one meter of sea level rise given the specified levee height upgrade without any nonstructural investment, and f_{cost} refers to the cost of said levee height upgrade. Both f_{EAD} and f_{cost} are calculated with the SWaMPS model, the complexity of which (due to the complexity of overtopping processes, rainfall and pumping, and aggregation of

flood depths calculated at the level of individual storms) precludes presentation of closed-form parameterizations here. The structure of the SWaMPS model is described in greater detail in Appendix A below, as well as in its original publication by Johnson et al. [48]. All stochastic components of the SWaMPS model are reduced to a deterministic sample path problem with 25 Monte Carlo iterates from a fixed initial random seed. Here $\boldsymbol{c}$ refers to climatological parameters, in this case set to reflect no change in intensity and frequency of tropical cyclones compared to the historic record and one meter of sea level rise compared values observed in the year 2015. Here n refers to a fixed nonstructural mitigation standard reflecting no nonstructural mitigation, equal to the -13 feed NAVD88 below which no structures exist in the system. Multi-objective optimization is caried out via a mixed-lever implementation of the NSGA-II genetic algorithm in the PyMoo package for Python [64]. Bounds on levee height upgrades are imposed for physical realism and implemented as "bounds" as described in PyMoo documentation, with initial population sampling and mutation implemented in PyMoo so as to respect bounds. As described above, the resulting Pareto frontier is used to construct a transformation $\boldsymbol{M}(\cdot)$ which converts a dimensionless scalar parameter governing structural upgrade magnitude to a twelve-dimensional real vector of levee upgrade heights in meters. Using a single continuous index along the resulting Pareto frontier to parameterize structural investment intensity, the optimization problem posed by the adaptive experiment can be expressed:

$$Minimize < f_1(\boldsymbol{x}), f_2(\boldsymbol{x}), f_3(\boldsymbol{x}) >, \boldsymbol{x} = (x_{sb,0}, x_{sm,0}, x_{ns,0}, \dots, x_{sb,3}, x_{sm,3}, x_{ns,3})$$

$$s.t.$$

$$x_{sb,i} \in \{0,1\}\ \forall\ i \in \{0,1,2,3\}$$

$$x_{sm,i} \in [0,99]\ \forall\ i \in \{0,1,2,3\}$$

$$x_{ns,i} \in [-13, 15], \forall\ i \in \{0,1,2,3\}$$

$$x_{sm,i} \geq x_{si,0}\ \forall\ i \in \{1,2,3\}$$

$$x_{ns,i} \geq x_{ns,0} - 0.5\ \forall\ i \in \{1,2,3\}$$

Here f_1, f_2, and f_3 refer to the average net present value of strategy costs under low, medium, and high sea level rise futures respectively, as expressed below:

$$f_s(\boldsymbol{x}) = \frac{1}{|K(s)|}\sum_{k\in K(s)}\left(\left(\sum_{t\in T_1} f_{ead}\left(\boldsymbol{M}(x_{si,0}), \boldsymbol{c}_{k,t}, x_{ns,0}\right) * \left(\frac{1}{1-\delta}\right)^t * \tau \right.\right.$$
$$\left. + \sum_{t\in T_2} f_{ead}\left(\boldsymbol{M}(x_{si,s}), \boldsymbol{c}_{k,t}, x_{ns,s}\right) * \tau * \left(\frac{1}{1-\delta}\right)^t \right) + f_{cost}\left(\boldsymbol{M}(x_{si,0}), x_{ns,0}\right)$$
$$\left. + f_{cost}\left(\boldsymbol{M}(x_{si,s}), x_{ns,s}\right)\left(\frac{1}{1-\delta}\right)^{t^*} \right)$$

Here, k is used as an index over the set of sea level rise trajectories $K(s)$ associated with a given set of sea level rise trajectories (low, medium, or high). T_1 is the set of years into the planning horizon at five year intervals prior to the second decision period, in this case {0, 10, 20, 30} while T_2 is the set of years into the planning horizon at five year intervals following the second decision period, in this case {40, 50, 60, 70}, and τ is equal to 10, the time resolution at which discounted expected annual damage is calculated. The parameter t^* refers to number of years into the planning horizon at which the second decision stage occurs, in this case 40. Discounted expected annual damages at each time t, which are then multiplied by the time resolution t is meant to approximate the total damages over the time period $[t, t+\tau]$. Discounted expected annual damage values $f_{ead}\left(x_{si,0}, \boldsymbol{c}_{k,t}, x_{ns,0}\right) * \left(\frac{1}{1-\delta}\right)^t$ and $f_{ead}\left(x_{si,s}, \boldsymbol{c}_{k,t}, x_{ns,s}\right) * \left(\frac{1}{1-\delta}\right)^t$ are calculated at ten-year intervals, but the expected annual damage itself characterized by $f_{ead}\left(x_{si,0}, \boldsymbol{c}_{k,t}, x_{ns,0}\right)$ and $f_{ead}\left(x_{si,s}, \boldsymbol{c}_{k,t}, x_{ns,s}\right)$ is calculated at 20-year intervals such that e.g. $f_{ead}\left(x_{si,0}, \boldsymbol{c}_{k,10}, x_{ns,0}\right)$ resolves to the modelled expected annual damage $f_{ead}\left(x_{si,0}, \boldsymbol{c}_{k,20}, x_{ns,0}\right)$. As above, $\boldsymbol{c_{k,t}}$ refers to climatological parameters impacting flood risk, but is now indexed by sea level rise trajectory and year. The parameter δ refers to economic discount rate, in this case equal to 0.03. Decision variables $x_{sb,i}$ are binary decision variables indicating whether a structural upgrade will be made in a given decision period.

Decision variables $x_{sm,i}$ parameterize levee crest heights following an upgrade (or crest height upgrade magnitude) as described above. Decision variables $x_{ns,i}$ reflect nonstructural elevation standards as described above. Decision variables with index 0 parameterize protection decisions made at the initial decision point, and decision variables index 1, 2, and 3 respectively

parameterize protection decisions made at the second decision point under low, medium, and high sea level rise trajectories respectively. The bounds on nonstructural protection given by the third constraint are as shown because the lowest building in the modelled system sits at -13 feet NAVD88, and a nonstructural elevation of 15 feet NAVD88 would buy out every building in the system. Optimization was carried out using mixed-lever NSGA-II as implemented in the PyMoo Python package; bounds listed for individual decision variables are implemented via initial population sampling and mutation operators designed to respect such bounds. Inequality constraints were implemented using the repair operator interface present in PyMoo to implement a repair operator which snaps infeasible solutions into the feasible space by setting parameters on the left-hand side of violated constraints to the value of the right-hand side. Decision variables parameterizing nonstructural standards at the second decision point are permitted to be lower than the decision variable parameterizing the nonstructural standard at the initial decision point by up to 0.5 feet, as shown in the final constraint. When occurs, the system model does not apply additional nonstructural protection in the second decision stage; this is intended to permit solutions which do not apply nonstructural protection in the second decision stage without having to introduce an additional binary decision variable, and to permit such solutions to persist in the solution population as the genetic algorithm converges.

The optimization problem posed by the static problem formulation can be posed similarly:

$$Minimize < f_1(\boldsymbol{x}), f_2(\boldsymbol{x}), f_3(\boldsymbol{x}) >, \boldsymbol{x} = (x_{sb,0}, x_{sm,0}, x_{ns,0}, x_{sb,1}, x_{sm,1}, x_{ns,1})$$

$$s.t.$$

$$x_{sb,i} \in \{0,1\}\ \forall\ i \in \{0,1\}$$

$$x_{sm,i} \in [0,99]\ \forall\ i \in \{0,1\}$$

$$x_{ns,i} \in [-13, 15], \forall\ i \in \{0,1\}$$

$$x_{sm,1} \geq x_{si,0}$$

$$x_{ns,1} \geq x_{ns,0} - 0.5$$

It is the same in almost all respects. The only difference here is that there is only one set of decision variables for the second decision point which is applied irrespective of sea level rise trajectory. The three objective functions have a corresponding alteration, in that while the objectives continue to correspond to net present values of costs in low, medium, and high sea level rise futures

respectively, they admit a single parameterization of mitigation at the second decision-stage rather than one for each set of sea level rise futures.

$$f_s(\boldsymbol{x}) = \frac{1}{|K(s)|}\sum_{k\in K(s)}\left(\left(\sum_{t\in T_1} f_{ead}\left(\boldsymbol{M}(x_{si,0}), \boldsymbol{c}_{k,t}, x_{ns,0}\right) * \left(\frac{1}{1-\delta}\right)^{t} * \tau \right.\right.$$
$$\left. + \sum_{t\in T_2} f_{ead}\left(\boldsymbol{M}(x_{si,1}), \boldsymbol{c}_{k,t}, x_{ns,1}\right) * \tau * \left(\frac{1}{1-\delta}\right)^{t}\right) + f_{cost}\left(\boldsymbol{M}(x_{si,0}), x_{ns,0}\right)$$
$$\left. + f_{cost}\left(\boldsymbol{M}(x_{si,1}), x_{ns,1}\right)\left(\frac{1}{1-\delta}\right)^{t^*}\right)$$

3.2 Results: Initial Experiment

Optimization of both adaptive and static problem formulations converged approximately to their Pareto frontiers as shown in figures 9 and 10. The static context appears to converge slightly faster than the adaptive context due to its lower-dimensional lever space. Outside of cases where adaptive actions are governed by a parameterized decision rule such as in direct policy search [39], this reflects a fundamental limitation in optimization-driven generation of adaptive strategies. This bottleneck becomes more impactful as the dimensionality of the lever space of the adaptive strategy increases.

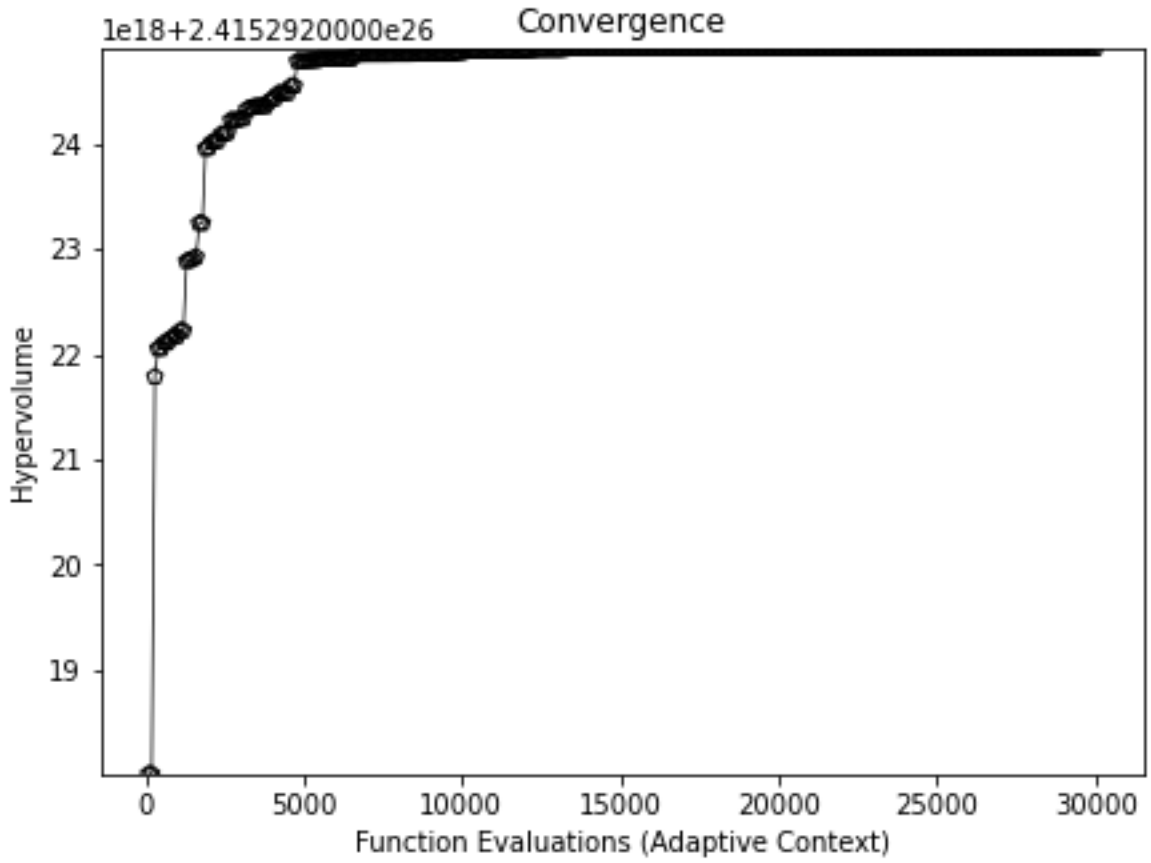

Figure 9: Dominated hypervolume of estimated Pareto frontier in the adaptive problem context versus number of function evaluations.

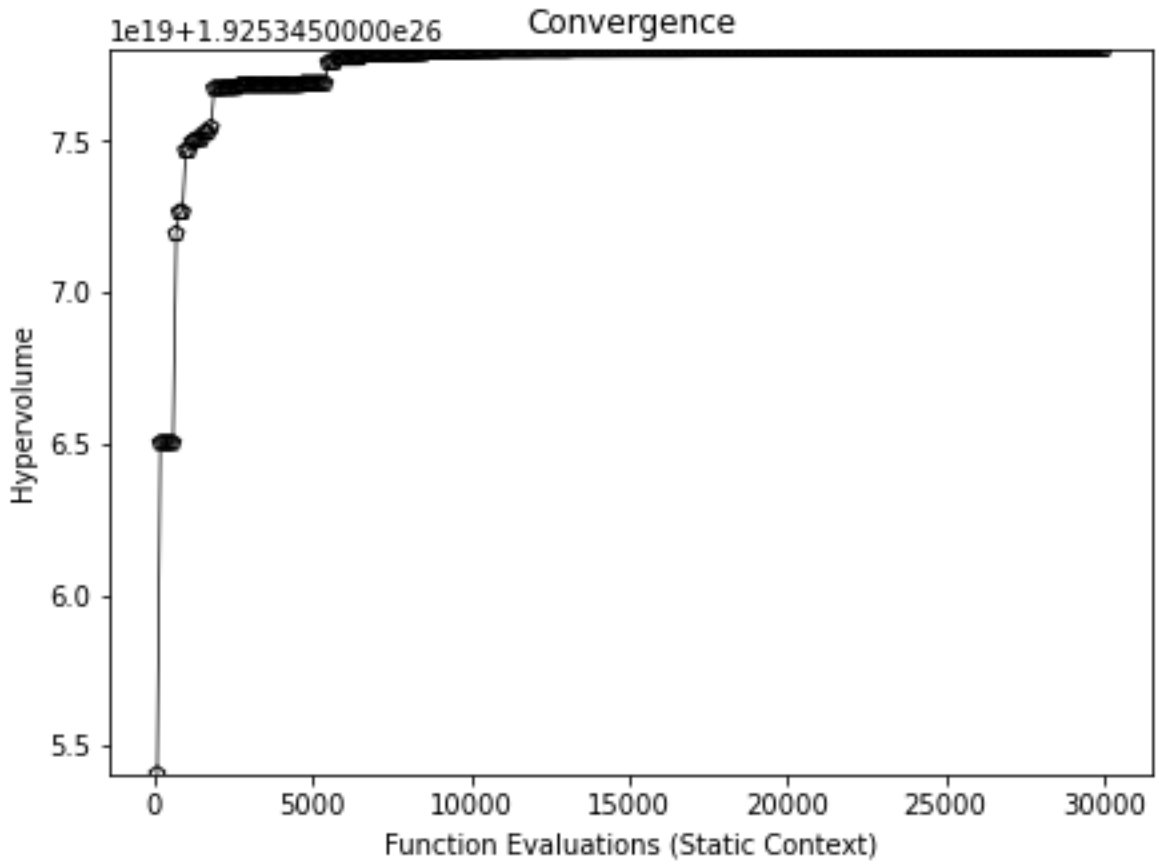

Figure 10: Dominated hypervolume of estimated Pareto frontier in the static problem context versus number of function evaluations.

Surprisingly, solutions found by the static problem context frequently outperform solutions found by the adaptive problem context as shown in Figures 11 and 12. Compared to the solutions found by the static context, adaptive solutions have a much narrower range of performance in the "high" set of sea level rise trajectories. Within this narrow performance range over high sea level rise trajectories, adaptive solutions exhibit a much wider range of performance in the adaptive context over the medium and low sea level rise trajectories.

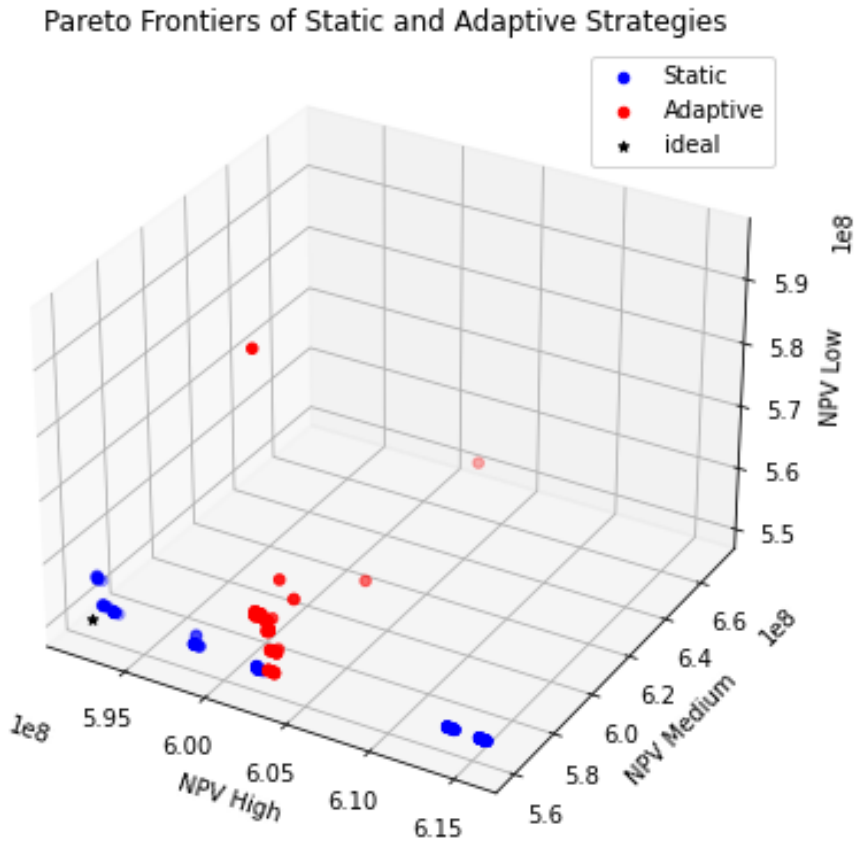

Figure 11: Pareto frontiers of flood risk mitigation strategies found by static and adaptive optimization contexts. Static strategies are shown in blue, and adaptive strategies are shown in red.

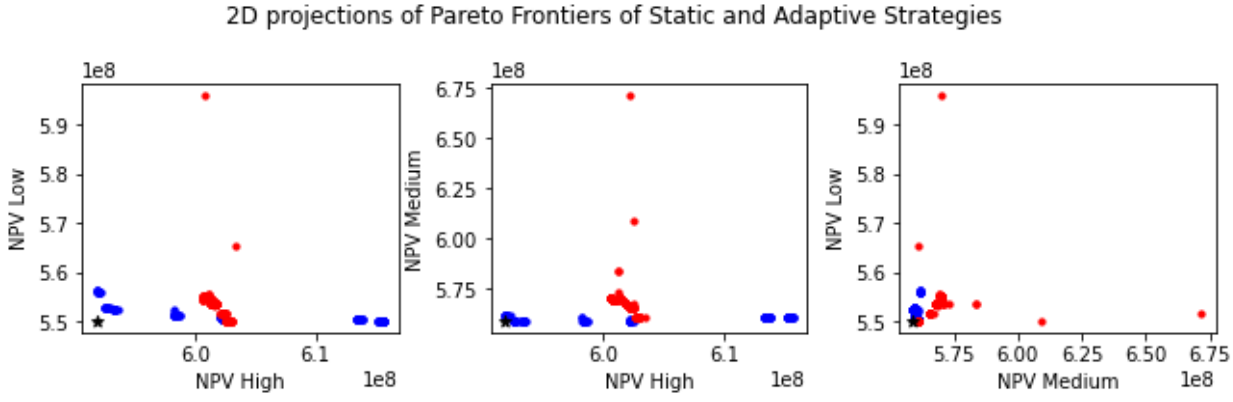

Figure 12: Two-dimensional projections of the Pareto frontiers of flood risk mitigation strategies found by static and adaptive optimization contexts. Static strategies are shown in blue, and adaptive strategies are shown in red.

The relatively poor performance in the adaptive optimization context relative to the static context is attributable to the lever structure of their respective solutions. The adaptive and static contexts both consistently produced strategies which invested in structural upgrades at the first decision point in 2015 and which did not invest in an additional structural upgrade at the second decision point in 2055, as shown in Figure 13. The few adaptive strategies which did invest in an additional structural upgrade in 2055 likely reflect a failure of the optimization to converge to the true Pareto frontier; the worst-performing solution in low sea level rise futures invests in a second structural upgrade in low futures, and the worst-performing solution in medium futures similarly invests in a second structural upgrade in medium futures. This is attributable to the substantial overhead costs associated with structural protection upgrades, which physically correspond to the need to partially break down and reconstruct the existing levee system. The magnitude of structural upgrades is similar between strategies discovered by the adaptive and static optimization contexts, as shown in Figure 14. This figure does show however that the adaptive decision context results have slightly greater intensity of structural upgrade at the first decision point than the static context. The most notable difference in solution behavior between the strategies produced by the adaptive and static optimization contexts is the nonstructural elevation standard as shown in Figure 15. This figure shows that the static context invests in nonstructural mitigation exclusively at the first decision point so as to avoid the excess expenditure of retrofitting the same building twice. The adaptive context invests less aggressively in nonstructural mitigation at the first decision point and tunes its investment at the second point based on observed sea level rise. While this does not lead to better performance compared to the strategies generated in the static context, it does appear to show that the most effective adaptive flood risk mitigation strategies make structural investments in the same way as non-adaptive strategies, and that the most effective adaptive behaviors in flood risk mitigation use nonstructural mitigation to respond to difficult-to-predict changes in flood hazard.

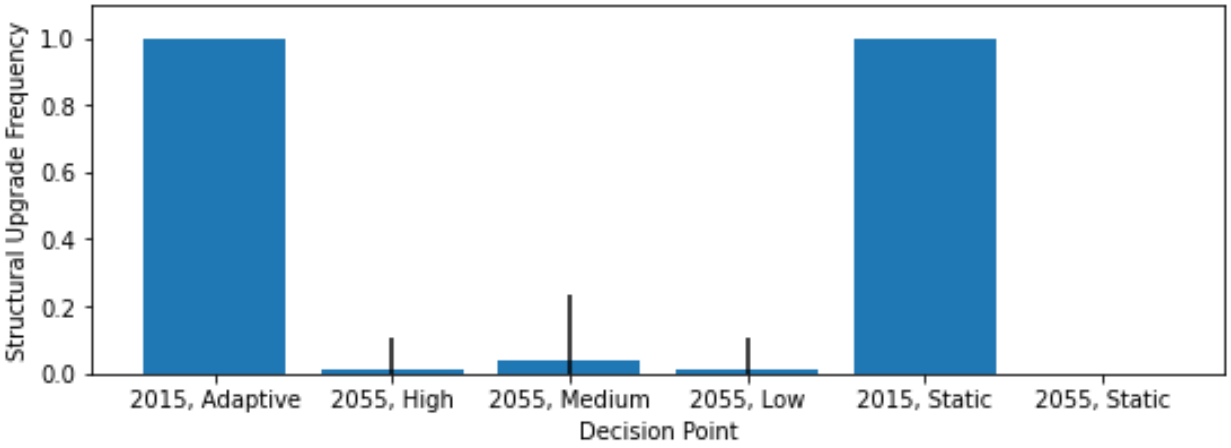

Figure 13: Frequency of structural upgrade at each decision point for adaptive and static optimization contexts. High, medium, and low values for the year 2055 are the conditional decisions made in the adaptive context in high, medium, and low sets of sea level rise trajectories respectively. Vertical black lines show standard deviations among solution sets.

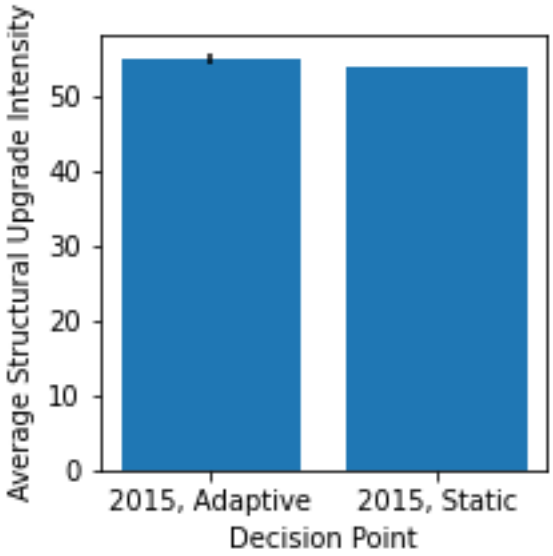

Figure 14: Average magnitude of structural upgrade the initial decision point for adaptive and static optimization contexts. Small vertical black lines show standard deviations among solution sets. Note that these values are a dimensionless index along a pre-computed structural upgrade path and do not directly reflect levee crest heights or levels of expenditure.

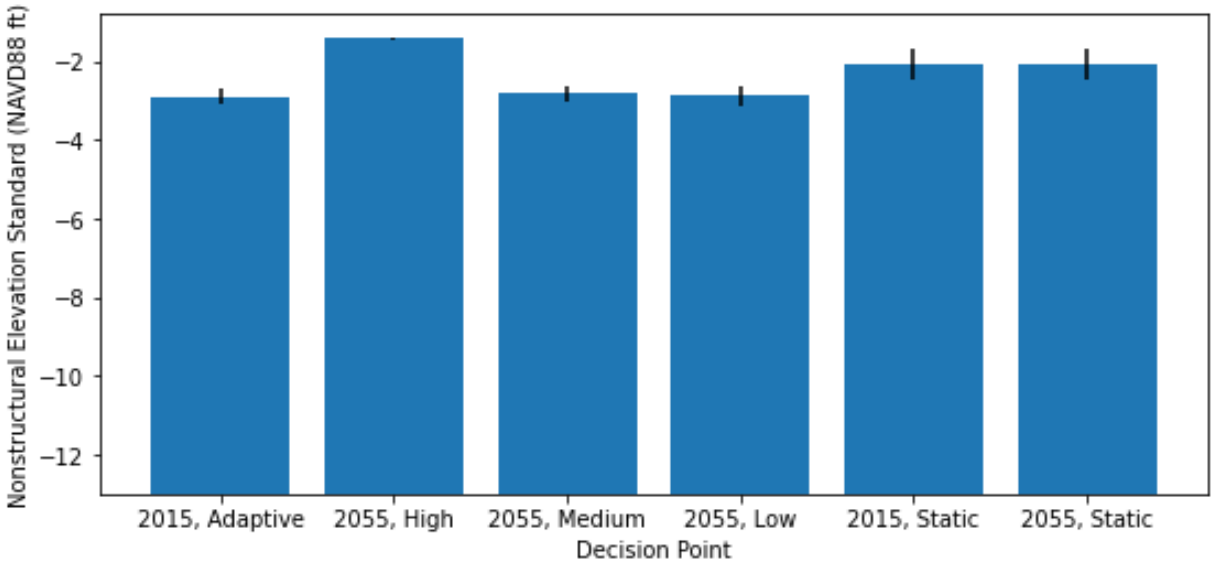

Figure 15: Average nonstructural standard elevation at each decision point for adaptive and static optimization contexts. High, medium, and low values for the year 2055 are the conditional decisions made in the adaptive context in high, medium, and low sets of sea level rise trajectories respectively. Nonstructural standard elevations permitted by the model range from -13.0 to 15.0 feet above sea level (NAVD88).

3.3 Discussion: Initial Experiment

The relatively poor performance of adaptive solutions in this experiment is best understood by looking back at its initial motivation. Static decision-making which does not explicitly support adaptivity is vulnerable to regret; a strategy which minimizes the net present value of protection and damages in low sea level rise futures will carry additional costs high sea level rise futures in the form of excess damage, and a strategy which minimizes the net present value of protection costs and damages in high sea level rise futures incurs additional costs in the form of unnecessary protection. Adaptivity is introduced into the decision context in order to mitigate these costs. As seen in Figures 11 and 12, this is successful in some respect—the performance of adaptive strategies in high sea level rise trajectories was more stable than that of non-adaptive strategies due to the tendency of adaptive strategies to tune nonstructural investment to the higher level of flood hazard in those trajectories. However, while static decision-making has associated costs associated with regret, adaptive decision-making carries its own costs in the form of overhead associated with repeated protection investments. As discussed above, repeated structural investment has considerable overhead associated with construction costs, and repeated

nonstructural investment has smaller but non-trivial overhead cost associated with repeated retrofits of the same buildings. As shown, given that modelled flood risk mitigation levers consist of structural and nonstructural upgrades in line with Louisiana's 2017 Coastal Master Plan, the cost of failing to adaptively protect against flood hazard changes driven by sea level rise alone is smaller than or similar to the additional overhead costs associated with adaptive planning.

In light of these findings, it is likely that more variability in deeply uncertain climate futures would lead to greater regret associated with static flood risk mitigation strategies and therefore a greater relative performance of adaptive mitigation strategies. In considering sea level rise only, the initial experiment underestimates plausible changes in flood hazard driven by climate change. Louisiana's 2023 Coastal Master Plan considers changes in the mean intensity of tropical cyclones in addition to sea level rise [66]. For this reason, an additional secondary experiment was implemented to evaluate whether the additional variability in climate induced flood hazard attributable to changes in tropical cyclone intensity as well as sea level rise leads to superior performance of the adaptive optimization context compared to the static context.

3.4 Methods: Secondary Experiment

Our secondary experiment incorporates changes in mean intensity of tropical cyclones in addition to sea level rise in order to characterize the value of adaptive decision-making in planning structural and nonstructural protection measures in plausible climate futures. The 2023 Master Plan uses a 50-year increase in mean intensity of tropical cyclones of 10% in its high-hazard future scenario based on a review by Knutson et al. [67] which finds increases in mean intensity ranging from 0 to 10% associated with a 2 degree Celsius increase in mean global surface temperature. As in the 2023 Master Plan, increases in tropical cyclone frequency are not considered because Knutson et al. [67] do not find evidence thereof. A longer planning horizon is used for this experiment, and the 2021 IPCC report [68] suggests a change in mean global surface temperature of up to 5 degrees Celsius by the year 2100 compared to the period from 1995 to 2014. Because Knutson et al. have not aggregated a similar set of projections of tropical cyclone intensification for greater increases in mean global surface temperature, and because more detailed accounting of future changes in flood hazard due to climate change is outside the scope of the research documented here, the secondary experiment uses a range of tropical cyclone intensification rates

of 0 to 25%, arriving at the latter by scaling up the results of [67] to the upper bounds of mean global surface temperature forecast by [68].

In order to avoid complicating the one-dimensional structure of climate hazard intensity modelled in our initial experiment, our secondary experiment varies tropical cyclone intensity as a function of sea level. Tropical cyclone intensification is assumed to scale linearly with sea level rise, such that the greatest value for sea level rise is associated with a mean intensity change of 25%, 0 sea level rise is associated with a 0% change in mean intensity. Note that the detailed formulations of the adaptive and static problem contexts presented above continue to hold; the vector $\boldsymbol{c}_{k,t}$ of climatological parameters at time t under trajectory k now reflects changes in sea level rise and mean storm intensity with a constant value for storm frequency, rather than changing sea level rise with constant mean storm intensity and frequency as it did in the initial experiment.

3.5 Results: Secondary Experiment

While the solution sets for our secondary experiment incorporating tropical cyclone intensification in addition to sea level converged in both the adaptive and static optimization contexts, it is notable that the adaptive context took substantially more function evaluations to converge than in the initial experiment (Figures 16 and 17). This is attributable to a more thorough exploration of adaptive combinations of decision levers; as flood damage increases under more extreme climate futures, the overhead cost of repeated structural upgrades becomes less prohibitive, leading to a deeper and correspondingly more computationally expensive exploration of solutions with repeated structural upgrades.

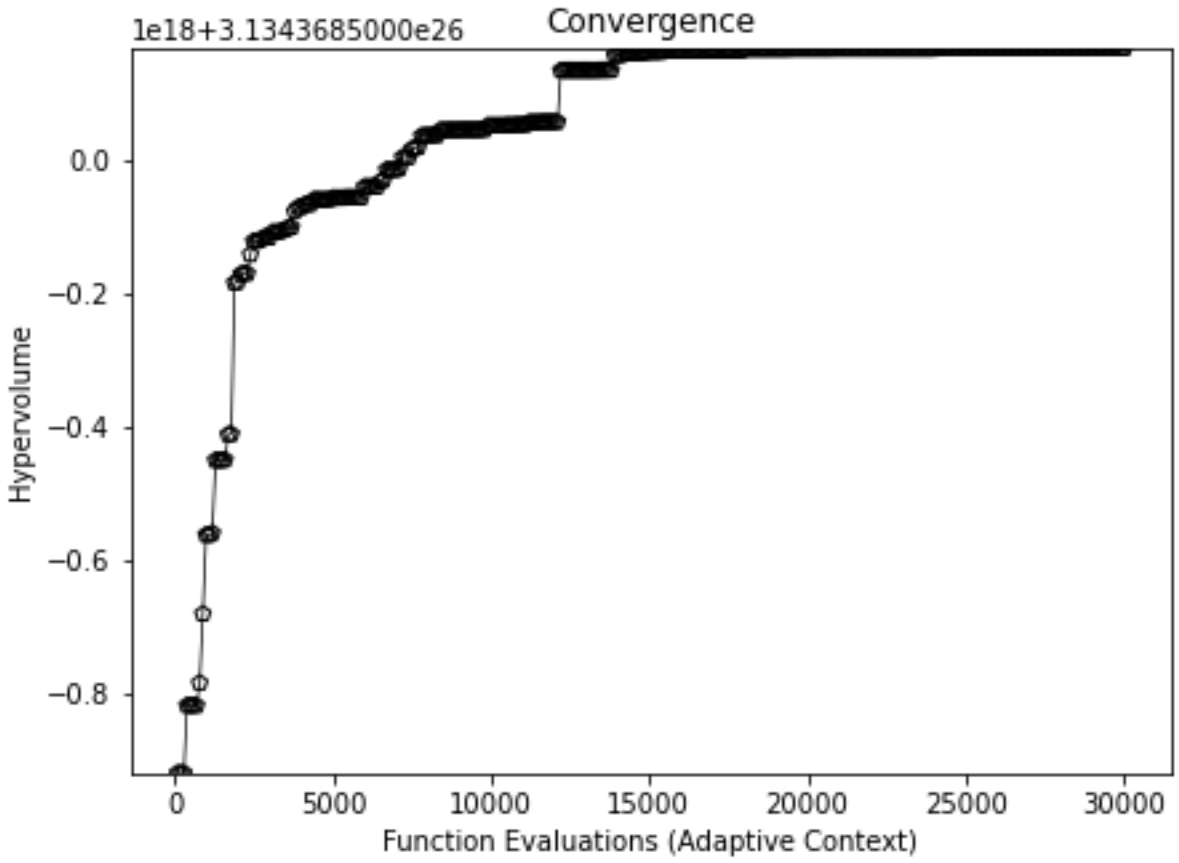

Figure 16: Dominated hypervolume of estimated Pareto frontier in the adaptive problem context versus number of function evaluations (secondary experiment).

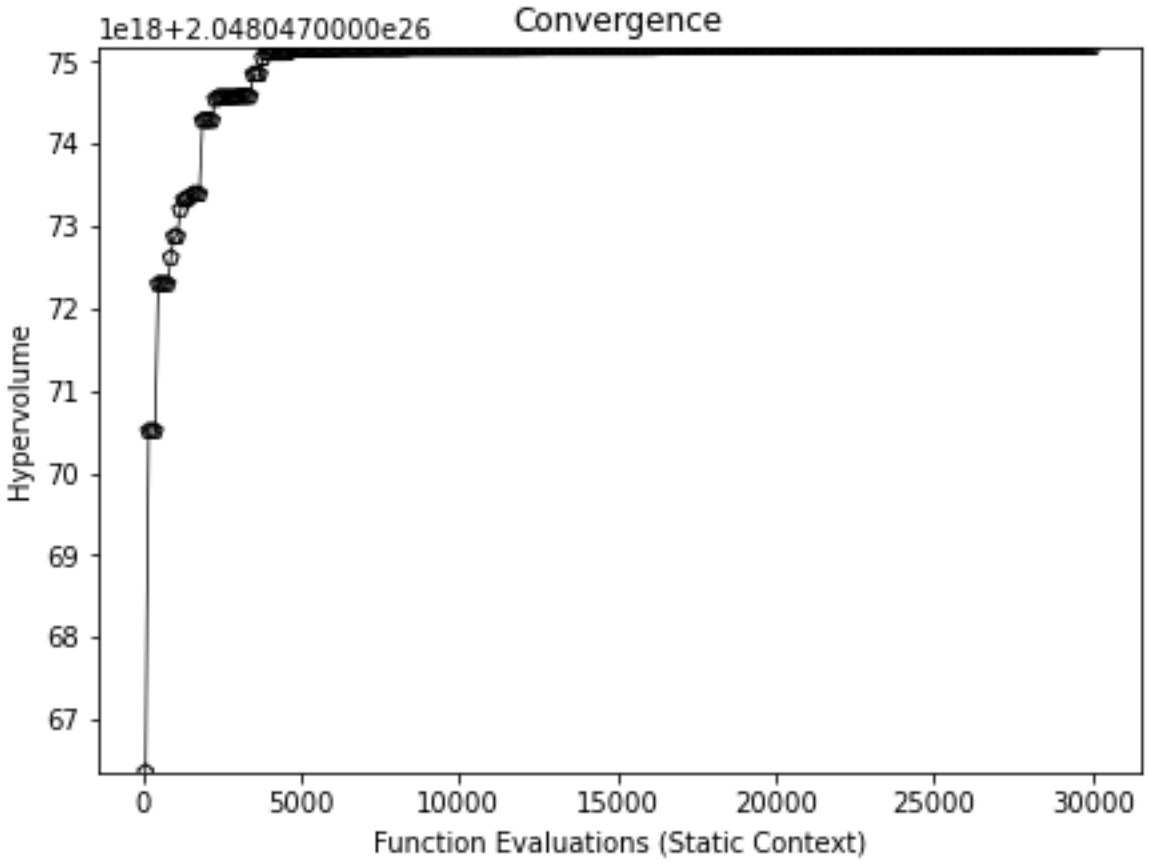

Figure 17: Dominated hypervolume of estimated Pareto frontier in the static problem context versus number of function evaluations (secondary experiment).

The performance of adaptive solutions in the secondary experiment are more in line with initial expectations than those of the initial experiment. As shown in Figures 18-21, most static strategies are dominated by adaptive strategies. The static strategies which remain non-dominated are those that perform best in high severity climate futures. This is unsurprising; strategies which directly target performance in high severity climate futures are expected to perform better in those futures than strategies which reward delaying investment for adaptive decision-making later. These figures also show that high-performing adaptive strategies only out-perform dominated static strategies by several million dollars of Net Present Value in any set of climate futures, which is a small margin considering that it is aggregated over 80 years.

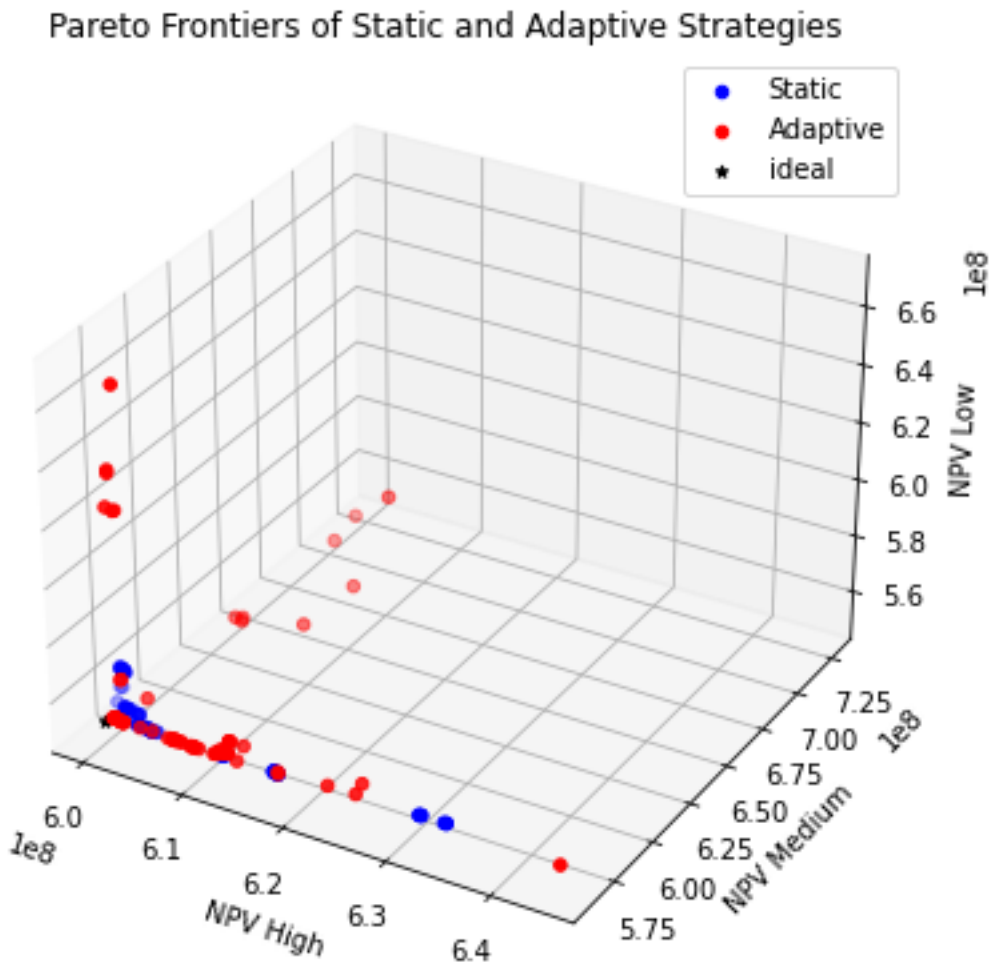

Figure 18: Pareto frontiers of flood risk mitigation strategies found by static and adaptive optimization contexts in the secondary experiment. Static strategies are shown in blue, and adaptive strategies are shown in red.

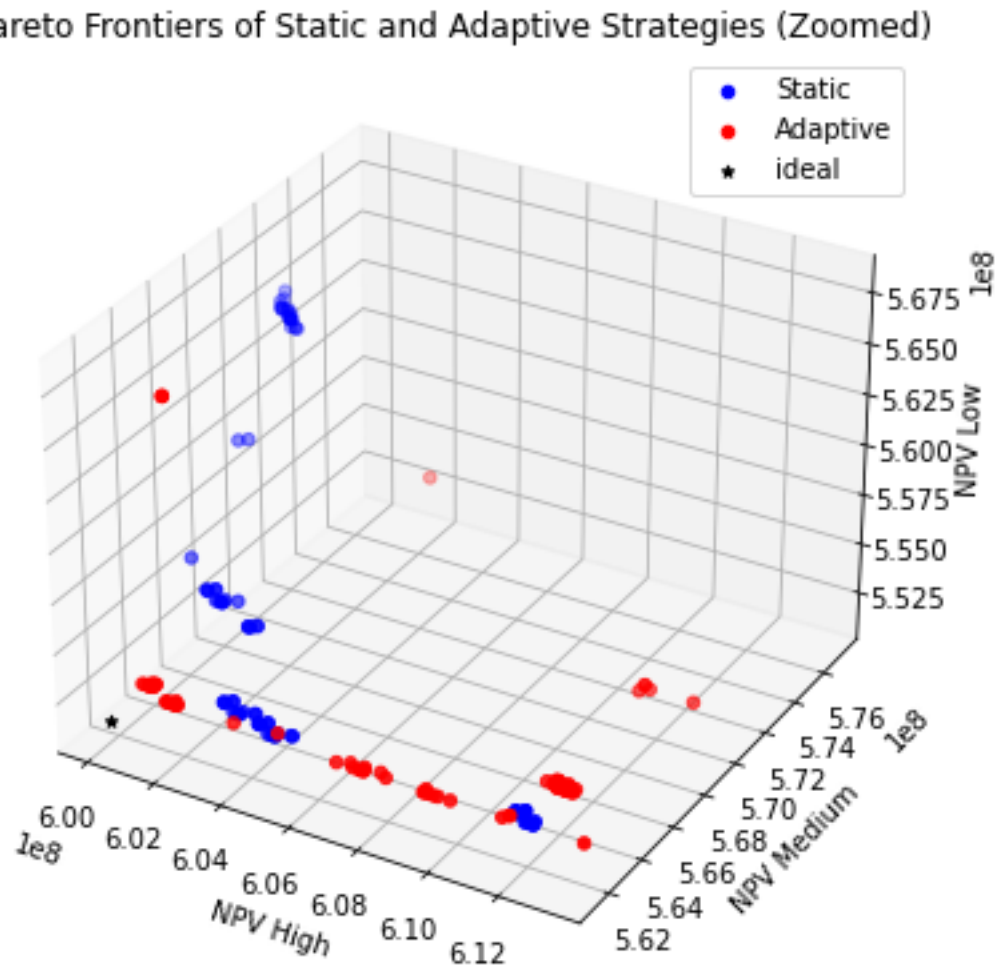

Figure 19: Subset Pareto frontiers of flood risk mitigation strategies found by static and adaptive optimization contexts nearest to the ideal point in the secondary experiment. Static strategies are shown in blue, and adaptive strategies are shown in red.

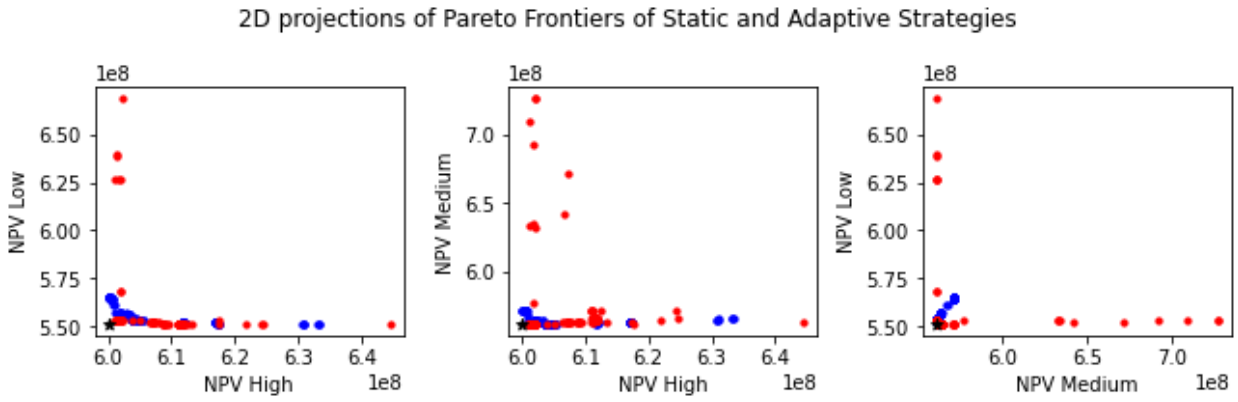

Figure 20: Two-dimensional projections of the Pareto frontiers of flood risk mitigation strategies found by static and adaptive optimization contexts in the secondary experiment. Static strategies are shown in blue, and adaptive strategies are shown in red.

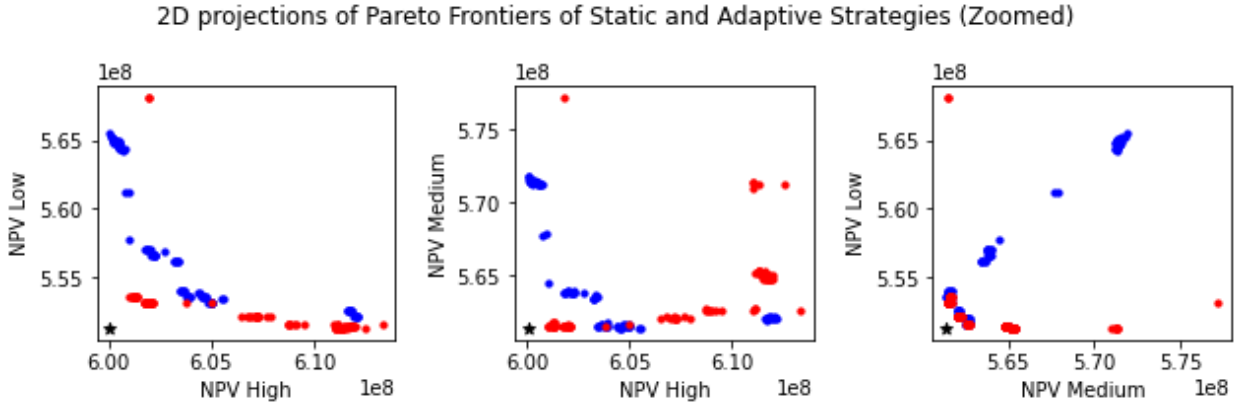

Figure 21: Two-dimensional projections of the Pareto frontiers of flood risk mitigation strategies found by static and adaptive optimization contexts in the secondary experiment, displaying solutions nearest the ideal point. Static strategies are shown in blue, and adaptive strategies are shown in red.

While the adaptive strategies identified in the first experiment largely use nonstructural mitigation to react adaptively to different climate futures, the strategies found in the second experiment use both structural and nonstructural mitigation adaptively (Figures 22-25).

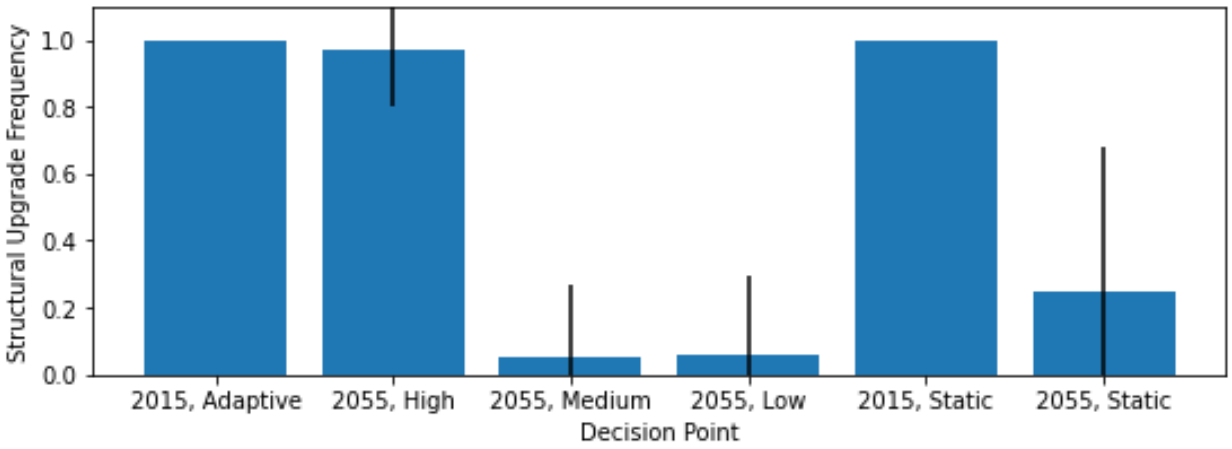

Figure 22: Frequency of structural upgrade at each decision point for adaptive and static optimization contexts in the secondary experiment. High, medium, and low values for the year 2055 are the conditional decisions made in the adaptive context in high, medium, and low sets of sea level rise trajectories respectively. Vertical black lines show standard deviations among solution sets.

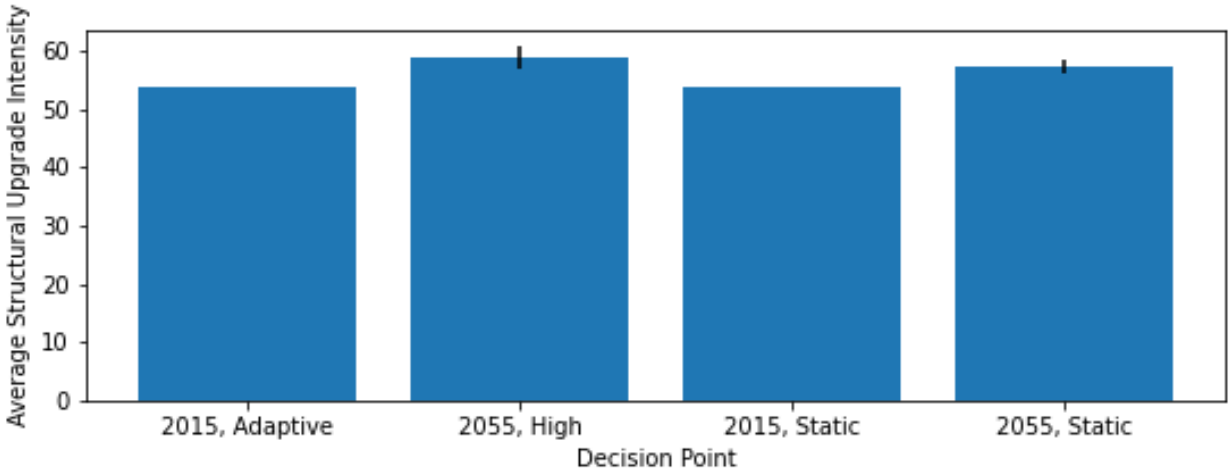

Figure 23: Average magnitude of structural upgrade the decision points with non-trivial structural upgrade frequency for adaptive and static optimization contexts in the secondary experiment. Small vertical black lines show standard deviations among solution sets. Note that these values are a dimensionless index along a pre-computed structural upgrade path and do not directly reflect levee crest heights or levels of expenditure.

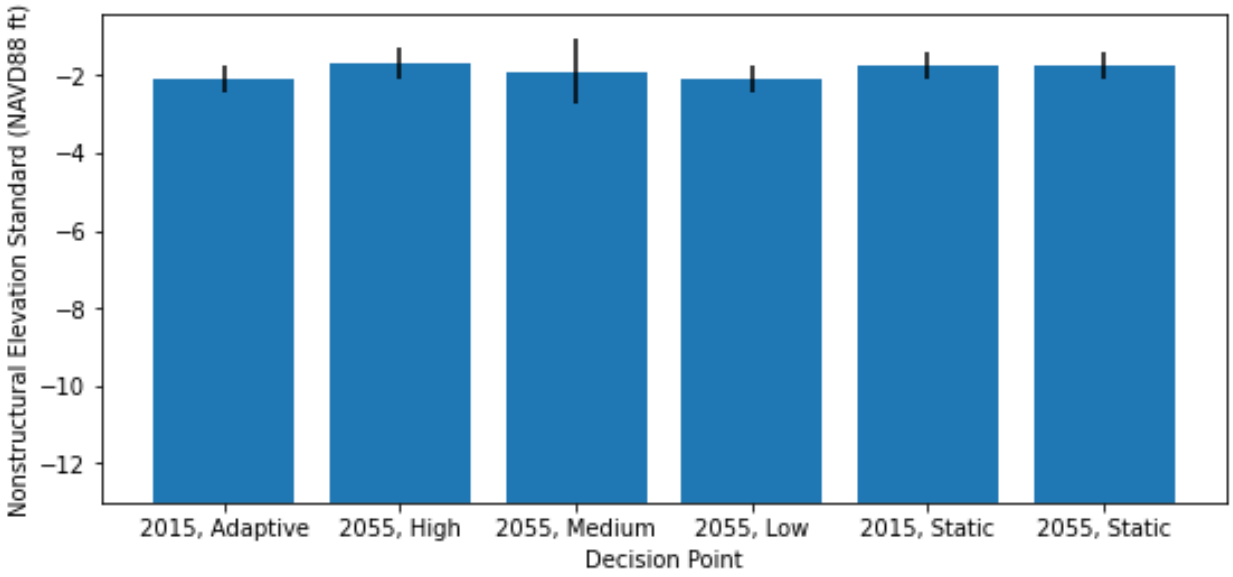

Figure 24: Average nonstructural standard elevation at each decision point for adaptive and static optimization contexts in the secondary experiment. High, medium, and low values for the year 2055 are the conditional decisions made in the adaptive context in high, medium, and low sets of sea level rise trajectories respectively. Nonstructural standard elevations permitted by the model range from -13.0 to 15.0 feet above sea level (NAVD88).

3.6 Discussion: Secondary Experiment

In this case study using the SWaMPS model, adaptive optimization-driven flood risk management strategies failed to substantively outperform non-adaptive strategies across deeply uncertain trajectories for sea level rise while other hazard drivers such as average intensity of tropical cyclones are held constant. When plausible ranges of increasing average intensity of tropical cyclones are considered, adaptive strategies began to outperform non-adaptive strategies, but not by a large margin. These counterintuitive results are driven by the overhead costs associated with repeated application of the protection measures considered in this study. The overhead cost of levee upgrades and the additional cost of retrofitting an individual building for the second time is close to or at times greater than the performance gains of adaptivity.

These experiments have also shown that when the range of climate futures considered is relatively narrow (which in this study occurs when considering only sea level rise and exclude other forms of climate-driven changes to flood hazard but in principle may generalize to projects with short planning horizons) adaptive strategies invest in structural protection similarly to non-adaptive strategies and use nonstructural mitigation to respond to extreme climate futures. As the

range of severity in considered climate futures increases, the regret associated with failing to protect against severe climate futures begins to outweigh the overhead associated with repeated structural protection investments, and adaptive strategies begin to apply structural protection adaptively in addition to nonstructural protection.

Taken together, these results show that adaptivity is unlikely to yield substantial performance improvements in decision-making for tropical flood risk mitigation when it is treated as a drop-in improvement. For adaptive strategies to outperform static ones, the set of available protection measures considered when generating mitigation strategies must include low overhead and low lock-in options that are incorporated into the decision context with adaptivity in mind. The overhead costs and lock-in associated with structural and nonstructural protection measures, designed to perform well in static planning context, substantially hinder their effectiveness in adaptive planning contexts.

From a methodological perspective, this study has shown that the Multi-Scenario Multi-Objective Robust Optimization framework proposed by Shavazipour et al. [57], [58] is highly effective in permitting analysts to generate and compare robust and efficient strategies in the context of decision-making under deep uncertainty, and enables straightforward and intuitive evaluation of tradeoffs between performance in alternate possible futures. While this comes at a significant computational cost, use and this framework may meaningfully improve the performance and interpretability of solution tradeoffs in future applications of decision-making under deep uncertainty in circumstances where it is computationally tractable to integrate and where optimization-driven strategy generation is desired.

4. EFFICIENT FLOOD RISK MITIGATION AND INTERSECTIONAL EQUITY IMPLICATIONS: A CASE STUDY IN NEW ORLEANS

To effectively mitigate climate related risks, decision-makers and analysts must first decide what metrics to consider as they rank and prioritize protection projects. Cost-benefit analysis (CBA) remains the predominant decision metric in applied mitigation efforts and governmental policy writ large [69]–[74]. Many institutions such as the Federal Emergency Management Agency and the U.S. Army Corps of Engineers are required to perform CBA for all flood risk mitigation projects and grants, although FEMA has recently relaxed their cost-effectiveness requirements for marginalized communities [75], [76]. Typical implementations of CBA treat the direct structure damages avoided over a project's lifetime as a key benefit, which implicitly prioritizes projects that protect expensive structures. This has the potential to worsen existing social inequity [77], [78]. Substantial differences already exist between racial and income groups in air quality [79]–[81], water quality [82]–[84], and flood risk [85]–[87]. This potential for harm is increasingly salient; the growing adoption of analyses which characterize flood risk at the level of individual structures [88]–[90] makes it possible to algorithmically design interventions at the level of individual homes which (inadvertently) selectively protect expensive structures, the effects of which may then be masked by data aggregation [91]–[93]. Practical efforts to reduce these outcome gaps are hindered by divergent perspectives on how to define and operationalize equity [94]–[98]. Action to mitigate these outcome gaps or otherwise prevent policy action from worsening them is required as a matter of standing federal policy; Federal agencies are required by executive order under the National Environmental Policy Act to consider environmental justice in mitigation measures [99].This research presents an alternative decision metric from those used in traditional const-benefit analysis which yields improved equity outcomes from the perspective of multiple competing definitions of equity. This decision metric is analytically straightforward and yields economic efficiency comparable to a cost-benefit maximizing approach.

Equity, as defined by the Federal government, refers to the consistent and systematic treatment of all individuals in a fair, just, and impartial manner [100]. However, the tension between conflicting ideals of equity and equality complicates the design of policies which seek to lessen gaps in environmental risk outcomes. While a wider range of perspectives exist, this work focuses on the tension between deontological egalitarianism and prioritarianism as described by

Parfit [101]. Deontological in this context refers to the belief that the morality of an action comes from the action itself rather than the consequences of the action. As a consequence, deontological egalitarianism (hereafter referred to as egalitarianism) mandates that each individual is granted equal consideration and individual treatment irrespective of outcomes. Egalitarianism does not hold equality of outcome, equality in final distribution of resources, or the resolution of unequal status quo resource allocations, as normative aims, although they may be desirable for other reasons. This egalitarian ideal of equity is strongly tied to the notion of procedural justice, which is broadly similar in mandating that people are treated equally [102] but relates more to whether individuals perceive allocation process as fair and impartial than to the allocation mechanism itself [103]. In contrast, prioritarianism mandates that we prioritize the welfare of the least advantaged members of society in pursuit of more equal outcomes [101]. Prioritarianism is strongly related to the notion of distributive justice, which requires that resources are distributed are seen as fair or just[104]. Also relevant is the notion of corrective justice, which holds that if some party is harmed for the gain of another party, the former party is entitled to restitution from the latter.

A natural prioritarian solution to unequal environmental risk outcomes would be to funnel risk mitigation resources directly to socially vulnerable groups, or alternatively to incorporate diminishing marginal utility and equity weights as recommended by some welfare economists [105]. However, this would violate the egalitarian ideal of equity which requires that individuals are given equal consideration irrespective of the *status quo* resource allocation. An egalitarian perspective may judge the elimination of outcome gaps as a valid instrumental good, but likely opposes explicit prioritization of disadvantaged groups as unfair and discriminatory. An egalitarian perspective might see explicit prioritization of disadvantaged groups as a justified measure of corrective justice under the premise that disadvantage experienced by those groups is the result of systemic violations of procedural justice and egalitarian equity. However, recent events in US politics suggest that many decision-makers in states threatened by coastal flooding would reject this premise[106]–[108]. A substantial body of research shows that ideological disagreement between prioritarians and those who seek corrective justice for systemic procedural injustices on one side and egalitarians who reject the premise that status quo disparities are driven by systemic injustice on the other side hinders the remediation of racial disparities along a wide range of issues [94]–[98]. In order to mitigate existing disparities in environmental risk outcomes or otherwise avoid worsening those disparities in a way that is politically feasible in the near term, it is therefore

necessary to develop risk mitigation strategies which do so without violating widely held notions of justice and equity. The research presented here seeks to do so, presupposing the normative assertion that the moral obligation of an analyst (in a professional capacity) in an applied context is not to privilege their preferred notions of equity and justice, but rather to seek solutions which are Pareto-improving when viewing valid notions of justice and equity from a multiobjective lens, and to ensure that decision-makers are fully informed of the implications of their decisions in relation to each such axis.

The capacity of CBA to worsen existing socioeconomic disparities and therefore violate prioritarian notions of equity is evident from its core premises. To perform a cost-benefit analysis one first compares the sum total of a project's expected discounted economic (i.e., monetary) benefits to the sum total of the project's discounted financial costs (and monetized disbenefits if applicable). One then measures the performance or cost-effectiveness of the project by the ratio of the benefits to the costs [72]. As benefits are expressed purely in economic terms, CBA places greater value on the protection of expensive assets than the protection of inexpensive assets. In light of well-studied geographical and racial disparities in wealth [109], it is clear that purely CBA-driven evaluation of protection projects has the potential to neglect historically marginalized communities with lower average building costs (discussed next).

The severity of economic and racial gaps in environmental risks is starkly highlighted in the flood risk faced by the city of New Orleans, Louisiana. Predominantly Black neighborhoods in New Orleans bore the brunt of lives lost during Hurricane Katrina in 2005 [87]. Laska and Morrow [110] argue that the severity of the outcome of Katrina was in part driven by inequality among racial and income groups. The speed and effectiveness of post-Katrina recovery efforts differed wildly between socio-economic and racial groups in ways that are only visible upon sufficiently granular analysis [111], [112]. There were substantial racial disparities in immediate post-Katrina relief investments even when controlling for socio-economic status [113]. The existing disparities in environmental risk across racial and economic groups and the possibility that CBA could worsen them raise two main questions addressed in this study: (1) How much can a CBA approach worsen these disparities? (2) Can one remedy these disparities in a way that satisfies both egalitarian and prioritarian ideals?

Few studies have compared the effectiveness of algorithmic cost-benefit maximizing flood risk interventions with equity-driven alternatives. Those which have either pertain to equity

between regions in the design of flood protection structures [77] or integrate a Gini coefficient or equity weighting of flood exposure in the design of urban drainage projects or ecological restoration [114]–[116]. Similar studies in the area of water resource management have also used a Gini coefficient or a coefficient of variation of benefits as an equity measure [117]–[119]. These studies consistently find that inclusion of equity criteria in decision-making changes the design of interventions and identify severities of trade-off between equity and efficiency depending on how equity is operationalized. To our knowledge however, no such study examines flood risk interventions which are made at the level of individual structures. Additionally, every such study examines an equity-oriented strategy which either because it explicitly prioritizes poorer residents or because it explicitly targets equal final resource distributions, violating egalitarian norm of granting equal procedural consideration to individuals irrespective of status quo or final resource allocation. This suggests a research gap; if CBA worsens socioeconomic disparities, then there is a pressing need for a method to design flood risk mitigation strategies at the level of individual structures that mitigates the potential harms of CBA and thereby improving distributive justice outcomes from a prioritarian perspective without violating egalitarian norms of procedural justice.

Recently available methods for characterizing flood risk at the level of individual buildings have increased the scope of potential harm from traditional cost-benefit analysis. Building-level interventions have proven effective in mitigating flood risk [120]–[123], and the Coastal Louisiana Risk Assessment Model (CLARA) is able to characterize changes to the flood risk faced by individual buildings across coastal Louisiana under a range of protection projects including those at the level of individual buildings [89]. It is therefore possible to generate, evaluate, and optimize mitigation projects at the level of individual buildings, although to our knowledge this has not yet been applied in practice. A mitigation project which discriminates between individual buildings in risk mitigation naturally carries a far greater potential for inequity than one designed at the level of spatially broad and socioeconomically diverse areas. While to our knowledge no applied flood risk mitigation policies have been implemented to date which were designed using building-level optimizations, it is necessary to investigate the potential equity impacts of such policies.

4.1 Methods

Flood risk mitigation projects are frequently selected on the basis of which project most cost-effectively minimizes direct economic damage (DED) in expectation. The Coastal Louisiana Risk

Assessment model used by Louisiana's Coastal Master Plan estimates a discretized annual exceedance probability distribution for flooding at a given structure's location for 22 return periods ranging from the 2-year flood (i.e., 1-in-2 chance of occurring or being exceeded in a given year) to the 2000-year flood. The damage suffered by a given structure as a fraction of its replacement cost is estimated by USACE methods as a piecewise linear function of flood depths. The replacement cost of an individual structure is estimated using structure-level attributes (e.g., square footage, number of stories) combined with valuation assumptions from the FEMA Hazus-MH model [124]. The expected discounted value of direct economic damage over time can therefore be expressed as shown below.

$$\mathbb{E}[DED] = \sum_{t}\sum_{s}\sum_{e} p_{s,t}(e)D_s(e)r(s)\gamma^t$$

Here, t denotes years elapsed, s indexes structures, e represents plausible flood elevations, $p_{s,t}(e)$ denotes the probability that structure s in year t will face flood elevation e, $D_s(e)$ denotes the damage suffered by structure s and its contents as a fraction of their replacement cost resulting from flood elevation e, $r(s)$ denotes the replacement cost of structure s and its contents, and γ denotes the economic discounting rate.

The use of direct economic damage as a characterization of risk carries a number of benefits – dimensionless cost/benefit ratios and net present values in units can be easily calculated and interpreted as the costs and benefits of a protection project are both expressed in units of dollars. Further, policies which minimize DED under a fixed budget appear to be socially optimal if we neglect distributional impacts or otherwise assume that *post-hoc* wealth transfers implicitly resolve problematic distributional impacts as is commonly assumed in resource economics [125]. The chief weakness of DED as a decision metric is clear from its formulation – if all residences are exposed to equal flood elevations and therefore suffer an equal amount of damage proportional to their replacement costs, a scenario which might reasonably be described as all residents having equal exposure to flood risk, a DED-minimizing strategy would preferentially protect more expensive residences. To address this weakness, I define a similar flood risk metric with a nearly identical formulation, but which explicitly does not consider the replacement value of a residence, which is referred to as Residence Loss Equivalents (RLE).

$$\mathbb{E}[RLE] = \sum_{t}\sum_{s}\sum_{e} p_{s,t}(e)D_s^*(e)\gamma^t$$

RLE here represents an attempt to capture similar information to DED without consideration for a residence's replacement value. The use of D_s^* rather than $D_s(e)$ denotes damage to the structure alone rather than the structure and its contents. In equalizing the weight applied to the total loss of a residence, this study effects a deontologically egalitarian metric which explicitly grants each household equal procedural consideration regardless of the financial value of the home. In addition to supporting decision processes compliant with egalitarian principles of equality and ultimately serving prioritarian aims to a greater extent than traditional CBA, RLE carries the distinct advantage that it can be calculated using a proper subset of the data and methods required to calculate RLE.

To investigate the distributional and therefore prioritarian-relevant implications of our risk measures, I identify coastal protection projects for the city of New Orleans which simultaneously minimize DED and RLE assuming flood hazard consistent with Louisiana's 2023 Comprehensive Master Plan for a Sustainable Coast [89]. These projects apply elevation-in-place retrofits to single-family residences within the New Orleans levee system under a fixed 100-million-dollar budget to minimize the weighted average of DED and RLE under varying relative weights. In the limiting case, the strategy which minimizes DED alone is referred to as the "cost" strategy, and that which minimizes RLE alone is referred to as the "housing" strategy. I first project changes to flood risk over a 50-year planning horizon resulting from elevation-in-place retrofits at one-foot intervals. Flood depth exceedances are modeled at decadal time slices according to the changing environmental conditions (e.g., sea level rise) of the Coastal Master Plan's Lower landscape scenario. Risk is calculated at 5-year intervals with linearly interpolated flood depth exceedances, treated as constant within those 5-year intervals, and aggregated with a discount rate of zero. The cost of retrofits are assumed to scale linearly with the square footage of residences with a cost coefficient which depends on the final elevation of the retrofit in line with the 2023 Coastal Master Plan [126].

Optimal elevation strategies are derived by sorting individual structure retrofits by cost-effectiveness in the relevant metric and iteratively removing redundant strategies, e.g., eliminating a 1-foot elevation-in-place of a given residence for which a 2-foot elevation-in-place is selected, until the fixed budget is fully expended without redundancy. Intermediate strategies are derived by scalarizing DED and RLE with arbitrary relative weights before calculating the cost-effectiveness of the resulting composite metric. I estimate the distributional impacts of these

strategies under the assumption that the occupants of these residences have the same demographic features of the Census block group in which they reside as given by the 2019 American Community Survey [127]. This assumption reflects a meaningful limitation; a portion of the results below are shown aggregated by racial or economic group, and error associated with this simplified downscaling of demographic information leads to some amount of error in demographic estimates shown below. For both the readability of mapped figures and the privacy of individual households, all mapped figures are shown as neighborhood averages based on neighborhood definitions published by the city of New Orleans [128], shown in Figure 25. Neighborhood characteristics including the number of single-family residences, the average replacement cost of single-family residences, the proportion of the population which is White or Black, and the poverty rate are shown in Appendix B, in Table 2 and Table 3.

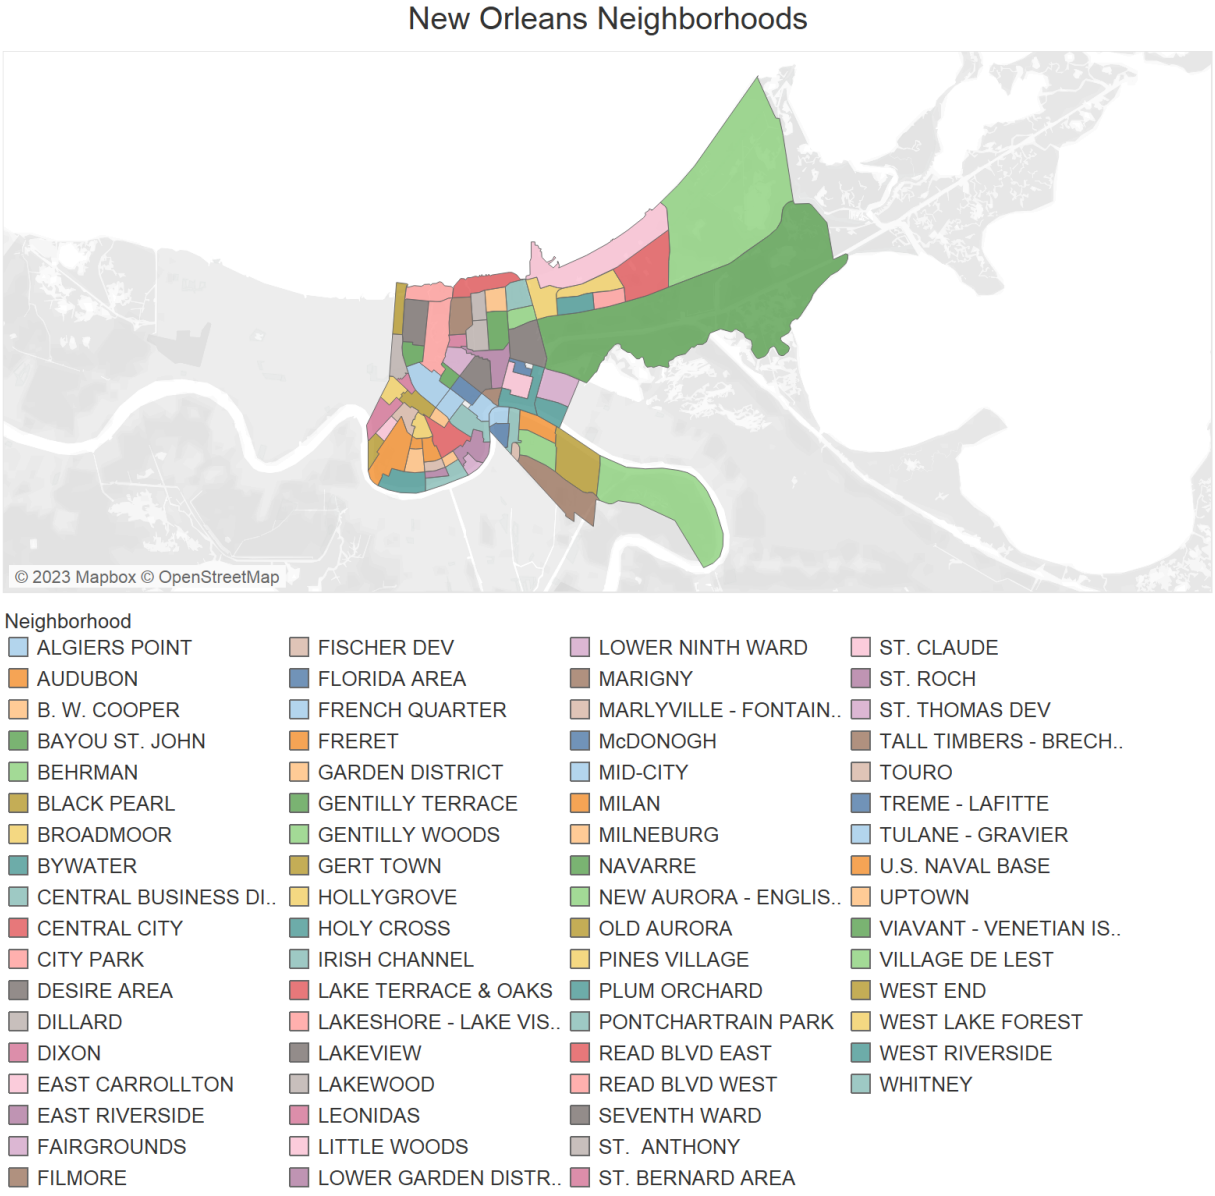

Figure 25: Map of New Orleans neighborhoods as provided by the city of New Orleans

4.2 Risk Outcomes by Race and Income

This study follows recent recommendations [129] to help mitigate inequitable CBA outcomes through the inclusion of alternative decision criteria. This study develops optimally cost-effective flood risk mitigation strategies targeting individual single-family residences for elevation-in-place (i.e., raising the elevation above grade of the first inhabited floor), respectively minimizing expected direct economic damage (DED) as per traditional CBA, and an alternative

egalitarian metric, denoted residence loss equivalents (RLE), which treats each home as if it had an identical value or replacement cost. In all cases, this study adopts a 50-year planning horizon when evaluating expected risks. Results show that strategies minimizing the egalitarian RLE are preferable to those minimizing DED from a prioritarian perspective, reducing income-related disparities in mitigation resource allocation and reducing nominal outcome gaps between Black and White New Orleans residents as measured by RLE. This article refers hereon to the DED-minimizing strategy as the "cost" strategy, and the RLE-minimizing strategy as the "housing" strategy.

Our analysis shows that Black residents receive the vast majority of mitigation resources under both the cost and housing strategies because they currently experience the overwhelming and disproportionate bulk of flood risk in New Orleans (Fig. 25). Black residents make up 59% of the population, but residences occupied by Black residents bear 79% of direct economic damage and 88% of residence loss equivalents. Black residents receive 81% and 90% of mitigation resources under the cost and housing strategy respectively. This is – perhaps – counterintuitive given the larger weight more expensive houses receive in the CBA framework and the current correlation between average house price and racial composition of a neighborhood. Our analysis reveals that the existing disparity in flood risks faced between Black and White residents in New Orleans is so overwhelming that Black neighborhoods bear the brunt of the flood risk burden within all considered strategies. This figure does show, however, that the housing strategy mitigates a large portion of the residence loss equivalent damage faced by New Orleans residents overall, thereby reducing the nominal size of the racial disparity in risk to housing (RLE).

Our analysis further reveals that the cost strategy principally benefits households with poverty-income ratio (i.e., the ratio of household income to the US poverty threshold) greater than 2, whereas the housing strategy provides greater benefits to lower-income households (Fig. 25). Taken together, the racial and income distributions of flood risk and mitigation investment show that the question of whether to target replacement cost (via DED) or cost-blind measures of housing risk (via RLE) becomes a question of how to allocate resources between wealthier or poorer predominantly Black neighborhoods, which is made more complicated in light of well-studied wealth disparities between Black and White populations more broadly [130]–[135].

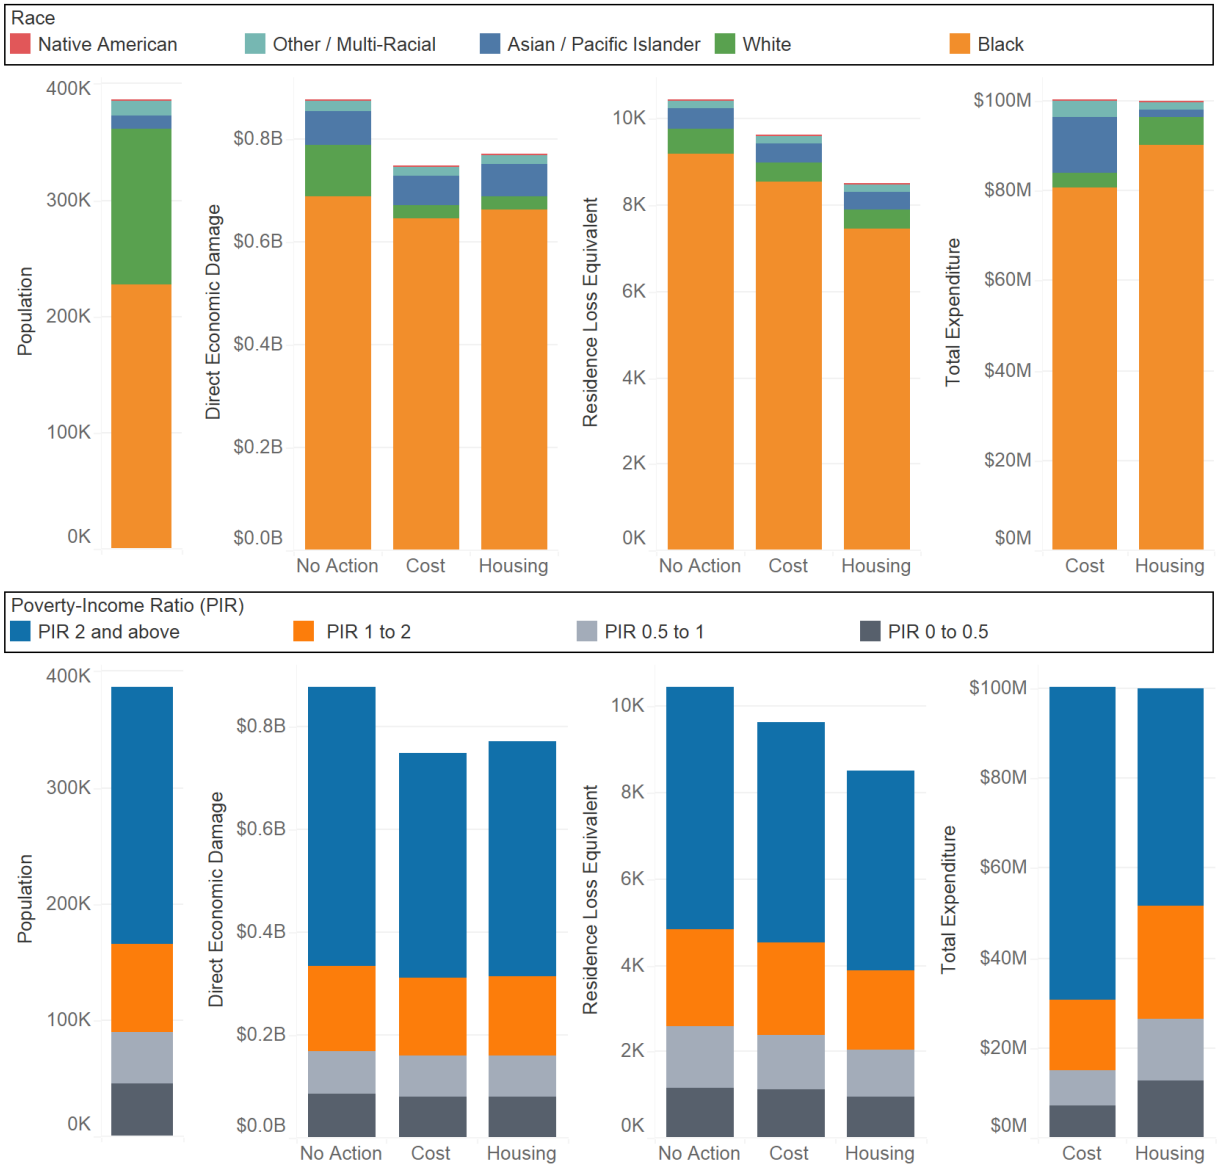

Figure 26: Total population, aggregate single-family residence flood risk estimates measured as direct economic damage and as residence loss equivalents, and total expenditure under no-action, cost, and housing strategies. Top row disaggregates measures by race and bottom row disaggregates measures by poverty-income ratio. Risk measures and spending are estimated at the individual building level and demographic estimates assume homogeneity within census block groups.

4.3 Risk Outcomes by Neighborhood

Assessed at the neighborhood level, predominantly Black neighborhoods visibly bear the brunt of flood risk in New Orleans (Fig. 26). Outside of the Lower 9th Ward, which receives substantial mitigation investment in both the cost and housing strategies, the cost strategy concentrates investment in the predominantly middle-class Black neighborhood of Read Boulevard East, whereas the housing strategy distributes investment across a series of poorer, predominantly Black neighborhoods (Fig. 26, 27). One such neighborhood, Read Boulevard West, carries an expected 0.35 residence loss equivalents per household under the status quo and receives approximately zero risk reduction along either metric from the cost strategy. Implementing the housing strategy reduces the residence loss equivalents per household by 40%, although its direct economic damage per household only decreases by roughly 2%. Conversely, Read Boulevard East sees a roughly 9% reduction in both direct economic damage and a reduction in residence loss equivalents per household under the cost strategy, and sees a negligible reduction in direct economic damage and 5% reduction in residence loss equivalents under the housing strategy. It is also worth noting that both predominantly White neighborhoods facing substantial flood risk as measured by direct economic damage, West End and Lakeshore, have 100% of their flood risk mitigated as measured by both metrics under both strategies. This indicates that the flood risk in these neighborhoods is concentrated in a small number of expensive structures that are inexpensive to mitigate. These neighborhoods are illustrative of the broader tendency for wealthier, predominantly White neighborhoods to face relatively little flood risk and for both mitigation strategies to target predominantly Black neighborhoods facing substantive flood risk. Among those neighborhoods, the cost strategy targets relatively more expensive homes in relatively higher-earning communities and for the housing strategy to target less expensive homes in lower-income communities. The distributional equity of these competing allocations remains complicated: each strategy helps to mitigate a facet of the overwhelming risk burden facing Black New Orleans residents.

The tendency of the cost strategy to protect a smaller number of more expensive homes than the housing strategy is observable within neighborhoods as well as between neighborhoods. In the Lower 9th Ward, implementing the cost strategy reduces direct economic damage by 41% and residence loss equivalents by 56%. In contrast, the housing strategy reduces direct economic damage by 32% and residence loss equivalents by 77% while spending fewer resources in the

neighborhood. This shows that even within the neighborhood, the cost strategy targets more expensive homes, and the housing strategy targets a greater number of less expensive homes.

Figure 27: Map showing salient neighborhood characteristics and flood risk measures across no action, cost, and housing strategies. A) shows the average replacement cost of single-family residences in each neighborhood. B) shows the proportional Black population of each neighborhood. C) shows the poverty rate. D), E), and F) show neighborhood average direct economic damage per household under no action, cost, and housing strategies respectively. G), H), and I) show neighborhood average residence loss equivalents respectively. Colored points on the map reflect the locations of individual single-family residences. Select neighborhoods are labelled in panel D).

Figure 27 continued

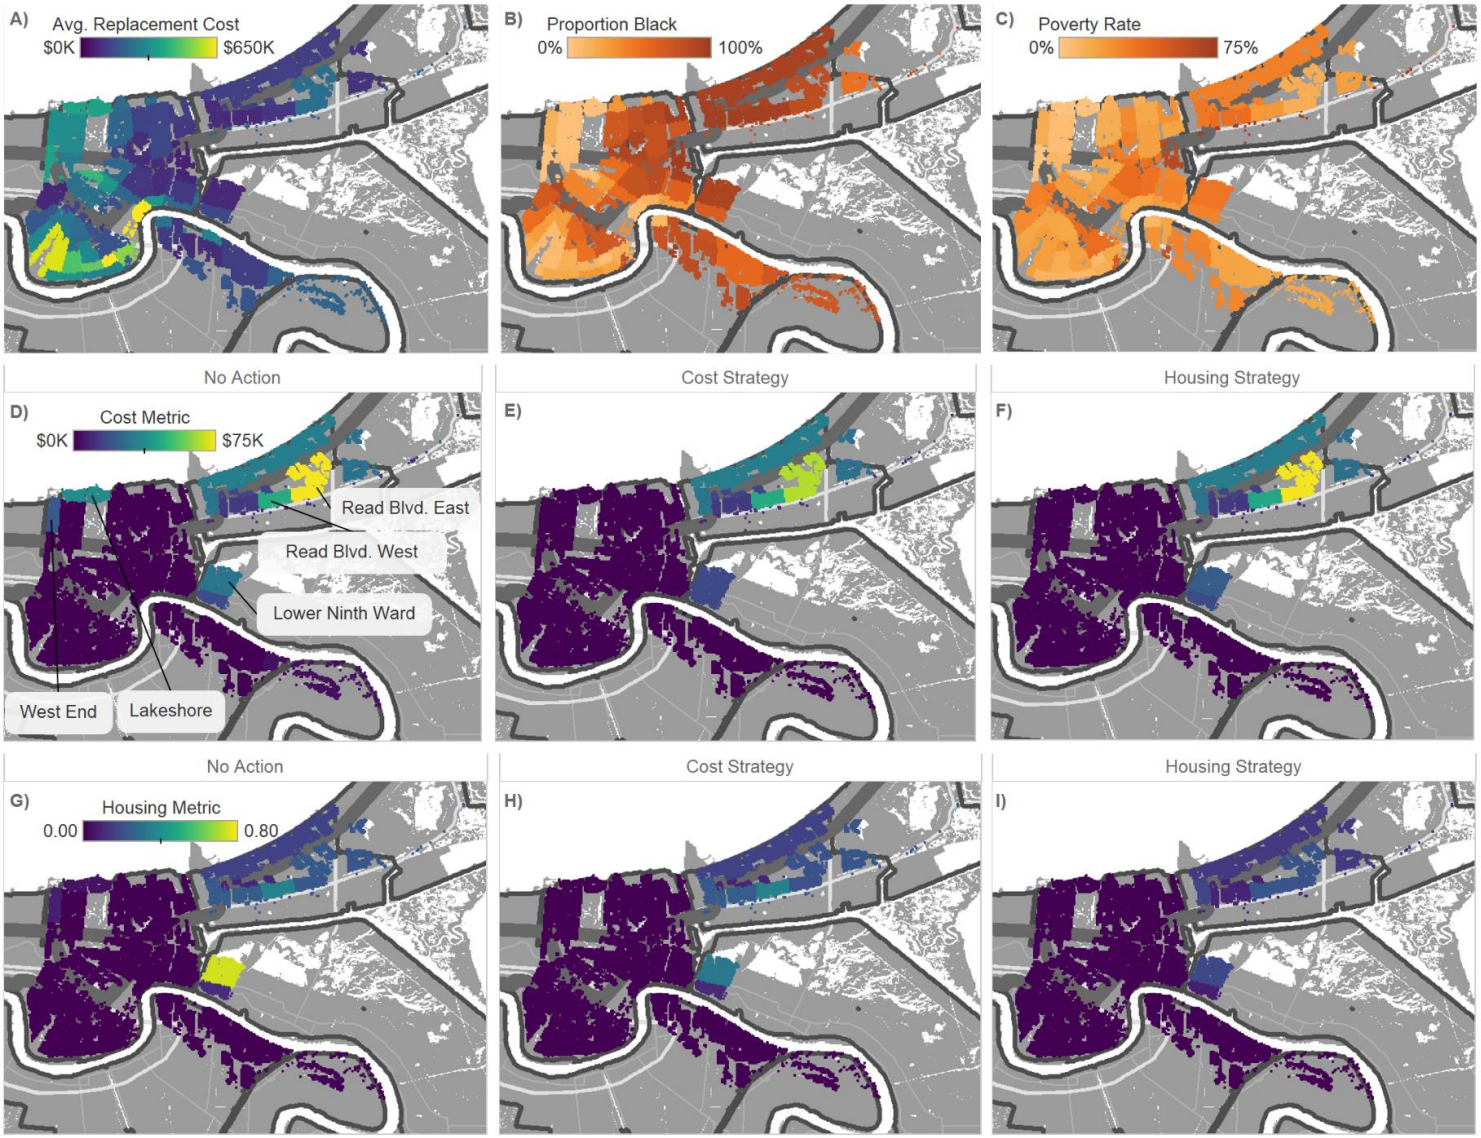

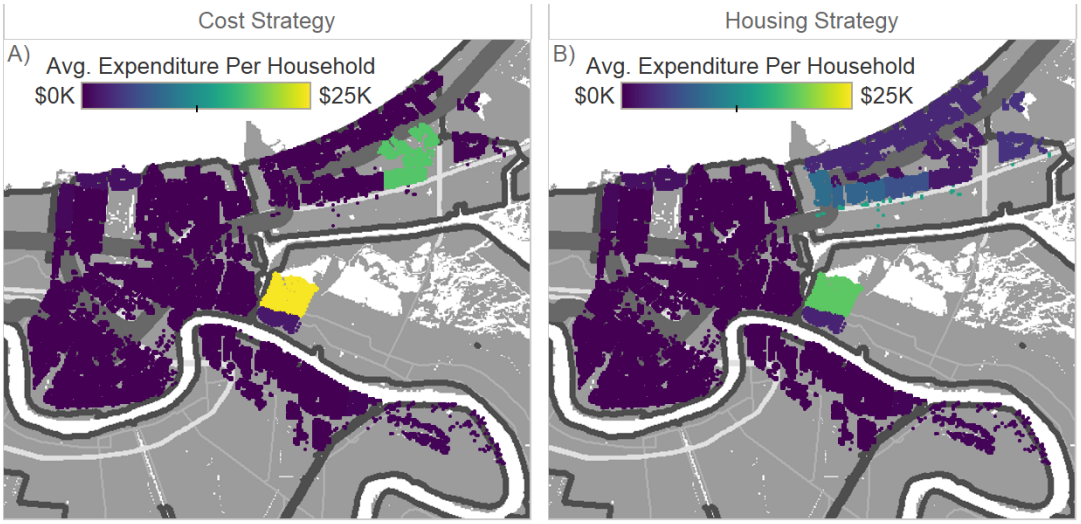

Figure 28: Map of average expenditure per household at the neighborhood level for cost and housing strategies. Panel A) shows cost strategy and panel B) shows housing strategy.

Intermediate strategies between the cost and housing strategies (i.e., strategies that place some weight on both DED and RLE reduction) perform well according to both metrics (Fig. 28). As preference for the housing strategy increases, residence loss equivalent reduction increases at the cost of direct economic damage reduction. Placing a small but non-zero weight on direct economic damage compared to the housing strategy reduces direct economic damage by an additional 0.5 percentage points at a cost of 0.02 percentage points of projected RLE. Placing a small weight on residence loss equivalents relative to the cost strategy reduces residence loss equivalents by an additional 4% at the cost of only a 0.04% decrease in direct economic damage reduction. As mentioned above, the question of how to balance protection of impoverished Black communities against protection of accrued wealth within Black communities is complicated considering broader racial wealth disparities. These results show that using deontologically egalitarian metrics in addition to traditional cost-benefit approaches can yield substantial benefits in the pursuit of prioritarian ideals of equity.

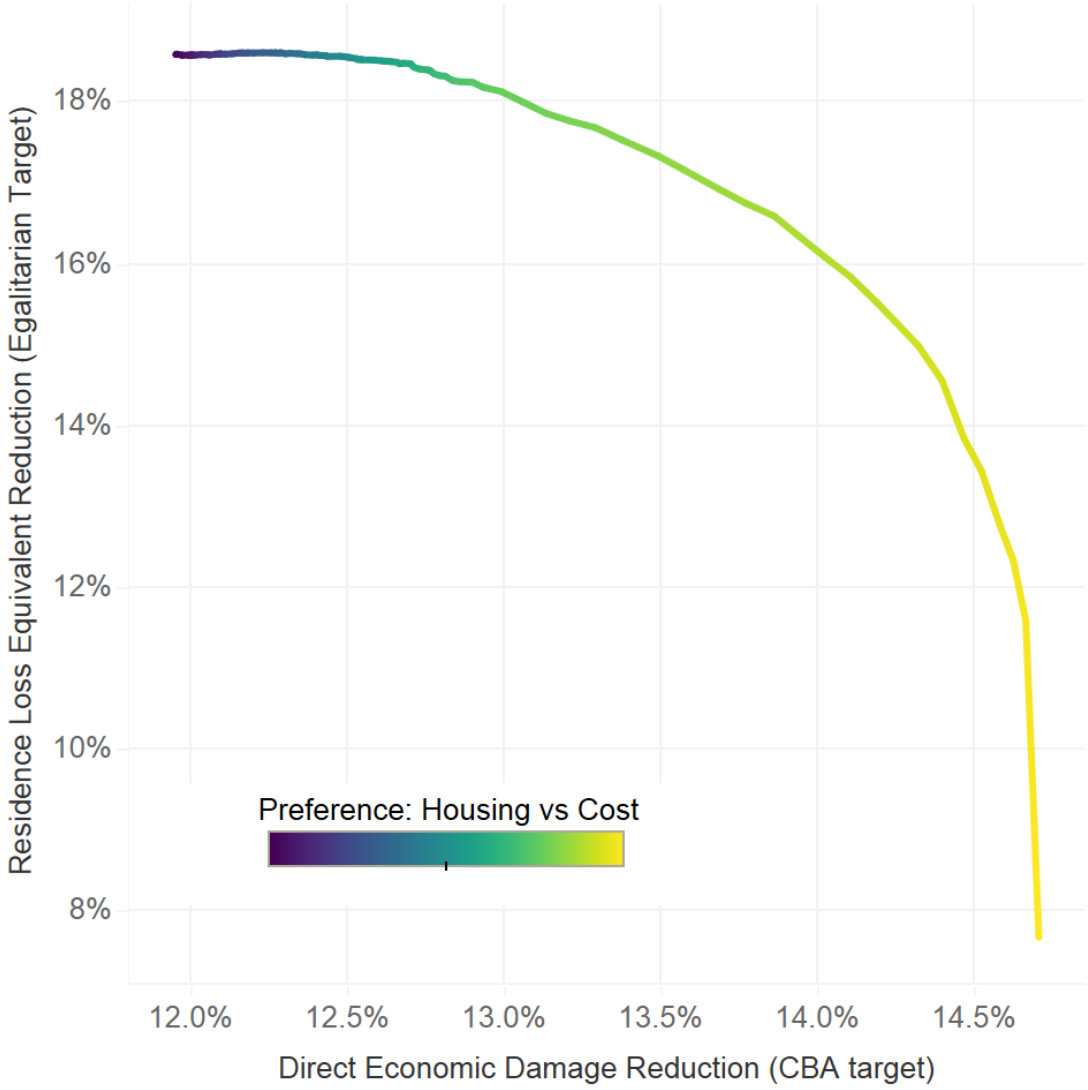

Figure 29: Pareto frontier between direct economic damage reduction and residence loss equivalents reduction for a fixed budget of 100 million US dollars. Segments in purple to the top-left reflect a high weight placed on residence loss equivalents reduction, as in the housing strategy, and segments in yellow towards the bottom-right reflect a high weight placed on direct economic damage reduction, as in the cost strategy.

4.4 Discussion

This case study illustrates how targeting aggregate economic efficiency in environmental risk mitigation projects can allocate few mitigation resources to impoverished areas that need them. This research proposes and demonstrates a potential remedy: targeting an alternative risk measure which deliberately places an equal weight on each residence irrespective of its economic value. This alternative measure satisfies deontological egalitarian notions of equity which other economically focused approaches may fall short of, and also improves distributional outcomes from a prioritarian perspective. While direct economic damage and residence loss equivalents are equally valid efficiency metrics, the use of residence loss equivalents to direct risk mitigation decisions can to some extent correct for historical underinvestment in risk mitigation to marginalized, less affluent communities without explicitly including socioeconomic demographics into the decision calculus.

Despite its evident advantages from a prioritarian perspective in dedicating mitigation resources towards impoverished communities, the implications of targeting residence loss equivalents as a risk metric from a corrective justice perspective are more nuanced than initially expected. It had been hypothesized that targeting direct economic damage would invest a large portion of mitigation resources towards wealthy, predominantly White communities, and that targeting residence loss equivalents would direct those resources to impoverished communities with larger Black populations. Instead, this study has found that the burden of flood risk by any metric in New Orleans falls so overwhelmingly on largely middle class and poor neighborhoods, all predominantly Black, that efficient flood risk mitigation targeting either metric invests the majority of mitigation resources towards predominantly Black neighborhoods. As a result, the choice between targeting direct economic damage and residence loss equivalents amounts to a choice between protecting middle class and poor neighborhoods. This complicates matters from a corrective justice perspective, as while prioritarian notions of distributive justice suggest prioritizing poorer neighborhoods and therefore targeting residence loss equivalents exclusively, doing so may deprive residents of middle class Black neighborhoods of the mitigation resources needed to maintain existing adaptive capacity in the face of flood risk.

Fortunately, this study identified compromise strategies that balance economic and egalitarian measures of cost effectiveness and perform well in both metrics. Compared to the status quo of seeking purely economic efficiency, substantial gains in these egalitarian measures can be achieved with only modest sacrifices in economic efficiency. This is highly useful to analysts and decision-makers interested in balancing the needs of poor and middle class neighborhoods facing severe flood risk in New Orleans; an intermediate strategy can efficiently provide protection to both poor and middle class neighborhoods which would be largely neglected by a strategy targeting one metric to the exclusion of the other. As the data elements required to calculate such egalitarian and hybrid risk measures are a proper subset of those used to calculate economic damages as per traditional CBA, their use can be integrated into existing analytical workflows in an operationally straightforward manner. This provides a tractable and parsimonious approach to better incorporate equity considerations into environmental risk management decisions, where budgets and schedules are tight and concerns for social equity are of increasing political salience. Hopefully, the proposed approach can help policy makers to design risk management strategies that better balance economic efficiency and equity outcomes when considering a range of perspectives on equity and justice.

5. CONCLUSION AND FUTURE WORK

This dissertation has presented meaningful, incremental advances in three distinct areas of coastal flood risk analysis, reflecting three distinct high-level questions that must be addressed in the process of applied coastal flood risk analysis in support of policy decision-making. 1) How do we characterize coastal flood hazard? 2) How do we identify high-performing strategies for mitigating the consequences of that hazard given specified performance metrics? 3) What performance metrics should we consider when evaluating mitigation strategies? In helping to advance the state of practice in answering these questions, this work makes contributions not in a particular facet of the process of applied coastal flood risk analysis for policy decision support, but holistically throughout the modeling process. It is my intention moving forward, in line with my previous work on Louisiana's 2023 Coastal Master Plan and the Louisiana Watershed Initiative, to collaborate with state and federal agencies as well as nonprofit and academic researchers to integrate all of these advances into a coherent applied modeling exercise to permit efficient, epistemically rigorous, and ethically sound coastal flood risk mitigation decision-making in the state of Louisiana and elsewhere to help manage flood risk driven by surge, waves, rainfall, and riverine dynamics.

The work presented in this dissertation has a number of limitations that I intend to address with future work. The optimal sampling discretization approach described in Chapter 2, as stated therein, is optimal in the sense of minimizing square error over boundary conditions but not in the sense of minimizing square error over maximum surface water elevations. Future work will evaluate the extent to which these differ; a large number of HEC-RAS outputs will be generated corresponding to previously generated boundary conditions, either with a more efficient HEC-RAS model combined with improved computing resources or with a metamodel, and comparisons will be made between the hazard distribution estimated by applying optimal sampling discretization to boundary conditions, the hazard distribution estimated by applying optimal sampling discretization to water surface elevations, and the hazard distribution generated when using the full set of HEC-RAS model outputs without optimal sampling discretization. This will clarify the amount of error induced in final hazard distributions by optimal sampling discretization.

Varying the number of clusters used in these comparisons will characterize this error as a function of the number of clusters used and therefore as a function of the number of HEC-RAS model runs which are computationally feasible when using optimal sampling discretization to accommodate tight computing budgets. Repeating these comparisons with random subsamples of available simulations will clarify the amount of error in final hazard curves as a function of the number of stochastic rainfall fields simulated. While these later questions were addressed in section 2.5.4 with respect to error in boundary condition distributions on a per-synthetic storm basis, performing similar experiments with respect to error in the final hazard distribution when sampling across rather than within synthetic storms would provide a more direct and actionable measure of modeling error and would better guide the allocation of computing resources in future modeling exercises.

The analysis presented in Chapter 3 was substantially constrained by computing resources. Future work with a more efficient implementation of the SWaMPS model or more computing resources will consider a wider range of sea level rise trajectories. Future work will also permit deeply uncertain climate parameters to vary independently, relaxing the restrictive assumption that climate trajectories can be straightforwardly rank-ordered by severity at a given time. Future work will also ensure that the modelled level of confidence in long-term climate expectations gained by near-term observations on the part of decision-makers (i.e., the granularity with which decision-makers can distinguish the severity of climate hazards which may follow the second decision stage) is informed by plausible mechanisms of observation and inference, rather than assumed as a matter of convenience. Future work will also permit the time at which the second decision stage is made to vary, governed by additional decision variables. Future work will treat the net present value of costs in each climate future as a distinct objective. Most importantly, future work will utilize an expanded implementation of SWaMPS which integrates deliberately adaptive, low-overhead mitigation strategies, in order to generate and compare robust and adaptive strategies which meaningfully outperform static strategies.

Future work building on Chapter 4 will quantify the uncertainty associated with downscaling census block group level demographic characteristics to the residence level. More sophisticated downscaling techniques will be investigated. Future work will expand the spatial scope of analysis to determine whether similar experiments in other locations

show the same tradeoff between protecting poor and middle-class neighborhoods or whether the tradeoff in some locations is between wealthy and poor neighborhoods as previously hypothesized. Expanding the spatial scope of the work will also help to clarify whether the favorable performance of intermediate strategies targeting both direct economic damage and residence loss equivalents generalizes to other locations. Future work may also investigate additional alternative risk measures such as expected loss in subjective utility [78].

Integration of the methods presented in Chapters 2, 3, and 4 represent another direction for future work, as the limitations addressed in each individual chapter remain present in the others.

APPENDIX A: TECHNICAL APPENDIX FOR SWAMPS MODEL

The following supplementary information was published alongside [48].

This Supplementary Information describes the design and methods used by the Surge and Waves Model for Protection Systems (SWaMPS), a process-based simulation model of storm surge-based flood risk on the interior of a levee/floodwall system. While SWaMPS could be configured to represent any real or proposed protection system subject to storm surge from tropical cyclones, our initial implementation is a representation of the Larose to Golden Meadow Hurricane Protection Project (LGM), located southwest of New Orleans, Louisiana.

The model runs a Monte Carlo simulation of 120 synthetic tropical cyclones whose hydrographs (i.e., storm surge elevation over time), significant wave heights, and wave periods are predicted as a function on the storm's parameters, location on the boundary of the protection system, and local mean sea level. The simulations include surge and wave overtopping, rainfall, pumping, and the potential for system failures in the form of catastrophic breaches. Each replicate calculates a resulting stillwater elevation (SWE) on the system interior. The annual exceedance probability (AEP) distribution function for flood elevations is calculated using joint probability methods to estimate the relative likelihood of each synthetic storm. The underlying mean arrival rate of storms is used to convert the cumulative distribution function for flooding, conditional upon a storm occurring, to the AEP curve. Direct economic damage is a function of a SWE, population, mitigation standard for nonstructural flood protection (e.g., elevating or floodproofing houses), and a rate of participation in the nonstructural practices. The resulting damage metrics include estimates of damage by return period (e.g., "100-year" damage exceedances) and the average annual losses (AAL) found by integrating over the damage AEP curve.

In creating SWaMPS, we were inspired by the Coastal Louisiana Risk Assessment (CLARA) model. CLARA is a peer-reviewed risk model used to evaluate the risk reduction impacts of flood protection measures considered by the State of Louisiana's *Comprehensive Master Plan for a Sustainable Coast* (hereafter, "Coastal Master Plan"). It

is described elsewhere in hundreds of pages of detail, so in this document, we provide a conceptual outline of the modeling framework while pointing to CLARA references for greater detail where methods are the same between models. We focus here on differences between the models and information about SWaMPS' calibration and validation. For example, one of the major efficiencies of SWaMPS is its representation of the LGM system using 12 reach segments, whereas CLARA uses 318.

Model Overview

SWaMPS implements a risk framework consisting of the following elements. Each is explained in more detail in the subsequent sections. Following a description of the risk framework, we also outline the cost model.

1) Hazard – the probability distribution of experiencing a tropical storm with specific characteristics, relating to the storm surge and waves experienced on the exterior boundary of the protection system
2) Vulnerability – the probability distribution of stillwater elevations (SWE) in the polder, expressed as an annual exceedance probability
3) Exposure – the collection of economic assets at risk of damage and their associated attributes, such as the topographic elevation at their location, asset type (e.g., single-family residence, industrial facility), foundation height, and square footage
4) Consequences – the direct economic damage caused by a given stillwater elevation, aggregated over the polder and expressed as a damage AEP curve or AAL, sometimes referred to as expected annual damage (EAD)

Risk in our representation of the LGM system is derived from Monte Carlo simulation of 120 synthetic storms. These correspond to the synthetic storms run through the CLARA model to estimate risk in future states of the world for Louisiana's 2017 Coastal Master Plan.

Surge and Wave Hazard. As noted in the main text, SWaMPS does not follow a standards-based design approach that requires calculation of a storm surge exceedance curve. However, in order to calculate the AEP curve for flooding on the system interior, we still need to know the relative likelihood of experiencing tropical cyclones with

characteristics like the 120 synthetic storms. We also need to predict the storm surge over time, significant wave height, and mean wave period for any synthetic storm. To allow for storms to be simulated in any future year, i.e., with any future assumptions about sea level rise, we developed a response surface to predict these characteristics as a function of storm parameters and mean local sea level.

Relative Likelihood of Synthetic Storms. SWaMPS adopts the joint probability distribution functions (PDF) recommended by previous coastal Louisiana flood risk studies using the joint probability method with optimal sampling (JPM-OS) [136]–[138]. This involves using the characteristics of historic observed tropical cyclones at landfall (as extracted from the National Hurricane Center's HURDAT2 dataset for storms from 1950 to present [139]) to fit a joint PDF over the storm parameters. Because the training storms do not vary by angle or forward velocity, this simplifies to fitting a joint distribution over the central pressure c_p, radius of maximum windspeed r, and landfall location in degrees longitude x:

$$\Lambda(c_p, r, x) = \Lambda_1 \cdot \Lambda_2 \cdot \Lambda_3$$

$$\Lambda_1 = f(c_p|x) = \frac{\partial}{\partial x}\left\{\exp\left\{-exp - \left[\frac{c_p - a_0(x)}{a_1(x)}\right]\right\}\right\}$$

$$\Lambda_2 = f(r|c_p) = \frac{1}{\sigma(\mathrm{c_p})\sqrt{2\pi}} e^{-\frac{(\bar{r}(c_p)-r)^2}{2\sigma^2(c_p)}}$$

$$\Lambda_3 = f(x) = \Phi(x)$$

$\Phi(x)$ is the empirical frequency distribution of longitudinal landfall locations (rounded to 1-degree bins). Using the HURDAT data, we construct $\Phi(x)$ and then fit $a_0(x)$, $a_1(x)$, $\bar{r}(c_p)$, and $\sigma^2(c_p)$ using maximum likelihood. This follows CLARA methodology outlined in more detail in other reports [138], [140].

The space of storm parameters (central pressure, radius of maximum windspeed, and landfall location) is partitioned into boxes defined by an interval for each parameter. The bounds of the intervals defining each box are set such that each box is occupied by a single synthetic storm with parameters equal to the average of the interval bounds. The joint PDF is then numerically integrated over each box to calculate a probability mass associated with each storm. The probability masses of each synthetic storm represent the relative

likelihoods of observing a storm with parameters contained within the box a synthetic storm inhabits.

Response Surface Predictions. The 2017 Coastal Master Plan ran hydrodynamic simulations of 60 synthetic storms using a coupled ADCIRC+SWAN model in a total of ten states of the world leveraged in this work. They consist of the current conditions (2015) landscape, and all combinations of the years 2025, 2040, and 2065 under three scenarios with different sea level conditions. Conveniently, this experimental design provided 10 distinct values for mean local sea level. Using the 10 sea levels and 600 synthetic storm simulations, we train a response surface for peak surge elevation, s_i, and significant wave height, w_i, at reach i, as a function of the storm parameters at landfall and mean local sea level:

$$\boldsymbol{s_i, w_i = \beta_0 + \beta_1 c_p + \beta_2 r + \beta_3 d_{il}^3 + \beta_4 d_{il}^2 + \beta_5 d_{il} + \beta_6 \sin \varphi_{il} + \beta_7 x + \beta_8 slr + \varepsilon_i}$$

where $\boldsymbol{c_p}$ is the central pressure, $\boldsymbol{r}$ the radius of maximum windspeed, $\boldsymbol{d_{il}}$ the distance between the centroid of reach i and the storm's landfall point l, $\boldsymbol{\varphi_{il}}$ the azimuthal angle between the landfall point l and reach i, $\boldsymbol{x}$ the location of landfall in degrees longitude, and $\boldsymbol{slr}$ the mean local sea level in a given state of the world.

This response surface is identical to that of CLARA, with the exception of adding the sea level rise term; CLARA fits a response surface independently for each state of the world using only the storm parameters and geospatial variables. Supplementary Figure 1 demonstrates that incorporating sea level rise into the response surface and training on all ten landscapes simultaneously does not meaningfully alter the predictive accuracy of the response surface. It shows the root mean squared error of predictions of storm surge from the 600 synthetic storms available in the training corpus. Colors indicate whether the response surface is fit independently on each of the ten states of the world (orange), as in the CLARA method; fitting with all states of the world simultaneously (blue); or applying leave-one-out cross-validation to a fit with all states of the world simultaneously (beige). The figure reports RMSE as a transect plot at each of the 318 points along the system centerlines in CLARA's representation of the LGM system. Performance is nearly identical, indicating that SWaMPS is justified in using the modified response surface to predict storm surge and waves; the leave-one-out cross-validation results suggest acceptable

performance for out-of-sample prediction, at least in states of the world with sea levels close to the range of scenarios modeled in the training set.

The hydrograph (i.e., runup and drawdown of surge over time) and mean wave period are fit as a function of a synthetic storm's peak surge elevation, using the 600 ADCIRC+SWAN simulations as training data and methods following the U.S. Army Corps of Engineers' Louisiana Coastal Protection and Restoration (LACPR) study [141]. SWaMPS uses a temporal resolution of 15 minutes and calculates surge for one day before peak surge and two days after.

Vulnerability. Once the surge and wave characteristics of each synthetic storm have been predicted, SWaMPS runs a Monte Carlo simulation of overtopping and levee/floodwall failure. Testing indicated that 100 replications of each synthetic storm generated virtually identical estimates of stillwater elevation (SWE) exceedances and expected annual damage (EAD) as 200, 500, or 1000 replications, so this analysis utilizes 100 replicates to minimize the model runtime. In each replicate, the total volume of water entering the polder is then added to rainfall volumes, pumping volumes are subtracted, and the resulting volume of water is translated to a stillwater elevation (SWE) using a stage-storage curve (i.e., a function, derived from a digital elevation model, that relates a volume of water within an enclosed spatial region to the SWE that results). To arrive at an AEP distribution for interior SWE, we describe the following steps of the model in greater detail: (i) calculating overtopping volumes for a given storm; (ii) calculating the probability and consequences of system failure(s); and (iii) aggregating the frequency distributions of SWE from the Monte Carlo simulations of each storm to the AEP curve.

Surge and Wave Overtopping. SWaMPS uses equations for surge and wave overtopping from the LACPR study that include coefficients defining stochastic variability in overtopping rates for a given surge elevation and significant wave height [141]. As described in the previous section, response surfaces to predict surge and wave characteristics were fit at each of CLARA's 318 points along the boundary of the LGM system. To simplify the system representation in SWaMPS wherein each of the 12 reaches

is treated as homogeneous, we needed to identify what surge and wave values should be assigned as representative of all CLARA points on a given reach.

Because overtopping dynamics are nonlinear, using a mean value for surge elevation, significant wave height, or wave period could introduce bias in the total overtopping volume from a storm. To calibrate a bias-correction factor β_i for each reach, we proposed to assign reach i a surge and wave characteristics for storm j of the form $\boldsymbol{\mu_{ij}} + \beta_i \boldsymbol{\sigma_{ij}}$, where $\boldsymbol{\mu_{ij}}$ and $\boldsymbol{\sigma_{ij}}$ are vectors containing the mean and standard deviation of the peak surge, significant wave height, and wave period values over the CLARA points located along reach i. We then minimized the sum of the squared error in overtopping volumes over all 600 training storms, as compared to the overtopping volumes calculated by CLARA. Where the CLARA overtopping volume for reach i and storm j is $\widetilde{V_{ij}}$ and the volume calculated by SWaMPS is $V(..)$, we thus identified an optimal bias-correction factor:

$$\beta_i^* = \underset{\beta_i}{\operatorname{argmin}} \sum_j \left(V\left(\boldsymbol{\mu_{ij}} + \beta_i \boldsymbol{\sigma_{ij}}\right) - \widetilde{V_{ij}}\right)^2$$

Upon finding bias-correction factors for every reach, we found little variation between them, so for simplicity, the model ultimately adopts a bias-correction factor found by minimizing over all reaches and storms simultaneously:

$$\beta^* = \underset{\beta}{\operatorname{argmin}} \sum_i \sum_j \left(V\left(\boldsymbol{\mu_{ij}} + \beta \boldsymbol{\sigma_{ij}}\right) - \widetilde{V_{ij}}\right)^2 = 0.03181$$

Figure 3 (explained in more detail below) indicates that SWaMPS does a good job of replicating statistical outcomes within a reasonable (i.e., policy-relevant) margin of error when compared to higher-fidelity models like CLARA. However, we caution that the model may not do as well at replicating inundation from individual observed storms, especially if they exhibit substantial variability in surge characteristics along reaches. For example, a CLARA-based analysis after Hurricane Ida in 2021 was consistent with an estimate that overtopping only occurred along 4% of the systems' levees. To replicate this accurately using SWaMPS, the system configuration may need to be adapted to have a larger number of shorter reaches.

System Fragility. The probability of system failure is expressed as a function of the overtopping rate, with the functional relationship taken from CLARA's IPET Low fragility curve [138]. However, SWaMPS departs from CLARA's fragility calculations in two key ways: (i) because we assume overtopping rates to be homogeneous at all points along one of the twelve reaches, the reach is treated as a series of characteristic lengths of 305 m (1000 ft), each of which can fail independently of other lengths; and (ii) failures are allowed to occur at any time during surge runup (whereas CLARA assumes failures occur at the time of peak surge). For each characteristic length, SWaMPS draws a uniform random variate. It then calculates the failure probability for each length given the overtopping rate at each timestep, and the number of breaches at a given time is taken to be the number of uniform variates that are exceeded by the failure probability up to and including that time. Failures are assumed to be catastrophic, i.e., full-length for a characteristic length, and full-depth. For breached regions, the free weir equations are used after setting the levee crest elevation to the elevation of the levee base. In each time period, the final overtopping volumes are calculated by calculating the length of the reach with and without breaches, then calculating overtopping assuming a constant rate over the 15-minute time step.

The possibility of failures occurring before the time of peak surge makes a modest but noticeable difference, as shown in Supplementary Figure 2. Blue indicates the new fragility implementation, while orange represents the CLARA approach to fragility (as implemented in SWaMPS, so that the only difference is whether failures are constrained to occur only at peak surge). Each trace of the same color represents a different macroreplication of the case (i.e., initializing the model with a different random seed). Results are shown in the top pane for current conditions and for Year 50 of the High environmental scenario in the bottom. Differences are larger in the latter case, where the SWaMPS approach generally increases SWE exceedances by 0.1-0.3 m over a range of return periods.

Annual Exceedance Probabilities. Once each storm in a given case has been simulated, SWaMPS generates the AEP curve by aggregating the results using a simplified version of the approach used by the CLARA model; full details are given in Fischbach, *et al* (2017) [138]. As described in the Hazard section, a joint PDF is estimated using HURDAT2 data.

The resulting probability masses associated with each synthetic storm are combined with the frequency distribution of SWE over the Monte Carlo replicates to form a cumulative distribution function (CDF), conditional upon a storm occurring. Storms are modeled as a Poisson process, and the observed frequency of historical tropical cyclones impacting coastal Louisiana is used to estimate the probability of observing $n \in \{0, \mathbb{N}\}$ storms in a given year. The law of total probability is then applied to convert the CDF of SWE for a single storm event to a CDF for the maximum annual SWE. SWaMPS follows this same approach but does not take CLARA's extra step of bootstrapping the historical record and repeating the process to generate confidence bounds around the point estimates of SWE exceedances.

Supplementary Figure 3 compares SWaMPS' estimates of SWE exceedances to CLARA's median estimates and 80% confidence bounds in current conditions and Year 50 of the High scenario. In this figure, both models use the IPET Low fragility curve, and SWaMPS is constrained to allow failures only at the time of peak surge (in order to make as close a comparison as possible). The results show that the bias-correction of surge and wave characteristics results in a well-calibrated model of vulnerability within the polder.

Exposure and Consequences. SWaMPS, like CLARA, models direct economic damage associated with storm-surge based flooding. In fact, the SWaMPS damage model consists of a lookup table of precomputed damage results from CLARA. Results were calculated at 0.1-foot intervals for SWE and can be modified using a population multiplier and assumptions about mitigation measures such as elevating houses, floodproofing, and buying out high-risk properties. All assets but roads and agricultural crops are assumed to scale with the population multiplier, which takes a value of 1 in the current conditions (2015) case as its baseline. Mitigation measures are specified as a mitigation standard corresponding to the NAVD88 elevation to which policy makers desire protection. Not all types of assets are eligible for each measure, but for eligible assets, floodproofing is applied if it would take 3 feet or less of elevation above grade to raise the foundation to the mitigation standard; elevation-in-place is applied where the mitigation standard can be achieved by foundation heights 3-12 feet above grade; and buyouts occur if more than 12

feet of elevation would be required. Another system parameter specifies the assumed rate of participation in these measures.

Within a given state of the world and time period, SWaMPS obtains the damage associated with the SWE exceedances for 22 return periods ranging from AEP values of 0.2 to 0.0005 (i.e., 5-year to 2,000-year). Average annual losses are calculated by numerically integrating (i.e., taking a weighted average) the damage exceedance curve.

The analysis presented in this paper assumes no additional mitigation measures are applied and that population remains at the current level. We made these choices to isolate the impact of climate change and to focus on the risk-informed levee/floodwall design framework.

Costs. To calculate the present value of costs for structural protection and the additional risk mitigation measures outlined in the previous section, we replicated the cost models and unit cost assumptions used by the State of Louisiana in their 2017 Coastal Master Plan. These equations and assumptions are detailed in the plan's Project Definitions appendix [142], but we used the actual costing spreadsheet templates for two proposed upgrades to the existing LGM protection system as a reference to ensure the accuracy of formulas and unit costs. Protection system costs include planning and design, mobilization, construction, and annual operations and maintenance over a user-specified planning horizon. The present value of this stream of costs depends on a user-specified discount rate.

The model is initialized by a configuration file that allows for an existing system to be in place on the landscape. Costs are therefore implicitly defined as the cost of upgrading an already-present system from its existing reach elevations to an upgraded set of design elevations. In this analysis, we designed a new system by configuring the model to assume the "existing" system had design heights of zero meters above grade.

We note that the cost model for our representation of the Larose to Golden Meadow system is built on a considerable amount of pre-existing knowledge and the design work that informed the existing system. Using SWaMPS to optimize design heights for a new system would still require substantial effort to assess factors like where floodwalls are preferable to levees, if pumping systems will be necessary, whether design costs will

include interior canals or other drainage structures, and recommended levee geometries as informed by soil boring samples.

Software Implementation and Portability. SWaMPS was developed in Python 3 using information about the LGM Hurricane Protection Project, but the major components are designed as modular input files and functions that could be modified to represent another system or to generate additional metrics. The protection system is defined using a reach configuration file that provides geometry and other characteristics for each reach (e.g., length, crown height and width, front and backside slope, height and thickness of T-walls on top of the levee [wherever present], etc.). This also includes some construction information relevant to costing the project, such as overbuild to account for soil compaction and the width of inspection corridors on either side that must be cleared. Cost calculations then also depend on another input file defining unit costs for turf, fill, concrete and piles; maintenance assumptions; planning, engineering and design costs; contingency, etc.
Another example of a modular components is the response surface mapping the tropical cyclone and environmental parameters to surge and wave characteristics. The LGM response surface was trained on ADCIRC+SWAN simulations supporting Louisiana's 2017 Coastal Master Plan which were conveniently available, but more budget- or computationally-constrained studies could reduce costs and complexity by using more simplistic models or by assuming that hazard is stationary over time (eliminating the need to run future sea level and land subsidence scenarios).

The damage model was also programmed here as a lookup table from results using the CLARA damage module, but this too could be replaced with any source of decision-relevant outcome metrics indicating system performance.

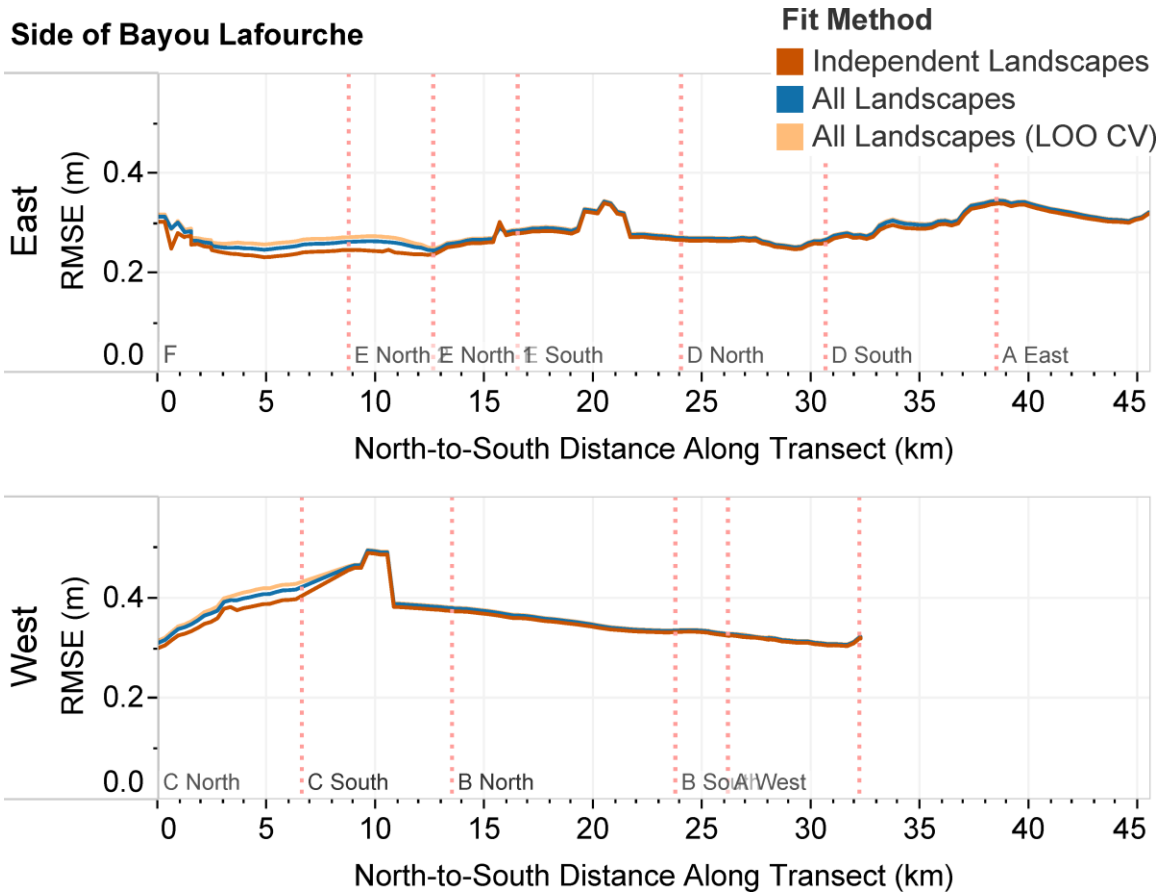

Supplementary Figure 1. Root mean squared error of response surface predictions of storm surge over all training storms. Colors indicate whether the response surface is fit independently on each of the ten states of the world (orange), fitting with all states of the world simultaneously (blue), or applying leave-one-out cross-validation to a fit with all states of the world simultaneously (beige).

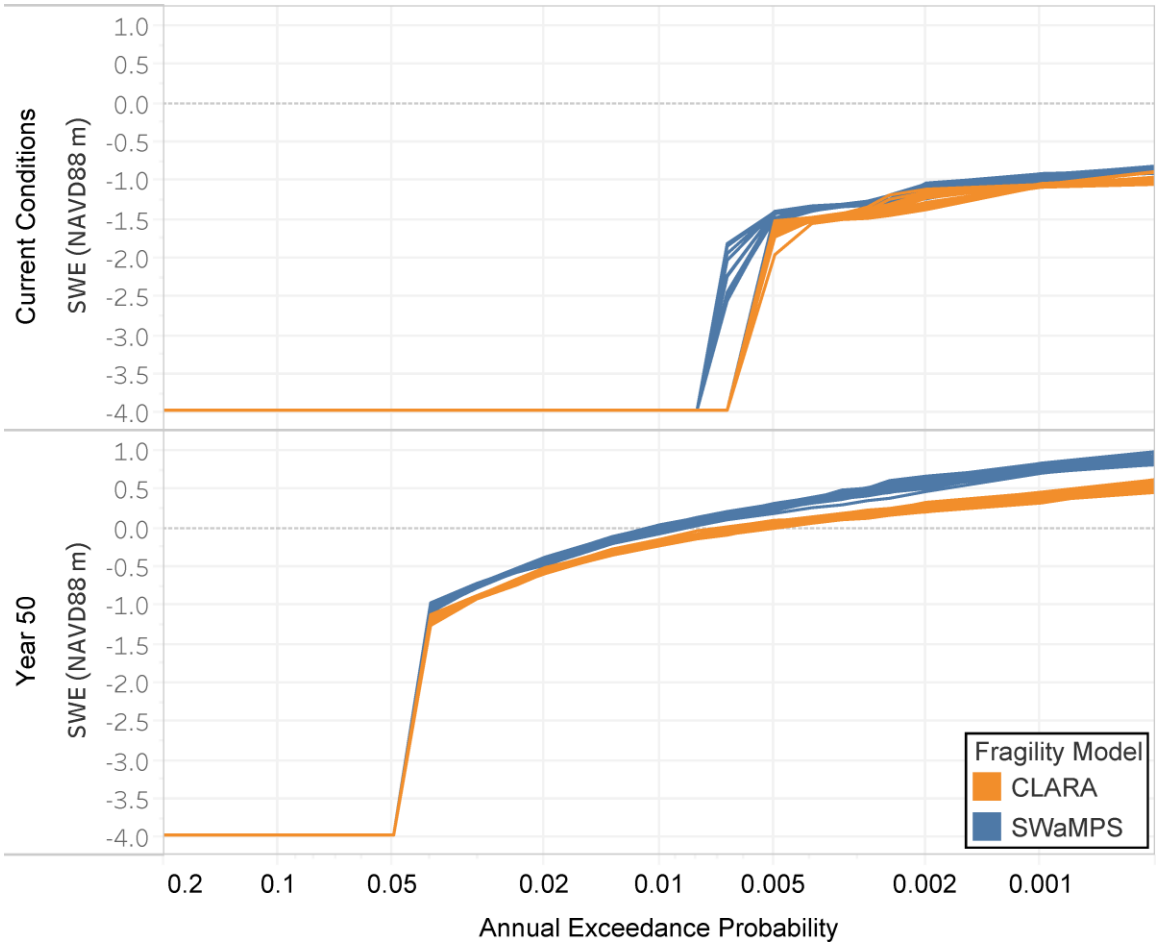

Supplementary Figure 2. Stillwater elevation by AEP under current conditions and Year 50 of the High environmental scenario, according to the fragility model used. Each trace of the same color represents a different macro replication of the case.

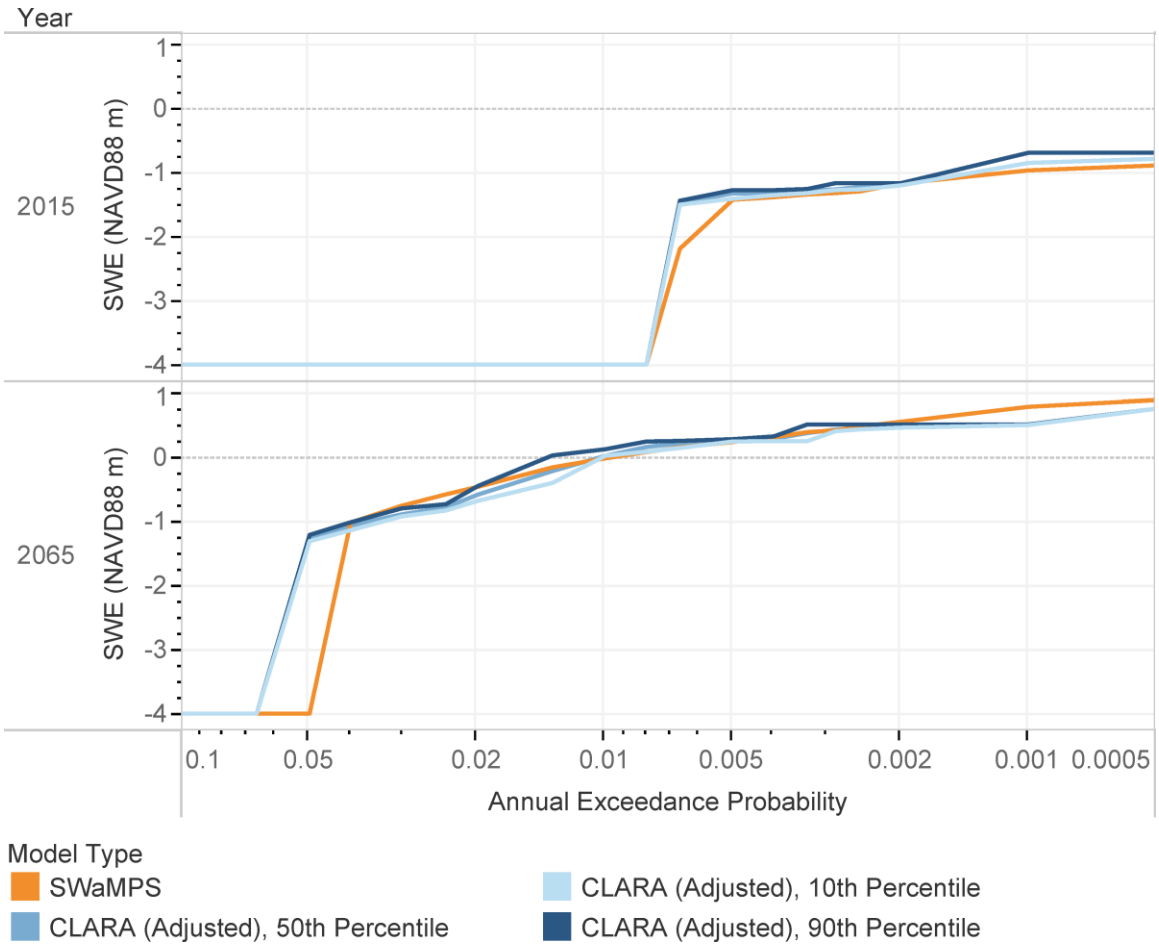

Supplementary Figure 3. Comparison of stillwater elevations by annual exceedance probability estimated by SWaMPS and CLARA, 2015 and 2065 (High scenario). Shades of blue indicate CLARA's median estimate and 80% confidence bounds; orange indicates the SWaMPS estimates.

Table S1. Standard deviation of design heights (as a percentage of the mean design height) among all strategies with lower cost and lower residual risk than the existing protection system.

	Future State of the World			
Reach	Stationary Risk	Low Scenario	Medium Scenario	High Scenario
A East	10.2%	2.4%	2.6%	1.7%
A West	3.7%	2.5%	3.5%	4.1%
B North	1.1%	6.1%	2.4%	1.5%
B South	1.8%	6.2%	3.0%	6.6%
C North	4.0%	3.7%	2.5%	2.3%
C South	1.7%	1.5%	2.8%	4.6%
D North	1.7%	2.4%	2.6%	1.8%
D South	9.3%	5.9%	2.5%	2.6%
E North 1	6.5%	5.7%	1.9%	2.9%
E North 2	3.9%	3.6%	4.5%	4.4%
E South	1.4%	1.4%	2.4%	3.1%
F	1.3%	2.6%	2.0%	2.2%

Table S2. Regression coefficients, statistical significance of covariates predicting risk-based design heights, and r^2 values for each reach.

Reach	β_0: $intercept$	β_1: $\mathcal{L}$	β_2: int	β_3: α	β_4: slr	β_5: $\mathcal{L} \times int$	β_6: $\mathcal{L} \times \alpha$	β_7: $\mathcal{L} \times slr$	r^2
A East	4.67E+0***	-5.03E-5***	2.00E-1	1.31E-1*	6.47E-2	-2.32E-6	1.79E-5***	2.01E-5***	0.584
A West	4.11E+0***	-4.18E-5***	8.45E-1*	1.26E-1*	3.97E-2	-1.76E-5	1.50E-5***	1.81E-5***	0.521
B North	4.18E+0***	-4.18E-5***	-2.02E-1	-9.23E-3	4.39E-1***	5.55E-6	2.06E-5***	1.07E-5***	0.457
B South	4.66E+0***	-5.25E-5***	5.01E-1	1.45E-1*	-2.33E-1***	-8.03E-6	1.70E-5***	2.66E-5***	0.512
C North	3.24E+0***	-2.17E-5***	6.37E-2	6.70E-2	7.63E-1***	-1.72E-6	1.30E-5***	-1.37E-6	0.583
C South	3.49E+0***	-2.90E-5***	6.70E-2	2.35E-1***	6.60E-1***	4.28E-6	8.87E-6**	4.76E-6*	0.457
D North	4.12E+0***	-3.93E-5***	4.12E-1	2.56E-2	2.89E-1***	-7.91E-6	1.81E-5***	1.28E-5***	0.628
D South	4.75E+0***	-5.24E-5***	2.77E-1	-2.64E-2	-2.66E-1***	-1.91E-6	2.09E-5***	2.70E-5***	0.544
E North 1	4.39E+0***	-4.14E-5***	-5.03E-1	-3.76E-2	2.68E-1***	3.61E-6	2.05E-5***	1.32E-5***	0.532
E North 2	3.65E+0***	-3.39E-5***	-2.10E-1	8.46E-2	5.29E-1***	7.03E-6	1.48E-5***	7.96E-6***	0.557
E South	4.57E+0***	-4.94E-5***	1.26E-1	-2.25E-2	-1.31E-1**	4.44E-6	1.90E-5***	2.36E-5***	0.592
F	3.36E+0***	-2.73E-5***	1.36E-1	4.06E-2	7.45E-1***	-2.93E-6	1.55E-5***	1.94E-6	0.555

* $\mathbf{p < 0.05}$; ** $\mathbf{p < 0.01}$; *** $\mathbf{p < 0.001}$; average $\boldsymbol{r^2}$ over all reaches = 0.543

$\boldsymbol{\mathcal{L}}$ = expected losses over 50 years; $\boldsymbol{int}$ = change to future mean storm intensity; $\boldsymbol{\alpha}$ = change to future storm frequency; $\boldsymbol{slr}$ = local sea level rise over 50 years

APPENDIX B: SUPPLEMENTAL TABLES AND FIGURES

Table 2: Aggregate single-family residential structure and demographic information for New Orleans neighborhoods (2 of 2)

	Average Replacement Cost	Number of Structures	Proportion Black	Proportion White	Poverty Rate
ALGIERS POINT	$250K	394	9%	90%	8%
AUDUBON	$617K	3,073	8%	82%	19%
B. W. COOPER	$80K	14	99%	0%	46%
BAYOU ST. JOHN	$350K	578	27%	71%	15%
BEHRMAN	$85K	1,993	79%	14%	30%
BLACK PEARL	$312K	297	19%	70%	16%
BROADMOOR	$231K	1,004	56%	40%	15%
BYWATER	$233K	616	26%	69%	19%
CENTRAL BUSINESS DI..	$1,607K	25	32%	63%	21%
CENTRAL CITY	$166K	877	70%	26%	39%
CITY PARK	$436K	445	4%	89%	7%
DESIRE AREA	$63K	550	98%	1%	64%
DILLARD	$91K	1,576	90%	8%	31%
DIXON	$68K	217	86%	10%	35%
EAST CARROLLTON	$303K	675	26%	62%	15%
EAST RIVERSIDE	$289K	505	30%	67%	7%
FAIRGROUNDS	$223K	1,155	53%	40%	15%
FILMORE	$218K	2,281	70%	25%	13%
FISCHER DEV	$70K	106	97%	3%	58%
FLORIDA AREA	$70K	513	95%	5%	28%
FRENCH QUARTER	$681K	439	5%	93%	13%
FRERET	$237K	255	49%	44%	25%
GARDEN DISTRICT	$781K	407	5%	94%	6%
GENTILLY TERRACE	$142K	3,423	77%	18%	18%
GENTILLY WOODS	$103K	1,052	70%	24%	16%
GERT TOWN	$135K	242	84%	13%	35%
HOLLYGROVE	$76K	1,265	69%	28%	31%
HOLY CROSS	$142K	888	65%	31%	30%
IRISH CHANNEL	$288K	441	22%	75%	11%
LAKE TERRACE & OAKS	$229K	845	29%	58%	10%
LAKESHORE - LAKE VI..	$378K	1,614	4%	92%	1%
LAKEVIEW	$314K	3,345	4%	91%	4%
LAKEWOOD	$380K	816	2%	96%	1%
LEONIDAS	$184K	1,298	52%	39%	27%
LITTLE WOODS	$123K	9,066	94%	3%	31%

Table 3: Aggregate single-family residential structure and demographic information for New Orleans neighborhoods (2 of 2)

	Average Replacement Cost	Number of Structures	Proportion Black	Proportion White	Poverty Rate
LOWER GARDEN DIST..	$538K	408	14%	77%	15%
LOWER NINTH WARD	$77K	1,392	91%	7%	35%
MARIGNY	$346K	318	13%	79%	14%
MARLYVILLE - FONTAI..	$264K	1,439	25%	68%	16%
McDONOGH	$117K	539	74%	23%	23%
MID-CITY	$247K	979	39%	54%	20%
MILAN	$306K	636	50%	46%	19%
MILNEBURG	$133K	1,420	87%	8%	15%
NAVARRE	$271K	854	6%	87%	8%
NEW AURORA - ENGLI..	$212K	1,778	71%	15%	15%
OLD AURORA	$124K	6,214	57%	32%	17%
PINES VILLAGE	$83K	928	97%	2%	34%
PLUM ORCHARD	$80K	1,359	97%	1%	39%
PONTCHARTRAIN PARK	$122K	874	97%	2%	25%
READ BLVD EAST	$241K	3,329	81%	3%	12%
READ BLVD WEST	$104K	1,820	95%	2%	21%
SEVENTH WARD	$103K	1,523	77%	20%	39%
ST. ANTHONY	$130K	1,054	80%	16%	27%
ST. BERNARD AREA	$134K	218	91%	3%	35%
ST. CLAUDE	$110K	1,635	62%	33%	30%
ST. ROCH	$84K	1,370	82%	14%	37%
ST. THOMAS DEV	$214K	87	77%	19%	50%
TALL TIMBERS - BREC..	$169K	2,478	82%	14%	28%
TOURO	$496K	370	13%	79%	19%
TREME - LAFITTE	$215K	495	54%	41%	29%
TULANE - GRAVIER	$166K	177	68%	24%	39%
U.S. NAVAL BASE	$294K	680	74%	20%	29%
UPTOWN	$469K	1,181	15%	80%	12%
VIAVANT - VENETIAN I..	$200K	91	81%	15%	73%
VILLAGE DE LEST	$97K	1,849	49%	3%	18%
WEST END	$326K	997	14%	77%	13%
WEST LAKE FOREST	$122K	674	93%	1%	31%
WEST RIVERSIDE	$345K	940	19%	77%	9%
WHITNEY	$106K	631	79%	18%	15%

REFERENCES

[1] K. Hall, "Expected costs of damage from hurricane winds and storm-related flooding," *Congressional Budget Office: Washington, DC, USA*, pp. 1–48, 2019.

[2] B. E. Montz and G. A. Tobin, "Livin'large with levees: Lessons learned and lost," *Natural Hazards Review*, vol. 9, no. 3, pp. 150–157, 2008.

[3] F. E. M. A. M. D. R. A. Branch, *The national levee challenge: Levees and the FEMA flood map modernization initiative*. Interagency Levee Policy Review Commitee, 2006.

[4] J. R. Fischbach, D. R. Johnson, M. T. Wilson, N. B. Geldner, and C. Stelzner, "Draft Coastal Master Plan: Risk Assessment Model Improvements," *Version I*, pp. 1–78, 2023.

[5] D. R. Johnson, J. R. Fischbach, and D. S. Ortiz, "Estimating surge-based flood risk with the coastal Louisiana risk assessment model," *Journal of Coastal Research*, no. 67 (10067), pp. 109–126, 2013.

[6] D. R. Johnson, J. R. Fischbach, N. B. Geldner, M. T. Wilson, C. Story, and J. Wang, "Draft Coastal Master Plan: Attachment C11: 2023 Risk Model," *Version*, vol. 2, p. 33, 2023.

[7] J. Zscheischler *et al.*, "Future climate risk from compound events," *Nature Climate Change*, vol. 8, no. 6, pp. 469–477, 2018.

[8] H. Moftakhari, D. F. Muñoz, J. Y. Song, A. Alipour, and H. Moradkhani, "Challenges for Appropriate Characterization of Compound Coastal Hazards," in *Geo-Extreme 2021*, 2021, pp. 58–68.

[9] A. AghaKouchak *et al.*, "Climate extremes and compound hazards in a warming world," *Annual Review of Earth and Planetary Sciences*, vol. 48, pp. 519–548, 2020.

[10] A. Sebastian, "Compound flooding," in *Coastal Flood Risk Reduction*, Elsevier, 2022, pp. 77–88.

[11] F. L. Santiago-Collazo, M. V. Bilskie, and S. C. Hagen, "A comprehensive review of compound inundation models in low-gradient coastal watersheds," *Environmental Modelling & Software*, vol. 119, pp. 166–181, 2019.

[12] N. C. Nadal-Caraballo, V. M. Gonzalez, and L. Chouinard, "Storm Recurrence Rate Models for Tropical Cyclones: Report 1," ENGINEER RESEARCH AND DEVELOPMENT CENTER VICKSBURG MSMCGILL UNIV MONTREAL …, 2019.

[13] M. Bensi and T. Weaver, “Evaluation of tropical cyclone recurrence rate: factors contributing to epistemic uncertainty,” *Natural Hazards*, vol. 103, no. 3, pp. 3011–3041, 2020.

[14] H. Kim, G. Villarini, R. Jane, T. Wahl, S. Misra, and A. Michalek, “On the generation of high-resolution probabilistic design events capturing the joint occurrence of rainfall and storm surge in coastal basins,” *International Journal of Climatology*, 2022.

[15] A. Couasnon, A. Sebastian, and O. Morales-Nápoles, “A copula-based Bayesian network for modeling compound flood hazard from riverine and coastal interactions at the catchment scale: An application to the Houston Ship Channel, Texas,” *Water*, vol. 10, no. 9, p. 1190, 2018.

[16] Z. Hao and V. P. Singh, “Compound events under global warming: a dependence perspective,” *Journal of Hydrologic Engineering*, vol. 25, no. 9, p. 03120001, 2020.

[17] M. Hofert, I. Kojadinovic, M. Maechler, J. Yan, J. G. Nešlehová (evTestK()), and R. M. (fitCopula ml(): code for free mixCopula weight parameters), “copula: Multivariate Dependence with Copulas.” Jan. 25, 2023. Accessed: Feb. 08, 2023. [Online]. Available: https://CRAN.R-project.org/package=copula

[18] A. Gori, N. Lin, and D. Xi, “Tropical cyclone compound flood hazard assessment: From investigating drivers to quantifying extreme water levels,” *Earth’s Future*, vol. 8, no. 12, p. e2020EF001660, 2020.

[19] A. Gori and N. Lin, “Projecting compound flood hazard under climate change with physical models and joint probability methods,” *Earth’s Future*, p. e2022EF003097, 2022.

[20] G. R. Toro, D. T. Resio, D. Divoky, A. W. Niedoroda, and C. Reed, “Efficient joint-probability methods for hurricane surge frequency analysis,” *Ocean Engineering*, vol. 37, no. 1, pp. 125–134, 2010.

[21] Y. Peng, K. Chen, H. Yan, and X. Yu, “Improving flood-risk analysis for confluence flooding control downstream using Copula Monte Carlo method,” *Journal of Hydrologic Engineering*, vol. 22, no. 8, p. 04017018, 2017.

[22] G. R. Toro, A. W. Niedoroda, C. W. Reed, and D. Divoky, “Quadrature-based approach for the efficient evaluation of surge hazard,” *Ocean Engineering*, vol. 37, no. 1, pp. 114–124, 2010.

[23] K. Yin, S. Xu, and W. Huang, “Estimating extreme sea levels in Yangtze Estuary by quadrature joint probability optimal sampling method,” *Coastal Engineering*, vol. 140, pp. 331–341, 2018.

[24] M. Bartles *et al.*, “Hydrologic Modeling System HEC-HMS User’s Manual,” CPD-74A. [Online]. Available: https://www.hec.usace.army.mil/confluence/hmsdocs/hmsum/latest/report-documentation-page

[25] Z. Cobell and H. Roberts, “Storm Surge and Waves Model Updatesd for the 2023 Coastal Master Plan,” Coastal Protection and Restoration Authority, 2021. [Online]. Available: https://coastal.la.gov/wp-content/uploads/2021/03/StormSurge_Waves_Report_Jan2021.pdf

[26] G. W. Brunner, “HEC-RAS, River Analysis System Hydraulic Reference Manual,” US Army Corps of Engineers Hydrologic Engineering Center (HEC), CPD-69, 2020.

[27] D. T. Resio *et al.*, “White paper on estimating hurricane inundation probabilities,” 2007.

[28] G. Villarini, W. Zhang, P. Miller, D. R. Johnson, L. E. Grimley, and H. J. Roberts, “Probabilistic rainfall generator for tropical cyclones affecting Louisiana,” *International journal of climatology*, vol. 42, no. 3, pp. 1789–1802, 2022.

[29] N. C. Nadal-Caraballo, M. O. Campbell, V. M. Gonzalez, M. J. Torres, J. A. Melby, and A. A. Taflanidis, “Coastal hazards system: a probabilistic coastal hazard analysis framework,” *Journal of Coastal Research*, vol. 95, no. SI, pp. 1211–1216, 2020.

[30] T. Chai and R. R. Draxler, “Root mean square error (RMSE) or mean absolute error (MAE)?–Arguments against avoiding RMSE in the literature,” *Geoscientific model development*, vol. 7, no. 3, pp. 1247–1250, 2014.

[31] D. Pollard, “Strong consistency of k-means clustering,” *The Annals of Statistics*, pp. 135–140, 1981.

[32] D. Johnson, N. Geldner, H. Roberts, B. Yuill, and N. Young, “Recurrence Analysis: Amite River Basin Pilot Study Technical Memorandum.” (unpublished), 2021.

[33] R. J. Lempert, “Robust decision making (RDM),” in *Decision making under deep uncertainty*, Springer, Cham, 2019, pp. 23–51.

[34] T. Erfani, K. Pachos, and J. J. Harou, “Real-options water supply planning: Multistage scenario trees for adaptive and flexible capacity expansion under probabilistic climate change uncertainty,” *Water Resources Research*, vol. 54, no. 7, pp. 5069–5087, 2018.

[35] J. Buurman and V. Babovic, “Adaptation Pathways and Real Options Analysis: An approach to deep uncertainty in climate change adaptation policies,” *Policy and Society*, vol. 35, no. 2, pp. 137–150, 2016.

[36] M. Haasnoot *et al.*, “Investments under non-stationarity: economic evaluation of adaptation pathways,” *Climatic change*, vol. 161, no. 3, pp. 451–463, 2020.

[37] D. Van Dantzig, “Economic decision problems for flood prevention,” *Econometrica: Journal of the Econometric Society*, pp. 276–287, 1956.

[38] S. N. Jonkman, M. Kok, M. k Van Ledden, and J. K. Vrijling, “Risk-based design of flood defence systems: a preliminary analysis of the optimal protection level for the New Orleans metropolitan area,” *Journal of Flood Risk Management*, vol. 2, no. 3, pp. 170–181, 2009.

[39] G. G. Garner and K. Keller, “Using direct policy search to identify robust strategies in adapting to uncertain sea-level rise and storm surge,” *Environmental Modelling & Software*, vol. 107, pp. 96–104, Sep. 2018, doi: 10.1016/j.envsoft.2018.05.006.

[40] J. M. Kind, “Economically efficient flood protection standards for the N etherlands,” *Journal of Flood Risk Management*, vol. 7, no. 2, pp. 103–117, 2014.

[41] K. A. H. Slijkhuis, P. Van Gelder, and J. K. Vrijling, “Optimal dike height under statistical-construction-and damage uncertainty,” *Structural Safety and Reliability*, vol. 7, pp. 1137–1140, 1997.

[42] C. J. Eijgenraam, “Optimal safety standards for dike-ring areas,” *Discussion paper nr. 62, ISBN 90-5833-267-5*, 2006.

[43] J. K. Vrijling, “Probabilistic design of water defense systems in The Netherlands,” *Reliability engineering & system safety*, vol. 74, no. 3, pp. 337–344, 2001.

[44] R. L. Sriver, R. J. Lempert, P. Wikman-Svahn, and K. Keller, “Characterizing uncertain sea-level rise projections to support investment decisions,” *PLoS One*, vol. 13, no. 2, p. e0190641, 2018.

[45] T. E. Wong and K. Keller, “Deep uncertainty surrounding coastal flood risk projections: A case study for New Orleans,” *Earth’s Future*, vol. 5, no. 10, pp. 1015–1026, 2017.

[46] T. E. Wong, A. M. Bakker, and K. Keller, “Impacts of Antarctic fast dynamics on sea-level projections and coastal flood defense,” *Climatic Change*, vol. 144, no. 2, pp. 347–364, 2017.

[47] J. R. Fischbach *et al.*, “2017 Coastal Master Plan: Attachment C3-25: Storm Surge and Risk Assessment.” 2016.

[48] D. R. Johnson, J. Wang, N. B. Geldner, and A. B. Zehr, “Rapid, risk-based levee design framework for greater risk reduction at lower cost than standards-based design,” *Journal of Flood Risk Management*, p. e12786, 2022, doi: 10.1111/jfr3.12786.

[49] J. H. Kwakkel, S. Eker, and E. Pruyt, “How robust is a robust policy? Comparing alternative robustness metrics for robust decision-making,” *Robustness analysis in decision aiding, optimization, and analytics*, pp. 221–237, 2016.

[50] M. Haasnoot, J. H. Kwakkel, W. E. Walker, and J. Ter Maat, “Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world,” *Global environmental change*, vol. 23, no. 2, pp. 485–498, 2013.

[51] M. Haasnoot, H. Middelkoop, E. Van Beek, and W. P. A. Van Deursen, “A method to develop sustainable water management strategies for an uncertain future,” *Sustainable Development*, vol. 19, no. 6, pp. 369–381, 2011.

[52] C. M. Regan *et al.*, “Real options analysis for land use management: Methods, application, and implications for policy,” *Journal of environmental management*, vol. 161, pp. 144–152, 2015.

[53] J. R. Kasprzyk, S. Nataraj, P. M. Reed, and R. J. Lempert, “Many objective robust decision making for complex environmental systems undergoing change,” *Environmental Modelling & Software*, vol. 42, pp. 55–71, 2013.

[54] S. Eker and J. H. Kwakkel, “Including robustness considerations in the search phase of Many-Objective Robust Decision Making,” *Environmental Modelling & Software*, vol. 105, pp. 201–216, 2018.

[55] E. H. Beh, F. Zheng, G. C. Dandy, H. R. Maier, and Z. Kapelan, “Robust optimization of water infrastructure planning under deep uncertainty using metamodels,” *Environmental Modelling & Software*, vol. 93, pp. 92–105, 2017.

[56] T. Roach, Z. Kapelan, R. Ledbetter, and M. Ledbetter, "Comparison of robust optimization and info-gap methods for water resource management under deep uncertainty," *Journal of Water Resources Planning and Management*, vol. 142, no. 9, p. 04016028, 2016.

[57] B. Shavazipour, J. H. Kwakkel, and K. Miettinen, "Multi-scenario multi-objective robust optimization under deep uncertainty: A posteriori approach," *Environmental Modelling & Software*, vol. 144, p. 105134, 2021.

[58] B. Shavazipour, D. Podkopaev, and K. Miettinen, "Interactive decision support and trade-off analysis for sustainable forest landscape planning under deep uncertainty," *Canadian Journal of Forest Research*, vol. 52, no. 11, pp. 1423–1438, 2022.

[59] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE transactions on evolutionary computation*, vol. 6, no. 2, pp. 182–197, 2002.

[60] R. J. Lempert, S. W. Popper, and S. C. Banks, "The challenge of long-term policy analysis," *Shaping the next one hundred years: new methods for quantitative, long-term policy analysis, 1st ed. RAND Corporation*, pp. 1–9, 2003.

[61] T. E. Wong, A. M. Bakker, K. Ruckert, P. Applegate, A. Slangen, and K. Keller, "BRICK v0. 2, a simple, accessible, and transparent model framework for climate and regional sea-level projections," *Geoscientific Model Development*, vol. 10, no. 7, pp. 2741–2760, 2017.

[62] N. Manocha and V. Babovic, "Real options, multi-objective optimization and the development of dynamically robust adaptive pathways," *Environmental science & policy*, vol. 90, pp. 11–18, 2018.

[63] O. Schutze, A. Lara, and C. A. C. Coello, "On the influence of the number of objectives on the hardness of a multiobjective optimization problem," *IEEE Transactions on Evolutionary Computation*, vol. 15, no. 4, pp. 444–455, 2010.

[64] J. Blank and K. Deb, "Pymoo: Multi-Objective Optimization in Python," *IEEE Access*, vol. 8, pp. 89497–89509, 2020, doi: 10.1109/ACCESS.2020.2990567.

[65] D. R. Johnson, J. R. Fischbach, and D. S. Ortiz, "Estimating surge-based flood risk with the coastal Louisiana risk assessment model," *Journal of Coastal Research*, no. 67 (10067), pp. 109–126, 2013.

[66] CPRA, "2023 DRAFT COASTAL MASTER PLAN APPENDIX B: SCENARIO DEVELOPMENT AND FUTURE CONDITIONS," *Version I*, pp. 1–18, 2023, [Online]. Available: https://coastal.la.gov/our-plan/2023-coastal-master-plan/2023-plan-appendices/#B

[67] T. Knutson *et al.*, "Tropical cyclones and climate change assessment: Part II: Projected response to anthropogenic warming," *Bulletin of the American Meteorological Society*, vol. 101, no. 3, pp. E303–E322, 2020.

[68] J.-Y. Lee *et al.*, "Future Global Climate: Scenario-Based Projections and Near-Term Information," in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, United Kingdom and New York NY, USA: Cambridge University Press, 2021, pp. 553–672.

[69] R. W. Hahn and P. M. Dudley, "How well does the US Government do benefit-cost analysis?," *Review of Environmental Economics and Policy*, 2020.

[70] B. Fischhoff, "The realities of risk-cost-benefit analysis," *Science*, vol. 350, no. 6260, 2015.

[71] J. Schofield, *Cost-benefit analysis in urban & regional planning*, 1st ed. London: Routledge, 1987.

[72] E. J. Mishan and E. Quah, *Cost-benefit analysis*, 6th ed. London: Routledge, 2020.

[73] R. Mechler, "Reviewing estimates of the economic efficiency of disaster risk management: opportunities and limitations of using risk-based cost–benefit analysis," *Natural Hazards*, vol. 81, no. 3, pp. 2121–2147, 2016.

[74] J. Van Alphen, "The Delta Programme and updated flood risk management policies in the Netherlands," *Journal of Flood Risk Management*, vol. 9, no. 4, pp. 310–319, 2016.

[75] "ER 1105-2-100: Planning Guidance Notebook." U.S. Army Corps of Engineers, Apr. 22, 2000.

[76] G. W. Bennett, "Memorandum For Building Resilient Infrastructure and Communities (BRIC) and Flood Mitigation Assistance (FMA) Grant Program Applicants and Subapplicants." Federal Emergency Management Agency, Oct. 06, 2022. [Online]. Available: https://www.fema.gov/sites/default/files/documents/fema_alternative-cost-effectiveness-methodology_102022.pdf

[77] A. Ciullo, J. H. Kwakkel, K. M. De Bruijn, N. Doorn, and F. Klijn, “Efficient or fair? Operationalizing ethical principles in flood risk management: A case study on the Dutch-German Rhine,” *Risk analysis*, vol. 40, no. 9, pp. 1844–1862, 2020.

[78] J. Kind, W. W. Botzen, and J. C. Aerts, “Social vulnerability in cost-benefit analysis for flood risk management,” *Environment and Development Economics*, vol. 25, no. 2, pp. 115–134, 2020.

[79] L. Downey and B. Hawkins, “Race, Income, and Environmental Inequality in the United States,” *Sociological Perspectives*, vol. 51, no. 4, pp. 759–781, Dec. 2008, doi: 10.1525/sop.2008.51.4.759.

[80] H. M. Lane, R. Morello-Frosch, J. D. Marshall, and J. S. Apte, “Historical redlining is associated with present-day air pollution disparities in US cities,” *Environmental science & technology letters*, vol. 9, no. 4, pp. 345–350, 2022.

[81] L. P. Clark, D. B. Millet, and J. D. Marshall, “Changes in transportation-related air pollution exposures by race-ethnicity and socioeconomic status: outdoor nitrogen dioxide in the United States in 2000 and 2010,” *Environmental health perspectives*, vol. 125, no. 9, p. 097012, 2017.

[82] D. Switzer and M. P. Teodoro, “Class, Race, Ethnicity, and Justice in Safe Drinking Water Compliance*,” *Social Science Quarterly*, vol. 99, no. 2, pp. 524–535, 2018, doi: 10.1111/ssqu.12397.

[83] U. I. Uche, S. Evans, S. Rundquist, C. Campbell, and O. V. Naidenko, “Community-level analysis of drinking water data highlights the importance of drinking water metrics for the state, federal environmental health justice priorities in the United States,” *International Journal of Environmental Research and Public Health*, vol. 18, no. 19, p. 10401, 2021.

[84] M. Allaire and S. Acquah, “Disparities in drinking water compliance: Implications for incorporating equity into regulatory practices,” *AWWA Water Science*, vol. 4, no. 2, p. e1274, 2022.

[85] J. Chakraborty, T. W. Collins, M. C. Montgomery, and S. E. Grineski, “Social and spatial inequities in exposure to flood risk in Miami, Florida,” *Natural Hazards Review*, vol. 15, no. 3, p. 04014006, 2014.

[86] A. Pallathadka, J. Sauer, H. Chang, and N. B. Grimm, “Urban flood risk and green infrastructure: Who is exposed to risk and who benefits from investment? A case study of three US Cities,” *Landscape and Urban Planning*, vol. 223, p. 104417, 2022.

[87] P. Sharkey, “Survival and death in New Orleans: An empirical look at the human impact of Katrina,” *Journal of Black Studies*, vol. 37, no. 4, pp. 482–501, 2007.

[88] O. M. Nofal and J. W. van de Lindt, “High-resolution approach to quantify the impact of building-level flood risk mitigation and adaptation measures on flood losses at the community-level,” *International Journal of Disaster Risk Reduction*, vol. 51, p. 101903, 2020.

[89] J. R. Fischbach, D. R. Johnson, M. T. Wilson, N. B. Geldner, and C. Stelzner, “2023 Coastal Master Plan: Model Improvement Report, Risk Assessment,” *Version I*, pp. 1–77, 2021.

[90] S. Armal, J. R. Porter, B. Lingle, Z. Chu, M. L. Marston, and O. E. Wing, “Assessing property level economic impacts of climate in the US, new insights and evidence from a comprehensive flood risk assessment tool,” *Climate*, vol. 8, no. 10, p. 116, 2020.

[91] S. Wachter, “Normative challenges of identification in the Internet of Things: Privacy, profiling, discrimination, and the GDPR,” *Computer law & security review*, vol. 34, no. 3, pp. 436–449, 2018.

[92] K. Taparra and K. Pellegrin, “Data aggregation hides Pacific Islander health disparities,” *The Lancet*, vol. 400, no. 10345, pp. 2–3, 2022.

[93] T. J. Kauh, J. G. Read, and A. J. Scheitler, “The critical role of racial/ethnic data disaggregation for health equity,” *Population research and policy review*, vol. 40, no. 1, pp. 1–7, 2021.

[94] C. A. Gallagher, “Color-blind egalitarianism as the new racial norm,” *Theories of race and ethnicity: Contemporary debates and perspectives*, pp. 40–56, 2015.

[95] T. Thaler, “Justice and Resilience in Flood Risk Management: What Are the Socio-Political Implications?,” in *Building resilience to natural hazards in the context of climate change*, Springer, 2021, pp. 41–54.

[96] B. Tolchin, S. C. Hull, and K. Kraschel, “Triage and justice in an unjust pandemic: ethical allocation of scarce medical resources in the setting of racial and socioeconomic disparities,” *Journal of Medical Ethics*, vol. 47, no. 3, pp. 200–202, 2021.

[97] J. L. Harrison, “Neoliberal environmental justice: mainstream ideas of justice in political conflict over agricultural pesticides in the United States,” *Environmental Politics*, vol. 23, no. 4, pp. 650–669, 2014.

[98] N. Daniels, “Reducing Health Disparities,” *Inequalities in health: Concepts, measures, and ethics*, p. 178, 2013.

[99] Council on Environmental Quality, “Environmental justice guidance under the national environmental policy act.” 1997.

[100] “Executive Order No. 14091,” No. 88 Fed. Reg. 10825, 2023.

[101] D. Parfit, “Equality or Priority?,” in *The Ideal of Equality*, 1st ed.Hampshire: Palgrave Macmillan, 2000, pp. 81–125.

[102] J. Walton *et al.*, “Talking culture? Egalitarianism, color-blindness and racism in Australian elementary schools,” *Teaching and Teacher Education*, vol. 39, pp. 112–122, 2014.

[103] M. A. Konovsky, “Understanding procedural justice and its impact on business organizations,” *Journal of management*, vol. 26, no. 3, pp. 489–511, 2000.

[104] K. S. Cook and K. A. Hegtvedt, “Distributive justice, equity, and equality,” *Annual review of sociology*, vol. 9, no. 1, pp. 217–241, 1983.

[105] J. Kind, W. J. Wouter Botzen, and J. C. Aerts, “Accounting for risk aversion, income distribution and social welfare in cost-benefit analysis for flood risk management,” *Wiley Interdisciplinary Reviews: Climate Change*, vol. 8, no. 2, p. e446, 2017.

[106] J. Finn, “Louisiana Republican Party wants to ban college study of diversity, equity, inclusion,” *nola.com*, New Orleans, LA, Apr. 15, 2023.

[107] P. Jankowski, “Senate passes critical race theory ban at Texas’ public universities,” *The Dallas Morning News*, Dallas, TX, Apr. 11, 2023.

[108] S. Blaskey, “Few complianed of ‘woke’ classes at Florida universities. Still, DeSantis pushed new bans,” *Miami Herald*, Miami, FL, Apr. 26, 2032.

[109] J. Gyourko, P. Linneman, and S. Wachter, “Analyzing the relationships among race, wealth, and home ownership in America,” *Journal of Housing Economics*, vol. 8, no. 2, pp. 63–89, 1999.

[110] S. Laska and B. H. Morrow, “Social vulnerabilities and Hurricane Katrina: an unnatural disaster in New Orleans,” *Marine technology society journal*, vol. 40, no. 4, 2006.

[111] R. M. Zoraster, “Vulnerable populations: Hurricane Katrina as a case study,” *Prehospital and disaster medicine*, vol. 25, no. 1, pp. 74–78, 2010.

[112] E. Fussell, “The long-term recovery of New Orleans’ population after Hurricane Katrina,” *American Behavioral Scientist*, vol. 59, no. 10, pp. 1231–1245, 2015.

[113] T. Craemer, “Evaluating racial disparities in Hurricane Katrina relief using direct trailer counts in New Orleans and FEMA records,” *Public Administration Review*, vol. 70, no. 3, pp. 367–377, 2010.

[114] J. D. Gourevitch *et al.*, “Spatial targeting of floodplain restoration to equitably mitigate flood risk,” *Global Environmental Change*, vol. 61, p. 102050, 2020.

[115] J. Villarreal-Rosas, A. L. Vogl, L. J. Sonter, H. P. Possingham, and J. R. Rhodes, “Trade-offs between efficiency, equality and equity in restoration for flood protection,” *Environmental Research Letters*, vol. 17, no. 1, p. 014001, 2021.

[116] O. Seyedashraf, A. Bottacin-Busolin, and J. J. Harou, “A design framework for considering spatial equity in sustainable urban drainage infrastructure,” *Sustainable Cities and Society*, vol. 85, p. 103960, 2022.

[117] G. Yang, M. Giuliani, and A. Castelletti, “Operationalizing equity in multipurpose water systems,” *Hydrology and Earth System Sciences*, vol. 27, no. 1, pp. 69–81, 2023.

[118] M. Kazemi, O. Bozorg-Haddad, E. Fallah-Mehdipour, and H. A. Loáiciga, “Inter-basin hydropolitics for optimal water resources allocation,” *Environmental Monitoring and Assessment*, vol. 192, pp. 1–10, 2020.

[119] A. Gullotta, D. Butler, A. Campisano, E. Creaco, R. Farmani, and C. Modica, “Optimal location of valves to improve equity in intermittent water distribution systems,” *Journal of Water Resources Planning and Management*, vol. 147, no. 5, p. 04021016, 2021.

[120] M. Amini and A. M. Memari, “Comparative review and assessment of various flood retrofit methods for low-rise residential buildings in coastal areas,” *Natural Hazards Review*, vol. 22, no. 3, p. 04021009, 2021.

[121] S. Xian, N. Lin, and H. Kunreuther, “Optimal house elevation for reducing flood-related losses,” *Journal of Hydrology*, vol. 548, pp. 63–74, 2017.

[122] J. R. Fischbach, *Managing New Orleans flood risk in an uncertain future using non-structural risk mitigation*. The Pardee RAND Graduate School, 2010.

[123] M. Zarekarizi, V. Srikrishnan, and K. Keller, “Neglecting uncertainties biases house-elevation decisions to manage riverine flood risks,” *Nature communications*, vol. 11, no. 1, pp. 1–11, 2020.

[124] “Hazus Hurricane Model Technical Manual: Hazus 5.1.” FEMA, Jul. 2022.

[125] S. Dekel and S. Fischer, “Potential pareto public goods,” *Journal of Public Economics*, vol. 146, pp. 87–96, 2017.

[126] M. T. Wilson *et al.*, “Draft Coastal Master Plan: Attachment E3: Nonstructural Protection Evaluation Results,” *Version I*, p. 36, 2023.

[127] U.S. Census Bureau, “2019 American Community Survey 5-Year Data Tables.” 2019. [Online]. Available: https://data.census.gov/cedsci/

[128] City of New Orleans, “Neighborhood Statistical Areas.” https://portal-nolagis.opendata.arcgis.com/maps/e7daa4c977d14e1b9e2fa4d7aff81e59/about. [Online]. Available: https://portal-nolagis.opendata.arcgis.com/maps/e7daa4c977d14e1b9e2fa4d7aff81e59/about

[129] P. Hudson and W. W. Botzen, “Cost–benefit analysis of flood-zoning policies: A review of current practice,” *Wiley Interdisciplinary Reviews: Water*, vol. 6, no. 6, p. e1387, 2019.

[130] M. L. Oliver and T. M. Shapiro, “Race and wealth,” *Rev Black Polit Econ*, vol. 17, no. 4, pp. 5–25, Mar. 1989, doi: 10.1007/BF02910814.

[131] L. A. Keister, “Race and Wealth Inequality: The Impact of Racial Differences in Asset Ownership on the Distribution of Household Wealth,” *Social Science Research*, vol. 29, no. 4, pp. 477–502, Dec. 2000, doi: 10.1006/ssre.2000.0677.

[132] H. D. Horton, “Race and Wealth: A Demographic Analysis of Black Homeownership*,” *Sociological Inquiry*, vol. 62, no. 4, pp. 480–489, 1992, doi: 10.1111/j.1475-682X.1992.tb00296.x.

[133] S. A. Bond Huie, P. M. Krueger, R. G. Rogers, and R. A. Hummer, “Wealth, Race, and Mortality*,” *Social Science Quarterly*, vol. 84, no. 3, pp. 667–684, 2003, doi: 10.1111/1540-6237.8403011.

[134] L. M. Kohm, K. Sumner, and P. Farley, “Empowering Black Wealth in the Shadow of the Tulsa Race Massacre,” *Tulsa L. Rev.*, vol. 57, p. 243, 2022 2021, [Online]. Available: https://heinonline.org/HOL/Page?handle=hein.journals/tlj57&id=249&div=&collection=

[135] F. R. Addo and W. A. Darity, “Disparate Recoveries: Wealth, Race, and the Working Class after the Great Recession,” *The ANNALS of the American Academy of Political and Social Science*, vol. 695, no. 1, pp. 173–192, May 2021, doi: 10.1177/00027162211028822.

[136] D. T. Resio, “White paper on estimating hurricane inundation probabilities,” in *Performance Evaluation of the New Orleans and Southeast Louisiana Hurricane Protection System. Appendix 8: Hazard Analysis*, U.S. Army Corps of Engineers, 2007.

[137] D. T. Resio, J. Irish, and M. Cialone, “A surge response function approach to coastal hazard assessment–part 1: basic concepts,” *Natural Hazards*, vol. 51, no. 1, pp. 163–182, 2009.

[138] J. R. Fischbach *et al.*, “Coastal Master Plan Modeling: Attachment C3-25 – Storm Surge and Risk Assessment,” Louisiana Coastal Protection and Restoration Authority, Baton Rouge, Louisiana, 2017.

[139] C. W. Landsea and J. L. Franklin, “Atlantic hurricane database uncertainty and presentation of a new database format,” *Monthly Weather Review*, vol. 141, no. 10, pp. 3576–3592, 2013.

[140] M. R. Meyer and D. R. Johnson, “Variability of Best-Estimate Flood Depth Return Periods in Coastal Louisiana,” *Journal of Marine Science and Engineering*, vol. 7, no. 5, p. 145, 2019.

[141] U.S. Army Corps of Engineers, “Louisiana Coastal Protection and Restoration Final Technical Report: Hydraulics and Hydrology Appendix,” U.S. Army Corps of Engineers, New Orleans, Louisiana, 2009.

[142] B. McMann, M. Schulze, H. Sprague, and K. Smyth, “Coastal master plan: Appendix A: Project definitions,” *Version I. Coastal Protection and Restoration Authority, Baton Rouge*, 2017.