Data-driven control reveals distributed flood adaptation priorities across large river networks

Jeil Oh[1] **and Matthew Bartos**[1]

[1]Fariborz Maseeh Department of Civil, Architectural and Environmental Engineering, The University of Texas at Austin, Austin, TX, USA

Key Points:

- New method for basin-scale flood adaptation planning uses optimal control to determine best reach-scale attenuation targets
- Network-level analysis identifies spatial priorities that local reach attributes alone cannot capture
- Effort, residual-risk, and ensemble-agreement maps together inform where effort concentrates, risk accumulates, and uncertainty remains

Corresponding author: Jeil Oh, `jeoh@utexas.edu`

Abstract

In the face of growing flood risks, decentralized adaptation measures like detention storage, enhanced infiltration, and floodplain reconnection have the potential to mitigate flood impacts at the river basin scale. However, optimal spatial allocation of flood control measures is complicated by the high dimensionality of hydrologic systems and the sensitivity of proposed strategies to climate uncertainty. To overcome these challenges, we propose a diagnostic framework that combines reduced-order data-driven modeling with optimal control to directly estimate reach-level attenuation targets without the need for iterative simulation and optimization. First, Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition with control (DMDc) are used to create a low-rank linear surrogate model of basin rainfall-runoff dynamics. A Linear Quadratic Regulator (LQR) is then applied to compute optimal reach-scale attenuation targets that mitigate flood impacts. Applied to a large river basin under a multi-model climate ensemble, the framework successfully determines distributed attenuation strategies that reduce bankfull discharge exceedances under varying adaptation budgets. Across scenarios, marginal increases to flow attenuation are found to yield diminishing returns to flood mitigation, while higher-emission scenarios retain substantially greater residual flood volume for the same effort level. Although attenuation allocation generally scales with mean flow, we identify tributary and transitional reaches where attenuation demand is disproportionate to local hydrologic size, showing that network-level flood dynamics produce spatial priorities that cannot be recovered from reach attributes alone. Taken together, the proposed framework provides a scalable approach for flood adaptation planning that is effective for basin- to continental-scale applications under climate uncertainty.

Plain Language Summary

As climate change intensifies flooding, communities face a core planning problem: where to deploy flood defenses across river networks, and how to size them. Because rivers are connected systems, reducing floodwaters in one place changes flood peaks downstream. Determining the best placement of flood interventions across thousands of river branches and many climate futures is thus computationally demanding. To overcome this challenge, we develop a flood control planning framework that compresses the complex behavior of a river network into a compact model that can be solved directly, bypassing trial-and-error simulations and producing requirement maps detailing the level of flood reduction needed at every stream segment across the basin. When applied to a large river basin, the framework identifies critical flood bottlenecks caused by interactions of tributary flows across the network that simple size-based rankings miss. The framework also shows how flood reduction benefits diminish as more attenuation is added, and identifies priority locations that remain robust across different climate projections. Together, these outputs diagnose basin-scale adaptation needs, helping planners decide where to act, how much capacity is realistic, and where to remain flexible as conditions evolve.

1 Introduction

Climate change is reshaping flood regimes worldwide, altering the magnitude, timing, and seasonality of extreme flows in regionally distinct ways (Blöschl et al., 2019; Tabari, 2020; Thober et al., 2018). In response, flood adaptation strategies are expanding beyond centralized defenses toward distributed, landscape-scale interventions such as detention basins, floodplain reconnection, wetland restoration, and enhanced infiltration (Dadson et al., 2017; Opperman et al., 2009). These measures attenuate flood peaks across the drainage network rather than concentrating protection at a few major dams. Catchment-scale evidence shows that such measures can substantially reduce flood peaks when deployed strategically (Black et al., 2021; Nicholson et al., 2020), and that combined port-

folios of nature-based measures outperform individual interventions in isolation (Cheng et al., 2017).

However, the effectiveness of distributed flood control depends on the coordination and spatial placement of interventions at the watershed scale. Flooding at any given point in the river network reflects accumulated and time-shifted flow contributions from many upstream tributaries; similarly, attenuation at any given location alters peak flows farther downstream. Thus, tributary timing, upstream–downstream coupling, and climate-driven shifts in flow generation jointly determine the locations where attenuation has the greatest leverage and the level of flow reduction that should be achieved in each part of the network (Ayalew et al., 2015; Thomas et al., 2016). This coupling raises a planning question that precedes detailed design and siting decisions: *where* in a river network should flood adaptation efforts concentrate, *how much* attenuation is needed at each location, and how do these spatial patterns *shift* under varying climate futures? Without spatially resolved guidelines, adaptation planning cannot distinguish reaches that require substantial intervention from reaches where additional investment yields diminishing returns.

Optimal planning of flood control measures at the scale of major river basins confronts three interlocking challenges. First, *hydrologic complexity*: the flood response at any river reach depends on a complex array of nonlinear runoff generation and routing processes that integrate precipitation forcing, antecedent soil conditions, and the timing and magnitude of inflows from all upstream tributaries. Second, *dimensionality*: river networks routinely comprise tens of thousands of interconnected reaches (David et al., 2011; Mizukami et al., 2016), each with distinct hydrologic characteristics, making exhaustive reach-by-reach optimization computationally intractable even for a single climate scenario. Third, *scenario dependence*: climate change rearranges precipitation patterns and flood-generating mechanisms in projection-dependent ways, and robust adaptation planning requires characterization of reach-level attenuation requirements not under a single design condition but across ensembles of general circulation models (GCMs) and emission pathways, multiplying the computational burden by orders of magnitude. Together, these challenges have severely limited the development of large watershed-scale flood adaptation plans.

On the sensitivity of flood adaptation planning to climate uncertainty, recent work has made substantial progress toward characterizing the effects of climate change on the magnitude and distribution of flood hazards. Global and regional projections have coupled GCMs with hydrologic and hydrodynamic models to quantify changes in flood magnitude and frequency (Hirabayashi et al., 2013; Alfieri et al., 2017), while high-resolution flood mapping efforts have identified locations where inundation risk concentrates at the landscape scale (Bates et al., 2010; Bentivoglio et al., 2022). Scenario-neutral and bottom-up frameworks complement these projections by stress-testing system performance across broad ranges of climatic perturbations and identifying thresholds at which current infrastructure becomes inadequate (Broderick et al., 2019; Prudhomme et al., 2010). Collectively, these tools equip planners with increasingly detailed pictures of *what may happen and where*. They do not, however, provide the complementary information that adaptation planning requires: the spatial distribution of attenuation measures needed to reduce flood exceedance across river networks.

On the response side, two bodies of work address the problem of optimal allocation of flood control measures, focusing on either the placement or operation of predetermined control assets like reservoirs or detention basins. The first body of work estimates cost-effective adaptation strategies at regional to continental scales. Cost–benefit analyses have compared portfolios of structural and nature-based measures under multiple warming levels (Dottori et al., 2023; Alfieri et al., 2016), and global frameworks have quantified economically optimal protection standards for urban river reaches (Ward et al., 2017). These studies show that distributed strategies such as detention areas can be highly cost-effective. However, the principal focus of these studies is to quantify the cost–

benefit trade-off of predefined portfolios of mitigation measures—typically at the level of administrative regions or coarse grid cells—and not to understand general spatial patterns of how interventions should be organized at the watershed scale. The second body of work applies control and optimization theory to management of water infrastructure, primarily for operations. Stochastic and multi-objective optimization have been used to derive reservoir operating policies (Giuliani et al., 2016; Bertoni et al., 2020; Giuliani et al., 2021); Model Predictive Control and reinforcement learning have been used for adaptive real-time management (Bowes et al., 2022; Ning et al., 2022; Castelletti et al., 2023; Oh & Bartos, 2023); linear quadratic feedback has been applied to real-time stormwater control (Wong & Kerkez, 2018; Lin et al., 2024); and data-driven surrogates have been shown to accelerate the simulation–optimization loop that these methods require (Razavi et al., 2012; Castelletti et al., 2012). Simulation–optimization methods have also been used to select cost-effective detention-basin layouts, but computational tractability typically limits these applications to $O(10^1)$–$O(10^2)$ subcatchments (Travis & Mays, 2008; Perez-Pedini et al., 2005). As a result, these approaches do not directly resolve the reach-level attenuation requirements of large river networks with $O(10^4)$ reaches.

These limitations in spatial extent and resolution not only reflect a computational barrier to overcome, but also reflect a more fundamental knowledge gap regarding the coordination of flood defenses in large, interconnected river systems: namely, how should decentralized flood attenuation measures generally be distributed within a river basin, and what flood mitigation benefits can be expected under this optimized placement? While approaches based on process-based models of flood defenses provide practical guidelines for siting a limited number of interventions, the general question of how interventions should be spatially organized at the watershed scale is obscured by the computational challenge of repeated simulation–optimization over thousands of potential control sites. Thus, rather than starting with a fixed portfolio of interventions, an alternative approach is to instead first determine the level of attenuation at each reach that minimizes flood impacts, and then work backwards to find the specific interventions that satisfy this target attenuation. In addition to revealing general patterns for the spatial organization of interventions within drainage networks, this approach ensures a system-scale adaptation strategy that minimizes flood impacts over the entire drainage network, accounting for the impacts of upstream interventions on peak flows farther downstream.

Using a similar line of reasoning, recent work has highlighted connections between control theory and climate uncertainty in adaptation planning (Herman et al., 2020). We build on this perspective by treating flood adaptation planning as an optimal control problem. However, rather than prescribing operations or investment timing, we use control theory *diagnostically*. In this context, our proposed control strategy specifies the reach-scale reductions in flow required to minimize exceedance of flood thresholds throughout the drainage network. We interpret the magnitude of this attenuation as a scenario-conditioned indicator of adaptation need, rather than as an operating rule for a specific structure. The resulting attenuation target informs detention storage, infiltration capacity, or floodplain reconnection goals depending on local conditions. The challenge is making this diagnostic tractable for networks of $O(10^4)$ reaches under an ensemble of climate futures.

We meet this challenge by proposing a new method for basin-scale planning of flood defenses that combines data-driven dimensionality reduction with closed-form optimal control. First, Proper Orthogonal Decomposition (POD) extracts the dominant spatial modes of streamflow variability across the river network, compressing $O(10^4)$ reaches into a low-rank basis that captures the main patterns of flood dynamics. Dynamic Mode Decomposition with control (DMDc) (Proctor et al., 2016; Kutz et al., 2016) then identifies a linear surrogate in this reduced space, describing how reach-level discharge evolves under precipitation forcing. The surrogate is trained directly on process-based hydrologic simulations rather than on simplified analytical representations of basin hydrology. Finally, because the surrogate operates in reduced coordinates, a Linear Quadratic Reg-

ulator (LQR) (Anderson & Moore, 2007) solves in closed form for the reach- and time-specific attenuation for any prescribed balance between flood reduction and total effort.

Applied across a range of effort-penalty weights, the POD–DMDc–LQR formulation produces three diagnostic outputs for basin-scale flood adaptation. First, effort–residual trade-off curves link total attenuation effort to residual flood risk and identify points of diminishing returns where additional adaptation becomes less effective. Second, reach-level maps show where attenuation effort is allocated spatially and where residual risk remains after optimized attenuation. Third, scenario-ensemble diagnostics distinguish spatial patterns of proposed adaptation measures that are robust across climate futures versus those that remain sensitive to climate forcing. Together, these outputs identify priority locations for flood adaptation and provide reach-level attenuation targets that inform subsequent design and feasibility analysis.

We demonstrate the framework on the Brazos River basin (Texas, USA) and address three questions:

1. How does the basin-wide trade-off between attenuation effort and residual flood risk vary across climate scenarios?
2. How does the framework distribute adaptation effort across the drainage network, where does residual risk concentrate, and how do these spatial patterns relate to network structure and climate forcing?
3. Which attenuation and residual-risk priorities are robust across the climate ensemble, and which remain sensitive to individual climate scenarios?

2 Methodology

We formulate distributed flood adaptation as a control problem in which projected streamflow and precipitation are used to determine the reach-level attenuation targets needed to avoid flood impacts. The flood control planning framework proceeds in three stages (Figure 1): (i) a reduced-order system identification scheme based on POD and DMDc is applied to yield a compact linear surrogate model of basin rainfall-runoff dynamics; (ii) an optimal control scheme based on LQR is derived to compute optimal reach-level flow attenuation targets for a range of effort-penalty weights; and (iii) using this LQR control scheme, a threshold-triggered closed-loop simulation is executed under a range of future climate scenarios to determine flood attenuation targets needed to mitigate exceedance of bankfull flows.

2.1 Study Area and Data

We demonstrate the framework on the Brazos River basin in Texas, USA (approximately 118,000 km^2), a large rainfall-dominated basin that drains from the semiarid Southern High Plains to the Gulf of Mexico. The basin contains 14,421 NHDPlusV2 flowlines and exhibits spatially heterogeneous flood generation driven by diverse storm types, making it a challenging test case for reach-level adaptation planning.

Streamflow simulations are drawn from the CMIP6-based daily streamflow projection dataset developed at Oak Ridge National Laboratory (Kao et al., 2022; Ghimire et al., 2023). The dataset provides routed daily discharge at NHDPlusV2 reaches across the conterminous United States using the RAPID routing model (David et al., 2011), without representing reservoir operations. We use two hydrologic models (VIC and PRMS), seven GCMs, and four SSPs, yielding 56 ensemble members spanning 1980–2099 (Kao et al., 2024).

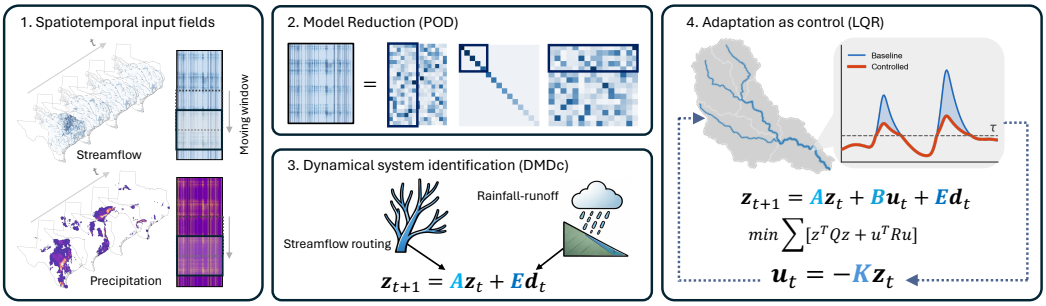

Figure 1. Overview of the diagnostic framework. (1) Spatiotemporal streamflow and precipitation fields from process-based hydrologic simulations are assembled in moving windows. (2) Proper Orthogonal Decomposition (POD) compresses the high-dimensional discharge field into a low-rank basis. (3) Dynamic Mode Decomposition with control (DMDc) identifies a linear surrogate that captures streamflow routing and rainfall–runoff dynamics in reduced coordinates. (4) A Linear Quadratic Regulator (LQR) solves for the spatially distributed attenuation that balances flood reduction against total effort, producing reach-level adaptation requirements across the river network.

2.2 Problem Formulation

Let $\mathbf{x}(t) \in \mathbb{R}^n$ denote daily streamflow across n river reaches at time t, $\mathbf{u}(t) \in \mathbb{R}^n$ represent a spatially distributed control signal associated with flood attenuation, and $\mathbf{d}(t) \in \mathbb{R}^n$ represent the precipitation input forcing mapped to each reach. The uncontrolled basin dynamics can be written generally as

$$\mathbf{x}(t+1) = \mathcal{F}(\mathbf{x}(t), \mathbf{d}(t)), \tag{1}$$

where $\mathcal{F}$ is a function representing relevant hydrologic processes, including rainfall–runoff generation at each subcatchment and flow routing across the river network. The controlled system augments Eq. (1) with a control input $\mathbf{Bu}(t)$

$$\mathbf{x}(t+1) = \mathcal{F}(\mathbf{x}(t), \mathbf{d}(t)) + \mathbf{Bu}(t), \tag{2}$$

where $\mathbf{B} = \mathbf{I}_n$, reflecting the assumption that every reach is in principle controllable. The control input $\mathbf{Bu}(t)$ represents the reach-scale flow adjustment associated with distributed flood attenuation. The more realistic case of intervention at selected sub-basin units is addressed in Section 4.2.

2.3 Reduced-Order System Identification

2.3.1 State-Space Reduction via POD

To make the control problem tractable, we project the high-dimensional streamflow state onto a low-dimensional subspace using Proper Orthogonal Decomposition (POD) (Berkooz et al., 1993). Let $\mathbf{X} = [\mathbf{x}(1), \ldots, \mathbf{x}(T)] \in \mathbb{R}^{n \times T}$ be the discharge snapshot matrix and $\bar{\mathbf{x}}$ its temporal mean. We retain the leading r left singular vectors of the mean-centered snapshots to form $\mathbf{U}_r \in \mathbb{R}^{n \times r}$, so that the reduced state is

$$\mathbf{z}(t) = \mathbf{U}_r^\top (\mathbf{x}(t) - \bar{\mathbf{x}}) \in \mathbb{R}^r. \tag{3}$$

For this case study, we fix the dimension of the reduced system as $r = 30$, which captures the dominant spatial modes of streamflow variability within the Brazos River basin while ensuring subsequent computations remain tractable. For each ensemble member

(GCM × hydrologic model × SSP), $\mathbf{U}_r$ and $\bar{\mathbf{x}}$ are computed once from its full 2020–2099 discharge matrix and held fixed across all subsequent DMDc windows and LQR designs within that member, ensuring the same coordinate system for system identification, control, and reconstruction.

2.3.2 DMDc System Identification

Within the POD basis, we identify the basin response to precipitation forcing using Dynamic Mode Decomposition with control (DMDc) (Proctor et al., 2016; Kutz et al., 2016). DMDc provides a reduced-order linear dynamical model that maps the current reduced state and precipitation forcing to the next reduced state, with separate terms for routing dynamics and rainfall-driven inputs. The precipitation forcing is first centered and normalized: $\tilde{\mathbf{d}}(t) = (\mathbf{d}(t) - \boldsymbol{\mu}_d)/\sigma_d$, where $\boldsymbol{\mu}_d$ is the reach-wise temporal mean and σ_d is the global standard deviation across all reaches and time steps. A single global scale preserves relative spatial differences in precipitation magnitude across the basin. In reduced coordinates, the uncontrolled surrogate takes the form

$$\mathbf{z}(t+1) = \mathbf{A}\mathbf{z}(t) + \mathbf{E}\tilde{\mathbf{d}}(t), \tag{4}$$

where $\mathbf{A} \in \mathbb{R}^{r\times r}$ is the reduced transition matrix encoding flow-routing dynamics across the network, and $\mathbf{E} \in \mathbb{R}^{r\times n}$ maps precipitation forcing to reduced-state responses, implicitly capturing rainfall–runoff generation. We estimate $\mathbf{G} = [\mathbf{A}, \mathbf{E}]$ jointly via Tikhonov-regularized regression:

$$\mathbf{G}^* = \arg\min_{\mathbf{G}} \left\{ \|\mathbf{Z}' - \mathbf{G}\boldsymbol{\Omega}\|_F^2 + \alpha \|\mathbf{G}\|_F^2 \right\}, \tag{5}$$

where $\boldsymbol{\Omega} = [\mathbf{Z}^\top, \tilde{\mathbf{D}}^\top]^\top$ concatenates state and disturbance snapshots (Proctor et al., 2016). The regularization parameter α is selected via grid search as the smallest value yielding spectral radius $\rho(\mathbf{A}) < 0.995$, promoting stable long-horizon simulations.

To accommodate nonstationarity in climate-driven streamflow dynamics, we estimate $(\mathbf{A}^{(w)}, \mathbf{E}^{(w)})$ in overlapping temporal windows of 10 years, advanced by 5 years. Each time step uses the dynamics from its containing window:

$$\mathbf{z}(t+1) = \mathbf{A}^{(w(t))}\mathbf{z}(t) + \mathbf{E}^{(w(t))}\tilde{\mathbf{d}}(t). \tag{6}$$

Because the POD basis is fixed, only the system matrices $\mathbf{A}^{(w)}$ and $\mathbf{E}^{(w)}$ vary across windows, capturing evolving climate-driven dynamics within a stationary coordinate frame.

2.3.3 Adding Control to the Identified System

Having identified a reduced-order surrogate model describing rainfall-runoff response within the river network, we augment this surrogate model with a control term representing the adaptation effort from Eq. (2), so that attenuation targets can be computed in reduced coordinates. In physical space, flow can be attenuated at each individual reach (i.e., $\mathbf{B} = \mathbf{I}_n$). However, to operate in reduced coordinates, we constrain the attenuation to lie in the column space of the POD basis, $\mathbf{u}(t) = \mathbf{U}_r\mathbf{u}_z(t)$. Projecting the control term into reduced space gives

$$\mathbf{U}_r^\top \mathbf{B}\, \mathbf{U}_r\mathbf{u}_z(t) = \mathbf{U}_r^\top \mathbf{I}_n\, \mathbf{U}_r\mathbf{u}_z(t) = \mathbf{I}_r\, \mathbf{u}_z(t), \tag{7}$$

so the reduced input matrix is $\mathbf{B}_r = \mathbf{I}_r$ and the controlled reduced system becomes

$$\mathbf{z}(t+1) = \mathbf{A}^{(w)}\mathbf{z}(t) + \mathbf{B}_r\mathbf{u}_z(t) + \mathbf{E}^{(w)}\tilde{\mathbf{d}}(t). \tag{8}$$

Equation (8) can be read as a reduced-order water budget: the homogeneous term $\mathbf{A}^{(w)}\mathbf{z}$ propagates internal basin dynamics (routing and storage release), while the disturbance

term $\mathbf{E}^{(w)}\tilde{\mathbf{d}}$ introduces the precipitation-driven inflow, and the control term $\mathbf{B}_r\mathbf{u}_z$ represents the flood adaptation effort, i.e., the spatially distributed attenuation that interventions must provide. Separating the disturbance component (learned from data) from the control component (imposed by design) allows the framework to quantify the degree of adaptation required at each reach under a given climate scenario without prescribing specific interventions.

2.4 Optimal Control Design

To determine the optimal allocation of flood control interventions at the river basin scale, we apply a Linear Quadratic Regulator (LQR) approach (Anderson & Moore, 2007). The LQR computes a feedback policy that penalizes high-flow states while limiting total attenuation effort. For a given window with dynamics $\mathbf{A}^{(w)}$ and $\mathbf{B}_r = \mathbf{I}_r$, we minimize the quadratic cost function

$$J = \sum_{t=0}^{T-1} \left[\mathbf{z}(t)^\top \mathbf{Q}\, \mathbf{z}(t) + \mathbf{u}_z(t)^\top \mathbf{R}\, \mathbf{u}_z(t)\right], \tag{9}$$

where $\mathbf{Q} = \mathbf{I}_r$ penalizes reduced-state magnitude and $\mathbf{R} = R_w\mathbf{I}_r$ penalizes control effort. The scalar weight R_w governs the trade-off between flood reduction and attenuation effort. Small R_w allows stronger attenuation when control is active, whereas large R_w penalizes attenuation more heavily and leaves greater residual exceedance. The optimal policy is given by the feedback law $\mathbf{u}_z(t) = -\mathbf{K}\mathbf{z}(t)$ with

$$\mathbf{K} = \left(\mathbf{R} + \mathbf{B}_r^\top \mathbf{P}\mathbf{B}_r\right)^{-1} \mathbf{B}_r^\top \mathbf{P}\mathbf{A}, \tag{10}$$

where $\mathbf{P}$ solves the discrete algebraic Riccati equation (DARE),

$$\mathbf{P} = \mathbf{Q} + \mathbf{A}^\top \mathbf{P}\mathbf{A} - \mathbf{A}^\top \mathbf{P}\mathbf{B}_r \left(\mathbf{R} + \mathbf{B}_r^\top \mathbf{P}\mathbf{B}_r\right)^{-1} \mathbf{B}_r^\top \mathbf{P}\mathbf{A}. \tag{11}$$

Since $\mathbf{B}_r = \mathbf{I}_r$, the optimal control gain reduces to $\mathbf{K} = (\mathbf{R} + \mathbf{P})^{-1}\mathbf{P}\mathbf{A}$, with $\mathbf{P}$ given by the standard DARE. Gains $\mathbf{K}^{(w)}$ are computed separately for each window using the corresponding $\mathbf{A}^{(w)}$.

Given that the optimal control policy depends on the weighting R_w that determines the emphasis applied to flood reduction versus attenuation effort, we sweep R_w over a logarithmic grid spanning 10^{-4} to 10^5 to trace the trade-off between these two objectives. For cross-scenario comparison at a fixed effort budget, we linearly interpolate reach-level metrics between the two bracketing R_w values whose basin-total efforts straddle the target. This strategy ensures that all scenarios and ensemble members are evaluated at identical total effort, isolating the effect of climate forcing from differences in effort allocation.

2.5 Flood Threshold Definition

Because the controller is intended to attenuate flood exceedance rather than modify ordinary within-bank flows, we require a reach-specific flood threshold $\boldsymbol{\tau} \in \mathbb{R}^n$ that determines when flood control activates. We derive thresholds from pooled annual maximum discharge series (AMS) over the period from 1980 to 2019. For each reach, we extract the maximum daily discharge in each year and pool these annual maxima across all four SSP-specific historical segments. Because the bias correction applies SSP-dependent adjustments, these segments are not identical even during the shared historical period; pooling therefore incorporates downscaling uncertainty into the threshold estimate. The return level for return period T is the empirical $(1-1/T)$ quantile of the pooled AMS:

$$\tau_i^{(T)} = \hat{F}_{\mathcal{A}_i}^{-1}\left(1 - \frac{1}{T}\right). \tag{12}$$

We use the T=2-year return level as the default threshold, approximating bankfull discharge (Wolman & Miller, 1960; Wilkerson, 2008) and targeting the onset of overbank flooding rather than rare extremes.

2.6 Threshold-Triggered Simulation

The LQR gain $\mathbf{K}^{(w)}$ is derived from the standard (ungated) linear system but applied in a threshold-modulated manner so that feedback acts only on flood exceedance. At each time step, we reconstruct discharge as $\hat{\mathbf{x}}(t) = \max(\mathbf{U}_r\mathbf{z}(t) + \bar{\mathbf{x}}, \mathbf{0})$ and define the exceedance vector

$$\mathbf{x}^+(t) = \max(\hat{\mathbf{x}}(t) - \boldsymbol{\tau},\, \mathbf{0})\,. \tag{13}$$

We project exceedance back to reduced space, $\mathbf{z}^+(t) = \mathbf{U}_r^\top \mathbf{x}^+(t)$, and compute flood-only feedback

$$\mathbf{u}_z(t) = -\mathbf{K}^{(w(t))}\mathbf{z}^+(t). \tag{14}$$

When all reaches are below the threshold, $\mathbf{z}^+(t) = \mathbf{0}$ and $\mathbf{u}_z(t) = \mathbf{0}$. Conversely, during floods, effort scales with exceedance magnitude. This gating mechanism ensures that control effort is zero during normal flows and activates proportionally only when discharge exceeds the flood threshold, producing attenuation signals that correspond to the effort required to bring flows back toward the flood threshold $\boldsymbol{\tau}$.

The closed-loop surrogate evolves according to Eq. (8) with $\mathbf{u}_z(t)$ given by Eq. (14), initialized at $\mathbf{z}(0) = \mathbf{U}_r^\top(\mathbf{x}(0) - \bar{\mathbf{x}})$. The implied physical-space attenuation is $\mathbf{u}(t) = \mathbf{U}_r\mathbf{u}_z(t)$. Because this attenuation is solved in reduced coordinates and lifted through the POD basis rather than applied independently at each reach, the recovered $u_i(t)$ is not sign-constrained; we therefore quantify effort by its magnitude $|u_i(t)|$ (Table 1); we verify the attenuation interpretation of this signal in Section 4.2.

2.7 Performance Metrics

We evaluate flood attenuation performance for the future period 2020–2099 using metrics computed from surrogate-reconstructed discharge trajectories. Baseline metrics are computed from open-loop surrogate simulations (no control), and controlled metrics are computed from closed-loop simulations under Eq. (8). This setup ensures consistent comparisons within the surrogate framework. Table 1 defines the core performance metrics. Volumetric metrics include excess volume, excess-volume reduction, total effort, and efficiency; together, these metrics quantify mean annual flood exposure above the threshold, the fraction of that exposure removed by control, the annual attenuation volume required at each reach, and the reduction achieved per unit effort. Peak discharge reduction measures the relative change in maximum daily flow. Return-level residuals compare controlled and uncontrolled annual-maximum quantiles for $T \in \{2, 5, 10, 20, 50, 100\}$ years. For basin-scale summaries, volumetric quantities are summed across reaches. Because the same water can pass through multiple downstream reaches, these sums represent network-integrated exceedance and attenuation burdens rather than unique basin water volumes at the outlet. Non-additive metrics are summarized using medians and interquartile ranges.

2.8 Surrogate Validation

We validate the reduced-order surrogate by iterating it forward from initial conditions over 2020–2099 without re-initialization and comparing simulated streamflows with those from the original process-based model. Performance is measured at each reach using Nash–Sutcliffe Efficiency (NSE), Kling–Gupta Efficiency (KGE), squared Pearson correlation (R^2), and absolute percent bias (|PBIAS|). To evaluate the benefit of time-varying system identification, we compare the windowed surrogate (10-year windows, 5-year step) against a single full-period fit.

Table 1. Core performance metrics. Baseline (bl) and controlled (cl) refer to open-loop and closed-loop surrogate simulations. Δt denotes the daily time step, N_y is the number of evaluation years, $Q_i^{\max}$ is the maximum daily discharge at reach i, and ϵ is a small constant that avoids division by zero.

Metric	Definition
Excess volume	$V_i = \frac{1}{N_y} \sum_t \max(x_i(t) - \tau_i, 0)\, \Delta t$
Excess-volume reduction	$\phi_i = \frac{V_i^{(\mathrm{bl})} - V_i^{(\mathrm{cl})}}{V_i^{(\mathrm{bl})} + \epsilon}$
Total effort	$E_i = \frac{1}{N_y} \sum_t \lvert u_i(t) \rvert\, \Delta t, \quad u_i(t) = [\mathbf{U}_r \mathbf{u}_z(t)]_i$
Efficiency	$\eta_i = \frac{V_i^{(\mathrm{bl})} - V_i^{(\mathrm{cl})}}{E_i + \epsilon}$
Peak discharge reduction	$\delta_i = \frac{Q_i^{\max,(\mathrm{bl})} - Q_i^{\max,(\mathrm{cl})}}{Q_i^{\max,(\mathrm{bl})} + \epsilon}$, where $Q_i^{\max} = \max_t\, x_i(t)$
Return-level residual	$\mathrm{RL}_{T,i}^{(\mathrm{cl})} / \mathrm{RL}_{T,i}^{(\mathrm{bl})}$, where $\mathrm{RL}_{T,i}$ is the empirical $(1 - 1/T)$ quantile of annual maxima

3 Results

We first assess surrogate fidelity (Section 3.1) and then examine the basin-wide effort–residual trade-off (Section 3.2). We then analyze the spatial structure of optimized attenuation in three steps: where effort is allocated across the network (Section 3.3), where residual risk remains and how it amplifies across climate scenarios (Section 3.4), and how fixed-budget controlled outcomes vary across the climate ensemble (Section 3.5).

3.1 Surrogate Fidelity

The windowed DMDc surrogate reproduces flood dynamics across the Brazos River basin, with the strongest performance along the mainstem and major tributaries that dominate basin-integrated flood volume (Figure 2a). Weaker performance is concentrated in low-order headwater reaches in the western basin, where intermittent flow regimes are more difficult to approximate linearly. These reaches contribute minimally to basin-integrated flood volume. Compared with a single full-period fit, the windowed approach (10-year windows, 5-year step) improves all validation metrics: ensemble-median NSE increases from 0.54 to 0.69, KGE increases from 0.52 to 0.71, and |PBIAS| decreases from 0.15 to 0.08. This improvement is consistent across all 14 GCM–hydrologic model combinations (Figure 2b–d). A single operator must represent 80 years of evolving dynamics with one set of matrices, which can conflate distinct hydrologic regimes. The windowed formulation reduces this limitation by estimating separate linear operators over successive temporal windows. Within each window, rainfall–runoff dynamics are treated as approximately stationary; re-identification across windows then updates the surrogate as projected hydroclimatic conditions evolve. Because the controller acts only during exceedance events, we also evaluate surrogate fidelity at high flows in the Supporting Information (Figure S1). Basin-total return-level bias remains within ±5% across all return periods (Figure S1b), and bias converges toward zero with increasing stream order (Figure S1c). These results indicate that the high-order reaches that dominate basin-total flood volume are modeled most accurately when the controller is active. Basin-aggregate mass conservation is further verified in Figure S2.

3.2 Effort–Residual Trade-off

Figure 3 shows the temporal evolution of basin-integrated flood exceedance when each ensemble member is evaluated under a common fixed attenuation budget of 10 $\mathrm{Gm^3\,yr^{-1}}$.

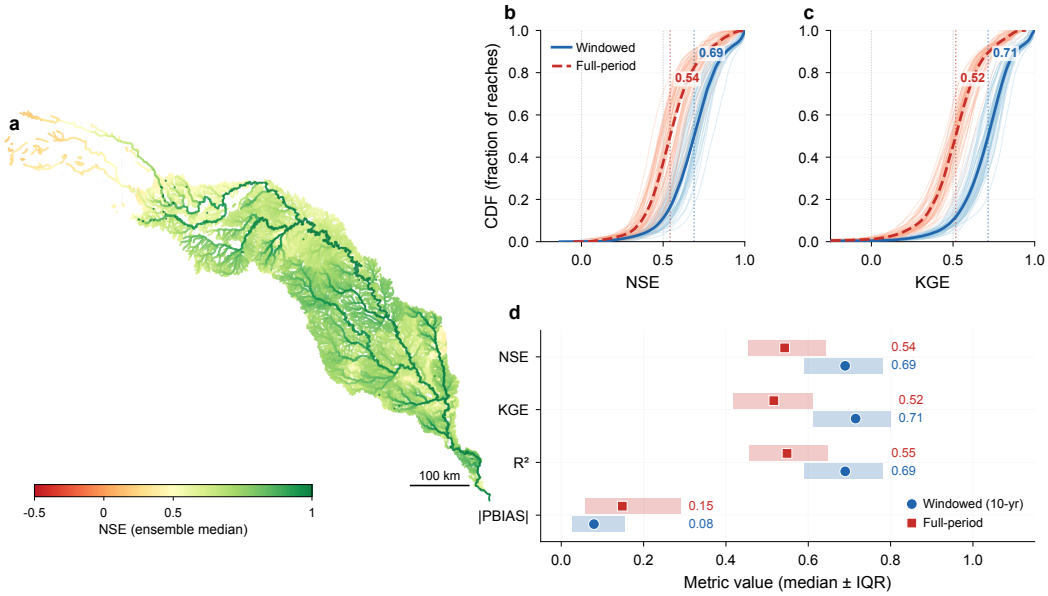

Figure 2. Surrogate validation comparing windowed (10-year windows, 5-year step) and full-period DMDc fitting. (a) Spatial distribution of ensemble-median NSE under the windowed approach. (b,c) Cumulative distributions of NSE and KGE; solid blue lines show the windowed approach and dashed red lines show the full-period fit; light shading shows individual ensemble members (7 GCMs × 2 hydrologic models). (d) Multi-metric comparison showing median (dots) and interquartile range (bars) across all reaches.

Under all four SSPs, controlled exceedance trajectories remain below the baseline throughout the 80-year projection period (Figure 3a–d), indicating sustained flood attenuation under closed-loop control. The flow duration curves show that attenuation is concentrated in the upper tail of the discharge distribution, with the largest absolute reductions at the lowest exceedance probabilities (Figure 3e). This pattern follows from the threshold-triggered controller design: attenuation activates only when discharge exceeds the 2-year return level, so ordinary flows pass unmodified while flood peaks are preferentially suppressed.

Because the same R_w can produce different total attenuation volumes across scenarios, direct cross-scenario comparison at a single penalty value is ambiguous. We therefore sweep R_w across the full penalty range and compare scenarios using the trade-off between total attenuation and residual flood volume (Figure 4a). All four SSPs exhibit a common diminishing-return structure: marginal reductions in residual flood volume decline smoothly as total attenuation increases. Higher-emission scenarios show systematically higher flood exceedances above SSP1-2.6 across the full effort range. At the reference level of 10 $\mathrm{Gm^3\,yr^{-1}}$, residual flood volume under SSP5-8.5 is roughly 60% higher ($\sim$1.6×) than under SSP1-2.6 (Figure 4a,b). This separation is consistent with stronger high-flow forcing under higher-emission pathways and implies that closing the gap would require substantially more annualized attenuation than the reference level.

Fixing total attenuation at 10 $\mathrm{Gm^3\,yr^{-1}}$ enables a consistent multi-metric comparison across scenarios (Figure 4b,c). The residual-risk ratios show that climate-driven amplification of flood impacts depends on the flood metric considered. Excess flood volume exhibits the strongest scenario dependence, with the upper end of the inter-model range exceeding 2× the SSP1-2.6 reference under SSP5-8.5. Annual peak discharge and return-

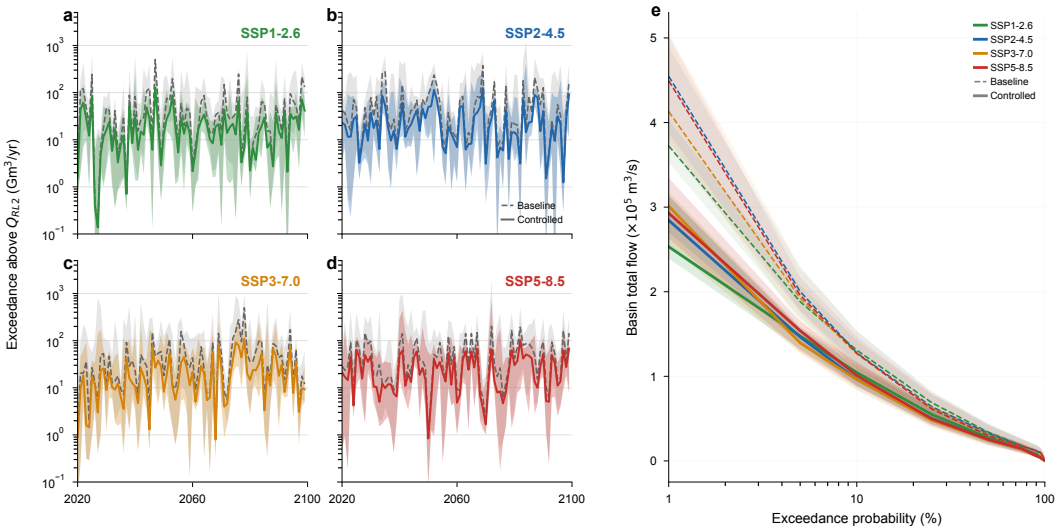

Figure 3. Temporal evolution of basin-integrated flood exceedance. (a–d) Annual exceedance volume above the 2-year return level for SSP1-2.6 through SSP5-8.5; dashed lines show baseline (open-loop surrogate), solid lines show controlled (closed-loop). Shading spans the inter-model range (7 GCMs × 2 hydrologic models). (e) Flow duration curves of basin total flow under baseline (dashed) and controlled (solid) conditions, with shading indicating inter-model spread.

level metrics increase more modestly (Figure 4b). Controlled flood-frequency curves show that scenario separation increases with return period and becomes more pronounced beyond ∼10 years (Figure 4c). Thus, even after optimized attenuation, the upper tail of the flood-frequency distribution remains the most sensitive to forcing pathway and the most uncertain across climate models.

The trade-off curve can be interpreted as the progressive exhaustion of readily attenuated flood exceedance. The steep initial decline reflects the removal of moderate, frequently recurring floods that can be reduced with relatively small amounts of distributed attenuation. As effort increases, the remaining exceedance is increasingly dominated by larger events, so additional attenuation produces diminishing marginal gains. The vertical separation among SSP curves indicates that stronger forcing leaves a larger residual exceedance burden at the same total effort level. Additional effort reduces this burden, but scenario separation persists over the range of budgets evaluated. In the following sections, we use the 10 $\mathrm{Gm^3\,yr^{-1}}$ reference budget to examine how attenuation effort is distributed across the network and where residual risk remains.

3.3 Effort Allocation Structure

Given a fixed basin-wide adaptation budget, a natural question is whether the effort allocation simply tracks a hydrologic size metric, reproducing what a scaling argument would yield without the control framework. In other words, are flood attenuation measures simply applied to the largest river reaches, where the largest potential for discharge attenuation exists? Figure 5 investigates this question by benchmarking the effort allocation against mean annual discharge ($\overline{Q}$) and drainage area, then characterizing the spatial structure of the departures under SSP2-4.5 at the reference budget of 10 $\mathrm{Gm^3\,yr^{-1}}$.

Raw effort (Figure 5a), measured in terms of the volume of water attenuated, spans more than two orders of magnitude across reaches, with the largest volumes concentrated along the mainstem and major tributaries. A log–log regression against mean annual dis-

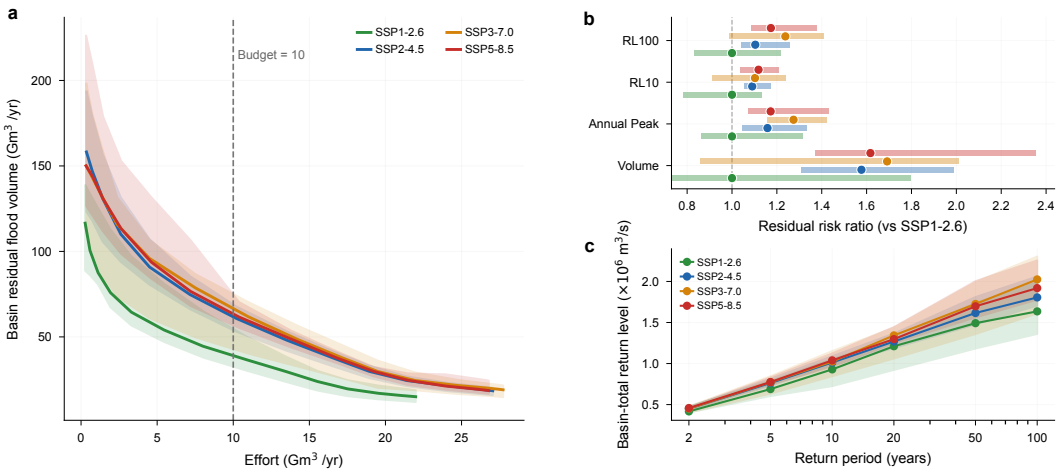

Figure 4. Basin-scale effort–residual trade-off. (a) Trade-off curve between total attenuation and residual flood volume for four SSPs; the dashed line marks a reference level of 10 $\mathrm{Gm}^3\,\mathrm{yr}^{-1}$. Shading indicates inter-model spread (7 GCMs × 2 hydrologic models). (b) Residual risk ratios (relative to SSP1-2.6) for multiple flood metrics at the reference level; dots show ensemble-median values and bars span the inter-model range. (c) Controlled flood frequency curves (flood magnitude vs. return period) with inter-model uncertainty.

charge yields $R^2 = 0.90$ with a sub-linear slope of 0.81 (Figure S3a), indicating that mean flow magnitude is the dominant predictor of applied control effort. By contrast, drainage area alone explains only 33% of the variance (Figure S3b), reflecting the strong east–west precipitation gradient across the Brazos River basin. Here, arid western headwaters drain large areas but generate modest discharge, while more humid eastern tributaries carry disproportionately large flows relative to their contributing area. Because flow magnitude, not contributing area, governs the allocation of adaptation measures, the proposed flood control planning framework encodes the basin's hydroclimatic structure rather than simply its network topology.

Normalizing effort by mean annual discharge (Figure 5b) removes the dominant flow-magnitude dependence and reveals the effort intensity, i.e., how much attenuation the optimal adaptation strategy demands per unit mean discharge. The resulting map shows a qualitatively different spatial pattern from the raw effort. Specific mid-basin tributaries and transitional reaches between headwaters and the mainstem display the highest intensities, indicating that the optimized adaptation scheme prioritizes these locations well beyond what their flow alone would warrant. Mainstem reaches, despite receiving the largest absolute effort, have relatively low intensity of effort when normalized by mean discharge. This low relative intensity can be explained by the fact that adaptation efforts applied to mainstem reaches have limited downstream impact. While flow attenuation achieved at headwater streams will accrue 'bonus' flood mitigation effects in receiving reaches, this bonus effect is limited for flow attenuation measures applied at mainstem reaches that are located near the terminus of the drainage network.

To isolate these departures formally, we regress log-transformed effort against log mean discharge and map the residuals (Figure 5c). Positive residuals (red) identify reaches receiving more effort than the mean-flow power law predicts; negative residuals (blue) mark those receiving less. The residual map reveals a coherent, physically interpretable spatial structure rather than random scatter. This pattern indicates that the LQR control scheme does not merely scale effort with flow magnitude but redistributes it within

the network, allocating disproportionate effort to specific mid-basin reaches where attenuation demand exceeds what local flow magnitude alone would suggest. Conversely, reaches that benefit from cumulative upstream attenuation receive less effort per unit discharge than the basin-wide power law would predict.

The priority-mismatch map (Figure 5d) tests whether the high-attenuation reaches identified in panel c could have been flagged using local reach attributes alone. We rank all reaches independently by three criteria—LQR effort, mean annual discharge ($\overline{Q}$), and drainage area (DA)—and compare the top 10% under each ranking. Red, blue, and green segments mark reaches that appear in only one ranking (LQR-unique, $\overline{Q}$-unique, and DA-unique, respectively), while gray indicates reaches flagged by two or more. The 232 LQR-unique reaches are predominantly mid-order tributaries (median stream order 3, median drainage area 256 km^2), too small to enter the top decile of either flow or area rankings but identified by the framework as requiring disproportionate attenuation relative to their hydrologic size. That these reaches cannot be recovered from local attributes confirms that the spatial allocation encodes network-level interactions—upstream–downstream coupling and confluence timing—that system-level analysis resolves but reach-attribute-based reasoning cannot.

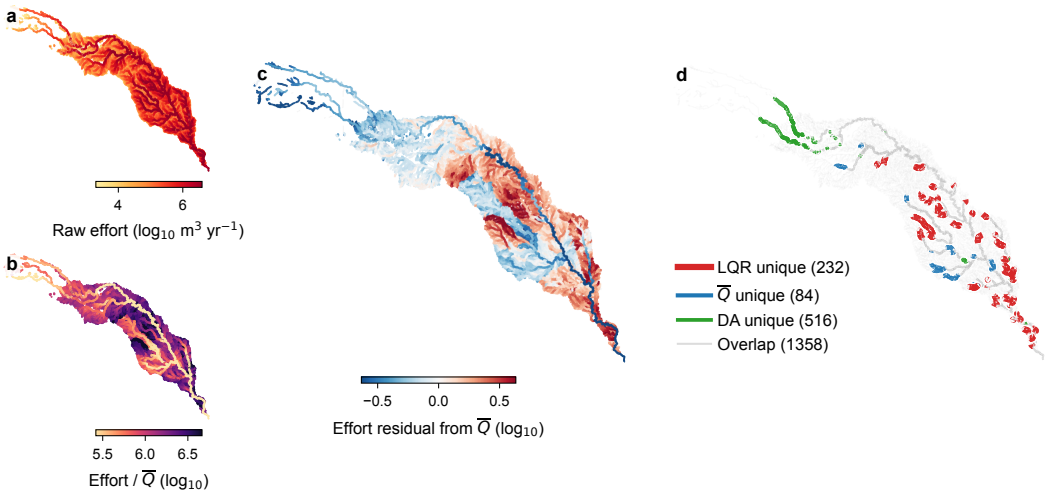

Figure 5. Spatial structure of effort allocation under SSP2-4.5 at the reference attenuation level of 10 Gm^3 yr^{-1}. (a) Raw effort at each reach. (b) Effort normalized by mean annual discharge ($\overline{Q}$). (c) Residual from regressing log effort against log $\overline{Q}$; red indicates more effort than $\overline{Q}$ scaling predicts, blue indicates less. (d) Priority-mismatch map: reaches are ranked independently by LQR effort, $\overline{Q}$, and DA, and the top 10% under each ranking are compared. Colored segments appear in only one ranking; gray appears in two or more.

3.4 Spatial Structure of Residual Risk

The preceding section showed that optimized effort allocation largely tracks mean-flow scaling, with systematic departures at selected tributary and transitional reaches. We now examine where residual exceedance remains after attenuation is implemented and how this residual varies with climate forcing. Figure 6a–c maps reach-level residual exceedance volume relative to SSP1-2.6 (the lowest-emission scenario) for the three higher-emission pathways. The amplification of residual flood risk due to climate change is not spatially uniform: mainstem and high-order tributary reaches show ratios of 2–4×, indicating that major rivers show significantly greater residual flood risk under higher-

emission scenarios, while most headwater reaches remain near 1:1. Thus, higher-emission scenarios produce the largest residual increases along the largest streams in the network.

The stream-order decomposition (Figure 6d) clarifies the outsize impact of climate change on larger streams. High-order reaches (order 5+) integrate runoff from many upstream branches, so scenario-driven increases in runoff appear as larger and more spatially coherent residual exceedance along major channels. Low-order reaches (order 1–2) have smaller absolute residual volumes but greater ensemble-to-ensemble variability, reflecting their stronger dependence on localized storm placement and intensity. From a flood management perspective, the major-channel residual pattern is therefore more stable across climate futures, whereas headwater residual risk remains more uncertain relative to its magnitude.

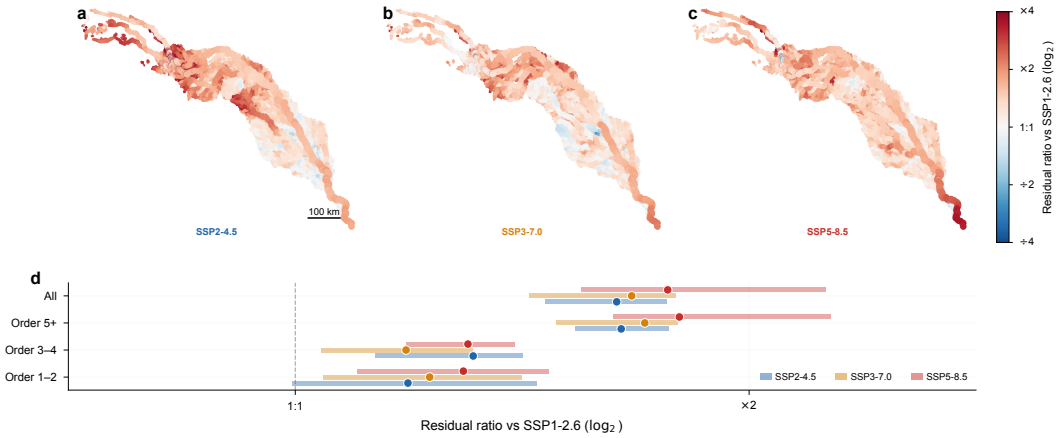

Figure 6. Spatial distribution of residual risk amplification at the reference attenuation level of 10 $Gm^3\,yr^{-1}$. (a–c) Reach-level ratio of controlled residual exceedance volume relative to SSP1-2.6 for SSP2-4.5, SSP3-7.0, and SSP5-8.5. Color scale is logarithmic (base 2); line widths are scaled by log drainage area. (d) Distribution of residual ratio by Strahler stream order; dots show flow-weighted medians, bars span flow-weighted interquartile ranges.

Tracing residual risk along the mainstem as a function of drainage area reveals how these reach-level patterns accumulate downstream (Figure 7). Both residual exceedance volume and 10-year return-level residual increase on average with drainage area, with some of the largest gains coinciding with major tributary junctions near 50,000, 75,000, and 100,000 km^2 where contributing area grows abruptly. The profiles are not strictly monotonic, because residual exceedance volume is a reach-specific metric dependent on local thresholds, control allocation, and event timing, rather than a conserved downstream flux.

The clearest longitudinal signal is the progressive widening of inter-scenario separation with drainage area: from roughly 1.5× for SSP5-8.5 relative to SSP1-2.6 in the upper basin to approximately 3× near the outlet. This widening is consistent with downstream area integration, as each additional tributary contributes scenario-sensitive residual exceedance from its own drainage area. Downstream reaches therefore face both larger absolute residual volumes and greater scenario separation.

This reach-to-reach variability is a direct consequence of treating the river network as a single coupled dynamical system. Because the DMDc surrogate is trained on network-wide discharge time series, the transition matrix $\mathbf{A}$ encodes how flow perturbations prop-

agate between reaches—including connectivity, travel times, and confluence interactions—and the LQR allocation inherits this system-level structure.

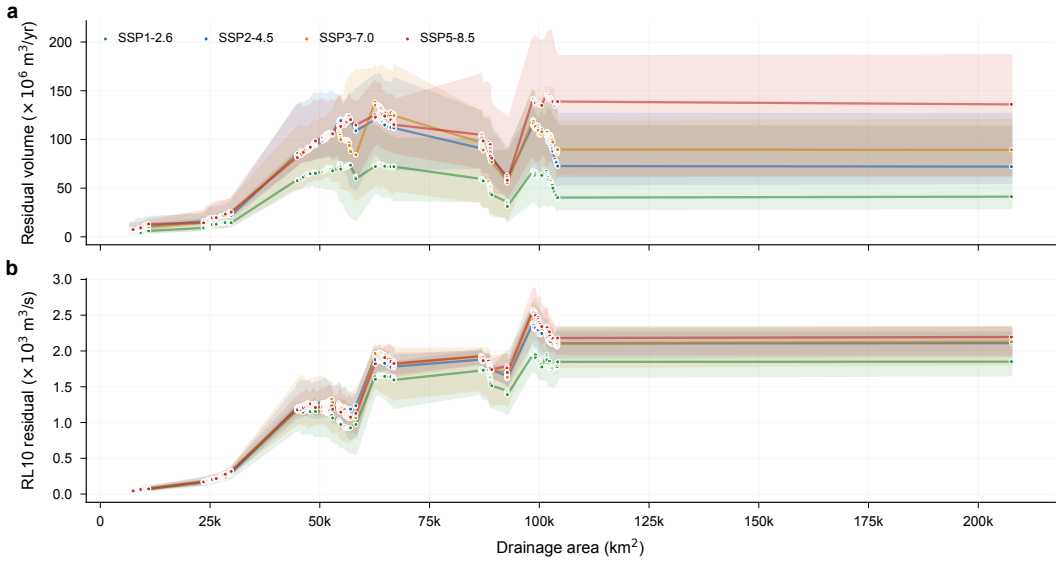

Figure 7. Mainstem profiles of residual risk at the reference attenuation level of 10 $Gm^3\,yr^{-1}$. (a) Residual exceedance volume and (b) 10-year return level residual under controlled conditions. Filled circles show ensemble-median values (7 GCMs × 2 hydrologic models); shading spans the inter-model range. Notable increases near 50,000, 75,000, and 100,000 km^2 coincide with major tributary confluences.

3.5 Scenario Sensitivity and Ensemble Agreement

The previous sections showed where attenuation effort is allocated and where residual exceedance remains after control. We next evaluate whether adaptation benefits remain stable under stronger forcing. To do so, we compare peak-discharge reduction under SSP2-4.5 with the corresponding reduction under SSP5-8.5 at the same 10 $Gm^3\,yr^{-1}$ attenuation budget.

Under SSP2-4.5, optimized attenuation produces spatially uneven peak-discharge reductions across the river network (Figure 8a). The comparison with SSP5-8.5 shows how these benefits change under stronger forcing. In some reaches, the same total attenuation budget produces a smaller peak-reduction benefit under SSP5-8.5; in others, the change is weak or the benefit increases. This heterogeneous response reflects the combined influence of scenario-dependent runoff changes, ensemble-specific hydroclimate projections, and routed network effects, including the way upstream attenuation alters downstream peak flows. The map is therefore best interpreted as a scenario-sensitivity diagnostic of the optimized control solution, rather than as reach-by-reach attribution to a single physical mechanism.

The ensemble agreement map adds a robustness layer to this sensitivity analysis (Figure 8b). For each reach, it counts how many ensemble members agree that peak-reduction benefit decreases from SSP2-4.5 to SSP5-8.5. High agreement identifies reaches where the direction of change is consistent across the ensemble, whereas mixed agreement identifies reaches where the result depends more strongly on climate or hydrologic model choice. In this application, roughly half of reaches show majority agreement on decreased con-

trol benefit under SSP5-8.5, about 30% show majority agreement on increased benefit, and the remaining reaches show mixed ensemble signals.

The stream-order summary shows that these changes are not confined to one part of the drainage hierarchy (Figure 8c). Sensitivity appears across multiple stream orders, indicating that changes in modeled control effectiveness are a network-wide property of the optimized solution rather than a feature of only headwater channels or mainstem reaches.

Together, these results add an uncertainty and confidence layer to the previous maps. The effort-allocation maps identify where control effort is placed (Section 3.3), the residual-risk maps identify where flood exceedance remains after control (Section 3.4), and the sensitivity and agreement maps show whether the modeled control benefits remain consistent across climate futures (Section 3.5). This latter sensitivity assessment makes the framework useful not only for locating priority reaches, but also for distinguishing reaches with robust ensemble support from reaches where adaptive or staged planning may be more appropriate.

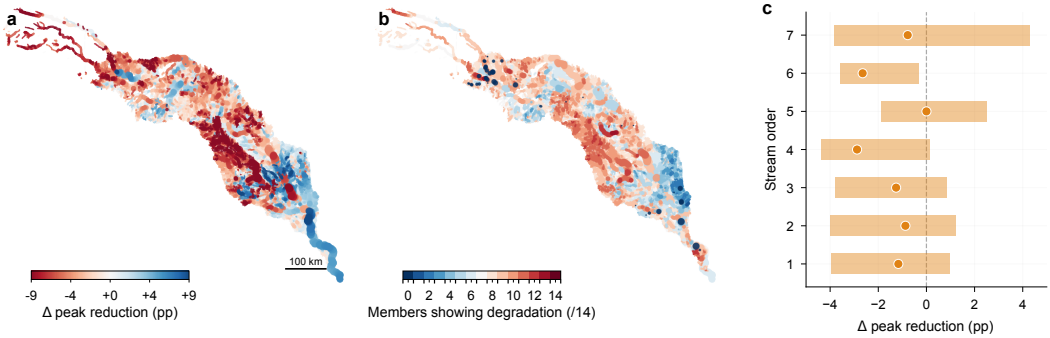

Figure 8. Control effectiveness and robustness across the ensemble at the reference attenuation level of 10 $Gm^3\,yr^{-1}$. (a) Change in ensemble-median peak reduction from SSP2-4.5 to SSP5-8.5 (percentage points; red = degradation, blue = improvement). (b) Map of ensemble agreement on degradation (0–14 scale); each ensemble member is a unique GCM–hydrologic model pair. (c) Distribution of peak reduction change by stream order; dots show medians and bars span the interquartile range.

4 Discussion

4.1 Planning Implications

The results suggest that basin-scale flood adaptation planning can be organized around the combined assessment of flood control capacity, placement, residual risk, and robustness. At a basic level, flood planners need to know the level of investment required to satisfactorily mitigate flood impacts, the spatial locations where flood defenses should be allocated, the level of flooding risk that remains after implementation, and the robustness of the given adaptation strategy in the face of climate uncertainty. The proposed planning framework addresses each of these four criteria in turn.

The effort–residual trade-off analysis (Section 3.2) first defines the level of flood mitigation benefit that is expected for a given level of adaptation effort, and provides water managers with a means to decide an appropriate level of investment. Our analysis shows that increasing attenuation effort reduces flood exceedance monotonically across all climate scenarios considered; however, marginal flood control benefits decline as the remaining risk becomes increasingly associated with large events (Figure 4). The trade-

off curve therefore provides a screening basis for evaluating the scale of adaptation required to meet flood mitigation goals.

The effort-allocation analysis (Section 3.3) reveals the optimal spatial distribution of attenuation needed to reduce flood exceedance, and thus provides water managers with spatially explicit guidance for coordination of flood defenses. In general, raw attenuation effort concentrates along the mainstem and major tributaries, consistent with their larger flood volumes (Figure 5a). These results confirm that major dams located on large rivers are effective targets for basin-scale flood control. However, optimized allocation of attenuation does not merely reflect hydrologic size; in particular, selected mid-order tributaries and transitional reaches receive disproportionate priority even when they are not the largest reaches by discharge or drainage area (Figure 5b–d). These prioritized regions arise because network position—how far a reach sits above the outlet, and how the timing of its hydrograph aligns with those of converging tributaries—governs how much downstream benefit a given attenuation produces. While additional research is needed to contextualize these findings with respect to real-world implementation constraints, these results suggest that decentralized interventions (e.g. small-scale detention storage) may offer greater flood mitigation benefits per unit volume of water attenuated when they are located at these critical locations in the drainage network. Consistent with this distributed-adaptation framing, a stream-order constrained experiment confirms that the bulk of the achievable flood-volume reduction can be realized through distributed control across lower-order portions of the drainage hierarchy, rather than concentrated at the largest channels alone (Figure S4).

The residual-risk analysis (Section 3.4) indicates the level of remaining flood risk that is expected after flood defenses are implemented, and the degree to which this risk varies by spatial location and climate scenario. Generally, we find that higher-emissions scenarios and major rivers exhibit greater flooding risks. Even after optimized attenuation, residual exceedance accumulates along high-order channels and near major confluences, and inter-scenario separation widens downstream as scenario-sensitive tributary inflows are routed through the mainstem (Figures 6 and 7). This distinction is important because the reaches prioritized for attenuation and the reaches where residual flood burden remains concentrated are related but not identical. Mid-basin tributaries with high attenuation priority may be candidate locations for distributed retention, infiltration enhancement, or floodplain reconnection, whereas downstream mainstem reaches may require complementary measures to manage risk, such as improved emergency response or removal of structures from floodplains (Tanoue et al., 2021). These spatial priorities are also broadly stable across the projection period and across imposed total effort levels (Figure S5). Although absolute residual volumes increase in the later period under higher-emission pathways, reach-level effort rankings remain similar, suggesting that rank-based priority maps provide robust starting points even as the required magnitude of attenuation capacity grows.

The scenario-robustness analysis (Section 3.5) qualifies the consistency of the predicted spatial adaptation plan with respect to climate uncertainty, thereby providing water managers with guidance for planning confidence and sequencing. Reaches where peak-reduction benefits degrade consistently across the ensemble indicate locations where maintaining control benefits under stronger forcing may require additional capacity or complementary measures (Figure 8). Reaches with mixed ensemble signals are not necessarily lower-priority; rather, they identify places where flexible, low-regret measures (Hallegatte, 2009) and staged implementation are appropriate.

Together, these diagnostics translate the reach-level optimization results into a practical planning sequence: evaluate the appropriate scale of basin-wide flood control capacity, locate reaches where attenuation provides the greatest network-level flood mitigation potential, identify corridors where residual risk accumulates after adaptation, and

apply ensemble agreement and temporal stability analyses to distinguish robust near-term priorities from less certain interventions that should be implemented adaptively.

4.2 Methodological Considerations and Outlook

Several methodological assumptions define the scope of the diagnostic framework and point to natural extensions. The most immediate extension concerns the surrogate itself. The DMDc operator approximates nonlinear basin dynamics with a linear state-space model in reduced coordinates, a trade-off that enables closed-form LQR solutions across the full network but may reduce fidelity during strongly nonlinear rainfall–runoff episodes. Nonlinear alternatives such as Koopman-based lifting (Kutz et al., 2016) or kernel-based extended DMD (Williams et al., 2015) could improve fidelity for extreme events while preserving compatibility with structured control design. Separately, parametric data-driven representations of nonautonomous systems (Lu & Tartakovsky, 2024) could replace the sequential windowing adopted here with a more systematic treatment of nonstationarity in the external forcing.

A second extension involves the control structure. The assumption $\mathbf{B} = \mathbf{I}_n$ treats every reach as independently controllable, but flood adaptation planning typically operates at the scale of sub-basins or administrative units. The framework accommodates this naturally through the control input matrix $\mathbf{B}$: replacing $\mathbf{B} = \mathbf{I}_n$ with configurations where control acts only at sub-basin outlets reduces the number of independent decision points while preserving the same surrogate dynamics and POD basis. The reduced input matrix $\mathbf{B}_r = \mathbf{U}_r^\top \mathbf{B}$ enters the DARE directly, so the LQR is re-solved in closed form at any planning resolution without re-running process-based simulations. To demonstrate this capability, we tested outlet-only configurations at HUC12, HUC10, and HUC8 resolution (Figure S6), showing that the framework can operate across a range of institutional planning scales. The same mechanism extends to excluding infeasible reaches, weighting columns by local intervention capacity, or matching control units to administrative boundaries—each requiring only a re-solve of the closed-form LQR rather than new simulations. Requirement maps produced at alternative resolutions could also inform related placement problems, such as siting hydraulic control structures within the network (Bartos & Kerkez, 2019) or prioritizing hydrological monitoring (Oh & Bartos, 2025).

A related consideration is the physical interpretation of the lifted control signal. Because the LQR is solved in reduced coordinates and lifted back to physical space through the POD basis, the reach-level signal is not constrained to act in the attenuation direction at every reach and time step. We therefore evaluated the lifted signal a posteriori to test whether this unconstrained representation compromises its interpretation as flood attenuation. At the 10 $\mathrm{Gm^3\,yr^{-1}}$ reference budget, the recovered physical-space signal acts in the intended attenuation direction for the overwhelming majority of exceedance events, with an ensemble-median attenuation fraction of 99%; at the basin scale, attenuation accounts for 95–97% of the absolute control effect across SSPs (Figure S7a,b). The magnitude of the lifted signal is also small relative to local flow: the median per-reach attenuation during exceedance events is 12% of channel flow and never exceeds local discharge in any ensemble member (Figure S7c). These checks support interpreting $|u_i(t)|$ as a diagnostic attenuation-effort metric rather than as a sign-constrained operating rule.

Because the process-based simulations used to train the surrogate omit reservoir operations, existing reservoir capacity is not credited in the control analysis. The resulting targets therefore provide a conservative estimate of the distributed attenuation required across the basin. Although annualized attenuation volume is not directly equivalent to static storage capacity, the reference target of 10 $\mathrm{Gm^3\,yr^{-1}}$ (Figure 4a) is broadly comparable in order of magnitude to the approximately 4.9 $\mathrm{Gm^3}$ of designated flood-pool storage across the nine USACE reservoirs in the Brazos basin (Brazos River Authority

& Halff Associates, 2019). Incorporating reservoir dynamics explicitly, by augmenting the state vector with storage volumes and treating release schedules as constrained control inputs, would yield a coupled channel–reservoir system that partitions projected flood risk into the portion manageable by existing capacity and the residual requiring distributed interventions. Pairing such an augmented formulation with active control strategies such as model predictive control would allow existing reservoir capacity to be adaptively coordinated with distributed measures as climate conditions evolve.

The framework currently uses precipitation as the sole exogenous forcing, appropriate for a rainfall-dominated basin such as the Brazos but insufficient where snowmelt, glacier retreat, or rain-on-snow events (Musselman et al., 2018) drive flood generation. The modeling pipeline can accommodate additional forcing channels: expanding $\tilde{\mathbf{d}}(t)$ to include temperature or snow water equivalent requires re-estimating the disturbance matrix $\mathbf{E}$ without altering the control design. Applying this to snowmelt-driven or mixed-regime basins would test whether the capacity, placement, and robustness diagnostics generalize beyond the rainfall-dominated setting demonstrated here. Similarly, including hydrologic model structures beyond VIC and PRMS would broaden the ensemble's coverage of land-surface process uncertainty. Such an expanded ensemble would enable a formal decomposition of the relative contributions of climate model choice, emission pathway, hydrologic model structure, and surrogate model error to reach-level uncertainty (Steinschneider et al., 2023).

A further advantage follows from how the operators are obtained: because they are identified from input–output data rather than specified a priori, the framework can be re-estimated as new simulation outputs or observational records become available. The windowed identification strategy already demonstrates this property within a single projection: surrogate dynamics are updated every five years to track evolving hydroclimatic conditions (Section 2.3). The same mechanism extends to incorporating updated climate projections, additional GCM ensembles, or monitoring records by re-running the identification and control steps without modifying the framework itself, making the pipeline compatible with adaptive replanning in which requirement maps are periodically revised as new information becomes available (Cohen & Herman, 2021; Fletcher et al., 2019). This re-estimability positions the framework not as a one-time diagnostic but as a recurring component of long-term adaptation planning. Future work should integrate these spatial diagnostic targets into deeply uncertain pathway simulation frameworks (Trindade et al., 2019) to evaluate the robust sequencing of distributed flood interventions over multi-decadal planning horizons.

5 Conclusions

We present a flood-control planning framework for coordinating the placement of decentralized interventions, such as detention storage, floodplain reconnection, wetland restoration, and infiltration enhancement, at the river-basin scale. The framework combines multi-model hydrologic simulations with reduced-order dynamics learning and optimal control to estimate reach-level flow-attenuation requirements under future climate scenarios. By combining POD, DMDc, and LQR, the framework transforms a high-dimensional network-scale problem into a tractable sequence of reduced-order identification and control steps, producing reach-level attenuation targets, effort–residual trade-off curves, and robustness diagnostics. We demonstrate the framework on the Brazos River basin and address three research questions.

How does the effort–residual trade-off vary across scenarios? The trade-off curves share a common diminishing-return shape across all emission pathways. Higher-emission scenarios retain substantially greater residual flood volume at comparable effort levels, implying that closing the gap would require substantially greater attenuation capacity or complementary adaptation measures.

How does the framework distribute effort, where does residual risk concentrate, and how do these patterns relate to network structure? The bulk of the effort allocation tracks mean-flow scaling, but the framework identifies mid-order tributary reaches where attenuation demand substantially exceeds what local hydrologic size would suggest. This result indicates that system-level analysis identifies spatial patterns in attenuation demand that reach-attribute-based screening alone cannot capture. Residual risk follows the drainage-network hierarchy, with inter-scenario separation widening progressively downstream as each tributary compounds climate-sensitive runoff.

How robust are fixed-budget controlled outcomes to inter-model uncertainty? The ensemble agreement map identifies where controlled performance consistently degrades under stronger forcing and where the direction of change remains mixed across GCM–hydrologic model members. This diagnostic provides a confidence layer on the effort-allocation and residual-risk results: high-priority reaches with strong agreement provide stronger support for near-term capacity planning, whereas high-priority reaches with mixed agreement motivate flexible or staged adaptation measures.

More broadly, the proposed framework provides a computationally tractable tool for translating climate projection ensembles into spatially explicit flood adaptation strategies. By identifying where attenuation demand is robust, where residual risk remains concentrated, and where flexibility should be preserved as new information becomes available, the proposed framework supports adaptive flood-planning decisions across large river networks.

Acknowledgments

This study was supported in part by the United States National Science Foundation under Grant No. 2340176.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Open Research

All replication code and intermediate pipeline outputs supporting this study are archived on Zenodo (Oh & Bartos, 2026). The source repository is at `https://github.com/future-water/flood-adaptation-control`. Climate-forced streamflow and precipitation projections used as forcing inputs are publicly available from the ORNL HydroSource record (Kao et al., 2024).

References

Alfieri, L., Bisselink, B., Dottori, F., Naumann, G., de Roo, A., Salamon, P., ... Feyen, L. (2017). Global projections of river flood risk in a warmer world. *Earth's Future*, *5*(2), 171–182. doi: 10.1002/2016EF000485

Alfieri, L., Feyen, L., & Di Baldassarre, G. (2016). Increasing flood risk under climate change: a pan-European assessment of the benefits of four adaptation strategies. *Climatic Change*, *136*(3), 507–521. doi: 10.1007/s10584-016-1641-1

Anderson, B. D., & Moore, J. B. (2007). *Optimal control: linear quadratic methods*. Courier Corporation.

Ayalew, T. B., Krajewski, W. F., & Mantilla, R. (2015). Insights into expected changes in regulated flood frequencies due to the spatial configuration of flood retention ponds. *Journal of Hydrologic Engineering*, *20*(10), 04015010. doi: 10.1061/(ASCE)HE.1943-5584.0001173

Bartos, M., & Kerkez, B. (2019). Hydrograph peak-shaving using a graph-theoretic algorithm for placement of hydraulic control structures. *Advances in Water Resources*, *127*, 167–179. doi: 10.1016/j.advwatres.2019.03.016

Bates, P. D., Horritt, M. S., & Fewtrell, T. J. (2010). A simple inertial formulation of the shallow water equations for efficient two-dimensional flood inundation modelling. *Journal of Hydrology*, *387*(1-2), 33–45. doi: 10.1016/j.jhydrol.2010.03.027

Bentivoglio, R., Isufi, E., Jonkman, S. N., & Taormina, R. (2022). Deep learning methods for flood mapping: a review of existing applications and future research directions. *Hydrology and Earth System Sciences*, *26*(16), 4345–4378. doi: 10.5194/hess-26-4345-2022

Berkooz, G., Holmes, P., & Lumley, J. L. (1993). The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics*, *25*(1), 539–575. doi: 10.1146/annurev.fl.25.010193.002543

Bertoni, F., Giuliani, M., & Castelletti, A. (2020). Integrated design of dam size and operations via reinforcement learning. *Journal of Water Resources Planning and Management*, *146*(4), 04020010. doi: 10.1061/(ASCE)WR.1943-5452.0001182

Black, A., Peskett, L., MacDonald, A., Young, A., Spray, C., Ball, T., ... Werritty, A. (2021). Natural flood management, lag time and catchment scale: Results from an empirical nested catchment study. *Journal of Flood Risk Management*, *14*(3), e12717. doi: 10.1111/jfr3.12717

Blöschl, G., Hall, J., Viglione, A., Perdigão, R. A., Parajka, J., Merz, B., ... others (2019). Changing climate both increases and decreases European river floods. *Nature*, *573*(7772), 108–111. doi: 10.1038/s41586-019-1495-6

Bowes, B. D., Wang, C., Ercan, M. B., Culver, T. B., Beling, P. A., & Goodall, J. L. (2022). Reinforcement learning-based real-time control of coastal urban stormwater systems to mitigate flooding and improve water quality. *Environmental Science: Water Research & Technology*, *8*(10), 2065–2086. doi: 10.1039/d1ew00582k

Brazos River Authority, & Halff Associates. (2019, March). *Lower Brazos floodplain protection planning study* (Tech. Rep.). Brazos River Authority. Retrieved from `https://brazos.org/Portals/0/Documents/LowerBrazosFloodplainStudy/Appendix-A-Previous-Studies-and-Data-Collection.pdf`

Broderick, C., Murphy, C., Wilby, R. L., Matthews, T., Prudhomme, C., & Adamson, M. (2019). Using a scenario-neutral framework to avoid potential maladaptation to future flood risk. *Water Resources Research*, *55*(2), 1079–1104. doi: 10.1029/2018WR023623

Castelletti, A., Ficchì, A., Cominola, A., Segovia, P., Giuliani, M., Wu, W., ... Maestre, J. M. (2023). Model predictive control of water resources systems: A review and research agenda. *Annual Reviews in Control*, *55*, 442–465. doi: 10.1016/j.arcontrol.2023.03.013

Castelletti, A., Galelli, S., Ratto, M., Soncini-Sessa, R., & Young, P. C. (2012). A general framework for dynamic emulation modelling in environmental problems. *Environmental Modelling & Software*, *34*, 5–18. doi: 10.1016/j.envsoft.2012.01.002

Cheng, C., Yang, Y. E., Ryan, R., Yu, Q., & Brabec, E. (2017). Assessing climate change-induced flooding mitigation for adaptation in Boston's Charles River watershed, USA. *Landscape and Urban Planning*, *167*, 25–36. doi: 10.1016/j.landurbplan.2017.05.019

Cohen, J. S., & Herman, J. D. (2021). Dynamic adaptation of water resources systems under uncertainty by learning policy structure and indicators. *Water Resources Research*, *57*(11), e2021WR030433. doi: 10.1029/2021WR030433

Dadson, S. J., Hall, J. W., Murgatroyd, A., Acreman, M., Bates, P., Beven, K.,

... others (2017). A restatement of the natural science evidence concerning catchment-based 'natural' flood management in the UK. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *473*(2199), 20160706. doi: 10.1098/rspa.2016.0706

David, C. H., Maidment, D. R., Niu, G.-Y., Yang, Z.-L., Habets, F., & Eijkhout, V. (2011). River network routing on the NHDPlus dataset. *Journal of Hydrometeorology*, *12*(5), 913–934. doi: 10.1175/2011JHM1345.1

Dottori, F., Mentaschi, L., Bianchi, A., Alfieri, L., & Feyen, L. (2023). Cost-effective adaptation strategies to rising river flood risk in Europe. *Nature Climate Change*, *13*(2), 196–202. doi: 10.1038/s41558-022-01540-0

Fletcher, S., Lickley, M., & Strzepek, K. (2019). Learning about climate change uncertainty enables flexible water infrastructure planning. *Nature Communications*, *10*(1), 1782. doi: 10.1038/s41467-019-09677-x

Ghimire, G. R., Kao, S.-C., & Gangrade, S. (2023). *ORNL hydroclimate projections: Technical documentation* (Tech. Rep. No. ORNL/TM-2023/2912). Oak Ridge National Laboratory. doi: 10.2172/1887712

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., & Reed, P. M. (2016). Curses, tradeoffs, and scalable management: Advancing evolutionary multiobjective direct policy search to improve water reservoir operations. *Journal of Water Resources Planning and Management*, *142*(2), 04015050. doi: 10.1061/(ASCE)WR.1943-5452.0000570

Giuliani, M., Lamontagne, J. R., Reed, P. M., & Castelletti, A. (2021). A state-of-the-art review of optimal reservoir control for managing conflicting demands in a changing world. *Water Resources Research*, *57*(12), e2021WR029927. doi: 10.1029/2021WR029927

Hallegatte, S. (2009). Strategies to adapt to an uncertain climate change. *Global Environmental Change*, *19*(2), 240–247. doi: 10.1016/j.gloenvcha.2008.12.003

Herman, J. D., Quinn, J. D., Steinschneider, S., Giuliani, M., & Fletcher, S. (2020). Climate adaptation as a control problem: Review and perspectives on dynamic water resources planning under uncertainty. *Water Resources Research*, *56*(2), e24389. doi: 10.1029/2019WR025502

Hirabayashi, Y., Mahendran, R., Koirala, S., Konoshima, L., Yamazaki, D., Watanabe, S., ... Kanae, S. (2013). Global flood risk under climate change. *Nature Climate Change*, *3*(9), 816–821. doi: 10.1038/nclimate1911

Kao, S.-C., Ghimire, G. R., DeNeale, S. T., Gangrade, S., Rastogi, D., Ashfaq, M., & Uria Martinez, R. (2022). CMIP6 downscaling and hydrologic projection for the conterminous United States. *Earth's Future*, *10*(12), e2022EF002734. doi: 10.1029/2022EF002734

Kao, S.-C., Ghimire, G. R., & Gangrade, S. (2024). *CMIP6-based multi-model streamflow projections over the conterminous US, version 1.1.* Oak Ridge National Laboratory. doi: 10.13139/OLCF/2318650

Kutz, J. N., Brunton, S. L., Brunton, B. W., & Proctor, J. L. (2016). *Dynamic mode decomposition: data-driven modeling of complex systems.* SIAM. doi: 10.1137/1.9781611974508

Lin, C.-Y., Yang, Y.-C. E., & Moazeni, F. (2024). Flood risks of cyber-physical attacks in a smart storm water system. *Water Resources Research*, *60*(1), e2023WR034827. doi: 10.1029/2023WR034827

Lu, H., & Tartakovsky, D. M. (2024). Data-driven models of nonautonomous systems. *Journal of Computational Physics*, *507*, 112976. doi: 10.1016/j.jcp.2024.112976

Mizukami, N., Clark, M. P., Sampson, K., Nijssen, B., Mao, Y., McMillan, H., ... others (2016). mizuRoute version 1: a river network routing tool for a continental domain water resources applications. *Geoscientific Model Development*, *9*(6), 2223–2238. doi: 10.5194/gmd-9-2223-2016

Musselman, K. N., Lehner, F., Ikeda, K., Clark, M. P., Prein, A. F., Liu, C., ...

Rasmussen, R. (2018). Projected increases and shifts in rain-on-snow flood risk over western North America. *Nature Climate Change*, *8*(9), 808–812. doi: 10.1038/s41558-018-0236-4

Nicholson, A. R., O'Donnell, G. M., Wilkinson, M. E., & Quinn, P. F. (2020). The potential of runoff attenuation features as a natural flood management approach. *Journal of Flood Risk Management*, *13*, e12565. doi: 10.1111/jfr3.12565

Ning, J., Bowes, B. D., Goodall, J. L., & Behl, M. (2022). Data-driven model predictive control for real-time stormwater management. In *2022 American Control Conference (ACC)* (pp. 1438–1443). doi: 10.23919/ACC53348.2022.9867410

Oh, J., & Bartos, M. (2023). Model predictive control of stormwater basins coupled with real-time data assimilation enhances flood and pollution control under uncertainty. *Water Research*, *235*, 119825. doi: 10.1016/j.watres.2023.119825

Oh, J., & Bartos, M. (2025). Scalable, adaptive and risk-informed design of hydrological sensor networks. *Nature Water*, *3*(10), 1144–1154. doi: 10.1038/s44221-025-00496-7

Oh, J., & Bartos, M. (2026). *Replication code and intermediate data for: Data-driven control reveals distributed flood adaptation priorities across large river networks.* Zenodo. Retrieved from `https://doi.org/10.5281/zenodo.20576488` doi: 10.5281/zenodo.20576488

Opperman, J. J., Galloway, G. E., Fargione, J., Mount, J. F., Richter, B. D., & Secchi, S. (2009). Sustainable floodplains through large-scale reconnection to rivers. *Science*, *326*(5959), 1487–1488. doi: 10.1126/science.1178256

Perez-Pedini, C., Limbrunner, J. F., & Vogel, R. M. (2005). Optimal location of infiltration-based best management practices for storm water management. *Journal of Water Resources Planning and Management*, *131*(6), 441–448. doi: 10.1061/(ASCE)0733-9496(2005)131:6(441)

Proctor, J. L., Brunton, S. L., & Kutz, J. N. (2016). Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, *15*(1), 142–161. doi: 10.1137/15M1013857

Prudhomme, C., Wilby, R. L., Crooks, S., Kay, A. L., & Reynard, N. S. (2010). Scenario-neutral approach to climate change impact studies: application to flood risk. *Journal of Hydrology*, *390*(3-4), 198–209. doi: 10.1016/j.jhydrol.2010.06.043

Razavi, S., Tolson, B. A., & Burn, D. H. (2012). Review of surrogate modeling in water resources. *Water Resources Research*, *48*(7). doi: 10.1029/2011WR011527

Steinschneider, S., Herman, J. D., Kucharski, J., Abellera, M., & Ruggiero, P. (2023). Uncertainty decomposition to understand the influence of water systems model error in climate vulnerability assessments. *Water Resources Research*, *59*(1), e2022WR032349. doi: 10.1029/2022WR032349

Tabari, H. (2020). Climate change impact on flood and extreme precipitation increases with water availability. *Scientific Reports*, *10*(1), 13768. doi: 10.1038/s41598-020-70816-2

Tanoue, M., Taguchi, R., Alifu, H., & Hirabayashi, Y. (2021). Residual flood damage under intensive adaptation. *Nature Climate Change*, *11*(10), 823–826. doi: 10.1038/s41558-021-01158-8

Thober, S., Kumar, R., Wanders, N., Marx, A., Pan, M., Rakovec, O., ... Zink, M. (2018). Multi-model ensemble projections of European river floods and high flows at 1.5, 2, and 3 degrees global warming. *Environmental Research Letters*, *13*(1), 014003. doi: 10.1088/1748-9326/aa9e35

Thomas, N. W., Amado, A. A., Schilling, K. E., & Weber, L. J. (2016). Evaluating the efficacy of distributed detention structures to reduce downstream flooding under variable rainfall, antecedent soil, and structural storage conditions. *Advances in Water Resources*, *96*, 74–87. doi: 10.1016/j.advwatres.2016.07.002

Travis, Q. B., & Mays, L. W. (2008). Optimizing retention basin networks. *Journal of Water Resources Planning and Management*, *134*(5), 432–439. doi: 10.1061/(ASCE)0733-9496(2008)134:5(432)

Trindade, B., Reed, P., & Characklis, G. (2019). Deeply uncertain pathways: Integrated multi-city regional water supply infrastructure investment and portfolio management. *Advances in Water Resources*, *134*, 103442. doi: 10.1016/j.advwatres.2019.103442

Ward, P. J., Jongman, B., Aerts, J. C., Bates, P. D., Botzen, W. J., Diaz Loaiza, A., ... others (2017). A global framework for future costs and benefits of river-flood protection in urban areas. *Nature Climate Change*, *7*(9), 642–646. doi: 10.1038/nclimate3350

Wilkerson, G. V. (2008). Improved bankfull discharge prediction using 2-year recurrence-period discharge. *JAWRA Journal of the American Water Resources Association*, *44*(1), 243–257. doi: 10.1111/j.1752-1688.2007.00151.x

Williams, M. O., Rowley, C. W., & Kevrekidis, I. G. (2015). A kernel-based method for data-driven Koopman spectral analysis. *Journal of Computational Dynamics*, *2*(2), 247–265. doi: 10.3934/jcd.2015005

Wolman, M. G., & Miller, J. P. (1960). Magnitude and frequency of forces in geomorphic processes. *The Journal of Geology*, *68*(1), 54–74. doi: 10 .1086 / 626637

Wong, B. P., & Kerkez, B. (2018). Real-time control of urban headwater catchments through linear feedback: Performance, analysis, and site selection. *Water Resources Research*, *54*(10), 7309–7330. doi: 10.1029/2018WR022657

Supporting Information for "Data-driven control reveals distributed flood adaptation priorities across large river networks"

Jeil Oh[1], and Matthew Bartos[1]

[1]Fariborz Maseeh Department of Civil, Architectural and Environmental Engineering,

The University of Texas at Austin, Austin, TX, USA

Contents of this file

1. Figures S1 to S7

Introduction

This supporting information provides supplementary analyses that document the fidelity of the reduced-order surrogate model and the physical interpretability of the optimized control signals reported in the main text. Figure S1 evaluates surrogate fidelity during extreme flows, and Figure S2 verifies basin-aggregate mass conservation. Figure S3 characterizes the scaling of optimal control effort with hydrologic size, and Figure S4 reports stream-order controllability. Figure S5 assesses the robustness of the spatial allocation priorities across climate scenarios and budget levels, and Figure S6 demonstrates the flexibility of the framework across institutional planning resolutions. Figure S7 documents the physical feasibility of the recovered control signals. All figures use the 56-member climate ensemble (7 GCMs $\times$ 2 hydrological models $\times$ 4 SSPs) described in the main text; figures based on a single representative ensemble member identify that member in the caption.

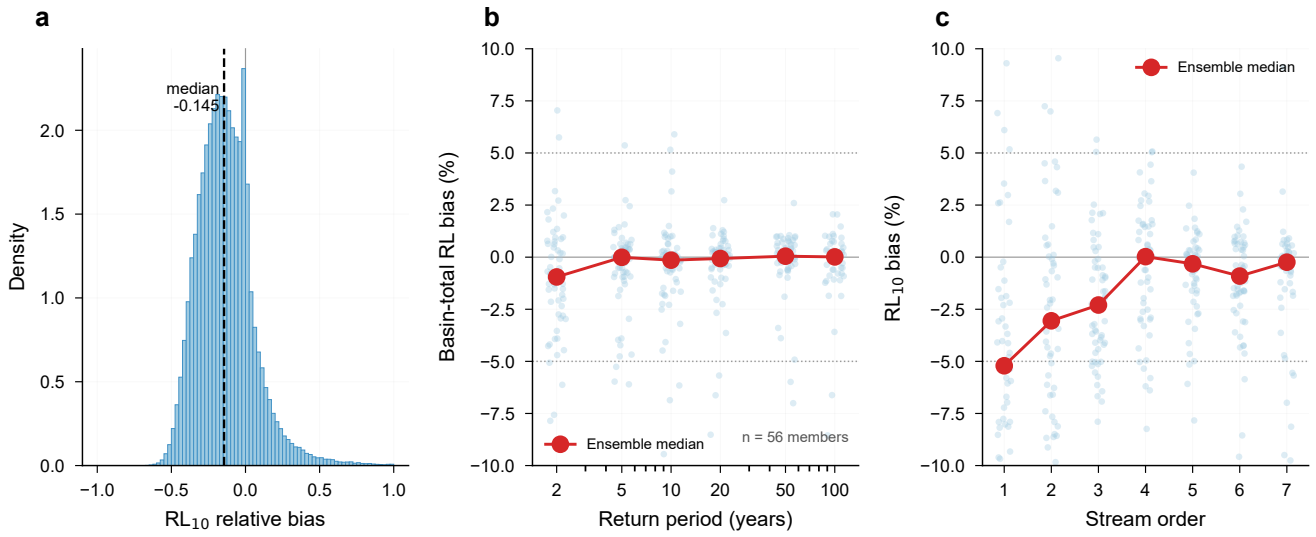

Figure S1. Surrogate-model fidelity during extreme flows, evaluated against reference-model annual maxima across all 56 ensemble members. (a) Distribution of per-reach 10-year return-level relative bias (NSE $>$ 0.5 reaches only); ensemble-median bias is -0.145. (b) Basin-total return-level bias by return period. Daily discharge is summed across all reaches before computing annual maxima and return levels. Each dot shows one ensemble member; the red line connects ensemble medians. Bias remains within $\pm 5\%$ across return periods from 2 to 100 years. (c) The same basin-total return-level bias at the 10-year level, disaggregated by stream-order class.

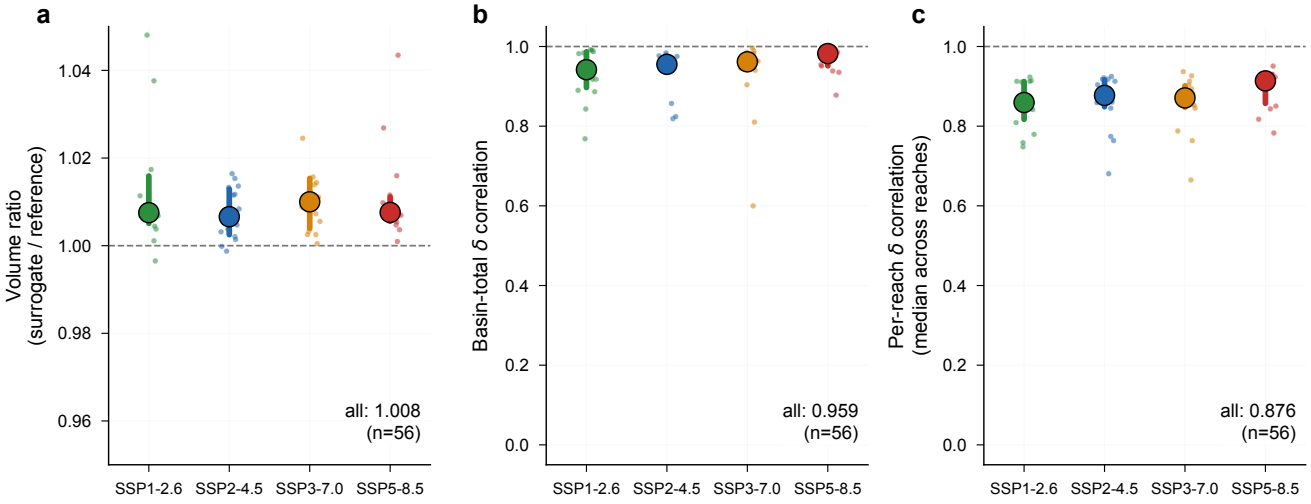

Figure S2. Surrogate mass fidelity and control-signal consistency across all 56 ensemble members (7 GCMs × 2 hydrological models × 4 SSPs). (a) Ratio of surrogate-predicted to reference-model basin-total annual discharge. The dashed line marks unity; ensemble median is 1.008. (b) Correlation r between the reduced-space prediction ($\mathbf{w}_r^\top \delta\mathbf{z}$) and the physical-space reconstruction ($\sum \max(x_{\mathrm{cl}}, 0) - \sum \max(x_{\mathrm{ol}}, 0)$) of the basin-total control effect $\delta(t)$. Ensemble-median correlation is 0.959; the corresponding ensemble-median R^2 is 0.872. (c) Per-reach correlation between the reduced-space and physical-space control effects, summarized as the median across reaches within each ensemble member. Ensemble-median value is 0.876. In all panels, small dots show individual ensemble members, large symbols mark the scenario median, and bars span the interquartile range.

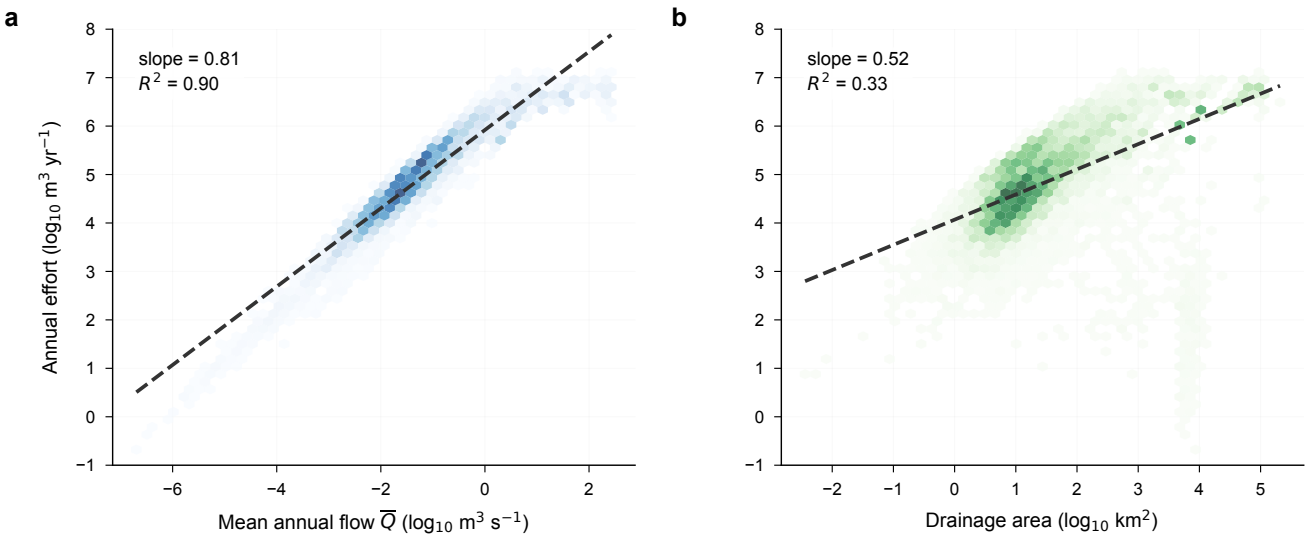

Figure S3. Log–log scaling of optimal control effort with hydrologic size, for SSP2-4.5 at a basin-wide budget of 10 $\mathrm{Gm^3\,yr^{-1}}$. (a) Annual effort versus mean annual flow $\overline{Q}$: slope = 0.81, $R^2 = 0.90$. (b) Annual effort versus drainage area: slope = 0.52, $R^2 = 0.33$. Each hexagon represents the ensemble-median effort for a reach; color intensity indicates point density. Dashed lines show ordinary least-squares fits in log–log space.

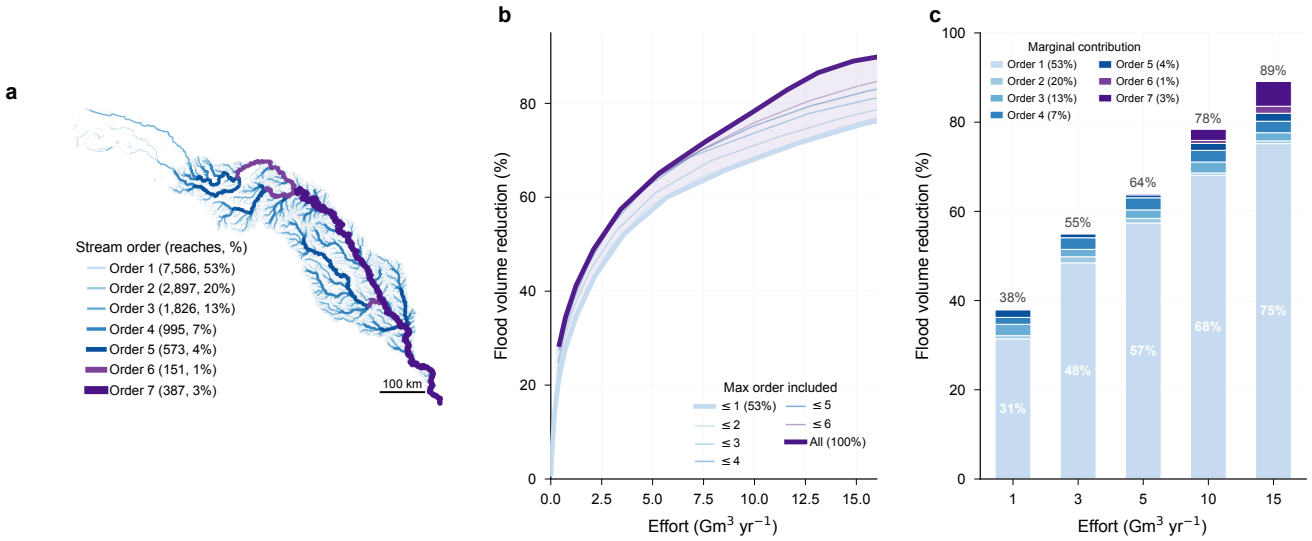

Figure S4. Stream-order controllability for a representative ensemble member (VIC5/MPI-ESM1-2-HR, SSP2-4.5). Restricting the control input matrix $\mathbf{B}$ to reaches of selected stream orders and re-solving the LQR in closed form yields the best achievable flood-volume reduction under that constraint. (a) Brazos river network colored by Strahler stream order; reach counts and percentages in the legend. (b) Effort–reduction trade-off curves for cumulative order subsets. Each curve gives the flood-volume reduction as a function of basin-total effort when control is restricted to reaches of order $\leq N$; the unrestricted case ("All", 100% of reaches) is shown in bold for reference. (c) Marginal contribution of each stream-order class to total flood-volume reduction at fixed effort levels (1, 3, 5, 10, 15 $\mathrm{Gm^3\,yr^{-1}}$). Values inside bars indicate the order-1 contribution; values above bars indicate the cumulative reduction.

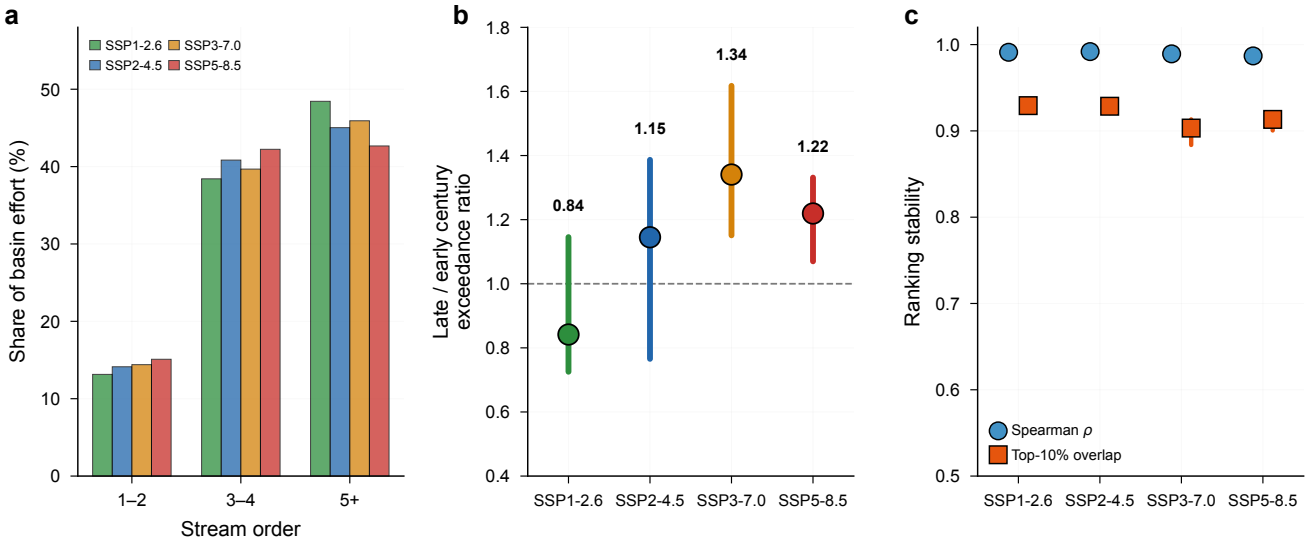

Figure S5. Robustness of spatial allocation priorities. (a) Share of basin-total effort ($\sum_i |u_i|$) by stream-order class, across SSPs. High-order reaches (5+) receive 43–49% of effort. This quantifies where attenuation capacity is installed in physical space. (b) Ratio of late-century (2060–2099) to early-century (2020–2059) controlled exceedance volumes at fixed effort. Values above 1.0 indicate higher residual flood volume in the late century. Medians: 0.84 (SSP1-2.6), 1.15 (SSP2-4.5), 1.34 (SSP3-7.0), 1.22 (SSP5-8.5). (c) Ranking stability across budget levels: Spearman ρ (blue circles) between full reach-level priority rankings at different budgets, and top-10% overlap (orange squares) for the highest-priority reaches. Both metrics exceed 0.9 across all scenarios. In (b) and (c), dots show the ensemble median, and bars span the interquartile range.

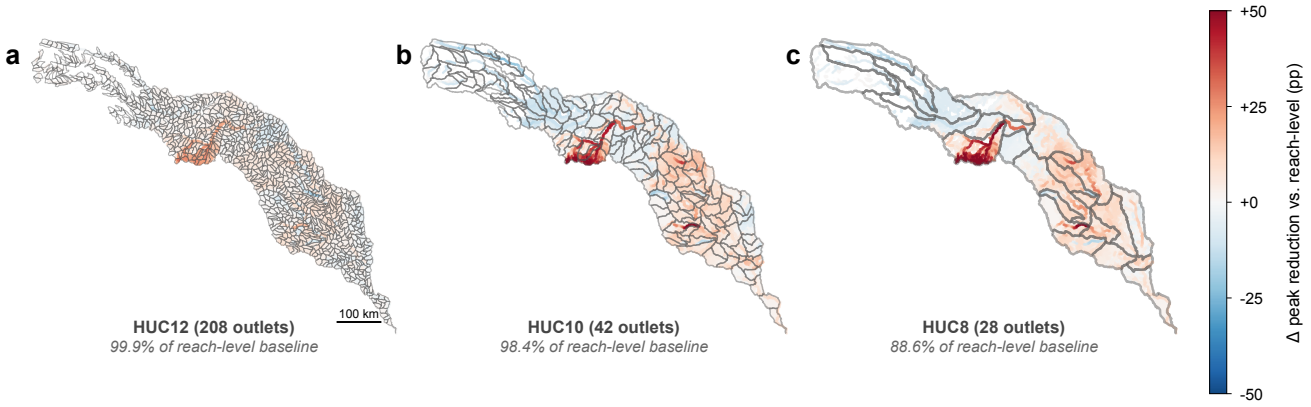

Figure S6. Change in peak discharge reduction when the control input matrix $\mathbf{B}$ is restricted to sub-basin outlet reaches, relative to reach-level control ($\mathbf{B} = \mathbf{I}_n$), at a basin-wide budget of 10 $\mathrm{Gm^3\,yr^{-1}}$ under SSP2-4.5 (VIC5/MPI-ESM1-2-HR). (a) HUC12 outlet control (208 control points). (b) HUC10 outlet control (42 points). (c) HUC8 outlet control (28 points). Red indicates reaches where coarser control reduces peak attenuation relative to the reach-level baseline. Gray lines show hydrologic unit boundaries at each resolution. Percentages beneath each map denote basin-total flood volume reduction as a fraction of the reach-level baseline.

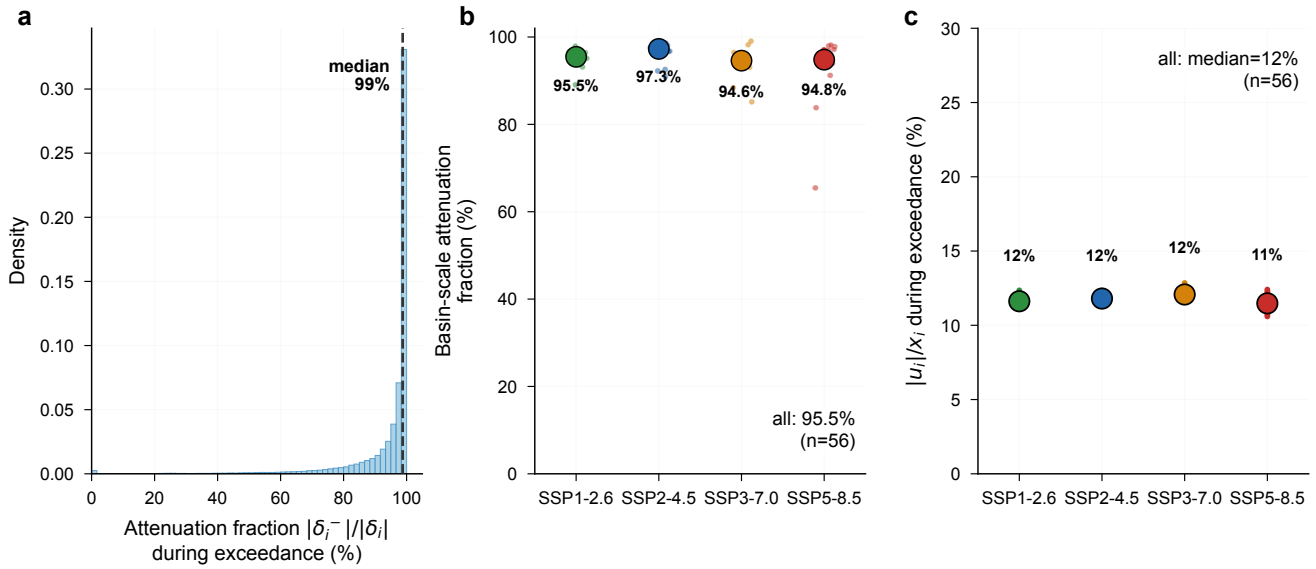

Figure S7. Physical feasibility of the control signal, pooled across all 56 ensemble members at a basin-wide budget of 10 $\mathrm{Gm^3\,yr^{-1}}$. (a) Distribution of the physical-space attenuation fraction $|\delta_i^-|/|\delta_i|$ during flood exceedance events, where $\delta_i(t) = \max(x_{\mathrm{cl},i},\, 0) - \max(x_{\mathrm{ol},i},\, 0)$ is the change in discharge at reach i between the controlled and uncontrolled simulations. The dashed line marks the ensemble median (99%). (b) Basin-scale attenuation fraction per SSP, computed as $\sum|\delta^-|/\sum|\delta|$ across all reaches and time steps. Ensemble medians: 95.5% (SSP1-2.6), 97.3% (SSP2-4.5), 94.6% (SSP3-7.0), 94.8% (SSP5-8.5). (c) Control-input-to-flow ratio $|u_i|/x_i$ during exceedance events; ensemble-median across scenarios is 12%. In (b) and (c), small dots show individual ensemble members, large circles mark the scenario median, and bars span the interquartile range.