FIDE: FIBER-INFORMED DIFFERENTIABLE EXPERIMENTAL DESIGN FOR REFERENCE-RELATIVE LAW DYNAMICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Aggregate measurements generally leave an evolving population law underdetermined, so different measurement systems can be equally adequate for a scientific task while constraining different aspects of the law. We study how to resolve this residual design freedom in baseline-preserving settings, where an existing dynamical model remains the working baseline outside quantities explicitly valued by the task. **F**iber-**I**nformed **D**ifferentiable **E**xperimental Design (**FIDE**) completes each candidate's aggregate constraints by information projection relative to a common frozen reference, then ranks scientifically adequate designs by the dynamical revision required to realize the resulting completed paths. The resulting *Full action* is the minimum mean-squared velocity correction needed to realize a completed path; equivalently, its square root is the sharp reference-relative expectation-rate discrepancy over a gradient-normalized Sobolev class. This defines a conditional minimum-revision principle for baseline-preserving design rather than a generic truth-recovery objective. Experiments show that this preference changes sensor selection, can help or harm withheld downstream predictions depending on baseline reliability, can be used prospectively from aggregate predictions, and remains computationally usable for a 32-dimensional configuration law.

1 INTRODUCTION

Consider an experiment that deploys a small number of aggregate sensors to answer a prespecified scientific question about an evolving population. Several sensor layouts may achieve essentially the same task performance while constraining different aspects of the intermediate law. Because aggregate measurements generally do not identify the full law, these task-equivalent designs can imply different conclusions outside the quantities used to evaluate the task. The primary objective tells us which designs are scientifically adequate, but not how to choose among these remaining law-level differences. Resolving that choice requires an explicit secondary design preference.

Classical and modern optimal experimental design typically chooses measurements to improve information gain, posterior uncertainty, reconstruction quality, or a task-specific utility (Chaloner & Verdinelli, 1995; Foster et al., 2019; Attia et al., 2018). Our question begins after such a primary objective has already identified several scientifically adequate designs. We focus on the baseline-preserving regime: an existing dynamical model is the declared working baseline, and the present experiment is scoped to improve or assess prescribed quantities. Among designs that satisfy that task, we prefer not to introduce additional departures from the baseline unless those departures are valued by another objective or supported by independent evidence.

This regime is natural when new sensors augment an existing modeling or decision pipeline rather than replace it. Related model-credibility and observing-system frameworks likewise evaluate models and measurements within a specified context of use (Morrison et al., 2019; Viceconti et al., 2021; Hoffman & Atlas, 2016; Mourre & Alvarez, 2012). Once the scientific task is already satisfied, leaving the residual choice unspecified means that it may be determined by incidental details of the parameterization, optimization procedure, or sensor family rather than by a declared design preference. Our formulation makes that preference explicit: preserve the working baseline in directions

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

left unresolved by the task unless another objective provides a reason to depart from it. Departures that are themselves scientifically valuable should therefore enter the primary risk.

Within this baseline-preserving setting, we propose **F**iber-**I**nformed **D**ifferentiable **E**xperimental Design (**FIDE**) to resolve the residual design choice. Aggregate measurements do not determine a unique population law, so FIDE first information-projects each candidate's constraints relative to the same frozen reference (Kullback & Leibler, 1951; Jaynes, 1957; Csiszár, 1975). This gives every candidate a completed law path under the same reference-relative rule before any dynamical score is computed.

The completed path need not follow the frozen reference dynamics, so FIDE asks how much those dynamics must change to realize it. The *Full action* is the minimum mean-squared velocity correction required to do so. Equivalently, its square root is the largest reference-relative error in expectation rate, per unit gradient energy, over a normalized class of smooth observables. Thus, after scientific adequacy has been enforced, Full measures dynamical revision of the declared baseline across a broad observable class.

Our contributions are:

- **Residual experimental design.** We formulate baseline-preserving measurement selection as a hierarchical ε-constraint problem in which scientific risk determines adequacy and minimum dynamical revision resolves only the remaining design freedom.
- **Reference-relative dynamical revision.** We characterize Full as the minimum velocity-correction energy required to realize the completed law path, with its square root equivalently giving a sharp expectation-rate discrepancy over a normalized class of smooth observables, and derive its Full–Tangent–Hidden decomposition and matched-initial stability bound.
- **Differentiable design optimization.** For continuous designs, we establish local end-to-end differentiability of the finite-dimensional FIDE objective under explicit nonsingularity conditions and differentiate through both information-projection calibration and the Full correction solve.
- **Empirical characterization and scope.** We study criterion-dependent sensor selection, downstream effects of baseline preservation, prospective selection from aggregate predictions, and a 32-dimensional configuration-law application.

2 Problem Formulation

Let $t \in [0,1]$ denote normalized experimental time and let $X_t \in \mathcal{X}$ have unknown population law P_t, where $\mathcal{X}$ is a domain in $\mathbb{R}^d$ or a flat periodic domain. A measurement design $\boldsymbol{\eta} \in \mathcal{D}$ specifies a time-independent observable map $\boldsymbol{\Phi}_{\boldsymbol{\eta}}$ and is paired, for design purposes, with an aggregate moment trajectory $\boldsymbol{c}_{\boldsymbol{\eta}}(t) \in \mathbb{R}^R$. At each time, these moments define the fiber

$$\boldsymbol{\Phi}_{\boldsymbol{\eta}} : \mathcal{X} \to \mathbb{R}^R, \qquad \mathcal{F}_{\boldsymbol{\eta}}(\boldsymbol{c}) := \{Q \in \mathcal{P}(\mathcal{X}) : \mathbb{E}_Q[\boldsymbol{\Phi}_{\boldsymbol{\eta}}] = \boldsymbol{c}\}\,. \tag{1}$$

In the ideal population formulation, $\boldsymbol{c}_{\boldsymbol{\eta}}(t) = \mathbb{E}_{P_t}[\boldsymbol{\Phi}_{\boldsymbol{\eta}}(X_t)]$. In applications, $\boldsymbol{c}_{\boldsymbol{\eta}}(t)$ denotes the candidate-specific aggregate trajectory supplied to FIDE; Section 3 specifies how these inputs are obtained retrospectively or prospectively. The observable coordinates may represent sensor responses, projections, localized windows, or coefficients of an observable subspace. Because finitely many expectations generally do not identify a full law, changing the design changes both what is constrained and which law directions remain unresolved. Here $\mathcal{P}(\mathcal{X})$ denotes probability laws for which the displayed expectations exist.

Before design optimization, we fix a reference law path $\widetilde{Q}_t$ with density $\widetilde{q}_t$ and velocity u_t satisfying

$$\partial_t \widetilde{q}_t + \nabla \cdot (\widetilde{q}_t u_t) = 0. \tag{2}$$

The equation is understood weakly under the boundary or decay conditions of Section 4. The pair $(\widetilde{Q}_t, u_t)$ is held fixed as $\boldsymbol{\eta}$ varies and is not assumed to reproduce the unknown intermediate population dynamics. Full is defined relative to this entire law–velocity pair, since different velocity fields realizing the same reference marginals can induce different dynamical revision costs.

Scientific adequacy is specified independently by an application-dependent design-stage risk $\mathcal{R} : \mathcal{D} \to \mathbb{R}$, fixed before design optimization. This risk determines which designs are adequate; Full is used only to rank the design freedom that remains after adequacy has been enforced. Evaluating $\mathcal{R}(\boldsymbol{\eta})$ may require additional candidate-specific summaries beyond $\boldsymbol{c}_{\boldsymbol{\eta}}(t)$, such as response covariances or predicted scientific quantities, which are treated as design-stage inputs. Any departure that should affect the design choice must therefore enter the scientific risk or an additional objective. Transfer of scientific adequacy to post-freeze observations is evaluated separately.

3 METHODOLOGY

FIDE takes two conceptually distinct inputs. First, each candidate design $\boldsymbol{\eta}$ is paired with an aggregate trajectory $\boldsymbol{c}_{\boldsymbol{\eta}}(t)$ and the summaries needed to evaluate its scientific risk $\mathcal{R}(\boldsymbol{\eta})$. Retrospectively these may come from existing data or simulation, while prospectively they are supplied by a predictive interface fixed before optimization. Second, the frozen pair $(\widetilde{Q}_t, u_t)$ supplies the common dynamical baseline against which every candidate is evaluated. The predictive interface therefore answers *what would this measurement system constrain?*, whereas the frozen reference answers *how much dynamical revision would those constraints require?* Given these inputs, FIDE applies the same three steps to every candidate: complete its aggregate constraints, measure the resulting dynamical revision, and minimize that revision among scientifically adequate designs.

Step 1: complete the moment constraints. For each $(t, \boldsymbol{\eta})$ admitting the projection, define

$$Q_t^{\boldsymbol{\eta}} \in \operatorname*{arg\,min}_{Q \ll \widetilde{Q}_t} \left\{ D_{\mathrm{KL}}(Q \| \widetilde{Q}_t) : \mathbb{E}_Q[\boldsymbol{\Phi}_{\boldsymbol{\eta}}] = \boldsymbol{c}_{\boldsymbol{\eta}}(t) \right\}. \tag{3}$$

The aggregate trajectory fixes the measured expectations, while information projection fills in the remaining law directions relative to the common reference. Thus $\widetilde{Q}_t$ is fixed across designs and $Q_t^{\boldsymbol{\eta}}$ is the completed law associated with design $\boldsymbol{\eta}$. The ordering of these operations is deliberate: FIDE first completes every candidate by the same rule and only then asks how dynamically costly that completed path is. This keeps the completion rule common across designs rather than allowing the dynamical score itself to choose unresolved law directions. A path-level minimum-action completion defines the corresponding coupled alternative; Appendix D.2.1 quantifies the difference on a controlled example. Theorem 1 gives the corresponding existence, uniqueness, differentiability, and exponential-tilt results.

Step 2: measure dynamical revision. Let $q_t^{\boldsymbol{\eta}}$ denote the projected density. The information projection fixes a law path, but the frozen reference velocity u_t need not transport that path. FIDE therefore asks for the smallest velocity correction δ_t such that $u_t + \delta_t$ does. We define

$$a_{\mathrm{full}}(t, \boldsymbol{\eta}) := \inf_{\delta \in L^2(Q_t^{\boldsymbol{\eta}}; \mathbb{R}^d)} \left\{ \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\|\delta\|^2] : \partial_t q_t^{\boldsymbol{\eta}} + \nabla \cdot \left(q_t^{\boldsymbol{\eta}} (u_t + \delta) \right) = 0 \right\}, \tag{4}$$

$$\mathcal{A}(\boldsymbol{\eta}) := \int_0^1 a_{\mathrm{full}}(t, \boldsymbol{\eta})\, \rho(dt), \tag{5}$$

where ρ is a fixed probability measure on $[0, 1]$, the continuity constraint is understood weakly under the prescribed boundary conditions, and $\inf \varnothing = +\infty$. We call $a_{\mathrm{full}}(t, \boldsymbol{\eta})$ the instantaneous Full action density and $\mathcal{A}(\boldsymbol{\eta})$ the time-integrated Full action. Only the correction δ_t is optimized; the reference and completed path remain fixed. For later analysis, whenever the required density derivatives exist on the positive support of $q_t^{\boldsymbol{\eta}}$, define the normalized continuity forcing $h_{t,\boldsymbol{\eta}} := [\partial_t q_t^{\boldsymbol{\eta}} + \nabla \cdot (q_t^{\boldsymbol{\eta}} u_t)]/q_t^{\boldsymbol{\eta}}$. The continuity constraint is then equivalent to $\nabla \cdot (q_t^{\boldsymbol{\eta}} \delta_t) = -q_t^{\boldsymbol{\eta}} h_{t,\boldsymbol{\eta}}$. Theorem 2 gives the corresponding weak and weighted-Poisson characterizations.

Step 3: select among scientifically adequate designs. The observable family $\mathcal{D}$, scientific risk $\mathcal{R}$, tolerance ε, frozen reference $(\widetilde{Q}_t, u_t)$, information-projection rule, and Full correction geometry are fixed before optimizing $\boldsymbol{\eta}$. Let $\mathcal{R}^\star := \inf_{\boldsymbol{\eta} \in \mathcal{D}} \mathcal{R}(\boldsymbol{\eta})$ and assume $0 < \mathcal{R}^\star < \infty$. For $\varepsilon \geq 0$, define

$$\mathcal{D}_\varepsilon := \{ \boldsymbol{\eta} \in \mathcal{D} : \mathcal{R}(\boldsymbol{\eta}) \leq (1 + \varepsilon) \mathcal{R}^\star \}, \qquad \boldsymbol{\eta}_{\mathrm{Full}} \in \operatorname*{arg\,min}_{\boldsymbol{\eta} \in \mathcal{D}_\varepsilon} \mathcal{A}(\boldsymbol{\eta}). \tag{6}$$

Scientific risk therefore determines feasibility, and Full ranks only admissible designs. This gives FIDE a hierarchical structure: scientific adequacy is enforced first through the ε-constraint, and Full then resolves the remaining design freedom. The following elementary property makes the reference-relative meaning of this secondary criterion explicit.

Proposition 1 (Reference consistency and scope)**.** *Fix a design $\boldsymbol{\eta}$ for which the information projection is well defined. If $\boldsymbol{c}_{\boldsymbol{\eta}}(t) = \mathbb{E}_{\widetilde{Q}_t}[\boldsymbol{\Phi}_{\boldsymbol{\eta}}]$ for every t, then $Q_t^{\boldsymbol{\eta}} = \widetilde{Q}_t$ and $a_{\mathrm{full}}(t, \boldsymbol{\eta}) = 0$ for every t, hence $\mathcal{A}(\boldsymbol{\eta}) = 0$. If an alternative population law $P_t \neq \widetilde{Q}_t$ also satisfies $\mathbb{E}_{P_t}[\boldsymbol{\Phi}_{\boldsymbol{\eta}}] = \boldsymbol{c}_{\boldsymbol{\eta}}(t)$ for every t, this zero action does not imply $Q_t^{\boldsymbol{\eta}} = P_t$. Thus zero Full certifies consistency with the declared reference under the supplied constraints; it does not identify among alternative laws that share those constraints.*

The first claim follows because the reference itself satisfies the moment constraints and attains zero relative entropy, so its frozen velocity realizes the projected path with zero correction. Intuitively, the candidate measurements then ask for nothing that the reference does not already predict, so neither the completed law nor its dynamics need to change. Existence qualifications for the outer minimizer and the finite-sample Law baseline are collected in Appendix B.3.

Comparator criteria. For empirical comparison, we also select designs using Tangent, residual L^2, forward KL, and static W_2 as secondary criteria, always on the same information-projected law path and within the same admissible set $\mathcal{D}_\varepsilon$. These criteria probe different properties of the completed path and their definitions are collected in Appendix C.1, while Section 4 summarizes how they differ from Full.

Numerical implementation and design gradients. In finite-sample experiments, we use *Law* to denote selection by scientific risk alone. We replace the population optimum $\mathcal{R}^\star$ by the selection risk R_{Law} of the numerically optimized Law design and use the operational admissible set $R_{\mathrm{sel}}(\boldsymbol{\eta}) \leq (1+\varepsilon)R_{\mathrm{Law}}$. This is a finite-search approximation rather than a certificate of global risk optimality. For each candidate, the numerical pipeline calibrates the empirical information projection, evaluates scientific risk and the chosen secondary criterion, and optimizes over the resulting admissible designs. All candidate-specific design-stage inputs are fixed before optimization; post-freeze target observations are used only for validation and are never fed back into candidate scoring. For an empirical reference $\sum_{i=1}^N \bar{w}_i \delta_{x_i}$ with $z_i = \boldsymbol{\Phi}_{\boldsymbol{\eta}}(x_i)$, exact calibration requires $\widehat{c} \in \mathrm{conv}\{z_1, \ldots, z_N\}$. We distinguish this mathematical *calibration validity* from *support certification*: a feasible calibrated law may still be poorly supported for estimating Full action. We therefore monitor effective sample size and covariance conditioning separately. Proposition 3 gives the precise feasibility result, while benchmark-specific discretizations and numerical audits are deferred to the corresponding experimental appendices.

FIDE allows both discrete and continuous designs. For continuous designs, the fixed finite-dimensional objective is locally differentiable end to end whenever the calibrated covariance and reduced Full systems remain nonsingular and vary smoothly with the design. In particular, if the reduced Full solve is $\mathsf{K}_{\boldsymbol{\eta}}\beta_{\boldsymbol{\eta}} = b_{\boldsymbol{\eta}}$, then for any design perturbation ξ the instantaneous action satisfies $D_{\boldsymbol{\eta}}\widehat{a}_{\mathrm{full}}[\xi] = 2\beta_{\boldsymbol{\eta}}^\top D_{\boldsymbol{\eta}} b_{\boldsymbol{\eta}}[\xi] - \beta_{\boldsymbol{\eta}}^\top D_{\boldsymbol{\eta}}\mathsf{K}_{\boldsymbol{\eta}}[\xi]\beta_{\boldsymbol{\eta}}$. Thus gradients propagate through the information-projection calibration, physical-time multiplier dynamics, and Full correction solve without differentiating through the iterations of the underlying solvers. Under these local conditions, continuous sensor locations or other differentiable design parameters can therefore be optimized directly with gradient-based methods. Proposition 9 gives the full conditions and integrated-action result.

4 Theoretical Analysis

This section develops the properties of Full needed for the residual-design interpretation. We first derive the information-projected path and its reference-relative continuity forcing. We then show that the minimum quadratic correction used in Section 3 is exactly a sharp Sobolev expectation-rate discrepancy, identify Tangent as its restriction to the measured observable span, and finally relate Full to the remaining comparators and to matched-initial stability.

Information-projected path and continuity forcing For a fixed design, write $\phi = \mathbf{\Phi}_{\boldsymbol{\eta}}$, $Q_t = Q_t^{\boldsymbol{\eta}}$, $q_t = q_t^{\boldsymbol{\eta}}$, $c_t = \boldsymbol{c}_{\boldsymbol{\eta}}(t)$, $\lambda_t = \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}$, and $h_t = h_{t,\boldsymbol{\eta}}$, and define $A_t(\lambda) = \log \mathbb{E}_{\widetilde{Q}_t}[e^{\lambda^\top \phi}]$, $m_t = J_\phi u_t$, and $C_t = \mathrm{Cov}_{Q_t}(\phi, \phi)$.

Theorem 1 (Information-projected lift and forcing)**.** *Assume that c_t lies in the relative interior of the feasible moment set, that the dual calibration objective $A_t(\lambda) - \lambda^\top c_t$ attains a minimizer λ_t and that A_t is finite in a neighborhood of λ_t. Under the conditions of Appendices B.4–B.6, the information projection exists, is unique at the law level, and, in a minimal observable representation with $C_t \succ 0$, satisfies*

$$q_t(x) = \widetilde{q}_t(x) \exp\big(\lambda_t^\top \phi(x) - A_t(\lambda_t)\big), \qquad \mathbb{E}_{Q_t}[\phi] = c_t, \tag{7}$$

$$C_t \dot{\lambda}_t = \dot{c}_t - \mathbb{E}_{Q_t}[m_t] - \mathrm{Cov}_{Q_t}(\phi, \lambda_t^\top m_t), \tag{8}$$

$$h_t = \dot{\lambda}_t^\top (\phi - c_t) + \lambda_t^\top (m_t - \mathbb{E}_{Q_t}[m_t]), \qquad \mathbb{E}_{Q_t}[h_t] = 0. \tag{9}$$

Thus $\partial_t q_t + \nabla \cdot (q_t u_t) = q_t h_t$, and the forcing can be evaluated from tilted expectations and an $R \times R$ covariance solve without differentiating the reference density in physical time.

Reference-relative expectation-rate robustness For a sufficiently regular time-independent observable f, define

$$\mathcal{E}_t(f) := \frac{d}{dt}\mathbb{E}_{Q_t}[f] - \mathbb{E}_{Q_t}[\nabla f \cdot u_t] = \mathbb{E}_{Q_t}[h_t f]. \tag{10}$$

Because $\mathbb{E}_{Q_t}[h_t] = 0$, this functional is invariant to additive constants. Let H_{q_t} denote the mean-zero weighted Sobolev energy space defined in Appendix B.1.

Theorem 2 (Sharp expectation-rate characterization of instantaneous Full action)**.** *Assume the weighted Dirichlet form is coercive on H_{q_t} and that $f \mapsto \mathbb{E}_{Q_t}[h_t f]$ defines a bounded linear functional on H_{q_t}, under the boundary conventions of Appendix B.1. Then*

$$a_{\mathrm{full}}(t, \boldsymbol{\eta}) = \sup_{\substack{f \in H_{q_t} \\ \mathbb{E}_{Q_t}[\|\nabla f\|^2] \le 1}} |\mathcal{E}_t(f)|^2, \tag{11}$$

$$\nabla \cdot (q_t \nabla \psi_t) = q_t h_t, \qquad \mathbb{E}_{Q_t}[\psi_t] = 0, \qquad \delta_t^\star = -\nabla \psi_t, \tag{12}$$

$$a_{\mathrm{full}}(t, \boldsymbol{\eta}) = \mathbb{E}_{Q_t}[\|\nabla \psi_t\|^2] = \|h_t\|^2_{H^{-1}(q_t)}. \tag{13}$$

Here $\delta_t^\star$ is the unique minimum-energy correction in equation 4.

In this sense, Full asks not how far the completed law is from the reference at one time, but how difficult it is for the reference dynamics to realize its evolution. The H^{-1} representation gives a useful transport interpretation of this cost. Full does not penalize continuity mismatch only by its pointwise amplitude; it weights that mismatch by the spatial effort required to correct it. Rapidly varying local discrepancies can often be repaired by moving mass over short distances and therefore have relatively small Full cost, whereas coherent low-frequency discrepancies requiring transport over longer distances are more expensive. Equivalently, in flux variables the energy weights a required correction flux by $1/q_t$, making transport of a fixed flux through poorly populated regions costly. This transport-scale sensitivity is precisely what distinguishes Full from a plain L^2 norm of the same continuity forcing; Appendix C.4 gives exact Fourier examples.

The three equivalent views of Full play complementary roles. The minimum-correction form gives it a physical interpretation as transport effort, the H^{-1} form exposes the spatial geometry of that effort, and the dual form gives the interpretation most directly relevant to experimental design. In particular, $\sqrt{a_{\mathrm{full}}(t, \boldsymbol{\eta})}$ is the smallest uniform constant controlling reference-relative expectation-rate discrepancy per unit Sobolev gradient energy. Consequently, once scientific adequacy has been enforced, minimizing $\mathcal{A}(\boldsymbol{\eta})$ minimizes the time-integrated squared worst-case discrepancy over this smooth observable class rather than only over the quantities explicitly protected by the primary risk. Equivalent sharp-constant and operator formulations are given in Appendices B.8 and C.2 (Ambrosio & Serfaty, 2008; Mroueh, 2018; Neklyudov et al., 2023; Chemseddine et al., 2025).

Full transports the entire completed law, whereas the measured observables constrain only some of its expectation rates. This suggests separating the required correction into the part already forced by those measured rates and the additional law-level correction they leave unresolved.

Proposition 2 (Full–Tangent–Hidden decomposition)**.** *Assume the conditions of Theorem 2, $J_\phi \in L^2(Q_t)$, $J_\phi u_t \in L^1(Q_t)$, and $\phi_j - c_{t,j} \in H_{q_t}$ for each observable coordinate, with moment differentiation and weak testing justified. Define*

$$r_t = \mathbb{E}_{Q_t}[J_\phi u_t] - \dot{c}_t, \qquad G_t = \mathbb{E}_{Q_t}[J_\phi J_\phi^\top]. \tag{14}$$

The Tangent correction is $\delta_t^{\tan} = -J_\phi^\top G_t^\dagger r_t$, with $a_{\tan}(t,\boldsymbol{\eta}) = r_t^\top G_t^\dagger r_t$. Writing $\delta_t^{\mathrm{hid}} = \delta_t^\star - \delta_t^{\tan}$ and $a_{\mathrm{hid}}(t,\boldsymbol{\eta}) = \mathbb{E}_{Q_t}[\|\delta_t^{\mathrm{hid}}\|^2]$, the Tangent and Hidden components are orthogonal and

$$a_{\mathrm{full}}(t,\boldsymbol{\eta}) = a_{\tan}(t,\boldsymbol{\eta}) + a_{\mathrm{hid}}(t,\boldsymbol{\eta}). \tag{15}$$

Integrating against ρ gives the corresponding design-level decomposition equation 150, and hence $\mathcal{A}_{\tan}(\boldsymbol{\eta}) \le \mathcal{A}(\boldsymbol{\eta})$.

Tangent is therefore the restriction of the Full dual problem to the finite-dimensional span of the centered measured observables, equivalently a Galerkin restriction in the weighted Sobolev energy space, while Hidden is the corresponding Galerkin energy gap (Mroueh, 2018). For experimental design, this means that two sensor systems can be equally adequate for the scientific task while inducing very different dynamical corrections elsewhere in the completed law. Full distinguishes precisely this residual difference. The exact Tangent–Full equality condition and examples showing that Tangent does not generally control Full are given in Appendices B.11 and C.3.

Criterion distinctions and matched-initial stability. The remaining secondary criteria probe different properties of the same information-projected law path. Residual L^2 measures the amplitude of the complete normalized continuity forcing, whereas Full weights that forcing through inverse transport geometry. Forward KL and static W_2 instead measure marginal displacement from the frozen reference and do not encode the temporal realization cost of the prescribed path. Appendix C establishes these distinctions through operator identities, exact information-projected separations, and ranking reversals.

Full also yields a positive stability statement in a stronger regular-flow regime. Let $\mathsf{T}_{s,t}^u$ denote the characteristic flow of u_t and define the matched-initial reference rollout $R_t^{\boldsymbol{\eta}} := (\mathsf{T}_{0,t}^u)_\# Q_0^{\boldsymbol{\eta}}$. Assuming finite second moments and Full action, classical characteristic flows for u_t and $u_t + \delta_t^\star$, and an integrable spatial Lipschitz bound $\|u_t(x) - u_t(y)\| \le L(t)\|x - y\|$, Proposition 14 gives, for $\rho(dt) = dt$ and $\Lambda = \int_0^1 L(t)\,dt$,

$$W_2(Q_t^{\boldsymbol{\eta}}, R_t^{\boldsymbol{\eta}}) \le \int_0^t \exp\!\left(\int_s^t L(\tau)\,d\tau\right) \sqrt{a_{\mathrm{full}}(s,\boldsymbol{\eta})}\,ds \le e^\Lambda \sqrt{\mathcal{A}(\boldsymbol{\eta})}, \qquad t \in [0,1]. \tag{16}$$

Hence every globally 1-Lipschitz observable f satisfies

$$\left|\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[f] - \mathbb{E}_{R_t^{\boldsymbol{\eta}}}[f]\right| \le e^\Lambda \sqrt{\mathcal{A}(\boldsymbol{\eta})}.$$

When $Q_0^{\boldsymbol{\eta}} = \widetilde{Q}_0$, one has $R_t^{\boldsymbol{\eta}} = \widetilde{Q}_t$ and $\mathcal{A}_{W_2}(\boldsymbol{\eta}) \le e^{2\Lambda}\mathcal{A}(\boldsymbol{\eta})$. More generally, the preceding bound gives a downstream-agnostic stability interpretation: after initial-law mismatch is removed, small Full action controls the departure of every later-chosen globally 1-Lipschitz observable from the corresponding frozen-reference rollout, even if that observable was not part of the scientific risk used to select the design. This is a one-sided reference-relative guarantee, not a guarantee of recovery of the unknown intermediate population law. Appendix C.6 shows that no reverse bound holds in general. Proofs, technical conditions, and additional analytical results are collected in Appendices B and C.

5 Experiments

The experiments follow a progression from mechanism to deployment. Gaussian-mixture transport first asks whether the choice of secondary criterion actually changes the measurement design once scientific adequacy is fixed. A controlled downstream experiment then asks when the baseline-preserving preference improves prediction of a quantity that was never used for selection and when that benefit fails. Prospective double-gyre transport removes access to intermediate microscopic

Selector	Selection risk increase (%)	Held-out Full (cross-eval.)	Truth-relative hidden-law SW_2^2
Law	0	29.4178	0.096460
Tangent	0.491	37.3317	0.094144
Residual L^2	0.441	36.0548	0.094173
Forward KL	0.303	33.0647	0.094352
Static W_2	0.294	32.9924	0.094357
Full	0.604	19.5359	0.102205

Table 1: Criterion-specific Gaussian designs at the 1% allowance, selected under the same scientific-risk restriction and cross-evaluated using held-out Full action and truth-relative hidden-law squared sliced-Wasserstein distance. The two diagnostics induce different orderings; values are comparable within columns but not across columns.

target states during design, and interacting skyrmions finally test the same construction on a higher-dimensional joint law. Across these studies we evaluate three distinct quantities when available: held-out task risk tests whether the scientific requirement transfers, held-out Full tests whether the selected minimum-revision preference persists on fresh data, and truth-relative downstream diagnostics test consequences outside the selection objective. Section 5.1 collects supporting analyses of the modeling assumptions and numerical approximations behind these results. The source code is available on GitHub.[1]

(1) Secondary criteria induce different designs: Gaussian-mixture transport The Gaussian benchmark isolates the effect of the secondary criterion in a transparent two-dimensional setting: the endpoint laws are fixed, the orientation of the intermediate mixture varies, and two localized Gaussian sensors are placed on a ring. Law, Tangent, residual L^2, forward KL, static W_2, and Full select different sensor geometries despite closely matched scientific risk. Table 1 cross-evaluates these designs using reference-relative Full action and truth-relative hidden-law discrepancy. The orderings disagree: the Full-selected geometry has the smallest held-out Full action but the largest hidden-law discrepancy. Thus the secondary criteria resolve the same scientifically admissible design freedom differently rather than serving as interchangeable approximations to a common objective. This disagreement motivates the next question: if minimum revision is not a generic truth-recovery principle, under what conditions does preserving unresolved baseline structure nevertheless improve a quantity that was never used to select the design? Complete criterion comparisons and allowance sweeps are reported in Appendices D.2.4 and E.2.

(2) When does baseline preservation help and when does it hurt? We next ask when the baseline-preserving preference has a measurable consequence for a quantity never used in design. We construct a product-law problem with one task-relevant coordinate and one task-unprotected coordinate. Each candidate sensor contains the same task-relevant signal but mixes it with a different component of the unprotected coordinate, so the primary scientific risk is exactly tied across five sensor geometries. The frozen baseline is initially correct in the unprotected coordinate. As an independent downstream test, we measure truth-relative local-occupancy RMSE on 16 smooth observables of that coordinate, none of which enters either the scientific risk or the secondary design criteria. Because Law is indifferent among all five risk-optimal geometries, we report its expected downstream error under uniform tie-breaking. When the baseline is correct, Full and static W_2 select the same geometry, reducing expected occupancy error by about half relative to Law, whereas the geometry selected by Tangent, residual L^2, and forward KL performs worse.

We then freeze these selections and the reference and introduce an increasingly strong true departure in the unprotected coordinate. Figure 1 shows the resulting transition: the advantage of the Full/W_2 geometry steadily erodes and eventually reverses relative to Law. This isolates when the baseline-preserving inductive bias is useful: preserving unresolved baseline structure improves prediction of downstream quantities withheld from design when that structure is reliable, but becomes harmful once the baseline is sufficiently wrong. In this construction static W_2 makes the same selection, so the result identifies the value of baseline preservation rather than an advantage unique to Full. This controlled setting lets us vary baseline reliability while retaining access to ground truth; we

[1] https://github.com/REPOSITORY-TO-BE-RELEASED

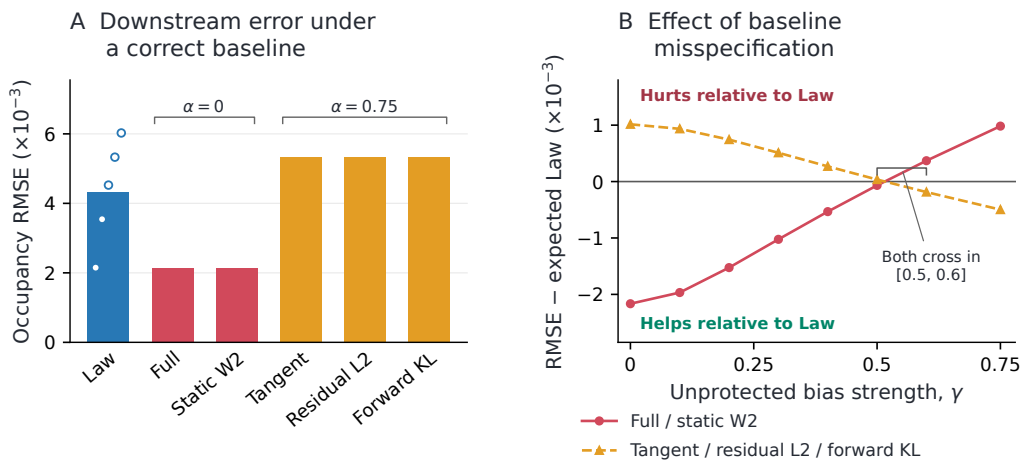

Figure 1: **Downstream consequences of baseline-preserving selection. A:** At signal amplitude $\epsilon = 0.10$ and occupancy localization scale $\kappa = 4$, we evaluate truth-relative local-occupancy RMSE on 16 smooth observables that are withheld from both the scientific risk and all secondary design criteria. Scientific risk is identical across the five candidate geometries. Full and static W_2 select $\alpha = 0$, while Tangent, residual L^2, and forward KL select $\alpha = 0.75$. The Law bar is the expected RMSE under uniform tie-breaking over its five risk-minimizing geometries; open circles show the corresponding geometry-specific errors. **B:** The selected geometries are then frozen while a true departure of strength γ is introduced in the task-unprotected coordinate. The ordinate is occupancy RMSE minus the expected Law RMSE, so negative values indicate improvement over Law and positive values indicate harm. The Full/static-W_2 geometry changes from lower to higher error than Law as misspecification grows, while the Tangent/residual-L^2/forward-KL geometry reverses in the opposite direction; both crossings occur between $\gamma = 0.5$ and 0.6 in this representative setting.

next remove that luxury and ask whether the residual design choice can be made prospectively before intermediate microscopic target states are available. Appendix D.2.2 gives the complete construction, robustness checks, and full misspecification sweep.

(3) Prospective selection without intermediate target states: double-gyre transport The double-gyre provides a nonautonomous transport benchmark with stretching, exchange, and nontrivial passive-tracer motion (Shadden et al., 2005). We choose four localized sensors for a heterogeneous tracer population while withholding all intermediate microscopic target states from the optimizer. Selection uses only endpoint ensembles, a frozen endpoint-trained reference, and a frozen aggregate predictive interface constructed from a separate simulation bank, which supplies the candidate moment trajectories and summaries needed to evaluate scientific risk. After the sensor geometries are frozen, we introduce a separately seeded validation reference and fresh microscopic validation trials. On these held-out data, the 2%–5% Full-selected designs satisfy their corresponding scientific-risk caps, while the 1% point remains unresolved; across the tested frontier, the Full-selected geometries also retain lower Full action than Law. Thus both scientific adequacy, where statistically resolved, and the minimum-revision ordering transfer after design freeze without access to intermediate target states during selection. Complete results and uncertainty analyses are in Appendix E.3, and Figure 2 summarizes the resulting risk–action tradeoff.

(4) Configuration-space design in 32 dimensions: interacting skyrmions The prospective study changes the information available at design time; our final experiment instead changes the computational regime. Motivated by interacting magnetic skyrmions (Nagaosa & Tokura, 2013; Reichhardt et al., 2015), we use a reduced stochastic model of 16 interacting planar particles, giving a 32-dimensional microscopic state. Four aggregate sensors observe averages over 512 configurations, while scientific risk is defined through nine permutation-invariant features of pair structure and collective order. Full is approximated directly in configuration space using a fixed $K = 280$ Galerkin

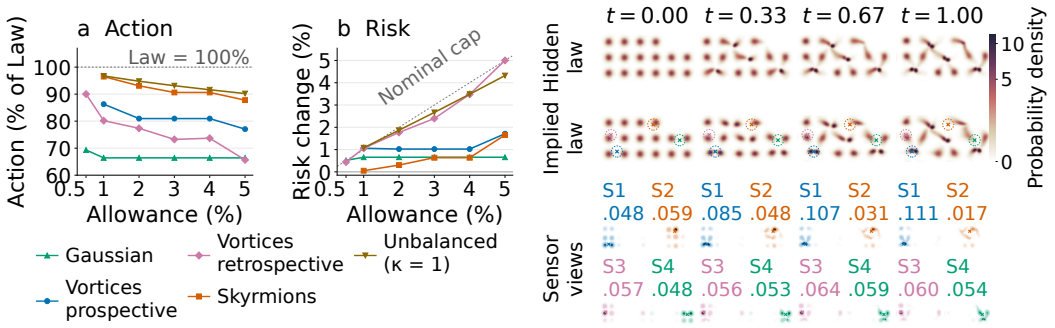

Figure 2: **Empirical overview.** *Left:* Held-out scientific risk and Full action across the tested risk allowances in the prospective double-gyre study. *Right:* Interacting-skyrmion setup, where aggregate sensor views constrain a 32-dimensional configuration law without determining it.

potential dictionary, so the reported values assess the finite-dimensional construction rather than continuum convergence. On fresh paired validation trials, the Full-selected designs satisfy their scientific-risk caps, differ from the corresponding Law selections, and retain lower fitted action. The experiment therefore shows that the same residual-design construction remains operational on a 32-dimensional joint law; complete results and numerical audits are in Appendix E.4.

5.1 Supporting analyses

Supporting studies examine the assumptions and numerical approximations behind the main experiments. Baseline-correctness and reference-perturbation analyses characterize when the minimum-revision preference helps or fails under misspecification (Appendices D.2.2 and D.2.5). Allowing an unresolved mode to adapt dynamically reduces Full action by roughly 10–26%, quantifying the effect of fixing completion before scoring, while retrospective double-gyre and prospective predictor perturbations assess repeatability and sensitivity to the aggregate interface (Appendices D.2.1, D.1.1, and D.2.8). Numerical checks further test raster, smoothing, and solver choices, manufactured-problem convergence, richer skyrmion potential spaces, and leave-one-out stability (Appendices D.2.6, D.2.7, and D.2.9). Together, these analyses distinguish sensitivity to the modeling assumptions from sensitivity to the numerical implementation.

6 Conclusion

FIDE separates two questions in experimental design: what must a measurement system accomplish for the scientific task, and how should the remaining law-level freedom be resolved once that task is satisfied? Scientific risk answers the first; FIDE addresses the second by completing each design's aggregate constraints relative to a common frozen reference and measuring the dynamical revision required to realize the resulting path. Full action makes this revision precise through equivalent minimum-correction and expectation-rate interpretations, while the Full–Tangent–Hidden decomposition identifies the part of the required correction that is invisible to the measured moment rates.

The experiments show that this residual choice matters. Different secondary criteria select different scientifically admissible measurements; preserving unresolved baseline structure can improve prediction of downstream quantities withheld from design when that structure is reliable and can hurt when it is sufficiently misspecified; prospective selection can transfer after design freeze without intermediate microscopic target states; and the finite-dimensional construction can be instantiated on a 32-dimensional joint law. Minimum Full is therefore most naturally viewed as a design principle for settings in which preserving a declared baseline outside task-valued departures is itself part of the experimental specification. Extending FIDE to propagate uncertainty over that baseline is a natural next step. A complementary direction is scalable optimization with discrete or combinatorial experimental decisions, where smooth surrogate objectives or continuous relaxations could retain gradient-based optimization while ultimately selecting feasible discrete designs.

AI USE STATEMENT

Generative AI tools were used to assist with code development and debugging and with language-level polishing of the manuscript. The authors reviewed and validated the resulting code and manuscript text, and take full responsibility for the final submission.

REPRODUCIBILITY STATEMENT

The code required to reproduce all experiments will be made publicly available on GitHub upon publication. Detailed experimental protocols, information boundaries, numerical procedures, implementation details, hardware and software specifications, and benchmark-specific settings are provided in Appendix E. The appendix also reports numerical audits, held-out validation procedures, and sensitivity analyses used to assess reproducibility and stability of the reported results.

REFERENCES

Michael S. Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=li7qeBbCR1t.

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025. URL https://www.jmlr.org/papers/v26/23-1605.html.

Alen Alexanderian, Noemi Petra, Georg Stadler, and Omar Ghattas. A-optimal design of experiments for infinite-dimensional bayesian linear inverse problems with regularized \ell_0-sparsification. *SIAM Journal on Scientific Computing*, 36(5):A2122–A2148, 2014.

Luigi Ambrosio and Sylvia Serfaty. A gradient flow approach to an evolution problem arising in superconductivity. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 61(11):1495–1539, 2008.

Anthony C Atkinson and VV370955 Fedorov. The design of experiments for discriminating between two rival models. *Biometrika*, 62(1):57–70, 1975.

Ahmed Attia, Alen Alexanderian, and Arvind K. Saibaba. Goal-oriented optimal design of experiments for large-scale bayesian linear inverse problems. *Inverse Problems*, 34(9):095009, 2018. doi: 10.1088/1361-6420/aad210.

Ahmed Attia, Sven Leyffer, and Todd S Munson. Robust a-optimal experimental design for sensor placement in bayesian linear inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, 13(2):744–774, 2025.

Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge–Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000. doi: 10.1007/s002110050002.

Raktim Bhattacharya. Sensor design for accuracy-bounded estimation via maximum-entropy likelihood synthesis. *arXiv preprint arXiv:2605.11120*, 2026.

George EP Box and WILLIAM J Hill. Discrimination among mechanistic models. *Technometrics*, 9(1):57–71, 1967.

Kristian Bredies and Silvio Fanzon. An optimal transport approach for solving dynamic inverse problems in spaces of measures. *ESAIM: Mathematical Modelling and Numerical Analysis*, 54 (6):2351–2382, 2020.

Damiano Brigo and Giovanni Pistone. Optimal approximations of the Fokker–Planck–Kolmogorov equation: projection, maximum likelihood eigenfunctions and Galerkin methods, 2017. URL https://arxiv.org/abs/1603.04348v2. Version 2; first posted in 2016.

Damiano Brigo, Bernard Hanzon, and François Le Gland. Approximate nonlinear filtering by projection on exponential manifolds of densities. *Bernoulli*, 5(3):495–534, 1999. doi: 10.2307/3318714.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995. doi: 10.1214/ss/1177009939.

Jannis Chemseddine, Christian Wald, Richard Duong, and Gabriele Steidl. Neural sampling from Boltzmann densities: Fisher–Rao curves in the Wasserstein geometry. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/1d1cea122b9ec9f78acc21510659e500-Abstract-Conference.html.

Yifan Chen and Wuchen Li. Optimal transport natural gradient for statistical manifolds with continuous sample space. *Information Geometry*, 3(1):1–32, 2020.

Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution, part i. *IEEE Transactions on Automatic Control*, 61(5):1158–1169, 2016. doi: 10.1109/TAC.2015.2457784.

Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. Optimal transport over a linear dynamical system. *IEEE Transactions on Automatic Control*, 62(5):2137–2152, 2017. doi: 10.1109/TAC.2016.2602103.

Austin H. Cheng, Alston Lo, Kin Long Kelvin Lee, Santiago Miret, and Alán Aspuru-Guzik. Stiefel flow matching for moment-constrained structure elucidation. In *International Conference on Learning Representations*, 2025.

Imre Csiszár. I-Divergence Geometry of Probability Distributions and Minimization Problems. *The Annals of Probability*, 3(1):146–158, 1975. doi: 10.1214/aop/1176996454.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/940392f5f32a7ade1cc201767cf83e31-Abstract.html.

Purushottam D Dixit, Jason Wagoner, Corey Weistuch, Steve Pressé, Kingshuk Ghosh, and Ken A Dill. Perspective: Maximum caliber is a general variational principle for dynamical systems. *The Journal of chemical physics*, 148(1), 2018.

Adam Foster, Martin Jankowiak, Eli Bingham, Paul Horsfall, Yee Whye Teh, Tom Rainforth, and Noah Goodman. Variational bayesian optimal experimental design. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/d55cbf210f175f4a37916eafe6c04f0d-Abstract.html.

Adam Foster, Desi R Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential bayesian experimental design. In *International conference on machine learning*, pp. 3384–3395. PMLR, 2021.

Samar Hadou, Yigit Berkay Uslu, and Alejandro Ribeiro. Constrained diffusion models with primal-dual inference. *arXiv preprint arXiv:2606.17192*, 2026.

Dion Häfner and Alexander Lavin. Tesseract core: Universal, autodiff-native software components for simulation intelligence. *Journal of Open Source Software*, 10(111):8385, 2025.

Yacov Haimes. On a bicriterion formulation of the problems of integrated system identification and system optimization. *IEEE transactions on systems, man, and cybernetics*, (3):296–297, 1971.

Markus Hainy, David J Price, Olivier Restif, and Christopher Drovandi. Optimal bayesian design for model discrimination via classification. *Statistics and Computing*, 32(2):25, 2022.

Ross N Hoffman and Robert Atlas. Future observing system simulation experiments. *Bulletin of the American Meteorological Society*, 97(9):1601–1616, 2016.

Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957. doi: 10.1103/PhysRev.106.620.

Ruhui Jin, Martin Guerra, Qin Li, and Stephen J. Wright. Optimal design for linear models via gradient flow. *Communications on Pure and Applied Analysis*, 32:282–311, 2026. doi: 10.3934/cpaa.2026041.

Shervin Khalafi, Dongsheng Ding, and Alejandro Ribeiro. Constrained diffusion models via dual training. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008. URL https://jmlr.org/papers/v9/krause08a.html.

Solomon Kullback and Richard A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951. doi: 10.1214/aoms/1177729694.

Richard S Laugesen, Prashant G Mehta, Sean P Meyn, and Maxim Raginsky. Poisson's equation in nonlinear filtering. *SIAM Journal on Control and Optimization*, 53(1):501–525, 2015.

Etienne Lempereur, Nathanaël Cuvelle-Magar, Florentin Coeurdoux, Stéphane Mallat, and Eric Vanden-Eijnden. MGD: Moment guided diffusion for maximum entropy generation. *arXiv preprint arXiv:2602.17211*, 2026.

Wuchen Li and Jiaxi Zhao. Wasserstein information matrix: W. li, j. zhao. *Information Geometry*, 6(1):203–255, 2023.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Liyao Lyu, Xinyue Yu, and Hayden Schaeffer. Multiscale nudging: From macroscopic observations to microscopic dynamics. *arXiv preprint arXiv:2606.06809*, 2026a.

Liyao Lyu, Xinyue Yu, David Schneidinger, and Hayden Schaeffer. Wasserstein moment nudging for vlasov-poisson data assimilation. *arXiv preprint arXiv:2608.21725*, 2026b.

Yuezhou Ma, Haixu Wu, Hang Zhou, Huikun Weng, Jianmin Wang, and Mingsheng Long. PhySense: Sensor placement optimization for accurate physics sensing. In *Advances in Neural Information Processing Systems*, volume 38, pp. 40146–40174, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/332b4fbe322e11a71fa39d91c664d8fa-Abstract-Conference.html.

Madhusudan Madhavan, Alen Alexanderian, Arvind K Saibaba, Bart van Bloemen Waanders, and Rebekah D White. A control-oriented approach to optimal sensor placement. *SIAM/ASA Journal on Uncertainty Quantification*, 14(2):711–741, 2026.

Krithika Manohar, Bingni W. Brunton, J. Nathan Kutz, and Steven L. Brunton. Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns. *IEEE Control Systems Magazine*, 38(3):63–86, 2018. doi: 10.1109/MCS.2018.2810460.

Aimee Maurais and Youssef Marzouk. Sampling in unit time with kernel fisher-rao flow. *arXiv preprint arXiv:2401.03892*, 2024.

Aimee Maurais, Bamdad Hosseini, and Youssef Marzouk. Learning paths for dynamic measure transport: A control perspective. *arXiv preprint arXiv:2511.03797*, 2025.

Tina M Morrison, Prasanna Hariharan, Chloe M Funkhouser, Payman Afshari, Mark Goodin, and Marc Horner. Assessing computational model credibility using a risk-based framework: application to hemolysis in centrifugal blood pumps. *Asaio Journal*, 65(4):349, 2019.

Baptiste Mourre and Alberto Alvarez. Benefit assessment of glider adaptive sampling in the ligurian sea. *Deep Sea Research Part I: Oceanographic Research Papers*, 68:68–78, 2012.

Olga Movilla Miangolarra, Asmaa Eldesoukey, and Tryphon T Georgiou. Inferring potential landscapes: A schrödinger bridge approach to maximum caliber. *Physical Review Research*, 6(3): 033070, 2024.

Youssef Mroueh. Regularized finite dimensional kernel sobolev discrepancy. *arXiv preprint arXiv:1805.06441*, 2018.

Youssef Mroueh and Mattia Rigotti. Unbalanced sobolev descent. *Advances in Neural Information Processing Systems*, 33:17034–17043, 2020.

Youssef Mroueh, Tom Sercu, and Anant Raj. Sobolev descent. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2976–2985. PMLR, 2019.

Naoto Nagaosa and Yoshinori Tokura. Topological properties and dynamics of magnetic skyrmions. *Nature nanotechnology*, 8(12):899–911, 2013.

Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*, pp. 25858–25889. PMLR, 2023.

Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A computational framework for solving Wasserstein lagrangian flows. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 37461–37485. PMLR, 2024. URL https://proceedings.mlr.press/v235/neklyudov24a.html.

Michele Pavon and Augusto Ferrante. On the geometry of maximum entropy problems. *SIAM Review*, 55(3):415–439, 2013. doi: 10.1137/120862843.

C Reichhardt, D Ray, and CJ Olson Reichhardt. Collective transport properties of driven skyrmions with random disorder. *Physical review letters*, 114(21):217202, 2015.

Shawn C Shadden, Francois Lekien, and Jerrold E Marsden. Definition and properties of lagrangian coherent structures from finite-time lyapunov exponents in two-dimensional aperiodic flows. *Physica D: Nonlinear Phenomena*, 212(3-4):271–304, 2005.

Louis Sharrock. Wasserstein gradient flows for batch Bayesian optimal experimental design, 2026. URL https://arxiv.org/abs/2603.12102. arXiv preprint.

Louis Sharrock and Nikolas Kantas. Joint online parameter estimation and optimal sensor placement for the partially observed stochastic advection-diffusion equation. *SIAM/ASA Journal on Uncertainty Quantification*, 10(1):55–95, 2022. doi: 10.1137/20M1375073.

Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. TrajectoryNet: A dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9526–9536. PMLR, 2020. URL https://proceedings.mlr.press/v119/tong20a.html.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=CD9Snc73AW.

Bruce Turkington. An optimization principle for deriving nonequilibrium statistical models of hamiltonian dynamics. *Journal of Statistical Physics*, 152(3):569–597, 2013.

Umesh Vaidya, Rajeev Rajaram, and Sambarta Dasgupta. Actuator and sensor placement in linear advection PDE with building system application. *Journal of Mathematical Analysis and Applications*, 394(1):213–224, 2012. doi: 10.1016/j.jmaa.2012.03.046.

Marco Viceconti, Francesco Pappalardo, Blanca Rodriguez, Marc Horner, Jeff Bischoff, and Flora Musuamba Tshinanu. In silico trials: Verification, validation and uncertainty quantification of predictive models used in the regulatory evaluation of biomedical products. *Methods*, 185: 120–127, 2021.

Di Wu, Ling Liang, and Haizhao Yang. Beyond expected information gain: Stable bayesian optimal experimental design with integral probability metrics and plug-and-play extensions. *arXiv preprint arXiv:2604.21849*, 2026.

A Related Work

FIDE sits at the intersection of experimental design under partial information, statistical completion of aggregate constraints, and dynamical transport relative to a reference model. The individual ingredients have substantial histories; the distinction here is how they are assembled at the experimental-design level. A candidate measurement system determines a finite set of observable constraints and therefore a design-dependent family of compatible laws. FIDE first completes that ambiguity relative to a common frozen reference, then asks how much the resulting completed dynamics must depart from the reference, but only after the declared scientific objective has determined which designs are acceptable.

Experimental design and sensor placement Bayesian optimal experimental design provides a decision-theoretic framework for selecting measurements by their expected utility (Chaloner & Verdinelli, 1995). Modern developments include variational estimators of expected information gain (Foster et al., 2019) and amortized sequential policies (Foster et al., 2021). Goal-oriented design instead concentrates measurement effort on prescribed quantities of interest (Attia et al., 2018), while control-oriented formulations target uncertainty that matters for subsequent control decisions (Madhavan et al., 2026). Closely related sensor-placement objectives include mutual information (Krause et al., 2008), posterior covariance (Alexanderian et al., 2014), and data-driven reconstruction quality (Manohar et al., 2018). Continuous sensor locations can also be learned from reconstruction feedback (Ma et al., 2025), and robust OED explicitly accounts for misspecification in priors or measurement uncertainty (Attia et al., 2025). In transport settings, placement has been studied through observability and controllability Gramians for linear advection (Vaidya et al., 2012) and jointly with state and parameter estimation for stochastic advection–diffusion systems (Sharrock & Kantas, 2022).

Recent work has also used transport geometry directly in the outer design optimization. Wasserstein gradient flows can optimize probability measures over experimental designs for classical linear-model objectives (Jin et al., 2026) and for batch expected information gain (Sharrock, 2026), while alternative BOED objectives replace KL-based information gain by integral probability metrics such as Wasserstein distance, maximum mean discrepancy, and energy distance (Wu et al., 2026). These approaches differ in the utility they optimize, but they share the usual objective of making the measurement system itself maximally informative, predictive, or useful for a specified downstream task. A particularly relevant contrast is optimal model-discrimination design, where disagreement among candidate models is itself informative and experiments are chosen to distinguish rival mechanisms or reduce uncertainty about model identity (Box & Hill, 1967; Atkinson & Fedorov, 1975; Hainy et al., 2022). FIDE addresses the complementary baseline-preserving regime: the scientific risk first determines an admissible set of comparably adequate designs, and residual disagreement with the frozen reference is not assigned additional scientific utility unless it enters that risk. Only within this declared scope does FIDE use a second criterion to resolve residual law-level freedom by minimum dynamical revision. The resulting ε-constraint ordering is classical in multiobjective optimization (Haimes, 1971), the contribution here is the reference-relative dynamical criterion used to compare the admissible designs.

Risk-informed model credibility A complementary literature treats model credibility as conditional on a specified context of use rather than as a global property of a model. Risk-informed frameworks make the required credibility evidence commensurate with the consequences of relying on a computational model for that use, and organize verification, validation, and uncertainty assessment around the question the model is intended to support (Morrison et al., 2019; Viceconti et al., 2021). This literature motivates the scoped-use setting considered here but does not prescribe FIDE's secondary criterion: baseline preservation is an additional design specification, and FIDE provides a dynamical rule for resolving the residual measurement-design freedom left by the primary task.

Aggregate information and statistical completion A separate literature asks what can be inferred when observations constrain only aggregate statistics rather than microscopic states. Maximum entropy and information projection provide classical ways to complete such partial statistical information (Kullback & Leibler, 1951; Jaynes, 1957; Csiszár, 1975). Related work develops geometric views of entropy optimization under linear constraints (Pavon & Ferrante, 2013) and generative procedures for sampling moment-constrained maximum-entropy laws (Lempereur et al., 2026). Particularly close in sensor-design motivation, Bhattacharya (2026) constructs a posterior and induced likelihood by minimizing KL divergence from a dynamical prior subject to distributional-accuracy constraints, including moment constraints. FIDE also uses KL projection relative to a dynamical baseline, but the role of the projection is different: for each candidate observable map, it provides a canonical completion of that design's moment fiber before the design is scored dynamically.

There are also dynamic entropy- and projection-based approaches to incomplete information. Maximum Caliber extends entropy principles to trajectory distributions (Dixit et al., 2018), and Schrödinger-bridge formulations can incorporate distributional moments and average currents (Movilla Miangolarra et al., 2024). Projection filters approximate nonlinear filtering dynamics on finite-dimensional exponential families (Brigo et al., 1999), with related Fokker–Planck approximations connecting tangent projection, moment matching, and Galerkin methods (Brigo & Pistone, 2017). Best-fit statistical closure takes another route, choosing paths in quasi-equilibrium families by minimizing a time-integrated weighted squared Liouville residual with a resolved/unresolved decomposition (Turkington, 2013). FIDE does not assume that the measured moments close the dynamics or identify the underlying law. Instead, the observable map itself is a design variable, so different candidate designs induce different moment fibers and different reference-relative completed paths. The completion is fixed before its dynamical correction cost is evaluated.

Reference-relative probability dynamics The frozen reference used by FIDE is closely related to modern learned probability-flow models. Flow matching and stochastic interpolants learn probability paths together with velocity fields (Lipman et al., 2023; Albergo & Vanden-Eijnden, 2023; Albergo et al., 2025), and optimal-transport couplings can be used to simplify the learned flows (Tong et al., 2024). In our experiments, endpoint-trained flow matching is one way to instantiate the frozen pair $(\widetilde{Q}_t, u_t)$. The theoretical construction itself is more general: it requires only a reference law path and velocity satisfying the continuity equation, and neither is assumed to coincide with the unknown intermediate population dynamics.

Other approaches modify or infer probability dynamics relative to prescribed constraints or prior processes. Schrödinger bridges construct path laws close to a reference stochastic process under marginal constraints (De Bortoli et al., 2021); TrajectoryNet infers continuous dynamics from cross-sectional populations (Tong et al., 2020); and constrained generative methods impose distributional (Khalafi et al., 2024), average (Hadou et al., 2026), or exact structural constraints through constrained flow geometry (Cheng et al., 2025). Dynamic optimal transport supplies the continuity-constrained kinetic-energy geometry underlying many such constructions (Benamou & Brenier, 2000; Ambrosio & Serfaty, 2008), with related formulations in minimum-energy distribution steering (Chen et al., 2016), optimal transport with prior dynamics (Chen et al., 2017), and dynamic-transport regularization of measure-valued reconstruction from indirect time-dependent observations (Bredies & Fanzon, 2020).

Measure-based nudging provides another related perspective. Lyu et al. (2026a) map macroscopic observations to microscopic mean-field particle corrections through Wasserstein-gradient feedback, while Lyu et al. (2026b) construct Wasserstein feedback from observed hydrodynamic moments for Vlasov–Poisson dynamics, where the moments leave the underlying phase-space law underdetermined. These methods seek corrections for a specified observation mechanism. In FIDE, the observation mechanism itself varies across candidate experimental designs, and the question is which scientifically adequate design induces the least reference-relative dynamical revision after its aggregate information has been canonically completed. More general Wasserstein Lagrangian formulations optimize density paths through combinations of kinetic and potential energies (Neklyudov et al., 2024); by contrast, FIDE does not reoptimize the projected path to reduce its action.

Minimum-energy realization and Sobolev geometry The dynamical criterion used by FIDE is most directly connected to work on minimum-energy realization of prescribed law paths. Action Matching learns minimum-kinetic-energy gradient velocities from temporally unpaired samples of

a given marginal evolution (Neklyudov et al., 2023). For Boltzmann density curves, Chemseddine et al. (2025) establish weighted-Poisson conditions for Wasserstein absolute continuity and minimum-norm realization and analyze velocity blow-up for particular interpolations. Related weighted-Poisson constructions connect variational Bayesian updating to feedback particle filtering (Laugesen et al., 2015), transport samples along geometric reference–target density interpolations in kernel Fisher–Rao flow (Maurais & Marzouk, 2024), and jointly optimize tilted density paths and their transport dynamics to obtain smoother velocities (Maurais et al., 2025). FIDE uses the same general realization geometry in a different outer problem: the moment-constrained path is determined by the measurement design and the common frozen reference before its realization cost is computed. Neither that path nor the reference velocity is subsequently optimized to make the action small.

The dual characterization of this correction is also related to Sobolev discrepancies. Finite-feature Sobolev discrepancy uses derivative Gram matrices and identifies the finite-feature approximation gap as an energy-norm error (Mroueh, 2018), while Sobolev Descent uses closely related geometry for particle transport (Mroueh et al., 2019). In FIDE, restricting the full expectation-rate discrepancy to the span of the measured observables produces the Tangent criterion, and the remaining energy defines the Hidden component. Wasserstein natural gradients and Wasserstein information matrices likewise pull transport geometry back to parametrized statistical families (Chen & Li, 2020; Li & Zhao, 2023), which is closely related to the inverse-transport Gram structure that appears in our comparisons.

Taken together, these literatures provide the main ingredients used by FIDE: information projection to complete finite statistical constraints, reference-relative transport geometry to quantify the effort required to realize a prescribed law path, and Sobolev duality to interpret that effort through observable expectation rates. Our contribution is not a new information-projection principle or a new underlying transport metric. It is to place these ingredients inside a residual experimental-design problem in which each candidate measurement map induces its own moment fiber and completed law path, scientific risk determines admissibility, and the remaining design freedom is resolved by the dynamical revision required relative to a common frozen reference.

B TECHNICAL STATEMENTS AND PROOFS

This appendix develops the technical results underlying Section 4. We first state the notation, weak-form conventions, and functional-analytic preliminaries used throughout the proofs. We then give the assumptions and proofs for the information-projected lift, physical-time multiplier dynamics, density-free continuity forcing, the minimum-energy and weighted-H^{-1} characterizations of Full, the sharp uniform certificate, Tangent, and the Full–Tangent–Hidden decomposition and equality condition. We also establish finite-reference information-projection feasibility, calibration sensitivity, and end-to-end differentiability of the fixed finite-dimensional Full objective. The finite-measure transport–reaction extension is developed alongside the balanced theory and evaluated separately in Appendix D.1.2. Operator-level comparisons and exact information-projected examples separating Full from Tangent, residual L^2, forward KL, and static Wasserstein displacement, together with the matched-initial stability result, are collected in Appendix C.

B.1 PRELIMINARIES AND STANDING CONVENTIONS

State space, laws, and boundary conventions. Whenever spatial derivatives are used, we identify the state space with a domain $\Omega \subseteq \mathbb{R}^d$, a flat torus, or $\mathbb{R}^d$. Bounded-domain statements use boundary conditions compatible with the corresponding continuity equation and variational problem; throughout the balanced theory these are periodic, sufficient-decay, or no-flux conditions chosen so that the stated integration-by-parts identities have no boundary contribution. Additional geometric hypotheses, such as connectedness or strict positivity of a density, are imposed only where explicitly required. Probability laws are denoted by capital letters and, when absolutely continuous with respect to Lebesgue measure, their densities by the corresponding lower-case letters. In particular, $\widetilde{Q}_t$ and $\widetilde{q}_t$ denote the frozen reference law and density, while $Q_t^{\boldsymbol{\eta}}$ and $q_t^{\boldsymbol{\eta}}$ denote the information-projected law and density associated with design $\boldsymbol{\eta}$. Expressions involving division by $q_t^{\boldsymbol{\eta}}$ are understood on its positive support. Expectations are written $\mathbb{E}_Q[\cdot]$ or $\mathbb{E}_q[\cdot]$ interchangeably when q is the density of Q.

Weak continuity equations. Unless the displayed densities and vector fields are regular enough to justify a strong formulation, all continuity equations are understood in the distributional sense. Specifically, for

$$\partial_t q_t + \nabla \cdot (q_t v_t) = s_t, \tag{17}$$

the weak formulation is

$$\int_0^1 \int_\Omega \left(q_t \, \partial_t \zeta + q_t v_t \cdot \nabla \zeta + s_t \zeta \right) \, dx \, dt = 0 \tag{18}$$

for every admissible smooth test function ζ with compact support in physical time and satisfying the corresponding spatial boundary convention. Whenever the map $t \mapsto \int f q_t \, dx$ is absolutely continuous, this implies

$$\frac{d}{dt} \mathbb{E}_{q_t}[f] = \mathbb{E}_{q_t}[\nabla f \cdot v_t] + \int_\Omega f(x) s_t(x) \, dx \tag{19}$$

for admissible time-independent scalar observables f. In the conservative case $s_t = 0$, applying this identity coordinatewise to a sufficiently regular time-independent observable vector $\mathbf{\Phi}_{\boldsymbol{\eta}} : \Omega \to \mathbb{R}^R$ gives

$$\frac{d}{dt} \mathbb{E}_{q_t}[\mathbf{\Phi}_{\boldsymbol{\eta}}] = \mathbb{E}_{q_t}[J_{\mathbf{\Phi}_{\boldsymbol{\eta}}} v_t]. \tag{20}$$

Weighted L^2 geometry. For scalar or vector fields f and g with the required square integrability, write

$$\langle f, g \rangle_{L^2(q)} := \mathbb{E}_q[f^\top g], \qquad \|f\|^2_{L^2(q)} := \mathbb{E}_q[\|f\|^2]. \tag{21}$$

For an admissible scalar function φ, define the weighted gradient-energy seminorm

$$\|\varphi\|^2_{\dot{H}^1(q)} := \mathbb{E}_q[\|\nabla \varphi\|^2]. \tag{22}$$

Whenever the weighted Dirichlet form is coercive modulo constants, as assumed explicitly in Appendix B.7, we use the mean-zero representative $\mathbb{E}_q[\varphi] = 0$ and denote by H_q the corresponding completed Hilbert energy space, equipped with the norm $\|\varphi\|_{H_q} := \|\varphi\|_{\dot{H}^1(q)}$. Thus the mean-zero condition fixes the constant gauge in the balanced Poisson problem; it is not a global assumption on every function appearing in the appendix. The screened space $H_{q,\kappa}$ used in the finite-measure extension is defined separately in Corollary 4, where its zeroth-order term removes the constant null direction.

Observable and Jacobian conventions. For a design $\boldsymbol{\eta}$, the observable map is $\mathbf{\Phi}_{\boldsymbol{\eta}} : \Omega \to \mathbb{R}^R$. Its Jacobian is taken in the row convention

$$J_{\mathbf{\Phi}_{\boldsymbol{\eta}}}(x) \in \mathbb{R}^{R \times d}, \tag{23}$$

so that $J_{\mathbf{\Phi}_{\boldsymbol{\eta}}} v \in \mathbb{R}^R$ for $v \in \mathbb{R}^d$ and $J^\top_{\mathbf{\Phi}_{\boldsymbol{\eta}}} a \in \mathbb{R}^d$ for $a \in \mathbb{R}^R$. Covariance matrices are

$$\mathrm{Cov}_Q(Y, Z) := \mathbb{E}_Q\left[(Y - \mathbb{E}_Q[Y])(Z - \mathbb{E}_Q[Z])^\top \right]. \tag{24}$$

Linear-algebra conventions. For a matrix M, $\mathrm{Range}(M)$ denotes its column space, $\ker(M)$ its nullspace, and $M^\dagger$ its Moore–Penrose pseudoinverse. All covariance and Gram matrices appearing below are symmetric positive semidefinite. Positive definiteness is imposed only when invertibility is required. In particular, positive definiteness of the projected covariance is needed for local multiplier identifiability and implicit differentiation, but not for uniqueness of the information-projected law itself.

Time integration and equality conventions. The design-level actions are integrated against the same fixed probability measure ρ on $[0, 1]$, unless a later comparison explicitly specializes to $\rho(dt) = dt$. Equalities between densities, scalar potentials, and vector fields are understood almost everywhere with respect to the relevant base measure. Integrated identities involving nonnegative action densities are valid in the extended-valued sense unless a finite-action assumption is explicitly needed; in particular, conclusions obtained by subtracting two integrated actions are invoked only when the corresponding quantities are finite.

Compact notation. For compactness, when $(t, \boldsymbol{\eta})$ is fixed we write

$$\widetilde{q} = \widetilde{q}_t, \qquad q = q_t^{\boldsymbol{\eta}}, \qquad \phi = \boldsymbol{\Phi}_{\boldsymbol{\eta}}, \qquad c = \boldsymbol{c}_{\boldsymbol{\eta}}(t), \qquad \lambda = \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}, \qquad u = u_t. \tag{25}$$

Expectations without an explicit subscript in a proof are taken under the law indicated by the surrounding expression. The conventions above are deliberately weaker than the hypotheses of individual results: existence of an information projection, differentiability of the calibrated multiplier, coercivity of the weighted Poisson problem, positivity of the projected density, and rank conditions are stated locally where they are used rather than imposed globally.

B.2 Finite-measure transport–reaction setup

The finite-measure results below require additional notation beyond the balanced probability-law construction. We write a finite positive measure as

$$\mu_t(dx) = M_t q_t(x)\, dx, \qquad M_t > 0, \qquad \int q_t(x)\, dx = 1, \tag{26}$$

where M_t is the total mass and q_t is its normalized shape. The normalized shape is information-projected relative to the frozen normalized reference shape $\widetilde{q}_t$, while the mass is treated as a separate scalar channel. We assume that the frozen reference shape satisfies

$$\partial_t \widetilde{q}_t + \nabla \cdot (\widetilde{q}_t u_t) = 0, \tag{27}$$

and that the reference mass evolution has a spatially constant logarithmic source rate $g_{\mathrm{ref}}(t)$. For a time-independent observable vector ϕ, define the raw and normalized moments by

$$\bar{c}_t := \int \phi\, d\mu_t = M_t \mathbb{E}_{q_t}[\phi], \qquad c_t := \frac{\bar{c}_t}{M_t} = \mathbb{E}_{q_t}[\phi]. \tag{28}$$

A corrected finite-measure evolution is taken to satisfy

$$\partial_t(M_t q_t) + \nabla \cdot (M_t q_t (u_t + \delta_t)) = M_t q_t \left(g_{\mathrm{ref}}(t) + \alpha_t\right), \tag{29}$$

where δ_t is the transport correction and α_t is the reaction-rate correction. Define

$$h_{\mathrm{shape}} := \frac{\partial_t q + \nabla \cdot (qu)}{q}, \qquad h_{\mathrm{ub}} := h_{\mathrm{shape}} + \frac{\dot{M}}{M} - g_{\mathrm{ref}}. \tag{30}$$

Equation equation 29 is then equivalent to

$$-\nabla \cdot (q\delta) + q\alpha = q h_{\mathrm{ub}}. \tag{31}$$

For a fixed reaction penalty $\kappa > 0$, define the instantaneous unbalanced Full action by

$$a^{\mathrm{ub}}_{\mathrm{full}}(t, \boldsymbol{\eta}) := M_t \inf_{\delta, \alpha} \left\{ \mathbb{E}_{q_t}\left[\|\delta\|^2 + \kappa \alpha^2\right] : -\nabla \cdot (q_t \delta) + q_t \alpha = q_t h_{\mathrm{ub}} \right\}. \tag{32}$$

The separation of this forcing and action into normalized-shape and total-mass components is established in Corollary 5.

Remark 1 (Mass is separate from the normalized-shape exponential family)**.** *The mass M_t should not be encoded by appending the constant observable 1 to the normalized-shape information projection. Normalization already imposes $\mathbb{E}_q[1] = 1$, so doing so introduces an exactly redundant coordinate and makes the corresponding covariance matrix singular. The constant observable may, however, be included in the unbalanced Tangent observable vector after the shape has been calibrated. In that setting its spatial Jacobian vanishes, and its moment-rate equation isolates the total-mass reaction channel.*

B.3 Finite-sample Law baseline

Equation equation 6 defines the population ideal using the exact risk infimum. In the experiments, this value is approximated by the selection risk of a numerically optimized Law design. The resulting baseline defines a finite-search approximation to the admissible set rather than a certificate of global risk optimality. Risk tolerances apply to the selection distribution; their transfer to fresh observations is evaluated separately.

B.4 CANONICAL INFORMATION PROJECTION

Technical conditions. Fix $(t, \boldsymbol{\eta})$. Assume that $\boldsymbol{c}_{\boldsymbol{\eta}}(t)$ lies in the relative interior of the feasible moment set, that the dual objective $A_{t,\boldsymbol{\eta}}(\lambda) - \lambda^\top \boldsymbol{c}_{\boldsymbol{\eta}}(t)$ attains a minimizer $\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}$, and that $A_{t,\boldsymbol{\eta}}$ is finite on an open neighborhood of that minimizer. If $C_{t,\boldsymbol{\eta}} \succ 0$ in a minimal observable representation, the calibrated multiplier is locally unique; without this rank condition, the projected law remains unique but its coordinates need not be.

Proof of the information-projection part of Theorem 1. Let

$$\mathcal{Q}_c := \left\{ Q \ll \widetilde{Q}_t : \mathbb{E}_Q[\phi] = c \right\} \tag{33}$$

with the implicit finiteness conditions from equation 3. For any λ at which $A_{t,\boldsymbol{\eta}}(\lambda) < \infty$, define the exponentially tilted law

$$\frac{dQ_\lambda}{d\widetilde{Q}_t}(x) = \exp\!\big(\lambda^\top \phi(x) - A_{t,\boldsymbol{\eta}}(\lambda)\big)\,. \tag{34}$$

The normalization follows directly from the definition of $A_{t,\boldsymbol{\eta}}$.

By assumption, the dual objective

$$D(\lambda) := A_{t,\boldsymbol{\eta}}(\lambda) - \lambda^\top c \tag{35}$$

attains its minimum at $\lambda = \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}$, and $A_{t,\boldsymbol{\eta}}$ is finite on an open neighborhood of that point. On this neighborhood the standard differentiation formula for a log moment-generating function gives

$$\nabla_\lambda A_{t,\boldsymbol{\eta}}(\lambda) = \mathbb{E}_{Q_\lambda}[\phi], \tag{36}$$

and

$$\nabla^2_\lambda A_{t,\boldsymbol{\eta}}(\lambda) = \mathrm{Cov}_{Q_\lambda}(\phi, \phi). \tag{37}$$

Since D is differentiable at an unconstrained minimizer,

$$0 = \nabla D(\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}) = \mathbb{E}_{Q_{\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}}}[\phi] - c. \tag{38}$$

Hence the tilted law $Q_{\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}}$ belongs to $\mathcal{Q}_c$. In the shorthand of equation 25, the calibration identity is

$$\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\phi] = c. \tag{39}$$

Set $Q^\star = Q_{\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}}$. For any feasible $Q \in \mathcal{Q}_c$ for which the displayed relative entropies are finite, the Radon–Nikodym chain rule and equation 34 give

$$\begin{aligned} D_{\mathrm{KL}}(Q\|\widetilde{Q}_t) &= D_{\mathrm{KL}}(Q\|Q^\star) + \mathbb{E}_Q\left[\log \frac{dQ^\star}{d\widetilde{Q}_t}\right] \\ &= D_{\mathrm{KL}}(Q\|Q^\star) + \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}^\top \mathbb{E}_Q[\phi] - A_{t,\boldsymbol{\eta}}(\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}) \\ &= D_{\mathrm{KL}}(Q\|Q^\star) + \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}^\top c - A_{t,\boldsymbol{\eta}}(\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}). \end{aligned} \tag{40}$$

For $Q = Q^\star$ the final two terms equal $D_{\mathrm{KL}}(Q^\star\|\widetilde{Q}_t)$. Therefore

$$D_{\mathrm{KL}}(Q\|\widetilde{Q}_t) - D_{\mathrm{KL}}(Q^\star\|\widetilde{Q}_t) = D_{\mathrm{KL}}(Q\|Q^\star) \geq 0. \tag{41}$$

Equality holds only when $Q = Q^\star$ almost everywhere. Thus the I-projection is unique at the law level and has the density in equation 7.

Finally, equation 37 evaluated at the calibrated multiplier gives

$$\nabla^2_\lambda A_{t,\boldsymbol{\eta}}(\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}) = C_{t,\boldsymbol{\eta}}. \tag{42}$$

If $C_{t,\boldsymbol{\eta}}$ is positive definite in a minimal observable representation, the dual Hessian is nonsingular at the optimizer. By continuity of the Hessian, the dual objective is strictly convex on a sufficiently small neighborhood, and the multiplier is locally unique. If the observable coordinates are redundant, different multipliers can generate the same log-density up to an additive constant on the support of $\widetilde{Q}_t$, so multiplier non-uniqueness does not contradict law-level uniqueness. □

Proposition 3 (Finite-reference information-projection feasibility)**.** *Fix $(t, \boldsymbol{\eta})$ and let the empirical reference be $\widehat{\widetilde{Q}}_N = \sum_{i=1}^N \bar{w}_i \delta_{x_i}$ with $\bar{w}_i > 0$ and $\sum_i \bar{w}_i = 1$. Write $z_i = \boldsymbol{\Phi}_{\boldsymbol{\eta}}(x_i)$ and $\mathcal{K}_N = \mathrm{conv}\{z_1, \dots, z_N\}$. For a target moment $\widehat{c}$, the empirical constraint $\sum_i p_i z_i = \widehat{c}$, $p_i \geq 0$, $\sum_i p_i = 1$, is feasible if and only if $\widehat{c} \in \mathcal{K}_N$, and whenever it is feasible the empirical KL projection relative to $(\bar{w}_i)_{i=1}^N$ is unique. If $\widehat{c} \in \mathrm{relint}(\mathcal{K}_N)$, the minimizer has strictly positive weights of the form*

$$p_i(\lambda) = \frac{\bar{w}_i \exp(\lambda^\top z_i)}{\sum_{j=1}^N \bar{w}_j \exp(\lambda^\top z_j)}, \qquad \sum_{i=1}^N p_i(\lambda) z_i = \widehat{c}. \tag{43}$$

After removal of affine redundancies in the observable coordinates, the corresponding finite multiplier is unique and the projected covariance $\widehat{C} = \sum_i p_i(\lambda)(z_i - \widehat{c})(z_i - \widehat{c})^\top$ is positive definite. If $\widehat{c}$ lies on the relative boundary of $\mathcal{K}_N$, the unique empirical KL projection still exists but cannot be represented by a finite full-support exponential tilt; it is obtained only as a limiting tilt with vanishing weights on part of the empirical support. If $\widehat{c} \notin \mathcal{K}_N$, no exact empirical information projection supported on the reference bank exists.

Proof. The moments attainable by probability weights on the empirical support are exactly the convex combinations of $z_1, \dots, z_N$, which proves the feasibility statement. On the feasible simplex the objective $\sum_i p_i \log(p_i/\bar{w}_i)$ is strictly convex, so its minimizer is unique. For a target in $\mathrm{relint}(\mathcal{K}_N)$, the equality-constrained first-order conditions give the exponential weights in equation 43; conversely, every finite multiplier gives strictly positive weights and therefore a moment in $\mathrm{relint}(\mathcal{K}_N)$. In a minimal affine representation, zero variance of a nonzero linear combination of the features would make that combination constant on the full empirical support, contradicting minimality, so $\widehat{C} \succ 0$. A target on a proper relative face requires the minimizing distribution to place zero mass outside that face and hence cannot be produced by finite strictly positive exponential weights. Outside $\mathcal{K}_N$ the moment constraint is infeasible. □

Corollary 1 (Finite-measure information projection with prescribed mass)**.** *Fix $M > 0$ and a raw finite-measure moment vector $\bar{c}$. Let $c = \bar{c}/M$, and suppose the conditions of Theorem 1 hold for this normalized moment vector. Then the finite measure $\mu^\star = MQ^\star$, where $Q^\star$ is the information projection in Theorem 1, is the unique finite measure of mass M whose normalized shape minimizes $D_{\mathrm{KL}}(Q\|\widetilde{Q}_t)$ subject to $\int \phi \, d(MQ) = \bar{c}$. Equivalently, $\mu^\star$ uniquely minimizes $M D_{\mathrm{KL}}(Q\|\widetilde{Q}_t)$ over the same fixed-mass fiber, and its normalized density has the same exponential-tilt representation equation 7.*

Proof. Every finite measure μ of mass M can be written uniquely as $\mu = MQ$ with Q a probability law, and the raw constraint $\int \phi \, d\mu = \bar{c}$ is equivalent to $\mathbb{E}_Q[\phi] = \bar{c}/M = c$. Since multiplication of the objective by the positive constant M does not change its minimizer, Theorem 1 applies directly to the normalized shape. □

B.5 Physical-time multiplier dynamics

Proposition 4 (Multiplier dynamics)**.** *Under the regularity assumptions stated below and in a minimal observable representation with $C_{t,\boldsymbol{\eta}} \succ 0$, the calibrated multiplier satisfies equation 8.*

Technical conditions. Fix $\boldsymbol{\eta}$ and a time t. Assume that $\boldsymbol{c}_{\boldsymbol{\eta}}(s)$ is continuously differentiable in s near t, that the reference density and velocity are regular enough for the reference continuity equation and the integrations by parts below, and that the normalized exponential family $q_{\lambda,s}$ defined below is continuously differentiable in (λ, s) near $(\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}, t)$ with differentiation and integration interchangeable. In a minimal observable representation, assume $C_{t,\boldsymbol{\eta}} \succ 0$.

Proof of Proposition 4. Fix $\boldsymbol{\eta}$ and suppress it from the notation. For λ and t in the neighborhood specified above, define

$$q_{\lambda,t}(x) := \widetilde{q}_t(x) \exp\big(\lambda^\top \phi(x) - A_t(\lambda)\big), \tag{44}$$

and the calibration residual

$$\mathcal{F}(\lambda, t) := \mathbb{E}_{q_{\lambda,t}}[\phi] - c_t. \tag{45}$$

By the stated regularity assumptions, $\mathcal{F}$ is continuously differentiable near (λ_t, t), and

$$\partial_\lambda \mathcal{F}(\lambda_t, t) = \mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}(\phi, \phi) = C_{t,\boldsymbol{\eta}}. \tag{46}$$

Since $C_{t,\boldsymbol{\eta}} \succ 0$, the implicit-function theorem gives a locally unique continuously differentiable calibrated path $t \mapsto \lambda_t$ satisfying $\mathcal{F}(\lambda_t, t) = 0$. In particular, $\dot{\lambda}_t$ exists, so the differentiations below are justified. Let

$$s_t(x) := \partial_t \log \widetilde{q}_t(x) \tag{47}$$

where the reference density is positive. From the exponential representation,

$$\log q_t(x) = \log \widetilde{q}_t(x) + \lambda_t^\top \phi(x) - A_t(\lambda_t). \tag{48}$$

At fixed λ, differentiation of the log-partition function in physical time gives

$$\partial_t A_t(\lambda) = \mathbb{E}_{q_{\lambda,t}}[s_t]. \tag{49}$$

Along the calibrated path $t \mapsto \lambda_t$,

$$\frac{d}{dt} A_t(\lambda_t) = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[s_t] + \dot{\lambda}_t^\top c_t, \tag{50}$$

where $c_t = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\phi]$.

Differentiating the calibration identity $c_t = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\phi]$ and using that ϕ has no explicit physical-time dependence yields

$$\begin{aligned} \dot{c}_t &= \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\phi\, \partial_t \log q_t] \\ &= \mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}(\phi, s_t) + C_{t,\boldsymbol{\eta}} \dot{\lambda}_t. \end{aligned} \tag{51}$$

It remains to express the covariance with s_t without differentiating the reference density explicitly.

Let

$$w_t(x) := \frac{q_t(x)}{\widetilde{q}_t(x)} = \exp\big(\lambda_t^\top \phi(x) - A_t(\lambda_t)\big). \tag{52}$$

The reference continuity equation gives $\partial_t \widetilde{q}_t = -\nabla \cdot (\widetilde{q}_t u_t)$. Therefore, using the assumed boundary conditions,

$$\begin{aligned} \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[s_t] &= \int w_t\, \partial_t \widetilde{q}_t\, dx \\ &= -\int w_t\, \nabla \cdot (\widetilde{q}_t u_t)\, dx \\ &= \int \widetilde{q}_t u_t \cdot \nabla w_t\, dx. \end{aligned} \tag{53}$$

Since

$$\nabla w_t = w_t J_\phi^\top \lambda_t, \tag{54}$$

we obtain

$$\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[s_t] = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}\big[\lambda_t^\top J_\phi u_t\big] = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\lambda_t^\top m_t]. \tag{55}$$

Similarly, componentwise integration by parts gives

$$\begin{aligned} \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\phi s_t] &= \int w_t \phi\, \partial_t \widetilde{q}_t\, dx \\ &= \int \widetilde{q}_t u_t \cdot \nabla(w_t \phi)\, dx \\ &= \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[m_t] + \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\phi\, \lambda_t^\top m_t]. \end{aligned} \tag{56}$$

Subtracting $\mathbb{E}[\phi]\mathbb{E}[s_t] = c_t \mathbb{E}[\lambda_t^\top m_t]$ yields

$$\mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}(\phi, s_t) = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[m_t] + \mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}(\phi, \lambda_t^\top m_t). \tag{57}$$

Substitution into equation 51 gives

$$C_{t,\boldsymbol{\eta}} \dot{\lambda}_t = \dot{c}_t - \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[m_t] - \mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}(\phi, \lambda_t^\top m_t), \tag{58}$$

which is equation 8. Positive definiteness of $C_{t,\boldsymbol{\eta}}$ in the stated minimal representation makes this a uniquely solvable $R \times R$ linear system. This proves the multiplier-dynamics clause of Theorem 1. □

Corollary 2 (Multiplier dynamics for a finite-measure path)**.** *Suppose the assumptions of Proposition 4 hold for the normalized shape q_t, and suppose $M_t > 0$ and the raw moment path $\bar{c}_t$ in equation 28 are continuously differentiable. Then the same covariance equation holds with the normalized target*

$$c_t = \frac{\bar{c}_t}{M_t}, \qquad \dot{c}_t = \frac{\dot{\bar{c}}_t}{M_t} - \frac{\dot{M}_t}{M_t} c_t, \tag{59}$$

namely

$$C_{t,\boldsymbol{\eta}} \dot{\lambda}_{t,\boldsymbol{\eta}} = \frac{\dot{\bar{c}}_t}{M_t} - \frac{\dot{M}_t}{M_t} c_t - \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[m_{t,\boldsymbol{\eta}}] - \mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}\left(\boldsymbol{\Phi}_{\boldsymbol{\eta}}, \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}^\top m_{t,\boldsymbol{\eta}}\right). \tag{60}$$

Proof. The information projection acts only on the normalized shape, so Proposition 4 applies with target $c_t = \bar{c}_t / M_t$. Equation equation 59 is the quotient rule, and substitution into equation 8 gives equation 60. □

B.6 Density-free continuity forcing

Proposition 5 (Density-free continuity forcing)**.** *Under the assumptions of Proposition 4, the reference-relative continuity residual $h_{t,\boldsymbol{\eta}}$ defined in Section 3 is given by equation 9 and has zero projected-law mean:*

$$\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[h_{t,\boldsymbol{\eta}}] = 0. \tag{61}$$

Proof of Proposition 5. Again fix $\boldsymbol{\eta}$ and write $q_t = q_t^{\boldsymbol{\eta}}$. For any positive sufficiently regular density q and vector field u,

$$\partial_t q + \nabla \cdot (qu) = q\left[(\partial_t + u \cdot \nabla) \log q + \nabla \cdot u\right]. \tag{62}$$

The reference continuity equation implies

$$(\partial_t + u_t \cdot \nabla) \log \widetilde{q}_t + \nabla \cdot u_t = 0. \tag{63}$$

Using equation 48, the time independence of ϕ, and $m_t = J_\phi u_t$,

$$(\partial_t + u_t \cdot \nabla) \log q_t + \nabla \cdot u_t = \dot{\lambda}_t^\top \phi + \lambda_t^\top m_t - \frac{d}{dt} A_t(\lambda_t). \tag{64}$$

Equations equation 50 and equation 55 give

$$\frac{d}{dt} A_t(\lambda_t) = \dot{\lambda}_t^\top c_t + \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\lambda_t^\top m_t]. \tag{65}$$

Hence the bracket in equation 62 equals

$$\dot{\lambda}_t^\top (\phi - c_t) + \lambda_t^\top \left(m_t - \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[m_t]\right) = h_{t,\boldsymbol{\eta}}. \tag{66}$$

Substituting into equation 62 gives $\partial_t q_t + \nabla \cdot (q_t u_t) = q_t h_{t,\boldsymbol{\eta}}$ and proves the explicit formula equation 9.

Finally, calibration gives $\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\phi - c_t] = 0$, and the second centered term also has zero expectation, proving equation 61. Equivalently, the zero-mean property follows by integrating the reference-relative continuity identity over space and using conservation of total mass. Together with Proposition 4 and the information-projection argument above, this completes the proof of Theorem 1. □

Corollary 3 (Density-free unbalanced forcing)**.** *Under the assumptions of Corollary 2, the normalized-shape residual satisfies*

$$h_{\mathrm{shape}}(x) = \dot{\lambda}_t^\top (\phi(x) - c_t) + \lambda_t^\top \left(m_t(x) - \mathbb{E}_{q_t}[m_t]\right), \tag{67}$$

and therefore the finite-measure reference-relative forcing in equation 30 is

$$h_{\mathrm{ub}}(x) = \dot{\lambda}_t^\top (\phi(x) - c_t) + \lambda_t^\top \left(m_t(x) - \mathbb{E}_{q_t}[m_t]\right) + \frac{\dot{M}_t}{M_t} - g_{\mathrm{ref}}(t). \tag{68}$$

Moreover,

$$\mathbb{E}_{q_t}[h_{\mathrm{shape}}] = 0, \qquad \mathbb{E}_{q_t}[h_{\mathrm{ub}}] = \frac{\dot{M}_t}{M_t} - g_{\mathrm{ref}}(t). \tag{69}$$

Proof. Equation equation 67 is Proposition 5 applied to the normalized shape. Adding the scalar mass-rate discrepancy from equation 30 gives equation 68. Taking expectation and using the centering in equation 67 gives equation 69. □

B.7 MINIMUM-ENERGY FULL CORRECTION

Fix $(t, \boldsymbol{\eta})$ and write $q = q_t^{\boldsymbol{\eta}}$ and $h = h_{t,\boldsymbol{\eta}}$. Let H_q be the mean-zero weighted energy space defined in Appendix B.1. Assume that the weighted Dirichlet form is coercive on this space, for example through a weighted Poincaré inequality, and that $\varphi \mapsto \mathbb{E}_q[h\varphi]$ defines a bounded linear functional on H_q. The weak Poisson equation is

$$\mathbb{E}_q[\nabla\psi \cdot \nabla\varphi] = -\mathbb{E}_q[h\varphi] \qquad \text{for every } \varphi \in H_q. \tag{70}$$

Proof of the minimum-energy and Poisson parts of Theorem 2. Define the bilinear form

$$a(\psi, \varphi) := \mathbb{E}_q[\nabla\psi \cdot \nabla\varphi] \tag{71}$$

on the mean-zero energy space H_q. By construction of the energy norm,

$$|a(\psi, \varphi)| \le \|\psi\|_{H_q} \|\varphi\|_{H_q}, \qquad a(\psi, \psi) = \|\psi\|_{H_q}^2. \tag{72}$$

The weighted Poincaré/coercivity assumption ensures that the mean-zero gauge removes the constant null direction and that this energy space is the appropriate Hilbert space for the problem. By assumption,

$$\ell(\varphi) := -\mathbb{E}_q[h\varphi] \tag{73}$$

is a bounded linear functional on H_q. Lax–Milgram therefore gives a unique $\psi \in H_q$ satisfying

$$a(\psi, \varphi) = \ell(\varphi) \qquad \text{for every } \varphi \in H_q, \tag{74}$$

which is exactly equation 70. Since $\mathbb{E}_q[h] = 0$, the restriction to mean-zero test functions causes no loss: for any admissible φ, setting $\varphi_0 := \varphi - \mathbb{E}_q[\varphi]$ gives $\nabla\varphi_0 = \nabla\varphi$ and $\mathbb{E}_q[h\varphi_0] = \mathbb{E}_q[h\varphi]$, so the same weak identity holds for arbitrary admissible test functions. Under the stated periodic, decay, or no-flux boundary conditions, integration by parts converts this weak identity into

$$\nabla \cdot (q\nabla\psi) = qh \tag{75}$$

in the distributional sense, together with the gauge $\mathbb{E}_q[\psi] = 0$.

Set

$$\delta^\star := -\nabla\psi. \tag{76}$$

Then $\delta^\star \in L^2(q; \mathbb{R}^d)$. The weak Poisson equation gives, for every $\varphi \in H_q$,

$$\mathbb{E}_q[\delta^\star \cdot \nabla\varphi] = \mathbb{E}_q[h\varphi]. \tag{77}$$

This is precisely the weak form of

$$\nabla \cdot (q\delta^\star) = -qh, \tag{78}$$

so $u_t + \delta^\star$ satisfies the projected-law continuity equation.

Now let $\delta \in L^2(q; \mathbb{R}^d)$ be any other feasible correction and define $w := \delta - \delta^\star$. Subtracting the two weak constraints gives the identity first on the admissible test class. Because $w \in L^2(q; \mathbb{R}^d)$, the map $\varphi \mapsto \mathbb{E}_q[w \cdot \nabla\varphi]$ is continuous in the H_q energy norm; density therefore extends the identity to all of H_q:

$$\mathbb{E}_q[w \cdot \nabla\varphi] = 0 \qquad \text{for every } \varphi \in H_q. \tag{79}$$

Taking $\varphi = \psi$ yields

$$\mathbb{E}_q[w \cdot \nabla\psi] = 0, \tag{80}$$

and therefore, because $\delta^\star = -\nabla\psi$,

$$\mathbb{E}_q[\delta^\star \cdot w] = 0. \tag{81}$$

Consequently

$$\begin{aligned} \mathbb{E}_q[\|\delta\|^2] &= \mathbb{E}_q[\|\delta^\star + w\|^2] \\ &= \mathbb{E}_q[\|\delta^\star\|^2] + \mathbb{E}_q[\|w\|^2] \\ &\ge \mathbb{E}_q[\|\delta^\star\|^2]. \end{aligned} \tag{82}$$

Equality holds if and only if $w = 0$ q-almost everywhere. Thus $\delta^\star$ is the unique minimum-energy feasible correction, and

$$a_{\mathrm{full}}(t, \boldsymbol{\eta}) = \mathbb{E}_q[\|\delta^\star\|^2] = \mathbb{E}_q[\|\nabla\psi\|^2]. \tag{83}$$

This proves the minimum-energy and Poisson assertions in Theorem 2. The weighted H^{-1} identity in equation 13 is proved in Appendix B.8. □

The screened operator below is the reference-relative analogue of the transport–reaction operator in Sobolev–Fisher discrepancy and Unbalanced Sobolev Descent (Mroueh & Rigotti, 2020). Its gradient and zeroth-order terms correspond to transport and reaction costs, respectively; the present forcing comes from the prescribed finite-measure path, with the boundary conditions stated in this appendix.

Corollary 4 (Minimum-energy unbalanced move–reaction correction)**.** *Fix $(t, \boldsymbol{\eta})$, write $q = q_t^{\boldsymbol{\eta}}$, $M = M_t$, and $h_{\mathrm{ub}} = h_{\mathrm{ub},t,\boldsymbol{\eta}}$, and let $\kappa > 0$. Let $H_{q,\kappa}$ be the completion of the admissible test functions under*

$$\|\psi\|^2_{H_{q,\kappa}} := \mathbb{E}_q\left[\|\nabla\psi\|^2 + \frac{\psi^2}{\kappa}\right]. \tag{84}$$

Assume that $\varphi \mapsto \mathbb{E}_q[h_{\mathrm{ub}}\varphi]$ is a bounded linear functional on $H_{q,\kappa}$. Then there exists a unique $\psi \in H_{q,\kappa}$ satisfying

$$\mathbb{E}_q\left[\nabla\psi\cdot\nabla\varphi + \frac{1}{\kappa}\psi\varphi\right] = \mathbb{E}_q[h_{\mathrm{ub}}\varphi] \qquad \textit{for every } \varphi \in H_{q,\kappa}. \tag{85}$$

Under the stated boundary and regularity assumptions, this is the weak form of

$$-\nabla\cdot(q\nabla\psi) + \frac{q}{\kappa}\psi = q h_{\mathrm{ub}}. \tag{86}$$

The unique minimum-energy correction in equation 32 is

$$\delta^\star = \nabla\psi, \qquad \alpha^\star = \frac{\psi}{\kappa}, \tag{87}$$

and

$$a^{\mathrm{ub}}_{\mathrm{full}}(t,\boldsymbol{\eta}) = M\mathbb{E}_q\left[\|\nabla\psi\|^2 + \frac{\psi^2}{\kappa}\right]. \tag{88}$$

Proof. Define

$$a_\kappa(\psi,\varphi) := \mathbb{E}_q\left[\nabla\psi\cdot\nabla\varphi + \frac{1}{\kappa}\psi\varphi\right]. \tag{89}$$

This bilinear form is continuous and coercive on $H_{q,\kappa}$ because $a_\kappa(\psi,\psi) = \|\psi\|^2_{H_{q,\kappa}}$. Lax–Milgram therefore gives a unique solution of equation 85. Integration by parts yields equation 86 whenever the strong form is justified.

Set $(\delta^\star, \alpha^\star) = (\nabla\psi, \psi/\kappa)$. Equation equation 85 gives

$$\mathbb{E}_q[\delta^\star\cdot\nabla\varphi + \alpha^\star\varphi] = \mathbb{E}_q[h_{\mathrm{ub}}\varphi], \tag{90}$$

which is the weak form of equation 31, so $(\delta^\star, \alpha^\star)$ is feasible. Let (δ, α) be any other feasible pair and write $w = \delta - \delta^\star$ and $\gamma = \alpha - \alpha^\star$. Subtracting the two weak constraints gives

$$\mathbb{E}_q[w\cdot\nabla\varphi + \gamma\varphi] = 0 \tag{91}$$

for every $\varphi \in H_{q,\kappa}$. Taking $\varphi = \psi$ yields

$$\mathbb{E}_q[w\cdot\delta^\star + \kappa\gamma\alpha^\star] = \mathbb{E}_q[w\cdot\nabla\psi + \gamma\psi] = 0. \tag{92}$$

Hence, in the product norm induced by $\mathbb{E}_q[\|\delta\|^2 + \kappa\alpha^2]$,

$$\begin{aligned}\mathbb{E}_q[\|\delta\|^2 + \kappa\alpha^2] &= \mathbb{E}_q[\|\delta^\star\|^2 + \kappa(\alpha^\star)^2] + \mathbb{E}_q[\|w\|^2 + \kappa\gamma^2] \\ &\geq \mathbb{E}_q[\|\delta^\star\|^2 + \kappa(\alpha^\star)^2].\end{aligned} \tag{93}$$

Equality holds only when $w = 0$ and $\gamma = 0$ in the corresponding $L^2(q)$ spaces. Multiplication by $M > 0$ proves both uniqueness and equation 88. □

Corollary 5 (Mass–shape decomposition of unbalanced Full action)**.** *Under the assumptions of Corollary 4, define the spatially constant logarithmic mass-rate discrepancy*

$$b_t := \frac{\dot{M}_t}{M_t} - g_{\mathrm{ref}}(t), \tag{94}$$

so that

$$h_{\mathrm{ub}} = h_{\mathrm{shape}} + b_t, \qquad \mathbb{E}_{q_t}[h_{\mathrm{shape}}] = 0. \tag{95}$$

Let $\psi_{\mathrm{shape}} \in H_{q_t,\kappa}$ denote the unique solution of

$$\mathbb{E}_{q_t}\left[\nabla\psi_{\mathrm{shape}} \cdot \nabla\varphi + \frac{1}{\kappa}\psi_{\mathrm{shape}}\varphi\right] = \mathbb{E}_{q_t}[h_{\mathrm{shape}}\varphi] \qquad \textit{for every } \varphi \in H_{q_t,\kappa}. \tag{96}$$

Then

$$\mathbb{E}_{q_t}[\psi_{\mathrm{shape}}] = 0, \qquad \psi = \psi_{\mathrm{shape}} + \kappa b_t, \tag{97}$$

where ψ is the Full screened potential from Corollary 4. Consequently,

$$\delta^\star = \nabla\psi_{\mathrm{shape}}, \qquad \alpha^\star = b_t + \frac{\psi_{\mathrm{shape}}}{\kappa}, \tag{98}$$

and the instantaneous unbalanced Full action decomposes orthogonally as

$$a^{\mathrm{ub}}_{\mathrm{full}}(t,\boldsymbol{\eta}) = M_t\kappa b_t^2 + M_t\mathbb{E}_{q_t}\left[\|\nabla\psi_{\mathrm{shape}}\|^2 + \frac{\psi^2_{\mathrm{shape}}}{\kappa}\right]. \tag{99}$$

The first term is the cost of the spatially uniform reaction required to correct the total-mass rate, while the second is the screened correction cost of the mean-zero shape forcing.

Proof. Because $\mathbb{E}_{q_t}[h_{\mathrm{shape}}] = 0$, testing equation 96 with $\varphi = 1$ gives

$$\frac{1}{\kappa}\mathbb{E}_{q_t}[\psi_{\mathrm{shape}}] = 0, \tag{100}$$

and hence $\mathbb{E}_{q_t}[\psi_{\mathrm{shape}}] = 0$.

Since constants have zero spatial gradient,

$$-\nabla\cdot(q_t\nabla(\kappa b_t)) + \frac{q_t}{\kappa}(\kappa b_t) = q_t b_t. \tag{101}$$

Adding this identity to the screened equation for ψ_{shape} shows that

$$\psi = \psi_{\mathrm{shape}} + \kappa b_t \tag{102}$$

solves

$$-\nabla\cdot(q_t\nabla\psi) + \frac{q_t}{\kappa}\psi = q_t(h_{\mathrm{shape}} + b_t) = q_t h_{\mathrm{ub}}. \tag{103}$$

Uniqueness in Corollary 4 therefore identifies this sum with the Full screened potential. Equation equation 98 follows immediately from equation 87.

Finally,

$$\begin{aligned}
\mathbb{E}_{q_t}\left[\|\nabla\psi\|^2 + \frac{\psi^2}{\kappa}\right] &= \mathbb{E}_{q_t}\left[\|\nabla\psi_{\mathrm{shape}}\|^2 + \frac{(\psi_{\mathrm{shape}} + \kappa b_t)^2}{\kappa}\right] \\
&= \mathbb{E}_{q_t}\left[\|\nabla\psi_{\mathrm{shape}}\|^2 + \frac{\psi^2_{\mathrm{shape}}}{\kappa}\right] + 2b_t\mathbb{E}_{q_t}[\psi_{\mathrm{shape}}] + \kappa b_t^2 \\
&= \mathbb{E}_{q_t}\left[\|\nabla\psi_{\mathrm{shape}}\|^2 + \frac{\psi^2_{\mathrm{shape}}}{\kappa}\right] + \kappa b_t^2.
\end{aligned} \tag{104}$$

Multiplying by M_t and using equation 88 proves equation 99. □

Remark 2 (Operator form of the mass–shape decomposition)**.** *Suppose additionally that the weighted Dirichlet form admits the self-adjoint operator representation*

$$\mathsf{L}_{q_t} = -q_t^{-1}\nabla\cdot(q_t\nabla) \tag{105}$$

on $L^2(q_t)$, so that $\mathsf{L}_{q_t} + \kappa^{-1}I$ is invertible. Then

$$\psi_{\mathrm{shape}} = \left(\mathsf{L}_{q_t} + \kappa^{-1}I\right)^{-1} h_{\mathrm{shape}}, \tag{106}$$

and Corollary 5 becomes

$$a^{\mathrm{ub}}_{\mathrm{full}}(t,\boldsymbol{\eta}) = M_t\kappa\left(\frac{\dot{M}_t}{M_t} - g_{\mathrm{ref}}(t)\right)^2 + M_t\left\langle h_{\mathrm{shape}}, \left(\mathsf{L}_{q_t} + \kappa^{-1}I\right)^{-1} h_{\mathrm{shape}}\right\rangle_{L^2(q_t)}. \tag{107}$$

Thus the spatially constant mass-rate mismatch and the mean-zero shape forcing are orthogonal in the screened energy geometry.

Remark 3 (Interpretation of the reaction penalty)**.** *The parameter $\kappa > 0$ fixes the relative cost of reaction versus transport and is chosen before measurement-design optimization. Since δ has units of state-distance per unit time and α has units of inverse time, κ has units of squared state distance, or is dimensionless after the corresponding nondimensionalization. Larger κ makes reaction more expensive, while smaller κ allows the forcing to be realized more cheaply through reaction. On the mean-zero shape subspace, $(\mathsf{L}_q + \kappa^{-1}I)^{-1} \to \mathsf{L}_q^{-1}$ as $\kappa \to \infty$ whenever the balanced inverse is well defined. Thus, if $b_t = \dot{M}_t/M_t - g_{\mathrm{ref}}(t) = 0$, the unbalanced shape cost approaches the balanced transport cost; if $b_t \neq 0$, the mass term $M_t\kappa b_t^2$ diverges, so no finite balanced limit exists. We therefore treat κ as a fixed modeling choice rather than tuning it separately across competing designs.*

Remark 4 (Screening removes the balanced compatibility and gauge conditions)**.** *The reaction term removes the constant null direction present in the balanced Dirichlet seminorm. In the balanced problem one therefore imposes a mean-zero gauge and requires $\mathbb{E}_q[h] = 0$. For finite κ, the term $\kappa^{-1}\mathbb{E}_q[\psi^2]$ controls constants directly, so neither a gauge condition nor zero-mean forcing is required. Testing equation 85 with $\varphi = 1$ gives $\mathbb{E}_q[\psi]/\kappa = \mathbb{E}_q[h_{\mathrm{ub}}]$. By Corollary 5, this mean component corresponds exactly to the spatially uniform reaction $b_t = \dot{M}_t/M_t - g_{\mathrm{ref}}(t)$, while the remaining screened potential has zero q_t-mean.*

Remark 5 (Scope of regularized density–source evaluations)**.** *At a fixed time, the weighted-Poisson construction applies to a probability density q and a density-level source $s = qh$ satisfying $\mathbb{E}_q[h] = 0$, whether or not this pair has been derived from a prescribed law path. Under the assumptions of Appendix B.7 for $h = s/q$, define*

$$\mathfrak{a}(q,s) := \inf_{\delta \in L^2(q;\mathbb{R}^d)} \left\{\mathbb{E}_q[\|\delta\|^2] : \nabla\cdot(q\delta) = -s\right\} = \|s/q\|^2_{H^{-1}(q)}. \tag{108}$$

This recovers the Full action of a specified path q_t relative to u_t when $s_t = \partial_t q_t + \nabla\cdot(q_t u_t)$. The Poisson variational identity alone does not establish this source identity.

For example, let S be a fixed convolution with a probability kernel on $\mathbb{R}^d$ or a flat torus, and assume the derivatives below commute with convolution. If $s_t = \partial_t q_t + \nabla\cdot(q_t u_t)$, then

$$\partial_t(Sq_t) + \nabla\cdot((Sq_t)u_t) = Ss_t + \nabla\cdot\big((Sq_t)u_t - S(q_t u_t)\big). \tag{109}$$

Thus separately smoothing the density and source generally produces the represented cost $\mathfrak{a}(Sq_t, Ss_t)$, not automatically the Full action of Sq_t under the unchanged velocity u_t. Moreover, with S^* denoting the adjoint and the integrals understood coordinatewise,

$$\int \mathbf{\Phi}_{\boldsymbol{\eta}}(x)(Sq_t)(x)\,dx = \int (S^*\mathbf{\Phi}_{\boldsymbol{\eta}})(x) q_t(x)\,dx,$$

so smoothing need not preserve calibration against the original observables. The continuum theorems therefore do not by themselves establish consistency or convergence of an empirical or spatially regularized evaluator. Discrete minimum-norm and orthogonality identities concern the specified discrete source, operators, and inner product; identifying their value with the population Full action requires a separate consistency argument.

B.8 Weighted negative-Sobolev interpretation

Define the weighted negative-Sobolev norm

$$\|h\|_{H^{-1}(q)} := \sup_{\varphi \in H_q:\, \mathbb{E}_q[\|\nabla\varphi\|^2]\le 1} |\mathbb{E}_q[h\varphi]|. \tag{110}$$

Lemma 1 (Weighted negative-Sobolev identity)**.** *Under the coercivity and regularity assumptions of Appendix B.7,*

$$a_{\mathrm{full}}(t,\boldsymbol{\eta}) = \|h_{t,\boldsymbol{\eta}}\|^2_{H^{-1}(q_t^{\boldsymbol{\eta}})}. \tag{111}$$

Proof of Lemma 1. On H_q use the inner product

$$\langle \psi, \varphi \rangle_{H_q} := \mathbb{E}_q[\nabla\psi \cdot \nabla\varphi]. \tag{112}$$

Theorem 2 gives

$$\mathbb{E}_q[h\varphi] = -\langle \psi, \varphi \rangle_{H_q} \tag{113}$$

for every $\varphi \in H_q$. Hence, by Cauchy–Schwarz,

$$|\mathbb{E}_q[h\varphi]| \le \|\psi\|_{H_q} \|\varphi\|_{H_q}, \tag{114}$$

so $\|h\|_{H^{-1}(q)} \le \|\psi\|_{H_q}$. If $\psi \neq 0$, choosing $\varphi = -\psi/\|\psi\|_{H_q}$ attains equality; if $\psi = 0$, both sides vanish. Therefore

$$\|h\|_{H^{-1}(q)} = \|\psi\|_{H_q} = \left(\mathbb{E}_q[\|\nabla\psi\|^2]\right)^{1/2}. \tag{115}$$

Squaring and using the minimum-energy identity proved in Appendix B.7 gives equation 111, which is the second equality in equation 13 and completes the proof of Theorem 2. □

Corollary 6 (Sharp uniform certificate)**.** *Under the assumptions of Theorem 2,*

$$a_{\mathrm{full}}(t, \boldsymbol{\eta}) = \inf\left\{B \ge 0 : \ |\mathcal{E}_t(f)|^2 \le B\, \mathbb{E}_{Q_t}[\|\nabla f\|^2] \text{ for every } f \in H_{q_t}\right\}. \tag{116}$$

Hence Full is the smallest uniform constant controlling squared reference-relative expectation-rate discrepancy over the admissible Sobolev observable class.

Proof. By Theorem 2 and Lemma 1, $a_{\mathrm{full}}(t, \boldsymbol{\eta})$ is the squared operator norm of the linear functional $f \mapsto \mathcal{E}_t(f)$ on H_{q_t} equipped with the gradient-energy norm. The displayed infimum is precisely that squared norm. □

Proposition 6 (Galerkin lower bound, monotonicity, and convergence)**.** *Fix $(t, \boldsymbol{\eta})$ under the assumptions of Theorem 2, write $q = q_t^{\boldsymbol{\eta}}$ and $h = h_{t,\boldsymbol{\eta}}$, and let $V \subset H_q$ be a finite-dimensional linear subspace. Define*

$$a_V := \sup_{f \in V}\left\{2\mathbb{E}_q[hf] - \mathbb{E}_q[\|\nabla f\|^2]\right\}. \tag{117}$$

If $\psi = \psi_{t,\boldsymbol{\eta}}$ is the Full Poisson potential, then

$$a_V = a_{\mathrm{full}} - \mathrm{dist}^2_{H_q}(-\psi, V), \qquad 0 \le a_V \le a_{\mathrm{full}}. \tag{118}$$

Hence $V_1 \subseteq V_2$ implies $a_{V_1} \le a_{V_2} \le a_{\mathrm{full}}$. If $V_1 \subseteq V_2 \subseteq \cdots$ and $\overline{\bigcup_n V_n}^{H_q} = H_q$, then $a_{V_n} \uparrow a_{\mathrm{full}}$.

Proof. The weak Poisson identity gives $\mathbb{E}_q[hf] = -\langle \psi, f \rangle_{H_q}$, and therefore

$$2\mathbb{E}_q[hf] - \|f\|^2_{H_q} = \|\psi\|^2_{H_q} - \|f + \psi\|^2_{H_q}.$$

Since $\|\psi\|^2_{H_q} = a_{\mathrm{full}}$, taking the supremum over $f \in V$ gives equation 118. Enlarging V can only decrease the distance to $-\psi$, which proves monotonicity; density of the nested union makes this distance converge to zero. □

Remark 6 (Tangent as a Galerkin restriction)**.** *Fix $(t, \boldsymbol{\eta})$, write $r = r_{t,\boldsymbol{\eta}}$, $G = G_{t,\boldsymbol{\eta}}$, and $\overline{\phi} = \boldsymbol{\Phi}_{\boldsymbol{\eta}} - \boldsymbol{c}_{\boldsymbol{\eta}}(t)$, and assume that each coordinate of $\overline{\phi}$ belongs to H_q. Testing the reference-relative continuity identity gives $\mathbb{E}_q[h\overline{\phi}] = -r$, while $\|\beta^\top\overline{\phi}\|^2_{H_q} = \beta^\top G\beta$. Taking $V = \mathrm{span}\{\overline{\phi}_1, \ldots, \overline{\phi}_R\}$ in Proposition 6 therefore gives $a_V = r^\top G^\dagger r = a_{\tan}$ and*

$$a_{\mathrm{full}} - a_{\tan} = \mathrm{dist}^2_{H_q}(-\psi, V) = a_{\mathrm{hid}}.$$

Thus Tangent is the measured-observable Galerkin restriction of Full, while Proposition 6 applies to arbitrary finite potential dictionaries such as those used in the skyrmion evaluator. This is the same finite-feature energy-gap viewpoint underlying Sobolev discrepancy (Mroueh, 2018).

Corollary 7 (Screened negative-Sobolev interpretation)**.** *Define*

$$\|h\|_{H_\kappa^{-1}(q)} := \sup_{\varphi \in H_{q,\kappa}:\, \mathbb{E}_q[\|\nabla\varphi\|^2 + \kappa^{-1}\varphi^2] \le 1} |\mathbb{E}_q[h\varphi]|. \tag{119}$$

Under the conditions of Corollary 4,

$$a^{\mathrm{ub}}_{\mathrm{full}}(t, \boldsymbol{\eta}) = M_t \, \|h_{\mathrm{ub},t,\boldsymbol{\eta}}\|^2_{H_\kappa^{-1}(q_t^{\boldsymbol{\eta}})}. \tag{120}$$

Proof. Equation equation 85 states that the functional $\varphi \mapsto \mathbb{E}_q[h_{\mathrm{ub}}\varphi]$ is represented in the Hilbert space $H_{q,\kappa}$ by $\psi^\star$. Cauchy–Schwarz gives an upper bound by $\|\psi^\star\|_{H_{q,\kappa}}$, and choosing $\varphi = \psi^\star/\|\psi^\star\|_{H_{q,\kappa}}$ when $\psi^\star \neq 0$ attains equality. Corollary 4 then gives equation 120. □

B.9 MINIMUM-ENERGY TANGENT CORRECTION

For the minimum-energy Tangent problem, exact feasibility is equivalent to

$$r_{t,\boldsymbol{\eta}} \in \mathrm{Range}(G_{t,\boldsymbol{\eta}}). \tag{121}$$

Proposition 7 (General minimum-energy tangent correction)**.** *Assume $J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} \in L^2(Q_t^{\boldsymbol{\eta}}; \mathbb{R}^{R\times d})$. Exact moment-rate matching is feasible if and only if equation 121 holds. In that case the unique minimum-energy correction is*

$$\delta^{\tan}_{t,\boldsymbol{\eta}}(x) = -J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}}(x)^\top G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}}, \tag{122}$$

and

$$a_{\tan}(t,\boldsymbol{\eta}) = r^\top_{t,\boldsymbol{\eta}} G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}}. \tag{123}$$

If equation 121 fails, the exact moment-rate constraint is infeasible; the pseudoinverse then has only a least-squares interpretation.

Proof of Proposition 7. Fix $(t,\boldsymbol{\eta})$ and let

$$\mathcal{H} := L^2(Q_t^{\boldsymbol{\eta}}; \mathbb{R}^d) \tag{124}$$

with its usual inner product. Define the bounded linear operator

$$L : \mathcal{H} \to \mathbb{R}^R, \qquad L\delta := \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}}\delta]. \tag{125}$$

Boundedness follows from Cauchy–Schwarz and the assumption $J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} \in L^2(Q_t^{\boldsymbol{\eta}})$. Its adjoint is

$$L^* a = J^\top_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} a, \qquad a \in \mathbb{R}^R, \tag{126}$$

because

$$\langle L\delta, a\rangle_{\mathbb{R}^R} = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}\left[\delta^\top J^\top_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} a\right]. \tag{127}$$

Moreover,

$$LL^* a = G_{t,\boldsymbol{\eta}} a. \tag{128}$$

Since the codomain is finite dimensional,

$$\mathrm{Range}(L) = \mathrm{Range}(LL^*) = \mathrm{Range}(G_{t,\boldsymbol{\eta}}). \tag{129}$$

Thus the condition $r_{t,\boldsymbol{\eta}} \in \mathrm{Range}(G_{t,\boldsymbol{\eta}})$ is exactly the feasibility condition for $L\delta = -r_{t,\boldsymbol{\eta}}$.

The feasible set is a closed affine subspace of the Hilbert space $\mathcal{H}$, so it contains a unique element of minimum norm. This minimum-norm element is orthogonal to $\ker L$ and therefore belongs to

$$(\ker L)^\perp = \overline{\mathrm{Range}(L^*)}. \tag{130}$$

Because L^* has finite-dimensional domain, its range is closed, so the minimizer has the form $\delta = L^* a$. The constraint becomes

$$G_{t,\boldsymbol{\eta}} a = -r_{t,\boldsymbol{\eta}}. \tag{131}$$

The minimum-norm coefficient may be chosen as

$$a = -G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}}, \tag{132}$$

which gives equation 122. Its squared norm is

$$\begin{aligned}\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\|\delta^{\tan}_{t,\boldsymbol{\eta}}\|^2] &= (G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}})^\top G_{t,\boldsymbol{\eta}} (G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}}) \\ &= r^\top_{t,\boldsymbol{\eta}} G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}},\end{aligned} \tag{133}$$

where the last equality uses symmetry of $G_{t,\boldsymbol{\eta}}$ and $r_{t,\boldsymbol{\eta}} \in \mathrm{Range}(G_{t,\boldsymbol{\eta}})$. This proves equation 123. If $r_{t,\boldsymbol{\eta}} \notin \mathrm{Range}(G_{t,\boldsymbol{\eta}}) = \mathrm{Range}(L)$, no exact solution exists, establishing the remaining feasibility statement. □

Corollary 8 (Minimum-energy unbalanced tangent correction)**.** *Let $\bar{c}_t = \int \phi \, d\mu_t$ denote the raw finite-measure moment vector, and assume $\phi \in L^2(q_t; \mathbb{R}^R)$ and $J_\phi \in L^2(q_t; \mathbb{R}^{R\times d})$. Define*

$$r^{\mathrm{ub}}_{t,\boldsymbol{\eta}} := \mathbb{E}_{q_t}[J_\phi u_t + g_{\mathrm{ref}}\phi] - \frac{\dot{\bar{c}}_t}{M_t} \tag{134}$$

and

$$G^{\mathrm{ub}}_{t,\boldsymbol{\eta}} := \mathbb{E}_{q_t}\left[J_\phi J_\phi^\top + \frac{1}{\kappa}\phi\phi^\top\right]. \tag{135}$$

If $r^{\mathrm{ub}}_{t,\boldsymbol{\eta}} \in \mathrm{Range}(G^{\mathrm{ub}}_{t,\boldsymbol{\eta}})$, then the minimum of

$$M_t \mathbb{E}_{q_t}[\|\delta\|^2 + \kappa\alpha^2] \tag{136}$$

subject to exact raw-moment-rate matching

$$\mathbb{E}_{q_t}[J_\phi\delta + \phi\alpha] = -r^{\mathrm{ub}}_{t,\boldsymbol{\eta}} \tag{137}$$

is attained uniquely by

$$\delta^{\tan}_{\mathrm{ub}} = -J_\phi^\top (G^{\mathrm{ub}})^\dagger r^{\mathrm{ub}}, \qquad \alpha^{\tan}_{\mathrm{ub}} = -\frac{1}{\kappa}\phi^\top (G^{\mathrm{ub}})^\dagger r^{\mathrm{ub}}, \tag{138}$$

and

$$a^{\mathrm{ub}}_{\tan}(t,\boldsymbol{\eta}) = M_t (r^{\mathrm{ub}}_{t,\boldsymbol{\eta}})^\top (G^{\mathrm{ub}}_{t,\boldsymbol{\eta}})^\dagger r^{\mathrm{ub}}_{t,\boldsymbol{\eta}}. \tag{139}$$

If $r^{\mathrm{ub}} \notin \mathrm{Range}(G^{\mathrm{ub}})$, exact raw-moment-rate matching is infeasible.

Proof. For any sufficiently regular finite-measure path satisfying

$$\partial_t \mu + \nabla\cdot(\mu v) = \mu g, \tag{140}$$

integration by parts gives

$$\frac{d}{dt}\int \phi \, d\mu = \int (J_\phi v + \phi g)\, d\mu. \tag{141}$$

Applying this with $v = u + \delta$, $g = g_{\mathrm{ref}} + \alpha$, and $\mu = Mq$ yields equation 137. Now equip

$$\mathcal{Z} := L^2(q; \mathbb{R}^d) \times L^2(q) \tag{142}$$

with

$$\langle(\delta,\alpha),(\delta',\alpha')\rangle_{\mathcal{Z}} := \mathbb{E}_q[\delta\cdot\delta' + \kappa\alpha\alpha']. \tag{143}$$

Define $L_{\mathrm{ub}} : \mathcal{Z} \to \mathbb{R}^R$ by $L_{\mathrm{ub}}(\delta,\alpha) = \mathbb{E}_q[J_\phi\delta + \phi\alpha]$. Its adjoint is

$$L^*_{\mathrm{ub}} a = \left(J_\phi^\top a, \frac{1}{\kappa}\phi^\top a\right), \tag{144}$$

since

$$\langle(\delta,\alpha), L^*_{\mathrm{ub}} a\rangle_{\mathcal{Z}} = a^\top \mathbb{E}_q[J_\phi\delta + \phi\alpha]. \tag{145}$$

Consequently,

$$L_{\mathrm{ub}} L^*_{\mathrm{ub}} a = G^{\mathrm{ub}} a. \tag{146}$$

The Hilbert-space minimum-norm argument from Proposition 7 therefore applies verbatim, giving $(\delta^{\tan}_{\mathrm{ub}}, \alpha^{\tan}_{\mathrm{ub}}) = -L^*_{\mathrm{ub}}(G^{\mathrm{ub}})^\dagger r^{\mathrm{ub}}$ and squared product norm $(r^{\mathrm{ub}})^\top (G^{\mathrm{ub}})^\dagger r^{\mathrm{ub}}$. Multiplication by M_t gives equation 139. □

Remark 7 (Making total-mass dynamics tangent-visible)**.** *If the raw tangent observable vector is augmented by the constant coordinate $\phi_0 \equiv 1$, then its transport Jacobian vanishes and its tangent equation reduces to the mass-rate condition. The corresponding Gram entry is κ^{-1}, so total-mass mismatch is represented entirely through reaction. This augmentation is appropriate for the tangent criterion even though, by Remark 1, the constant coordinate should not be inserted into the normalized-shape information-projection covariance.*

B.10 FULL–TANGENT–HIDDEN DECOMPOSITION

Fix $(t, \boldsymbol{\eta})$ and assume the hypotheses of Theorem 2. Assume additionally that $J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} \in L^2(Q_t^{\boldsymbol{\eta}})$, $J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} u_t \in L^1(Q_t^{\boldsymbol{\eta}})$, and $\boldsymbol{\Phi}_{\boldsymbol{\eta},j} - \boldsymbol{c}_{\boldsymbol{\eta},j}(t) \in H_{q_t^{\boldsymbol{\eta}}}$ for every observable coordinate. We further assume that differentiation of the calibrated moments and weak testing with these coordinates are justified and that the relevant boundary terms vanish. These conditions ensure that the Full correction satisfies the Tangent moment-rate constraints and that the Hilbert-space projections below are well defined.

The pointwise identities underlying the balanced decomposition are

$$\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}\left[J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} \delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}}\right] = 0, \tag{147}$$

and

$$\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}\left[\delta_{t,\boldsymbol{\eta}}^{\tan} \cdot \delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}}\right] = 0. \tag{148}$$

The hidden action density is

$$a_{\mathrm{hid}}(t, \boldsymbol{\eta}) := \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}\left[\left\|\delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}}\right\|^2\right] \geq 0. \tag{149}$$

For measurable action densities, define $\mathcal{A}_{\mathrm{hid}}(\boldsymbol{\eta}) := \int_0^1 a_{\mathrm{hid}}(t, \boldsymbol{\eta})\, \rho(dt)$. The design-level counterpart of equation 15 is

$$\mathcal{A}(\boldsymbol{\eta}) = \mathcal{A}_{\tan}(\boldsymbol{\eta}) + \mathcal{A}_{\mathrm{hid}}(\boldsymbol{\eta}), \qquad \mathcal{A}_{\tan}(\boldsymbol{\eta}) \leq \mathcal{A}(\boldsymbol{\eta}). \tag{150}$$

Proof of Proposition 2. Fix $(t, \boldsymbol{\eta})$ and write $Q = Q_t^{\boldsymbol{\eta}}$. Let

$$L\delta := \mathbb{E}_Q[J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} \delta] \tag{151}$$

be the moment-rate operator introduced implicitly in the proof of Proposition 7. Let $\delta_t^\star$ denote the minimum-energy Full correction. Because $u_t + \delta_t^\star$ transports the projected law and the calibration identity gives $\mathbb{E}_Q[\boldsymbol{\Phi}_{\boldsymbol{\eta}}] = \boldsymbol{c}_{\boldsymbol{\eta}}(t)$, the moment-rate identity equation 20 yields

$$L\delta_t^\star = -\left(\mathbb{E}_Q[J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} u_t] - \dot{\boldsymbol{c}}_{\boldsymbol{\eta}}(t)\right) = -r_{t,\boldsymbol{\eta}}. \tag{152}$$

Hence the tangent constraint is feasible. By Proposition 7, its unique minimum-norm solution is $\delta_{t,\boldsymbol{\eta}}^{\tan}$.

Define

$$\delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}} := \delta_t^\star - \delta_{t,\boldsymbol{\eta}}^{\tan}. \tag{153}$$

Both $\delta_t^\star$ and $\delta_{t,\boldsymbol{\eta}}^{\tan}$ satisfy $L\delta = -r_{t,\boldsymbol{\eta}}$, so

$$L\delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}} = 0. \tag{154}$$

Thus $\delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}} \in \ker L$, which proves equation 147.

The proof of Proposition 7 also shows that the minimum-norm tangent solution belongs to $(\ker L)^\perp$. Therefore

$$\mathbb{E}_Q\left[\delta_{t,\boldsymbol{\eta}}^{\tan} \cdot \delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}}\right] = 0, \tag{155}$$

which is equation 148. Applying the Pythagorean identity in $L^2(Q; \mathbb{R}^d)$ gives

$$\begin{aligned} a_{\mathrm{full}}(t, \boldsymbol{\eta}) &= \mathbb{E}_Q[\|\delta_t^\star\|^2] \\ &= \mathbb{E}_Q\left[\left\|\delta_{t,\boldsymbol{\eta}}^{\tan} + \delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}}\right\|^2\right] \\ &= \mathbb{E}_Q[\left\|\delta_{t,\boldsymbol{\eta}}^{\tan}\right\|^2] + \mathbb{E}_Q[\left\|\delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}}\right\|^2] \\ &= a_{\tan}(t, \boldsymbol{\eta}) + a_{\mathrm{hid}}(t, \boldsymbol{\eta}), \end{aligned} \tag{156}$$

proving equation 15. Integrating the pointwise identity against the common nonnegative measure ρ gives equation 150 whenever the terms are measurable; the stated pointwise and integrated inequalities follow from nonnegativity of the hidden action. At fixed t, equality $a_{\tan}(t, \boldsymbol{\eta}) = a_{\mathrm{full}}(t, \boldsymbol{\eta})$ holds if and only if $\delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}} = 0$ in $L^2(Q; \mathbb{R}^d)$, equivalently $\delta_t^\star = \delta_{t,\boldsymbol{\eta}}^{\tan}$ Q-almost everywhere. For the integrated actions, if $\mathcal{A}(\boldsymbol{\eta}) < \infty$, then $\mathcal{A}_{\tan}(\boldsymbol{\eta}) = \mathcal{A}(\boldsymbol{\eta})$ holds if and only if $a_{\mathrm{hid}}(t, \boldsymbol{\eta}) = 0$ for ρ-almost every t, equivalently $\delta_{t,\boldsymbol{\eta}}^{\mathrm{hid}} = 0$ in $L^2(Q_t^{\boldsymbol{\eta}}; \mathbb{R}^d)$ for ρ-almost every t. The decomposition remains valid for measurable nonnegative extended-valued actions, but equality of two infinite actions does not imply that the hidden action vanishes. □

Corollary 9 (Unbalanced Full–Tangent–Hidden decomposition)**.** *Under the assumptions of Corollaries 4 and 8, assume additionally that each coordinate of $\phi = \mathbf{\Phi}_{\boldsymbol{\eta}}$ belongs to $H_{q_t,\kappa}$, that $J_\phi u_t \in L^1(q_t;\mathbb{R}^R)$, and that raw-moment differentiation and weak testing with these coordinates are justified, with vanishing boundary terms. Let*

$$z^\star_{\mathrm{ub}} := (\delta^\star, \alpha^\star), \qquad z^{\tan}_{\mathrm{ub}} := (\delta^{\tan}_{\mathrm{ub}}, \alpha^{\tan}_{\mathrm{ub}}), \qquad z^{\mathrm{hid}}_{\mathrm{ub}} := z^\star_{\mathrm{ub}} - z^{\tan}_{\mathrm{ub}}. \tag{157}$$

Then

$$L_{\mathrm{ub}} z^{\mathrm{hid}}_{\mathrm{ub}} = 0 \tag{158}$$

and

$$\mathbb{E}_q\big[\delta^{\tan}_{\mathrm{ub}} \cdot \delta^{\mathrm{hid}}_{\mathrm{ub}} + \kappa \alpha^{\tan}_{\mathrm{ub}} \alpha^{\mathrm{hid}}_{\mathrm{ub}}\big] = 0. \tag{159}$$

Defining

$$a^{\mathrm{ub}}_{\mathrm{hid}}(t,\boldsymbol{\eta}) := M_t \mathbb{E}_q\big[\|\delta^{\mathrm{hid}}_{\mathrm{ub}}\|^2 + \kappa(\alpha^{\mathrm{hid}}_{\mathrm{ub}})^2\big], \tag{160}$$

one has

$$a^{\mathrm{ub}}_{\mathrm{full}}(t,\boldsymbol{\eta}) = a^{\mathrm{ub}}_{\tan}(t,\boldsymbol{\eta}) + a^{\mathrm{ub}}_{\mathrm{hid}}(t,\boldsymbol{\eta}), \qquad a^{\mathrm{ub}}_{\tan}(t,\boldsymbol{\eta}) \le a^{\mathrm{ub}}_{\mathrm{full}}(t,\boldsymbol{\eta}). \tag{161}$$

Whenever the three terms are integrable against the common nonnegative measure ρ, the same decomposition and inequality hold after integration in time.

Proof. The Full correction satisfies the complete finite-measure continuity equation, so the raw moment identity equation 141 implies that $z^\star_{\mathrm{ub}}$ satisfies the tangent constraint. The tangent correction is the unique minimum-norm point of that affine constraint in the product Hilbert space equation 143; therefore $z^{\tan}_{\mathrm{ub}} \in (\ker L_{\mathrm{ub}})^\perp$. Since both $z^\star_{\mathrm{ub}}$ and $z^{\tan}_{\mathrm{ub}}$ have the same image under L_{ub}, their difference lies in $\ker L_{\mathrm{ub}}$. This proves equation 158 and equation 159. Pythagoras in the product Hilbert space, followed by multiplication by M_t, gives equation 161; integration against ρ preserves the identity and inequality. □

When $0 < \mathcal{A}(\boldsymbol{\eta}) < \infty$, we summarize the relative size of the tangent-invisible component by the hidden-action fraction

$$\Gamma(\boldsymbol{\eta}) := \frac{\mathcal{A}_{\mathrm{hid}}(\boldsymbol{\eta})}{\mathcal{A}(\boldsymbol{\eta})} = 1 - \frac{\mathcal{A}_{\tan}(\boldsymbol{\eta})}{\mathcal{A}(\boldsymbol{\eta})} \in [0,1]. \tag{162}$$

Thus $\Gamma(\boldsymbol{\eta})$ near one indicates that most of the minimum law-level correction lies outside the directions visible through the measured moment rates. In the discretized experimental audits we denote the corresponding common-raster quantity by Γ_h.

B.11 Exact condition for Tangent–Full equality

Corollary 10 (When Tangent and Full corrections coincide)**.** *Under the conditions of Proposition 2, suppose additionally that $Q^{\boldsymbol{\eta}}_t$ has a density positive almost everywhere on a connected open domain Ω, and that $\psi_{t,\boldsymbol{\eta}}$ and the coordinates of $\mathbf{\Phi}_{\boldsymbol{\eta}}$ belong to $H^1_{\mathrm{loc}}(\Omega)$. Define*

$$\beta_{t,\boldsymbol{\eta}} := G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}}. \tag{163}$$

Then

$$a_{\tan}(t,\boldsymbol{\eta}) = a_{\mathrm{full}}(t,\boldsymbol{\eta}) \tag{164}$$

if and only if

$$\psi_{t,\boldsymbol{\eta}}(x) = \beta^\top_{t,\boldsymbol{\eta}}(\mathbf{\Phi}_{\boldsymbol{\eta}}(x) - \boldsymbol{c}_{\boldsymbol{\eta}}(t)) \qquad Q^{\boldsymbol{\eta}}_t\text{-a.e.} \tag{165}$$

Proof of Corollary 10. Let

$$\beta_{t,\boldsymbol{\eta}} = G^\dagger_{t,\boldsymbol{\eta}} r_{t,\boldsymbol{\eta}}. \tag{166}$$

By Proposition 7,

$$\delta^{\tan}_{t,\boldsymbol{\eta}} = -J^\top_{\mathbf{\Phi}_{\boldsymbol{\eta}}} \beta_{t,\boldsymbol{\eta}} = -\nabla\big(\beta^\top_{t,\boldsymbol{\eta}} \mathbf{\Phi}_{\boldsymbol{\eta}}\big) \tag{167}$$

where the final identity is understood in the weak sense. Theorem 2 gives

$$\delta^\star_t = -\nabla \psi_{t,\boldsymbol{\eta}}. \tag{168}$$

By Proposition 2, equality of tangent and full action is equivalent to $\delta_t^\star = \delta_{t,\boldsymbol{\eta}}^{\tan}$ $Q_t^{\boldsymbol{\eta}}$-almost everywhere. Because the density of $Q_t^{\boldsymbol{\eta}}$ is positive almost everywhere on Ω, this is equivalent to

$$\nabla\left(\psi_{t,\boldsymbol{\eta}} - \beta_{t,\boldsymbol{\eta}}^\top \boldsymbol{\Phi}_{\boldsymbol{\eta}}\right) = 0 \tag{169}$$

almost everywhere on Ω. A Sobolev function with zero weak gradient on a connected open domain is almost everywhere constant. Hence

$$\psi_{t,\boldsymbol{\eta}}(x) = \beta_{t,\boldsymbol{\eta}}^\top \boldsymbol{\Phi}_{\boldsymbol{\eta}}(x) + C \tag{170}$$

for some constant C. The mean-zero gauge for the Poisson potential and the calibration identity imply

$$0 = \mathbb{E}_{Q_t^\eta}[\psi_{t,\boldsymbol{\eta}}] = \beta_{t,\boldsymbol{\eta}}^\top \boldsymbol{c}_{\boldsymbol{\eta}}(t) + C, \tag{171}$$

so $C = -\beta_{t,\boldsymbol{\eta}}^\top \boldsymbol{c}_{\boldsymbol{\eta}}(t)$, giving equation 165.

Conversely, if equation 165 holds, then

$$-\nabla\psi_{t,\boldsymbol{\eta}} = -J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}}^\top \beta_{t,\boldsymbol{\eta}} = \delta_{t,\boldsymbol{\eta}}^{\tan} \tag{172}$$

$Q_t^{\boldsymbol{\eta}}$-almost everywhere. Thus the hidden component vanishes, and Proposition 2 gives $a_{\mathrm{full}}(t,\boldsymbol{\eta}) = a_{\tan}(t,\boldsymbol{\eta})$. □

Corollary 11 (Exact unbalanced tangent–Full equality condition)**.** *Under the assumptions of Corollary 9, assume that the screened-Poisson potential ψ and the coordinates of ϕ belong to $H_{q,\kappa}$, and define*

$$\beta^{\mathrm{ub}} := (G^{\mathrm{ub}})^\dagger r^{\mathrm{ub}}. \tag{173}$$

Then

$$a_{\tan}^{\mathrm{ub}}(t,\boldsymbol{\eta}) = a_{\mathrm{full}}^{\mathrm{ub}}(t,\boldsymbol{\eta}) \tag{174}$$

if and only if

$$\psi = -(\beta^{\mathrm{ub}})^\top \phi \qquad \text{in } H_{q,\kappa}, \tag{175}$$

equivalently, the screened Full potential belongs to the uncentered linear span of the tangent observables with the coefficient fixed by the tangent normal equations.

Proof. Corollary 9 gives equality if and only if the Full and Tangent correction pairs coincide. By Corollaries 4 and 8,

$$z_{\mathrm{ub}}^\star = \left(\nabla\psi, \frac{\psi}{\kappa}\right), \qquad z_{\mathrm{ub}}^{\tan} = \left(-J_\phi^\top \beta^{\mathrm{ub}}, -\frac{1}{\kappa}\phi^\top \beta^{\mathrm{ub}}\right). \tag{176}$$

Equality of the reaction components gives $\psi = -(\beta^{\mathrm{ub}})^\top \phi$ in $L^2(q)$; because both sides lie in $H_{q,\kappa}$, the corresponding weak gradients then agree as well, giving equality of the transport components. Conversely, equation 175 implies equality of both components and hence zero hidden action. Unlike the balanced case, no additive constant remains undetermined because the reaction component contains the potential itself. □

B.12 Differentiability with respect to the measurement design

For continuous designs, the numerical objective depends on the design through the calibrated information projection, its physical-time multiplier derivative, the resulting continuity forcing, and finally the finite-dimensional Full correction solve. We first record the sensitivity of the calibration multiplier and then use it to establish local differentiability of the complete fixed numerical objective.

Define the calibration residual

$$F(\boldsymbol{\lambda}, \boldsymbol{\eta}, t) := \mathbb{E}_{q_{\boldsymbol{\lambda},\boldsymbol{\eta},t}}[\boldsymbol{\Phi}_{\boldsymbol{\eta}}] - \boldsymbol{c}_{\boldsymbol{\eta}}(t), \tag{177}$$

where

$$q_{\boldsymbol{\lambda},\boldsymbol{\eta},t}(x) \propto \widetilde{q}_t(x) \exp\left(\boldsymbol{\lambda}^\top \boldsymbol{\Phi}_{\boldsymbol{\eta}}(x)\right). \tag{178}$$

Proposition 8 (Design sensitivity of the information projection)**.** *Fix t and express the design in local coordinates on an open set. Suppose F is continuously differentiable in $(\boldsymbol{\lambda}, \boldsymbol{\eta})$ on an open neighborhood of a calibrated point satisfying $F(\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}, \boldsymbol{\eta}, t) = 0$. Assume that the observable and target maps are continuously differentiable in the design coordinates and that the differentiations of the tilted expectations below can be interchanged with integration throughout this neighborhood. If $C_{t,\boldsymbol{\eta}} = \partial_{\boldsymbol{\lambda}} F(\boldsymbol{\lambda}_{t,\boldsymbol{\eta}}, \boldsymbol{\eta}, t)$ is positive definite, then the calibrated multiplier defines a locally unique continuously differentiable map of the design. For every design perturbation ξ in these coordinates,*

$$D_{\boldsymbol{\eta}} \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}[\xi] = -C_{t,\boldsymbol{\eta}}^{-1} D_{\boldsymbol{\eta}} F[\xi], \tag{179}$$

where

$$\begin{aligned} D_{\boldsymbol{\eta}} F[\xi] = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[D_{\boldsymbol{\eta}} \boldsymbol{\Phi}_{\boldsymbol{\eta}}[\xi]] + \mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}\left(\boldsymbol{\Phi}_{\boldsymbol{\eta}}, \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}^{\top} D_{\boldsymbol{\eta}} \boldsymbol{\Phi}_{\boldsymbol{\eta}}[\xi]\right) \\ - D_{\boldsymbol{\eta}} \boldsymbol{c}_{\boldsymbol{\eta}}(t)[\xi]. \end{aligned} \tag{180}$$

Proof of Proposition 8. Fix t and abbreviate

$$q_{\lambda,\boldsymbol{\eta}}(x) = \widetilde{q}_t(x) \exp\left(\lambda^{\top} \boldsymbol{\Phi}_{\boldsymbol{\eta}}(x) - A_t(\lambda, \boldsymbol{\eta})\right), \tag{181}$$

where

$$A_t(\lambda, \boldsymbol{\eta}) = \log \mathbb{E}_{\widetilde{Q}_t}\left[\exp(\lambda^{\top} \boldsymbol{\Phi}_{\boldsymbol{\eta}}(X))\right]. \tag{182}$$

The reference law is frozen with respect to $\boldsymbol{\eta}$. At a calibrated solution,

$$F(\lambda_{t,\boldsymbol{\eta}}, \boldsymbol{\eta}, t) = 0. \tag{183}$$

Differentiation with respect to λ gives

$$\partial_{\lambda} F = \mathrm{Cov}_{q_{\lambda,\boldsymbol{\eta}}}(\boldsymbol{\Phi}_{\boldsymbol{\eta}}, \boldsymbol{\Phi}_{\boldsymbol{\eta}}), \tag{184}$$

so at the calibrated point

$$\partial_{\lambda} F = C_{t,\boldsymbol{\eta}}. \tag{185}$$

By assumption this matrix is positive definite and hence invertible. The implicit-function theorem therefore gives a locally differentiable multiplier map $\boldsymbol{\eta} \mapsto \lambda_{t,\boldsymbol{\eta}}$ and, for any perturbation direction ξ,

$$D_{\boldsymbol{\eta}} \lambda_{t,\boldsymbol{\eta}}[\xi] = -C_{t,\boldsymbol{\eta}}^{-1} D_{\boldsymbol{\eta}} F[\xi]. \tag{186}$$

It remains to compute the directional derivative of F while holding λ fixed.

For any scalar or vector test function $f_{\boldsymbol{\eta}}$, differentiation of a normalized exponential tilt gives the identity

$$D_{\boldsymbol{\eta}} \mathbb{E}_{q_{\lambda,\boldsymbol{\eta}}}[f_{\boldsymbol{\eta}}][\xi] = \mathbb{E}_{q_{\lambda,\boldsymbol{\eta}}}[D_{\boldsymbol{\eta}} f_{\boldsymbol{\eta}}[\xi]] + \mathrm{Cov}_{q_{\lambda,\boldsymbol{\eta}}}\left(f_{\boldsymbol{\eta}}, \lambda^{\top} D_{\boldsymbol{\eta}} \boldsymbol{\Phi}_{\boldsymbol{\eta}}[\xi]\right), \tag{187}$$

provided differentiation and integration can be interchanged. Applying equation 187 with $f_{\boldsymbol{\eta}} = \boldsymbol{\Phi}_{\boldsymbol{\eta}}$ and subtracting $D_{\boldsymbol{\eta}} \boldsymbol{c}_{\boldsymbol{\eta}}(t)[\xi]$ gives

$$\begin{aligned} D_{\boldsymbol{\eta}} F[\xi] = \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[D_{\boldsymbol{\eta}} \boldsymbol{\Phi}_{\boldsymbol{\eta}}[\xi]] + \mathrm{Cov}_{Q_t^{\boldsymbol{\eta}}}\left(\boldsymbol{\Phi}_{\boldsymbol{\eta}}, \boldsymbol{\lambda}_{t,\boldsymbol{\eta}}^{\top} D_{\boldsymbol{\eta}} \boldsymbol{\Phi}_{\boldsymbol{\eta}}[\xi]\right) \\ - D_{\boldsymbol{\eta}} \boldsymbol{c}_{\boldsymbol{\eta}}(t)[\xi], \end{aligned} \tag{188}$$

which is equation 180. Substitution into equation 186 proves equation 179. □

Proposition 9 (End-to-end differentiability of the finite-dimensional Full objective)**.** *Let the continuous design $\boldsymbol{\eta}$ vary in local coordinates on an open set, and fix numerical time nodes $t_1, \ldots, t_M$ with nonnegative quadrature weights $\omega_1, \ldots, \omega_M$. At each t_j, suppose the assumptions of Proposition 8 hold, and write the physical-time multiplier equation from equation 8 as*

$$C_{j,\boldsymbol{\eta}} \dot{\boldsymbol{\lambda}}_{j,\boldsymbol{\eta}} = g_{j,\boldsymbol{\eta}}, \qquad g_{j,\boldsymbol{\eta}} := \dot{\boldsymbol{c}}_{\boldsymbol{\eta}}(t_j) - \mathbb{E}_{Q_{t_j}^{\boldsymbol{\eta}}}[m_{j,\boldsymbol{\eta}}] - \mathrm{Cov}_{Q_{t_j}^{\boldsymbol{\eta}}}\left(\boldsymbol{\Phi}_{\boldsymbol{\eta}}, \boldsymbol{\lambda}_{j,\boldsymbol{\eta}}^{\top} m_{j,\boldsymbol{\eta}}\right). \tag{189}$$

Assume that $C_{j,\boldsymbol{\eta}}$ and $g_{j,\boldsymbol{\eta}}$ are continuously differentiable in $\boldsymbol{\eta}$ and that $C_{j,\boldsymbol{\eta}}$ remains nonsingular. Suppose further that, after the fixed gauge choice or basis restriction used by the numerical evaluator, the Full correction at t_j is represented by a reduced linear system

$$\mathsf{K}_{j,\boldsymbol{\eta}} \beta_{j,\boldsymbol{\eta}} = b_{j,\boldsymbol{\eta}}, \tag{190}$$

where $\mathsf{K}_{j,\boldsymbol{\eta}}$ is symmetric positive definite and both $\mathsf{K}_{j,\boldsymbol{\eta}}$ and $b_{j,\boldsymbol{\eta}}$ are continuously differentiable in $\boldsymbol{\eta}$. Define the corresponding numerical action by

$$\widehat{a}_{\mathrm{full}}(t_j,\boldsymbol{\eta}) := \beta_{j,\boldsymbol{\eta}}^{\top}\mathsf{K}_{j,\boldsymbol{\eta}}\beta_{j,\boldsymbol{\eta}} = b_{j,\boldsymbol{\eta}}^{\top}\mathsf{K}_{j,\boldsymbol{\eta}}^{-1}b_{j,\boldsymbol{\eta}}, \qquad \widehat{\mathcal{A}}(\boldsymbol{\eta}) := \sum_{j=1}^{M}\omega_j\widehat{a}_{\mathrm{full}}(t_j,\boldsymbol{\eta}). \tag{191}$$

Then $\boldsymbol{\lambda}_{j,\boldsymbol{\eta}}$, $\dot{\boldsymbol{\lambda}}_{j,\boldsymbol{\eta}}$, $\beta_{j,\boldsymbol{\eta}}$, $\widehat{a}_{\mathrm{full}}(t_j,\boldsymbol{\eta})$, and $\widehat{\mathcal{A}}(\boldsymbol{\eta})$ are locally continuously differentiable in $\boldsymbol{\eta}$. For every design perturbation ξ,

$$D_{\boldsymbol{\eta}}\dot{\boldsymbol{\lambda}}_{j,\boldsymbol{\eta}}[\xi] = C_{j,\boldsymbol{\eta}}^{-1}\Big(D_{\boldsymbol{\eta}}g_{j,\boldsymbol{\eta}}[\xi] - D_{\boldsymbol{\eta}}C_{j,\boldsymbol{\eta}}[\xi]\dot{\boldsymbol{\lambda}}_{j,\boldsymbol{\eta}}\Big), \tag{192}$$

and the instantaneous action derivative is

$$D_{\boldsymbol{\eta}}\widehat{a}_{\mathrm{full}}(t_j,\boldsymbol{\eta})[\xi] = 2\beta_{j,\boldsymbol{\eta}}^{\top}D_{\boldsymbol{\eta}}b_{j,\boldsymbol{\eta}}[\xi] - \beta_{j,\boldsymbol{\eta}}^{\top}D_{\boldsymbol{\eta}}\mathsf{K}_{j,\boldsymbol{\eta}}[\xi]\beta_{j,\boldsymbol{\eta}}. \tag{193}$$

Consequently,

$$D_{\boldsymbol{\eta}}\widehat{\mathcal{A}}(\boldsymbol{\eta})[\xi] = \sum_{j=1}^{M}\omega_j D_{\boldsymbol{\eta}}\widehat{a}_{\mathrm{full}}(t_j,\boldsymbol{\eta})[\xi]. \tag{194}$$

Proof. Proposition 8 gives local continuous differentiability of $\boldsymbol{\lambda}_{j,\boldsymbol{\eta}}$. Since $C_{j,\boldsymbol{\eta}}$ remains nonsingular, equation 189 gives $\dot{\boldsymbol{\lambda}}_{j,\boldsymbol{\eta}} = C_{j,\boldsymbol{\eta}}^{-1}g_{j,\boldsymbol{\eta}}$, which is locally continuously differentiable, and differentiation of this identity yields equation 192. The assumed smooth dependence of the remaining numerical quantities therefore makes $\mathsf{K}_{j,\boldsymbol{\eta}}$ and $b_{j,\boldsymbol{\eta}}$ continuously differentiable functions of the design. Positive definiteness of $\mathsf{K}_{j,\boldsymbol{\eta}}$ gives $\beta_{j,\boldsymbol{\eta}} = \mathsf{K}_{j,\boldsymbol{\eta}}^{-1}b_{j,\boldsymbol{\eta}}$ and hence local continuous differentiability of the reduced correction. Finally,

$$\widehat{a}_{\mathrm{full}} = b^{\top}\mathsf{K}^{-1}b.$$

Using $D(\mathsf{K}^{-1})[\xi] = -\mathsf{K}^{-1}D\mathsf{K}[\xi]\mathsf{K}^{-1}$ and $\beta = \mathsf{K}^{-1}b$ gives

$$D\widehat{a}_{\mathrm{full}}[\xi] = 2\beta^{\top}Db[\xi] - \beta^{\top}D\mathsf{K}[\xi]\beta,$$

which proves equation 193. The finite weighted sum in equation 191 then gives equation 194. □

Remark 8 (Scope of the differentiability claim)**.** *Proposition 9 concerns the fixed finite-dimensional objective differentiated during continuous-design optimization. It does not assert differentiability of the continuum action under arbitrary perturbations of the weighted Sobolev problem, nor of discrete decisions such as calibration-feasibility rejection, support-certification thresholds, changes of numerical basis or gauge, or membership in the risk-constrained admissible set. Within a fixed numerical representation and away from such transitions, the proposition gives the end-to-end derivative through calibration, physical-time multiplier dynamics, and the Full correction solve.*

Corollary 12 (Design sensitivity in the finite-measure unbalanced extension)**.** *Suppose the normalized-shape information projection in Corollary 1 is calibrated to*

$$c_{\boldsymbol{\eta}}(t) = \frac{\bar{c}_{\boldsymbol{\eta}}(t)}{M_{\boldsymbol{\eta}}(t)}, \tag{195}$$

and suppose the assumptions of Proposition 8 hold. Then

$$D_{\boldsymbol{\eta}}\lambda_{t,\boldsymbol{\eta}}[\xi] = -C_{t,\boldsymbol{\eta}}^{-1}D_{\boldsymbol{\eta}}F[\xi] \tag{196}$$

with the same expression equation 180, but with

$$D_{\boldsymbol{\eta}}c_{\boldsymbol{\eta}}(t)[\xi] = \frac{1}{M_{\boldsymbol{\eta}}(t)}D_{\boldsymbol{\eta}}\bar{c}_{\boldsymbol{\eta}}(t)[\xi] - \frac{c_{\boldsymbol{\eta}}(t)}{M_{\boldsymbol{\eta}}(t)}D_{\boldsymbol{\eta}}M_{\boldsymbol{\eta}}(t)[\xi]. \tag{197}$$

In particular, if the mass schedule is fixed with respect to the measurement design, only the first term in equation 197 remains.

C Analytical Comparisons and Exact Information-Projected Examples

The criterion distinctions summarized in Section 4 are developed here at the operator level and through exact information-projected examples. We first collect the comparator definitions used in the experiments, then interpret Full through the inverse weighted transport operator, characterize the Tangent-invisible component, compare residual L^2 with Full through the transport spectrum, and establish separations from forward KL and static W_2. The final subsection proves the matched-initial stability result used in the main text. The analytical comparisons need not assume equal scientific risk. Unless stated otherwise, $\mathbb{T} = \mathbb{R}/(2\pi\mathbb{Z})$, spatial boundary conditions are periodic, Wasserstein distance uses the geodesic ground metric, and integrated comparisons use $\rho(dt) = dt$ on $[0, 1]$.

C.1 Comparator definitions

In the experiments, all comparator selectors use the same information-projected path $Q_t^{\boldsymbol{\eta}}$ and the same admissible set $\mathcal{D}_\varepsilon$ as Full; only the secondary score changes. With $J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}}$ denoting the observable Jacobian, we define

$$a_{\tan}(t, \boldsymbol{\eta}) := \inf_{\delta \in L^2(Q_t^{\boldsymbol{\eta}}; \mathbb{R}^d)} \left\{ \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[\|\delta\|^2] : \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}}(u_t + \delta)] = \dot{\boldsymbol{c}}_{\boldsymbol{\eta}}(t) \right\}, \tag{198}$$

$$a_{\mathrm{res}}(t, \boldsymbol{\eta}) := \mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[h_{t,\boldsymbol{\eta}}^2], \tag{199}$$

$$d_{\mathrm{KL}}(t, \boldsymbol{\eta}) := D_{\mathrm{KL}}(Q_t^{\boldsymbol{\eta}} \| \widetilde{Q}_t), \tag{200}$$

$$\mathcal{A}_{W_2}(\boldsymbol{\eta}) := \int_0^1 W_2^2(Q_t^{\boldsymbol{\eta}}, \widetilde{Q}_t)\, \rho(dt). \tag{201}$$

The corresponding integrated scores are $\mathcal{A}_{\tan}(\boldsymbol{\eta}) := \int_0^1 a_{\tan}(t, \boldsymbol{\eta})\, \rho(dt)$, $\mathcal{A}_{\mathrm{res}}(\boldsymbol{\eta}) := \int_0^1 a_{\mathrm{res}}(t, \boldsymbol{\eta})\, \rho(dt)$, and $J_{\mathrm{KL}}(\boldsymbol{\eta}) := \int_0^1 d_{\mathrm{KL}}(t, \boldsymbol{\eta})\, \rho(dt)$. Residual L^2 is used when $h_{t,\boldsymbol{\eta}}$ is square integrable, and static W_2 when the marginals have finite second moments.

C.2 Weighted transport operator

Fix $(t, \boldsymbol{\eta})$ and write $Q = Q_t^{\boldsymbol{\eta}}$, $q = q_t^{\boldsymbol{\eta}}$, and $h = h_{t,\boldsymbol{\eta}}$. Let $L_0^2(Q) = \{f \in L^2(Q) : \mathbb{E}_Q[f] = 0\}$. For the L^2 operator representation, assume $h \in L_0^2(Q)$, a densely defined closed weighted Dirichlet form, and the weighted Poincaré inequality of Appendix B.7. The associated self-adjoint operator L_q on $L_0^2(Q)$ is characterized by

$$\langle \varphi, \mathsf{L}_q \psi \rangle_{L^2(Q)} = \mathbb{E}_Q[\nabla\varphi \cdot \nabla\psi], \qquad \psi \in \mathrm{Dom}(\mathsf{L}_q), \quad \varphi \in H_q. \tag{202}$$

When the strong representation is justified,

$$\mathsf{L}_q \psi = -q^{-1} \nabla \cdot (q \nabla \psi). \tag{203}$$

The mean-zero Poisson potential satisfies $\psi = -\mathsf{L}_q^{-1} h$. Testing the weak Poisson equation with ψ gives

$$a_{\mathrm{full}}(t, \boldsymbol{\eta}) = \left\langle h_t, \mathsf{L}_{q_t}^{-1} h_t \right\rangle_{L^2(Q_t)}. \tag{204}$$

Thus Full applies the inverse weighted transport operator to the continuity forcing before measuring its quadratic size. This representation requires the stated L^2 operator assumptions; the more general variational characterization of Theorem 2 only requires the forcing to define a bounded functional on H_{q_t}, in which case the inverse is understood from $H_{q_t}^*$ to H_{q_t} and the pairing is dual rather than an L^2 inner product.

C.3 Tangent-invisible correction and lack of universal control

Remark 9 (Finite-rank origin of hidden action)**.** *Let $\mathcal{H} = L^2(Q; \mathbb{R}^d)$ and $L\delta = \mathbb{E}_Q[J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}} \delta]$. Then* $\mathrm{rank}(L) \le R$ *and* $\mathrm{Range}(L^*) = \{J_{\boldsymbol{\Phi}_{\boldsymbol{\eta}}}^\top b : b \in \mathbb{R}^R\}$. *Proposition 2 identifies Tangent as the orthogonal projection of the Full correction onto this range and places Hidden in* $\ker L$. *If $\mathcal{H}$ is*

infinite dimensional, finite R leaves an infinite-dimensional kernel; with D discrete velocity degrees of freedom, $\dim\ker L \geq D - R$. *This dimension count identifies unresolved directions but does not imply nonzero hidden action in every instance. The Galerkin interpretation and exact equality conditions are given in Remark 6 and Corollary 10.*

Theorem 3 (No universal tangent control)**.** *There is no universal finite K such that $a_{\mathrm{full}}(t,\boldsymbol{\eta}) \leq K a_{\tan}(t,\boldsymbol{\eta})$ over all regular instances of the information-projected construction. There are smooth instances with $a_{\tan} = 0 < a_{\mathrm{full}}$, and smooth families with positive Tangent action for which $a_{\mathrm{full}}/a_{\tan} \to \infty$. The zero-versus-positive construction can be stationary, so no such universal bound holds for the integrated actions either.*

Example 1 (A stationary law with entirely hidden correction)**.** *Take $\Phi(x) = \cos x$, $a > 0$, $\vartheta \neq 0$, and*

$$I_0(z) = \frac{1}{2\pi}\int_0^{2\pi} e^{z\cos x}\,dx, \qquad \widetilde{q}(x) = \frac{e^{-\vartheta\cos x}}{2\pi I_0(\vartheta)}, \qquad u(x) = ae^{\vartheta\cos x}. \tag{205}$$

The reference flux $\widetilde{q}u = a/(2\pi I_0(\vartheta))$ is constant. Information projection onto $\mathbb{E}[\cos X] = 0$ gives the uniform law $q = (2\pi)^{-1}$, with

$$a_{\tan} = 0, \qquad a_{\mathrm{full}} = a^2\left[I_0(2\vartheta) - I_0(\vartheta)^2\right] > 0. \tag{206}$$

Verification and proof of Theorem 3. To include the positive-Tangent perturbation, fix $t_0 \in (0,1)$ and set $\lambda_t = \vartheta + 2\varepsilon(t - t_0)$. With $A(\alpha) = \log \mathbb{E}_{\widetilde{Q}}[e^{\alpha\cos X}]$, define

$$q_t(x) = \widetilde{q}(x)\exp(\lambda_t\cos x - A(\lambda_t)), \qquad c(t) = \mathbb{E}_{q_t}[\cos X]. \tag{207}$$

This is the exact information projection for its target moment. All densities are smooth and positive, and the scalar observable covariance is positive. At t_0, $A(\vartheta) = -\log I_0(\vartheta)$, so $q_{t_0} = (2\pi)^{-1}$, $c(t_0) = 0$, $C_{t_0} = G_{t_0} = 1/2$, and $\dot{c}(t_0) = \varepsilon$. Since $\mathbb{E}_{q_{t_0}}[-\sin X\, u(X)] = 0$, one has $r_{t_0} = -\varepsilon$ and $\delta^{\tan}_{t_0} = -2\varepsilon\sin x$. The continuity residual and Poisson equation are

$$h_{t_0}(x) = 2\varepsilon\cos x + u'(x), \qquad \psi''(x) = 2\varepsilon\cos x + u'(x). \tag{208}$$

Periodicity fixes the integration constant in

$$\psi'(x) = 2\varepsilon\sin x + u(x) - \mathbb{E}_{q_{t_0}}[u], \qquad \delta^{\star}_{t_0} = -\psi'(x). \tag{209}$$

The sine term is odd and $u - \mathbb{E}[u]$ is even under $x \mapsto 2\pi - x$, so their inner product vanishes. Consequently,

$$a_{\tan}(t_0) = 2\varepsilon^2, \qquad a_{\mathrm{full}}(t_0) = 2\varepsilon^2 + a^2\left[I_0(2\vartheta) - I_0(\vartheta)^2\right]. \tag{210}$$

The bracket is the uniform-law variance of $e^{\vartheta\cos X}$ and is strictly positive because $\vartheta \neq 0$. Setting $\varepsilon = 0$ gives Example 1 at every time. For $\varepsilon \neq 0$,

$$\frac{a_{\mathrm{full}}(t_0)}{a_{\tan}(t_0)} = 1 + \frac{a^2[I_0(2\vartheta) - I_0(\vartheta)^2]}{2\varepsilon^2} \longrightarrow \infty \qquad \text{as } \varepsilon \to 0. \tag{211}$$

□

C.3.1 No universal tangent control in the unbalanced geometry

Corollary 13 (No universal tangent control in the unbalanced geometry)**.** *Fix any $\kappa > 0$. There is no universal finite constant K such that*

$$a^{\mathrm{ub}}_{\mathrm{full}}(t,\boldsymbol{\eta}) \leq K\, a^{\mathrm{ub}}_{\tan}(t,\boldsymbol{\eta}) \tag{212}$$

for all regular finite-measure instances of the unbalanced construction. In fact, there are smooth instances with

$$a^{\mathrm{ub}}_{\tan}(t,\boldsymbol{\eta}) = 0 \qquad \textit{and} \qquad a^{\mathrm{ub}}_{\mathrm{full}}(t,\boldsymbol{\eta}) > 0. \tag{213}$$

Proof. Use the construction in the proof of Theorem 3, set $M_t \equiv 1$ and $g_{\mathrm{ref}} \equiv 0$, and evaluate at $t = t_0$, where q_{t_0} is uniform and

$$h_{\mathrm{ub}}(x) = 2\varepsilon\cos x + u'(x). \tag{214}$$

For the single observable $\Phi(x) = \cos x$, Corollary 8 gives

$$G^{\mathrm{ub}} = \mathbb{E}[\sin^2 X] + \frac{1}{\kappa}\mathbb{E}[\cos^2 X] = \frac{1}{2}\left(1 + \frac{1}{\kappa}\right), \tag{215}$$

and therefore

$$a^{\mathrm{ub}}_{\tan}(t_0) = \frac{2\kappa}{\kappa+1}\varepsilon^2. \tag{216}$$

For the Full problem, the screened operator on the uniform torus is

$$\mathcal{L}_\kappa := -\partial_{xx} + \kappa^{-1}. \tag{217}$$

The forcing $2\varepsilon\cos x$ is even, whereas $u'(x)$ is odd, and $\mathcal{L}_\kappa^{-1}$ preserves parity. Hence the two forcing components are orthogonal in the screened energy geometry. Since

$$\mathcal{L}_\kappa(\cos x) = \left(1 + \frac{1}{\kappa}\right)\cos x, \tag{218}$$

the even component contributes exactly

$$\frac{2\kappa}{\kappa+1}\varepsilon^2. \tag{219}$$

The odd component contributes

$$c_\kappa := \mathbb{E}_{q_{t_0}}\left[u'\,\mathcal{L}_\kappa^{-1}u'\right] > 0, \tag{220}$$

because $\mathcal{L}_\kappa$ is strictly positive and $u' \not\equiv 0$. Thus

$$a^{\mathrm{ub}}_{\mathrm{full}}(t_0) = \frac{2\kappa}{\kappa+1}\varepsilon^2 + c_\kappa. \tag{221}$$

At $\varepsilon = 0$, equation 216–equation 221 give equation 213. For $\varepsilon \neq 0$ their ratio diverges as $\varepsilon \to 0$, ruling out any universal finite K. □

C.4 Residual amplitude versus transport effort

The residual comparator uses the full normalized residual, $a_{\mathrm{res}} = \|h\|^2_{L^2(Q)}$, rather than a finite-dimensional projection. Its difference from Full is the missing inverse transport operator. Weighted squared Liouville residuals also appear in best-fit statistical closure (Turkington, 2013); the present comparator is the plain $L^2(Q)$ norm on an observation-prescribed path, not a weighted objective that selects the path itself.

Proposition 10 (Spectral comparison)**.** *Under Appendix C.2, suppose L_q admits a complete orthonormal eigenbasis $(e_j)_{j\geq 1}$ on $L^2_0(Q)$, with $\mathsf{L}_q e_j = \nu_j e_j$ and $0 < \nu_1 \leq \nu_2 \leq \cdots$. If $h = \sum_j c_j e_j$, then*

$$a_{\mathrm{res}} = \sum_j c_j^2, \qquad a_{\mathrm{full}} = \sum_j \frac{c_j^2}{\nu_j}, \qquad a_{\mathrm{full}} \leq \nu_1^{-1} a_{\mathrm{res}}. \tag{222}$$

For residuals supported in a spectral band $\nu_j \leq \Lambda$,

$$\nu_1 a_{\mathrm{full}} \leq a_{\mathrm{res}} \leq \Lambda a_{\mathrm{full}}. \tag{223}$$

If $\nu_j \to \infty$, no finite C satisfies $a_{\mathrm{res}} \leq C a_{\mathrm{full}}$ for every $h \in L^2_0(Q)$.

Proof. Parseval's identity gives the residual norm. Expanding $\mathsf{L}_q^{-1}h$ in the same eigenbasis gives the Full quadratic form. The inequalities follow by bounding the active eigenvalues above and below. A universal reverse bound would imply $1 \leq C/\nu_j$ upon taking $h = e_j$, contradicting $\nu_j \to \infty$. □

Corollary 14 (Pointwise ranking reversal)**.** *If $\nu_j > \nu_\ell$ and $1 < a_A/a_B < \sqrt{\nu_j/\nu_\ell}$, then $h_A = a_A e_j$ and $h_B = a_B e_\ell$ satisfy $a_{\mathrm{res}}(h_A) > a_{\mathrm{res}}(h_B)$ but $a_{\mathrm{full}}(h_A) < a_{\mathrm{full}}(h_B)$.*

Proof. The two comparisons are respectively $a_A^2 > a_B^2$ and $a_A^2/\nu_j < a_B^2/\nu_\ell$. □

Example 2 (Exact information-projected Fourier family)**.** *Let $\widetilde{q} = q_0 = (2\pi)^{-1}$ and $u = 0$. For an integer $k \geq 1$, define*

$$\Phi_k(x) = \cos(kx), \qquad q_{\lambda,k}(x) = \frac{e^{\lambda \cos(kx)}}{2\pi I_0(\lambda)}, \qquad c(\lambda) = \frac{d}{d\lambda} \log I_0(\lambda). \tag{224}$$

Periodicity makes the normalization and the moment $\mathbb{E}_{q_{\lambda,k}}[\Phi_k] = c(\lambda)$ independent of k. Hence $q_{\lambda,k}$ is exactly the information projection onto this moment fiber. Along $q_{k,t} = q_{\lambda_k(t),k}$,

$$h_{k,t}(x) = \dot{\lambda}_k(t)\big(\cos(kx) - c(\lambda_k(t))\big). \tag{225}$$

At $\lambda_k(t_0) = 0$ and $\dot{\lambda}_k(t_0) = a$, the density is uniform, $h_{k,t_0} = a\cos(kx)$, and the mean-zero Poisson potential is $\psi(x) = -a\cos(kx)/k^2$. Therefore

$$a_{\mathrm{res}}(t_0, k) = \frac{a^2}{2}, \qquad a_{\mathrm{full}}(t_0, k) = \frac{a^2}{2k^2}. \tag{226}$$

The pair $h_A = 2\cos(3x)$ and $h_B = \cos x$ gives the exact instantaneous ranking reversal $2 > 1/2$ for residual L^2 and $2/9 < 1/2$ for Full.

Transport metric along the Fourier family. Write $q_\lambda = q_{\lambda,1}$ and define

$$g(\lambda) = \|\cos x - c(\lambda)\|^2_{H^{-1}(q_\lambda)}. \tag{227}$$

This function is continuous, strictly positive, and satisfies $g(0) = 1/2$. An explicit verification avoids any assumption about a limiting discretization. Set $f_\lambda(x) = \cos x - c(\lambda)$ and

$$F_\lambda(x) = \int_0^x q_\lambda(y) f_\lambda(y)\, dy, \qquad b_\lambda = \frac{\int_0^{2\pi} F_\lambda(x)/q_\lambda(x)\, dx}{\int_0^{2\pi} 1/q_\lambda(x)\, dx}. \tag{228}$$

For a unit multiplier rate, $v_\lambda = (b_\lambda - F_\lambda)/q_\lambda$ satisfies $(q_\lambda v_\lambda)' = -q_\lambda f_\lambda$. Its spatial integral vanishes by the definition of b_λ, so it is a periodic gradient and the minimum-energy correction. Consequently,

$$g(\lambda) = \int_0^{2\pi} \frac{(b_\lambda - F_\lambda(x))^2}{q_\lambda(x)}\, dx. \tag{229}$$

On every compact multiplier interval, q_λ is uniformly positive and all terms vary smoothly with λ, proving continuity. The energy cannot vanish because f_λ is not identically zero. At $\lambda = 0$, $v_0 = -\sin x$, giving $g(0) = 1/2$. Finally, if χ_λ solves $\mathsf{L}_{q_{\lambda,1}}\chi_\lambda = f_\lambda$, then $k^{-2}\chi_\lambda(kx)$ solves the frequency-k problem. Thus the exact action along any differentiable multiplier path is

$$a_{\mathrm{full}}(t, k) = \frac{\dot{\lambda}_k(t)^2}{k^2} g(\lambda_k(t)). \tag{230}$$

C.5 Forward KL: local geometry and exact separations

Proposition 11 (Exact dual form and local information geometry)**.** *At fixed $(t, \boldsymbol{\eta})$, abbreviate $A = A_{t,\boldsymbol{\eta}}$ and $\phi = \boldsymbol{\Phi}_{\boldsymbol{\eta}}$, and let Q_λ be the exponential tilt of $\widetilde{Q}_t$. Write $\widetilde{c} = \mathbb{E}_{\widetilde{Q}_t}[\phi]$ and $\widetilde{C} = \mathrm{Cov}_{\widetilde{Q}_t}(\phi, \phi)$. Then*

$$D_{\mathrm{KL}}(Q_\lambda \| \widetilde{Q}_t) = \lambda^\top \nabla A(\lambda) - A(\lambda). \tag{231}$$

At the calibrated multiplier, $\nabla A(\lambda) = \boldsymbol{c}_{\boldsymbol{\eta}}(t)$. If A is C^3 near zero, then

$$D_{\mathrm{KL}}(Q_\lambda \| \widetilde{Q}_t) = \frac{1}{2}\lambda^\top \widetilde{C}\lambda + O(\|\lambda\|^3). \tag{232}$$

If also $\widetilde{C} \succ 0$, then, for $\Delta c = \mathbb{E}_{Q_\lambda}[\phi] - \widetilde{c}$,

$$D_{\mathrm{KL}}(Q_\lambda \| \widetilde{Q}_t) = \frac{1}{2}\Delta c^\top \widetilde{C}^{-1}\Delta c + O(\|\Delta c\|^3). \tag{233}$$

Proof. Integrating $\log(dQ_\lambda/d\widetilde{Q}_t) = \lambda^\top\phi - A(\lambda)$ gives the dual identity. Since $A(0) = 0$, Taylor expansion yields $A(\lambda) = \widetilde{c}^\top\lambda + \frac{1}{2}\lambda^\top\widetilde{C}\lambda + O(\|\lambda\|^3)$ and $\nabla A(\lambda) = \widetilde{c} + \widetilde{C}\lambda + O(\|\lambda\|^2)$. Substitution cancels the linear term. When $\widetilde{C} \succ 0$, the inverse-function theorem gives $\lambda = \widetilde{C}^{-1}\Delta c + O(\|\Delta c\|^2)$, proving the moment-coordinate expansion. □

Marginal tilt versus multiplier velocity. For a stationary reference with $u = 0$, let $\overline{\phi}_\lambda = \phi - \mathbb{E}_{Q_\lambda}[\phi]$. When the centered coordinates and inverse operator are well defined, set

$$[\mathsf{H}(\lambda)]_{ab} = \langle \overline{\phi}_{\lambda,a}, \mathsf{L}_{q_\lambda}^{-1} \overline{\phi}_{\lambda,b} \rangle_{L^2(Q_\lambda)}. \tag{234}$$

The identity $h_t = \dot{\lambda}_t^\top \overline{\phi}_{\lambda_t}$ implies

$$a_{\mathrm{full}}(t) = \dot{\lambda}_t^\top \mathsf{H}(\lambda_t) \dot{\lambda}_t. \tag{235}$$

In a regular parametrization admitting a Wasserstein pullback, H is the transport metric in natural-parameter coordinates (Chen & Li, 2020; Li & Zhao, 2023). It differs from both the covariance/Fisher matrix and the observable derivative Gram matrix. Marginal KL acts on tilt displacement, whereas Full acts on the velocity of the prescribed path in transport geometry.

Proposition 12 (No universal comparison between forward KL and Full)**.** *Even with the state space $\mathbb{T}$ and time interval $[0,1]$ fixed, neither $d_{\mathrm{KL}} \le C a_{\mathrm{full}}$ nor $a_{\mathrm{full}} \le C d_{\mathrm{KL}}$ holds with an instance-independent finite C over all regular instances of the information-projected construction. Likewise, for $\rho(dt) = dt$, neither $J_{\mathrm{KL}} \le C\mathcal{A}$ nor $\mathcal{A} \le C J_{\mathrm{KL}}$ holds universally.*

Example 3 (Pointwise separations)**.** *Under the uniform reference with $u = 0$, take the Fourier family in equation 224. A constant nonzero multiplier λ_0 gives a stationary projected law distinct from the reference: $h_t = 0$, so $a_{\mathrm{full}} = 0$, while $d_{\mathrm{KL}} > 0$. Conversely, at $\lambda_k(t_0) = 0$ and $\dot{\lambda}_k(t_0) = a \neq 0$, the two marginals coincide but their time derivatives need not, and*

$$d_{\mathrm{KL}}(t_0, k) = 0, \qquad a_{\mathrm{full}}(t_0, k) = \frac{a^2}{2k^2} > 0. \tag{236}$$

The constant-tilt construction also gives $J_{\mathrm{KL}} > 0 = \mathcal{A}$ after integration.

Example 4 (Rapid marginal oscillation)**.** *Fix $\epsilon > 0$, set $\lambda_n(t) = \epsilon \sin(2\pi n t)$, and take*

$$q_{n,t}(x) = \frac{\exp(\lambda_n(t) \cos x)}{2\pi I_0(\lambda_n(t))}, \qquad n \in \mathbb{N}, \quad n \ge 1. \tag{237}$$

Every path is an exact information projection for its $\cos x$ moment trajectory and satisfies $q_{n,0} = q_{n,1} = q_0$. For integer n,

$$J_{\mathrm{KL}}(n) = J_{\mathrm{KL}}(1), \qquad \mathcal{A}(n) = n^2 \mathcal{A}(1) > 0. \tag{238}$$

There exist constants $C_\epsilon < \infty$ and $m_\epsilon > 0$, independent of n, such that

$$J_{\mathrm{KL}}(n) \le C_\epsilon \epsilon^2, \qquad \mathcal{A}(n) \ge 2\pi^2 m_\epsilon \epsilon^2 n^2. \tag{239}$$

As a tie-breaking thought experiment, suppose the primary task does not distinguish these paths beyond constraints that they all satisfy, such as their common endpoints. Their time-averaged forward KL displacement from the reference is identical for every integer n, yet their required dynamical revision grows as n^2. Under the minimum-additional-commitment preference specified in the main text, Full therefore distinguishes a path that makes a modest reference-relative excursion from one that repeatedly moves back and forth through the same range, even though a static marginal criterion and the declared endpoint task do not.

Verification and completion of Proposition 12. Let $D(\lambda)$ be the forward KL of $q_{\lambda,1}$ relative to q_0. This function is smooth and vanishes to second order at zero, so $D(\lambda) \le C_\epsilon \lambda^2$ on $[-\epsilon, \epsilon]$. The function g in equation 227 has a positive minimum m_ϵ on this interval. Thus

$$J_{\mathrm{KL}}(n) = \int_0^1 D(\epsilon \sin(2\pi n t))\, dt, \qquad \mathcal{A}(n) = (2\pi n \epsilon)^2 \int_0^1 \cos^2(2\pi n t) g(\epsilon \sin(2\pi n t))\, dt. \tag{240}$$

The substitution $\theta = 2\pi n t$ integrates n identical periods, proving equation 238. Bounding D and g as above and using $\int_0^1 \cos^2(2\pi n t)\, dt = 1/2$ proves equation 239. Constants may be chosen uniformly for $0 < \epsilon \le \epsilon_0$ by taking bounds on a fixed compact interval $[-\epsilon_0, \epsilon_0]$. The marginal score is positive for $\epsilon > 0$, so $\mathcal{A}(n)/J_{\mathrm{KL}}(n) \to \infty$. Together with Example 3, this proves all claims of Proposition 12. □

Corollary 15 (Local KL/Full ranking reversal)**.** *Let $s \in C^1([0,1])$ be nonconstant and compare $q_{A,t} = q_{2\epsilon s(t),3}$ with $q_{B,t} = q_{\epsilon s(t),1}$. Then $J_{\mathrm{KL}}(A) > J_{\mathrm{KL}}(B)$ for every $\epsilon > 0$, whereas $\mathcal{A}(A) < \mathcal{A}(B)$ for all sufficiently small $\epsilon > 0$. More precisely,*

$$J_{\mathrm{KL}}(A) = \epsilon^2 \int_0^1 s(t)^2\, dt + o(\epsilon^2), \qquad J_{\mathrm{KL}}(B) = \frac{\epsilon^2}{4} \int_0^1 s(t)^2\, dt + o(\epsilon^2), \tag{241}$$

while

$$\mathcal{A}(A) = \frac{2\epsilon^2}{9} \int_0^1 \dot{s}(t)^2\, dt + o(\epsilon^2), \qquad \mathcal{A}(B) = \frac{\epsilon^2}{2} \int_0^1 \dot{s}(t)^2\, dt + o(\epsilon^2). \tag{242}$$

Proof. For $A(\lambda) = \log I_0(\lambda)$, $D(\lambda) = \lambda A'(\lambda) - A(\lambda)$ is even and $D'(\lambda) = \lambda \mathrm{Var}_{q_{\lambda,1}}(\cos X)$. Hence D is strictly increasing in $|\lambda|$ away from zero. Since s is nonconstant, it is nonzero on a set of positive measure, proving the strict KL ordering. The expansion $D(\lambda) = \lambda^2/4 + o(\lambda^2)$ gives equation 241. Equation equation 230 gives the exact actions $4\epsilon^2 \dot{s}(t)^2 g(2\epsilon s(t))/9$ and $\epsilon^2 \dot{s}(t)^2 g(\epsilon s(t))$. Continuity of g at zero, $g(0) = 1/2$, and boundedness of s give equation 242 after integration. Because $\int_0^1 \dot{s}(t)^2\, dt > 0$, the leading coefficients yield the opposite strict Full ranking for sufficiently small ϵ. □

C.6 Static Wasserstein displacement and initial matching

The comparator $\mathcal{A}_{W_2}$ in equation 201 averages static distances between projected and reference marginals. It is not a dynamic endpoint-transport objective and does not optimize or directly assess the time derivative of the prescribed projected path.

Proposition 13 (Static Wasserstein displacement and Full are non-equivalent)**.** *Over regular instances of the information-projected construction on $\mathbb{T}$, neither $\mathcal{A}_{W_2} \le C\mathcal{A}$ nor $\mathcal{A} \le C\mathcal{A}_{W_2}$ holds with an instance-independent finite C when $\rho(dt) = dt$. The analogous pointwise inequalities for $W_2^2(Q_t^{\boldsymbol{\eta}}, \widetilde{Q}_t)$ and $a_{\mathrm{full}}(t, \boldsymbol{\eta})$ also fail. In the rapid-oscillation family of Example 4,*

$$\mathcal{A}_{W_2}(n) = \mathcal{A}_{W_2}(1), \qquad \mathcal{A}(n) = n^2 \mathcal{A}(1), \tag{243}$$

and there is $M_\epsilon < \infty$, independent of n, such that

$$\mathcal{A}_{W_2}(n) \le M_\epsilon \epsilon^2, \qquad \mathcal{A}(n) \ge 2\pi^2 m_\epsilon \epsilon^2 n^2. \tag{244}$$

Proof. The constant nonzero tilt in Example 3 has zero Full action but positive Wasserstein distance from the reference, both pointwise and after integration. At a zero multiplier with nonzero multiplier derivative, the Wasserstein distance is zero while Full is positive. These give the pointwise failures and rule out $\mathcal{A}_{W_2} \le C\mathcal{A}$ without initial matching. For the reverse integrated bound, the time average of $W_2^2(q_{\epsilon \sin(2\pi n t),1}, q_0)$ is independent of integer n by periodicity, whereas Full has the exact n^2 scaling already proved.

For the small-amplitude bound, set $M_\epsilon = \max_{|\lambda| \le \epsilon} g(\lambda)$. Given $|\lambda| \le \epsilon$, the path $r \mapsto q_{r\lambda,1}$, $r \in [0,1]$, has minimum kinetic energy density $\lambda^2 g(r\lambda)$. Using this path as an admissible competitor in dynamic optimal transport gives

$$W_2^2(q_{\lambda,1}, q_0) \le \int_0^1 \lambda^2 g(r\lambda)\, dr \le M_\epsilon \lambda^2. \tag{245}$$

Substitution of $\lambda = \lambda_n(t)$ and integration prove equation 244. The constant can be chosen uniformly for $0 < \epsilon \le \epsilon_0$ by using the maximum over $[-\epsilon_0, \epsilon_0]$. □

Proposition 14 (Wasserstein stability relative to a matched-initial reference rollout)**.** *Let $\mathcal{X} = \mathbb{R}^d$ or a flat torus, take $\rho(dt) = dt$, and assume $Q_t^{\boldsymbol{\eta}}, R_t^{\boldsymbol{\eta}} \in \mathcal{P}_2(\mathcal{X})$ and $\mathcal{A}(\boldsymbol{\eta}) < \infty$. Let $\mathsf{T}_{0,t}^u$ denote the characteristic flow generated by the frozen reference velocity u_t, and define the matched-initial reference rollout*

$$R_t^{\boldsymbol{\eta}} := (\mathsf{T}_{0,t}^u)_{\#} Q_0^{\boldsymbol{\eta}}. \tag{246}$$

Suppose u_t and $u_t + \delta_t^\star$ admit classical characteristic flows realizing $R_t^{\boldsymbol{\eta}}$ and $Q_t^{\boldsymbol{\eta}}$, respectively, and assume $\|u_t(x) - u_t(y)\| \le L(t)\|x - y\|$ with $L \ge 0$ integrable, imposed on periodic lifts for a torus. With $\Lambda = \int_0^1 L(t)\,dt$,

$$W_2(Q_t^{\boldsymbol{\eta}}, R_t^{\boldsymbol{\eta}}) \le \int_0^t \exp\left(\int_s^t L(\tau)\,d\tau\right) \sqrt{a_{\mathrm{full}}(s, \boldsymbol{\eta})}\,ds, \tag{247}$$

and consequently

$$\int_0^1 W_2^2(Q_t^{\boldsymbol{\eta}}, R_t^{\boldsymbol{\eta}})\,dt \le e^{2\Lambda}\mathcal{A}(\boldsymbol{\eta}). \tag{248}$$

If additionally $Q_0^{\boldsymbol{\eta}} = \widetilde{Q}_0$, then $R_t^{\boldsymbol{\eta}} = \widetilde{Q}_t$ and the left-hand side equals $\mathcal{A}_{W_2}(\boldsymbol{\eta})$ as defined in equation 201.

Moreover, for every globally 1-Lipschitz observable f,

$$\left|\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[f] - \mathbb{E}_{R_t^{\boldsymbol{\eta}}}[f]\right| \le W_1(Q_t^{\boldsymbol{\eta}}, R_t^{\boldsymbol{\eta}}) \le W_2(Q_t^{\boldsymbol{\eta}}, R_t^{\boldsymbol{\eta}}). \tag{249}$$

Since $\int_0^t a_{\mathrm{full}}(s, \boldsymbol{\eta})\,ds \le \mathcal{A}(\boldsymbol{\eta})$ and $\int_s^t L(\tau)\,d\tau \le \Lambda$, the pointwise bound also gives

$$\sup_{\mathrm{Lip}(f)\le 1} \left|\mathbb{E}_{Q_t^{\boldsymbol{\eta}}}[f] - \mathbb{E}_{R_t^{\boldsymbol{\eta}}}[f]\right| \le e^{\Lambda}\sqrt{\mathcal{A}(\boldsymbol{\eta})}, \qquad t \in [0, 1]. \tag{250}$$

Thus accumulated Full action controls a one-sided reference-relative stability envelope for Lipschitz observables after removing initial-law mismatch.

Proof. Couple the two characteristic flows from a common initial random state $X_0 = Y_0 \sim Q_0^{\boldsymbol{\eta}}$, with $\dot{X}_t = u_t(X_t) + \delta_t^\star(X_t)$ and $\dot{Y}_t = u_t(Y_t)$. Then $\mathcal{L}(X_t) = Q_t^{\boldsymbol{\eta}}$ and $\mathcal{L}(Y_t) = R_t^{\boldsymbol{\eta}}$. On a torus, use lifts with the same initial representative; lifted Euclidean distance bounds geodesic distance. Gronwall's inequality gives

$$\|X_t - Y_t\| \le \int_0^t \exp\left(\int_s^t L(\tau)\,d\tau\right) \|\delta_s^\star(X_s)\|\,ds. \tag{251}$$

The L^2 integral Minkowski inequality and $\mathbb{E}[\|\delta_s^\star(X_s)\|^2] = a_{\mathrm{full}}(s, \boldsymbol{\eta})$ give equation 247. Cauchy–Schwarz yields $W_2^2(Q_t^{\boldsymbol{\eta}}, R_t^{\boldsymbol{\eta}}) \le e^{2\Lambda} t \int_0^t a_{\mathrm{full}}(s, \boldsymbol{\eta})\,ds \le e^{2\Lambda}\mathcal{A}(\boldsymbol{\eta})$, and integration in t proves equation 248. Finally, the Kantorovich–Rubinstein bound gives the first inequality in equation 249, while $W_1 \le W_2$ follows from the finite second-moment assumption. □

The constant-tilt example has $Q_0^{\boldsymbol{\eta}} \ne \widetilde{Q}_0$, so it does not contradict the special-case bound relative to the original frozen reference marginals. Relative to the matched-initial rollout $R_t^{\boldsymbol{\eta}}$, the constant-tilt path has zero Full action and zero matched-initial Wasserstein departure, consistently with Proposition 14. When $Q_0^{\boldsymbol{\eta}} = \widetilde{Q}_0$, one has $R_t^{\boldsymbol{\eta}} = \widetilde{Q}_t$, and the proposition specializes to a bound on static Wasserstein displacement from the original reference path. The oscillatory family satisfies this initial equality and still rules out a reverse estimate. Likewise, the spectral equivalence in equation 223 requires common positive lower and finite upper spectral bounds to obtain constants uniform across designs. These statements are independent of the scientific-risk constraint.

D Additional Experiments and Ablations

Shared experimental conventions and reproducibility details are provided in Appendix E. Supporting-study-specific protocols are given below where needed.

D.1 Additional Experiments

We report two additional supporting experiments. The retrospective double-gyre study tests repeatability of the Full-versus-Law contrast across fresh observations and independently trained fits within the same declared reference-learning regime. The finite-measure study extends FIDE beyond mass-preserving probability laws by allowing reference-relative correction through both transport and reaction.

Allowance	Full holdout risk	Risk increase (%)	Excess over allowance (pp)
0.5%	0.038777	0.4404	−0.0596
1%	0.039016	1.0598	+0.0598
2%	0.039289	1.7686	−0.2314
3%	0.039531	2.3945	−0.6055
4%	0.039950	3.4794	−0.5206
5%	0.040538	5.0023	+0.0023

Table 2: **Held-out scientific-risk transfer for the retrospective frontier.** The holdout Law risk is 0.038607. Negative excess indicates a point estimate below the nominal allowance. Values are descriptive because the retained summaries do not contain paired per-trial scientific losses needed for matching simultaneous risk-feasibility inference.

D.1.1 Retrospective double-gyre repeatability study

System and principal results. The study uses the double-gyre dynamics defined in Appendix E.3, with four Gaussian sensors observing aggregate particle responses at nine acquisition times. Selection uses a frozen microscopic target simulation bank together with three independently seeded endpoint-trained references fitted from the same physical endpoint dataset. The study therefore asks a narrow question: does the Full-versus-Law action reduction repeat across fresh observations and across independent fits within the same declared reference-learning regime? On 64 shared fresh trials per reference, the mean reference-wise reduction ranges from 9.96% at the 0.5% allowance to 34.34% at 5%, and simultaneous 95% paired-bootstrap intervals remain above zero across all 18 reference-by-allowance contrasts. Held-out scientific-risk point estimates remain close to their nominal allowances, although matching simultaneous risk-feasibility inference is unavailable because the retained evaluation summaries do not contain paired per-trial Law and Full scientific losses. These results support finite-sample and within-regime reference-fit repeatability; they do not establish transfer to an unseen reference family, arbitrary reference misspecification, or independent physical endpoint datasets.

Protocol and numerical evaluation. The truth bank contains 50,000 particles. Each of the three endpoint-trained flow-matching references uses four width-128 hidden layers, 12,000 Adam steps, and 32,768 rollout particles. Moment reconstruction uses endpoint-anchored bounded cubic penalized least squares with three internal knots, smoothing 10^{-4}, and relative ridge 10^{-10}. Scientific risk averages four MMD bandwidths across the three references; the frozen Law baseline is 0.0386 and the population-risk ceiling is 0.0634. The shared selection bank contains 128 trials, and the final numerical Full evaluator uses a 256×128 raster at all 21 scientific times. After the six-point frontier is frozen, evaluation uses 64 shared fresh trials per reference. Simultaneous 95% action intervals use 100,000 paired bootstrap resamples over all 18 reference-by-allowance contrasts.

D.1.2 Finite-measure transport–reaction experiment

We consider an unbalanced Gaussian transport problem in which the evolving object is a finite measure rather than a probability law. FIDE acts on the normalized shape through the same information-projection construction used in the balanced experiments, while discrepancies in total mass are represented through the reaction channel of the screened transport–reaction correction. As in the balanced setting, Law provides the scientific-risk reference design and Tangent optimizes only the correction visible through the measured raw moment rates. Full instead optimizes the unrestricted transport–reaction correction in equation 32. We additionally introduce GlobalReaction as an intermediate structural comparator. It retains an unrestricted transport correction but restricts the reaction field to a single spatially constant rate at each time:

$$a^{\mathrm{ub}}_{\mathrm{GR}}(t,\boldsymbol{\eta}) := M_t \inf_{\delta,\, a\in\mathbb{R}} \left\{ \mathbb{E}_{q_t}\left[\|\delta\|^2 + \kappa a^2\right] : -\nabla\cdot(q_t\delta) + q_t a = q_t h_{\mathrm{ub}} \right\}. \tag{252}$$

Integrating the constraint in equation 252 over space fixes $a = \mathbb{E}_{q_t}[h_{\mathrm{ub}}] = \dot{M}_t/M_t - g_{\mathrm{ref}}(t)$. GlobalReaction therefore permits reaction to correct the global mass-rate discrepancy, while any spatially heterogeneous mean-zero forcing must be realized through transport. Full removes this restriction by allowing the optimal reaction field to vary over space, whereas Tangent imposes the distinct restriction of matching only the measured raw moment rates.

	Mean Full action				Reduction vs. Law (%)			
Allowance	Law	Tangent	GlobalReaction	Full	Tangent	GlobalReaction	Full	Full interval (%)
1%	0.77523	0.76973	0.76652	0.74994	0.71	1.12	3.26	[1.50, 5.03]
2%	0.77523	0.76502	0.76205	0.73475	1.32	1.70	5.22	[3.46, 6.99]
3%	0.77523	0.75927	0.75660	0.72193	2.06	2.40	6.88	[5.11, 8.64]
4%	0.77523	0.75296	0.75043	0.71029	2.87	3.20	8.38	[6.61, 10.14]
5%	0.77523	0.74635	0.74397	0.69936	3.72	4.03	9.79	[8.02, 11.55]

Table 3: **Fresh finite-measure comparison at $\kappa = 1$.** Each frozen selector is evaluated using the same numerical Full-action functional on 128 paired validation trials. Reductions are computed relative to the Law action at the same κ. Tangent and GlobalReaction are descriptive structural comparators; the final column reports the paired interval for the prespecified Full-versus-Law action reduction.

The primary experiment fixes $\kappa = 1$ before validation and evaluates the frozen Law, Tangent, GlobalReaction, and Full geometries on the same 128 fresh paired trials across the five prespecified risk allowances. Table 3 cross-evaluates every selected geometry using the common Full-action functional. Tangent and GlobalReaction both reduce the common Full cross-score relative to Law, while the Full-selected geometry yields the lowest Full cross-score at every operating point. The Full-versus-Law reduction increases from 3.26% at the 1% allowance to 9.79% at the 5% allowance, with positive lower endpoints of the paired action intervals throughout. At the 5% operating point, the Tangent- and GlobalReaction-selected geometries reduce the same Full cross-score by 3.72% and 4.03%, respectively, compared with 9.79% for the Full-selected geometry. These cross-scores quantify the unrestricted transport–reaction robustness consequence of the different structural restrictions; they do not rank Tangent or GlobalReaction on their native objectives. Tangent and GlobalReaction are used here as descriptive structural comparators, and the reported action inference is for the prespecified Full-versus-Law contrast.

The mechanism ablations below examine the structural restrictions imposed by Tangent and GlobalReaction. Tangent retains only correction visible through the measured raw moment rates, whereas GlobalReaction allows unrestricted transport but restricts reaction to the global mass-rate discrepancy. Their common-Full cross-scores exceed those of the Full-selected geometries in Table 3. Section D.2.11 examines the corresponding Move–Reaction and Tangent–Hidden decompositions directly.

D.2 Ablations

We begin with two method-level ablations. The first examines FIDE's separation of completion from scoring by comparing independent per-time information projection with a restricted path-level completion at exactly fixed measured moments. The second asks when the baseline-preserving preference itself has a truth-relative downstream benefit: a controlled product-law construction makes the reference exactly correct in an unprotected factor, evaluates a held-out local-occupancy loss that is independent of Full's Sobolev geometry, and then deliberately violates baseline correctness. The remaining studies are benchmark-specific. The Gaussian common-pool analysis compares all secondary criteria on the same promoted candidates and separately perturbs the learned reference; held-out risk analyses assess scientific-risk transfer; prospective studies test the numerical evaluator and aggregate predictive interface; a manufactured solution tests continuum convergence of the periodic Full evaluator; skyrmion analyses examine richer potential spaces and support sensitivity; and the finite-measure analyses examine the transport–reaction extension. Unless stated otherwise, the relevant sensor geometries remain frozen throughout these diagnostics.

D.2.1 Static information projection versus a restricted path-level completion

FIDE deliberately separates completion from scoring: each moment fiber is first completed by information projection, after which Full measures the dynamical revision required by the resulting path. A natural alternative is to let unresolved law directions vary jointly over time so as to reduce Full itself. These constructions answer different questions, but the distinction raises a quantitative issue: how much of the Full action assigned to an information-projected path is attributable to the

measured moment trajectory, and how much can be changed by choosing a different completion of the same fibers? We examine this in a controlled one-dimensional example in which the measured trajectory is held exactly fixed while one otherwise unmeasured law direction is allowed to adapt dynamically.

Let the state space be the flat torus $\mathbb{T} = \mathbb{R}/(2\pi\mathbb{Z})$, with uniform frozen reference $\widetilde{q} = (2\pi)^{-1}$ and reference velocity $u = 0$. The measured observable is $\phi(x) = \cos x$, and we prescribe

$$c_n(t) = 0.4\sin(2\pi n t), \qquad n \in \{1, 2, 4\}. \tag{253}$$

The standard FIDE completion is the pointwise information projection

$$q_t^{\mathrm{IP}}(x) = \frac{\exp(\lambda_t \cos x)}{\int_{\mathbb{T}} \exp(\lambda_t \cos y)\, dy}, \qquad \mathbb{E}_{q_t^{\mathrm{IP}}}[\cos X] = c_n(t). \tag{254}$$

To construct a restricted path-level alternative, we enlarge the exponential family by one unmeasured Fourier coordinate,

$$q_t^{\mathrm{dyn}}(x) = \frac{\exp(\lambda_t \cos x + \gamma_t \cos 2x)}{\int_{\mathbb{T}} \exp(\lambda_t \cos y + \gamma_t \cos 2y)\, dy}, \qquad \mathbb{E}_{q_t^{\mathrm{dyn}}}[\cos X] = c_n(t), \tag{255}$$

and parameterize $\gamma_t = \sum_{k=1}^{5} a_k \sin(k\pi t)$. Thus $\gamma_0 = \gamma_1 = 0$, and both completions have the same measured moment trajectory and the same endpoint densities. For every candidate γ_t, the multiplier λ_t is recalibrated so that equation 253 holds exactly. We then optimize $a_1, \ldots, a_5 \in [-2, 2]$ to minimize the time-integrated Full action of this enlarged path. The static information projection is the feasible special case $a_1 = \cdots = a_5 = 0$, so this ablation changes only the completion of the unmeasured law direction; it does not change the observed moment trajectory.

Because $u = 0$, differentiation of the calibrated exponential family gives

$$\frac{\partial_t q_t}{q_t} = \dot{\lambda}_t\big(\cos x - c_n(t)\big) + \dot{\gamma}_t\big(\cos 2x - \mathbb{E}_{q_t}[\cos 2X]\big), \qquad \dot{\lambda}_t = \frac{\dot{c}_n(t) - \mathrm{Cov}_{q_t}(\cos X, \cos 2X)\dot{\gamma}_t}{\mathrm{Var}_{q_t}(\cos X)}. \tag{256}$$

For a periodic one-dimensional density path, let $F_t(x) = \int_0^x \partial_t q_t(y)\, dy$. The minimum-energy continuity flux has $q_t(x)\delta_t(x) = b_t - F_t(x)$, where

$$b_t = \frac{\int_{\mathbb{T}} F_t(x)/q_t(x)\, dx}{\int_{\mathbb{T}} 1/q_t(x)\, dx}, \qquad a_{\mathrm{full}}(t) = \int_{\mathbb{T}} \frac{(b_t - F_t(x))^2}{q_t(x)}\, dx. \tag{257}$$

This gives the Full action for both paths without finite-differencing the density in physical time.

The measured trajectory itself imposes a nonzero action independently of the completion rule. Indeed,

$$\dot{c}_n(t) = \mathbb{E}_{q_t}[-\sin(X)\delta_t(X)],$$

so Cauchy–Schwarz and $\mathbb{E}[\sin^2 X] \le 1 - c_n(t)^2$ imply

$$a_{\mathrm{full}}(t) \ge \frac{\dot{c}_n(t)^2}{1 - c_n(t)^2}. \tag{258}$$

We use the integral of equation 258 only as a universal moment-imposed lower bound; it need not be attainable by a feasible law path.

The completion choice is therefore consequential but does not account for all of the dynamical cost. At $n = 1$, allowing the additional unmeasured mode reduces action from 6.52 to 5.89, a 9.67% reduction, while both values remain well above the universal moment-imposed lower bound of 3.30. The reduction grows to 20.00% and 26.15% over the two more rapidly varying trajectories tested. This shows that part of the Full action assigned to the static information-projected path depends on how the unresolved fiber directions are completed rather than being forced solely by the measured moment trajectory.

This result does not make the path-level construction a replacement for the FIDE completion rule. The restricted dynamic path is optimized using the same Full criterion by which it is subsequently evaluated, whereas FIDE intentionally fixes the completion independently of Full and then asks how much dynamical revision that prescribed completion requires. The two procedures therefore encode

Frequency n	Moment lower bound	Static I-projection	Restricted dynamic	Reduction (%)
1	3.295850	6.521546	5.890663	9.67
2	13.183401	26.086185	20.867915	20.00
4	52.733603	104.344741	77.063457	26.15

Table 4: **Effect of the completion rule at fixed measured moments.** The static column evaluates Full on the independent per-time information projections in equation 254. The restricted dynamic column optimizes one additional unmeasured Fourier coordinate over time while preserving exactly the same trajectory $c_n(t)$. "Reduction" is $100(\mathcal{A}_{\mathrm{IP}} - \mathcal{A}_{\mathrm{dyn}})/\mathcal{A}_{\mathrm{IP}}$. The moment lower bound is the time integral of equation 258 and is not asserted to be attainable.

different modeling choices: the former uses unmeasured law directions to reduce dynamical action, while the latter uses a common local minimum-information-revision rule to resolve those directions before dynamical scoring. The ablation quantifies the price of that separation in this toy family rather than establishing that one completion is scientifically superior.

The numerical calculation uses 512 periodic spatial points and 121 physical-time nodes. Moment calibration, mass compatibility, and differentiated-moment compatibility hold to below 6×10^{-15} in the reported runs, and setting all dynamic coefficients to zero reproduces the independently evaluated static action to machine precision. Re-evaluation of the frozen solutions on successively finer spatial and temporal grids changes the reported completion gaps by less than 0.006 percentage points. For $n = 2$ and $n = 4$, the first coefficient reaches the prespecified lower bound -2; the increase in relative gap across the three tested frequencies should therefore be read descriptively rather than as a converged or universal frequency law. The dynamic construction is restricted to the family in equation 255 and is not claimed to attain the globally minimum-action path over all laws satisfying the moment constraints.

D.2.2 When does baseline preservation help? A controlled downstream ablation

The baseline-preserving specification is useful only if preserving the reference in directions left outside the scientific task has a meaningful consequence. We therefore construct a population-level ablation in which this assumption is exactly true and then violate it in a controlled way. The purpose is not to show that Full recovers truth in general (Proposition 1 rules out such a claim), but to ask two narrower questions: when the reference is correct in an unprotected part of the state, which secondary selectors protect a downstream quantity that is independent of their native objectives, and how do those consequences change as the same unprotected direction becomes genuinely misspecified?

A baseline-correct residual design problem. Let $(Y, Z) \in \mathbb{T}^2$ with uniform stationary reference $\widetilde{q} = (2\pi)^{-2}$ and $u = 0$, and write $e_k(x) = \sqrt{2}\cos(kx)$. At baseline correctness the true population is

$$p_{t,0}(y, z) = \frac{1 + \epsilon s(t) e_1(y)}{(2\pi)^2}, \qquad s(t) = \sin(2\pi t), \tag{259}$$

so the scientific signal evolves only in Y while the unprotected Z marginal is exactly the stationary reference. Candidate sensors deliberately mix the protected signal with different unresolved Z directions:

$$\phi_\alpha(y, z) = e_1(y) + \psi_\alpha(z), \qquad \psi_\alpha(z) = \sum_{k \in \{1,3,10\}} \sqrt{w_{\alpha,k}}\, e_k(z), \qquad \alpha \in \{0, \tfrac{1}{4}, \tfrac{1}{2}, \tfrac{3}{4}, 1\}, \tag{260}$$

where $w_\alpha = (1 - \alpha)(0, 90/91, 1/91) + \alpha(10/11, 0, 1/11)$. Every ψ_α has the same L^2 energy and the same gradient energy,

$$\sum_k w_{\alpha,k} = 1, \qquad \sum_k k^2 w_{\alpha,k} = 10,$$

but different inverse-transport energy

$$H_\alpha := \sum_{k \in \{1,3,10\}} \frac{w_{\alpha,k}}{k^2} \in \{0.11, 0.31, 0.51, 0.71, 0.91\}. \tag{261}$$

At $\gamma = 0$, every sensor has the same target moment $\mathbb{E}_{P_t}[\phi_\alpha] = \epsilon s(t)$ and exactly the same integrated sampling risk, $2 - \epsilon^2/2$. Scientific risk therefore leaves all five designs tied by construction; the ablation isolates only how a secondary criterion resolves that residual freedom.

The information projection factorizes as $Q_t^\alpha = Q_{Y,t}^\alpha \otimes Q_{Z,t}^\alpha$. For $C_\epsilon := 2\pi^2\epsilon^2$, the small-amplitude expansion gives

$$\mathcal{A}(\alpha) = \frac{C_\epsilon}{4}(1 + H_\alpha) + o(\epsilon^2). \tag{262}$$

To expose the corresponding hidden mechanism directly, we also evaluate a fixed-reference Sobolev rate error on the unprotected factor,

$$D_Z(\alpha) := \int_0^1 \sum_{m\geq 1} \frac{|\dot{m}_{m,c}^\alpha(t)|^2 + |\dot{m}_{m,s}^\alpha(t)|^2}{m^2}\, dt, \qquad D_Z(\alpha) = \frac{C_\epsilon}{4} H_\alpha + o(\epsilon^2), \tag{263}$$

where $m_{m,c}^\alpha$ and $m_{m,s}^\alpha$ are the cosine and sine moments of $Q_{Z,t}^\alpha$. Because the true Z marginal is stationary at baseline correctness, these are truth-relative spurious rates. Full and D_Z therefore have the same predicted strict ordering, with $\alpha = 0$ preferred. Tangent, residual L^2, and forward KL are independent of α at leading order because the corresponding L^2 and H^1 quantities were matched; static W_2 shares the local inverse-transport ordering and is therefore expected to agree with Full in this construction.

The numerical experiment uses the prespecified amplitudes $\epsilon \in \{0.10, 0.05, 0.025\}$ without parameter search. At every amplitude, Full and D_Z are strictly increasing in α and have Spearman rank correlation one. The maximum relative errors of the leading Full and D_Z expansions are respectively 2.88×10^{-4} and 9.83×10^{-4} at $\epsilon = 0.10$, and decrease to 1.8×10^{-5} and 6.14×10^{-5} at $\epsilon = 0.025$. Tangent, residual L^2, and KL have small finite-amplitude preferences for $\alpha = 0.75$, but their relative score ranges shrink with ϵ, consistently with their leading-order indifference. Static W_2 selects the same $\alpha = 0$ geometry as Full.

Held-out local occupancy: an independent downstream loss. The preceding D_Z diagnostic is intentionally aligned with the inverse-transport mechanism underlying Full, so it is not by itself an independent downstream test. We therefore freeze the five sensors and evaluate ordinary local population predictions that never enter scientific risk, information projection, or any secondary criterion. For 16 evenly spaced centers $\theta_j = 2\pi j/16$ and three prespecified localization scales $\kappa \in \{2, 4, 8\}$, define

$$g_{\kappa,j}(z) = \exp\big(\kappa[\cos(z - \theta_j) - 1]\big), \tag{264}$$

and measure

$$\mathrm{OccRMSE}_\kappa(\alpha) := \left[\frac{1}{16}\sum_{j=0}^{15} \int_0^1 \Big(\mathbb{E}_{Q_{Z,t}^\alpha}[g_{\kappa,j}] - \mathbb{E}_{P_{Z,t}}[g_{\kappa,j}]\Big)^2 dt\right]^{1/2}. \tag{265}$$

This loss is evaluated directly against the known true Z marginal. It contains no reference-relative rate discrepancy, H^{-1} norm, inverse Laplacian, correction energy, gradient normalization, or design-dependent weighting.

At baseline correctness, Full and static W_2 select $\alpha = 0$, while Tangent, residual L^2, and forward KL select $\alpha = 0.75$ in the finite-amplitude calculation. Law remains exactly indifferent among all five geometries. For the selector-level comparison, we therefore use an outcome-independent Law policy that draws uniformly from the complete set of five risk minimizers and then freezes the sampled geometry; its expected occupancy RMSE is the arithmetic mean of the five per-geometry RMSEs, not the RMSE of a mixture distribution and not a tie broken using the downstream loss. At $\epsilon = 0.10$ and $\kappa = 4$, the resulting RMSEs are 2.1492×10^{-3} for Full/static W_2, 4.3147×10^{-3} for the expected Law policy, and 5.3296×10^{-3} for Tangent/residual L^2/forward KL, corresponding to differences of -50.19% and $+23.52\%$ relative to expected Law. Across all three tested amplitudes and all three localization scales, the correct-baseline ordering remains Full/static W_2 < expected Law < Tangent/residual L^2/forward KL. Figure 3 retains the more detailed candidate-family view: occupancy error increases monotonically across the five frozen geometries, with $\alpha = 0 < 0.25 < 0.5 < 0.75 < 1$ at all three localization scales.

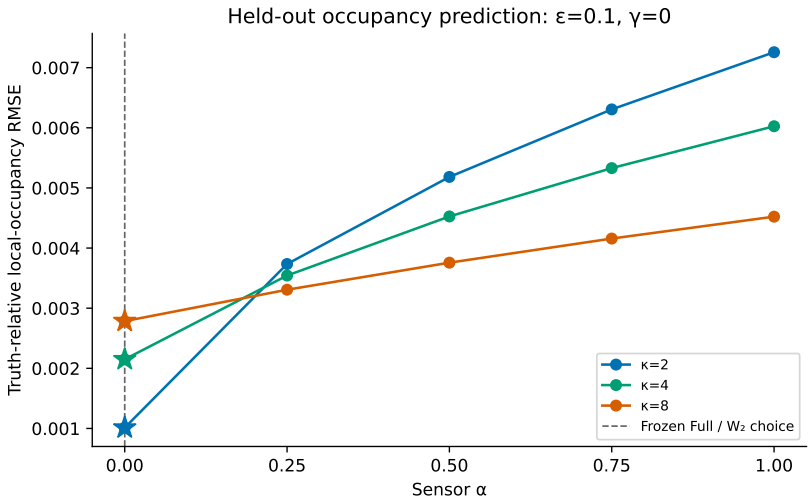

Figure 3: **Candidate-family occupancy error under baseline correctness.** At $\epsilon = 0.10$ and $\gamma = 0$, truth-relative local-occupancy RMSE is shown for all five frozen geometries and all three prespecified localization scales $\kappa \in \{2, 4, 8\}$. Scientific risk is identical across the five designs. Full and static W_2 select $\alpha = 0$, Tangent, residual L^2, and forward KL select $\alpha = 0.75$, and Law is indifferent among all five geometries. The occupancy loss is withheld from design and is not used by any selector.

The occupancy result is also predicted independently of Full. Writing I_k for the modified Bessel function,

$$\mathrm{OccMSE}_\kappa(\alpha) = \frac{\epsilon^2 e^{-2\kappa}}{8} \sum_{k\in\{1,3,10\}} w_{\alpha,k} I_k(\kappa)^2 + o(\epsilon^2). \tag{266}$$

This is a different spectral weighting from $H_\alpha = \sum_k w_{\alpha,k}/k^2$. The predicted squared-error ratios between $\alpha = 1$ and $\alpha = 0$ are 51.3866, 7.8609, and 2.6372 for $\kappa = 2, 4, 8$; at $\epsilon = 0.025$ the numerical ratios are 51.3868, 7.8610, and 2.6373. Thus the downstream agreement is not obtained by evaluating the Full objective under another name.

Deliberately violating baseline correctness. We next add a genuine unprotected departure to the true Z marginal while leaving the reference and sensor family unchanged:

$$p_{t,\gamma}(y, z) = \frac{1 + \epsilon s(t) e_1(y) + \gamma\epsilon r(t) e_1(z)}{(2\pi)^2}, \qquad r(t) = \sin(4\pi t). \tag{267}$$

The $\alpha = 0$ sensor has no $k = 1$ component in Z and is therefore blind to this new departure, whereas $\alpha = 1$ is strongly sensitive to it. No design is reselected as γ changes. The local expansion predicts that the $\alpha = 0$ versus $\alpha = 1$ ordering in D_Z reverses at $\gamma_\star = \sqrt{22/309} \approx 0.2668$; the numerical crossover estimates converge to this value as ϵ decreases.

The independently defined occupancy loss has its own, later, scale-dependent reversal. Its small-amplitude expansion is

$$\mathrm{OccMSE}_\kappa(w; \gamma) = \epsilon^2 e^{-2\kappa} \left[\frac{A_\kappa(w)}{8} + \frac{\gamma^2 B_\kappa(w)}{2} \right] + o(\epsilon^2), \tag{268}$$

with $A_\kappa(w) = \sum_{k\in\{1,3,10\}} w_k I_k(\kappa)^2$ and $B_\kappa(w) = I_1(\kappa)^2(w_1/2 - 1)^2 + \frac{w_1}{4}\{w_3 I_3(\kappa)^2 + w_{10} I_{10}(\kappa)^2\}$. This gives precomputed $\alpha = 0$ versus $\alpha = 1$ crossover predictions 0.5632, 0.5314, and 0.4482 for $\kappa = 2, 4, 8$. Using only the prespecified γ grid, the corresponding estimates at $\epsilon = 0.025$ are 0.56324, 0.53139, and 0.44817. Figure 4 shows the complete candidate-family sweep at $\epsilon = 0.10$: the baseline-preserving $\alpha = 0$ geometry is best near $\gamma = 0$, but its advantage erodes as the unprotected departure grows. Other frozen geometries eventually overtake it, and the $\alpha = 0$ versus $\alpha = 1$ ordering reverses at the separately predicted, scale-dependent thresholds.

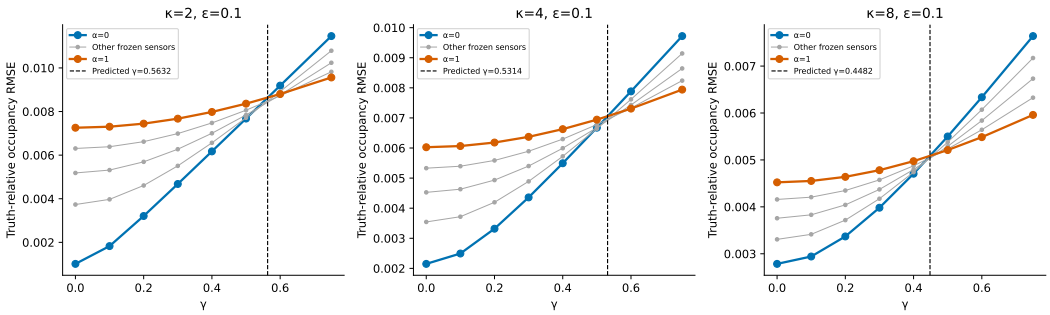

Figure 4: **Candidate-family occupancy error under misspecification.** At $\epsilon = 0.10$, truth-relative local-occupancy RMSE is shown across all five frozen geometries as the unprotected departure strength γ increases. Blue denotes $\alpha = 0$, orange denotes $\alpha = 1$, and gray denotes $\alpha \in \{0.25, 0.5, 0.75\}$. The baseline-preserving $\alpha = 0$ geometry is best near $\gamma = 0$ but eventually becomes worse than more departure-sensitive geometries. Dashed vertical lines mark the independently predicted small-amplitude $\alpha = 0$ versus $\alpha = 1$ occupancy crossovers from equation 268: $\gamma_\star = 0.5632$, 0.5314, and 0.4482 for $\kappa = 2, 4, 8$, respectively. No geometry is reselected as γ varies; Figure 1 gives the selector-level comparison emphasized in the main text.

These occupancy crossovers are substantially different from the D_Z crossover, confirming that the held-out occupancy loss probes a different downstream geometry. The main-text selector-level comparison in Figure 1 freezes the actual finite-amplitude selections, namely $\alpha = 0$ for Full/static W_2 and $\alpha = 0.75$ for Tangent/residual L^2/forward KL, and compares each with the expected Law tie-breaking policy. At $\epsilon = 0.10$ and $\kappa = 4$, both differences change sign between the evaluated points $\gamma = 0.5$ and 0.6; linear interpolation of the plotted segments gives 0.5156 for Full/static W_2 and 0.5160 for Tangent/residual L^2/forward KL. These interpolants are descriptive crossing locations rather than separately evaluated roots.

The same reversal directions hold for all nine combinations of $\epsilon \in \{0.10, 0.05, 0.025\}$ and $\kappa \in \{2, 4, 8\}$: Full/static W_2 change from lower to higher occupancy error than expected Law, while Tangent/residual L^2/forward KL change from higher to lower error. The crossing bracket is $[0.5, 0.6]$ for $\kappa \in \{2, 4\}$ and $[0.4, 0.5]$ for $\kappa = 8$ at every tested amplitude.

The main-text selector-level summary in Figure 1, together with the detailed candidate-family views in Figures 3–4, isolates the scope condition behind baseline-preserving residual design. When the task leaves a direction unprotected and the declared baseline is correct there, the geometry selected by Full and static W_2 preserves that structure and yields the smallest occupancy error among the selector groups considered here. As a genuine departure in the same unprotected direction grows, that advantage erodes and eventually reverses, while the geometry selected by Tangent, residual L^2, and forward KL exhibits the opposite change relative to the Law tie policy. No selector is therefore uniformly favored by this truth-relative downstream loss across baseline regimes. The result supports a conditional role for baseline preservation rather than a robustness guarantee under misspecification, and it does not establish a Full-specific advantage over static W_2 in this construction.

No parameter search is used. The three amplitudes, five sensors, three occupancy scales, and misspecification grid are fixed before evaluation. Across all nine amplitude–localization-scale combinations, the correct-baseline selector ordering Full/static W_2 < expected Law < Tangent/residual L^2/forward KL and the corresponding reversal directions under increasing misspecification are unchanged, although effect magnitudes and crossing locations depend on localization scale. All candidate-family rankings also persist over nine spatial–temporal refinements; the largest relative occupancy-metric change is 3.1×10^{-7}, direct quadrature and Fourier/Bessel evaluations agree within 10^{-12}, and the original baseline-ablation outputs are unchanged.

D.2.3 Do the held-out designs satisfy their scientific-risk allowances under sampling uncertainty?

Selection feasibility does not imply that the same relative risk allowance transfers to fresh observations. For studies whose current validation artifacts retain paired per-trial scientific losses, we assess

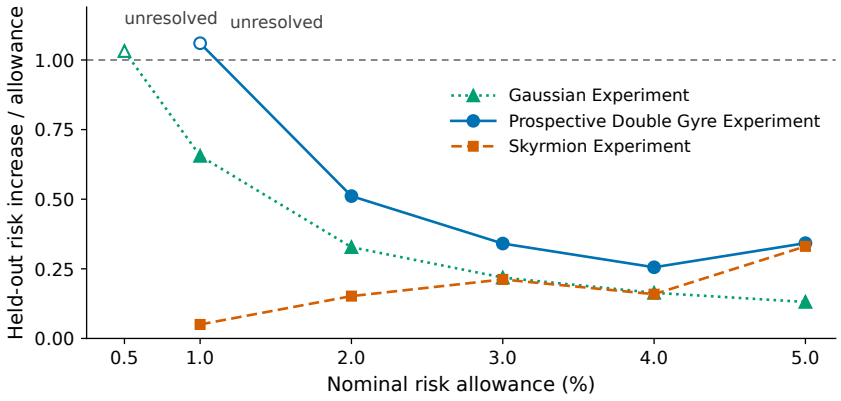

Figure 5: **Do the held-out designs satisfy their scientific-risk allowances under sampling uncertainty?** Descriptive held-out relative risk increase divided by the nominal allowance for the Gaussian-mixture, prospective double-gyre, and skyrmion validation studies. Gaussian and prospective double-gyre use their respective 64-trial independent holdouts, while skyrmion uses 32 fresh paired validation trials. The dashed line at one marks equality with the nominal allowance. Filled markers denote operating points whose simultaneous one-sided 95% upper bound for $\mathbb{E}[Z_i(p)]$ is below zero; open markers are unresolved. The plotted ratio is descriptive: inferential status is determined by the paired excess-loss statistic in equation 269, not by the ratio alone.

the held-out constraint directly. For paired evaluation losses, define

$$Z_i(p) = \ell_i(\boldsymbol{\eta}_{\mathrm{Full},p}) - \left(1 + \frac{p}{100}\right) \ell_i(\boldsymbol{\eta}_{\mathrm{Law}}), \tag{269}$$

so that $\mathbb{E}[Z_i(p)] < 0$ corresponds to satisfaction of the $p\%$ relative allowance on the evaluation distribution. For each benchmark family, we use 100,000 paired bootstrap resamples with common resampled trial indices across allowances. Writing $\bar{Z}_p$ for the observed paired mean and $\bar{Z}_p^*$ for its bootstrap replicate, the pointwise one-sided upper bound is $\bar{Z}_p + q_{0.95}(\bar{Z}_p^* - \bar{Z}_p)$, while the simultaneous bound replaces the scalar bootstrap deviation by its maximum over all allowances in that family. An operating point is supported below its held-out cap only when this simultaneous upper bound is negative; otherwise we report it as unresolved. Multiplicity control is applied separately within each benchmark family. The current paired-loss analysis covers the Gaussian, prospective double-gyre, and skyrmion validation studies. The current retrospective six-allowance frontier is treated separately because its retained evaluation summaries do not contain the paired per-trial Law and Full scientific-risk losses required for a matching bootstrap analysis.

Table 5 sharpens the held-out scientific-risk assessment for the Gaussian, prospective double-gyre, and skyrmion validation studies. In the Gaussian study, the 0.5% operating point is unresolved under simultaneous one-sided inference, whereas the 1%–5% points are supported below their Law-relative held-out caps. In the prospective double-gyre study, the 1% point likewise remains unresolved, with $\bar{Z} = 7.36 \times 10^{-4}$ and simultaneous upper bound 8.40×10^{-3}, whereas the 2%–5% operating points are supported below their held-out caps. All five fresh skyrmion operating points are supported below their corresponding caps. The current retrospective six-allowance frontier is not included in this bootstrap analysis because the retained evaluation summaries do not contain the paired per-trial scientific-risk losses required to construct matching excess-loss statistics; its verified holdout point estimates are reported in Appendix D.1.1.

The concentration of uncertainty at the tightest allowance is expected from the nested feasibility thresholds: for a fixed held-out loss pair, increasing p shifts $Z_i(p)$ downward by $(p/100)\ell_i(\boldsymbol{\eta}_{\mathrm{Law}})$ without comparably increasing its sampling variability. This is especially transparent for the prospective 2%–4% points, which use the same frozen Full geometry and have the same 1.022% held-out risk increase; the increasingly negative feasibility margin therefore reflects additional tolerance rather than a change in the validated design.

Study	Allowance	Risk increase (%)	Simultaneous upper	Conclusion
Gaussian	0.5%	0.516	1.61×10^{-4}	Unresolved
Gaussian	1%	0.656	-7.58×10^{-5}	Below cap
Gaussian	2%	0.656	-7.33×10^{-4}	Below cap
Gaussian	3%	0.656	-1.39×10^{-3}	Below cap
Gaussian	4%	0.656	-2.05×10^{-3}	Below cap
Gaussian	5%	0.656	-2.70×10^{-3}	Below cap
Prospective	1%	1.060	8.40×10^{-3}	Unresolved
Prospective	2%	1.022	-4.31×10^{-3}	Below cap
Prospective	3%	1.022	-1.66×10^{-2}	Below cap
Prospective	4%	1.022	-2.88×10^{-2}	Below cap
Prospective	5%	1.711	-3.26×10^{-2}	Below cap
Skyrmion	1%	0.050	-2.62×10^{-3}	Below cap
Skyrmion	2%	0.304	-6.52×10^{-3}	Below cap
Skyrmion	3%	0.636	-1.00×10^{-2}	Below cap
Skyrmion	4%	0.636	-1.52×10^{-2}	Below cap
Skyrmion	5%	1.653	-1.52×10^{-2}	Below cap

Table 5: **Held-out risk-feasibility inference.** Relative risk increase is $100(\bar{\ell}_{\mathrm{Full}}/\bar{\ell}_{\mathrm{Law}} - 1)$. The fourth column gives the simultaneous one-sided 95% upper bound for the paired excess-loss statistic $Z_i(p)$ within the corresponding benchmark family. A point is labeled below cap only when this bound is negative.

D.2.4 Common-pool comparison of Gaussian design criteria

To complement the criterion-specific optimization results, we compare all six objectives on a shared set of scientifically competitive designs. Let $\mathcal{C}$ be the union of the geometrically admissible candidates that reached the final evaluation stage under Law, Tangent, Full, residual L^2, forward KL, or static W_2. Across the six selectors and risk allowances, this gives 17 unique promoted candidates. The common-pool Full scores use the 64-trial selection bank, whereas Table 1 evaluates the same Law and Full geometries with the same numerical Full evaluator on an independent 64-trial holdout bank. Lower-fidelity prescreen candidates and raw optimization iterates are excluded so that every design in $\mathcal{C}$ has undergone the same final evaluation protocol.

At the 1% allowance, the operational feasibility threshold is

$$R_{\mathrm{sel}}(\boldsymbol{\eta}) \le 1.01 R_{\mathrm{Law}} = 0.06684759, \tag{270}$$

together with the population, support, and numerical requirements of the Gaussian experiment. All 17 candidates in $\mathcal{C}$ satisfy these conditions. For each criterion J, we define its common-pool selection by

$$\boldsymbol{\eta}_J^{\mathcal{C}} \in \underset{\boldsymbol{\eta} \in \mathcal{C}:\, \boldsymbol{\eta} \text{ feasible}}{\arg\min} \; J(\boldsymbol{\eta}). \tag{271}$$

Table 6 reports these selections and their cross-evaluated Full actions. Figure 6 provides the corresponding geometric view: panel (a) compares scientific risk and Full action across the common pool, while panel (b) locates the same designs in the two-angle sensor space.

Five distinct criterion-specific selections arise from the common pool, with only forward KL and static W_2 selecting the same candidate. Law and Full recover the geometries from the original comparison, while static W_2 moves to the candidate selected by forward KL. Under common Full cross-evaluation, the selected actions are 19.7193 for Full, 30.2016 for Law, 38.0653 for Tangent, 36.5877 for residual L^2, and 33.1209 for the shared KL/static-W_2 selection. These are selection-bank cross-scores, not values on a common native objective scale, and they should not be identified with the independent held-out scores in Table 1. The scientific risks remain closely matched across the promoted pool, while both the selected geometries and their Full cross-scores differ substantially.

D.2.5 Controlled reference misspecification in the Gaussian common pool

Appendix D.2.2 isolates the baseline-correctness mechanism in a controlled population setting with known truth and a predeclared failure threshold. Here we ask a complementary benchmark-specific

Criterion	Scientific risk	Own objective	Full action
Law	0.066186	0.066186	30.2016
Tangent	0.066511	2.10920	38.0653
Full	0.066586	19.71935	19.7193
Residual L^2	0.066478	52.50258	36.5877
Forward KL	0.066387	0.235235	33.1209
Static W_2	0.066387	0.134317	33.1209

Table 6: **Gaussian criterion-specific selections on a shared promoted candidate pool.** Each row gives the minimizer of one criterion over the same set of 17 scientifically feasible designs. The "own objective" column is reported on each criterion's native scale and is therefore not comparable across rows; the final column cross-evaluates every selected design using the common Full-action functional. Forward KL and static W_2 select the same candidate, whereas Law, Tangent, Full, and residual L^2 select distinct candidates.

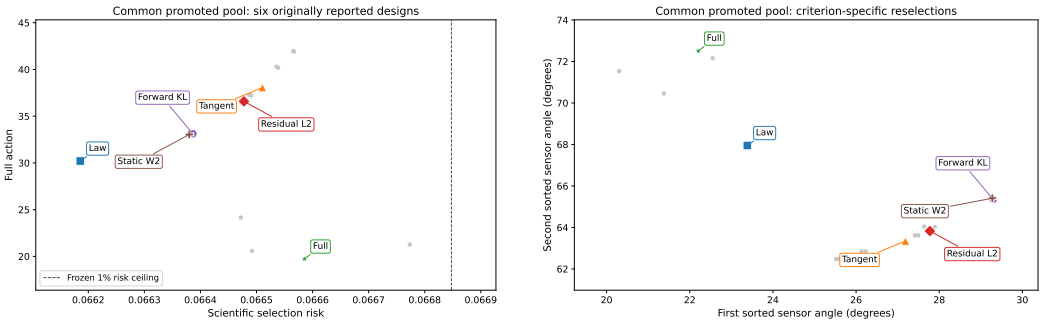

(a) Scientific risk versus Full action.

(b) Sensor geometries in the two-angle design space.

Figure 6: **Gaussian criterion-specific selections on the common promoted candidate pool.** Gray points denote the remaining promoted candidates and labeled markers denote the six criterion-specific minimizers obtained by reselecting each objective over the shared pool. **Left:** scientific risks occupy a narrow range while Full action varies substantially across candidate designs; the dashed line marks the frozen 1% operational risk ceiling. Forward KL and static W_2 select the same candidate and therefore overlap. **Right:** the corresponding sensor geometries occupy distinct regions of the design space, again with the forward-KL and static-W_2 selections coincident.

question: does the same reference sensitivity appear when the learned Gaussian reference itself is perturbed while the original scientific task and promoted candidate pool are held fixed? Starting from the frozen reference $(\widetilde{Q}_t, u_t)$, we construct five endpoint-preserving reference conditions by applying a smooth time-dependent rotation with $\gamma \in \{-\pi/6, -\pi/12, 0, \pi/12, \pi/6\}$, corresponding to $-30°, -15°, 0, +15°, +30°$. Let $\theta_\gamma(t) = \gamma \sin^2(\pi t)$ and let R_θ denote the planar rotation matrix. We define

$$\widetilde{Q}_t^\gamma = (R_{\theta_\gamma(t)})_\# \widetilde{Q}_t, \qquad u_t^\gamma(y) = \dot{\theta}_\gamma(t) J y + R_{\theta_\gamma(t)} u_t\big(R_{-\theta_\gamma(t)} y\big), \qquad J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

The transformed law–velocity pair therefore remains dynamically consistent, and because $\theta_\gamma(0) = \theta_\gamma(1) = 0$, every perturbation has the same endpoint reference laws as the original condition.

The purpose is to isolate reference dependence of the secondary selection rather than redefine the scientific task. We therefore keep the original 1% scientific-risk threshold, population screen, and 17 promoted candidates fixed across all five conditions. Every candidate remains admissible and numerically certified. Tangent, residual L^2, forward KL, static W_2, and Full re-rank this same pool after recalibration against the perturbed reference; Law remains the original action-independent risk minimizer. At $\gamma = 0$, all six selectors reproduce their original common-pool selections.

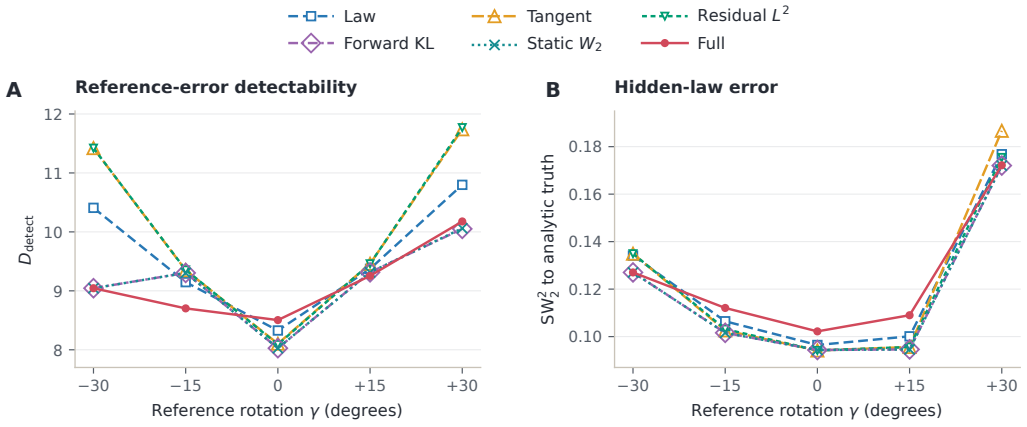

Figure 7: **Controlled Gaussian reference misspecification on the fixed 1% admissible pool. Left:** covariance-normalized detectability of the discrepancy between the analytic target and the perturbed reference; larger values mean that the selected sensors expose the reference error more strongly. **Right:** truth-relative hidden-law sliced-Wasserstein error of the corresponding information-projected completion; smaller values mean closer agreement with the analytic intermediate law. The two diagnostics need not induce the same ordering.

We measure sensitivity to the imposed reference error separately from each selector's native objective. For candidate $\boldsymbol{\eta}$, define

$$D_{\text{detect}}(\boldsymbol{\eta}, \gamma) = \frac{1}{B} \sum_{b=1}^{B} \sum_{q} \omega_q \, \Delta\mu_{bq}^{\top} V_{bq,\boldsymbol{\eta}}^{-1} \Delta\mu_{bq}, \qquad \Delta\mu_{bq} := \mathbb{E}_{P_{b,t_q}}[\boldsymbol{\Phi}_{\boldsymbol{\eta}}] - \mathbb{E}_{\widetilde{Q}_{t_q}^{\gamma}}[\boldsymbol{\Phi}_{\boldsymbol{\eta}}],$$

where $V_{bq,\boldsymbol{\eta}}$ is the existing measurement covariance and ω_q are the scientific-time quadrature weights. Larger D_{detect} means that the candidate measurements distinguish the analytic target from the perturbed reference more strongly. This diagnostic is distinct from truth-relative completion error, which we continue to measure using the original hidden-law sliced Wasserstein metric with 128 fixed directions.

The selectors separate clearly under misspecification. At all four nonzero rotations, the Full-selected geometry has lower D_{detect} than Law and ranks between 14th and 17th among the 17 admissible candidates by detectability. Forward KL and static W_2 are also less detectable than Law in three of the four conditions, with nonzero-condition ranks between 11th and 17th. Tangent and residual L^2, by contrast, are never less detectable than Law in this sweep and occupy ranks 4–8 and 1–8, respectively. Reduced sensitivity to a misspecified reference is therefore neither unique to Full nor a generic consequence of using the reference anywhere in the construction: it appears consistently for Full and often for the static reference-proximity criteria, but not for the two residual-rate criteria in this finite pool.

Detectability and truth fidelity remain separate. Forward KL and static W_2 have lower hidden-law SW_2^2 than Law at all four nonzero perturbations despite often selecting less detectable measurements. Full has higher hidden-law error than Law at the two $\pm 15^{\circ}$ conditions and lower error at the two $\pm 30^{\circ}$ conditions. Thus lower sensitivity to reference error does not by itself imply worse reconstruction of the hidden law, nor does higher detectability guarantee better truth-relative completion.

Finally, the rank changes should not be read as evidence that increasing misspecification causes Full to progressively move toward "blind" sensor geometries. Full selects the same geometry at 0°, $\pm 15^{\circ}$, and -30°, switching only to a nearby geometry at $+30^{\circ}$; much of the detectability-rank change occurs because the reference being compared against changes while the geometry does not. The ablation therefore supports a narrower conclusion: under the baseline-preserving preference, Full can favor scientifically admissible measurements that expose less of a misspecified baseline. This behavior is a limitation of the declared reference-preserving scope and should not be interpreted as a truth-recovery guarantee. The analysis is post hoc and conditional on this frozen reference, fixed

Evaluator	Law action	Full action	Reduction (%)	Paired Full−Law interval
128×64	1.50697	1.16113	22.95	$[-0.388, -0.307]$
192×96	1.50973	1.16290	22.97	$[-0.389, -0.308]$
256×128	1.51070	1.16352	22.98	$[-0.389, -0.308]$
Bandwidth $0.8\times$	1.63096	1.24506	23.66	$[-0.430, -0.344]$
Bandwidth $1.2\times$	1.38646	1.08121	22.02	$[-0.344, -0.269]$
PCG tolerance 10^{-10}	1.50697	1.16113	22.95	$[-0.388, -0.307]$

Table 7: **Is the prospective Full-action reduction stable to the numerical evaluator?** Mean numerical Full actions for the frozen Law and 5% Full geometries under spatial refinement, density–source bandwidth perturbation, and solver-tolerance perturbation. The reported paired intervals are for the Full-minus-Law action difference. No setting produces a numerical failure, and the action reduction remains between 22.02% and 23.66%.

promoted pool, and perturbation family, so it does not establish a universal robustness threshold or behavior under arbitrary reference misspecification.

D.2.6 Is the prospective Full-action reduction stable to the numerical evaluator?

The prospective result uses a spatially discretized density–source evaluator, so we test whether its Law–Full contrast depends materially on the spatial raster, smoothing bandwidth, or linear-solver tolerance. The Law and 5% Full geometries, observations, reference dynamics, moment reconstruction, and scientific-risk calculation remain frozen throughout; only the numerical action evaluator is changed.

As shown in Table 7, spatial refinement changes the Law–Full reduction only from 22.95% to 22.97% and 22.98% at 128×64, 192×96, and 256×128, respectively. Across the bandwidth and solver perturbations, the reduction remains between 22.02% and 23.66%, and every paired action-difference interval remains strictly below zero. No numerical failures occur, the minimum ESS fraction is 0.168, and the maximum Poisson residual is 1.02×10^{-10}. The reported Law–Full contrast is therefore stable across the tested numerical choices, although these checks do not by themselves establish convergence of the prospective density–source evaluator to the continuum Full functional. The manufactured-solution audit below separately tests continuum convergence for a periodic Full evaluator on an analytically tractable family without density estimation, smoothing, or observation noise.

D.2.7 Does the periodic Full evaluator converge to analytically known continuum action?

The preceding refinement study tests stability of an empirical Law–Full contrast but does not isolate discretization error from density–source preprocessing. We therefore perform a deterministic manufactured-solution audit on a periodic domain using the exact information-projected Fourier family of Appendix C.4. Let the frozen reference be uniform with $u = 0$, and for $k \in \{1, 2, 4\}$ define

$$q_{\lambda,k}(x) = \frac{\exp(\lambda \cos(kx))}{2\pi I_0(\lambda)}, \qquad \lambda(t) = \frac{1}{2}\sin(2\pi t). \tag{272}$$

The corresponding normalized continuity forcing is

$$h_{k,t}(x) = \dot{\lambda}(t)\left(\cos(kx) - \frac{I_1(\lambda(t))}{I_0(\lambda(t))}\right). \tag{273}$$

Writing $g(\lambda)$ for the exact transport metric in equation 227–equation 229, the continuum instantaneous action is

$$a_{\mathrm{full}}(t, k) = \frac{\dot{\lambda}(t)^2}{k^2} g(\lambda(t)). \tag{274}$$

The continuum benchmark is evaluated independently by adaptive one-dimensional quadrature and therefore does not use the discrete Poisson solver under test. Equation equation 274 also gives the exact integrated-action ratios $\mathcal{A}(1)/\mathcal{A}(2) = 4$ and $\mathcal{A}(1)/\mathcal{A}(4) = 16$.

Frequency k	Relative error at $N = 512$	Empirical spatial order
1	1.156×10^{-5}	2.00004
2	4.623×10^{-5}	2.00015
4	1.850×10^{-4}	2.00061

Table 8: **Manufactured-solution convergence of the periodic Full evaluator.** Relative errors compare the discrete time-integrated action with an independently quadrature-evaluated continuum benchmark. The empirical order is computed from the final two spatial refinements. At fixed N, the larger relative error at higher frequency is consistent with the expected dependence of a second-order spatial discretization on the dimensionless resolution kh.

The available periodic weighted-Poisson implementation is tested by extruding the one-dimensional family to the square torus as $q(x, y) = q(x)/(2\pi)$ and $s(x, y) = q(x)h(x)/(2\pi)$. Because the solution is independent of y, this extrusion leaves the continuum action unchanged. Exact cell-centered densities and sources are supplied directly, with no particle sampling, smoothing, density floor, source recentering, or observation noise. Spatial convergence is evaluated at $N \in \{32, 64, 128, 256, 512\}$ using 257 time nodes, while temporal convergence is evaluated separately on the finest spatial grid.

The evaluator exhibits approximately second-order spatial convergence for all three frequencies. At $N = 512$, the numerical action ratios are 3.999861 for $k = 1$ versus $k = 2$ and 15.997226 for $k = 1$ versus $k = 4$, compared with the exact values 4 and 16. The isolated temporal-quadrature error is approximately 6.7×10^{-13} already at 17 nodes and reaches floating-point roundoff under further refinement, so the finest-grid continuum errors are dominated by spatial discretization. The largest physical Poisson residual, 1.79×10^{-8}, occurs for the coarsest $N = 32$, $k = 4$ case, where only eight cells resolve each wavelength; for $N \geq 64$, the residual remains below 2.3×10^{-11} and no solver failures occur.

This manufactured-solution audit shows that the tested periodic Full evaluator converges to analytically known continuum action and reproduces the inverse-frequency transport scaling predicted by the theory.

D.2.8 Ablation: Which parts of the prospective predictive interface matter?

Prospective selection is conditional on three classes of predicted inputs: sensor-response means, single-particle response covariances, and scientific quantities used by the primary risk. This ablation isolates their roles using the 541 labeled sensor geometries discovered by the original prospective search, of which 406 satisfy the original geometric constraints, while holding the candidate pool fixed across perturbations. Fixing the pool removes variation from continuous design optimization and isolates how errors in the predictive interface change the ranking and admissibility of already-discovered measurement systems. The unperturbed condition is therefore a new selection restricted to this fixed pool rather than a reevaluation of the originally reported Law–Full pair, and its held-out action reduction need not equal that of the original prospective frontier. Under each perturbed interface, Law minimizes the corresponding scientific risk over the fixed pool and Full minimizes action within the condition-specific 5% Law-relative risk band and numerical-support constraints; all selected geometries are then frozen before evaluation with the common unperturbed post-freeze validation reference.

We consider three perturbation families. First, we add a fixed bias direction to the predicted sensor-response means at amplitudes 0.25 and 0.50 times the local response standard deviation. Second, we rescale the predicted single-particle sampling covariance by factors 0.75 and 1.25, leaving detector noise unchanged. Third, we perturb the predicted scientific-QoI targets by 0.25 and 0.50 frozen QoI scales along a fixed direction. Table 9 summarizes the resulting finite-pool selections and their held-out validation performance; the reported rank correlation is computed over the same 406 statically admissible geometries in every condition. Figure 8 complements this summary by showing how each perturbation reorganizes the design-time scientific-risk landscape over the common candidate pool as a function of minimum sensor separation.

Predictor condition	Feasible	Risk-rank ρ	Held-out risk change (%)	Held-out action reduction (%)
Unperturbed	246	1.000	2.030	24.324
Mean bias, 0.25 local SD	16	-0.294	-0.141	0.327
Mean bias, 0.50 local SD	1	-0.355	0.000	0.000
Sampling covariance $\times 0.75$	246	1.000	2.030	24.324
Sampling covariance $\times 1.25$	246	1.000	2.279	23.133
QoI-target bias, 0.25 scale	224	0.800	-1.749	26.747
QoI-target bias, 0.50 scale	46	0.746	0.641	25.642

Table 9: **Sensitivity of prospective selection to the aggregate predictive interface.** All conditions rank the same fixed pool of 541 candidate geometries. "Feasible" is the number of candidates satisfying the original numerical/support requirements and the condition-specific 5% Law-relative scientific-risk band. Risk-rank ρ is the Spearman correlation with the unperturbed condition over the 406 statically admissible geometries. Held-out quantities compare the Full and Law designs selected under the same predictor condition on the common post-freeze validation experiment.

As summarized in Table 9 and visualized in Figure 8, the covariance perturbations leave the design landscape essentially unchanged. At both covariance scales, the feasible set remains at 246 candidates, the scientific-risk rankings have correlation above 0.9999 with the unperturbed predictor, and the Full-selected geometry is unchanged. The corresponding held-out action reductions are 24.32% and 23.13%. Thus, within the tested range, moderate covariance rescaling has little effect on either admissibility or the secondary action ranking.

Systematic errors in the predicted sensor means have a markedly different effect. At the 0.25-SD perturbation, all 406 statically admissible candidates continue to pass the non-risk numerical and support requirements, but only 16 remain within the 5% scientific-risk band. The scientific-risk ranking over these valid geometries is negatively correlated with the unperturbed ranking, $\rho = -0.294$, and the selected Law and Full geometries become nearly identical, leaving only a 0.33% held-out action reduction. At the 0.50-SD perturbation, only one candidate remains feasible, so Law and Full necessarily coincide. This stronger perturbation also affects support: only 67 statically valid candidates satisfy the non-risk gates in both the perturbed and unperturbed conditions, yet their scientific-risk rankings are still strongly reversed, with $\rho = -0.761$. The sensitivity therefore persists after excluding both geometrically invalid designs and condition-dependent support failures.

The same mean perturbations also shift the apparent low-risk region of the fixed pool toward closely spaced sensors. Before applying the geometric constraint, none of the 246 unperturbed below-cap candidates violates the minimum-separation rule, compared with 15 of 31 candidates under the 0.25-SD perturbation and 30 of 31 under the 0.50-SD perturbation. Candidate-level decompositions show that the effect is not a uniform offset of the risk values: the perturbed sensor means change reconstructed moment trajectories and projected scientific quantities in a geometry-dependent way. For example, the original Law geometry changes from design-time risk 1.13 to 6.63 and 19.35 under the two mean perturbations, whereas the geometry selected under the larger perturbation changes from risk 8.52 in the unperturbed condition to 2.53. The dominant changes in the inspected candidates arise from squared mean QoI errors rather than trial-to-trial variance.

Perturbing the scientific-QoI targets produces a third, more localized mechanism. For a fixed geometry, these perturbations leave predicted sensor responses, reconstructed moments, projected law trajectories, and Full actions unchanged; only the scientific admissibility criterion changes. Consequently, the Full-action ranking is identical to the unperturbed ranking on the common pre-risk-valid pool, while the scientific-risk rank correlation decreases to 0.800 and 0.746 and the feasible set contracts to 224 and 46 candidates. The selected Full designs nevertheless retain held-out action reductions of 26.75% and 25.64%. This separates sensitivity of the scientific-risk constraint from sensitivity of the law-level action itself.

The held-out results also illustrate the scope of a relative risk constraint. Every Full design remains within 5% of the Law design selected under the same perturbed predictor, but this within-condition

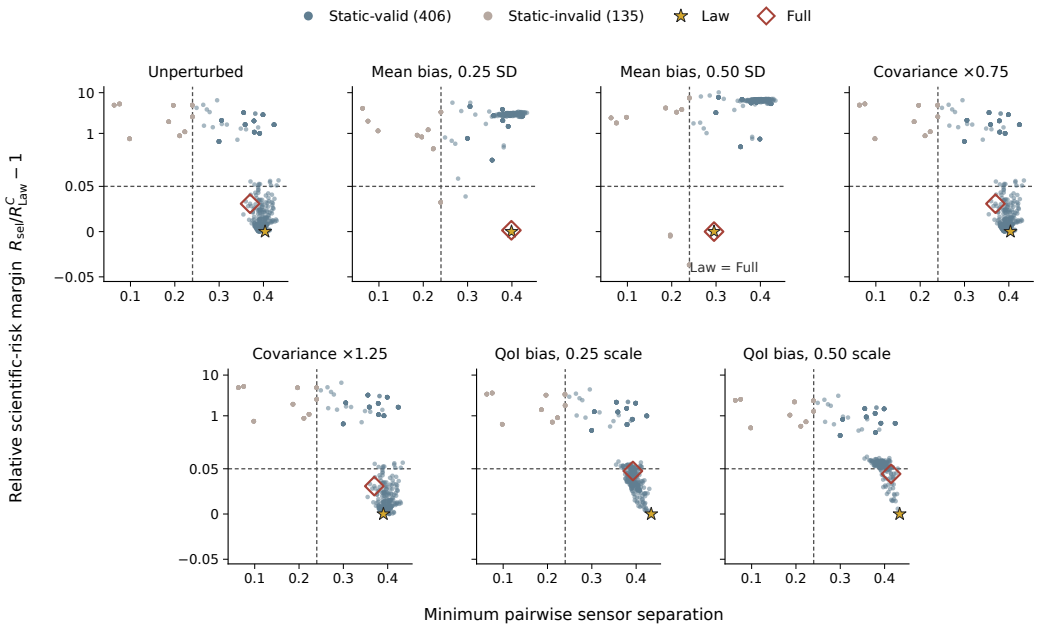

Figure 8: **Predictor perturbations reshape the prospective design-time risk landscape in different ways.** Each panel shows the same fixed pool of 541 candidate sensor geometries, with minimum pairwise sensor separation on the horizontal axis and scientific-risk margin $R_{\mathrm{sel}}/R^{\mathcal{C}}_{\mathrm{Law}} - 1$ relative to the condition-specific candidate-pool Law anchor on the vertical axis. The vertical dashed line marks the minimum allowed separation 0.24, and the horizontal dashed line marks the 5% Law-relative risk allowance. Highlighted markers denote the Law and Full designs selected from the fixed pool. The tested covariance rescalings leave the landscape nearly unchanged, whereas sensor-mean perturbations substantially reorganize the relative risk ordering and move an increasing fraction of the below-cap candidates into the separation-invalid region. QoI-target perturbations alter the scientific-risk landscape without changing the underlying reconstructed trajectories. Because each panel is normalized by its own Law anchor, vertical positions describe within-condition relative risk rather than absolute risk across perturbations. The two dashed thresholds alone do not define final feasibility, which also requires the condition-dependent numerical and support criteria.

comparison does not constrain absolute performance relative to the unperturbed design problem. In particular, the held-out Full risks under the two sensor-mean perturbations are 4.70 and 8.94, compared with 1.22 for the unperturbed Law design. Thus, the relative constraint continues to control the Full–Law tradeoff under a given predictive model, while systematic predictor error can shift the scientific-risk landscape shared by both methods.

The fixed-pool study identifies three forms of predictor sensitivity. The tested covariance rescalings leave both scientific and action rankings nearly unchanged; QoI-target shifts alter scientific admissibility while preserving the law-level action ranking; and sensor-mean bias changes the reconstructed moments themselves, substantially reorganizing both the scientific-risk landscape and the available Law–Full tradeoff. These results are conditional on the fixed discovered candidate pool, one selection reference and one separately trained post-freeze validation reference, and the specified perturbation directions.

D.2.9 Are the skyrmion action reductions robust to numerical approximation?

We test two possible explanations for the skyrmion action reductions: sensitivity to the potential dictionary used during design optimization, and disproportionate influence of a trial with concentrated importance weights. Sensor geometries remain frozen throughout, and all 32 paired validation trials remain in the primary analysis.

Risk allowance	Base	Fourier-400	Fourier-560	Interaction-416
1%	3.51	3.51	3.52	3.53
2%	6.91	6.92	6.93	6.95
3%	9.38	9.40	9.41	9.43
4%	9.38	9.40	9.41	9.43
5%	12.22	12.25	12.28	12.31

Table 10: **Does the action reduction persist under a richer potential space?** Percentage reductions in mean fitted action relative to Law under each evaluator. All 32 trials are retained, including failed-certification cases; the 3%/4% row represents one shared geometry. Fourier-400 and Fourier-560 refine the Fourier dictionary, while Interaction-416 adds interaction features. Fourier-560 has six held-out energy-check failures, and ESS support failures remain under every evaluator.

Design	Base ($K = 280$)	Fourier-560 ($K = 560$)	Increase (%)
Law	0.024433	0.024581	0.61
Full (5% allowance)	0.021449	0.021564	0.54

Table 11: **Absolute fitted action under nested Galerkin enrichment.** Mean fitted actions over the same 32 validation trials. In agreement with the monotonicity of the underlying restricted variational problem in Proposition 6, enlarging the Fourier potential space raises the fitted action for both frozen designs. The corresponding Full-versus-Law reduction changes only from 12.22% to 12.28%.

Does the action reduction persist under a richer potential space? A design optimized against a restricted Galerkin evaluator could appear favorable because the evaluator misses correction directions that matter more for that design. Proposition 6 makes this issue explicit: for a fixed projected law and continuity forcing, the exact restricted variational action is a lower bound on Full, and nested potential-space enrichments can only tighten that bound. This property alone does not show that the Base dictionary captures enough of the Full correction for reliable design comparison, so we re-evaluate the frozen frontier using four potential spaces: Base is the original dictionary of 248 one-particle Fourier functions and 32 radial pair functions, Fourier-400 and Fourier-560 retain these coordinates and append 120 and 280 Fourier functions respectively, while Interaction-416 instead appends 136 pre-specified distinct-particle products of low-frequency Fourier channels, testing pair-dependent potential directions rather than additional one-particle resolution. Within each trial and geometry, reconstructed moments, projected weights and continuity forcing are held fixed across evaluators, with fitting and audit banks kept separate. Original coordinates, normalization, solver settings and certification thresholds are unchanged; only the potential space is enriched.

For evaluator E, let $\overline{A}_E(\boldsymbol{\eta})$ denote mean fitted action over the 32 trials. We compare

$$\widehat{D}_{p,E} = 1 - \frac{\overline{A}_E(\boldsymbol{\eta}_{\mathrm{Full},p})}{\overline{A}_E(\boldsymbol{\eta}_{\mathrm{Law}})}, \qquad \Delta_{p,E} = 100\big(\widehat{D}_{p,E} - \widehat{D}_{p,\mathrm{Base}}\big), \tag{275}$$

where $100\widehat{D}_{p,E}$ is the percentage action reduction and $\Delta_{p,E}$ is its change in percentage points relative to Base. Reductions are ratios of mean actions, not means of trial-wise ratios. The 3% and 4% allowances select identical coordinates and are evaluated once.

The action reduction persists at every operating point under all four evaluators (Table 10), with a largest change of only 0.096 percentage points relative to Base. Full retains lower fitted action than Law in every individual trial under every evaluator, and the ordering across allowances is unchanged. Table 11 additionally shows the absolute-action change under the nested Base-to-Fourier-560 enrichment: both frozen designs acquire a small amount of additional fitted action while their relative contrast remains essentially unchanged.

Thus, the observed contrasts are stable under the tested Fourier and interaction enrichments, although these checks do not bound correction energy in untested potential directions.

Contrast stability is distinct from numerical certification. The positive, ordered frontier is reproduced by Fourier-400 and Interaction-416, both of which pass the reported held-out residual checks, so the enrichment result does not rely on the largest Fourier-560 system. Fourier-560 exceeds the held-out

Allowance	$\kappa = 0.25$	$\kappa = 1$	$\kappa = 4$
1%	3.02	3.26	2.97
2%	4.79	5.22	4.95
3%	6.27	6.88	6.59
4%	7.61	8.38	8.09
5%	8.85	9.79	9.49

Table 12: **Sensitivity of the Full-versus-Law contrast to the transport–reaction metric.** Entries are percentage reductions in mean Full action relative to Law, computed separately within each value of κ. Because κ changes the correction metric, absolute action values are not compared across columns.

energy threshold in six of 160 unique trial/design evaluations, identifying a numerical-certification limit of that larger basis rather than a reversal of the design ordering. Separately, the low-support validation pair fails one ESS support check for Law and all four distinct Full geometries under every evaluator. We therefore do not claim full numerical certification or continuum convergence; Proposition 6 characterizes what convergence of the underlying restricted variational problem would require, while the present experiment supports only that the observed Full-versus-Law contrast is stable under the tested one-particle and interaction-space enrichments.

Does the low-support trial drive the aggregate action reduction? We examine the low-support validation pair's influence using Base at the 5% allowance. Retaining all 32 pairs gives a 12.215% action reduction; excluding only this prespecified pair gives 12.380%, a change of $+0.165$ percentage points. Across all leave-one-pair-out calculations, the reduction ranges from 12.083% to 12.442%. The low-support validation pair ranks third in absolute deletion influence, but removing any single pair preserves the aggregate action reduction. Thus the low-support trial does not drive the reported contrast: every leave-one-pair-out analysis preserves an action reduction above 12%. These deletions are sensitivity checks rather than replacements for the primary analysis, and they do not resolve the trial's separate support-certification issue.

We additionally evaluate both frozen geometries on independent reference banks, preserving their saved moment trajectories and Base dictionary while recalibrating projected weights and forcing on each bank. Increasing the fitting bank from 65,536 to 131,072 and 262,144 configurations changes the low-support validation pair's fitted-action reduction from 9.685% to 8.421% and 8.877%, respectively. The trial-specific estimate varies with reference-bank size, but its influence on the aggregate contrast is small under this diagnostic: replacing only this pair's actions by the largest-bank estimates changes the overall reduction by -0.048 percentage points.

D.2.10 Sensitivity to the transport–reaction metric

In the unbalanced setup, the reaction penalty κ determines the relative geometry assigned to movement and reaction and is therefore a modeling choice rather than a parameter whose absolute action values can be compared across settings. We test whether the within-metric Full-versus-Law contrast depends strongly on this choice by evaluating the frozen construction at $\kappa \in \{0.25, 1, 4\}$ on the same shared physical validation bank. As shown in Table 12 and Figure 9, Full retains lower mean action than Law at every allowance under all three metrics. The reduction ranges from 2.97% to 3.26% at the 1% allowance and from 8.85% to 9.79% at the 5% allowance, with the intermediate allowances showing a similar pattern, indicating that the within-metric Full-versus-Law contrast is not tied to the default reaction penalty.

D.2.11 Unbalanced mechanism decompositions

The screened Full correction admits two complementary decompositions established in the finite-measure theory. First, the Full energy in equation 88 separates into Move and Reaction contributions, corresponding respectively to the transport and reaction terms of the optimal correction. Second, the orthogonal decomposition in equation 161 writes Full as Tangent plus Hidden, where Tangent is the minimum correction visible through the measured raw moment rates and Hidden is the remaining law-level component. GlobalReaction provides a separate restriction diagnostic: unlike Tangent,

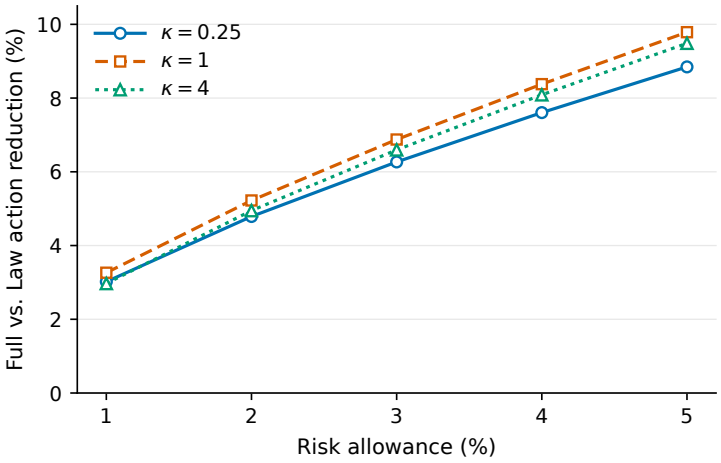

Figure 9: **The within-metric Full-versus-Law action reduction persists across reaction penalties.** Each curve reports the percentage reduction in mean Full action relative to Law at a fixed κ. Scientific observations are shared across the sweep, while κ changes only the transport–reaction correction metric; vertical comparisons of absolute action across κ are therefore not meaningful.

κ	Move / Full (%)	Reaction / Full (%)	Hidden / Full (%)	Full vs. Tangent (%)	Full vs. GlobalReaction (%)
0.25	21.99–22.65	77.35–78.01	43.15–43.80	1.95–6.12	2.23–6.28
1	23.97–24.21	75.79–76.03	64.58–65.51	2.57–6.30	2.16–6.00
4	38.79–39.33	60.67–61.21	73.51–74.62	2.27–5.62	1.76–5.20

Table 13: **Unbalanced mechanism summary across the 1%–5% allowances.** The first three quantities describe the Full-selected geometry and are reported as ranges over the five allowances. Move and Reaction form one decomposition of Full, while Tangent and Hidden form a separate orthogonal decomposition. The final two columns report the range of reductions in the common Full cross-score obtained by Full relative to the Tangent- and GlobalReaction-selected geometries.

which restricts the correction through the measured moment-rate equations, GlobalReaction restricts the admissible reaction field while retaining transport. Table 13 and Figure 10 show that the Tangent-invisible component remains substantial throughout the tested regime, while the common-Full cross-score of the Full-selected geometry is lower than those of both restricted comparators in every tested κ–allowance combination.

E Experimental details and reproducibility

Shared experimental conventions. For particle-based studies, a trial observes $\boldsymbol{y}_k = n^{-1}\sum_{i=1}^{n} \boldsymbol{\Phi}_{\boldsymbol{\eta}}(X_{t_k}^{(i)}) + \sigma_{\mathrm{obs}}\boldsymbol{\varepsilon}_k$ with $\boldsymbol{\varepsilon}_k \sim \mathcal{N}(0, I_R)$; one microscopic state is a tracer position in the transport studies and a complete interacting configuration in the skyrmion study. Reconstruction produces a differentiable moment trajectory that is information-projected relative to the frozen reference. Except for the stricter skyrmion development constraint stated below, allowance $p\%$ means $R_{\mathrm{sel}}(\boldsymbol{\eta}) \le (1 + p/100)R_{\mathrm{Law}}$, where R_{Law} is the selection risk of the numerically optimized Law design. For a fixed reference, $\widehat{D}_p = 1 - \overline{A}(\boldsymbol{\eta}_{\mathrm{Full},p})/\overline{A}(\boldsymbol{\eta}_{\mathrm{Law}})$, with $\overline{A}$ the trial-mean numerical Full action; when several frozen references are used, reference-wise effects are computed before aggregation as specified by the study. Selection feasibility, held-out scientific risk, numerical certification, and action improvement are reported separately. Numerical Full denotes the declared raster or Galerkin approximation, and absolute action values are not compared across systems or reference regimes. For normalized empirical projection weights $p_1, \dots, p_N$, $\mathrm{rESS} = (N\sum_i p_i^2)^{-1}$; the Gaussian and double-gyre studies require $\mathrm{rESS} \ge 0.03$ for their declared selection-time support checks

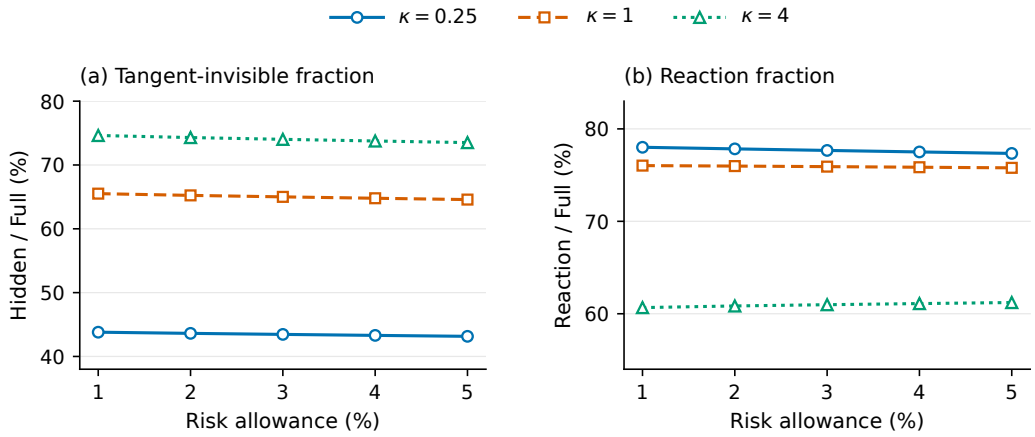

Figure 10: **Mechanism decomposition of the unbalanced Full correction. Left:** fraction of Full action lying in the Tangent-invisible Hidden component. **Right:** fraction of Full action allocated to reaction rather than spatial movement. Fractions are computed within each fixed transport–reaction metric; the figure does not compare absolute action values across κ.

and the skyrmion study requires 0.05. These thresholds diagnose support concentration rather than empirical moment feasibility, which is characterized separately by Proposition 3.

Reference flow matching. The Gaussian, retrospective and prospective double-gyre, and skyrmion studies use reference velocities trained by flow matching on endpoint pairs. For an experiment-specific endpoint coupling $(A_0, A_1) \sim \Pi$, bridge X_t^{br}, and $T \sim \mathrm{Unif}[0,1)$, we minimize

$$\mathcal{L}_{\mathrm{FM}}(\theta) = \mathbb{E}\left[\left\| f_\theta(X_T^{\mathrm{br}}, T) - \partial_t X_t^{\mathrm{br}}\big|_{t=T} \right\|_2^2\right]. \tag{276}$$

For the Gaussian benchmark, $X_t^{\mathrm{br}} = (1-t)A_0 + tA_1 + 0.15\sin(\pi t)Z$ with independently sampled endpoint states and $Z \sim \mathcal{N}(0, I)$. For both double-gyre studies, independently sampled empirical endpoints are first mapped to box-logit coordinates and connected by the noiseless linear bridge $X_t^{\mathrm{br}} = (1-t)A_0 + tA_1$; the learned latent velocity is mapped back to physical coordinates by the logistic Jacobian. For the skyrmion reference, endpoint configurations use the frozen particle-matched coupling and the periodic bridge $X_t^{\mathrm{br}} = [A_0 + t\Delta + 0.01\sin(\pi t)Z] \bmod L$, where Δ is the minimum-image displacement. After training, the reference is frozen and $\widetilde{Q}_t$ is represented by deterministic forward rollout of the learned velocity from the initial law; the training bridge itself is not used as the reference marginal path. The unbalanced experiment instead uses the fixed analytical reference described in Appendix E.5.

E.1 Information boundaries

The studies differ in the target information available during measurement design. Table 14 summarizes these regimes. In particular, the prospective double-gyre optimizer accesses intermediate target information only through a frozen aggregate interface: the interface is constructed offline from a separate microscopic simulation bank, but the underlying intermediate particles and target simulator are not available during sensor optimization.

E.2 Gaussian-mixture criterion comparison

Gaussian system and sensor family. For $d(\beta) = (\cos\beta, \sin\beta)^\top$, let

$$g_\beta(x) = \tfrac{1}{2}\mathcal{N}(x; 1.5d(\beta), 0.3^2 I_2) + \tfrac{1}{2}\mathcal{N}(x; -1.5d(\beta), 0.3^2 I_2)$$

and

$$\bar{p}_t^\alpha = (1-t)^2 g_0 + 2t(1-t)g_\alpha + t^2 g_{\pi/2}$$

Study	Information available during selection	Information introduced after design freeze
Gaussian mixture	Analytic target law, target-derived moments and covariances, noisy finite-sample observations, population-risk screen, and one frozen endpoint-trained reference.	Fresh target draws, particle sampling, and detector noise from the same analytic family. The reference is unchanged.
Retrospective double-gyre	Frozen microscopic target simulation bank and three endpoint-trained references; candidate observations and scientific-risk quantities are evaluated directly from the target bank.	Fresh particle resampling and detector noise, conditional on the same target bank and reference ensemble.
Prospective double-gyre	Endpoint ensembles, one frozen endpoint-trained reference, and fixed aggregate sensor-response means, cross-sensor covariances, and scientific-QoI predictions. These aggregate fields are precomputed from a separate microscopic simulation bank and its intermediate particles are not exposed to the optimizer.	A separately seeded evaluation reference trained from the same endpoint dataset, together with new microscopic validation trajectories, sampling, and detector noise for action and scientific-risk validation.
Skyrmions	A fixed development set of interacting configurations, one frozen endpoint-trained permutation-equivariant reference, aggregate sensor observations, and the nine permutation-invariant scientific features used to evaluate development risk.	A separate set of 32 fresh paired measurement trials for scientific-risk and action validation. The reference model and selected sensor geometries remain frozen. Richer potential dictionaries, leave-one-pair-out analyses, and larger reference banks are introduced only as post-selection numerical sensitivity checks.
Unbalanced Gaussian	A fixed analytical reference and finite-measure simulator, together with development observations and simulator truth used for scientific-risk evaluation.	A fresh bank of 128 paired trials with new physical nuisances, particle observations, detector noise, and log-mass noise. The same frozen bank is reused across the κ sensitivity analysis.

Table 14: **Information boundaries of the reported studies.** "After design freeze" denotes information unavailable when optimizing the reported sensor geometries. Fresh evaluation is conditional on the quantities explicitly retained in the corresponding row.

for $\alpha \in [\pi/6, \pi/3]$. Benchmark truth is the normalized mass of $\bar{p}_t^{\alpha}$ on a 51×51 grid in $[-3.2, 3.2]^2$. Two sensors have centers $z_j = 1.5d(\theta_j)$ on the radius-1.5 ring, angular separation at least $20°$, and responses $\Phi_{\boldsymbol{\eta},j}(x) = \exp(-\|x - z_j\|^2/(2\ell^2))$ with $\ell = 0.45$. Each trial observes 100 particles at eleven acquisition times with $\sigma_{\mathrm{obs}} = 0.01$, and endpoint-anchored quadratic generalized least squares reconstructs moments on 21 scientific times.

Reference, selection, and evaluation protocol. The reference is trained using equation 276 with four width-128 SiLU hidden layers for 12,000 Adam steps, batch size 2,048, initial learning rate 10^{-3} with cosine decay to 5×10^{-5}, and gradient clipping at 10. Scientific risk is Gaussian-kernel MMD^2 with bandwidth 0.55 on the unobserved time nodes. The frozen Law anchor is $R_{\mathrm{Law}} = 0.0662$ with population-risk ceiling 0.0692. Law, Tangent, Full, residual L^2, forward KL, and static W_2 are optimized at allowances $0.5, 1, 2, 3, 4, 5\%$ using the same information-projected construction; each receives seven common starts and 40 Adam steps per allowance with learning rate 0.012 and constraint penalty 10^4, together with the same number of final candidate evaluations. Selection uses 4 trials for gradient estimation, 32 for scientific-risk evaluation, 16 for preliminary action evaluation, and 64 for final objective evaluation. Full uses a 41^2 optimization raster and a 101^2 final evaluation raster with no added operator-density floor. Frozen geometries are evaluated on 64 paired holdout trials. Static W_2 was added after the original five-selector holdout had been inspected. It received the same search opportunity but reused that holdout and is therefore post hoc.

Allowance sweep and criterion-specific selections. Across the prescribed risk allowances, the secondary targets select different sensor geometries despite operating under the same scientific-risk restriction. At the 0.5% allowance, the Full-selected design reduces held-out Full action by 30.67% relative to Law; from 1% through 5%, Full selects the same geometry and gives a 33.59% reduction. The 1% cross-objective results are reported in Table 1. Appendix D.2.4 repeats the selection over a common promoted candidate pool and again produces criterion-specific geometries, showing that the separation is not solely a consequence of criterion-specific optimization trajectories. All reported designs satisfy their selection-time scientific and numerical requirements; held-out scientific-risk transfer is assessed separately in Appendix D.2.3, where the 0.5% point remains unresolved under simultaneous one-sided inference and the 1%–5% points are supported below their corresponding held-out caps.

E.3 Prospective double-gyre study

Double-gyre system. Passive tracers evolve on $\Omega = [0,2] \times [0,1]$ under the time-dependent double-gyre with $a(\tau) = 0.25\sin(2\pi\tau/10)$, $f(x_1,\tau) = a(\tau)x_1^2 + (1-2a(\tau))x_1$, $v(x,\tau) = 0.1\pi(-\sin(\pi f)\cos(\pi x_2), \cos(\pi f)\sin(\pi x_2)\partial_{x_1} f)^\top$, and $dX_t/dt = 10v(X_t, 10t)$. The initial population is a mixture of four truncated Gaussian components with coordinate standard deviation 0.07 and a 10% uniform background. The retrospective and prospective studies use this common transport system; their information boundaries and selection procedures differ as described below and in Appendix D.1.1.

Information boundary and scientific objective. The optimizer receives endpoint ensembles, one frozen endpoint-trained reference, and a frozen aggregate predictive interface constructed offline from a separate microscopic simulation bank; the underlying intermediate particles, hidden trajectories, and target simulator are not exposed during sensor optimization. On a 129×65 aggregate grid, the interface supplies the sensor-response mean $\widehat{\boldsymbol{\mu}}_{\boldsymbol{\eta}}(t)$, which provides the design-stage moment trajectory $\boldsymbol{c}_{\boldsymbol{\eta}}(t)$ used by FIDE, together with the full single-particle response covariance $\widehat{\Sigma}_{\boldsymbol{\eta}}(t)$ and predicted trajectories of the scientific quantities used by the design-stage risk. For $2{,}000$ sampled particles and detector-noise standard deviation 0.005, selection uses $\mathbb{E}[\boldsymbol{y}_t \mid \boldsymbol{\eta}] = \widehat{\boldsymbol{\mu}}_{\boldsymbol{\eta}}(t)$ and $\mathrm{Cov}(\boldsymbol{y}_t \mid \boldsymbol{\eta}) = \widehat{\Sigma}_{\boldsymbol{\eta}}(t)/2{,}000 + 0.005^2 I_4$, retaining cross-sensor sampling dependence. The scientific features are $\boldsymbol{\xi}(x_1,x_2) = (x_1/2, x_2, (x_1/2)^2, x_2^2, x_1x_2/2)^\top$. Writing $m_j(t)$ for their frozen aggregate predictions and $s_j = \max\{\mathrm{std}_t\, m_j(t), 0.02\}$, selection risk is $R_{\mathrm{sel}}(\boldsymbol{\eta}) = B^{-1}\sum_{b=1}^{B}\sum_{q=0}^{20}\omega_q\sum_{j=1}^{5}[(\mathbb{E}_{Q^{\boldsymbol{\eta}}_{b,t_q}}[\xi_j] - m_j(t_q))/s_j]^2$, where ω_q are normalized trapezoidal weights on the 21 scientific times. The final Law anchor is $R_{\mathrm{Law}} = 1.1344$ and is obtained entirely from selection-stage information.

Optimization, freeze, and validation. The selection reference uses $32{,}768$ rollout particles; Law and Tangent use 16 starts and 60 Adam steps, while Full uses 16 starts and 50 steps followed by higher-resolution evaluation of eight candidates and local refinement of three. Final selection uses 32 common trials and a 128×64 Full raster at all 21 scientific times. After the frontier for $p \in \{1,2,3,4,5\}$ is frozen, a separately seeded validation reference trained from the same endpoint dataset and 64 paired trials with fresh microscopic observations and hidden scientific targets are introduced, with no validation result used to alter the geometries or Law anchor. The Full-versus-Law action reductions are 13.71%, 19.05%, 19.05%, 19.05%, and 22.95%; the 2%–4% allowances select the same Full geometry and have simultaneous 95% interval $[16.72\%, 21.37\%]$, while the 5% interval is $[20.63\%, 25.27\%]$. At 5%, Full uses 2.834% additional selection risk and has mean Full action 1.16113, compared with 1.50697 for Law and descriptive Tangent cross-score 1.32591. Held-out scientific-risk increases are 1.060%, 1.022%, 1.022%, 1.022%, and 1.711%; simultaneous one-sided inference leaves the 1% point unresolved and supports the 2%–5% points below their corresponding caps.

E.4 Interacting-skyrmion study

System, observations, and scientific objective. One microscopic state is a configuration $X_t = (r_{1,t}, \ldots, r_{16,t}) \in \mathcal{X}_{\mathrm{sky}} := \Omega_{\mathrm{per}}^{16}$, where $\Omega_{\mathrm{per}} = (\mathbb{R}/2\mathbb{Z}) \times (\mathbb{R}/\mathbb{Z})$, so $\dim \mathcal{X}_{\mathrm{sky}} = 32$ and $\dot{r}_{i,t} = V_i(X_t, t)$ for $i = 1, \ldots, 16$. The implemented simulator uses 24 Euler–Maruyama substeps

per scientific interval, interaction strength 0.035 and range 0.16, pinning strength 0.055 and width 0.10, Magnus coefficient 0.32, and diffusion amplitude 0.006. Four width-0.12 aggregate sensors have periodic pairwise separation at least 0.20; each measurement averages 512 configurations and adds detector noise with standard deviation 0.002 at seven acquisition times, while reconstruction is evaluated at 13 scientific nodes. Each fresh validation truth contains $5{,}000$ configurations. Scientific risk compares nine permutation-invariant structural features—four radial-pair summaries, four structure factors, and one hexatic-order summary—using fixed covariance whitening and normalized trapezoidal time weights.

Reference, configuration-space approximation, and development protocol. The frozen reference is an endpoint-only permutation-equivariant flow trained for $6{,}000$ Adam steps with batch size 512; validation does not retrain or replace this reference. Full is approximated in configuration space by $\psi_{t,K}(X) = \sum_{a=1}^{K} \beta_a(t)\chi_a(X)$ and $\delta_{t,K}(X) = -\nabla_X \psi_{t,K}(X)$ with an immutable $K = 280$ dictionary containing 248 one-particle Fourier functions and 32 radial-pair functions. Risk, Galerkin fitting, and independent audit banks contain $32{,}768$, $65{,}536$, and $262{,}144$ complete configurations, respectively. The solver uses float64 arithmetic with Jacobi equilibration and relative eigenvalue cutoff 10^{-12}, and the prespecified support requirement is $\mathrm{rESS} \geq 0.05$. The reported frontier is optimized on 17 previously inspected development draws using the stricter paired development-risk ceiling $R_{\mathrm{dev}}(\boldsymbol{\eta}) \leq (1 + 0.8p/100)R_{\mathrm{Law}}$ for nominal allowances $p \in \{1, 2, 3, 4, 5\}$. Development action reductions are 4.04%, 7.94%, 10.70%, 10.70%, and 12.87%; the 3% and 4% operating points use the same selected geometry. These values describe the development frontier only and are separated from the fresh validation analysis.

Fresh validation and numerical sensitivity. After selection, the frozen frontier is evaluated on 32 new paired trials, with the reference model and sensor geometries unchanged. Full reduces mean fitted Galerkin action relative to Law by 3.51%, 6.91%, 9.38%, 9.38%, and 12.22% for nominal allowances 1% through 5%, while the corresponding mean scientific-risk increases are 0.050%, 0.304%, 0.636%, 0.636%, and 1.653%. Simultaneous one-sided held-out risk inference supports all five operating points below their nominal caps. Action uncertainty uses $100{,}000$ paired bootstrap resamples conditional on the fixed reference and shared quadrature samples; exploratory simultaneous 95% intervals have positive lower bounds throughout, including $[2.40\%, 4.61\%]$ at 1% and $[11.11\%, 13.32\%]$ at 5%. One validation pair falls below the prespecified $\mathrm{rESS} = 0.05$ support threshold, so complete numerical support certification is not claimed; all 32 pairs are retained, and every leave-one-pair-out analysis at 5% preserves an action reduction above 12%. Post-selection re-evaluation with richer Fourier and interaction potential spaces changes the reported reduction by at most 0.096 percentage points relative to the base evaluator; these checks assess stability under the specified numerical perturbations and do not establish Galerkin convergence.

E.5 Unbalanced Gaussian study

The finite-measure study uses a two-dimensional three-component Gaussian mixture on $\mathbb{R}^2$ over $t \in [0, 1]$. Writing $s(t) = \sin^2(\pi t)$ and $\theta_j(t) = 2\pi j/3 + \pi t/2$, the component centers are $u_j(t) = 1.35(\cos\theta_j, \sin\theta_j) + a\, s(t)(-\sin\theta_j, \cos\theta_j)$ and the common standard deviation is $\sigma(t) = 0.38 + 0.06s(t)$. Trial-level nuisances $a \sim U[-0.30, 0.30]$, $b_1, b_2 \sim U[-0.65, 0.65]$, and $c \sim U[-0.20, 0.20]$ perturb the intermediate shape and mass while vanishing at the endpoints. With $\beta = b_1(1, -1, 0) + b_2(1, 1, -2)/\sqrt{3}$, the component masses are $m_j(t) = e^{cs(t)}\,\mathrm{softmax}(s(t)\beta)_j$, so the total mass is $M_t = e^{c\sin^2(\pi t)}$. The frozen analytical reference sets all four nuisances to zero, giving equal component weights, unit mass, zero reference source, and the same rotating centers and time-dependent variance. No reference model is trained.

Two labeled Gaussian sensors of width 0.45 have centers in $[-2.2, 2.2]^2$ with minimum pairwise separation 0.45. Each trial observes 100 particles from the normalized shape at eleven acquisition times $j/10$ and adds independent detector noise with standard deviation 0.01; log mass is observed separately with noise standard deviation 0.02. Shape moments and log mass are reconstructed independently by endpoint-anchored generalized least squares using three $\sin^2(\pi t)$-weighted polynomial basis functions, after which the normalized shape is information-projected relative to the analytical reference and recombined with the reconstructed mass. Scientific risk is finite-measure Gaussian-

kernel MMD^2 with bandwidth 0.55, averaged over the ten non-acquisition interior times, so both shape and total-mass errors contribute to the admissibility criterion.

Law, Tangent, GlobalReaction, and Full are selected on 32 common development trials using constrained SLSQP with differentiable objectives and the full 1%–5% Law-relative scientific-risk allowances. Tangent uses the transport–reaction Galerkin correction generated by the two measured observables together with the constant mass coordinate. GlobalReaction, defined in equation 252, retains unrestricted transport but restricts reaction to the spatially uniform mass-rate correction, whereas Full solves the unrestricted screened transport–reaction problem. The authoritative action evaluator uses a cell-centered finite-volume discretization with no density smoothing or floor, float64 arithmetic, diagonally preconditioned conjugate gradients at relative tolerance 10^{-10}, and 193 trapezoidal action nodes. The primary $\kappa = 1$ evaluation uses a 97^2 spatial grid; the $\kappa = 0.25$ sensitivity evaluation uses 129^2, while $\kappa = 4$ again uses 97^2. Numerical certification checks moment calibration, projected effective sample size, captured Gaussian mass, projected covariance conditioning, screened and balanced PDE residuals, forcing centering, and the Move–Reaction and Tangent–Hidden energy identities.

After the five-point frontiers are frozen, validation uses one new bank of 128 paired trials shared across all methods, allowances, and $\kappa \in \{0.25, 1, 4\}$, with no pooling of development observations. The primary confirmatory analysis is the $\kappa = 1$ Full-versus-Law comparison. Action uncertainty uses $100{,}000$ paired bootstrap resamples with simultaneous intervals over the five allowances; held-out scientific-risk feasibility is assessed separately using simultaneous one-sided bounds for the paired excess-loss statistic. Tangent and GlobalReaction are retained as separate descriptive five-point comparator families. The κ sweep reuses the same physical validation bank so that changes across κ reflect the transport–reaction metric rather than different scientific observations; absolute action values are therefore compared only within a fixed κ.

E.6 Computing environment

Experiments were run locally on a system with an Intel Core Ultra 9 275HX processor, comprising 24 physical cores with maximum turbo frequency 5.4 GHz, and an NVIDIA GeForce RTX 5090 Laptop GPU with 24 GB of dedicated GDDR7 memory.

Execution routes differed across studies. JAX-based simulation, reference training, and selected optimization or evaluation stages used GPU acceleration where applicable; empirical information projection and several differentiable proxy calculations used native C++/OpenMP implementations.

The computational environment was hosted in Windows Subsystem for Linux 2 (WSL2), running Ubuntu 24.04.3 LTS with Linux kernel 6.6.87.2-microsoft-standard-WSL2. The principal implementation used Python 3.12.3 and JAX 0.8.3 with 64-bit numerical precision enabled. GPU acceleration was provided through the JAX CUDA 12 backend.

Empirical information projection used a C++ backend where specified by the study protocol. At a given trajectory node, for empirical reference particles x_i with base weights $\bar{w}_i$ and reconstructed target moment $\widehat{c}$, the double-precision Newton solver finds a finite multiplier λ satisfying $F(\lambda) = 0$, where $p_i(\lambda)$ is given by equation 43 and

$$F(\lambda) = \sum_{i=1}^{N} p_i(\lambda) \mathbf{\Phi}_{\boldsymbol{\eta}}(x_i) - \widehat{c}, \qquad \nabla_\lambda F(\lambda) = \mathrm{Cov}_{p(\lambda)}(\mathbf{\Phi}_{\boldsymbol{\eta}}, \mathbf{\Phi}_{\boldsymbol{\eta}}). \tag{277}$$

Thus failure to attain the requested moment and ill-conditioning of an otherwise calibrated interior solution are distinct numerical events: the former concerns empirical feasibility or solver convergence, whereas the latter concerns local stability of the multiplier coordinates and implicit derivatives. Multiplier estimates are warm-started between consecutive time nodes, while independent trajectories and linear-algebra operations are parallelized using OpenMP. The native solver is exposed to Python through pybind11 and integrated with JAX through Tesseract (Häfner & Lavin, 2025); Jacobian–vector and vector–Jacobian products use implicit solves with the moment covariance in equation 277 rather than differentiation through the individual Newton iterations.

The C++ extension was compiled in release mode using the GNU C++ 13.3 compiler, the C++17 language standard, and the optimization flags `-O3` and `-march=native`. Fast-math optimizations

were not enabled in order to preserve the numerical behavior of the double-precision projection solver. The complete software environment is summarized in Table 15.

Table 15: Hardware and software environment used for the experiments.

Component	**Configuration**
Operating system	Ubuntu 24.04.3 LTS under WSL2
Linux kernel	6.6.87.2-microsoft-standard-WSL2
Processor	Intel Core Ultra 9 275HX (24 cores, up to 5.4 GHz)
GPU	NVIDIA GeForce RTX 5090 Laptop GPU
GPU memory	24 GB GDDR7
NVIDIA driver	581.57
Python	3.12.3
JAX	0.8.3
jaxlib	0.8.3
Accelerator backend	CUDA 12 via the JAX CUDA plugin 0.8.3
CUDA runtime	12.9 (JAX-packaged runtime)
NumPy	2.5.1
SciPy	1.18.0
Tesseract Core	`tesseract-core` 1.11.0
Tesseract JAX interface	`tesseract-jax` 0.2.3
Projection backend	Native Tesseract C++/OpenMP batched trajectory solver
C++ language standard	C++17
C/C++ compiler	GCC/G++ 13.3.0
Build system	CMake 4.4.2, release configuration
Python/C++ bindings	pybind11 3.1.0
CPU parallelization	OpenMP 4.5
Native compilation	`-O3`, `-march=native`, without `-ffast-math`
Native numerical precision	IEEE 754 double precision (`float64`)
Neural optimization	Custom JAX implementation of Adam
JAX numerical precision	64-bit floating point enabled
Job scheduling	Local execution; no external scheduler