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

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

[†] Correspondence: jeoh@utexas.edu

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

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:

- Data-driven dynamics learning and optimal control infer reach-level flood attenuation requirements across large river networks
- The framework uncovers flood bottlenecks at tributary junctions that are invisible to conventional prioritization yet hardest to protect
- Ensemble agreement separates reaches needing confident near-term investment from those requiring flexible adaptive strategies

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

Abstract

Distributed flood adaptation requires knowing where in a river network attenuation effort should concentrate and how much each reach requires, but the spatial coupling, scenario dependence, and high dimensionality of real drainage networks have kept these requirements largely unresolved. We combine data-driven dynamics learning, reduced-order modeling, and optimal control theory into a diagnostic framework that infers reach-level attenuation targets directly from process-based hydrologic simulations without iterative simulation and optimization. Proper Orthogonal Decomposition compresses the network-wide discharge field into a low-rank basis, Dynamic Mode Decomposition with control identifies a linear surrogate of precipitation-driven flood dynamics, and a Linear Quadratic Regulator solves for the spatially distributed attenuation in closed form. Applied to a large river basin under a multi-model, multi-scenario climate ensemble, the effort–residual trade-off follows a common diminishing-return structure across emission pathways, but higher-emission scenarios retain substantially greater residual flood volume at comparable effort levels. The bulk of the allocation tracks mean-flow scaling, yet the framework identifies priority reaches at tributary junctions that neither drainage area nor mean discharge can flag; these reaches retain the highest residual-to-baseline exceedance ratio after optimal control, revealing structurally stubborn bottlenecks where flooding is hardest to attenuate. Inter-scenario separation in residual risk widens progressively downstream, and ensemble agreement on effectiveness degradation distinguishes reaches where investments can proceed with confidence from those requiring flexible, adaptive strategies.

Plain Language Summary

As climate change intensifies flooding, communities face a core planning problem: deciding where to deploy flood defenses across river networks, and sizing them correctly. Typically, planners prioritize the largest rivers or areas with past damage. But because rivers are highly interconnected, a simple ranking ignores how water combining at junctions affects the entire basin. Solving this problem across thousands of river branches and multiple unpredictable climate futures has been computationally overwhelming. To overcome this barrier, we developed a framework that compresses the complex behavior of a river network into a compact mathematical problem that is solvable directly, bypassing trial-and-error computer simulations and producing requirement maps showing how much flood reduction is needed at every location across the basin. When we applied this to a large river basin, we found that the network's most stubborn flood bottlenecks are hidden at specific tributary junctions—critical blind spots that size-based rules overlook. By comparing how different climate models project future flooding, the framework also produces an agreement map that shows planners where to confidently invest in permanent flood infrastructure today, and where to rely on flexible strategies that can be updated as new data become available.

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) that attenuate flood peaks across the drainage network rather than concentrating protection at a few downstream locations. Catchment-scale evidence confirms 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).

Yet the effectiveness of any distributed portfolio depends on how well it matches the spatially varying attenuation demands imposed by basin hydrology: tributary timing, upstream–downstream coupling, and climate-driven shifts in flow generation collectively determine where adaptation effort must concentrate and how much attenuation each part of the network requires (Ayalew et al., 2015; Thomas et al., 2016). The practical objective is to reduce flood peaks across the network so that discharge remains below damage-causing thresholds at every reach. This raises a fundamental planning question that precedes all design and siting decisions: *where* in a river network should flood adaptation effort concentrate, *how much* attenuation is needed at each location, and how do these spatial patterns shift across climate futures? Without spatially resolved answers, adaptation planning cannot distinguish reaches that require substantial intervention from those where investment yields diminishing returns.

Answering these questions at the scale of real river basins confronts three interlocking challenges. First, *hydrologic complexity*: process-based models couple rainfall–runoff generation at each subcatchment with flow routing across the drainage network, so that the flood response at any reach depends on precipitation forcing, antecedent soil conditions, and the timing and magnitude of inflows from all upstream tributaries. The attenuation required at a given location cannot be determined independently of what happens elsewhere in the network. Second, *dimensionality*: operational 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 reshuffles precipitation patterns and flood-generating mechanisms in projection-dependent ways, and robust adaptation planning requires characterizing reach-level requirements not under a single design condition but across ensembles of GCMs and emission pathways, multiplying the computational burden by orders of magnitude. Together, these challenges have left the spatial structure of reach-level adaptation requirements largely unresolved.

A prerequisite for spatially explicit flood adaptation is understanding how flood hazards will evolve under climate change, and here substantial progress has been made. Global and regional projections couple general circulation models 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 delineates 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, 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*. What they do not provide is the complementary information that flood adaptation demands: where attenuation effort should concentrate and how much is needed at each location to bring flood risk within acceptable bounds.

On the response side, two complementary bodies of work address aspects of this problem but leave the reach-level question open. The first focuses on selecting 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 demonstrate that distributed strategies such as detention areas can be highly cost-effective, but evaluate predefined flood mitigation measures at the scale of administrative regions or coarse grid cells, not individual reaches within a basin.

The second body of work applies control and optimization theory to water infrastructure, but in an *operational* capacity: stochastic and multi-objective optimization derive reservoir operating policies (Giuliani et al., 2016, 2021), Model Predictive Control and reinforcement learning enable adaptive real-time management (Bertoni et al., 2020; Castelletti et al., 2023; Oh & Bartos, 2023), linear quadratic feedback has been applied to real-time control of stormwater networks (Wong & Kerkez, 2018; Lin et al., 2024), and data-driven surrogates accelerate the simulation–optimization coupling that these methods require (Razavi et al., 2012; Castelletti et al., 2012). Simulation–optimization has also been used to select cost-effective layouts of detention basins, though tractability limits these approaches to $O(10^1)$–$O(10^2)$ subcatchments (Travis & Mays, 2008; Perez-Pedini et al., 2005), leaving large river networks unaddressed.

More fundamentally, both bodies of work address the question of *which measures to deploy*, but analyzing interventions individually cannot capture how attenuation demands compound across the network—the peak-reduction requirement at any reach depends on what every upstream tributary contributes, a system-level property that localized assessments cannot resolve. Diagnosing the basin as a unified dynamical system provides a top-down view that reveals structural bottlenecks invisible to localized assessments.

Recent work has highlighted connections between control theory and climate uncertainty in the context of adaptation planning (Herman et al., 2020). We build on this perspective, but rather than prescribing operations or investment timing, we use control theory *diagnostically*: to identify *where* in a river network flood discharge exceeds acceptable levels, and *how much* attenuation each reach requires to bring flows back below a prescribed flood threshold. Treated this way, the optimal control signal at each reach becomes a quantitative measure of adaptation need, an equivalent attenuation target that can later be translated into 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 combining data-driven dimensionality reduction with closed-form optimal control. Proper Orthogonal Decomposition (POD) extracts the dominant spatial modes of streamflow variability across the full river network, compressing $O(10^4)$ reaches into a low-rank basis that captures the essential 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 the system state evolves under precipitation forcing. The surrogate is built directly from the output of process-based hydrologic simulations rather than from simplified analytical representations of the underlying hydrology. Because the surrogate operates in reduced coordinates, a Linear Quadratic Regulator (LQR) (Anderson & Moore, 2007) can solve for the control signal in closed form: for any prescribed balance between flood reduction and total effort, the framework produces the spatially distributed attenuation that achieves that balance. Sweeping the effort-penalty parameter traces the effort–residual trade-off curve, and projecting each solution back to physical space through the POD basis yields reach-level requirement maps—spatiotemporally resolved estimates of where attenuation effort concentrates and how much is needed at each location under each climate scenario.

This construction produces three outputs that constitute the main contributions of the paper. First, reach-level requirement maps quantify how much attenuation each reach requires to keep discharge below the flood threshold, and show that, while the bulk of the allocation tracks mean-flow scaling, the framework identifies priority reaches that neither flow magnitude nor drainage area flags. These LQR-unique reaches are the most structurally stubborn bottlenecks in the network, retaining the highest residual-to-baseline exceedance ratio—revealing that proxy rankings overlook precisely the reaches where flooding is hardest to attenuate. Second, the trade-off curve connecting total effort to resid-

ual risk operates at both basin and reach scales, making diminishing-return thresholds explicit and providing a physical basis for staged capacity sizing under deep uncertainty. Third, scenario-ensemble diagnostics separate climate-driven hazard amplification from control-performance degradation, revealing where in the network effectiveness is most sensitive to the emission pathway and where it remains stable across the ensemble. Overall, the framework identifies spatial priorities and quantifies reach-scale flood adaptation requirements from multi-model hydrologic simulation output without iterative simulation–optimization.

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, and where do diminishing returns set in?
2. How does the framework distribute effort across the drainage network, where does residual risk concentrate, and how do these spatial patterns relate to network structure and climate forcing?
3. How robust is the spatial pattern of control effectiveness to inter-model uncertainty, and which reaches show consistent degradation across the ensemble?

2 Methodology

The framework proceeds in three stages (Figure 1): (i) reduced-order system identification, in which POD and DMDc yield a compact linear surrogate of basin dynamics; (ii) optimal feedback design, in which LQR produces optimal reach-level control gains for a range of effort-penalty weights; and (iii) threshold-triggered closed-loop simulation, in which feedback activates only during flood exceedance so that control effort maps directly to attenuation need.

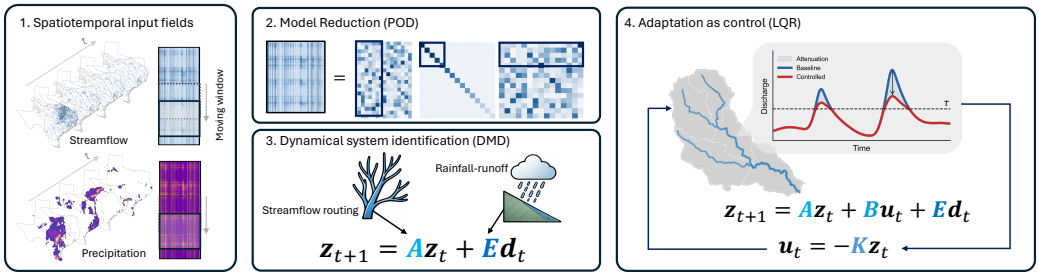

Figure 1. Overview of the diagnostic framework. (1) Spatiotemporal streamflow and precipitation fields from process-based hydrologic simulations are assembled in moving windows. (2) Proper Orthogonal Decomposition (POD) compresses the high-dimensional discharge field into a low-rank basis. (3) Dynamic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

2.1 Study Area and Data

We demonstrate the framework on the Brazos River basin in Texas, USA (∼118,000 km^2), a large, rainfall-dominated basin draining from the semi-arid Southern High Plains to

the Gulf of Mexico. The basin encompasses 14,421 NHDPlusV2 flowlines with spatially heterogeneous flood generation driven by diverse storm types, providing a challenging test case for the framework.

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), which provides routed daily discharge at all NHDPlusV2 reaches across the conterminous United States using the RAPID routing model (David et al., 2011) without 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).

2.2 Problem Formulation

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

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

where $\mathcal{F}$ represents the nonlinear hydrologic processes, including rainfall–runoff generation at each subcatchment and flow routing across the river network. The controlled system augments Eq. (1) with

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

where $\mathbf{B} = \mathbf{I}_n$, reflecting the assumption that every reach is in principle controllable, and $u_i(t)$ is the equivalent attenuation at reach i: the net flow reduction that any combination of interventions (detention storage, infiltration enhancement, floodplain reconnection) must collectively achieve at that location. The framework thus produces a spatially distributed requirement field, conceptually a water budget adjustment at each reach or subcatchment, and leaves the translation of these targets into implementable designs for subsequent engineering analysis.

2.3 Reduced-Order System Identification

2.3.1 State-Space Reduction via POD

We project the high-dimensional discharge field 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

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

We fix $r = 30$, which captures the dominant spatial modes of streamflow variability while keeping subsequent computations 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 a consistent coordinate system for system identification, control, and reconstruction.

2.3.2 DMDc System Identification

Within this fixed POD basis, we identify the basin's response to precipitation disturbance using Dynamic Mode Decomposition with control (DMDc) (Proctor et al., 2016; Kutz et al., 2016). The disturbance 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, $\mathbf{E} \in \mathbb{R}^{r\times n}$ 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 $\rho(\mathbf{A}) < 0.995$, ensuring long-horizon stability.

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

With the precipitation-driven flood dynamics identified across the river network, we augment the surrogate with the adaptation effort from Eq. (2). In physical space, $\mathbf{B} = \mathbf{I}_n$; 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: $\mathbf{A}^{(w)}\mathbf{z}$ propagates internal basin dynamics (routing and storage release), $\mathbf{E}^{(w)}\tilde{\mathbf{d}}$ introduces the precipitation-driven inflow, and $\mathbf{B}_r\mathbf{u}_z$ represents the flood adaptation effort, i.e. the spatially distributed attenuation that interventions must provide. Separating the disturbance channel (learned from data) from the control channel (imposed by design) allows the framework to quantify how much adaptation each reach requires under a given climate scenario without prescribing specific intervention technologies.

2.4 Optimal Control Design

We design feedback policies using the Linear Quadratic Regulator (LQR) (Anderson & Moore, 2007). For a given window with dynamics $\mathbf{A}^{(w)}$ and $\mathbf{B}_r = \mathbf{I}_r$, we minimize the quadratic cost

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

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: small R_w prioritizes state regulation (aggressive attenuation); large R_w limits control magnitude (conservative attenuation). The optimal policy is $\mathbf{u}_z(t) = -\mathbf{K}\mathbf{z}(t)$ 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 $\mathbf{B}_r = \mathbf{I}_r$, these reduce to $\mathbf{K} = (\mathbf{R} + \mathbf{P})^{-1}\mathbf{P}\mathbf{A}$ and the standard DARE. Gains $\mathbf{K}^{(w)}$ are computed separately for each window using the corresponding $\mathbf{A}^{(w)}$.

We sweep R_w over a logarithmic grid spanning 10^{-4} to 10^5 to trace the trade-off curve between flood reduction and total attenuation effort. 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 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

A reach-specific flood threshold $\boldsymbol{\tau} \in \mathbb{R}^n$ determines when control activates. We derive thresholds from pooled annual maximum series (AMS) over 1980–2019. For each reach, annual maximum daily discharge is pooled 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 field

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

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 threshold, $\mathbf{z}^+(t) = \mathbf{0}$ and $\mathbf{u}_z(t) = \mathbf{0}$; 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 $\boldsymbol{\tau}$.

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 $\mathbf{u}(t) = \mathbf{U}_r\mathbf{u}_z(t)$.

2.7 Performance Metrics

We evaluate performance over 2020–2099 using metrics computed from surrogate-reconstructed discharge trajectories. Baseline metrics are computed from open-loop surrogate simulations (no control), and controlled metrics from closed-loop simulations under Eq. (8), ensuring consistent comparisons within the surrogate framework. Table 1

defines the core metrics: excess volume quantifies mean annual flood exposure above the threshold, its fractional reduction measures control effectiveness, total effort quantifies the mean annual attenuation capacity required at each reach, efficiency measures flood reduction achieved per unit effort, peak discharge reduction captures the relative change in maximum flow, and return-level residuals quantify how controlled return levels compare with their uncontrolled counterparts for $T \in \{2, 5, 10, 20, 50, 100\}$ years. Volumetric quantities are summed across reaches for basin-scale totals; non-additive metrics are summarized using medians and interquartile ranges.

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	$\mathrm{RL}_{T,i}^{(\mathrm{cl})}$: empirical $(1 - 1/T)$ quantile of annual maxima under controlled conditions; reported as ratio $\mathrm{RL}_{T,i}^{(\mathrm{cl})} / \mathrm{RL}_{T,i}^{(\mathrm{bl})}$

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), coefficient of determination (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 overlap) against a single full-period fit.

3 Results

We first assess surrogate fidelity (Section 3.1), then examine the basin-wide effort–residual trade-off (Section 3.2). The core of the analysis addresses the spatial structure of the framework's outputs in three stages: how the controller distributes effort across the network (Section 3.3), where residual risk concentrates after optimal attenuation and how it amplifies across climate scenarios (Section 3.4), and which reaches the ensemble consistently identifies as climate-vulnerable (Section 3.5).

3.1 Surrogate Fidelity

The windowed DMDc surrogate faithfully reproduces flood dynamics across the river network, with the strongest performance along the main stem and major tributaries that dominate basin-wide flood generation (Figure 2a). Weaker performance is confined to low-order headwater reaches in the western basin where intermittent flow regimes are 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) substantially improves all metrics—ensemble-median NSE from 0.54 to 0.69, KGE from 0.52 to 0.71, |PBIAS| from 0.15 to 0.08—with consistent separation 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, inevitably conflating distinct hydrologic regimes. The windowed surrogate sidesteps this limitation: within each 10-year window, basin response is sufficiently stationary for a linear operator to capture the dominant precipitation–streamflow dynamics, while sequential re-identification across windows accommodates the nonstationarity of shifting climate regimes. Because the controller acts only during exceedance events, we separately evaluate surrogate fidelity at high flows in the Supporting Information (Figure S1); flow-weighted bias remains small across all return periods, confirming that the reaches dominating basin-total flood volume are modeled accurately when the controller is active. Basin-aggregate mass conservation is further verified in Figure S2.

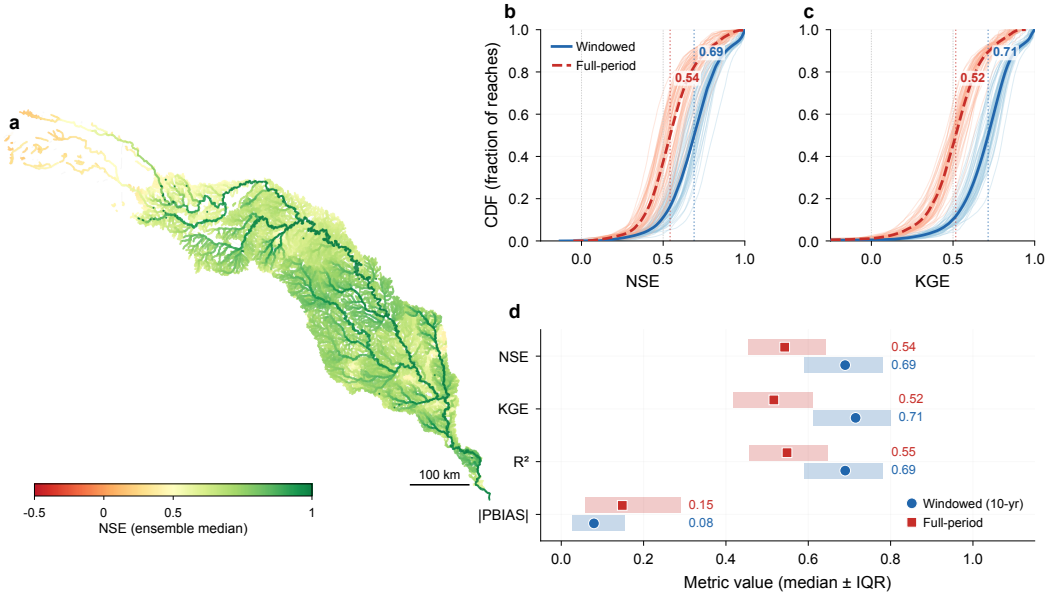

Figure 2. Surrogate validation comparing windowed (10-year windows, 5-year step) and full-period DMDc fitting. (a) Spatial distribution of ensemble-median NSE under the windowed approach. (b,c) Cumulative distributions of NSE and KGE; solid blue lines show the windowed approach and dashed red lines show 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.

3.2 Effort–Residual Trade-off

Figure 3 shows the temporal evolution of basin-integrated flood exceedance under a representative control penalty R_w. Under all four SSPs, controlled trajectories remain consistently below the baseline throughout the 80-year projection period (Figure 3a–d), confirming that the framework achieves sustained flood attenuation across multi-decadal climate projections. The flow duration curves (Figure 3e) show that attenuation is concentrated in the upper tail of the discharge distribution, with the largest absolute reductions at the lowest exceedance probabilities. This is a direct consequence of the threshold-triggered design: the controller activates only when discharge exceeds the 2-year return level, so normal flows pass unmodified while flood peaks are preferentially suppressed.

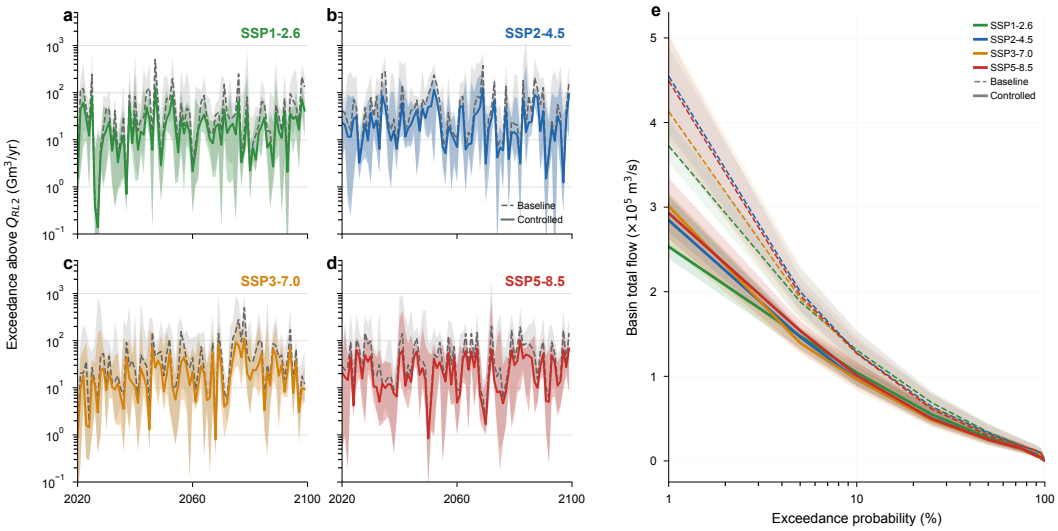

Figure 3. Temporal evolution of basin-integrated flood exceedance. (a–d) Annual exceedance volume above the 2-year return level for SSP1-2.6 through SSP5-8.5; dashed lines show baseline (open-loop surrogate), solid lines show controlled (closed-loop). Shading spans the inter-model range (7 GCMs × 2 hydrologic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

However, different R_w values and climate scenarios produce different total attenuation volumes, making direct cross-scenario comparison at a single R_w ambiguous. To isolate the effect of climate forcing from differences in total effort, we sweep R_w across the full penalty range and construct the trade-off curve between total attenuation and residual flood volume (Figure 4a). All four SSP curves share a convex shape: the first ∼5 Gm^3/yr of attenuation removes the bulk of achievable flood volume, with progressively smaller reductions beyond ∼10 Gm^3/yr. The transition to diminishing effectiveness occurs at a similar attenuation level across scenarios, but the residual flood volume at a given attenuation level differs: at 10 Gm^3/yr, the residual under SSP5-8.5 is roughly 60% higher (∼1.6×) than under SSP1-2.6 (Figure 4a,b). This separation reflects the greater precipitation intensity under higher-emission pathways; closing the gap would require a substantially higher annualized flow-reduction volume than the reference level.

Fixing the attenuation level at 10 Gm^3/yr enables a fair multi-metric comparison across scenarios (Figure 4b,c). The residual risk ratio analysis (Figure 4b) shows that climate-driven amplification varies by flood metric. Excess volume exhibits the strongest scenario dependence, with the upper end of the inter-model range exceeding 2× under SSP5-8.5, whereas annual peak discharge and return-level metrics increase more modestly. Uncertainty widens for metrics associated with rarer events, with the inter-model range increasing toward longer return periods. The controlled flood frequency curves (Figure 4c) show that scenario separation persists across return periods even after optimal attenuation, and becomes more pronounced beyond ∼10 years. The upper tail of the flood frequency distribution, the range most relevant to infrastructure design standards, is thus also the most sensitive to the forcing pathway and the most uncertain across climate models.

The shape of the trade-off curve carries a physical interpretation. The steep initial decline reflects the removal of moderate, frequently recurring exceedance that is readily attenuated by distributed measures. As effort increases, the remaining exceedance is increasingly dominated by the most intense precipitation events, whose peak runoff

volumes are large enough that further attenuation produces diminishing reductions. The vertical separation between SSP curves shows that stronger climate forcing amplifies these intense events, widening the gap between what attenuation can achieve and the flood volumes that precipitation delivers, a gap that persists regardless of how much attenuation is applied. The remainder of the analysis fixes the effort budget at 10 Gm^3/yr and examines how this budget is distributed across the network (Section 3.3) and what residual risk it leaves behind (Section 3.4).

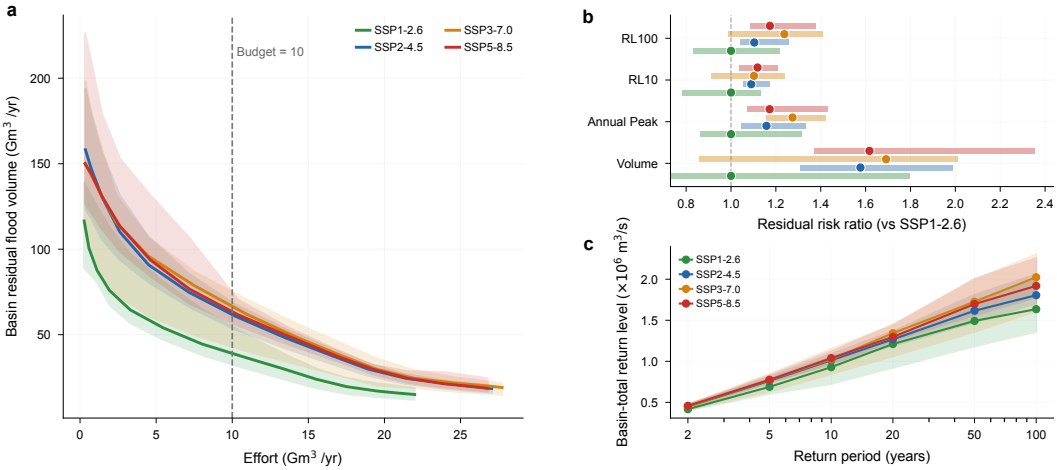

Figure 4. Basin-scale effort–residual trade-off. (a) Trade-off curve between total attenuation and residual flood volume for four SSPs; the dashed line marks a reference level of 10 Gm^3/yr. Shading indicates inter-model spread (7 GCMs × 2 hydrologic models). (b) Residual risk ratios (relative to SSP1-2.6) 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.

3.3 Effort Allocation Structure

Given a fixed basin-wide 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. Figure 5 investigates this question by benchmarking the 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 Gm^3/yr.

Raw effort (Figure 5a) spans more than two orders of magnitude across reaches, with the largest volumes concentrated along the main stem and major tributaries. A log–log regression against mean annual discharge yields $R^2 = 0.90$ with a sub-linear slope of 0.81 (Figure S3a), confirming that flow magnitude is the dominant predictor. By contrast, drainage area alone explains only 33% of the variance (Figure S3b), reflecting the strong precipitation gradient across this semi-arid basin: western headwaters drain large areas but generate modest discharge, while eastern tributaries carry disproportionately large flows relative to their contributing area. Because flow magnitude, not contributing area, governs the allocation, the 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 main stem display the highest intensities, indicating that the adaptation prioritizes these locations well beyond what their flow alone would warrant. Main-stem reaches, despite receiving the largest absolute effort, have moderate intensity because their large discharge dilutes the per-flow signal.

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 does not merely scale effort with flow magnitude but redistributes it within the network. Reaches at key tributary junctions, where flood peaks from multiple upstream branches converge, receive disproportionate effort because early attenuation at these locations prevents exceedance from compounding downstream. Conversely, reaches that benefit from cumulative upstream attenuation receive less effort per unit discharge than the basin-wide power law would predict. These systematic residuals, not the gross allocation to high-flow reaches, constitute the framework's distinctive contribution to spatial prioritization.

The priority-mismatch map (Figure 5d) tests whether these high-effort reaches could have been identified without the control framework. 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 LQR-unique reaches (red) cluster at the same tributary junctions highlighted in panel c, confirming that these are locations the framework prioritizes but that neither $\overline{Q}$ nor DA would flag. The inset bar chart reveals that these LQR-unique reaches are also the most structurally stubborn in the network: they retain the highest fraction of baseline flooding after optimal control (residual-to-baseline ratio of 0.84, versus 0.49 for $\overline{Q}$-unique and 0.46 for DA-unique reaches). Because a high ratio means flooding persists *despite* the controller's effort, the framework identifies network bottlenecks where attenuation is inherently difficult—precisely the reaches that conventional proxy-based prioritization would overlook.

3.4 Spatial Structure of Residual Risk

The preceding section established that effort allocation largely tracks mean-flow scaling, with systematic departures at key tributary junctions. We now examine the residual risk that this allocation leaves behind and how it varies with climate forcing. Figure 6a–c maps reach-level residual exceedance volume relative to SSP1-2.6 for the three higher-emission pathways. The amplification is not spatially uniform: main-stem and high-order tributary reaches show ratios of 2–4×, while most headwater reaches remain near 1:1.

The stream-order decomposition (Figure 6d) reveals a fundamental asymmetry in how the drainage network filters climate-driven changes in flood generation. High-order reaches (order 5+) integrate runoff from large contributing areas, producing high absolute attenuation demands that increase consistently with forcing intensity. Low-order reaches (order 1–2) generate smaller absolute demands but show proportionally greater sensitivity to which climate projection is considered, because their flood regime depends on local storm characteristics that vary more strongly across GCMs. From a flood management perspective, this means that attenuation requirements along major channels can be estimated with greater confidence across climate futures, whereas headwater requirements remain more uncertain relative to their magnitude.

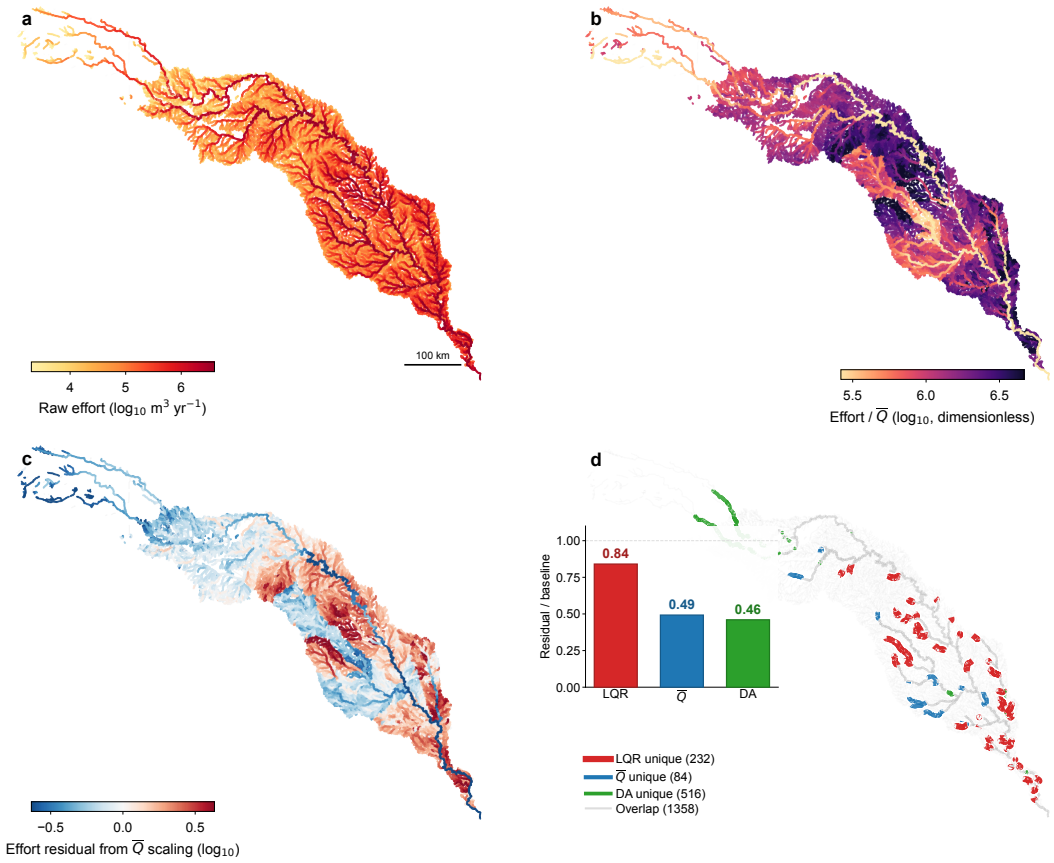

Figure 5. Spatial structure of effort allocation under SSP2-4.5 at the reference attenuation level of 10 Gm^3/yr. (a) Raw effort at each reach. (b) Effort normalized by mean annual discharge ($\overline{Q}$). (c) Residual from regressing log effort against log $\overline{Q}$; red indicates more effort 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. The inset shows the fraction of baseline flooding that persists after optimal control for each unique category.

Tracing residual risk along the main stem as a function of drainage area (Figure 7) places these reach-level patterns in a longitudinal context. Both residual volume and RL10 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 most robust signal is the progressive widening of inter-scenario separation with drainage area: from roughly 1.5× (SSP5-8.5 vs. SSP1-2.6) in the upper basin to approximately 3× near the outlet. This widening follows directly from area integration: each additional tributary compounds the scenario-dependent residual from its own contributing area, and implies that downstream reaches face both larger absolute residual volumes and greater scenario uncertainty. Distributed sub-basin measures should therefore be complemented by targeted interventions, such as off-channel detention or floodplain recon-

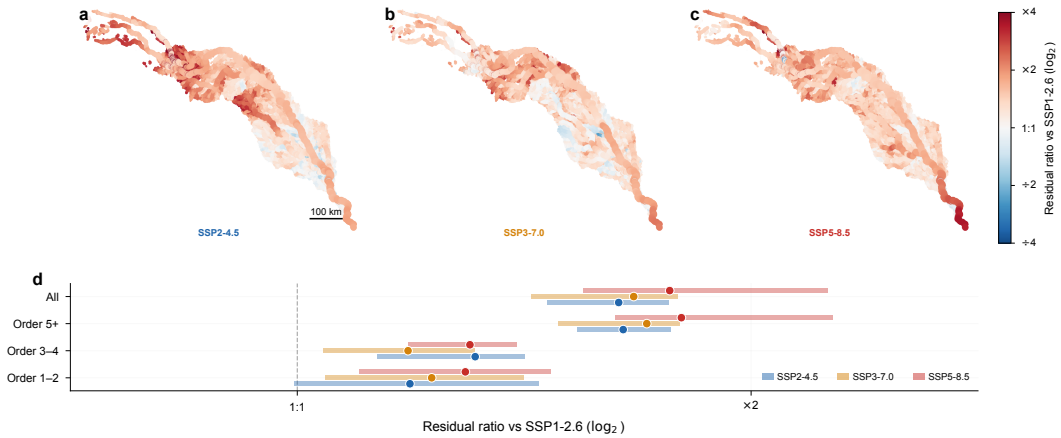

Figure 6. Spatial distribution of residual risk amplification at the reference attenuation level of 10 Gm^3/yr. (a–c) Reach-level ratio of controlled residual exceedance volume relative to SSP1-2.6 for SSP2-4.5, SSP3-7.0, and SSP5-8.5. Color scale is logarithmic (base 2); line widths are scaled by log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

nection, along major downstream channels, where the cumulative residual from upstream branches is greatest.

This reach-to-reach variability reflects a strength of the data-driven surrogate approach. Because the DMDc surrogate is trained on routed discharge rather than local runoff, the transition matrix **A** encodes network connectivity, travel times, and confluence timing, and the LQR allocation inherits this structure, resolving routing-mediated interactions that would be lost at coarser spatial aggregations.

3.5 Identifying Climate-Vulnerable Reaches

The preceding two sections established where effort concentrates and where residual flood exposure is amplified; this section asks which of those reaches become disproportionately more exposed under stronger forcing, and how consistently the ensemble identifies them.

Before examining changes in effectiveness, it is useful to note the baseline pattern. Under SSP2-4.5, ensemble-median peak discharge reduction ranges from less than 10% in headwaters to over 30% along major tributaries and the central main stem, broadly tracking the flow-magnitude gradient in raw effort (Section 3.3). The change in this reduction from SSP2-4.5 to SSP5-8.5 (Figure 8a) reveals where climate forcing erodes control performance: degradation (red) concentrates along the main stem and eastern tributaries, while localized improvement (blue) appears in western headwaters. This pattern reflects the role of contributing area: main-stem reaches integrate additional precipitation-driven runoff from the entire upstream network, so fixed-budget attenuation becomes progressively less effective as climate forcing intensifies.

The ensemble agreement map (Figure 8b) quantifies the confidence behind these spatial patterns. Approximately 50% of reaches exhibit majority agreement ($\geq$8/14 members) on degradation, concentrated along the main stem, while 30% show majority improvement, predominantly in western headwaters. Inter-model agreement varies systematically with stream order: high-order reaches show strong consensus on degradation, while headwater reaches show mixed signals reflecting greater divergence among GCM

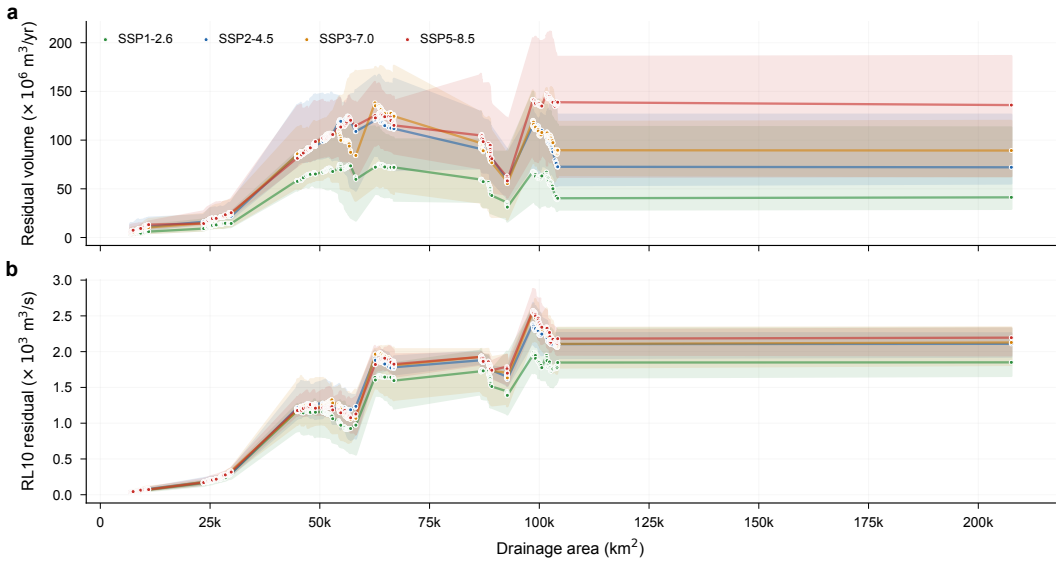

Figure 7. Main-stem profiles of residual risk at the reference attenuation level of 10 Gm^3/yr. (a) Residual exceedance volume and (b) 10-year return level residual under controlled conditions. Filled circles show ensemble-median values (7 GCMs × 2 hydrologic models); shading spans the inter-model range.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

projections at local scales. This separation distinguishes climate-vulnerable reaches, where attenuation needs increase under stronger forcing regardless of which climate model is considered, from reaches where requirements remain projection-dependent and the direction of change is not yet resolved by the ensemble.

The reach-level comparison (Figure 8c) confirms systematic degradation under stronger forcing: the OLS slope of 0.77 indicates that for every 10 percentage points of peak reduction achieved under SSP2-4.5, only ∼7.7 points are retained under SSP5-8.5. The largest departures from the 1:1 line occur at reaches with high baseline reduction, which correspond to high-order channels where larger contributing areas integrate more of the additional runoff generated under intensified precipitation. Low-order reaches cluster near the origin with smaller absolute departures, consistent with their smaller contributing areas and lower sensitivity to basin-integrated forcing changes.

4 Discussion

4.1 Planning Implications

The results collectively suggest a two-tier planning strategy in which the ensemble agreement map (Figure 8b) serves as the organizing criterion. The first tier encompasses reaches where the ensemble consistently agrees on degradation under stronger forcing. At these reaches, four findings converge: effort allocation concentrates along the drainage network hierarchy (43–49% of basin-total effort at order 5+ reaches; Figure 5a, Figure S4a), the allocation tracks mean flow rather than drainage area ($R^2 = 0.90$ vs. 0.33), the priority-mismatch map confirms that LQR identifies reaches missed by flow-based ranking (Figure 5d), and the ensemble agrees on degradation under higher forcing (Figure 8b). Because both the spatial allocation of effort and the identification of vulnerable reaches are robust across SSPs and GCMs, investments along these channels (off-channel detention,

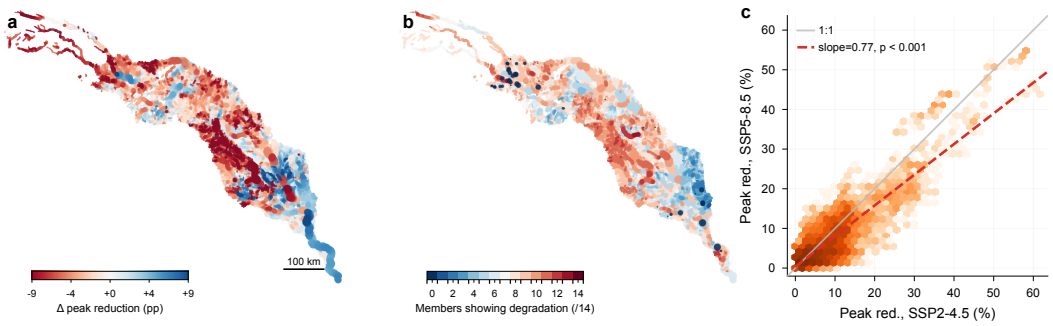

Figure 8. Control effectiveness and robustness across the ensemble at the reference attenuation level of 10 Gm^3/yr. (a) Change in ensemble-median peak reduction from SSP2-4.5 to SSP5-8.5 (percentage points; red = degradation, blue = improvement). (b) Map of ensemble agreement on degradation (0–14 scale); each ensemble member is a unique GCM–hydrologic model pair. (c) Hexbin density of peak reduction under SSP2-4.5 vs. SSP5-8.5; the gray line is 1:1 and the dashed red line shows the ordinary least-squares fit (slope = 0.77, $p < 0.001$).

floodplain reconnection, levee setbacks) can be planned with relatively high confidence regardless of which emission trajectory materializes. Global assessments of flood-risk drivers further support prioritizing high-exposure reaches where projected damage growth is steepest (Winsemius et al., 2016). The trade-off curve (Figure 4a) provides a quantitative basis for sizing this first-tier capacity: a reference budget of ∼10 Gm^3/yr is adopted to compare residual risk across SSPs at equal effort, a level at which the ensemble still shows meaningful marginal returns and the SSP-dependent divergence in residual flood volume becomes clearly visible (Figure 4b).

The second tier addresses reaches where ensemble agreement is mixed—where the 14 members do not converge on whether effectiveness degrades or improves under stronger forcing. These reaches tend to have smaller absolute attenuation demands but proportionally greater inter-model uncertainty (Figure 6d), and the tributary junctions immediately downstream—where the framework departs most from mean-flow scaling (Figure 5c)—depend on upstream inflows being managed effectively. This combination, high strategic value but unresolved ensemble signal, argues for flexible, low-regret measures (Hallegatte, 2009) such as infiltration enhancement, land-use management, and distributed retention that can be scaled or redirected as climate trajectories become clearer. Even with such measures, intensive adaptation leaves residual damage that must be managed through complementary instruments (Tanoue et al., 2021). The framework's ability to distinguish these two regimes quantitatively, using the ensemble agreement map rather than relying on qualitative judgment about where flexibility is needed, is a practical contribution that complements existing adaptive pathway approaches (Haasnoot et al., 2013).

The spatial priorities identified here are broadly stable between the first (2020–2059) and second (2060–2099) halves of the projection period, though absolute residual volumes increase in the later period under higher-emission pathways (Figure S4b). Reach-level effort rankings are similarly insensitive to the total budget imposed (Figure S4c). This temporal stability reinforces the case for near-term investment along high-priority reaches: the locations that most need attenuation capacity today are likely to remain priorities through mid-century, even as the required magnitude of that capacity grows. Because the requirement maps are agnostic to intervention technology, the same reach-level targets can inform detention basin sizing in alluvial valleys, infiltration enhancement on permeable soils, or floodplain reconnection along channelized reaches, depending on local conditions, bridging the gap between the basin-scale diagnostic presented

here and the site-specific design decisions that follow. Translating diagnostic targets into operational adaptation policies requires coupling reach-level requirements with decision rules that trigger implementation as monitored conditions evolve (Cohen & Herman, 2021; Quinn et al., 2017).

4.2 Methodological Considerations and Outlook

The formulation that underlies the framework accommodates several extensions that would broaden its applicability beyond the diagnostic mode demonstrated here. 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 at the cost of 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., 2016) could improve fidelity for extreme events while preserving compatibility with structured control design. Separately, parametric surrogate frameworks such as DRIPS (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 equally controllable, so the resulting effort allocation reflects purely hydrologic attenuation demand rather than local intervention capacity. In practice, $\mathbf{B}$ serves as a design lever: water managers can set columns to zero to exclude reaches where intervention is infeasible (e.g., heavily urbanized corridors or protected areas), scale the remaining columns to reflect spatially heterogeneous capacity (e.g., proportional to floodplain width or available detention area), and re-solve the LQR to test how physical constraints redistribute effort across the network. Overlaying the resulting maps with feasibility layers (land cover, soil infiltration capacity, cost surfaces) would further bridge the gap to site-level design. The requirement maps could also inform related placement problems, such as siting hydraulic control structures within the network (Bartos & Kerkez, 2019) or prioritizing hydrological monitoring to track adaptation performance (Oh & Bartos, 2025).

The process-based simulations from which the surrogate is trained do not represent reservoir operations, so the framework effectively assumes that all flood attenuation must be provided by new distributed interventions, representing a worst-case baseline. Even under this conservative assumption, the implied control magnitudes are physically plausible: the median per-reach attenuation ratio during exceedance events is 12% of channel flow (Figure S5), and in no ensemble member does the control signal exceed the local discharge. The aggregate annual attenuation of 10 Gm^3/yr (Figure 4a) can be compared with the approximately 4.0 million acre-feet ($\sim$4.9 Gm^3) of designated flood pool storage across the nine USACE reservoirs in the Brazos basin (Brazos River Authority & Halff Associates, 2019). Because flood pools can be drawn down and refilled across multiple events within a year, the effective annual throughput of existing reservoir capacity may exceed this static volume, and the diagnostic target of 10 Gm^3/yr thus falls within a physically plausible range relative to existing infrastructure scale. 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 new distributed interventions. Pairing such an augmented formulation with model predictive control would allow existing reservoir capacity to be adaptively coordinated with new 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 pipeline accommodates additional forcing channels naturally: expanding $\tilde{\mathbf{d}}(t)$ to include temperature or snow water equivalent requires only 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 two-tier planning structure and spatial robustness of effort allocation generalize beyond the setting demonstrated here. Similarly, including hydrologic model structures beyond VIC and PRMS would broaden the ensemble's coverage of land-surface process uncertainty.

More broadly, because the entire pipeline is estimated from data rather than from prescribed model structure, its components 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 conditions evolve (Cohen & Herman, 2021; Fletcher et al., 2023). This re-estimability positions the framework not as a one-time diagnostic but as a recurring component of long-term adaptation planning.

5 Conclusions

We presented a reduced-order control framework that infers reach-scale flood attenuation requirements from multi-model hydrologic simulation output. The framework combines POD, DMDc, and LQR to transform a computationally intractable network-scale problem into a sequence of standard linear algebra operations, producing reach-level requirement maps, trade-off curves, and robustness diagnostics without iterative simulation–optimization. We demonstrated the framework on the Brazos River basin and return to the 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, with the knee occurring at a similar effort level regardless of forcing intensity. However, higher-emission scenarios retain substantially greater residual flood volume at comparable effort levels, representing a gap that additional attenuation alone cannot close.

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 reveals flood-vulnerable reaches at tributary junctions that physical proxies alone—such as drainage area and mean discharge—cannot identify. These reaches retain the highest fraction of baseline flooding after optimal control, indicating structurally persistent chokepoints where supplementary adaptation strategies are most needed. Residual risk follows the drainage network hierarchy, with inter-scenario separation widening progressively downstream as each tributary compounds climate-sensitive runoff.

How robust is control effectiveness to inter-model uncertainty? The ensemble agreement map separates reaches where the majority of climate models agree on degradation under stronger forcing from those where the signal remains mixed, providing the organizing criterion for a two-tier strategy: confident investment where agreement is high, flexible measures where it is not.

More broadly, the framework provides water managers and policymakers with a computationally tractable tool for translating climate projection ensembles into spatially explicit adaptation priorities, identifying not only where to invest first, but how much capacity to provision and where flexibility should be preserved for future adjustment.

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

Conflict of Interest

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

Open Research

The CMIP6-based daily streamflow projections used in this study (Kao et al., 2022; Ghimire et al., 2023) are available from the Oak Ridge National Laboratory Constellation repository (https://doi.org/10.13139/OLCF/2318650). NHDPlusV2 flowline attributes are available from the U.S. Environmental Protection Agency (https://www.epa.gov/waterdata/nhdplus-national-hydrography-dataset-plus). Code and processed data will be made available in a public repository upon acceptance for publication.

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.

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.

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.

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.

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.

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 Discussions*, *2022*, 1–50.

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.

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.

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.

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.

Brazos River Authority, & Halff Associates. (2019, March). *Lower brazos floodplain protection planning study* (Tech. Rep.). Brazos River Au-

thority. 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.

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.

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.

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.

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.

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.

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.

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.

Fletcher, S., Zaniolo, M., Zhang, M., & Lickley, M. (2023). Climate oscillation impacts on water supply augmentation planning. *Proceedings of the National Academy of Sciences*, *121*(29), e2215681120.

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.

Giuliani, M., Lamontagne, J. R., Reed, P., & 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.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., & Ter Maat, J. (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global environmental change*, *23*(2), 485–498.

Hallegatte, S. (2009). Strategies to adapt to an uncertain climate change. *Global environmental change*, *19*(2), 240–247.

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.

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.

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.

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.

Lu, H., & Tartakovsky, D. M. (2024). Data-driven models of nonautonomous systems. *Journal of Computational Physics*, *507*, 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.

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.

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.

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.

Oh, J., & Bartos, M. (2025). Scalable, adaptive and risk-informed design of hydrological sensor networks. *Nature Water*, *3*(10), 1144–1154.

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.

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.

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.

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.

Quinn, J. D., Reed, P. M., & Keller, K. (2017). Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. *Environmental modelling & software*, *92*, 125–141.

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

Tabari, H. (2020). Climate change impact on flood and extreme precipitation increases with water availability. *Scientific reports*, *10*(1), 13768.

Tanoue, M., Taguchi, R., Alifu, H., & Hirabayashi, Y. (2021). Residual flood damage under intensive adaptation. *Nature Climate Change*, *11*(10), 823–826.

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.

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.

Travis, Q. B., & Mays, L. W. (2008). Optimizing retention basin networks. *Journal of Water Resources Planning and Management*, *134*(5), 432–439.

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.

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

Williams, M. O., Rowley, C. W., & Kevrekidis, I. G. (2016). A kernel-based method for data-driven koopman spectral analysis. *Journal of Computational Dynamics*, *2*(2), 247–265.

Winsemius, H. C., Aerts, J. C., Van Beek, L. P., Bierkens, M. F., Bouwman, A., Jongman, B., ... others (2016). Global drivers of future river flood risk. *Nature Climate Change*, *6*(4), 381–385.

Wolman, M. G., & Miller, J. P. (1960). Magnitude and frequency of forces in geomorphic processes. *The Journal of geology*, *68*(1), 54–74.

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.

Supporting Information

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

Contents

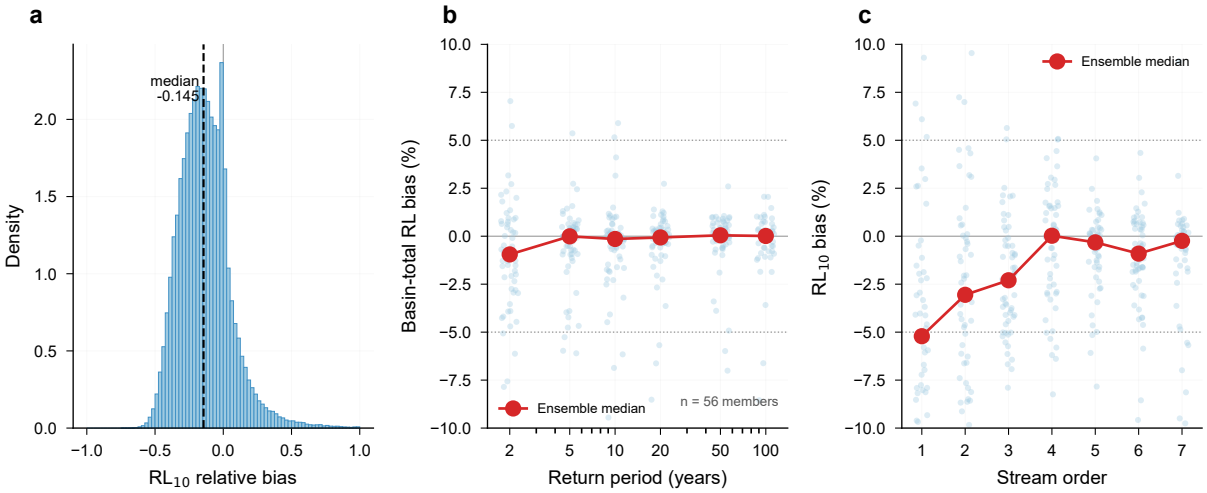

Figure S1: Surrogate-model fidelity during extreme flows. (a) Distribution of per-reach 10-year return-level relative bias (NSE > 0.5 reaches only); the negative median indicates the surrogate slightly underestimates return levels, making controlled flood-reduction estimates conservative. (b) Basin-total 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 (56 total); the red line connects ensemble medians. Bias remains within $\pm 5\%$ across all return periods. (c) Same aggregation applied within each stream-order class for the 10-year return level; bias converges toward zero with increasing stream order, confirming that the high-order channels driving basin-total flood risk are modeled most accurately.

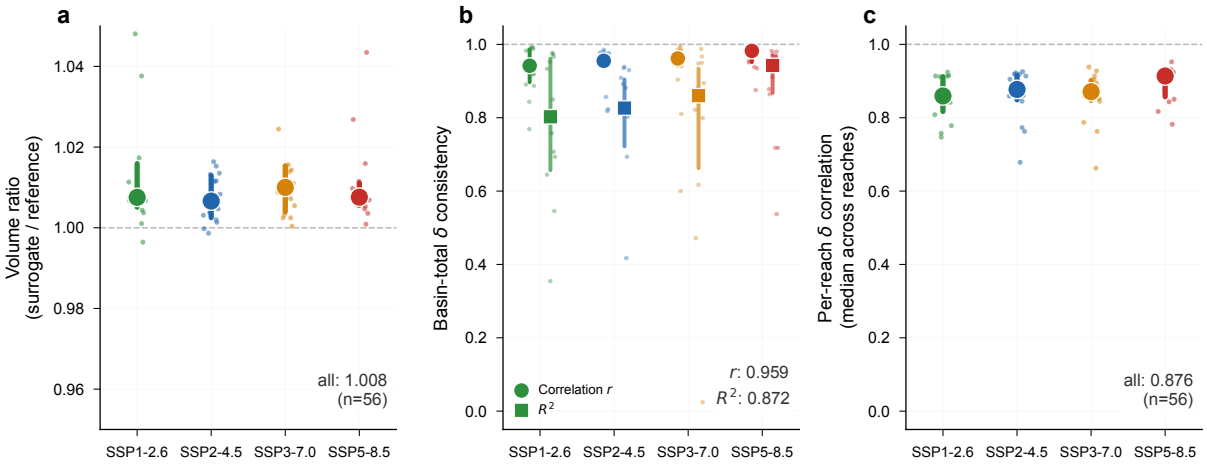

Figure S2: Surrogate fidelity across all 56 ensemble members (7 GCMs × 2 hydrological models × 4 SSPs). (a) Volume ratio of surrogate to reference-model basin-total annual discharge; the dashed line marks unity and the ensemble median is 1.008, confirming negligible mass drift. (b) Consistency of the basin-total control effect $\delta(t)$ between the reduced-space prediction ($\mathbf{w}_r^\top \delta\mathbf{z}$) and the physical-space reconstruction ($\sum \max(x_{\mathrm{cl}}, 0) - \sum \max(x_{\mathrm{ol}}, 0)$), measured by correlation r (circles) and R^2 (squares); the gap between the two metrics reflects residual distortion from non-negativity clipping. (c) Median per-reach correlation between the physical-space and reduced-space control effects, confirming that control signals predicted in reduced space translate accurately to individual reaches. In all panels, small dots show individual ensemble members, large symbols mark the scenario median, and bars span the interquartile range.

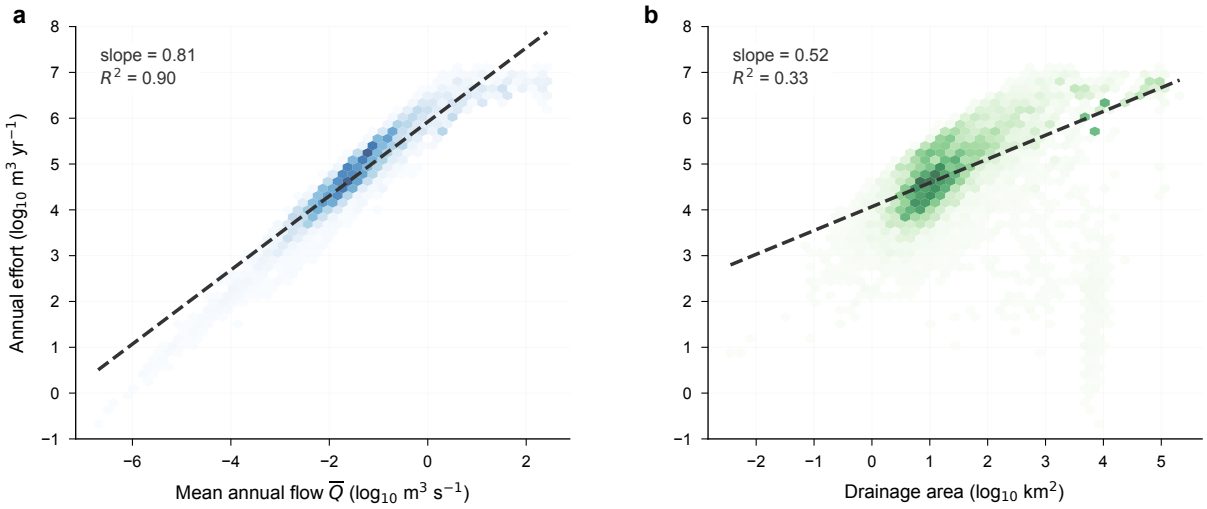

Figure S3: Log–log scaling of optimal control effort with hydrologic size. (a) Annual effort versus mean annual flow $\overline{Q}$ and (b) versus drainage area, both for SSP2-4.5 at a basin-wide budget of 10 $\mathrm{Gm^3\,yr^{-1}}$. Each hexagon represents the ensemble-median effort for a reach; color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

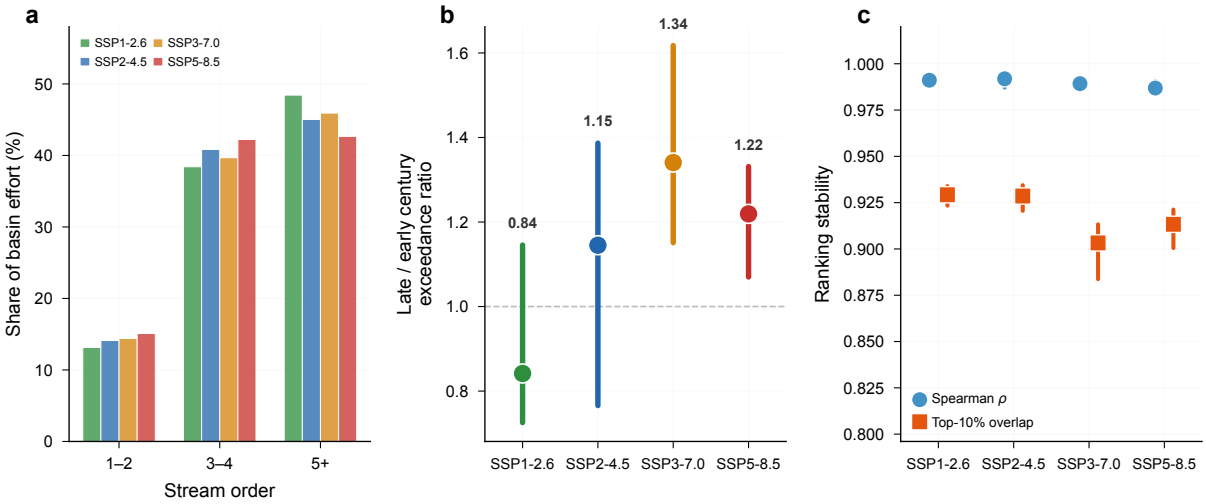

Figure S4: Robustness of spatial allocation priorities. (a) Share of basin-total effort by stream-order class across scenarios. (b) Late/early century controlled exceedance ratio; dots show the median, bars span the interquartile range. (c) Ranking stability across budget levels: Spearman ρ (blue) and top-10% overlap (orange).

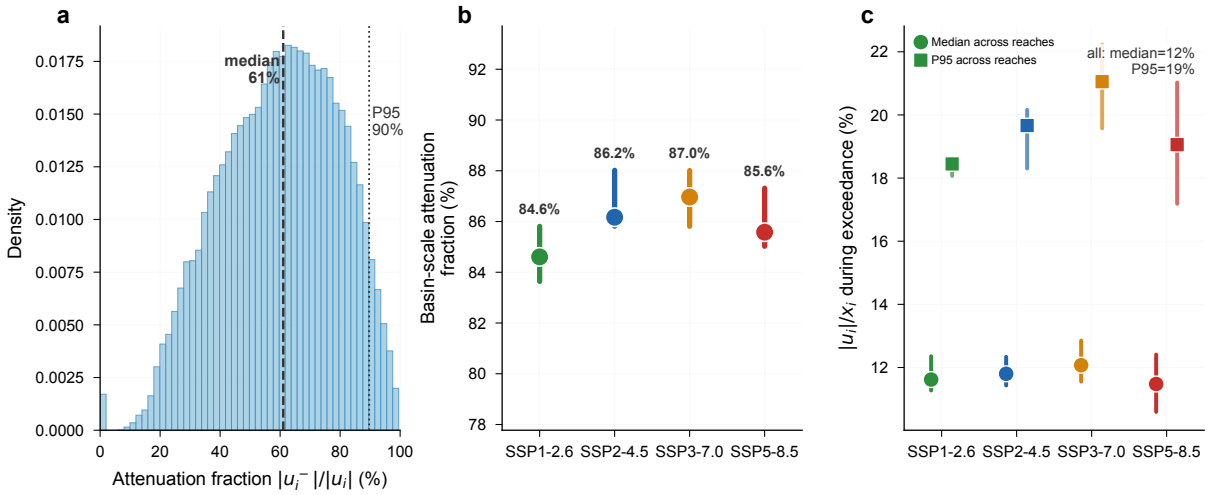

Figure S5: Physical feasibility of the control signal. (a) Distribution of the attenuation fraction $|u_i^-|/|u_i|$ across all reaches and ensemble members; dashed and dotted lines mark the median (61%) and 95th percentile (90%). (b) Basin-scale attenuation fraction per SSP; dots show ensemble-median values and bars span the interquartile range. (c) Control-to-flow ratio $|u_i|/x_i$ during exceedance events per SSP; circles show the median across reaches and squares show the 95th percentile, with whiskers spanning the inter-model interquartile range. The pooled median of 12% confirms that the implied attenuation is a small fraction of channel flow across all scenarios.