A reciprocal graph-deviation functional: structural identifiability from controlled equilibrium maps

Jonathan Washburn[1]‡ **and Megan Simons**[1]§

[1] Recognition Science, Recognition Physics Institute, 1401 Lavaca St #41100, Austin, TX 78701, USA

E-mail: `msimons@recognitionphysics.org`

Abstract. We constructed a reciprocal graph-deviation functional for positive ratios on a finite weighted graph. Its scalar generator $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ was characterized previously by ratio-composition and unit log-curvature. Our principal new result was structural identifiability: already on a two-site, one-edge system, the complete minimizer-location map obtained by sweeping the boundary and site-to-edge weight ratio determines any reciprocal generator whose log-coordinate representation is C^2 with everywhere-positive curvature, up to a positive scale; unit-curvature calibration fixes that scale. For the calibrated reciprocal family J_p we obtained the equilibrium response in closed form. A synthetic 15-condition sweep with seven simulated measurements per condition and log-scale noise $\sigma = 0.05$ selected J over $J_{1.1}$ with probability 0.997 under exact boundary and weight parameters; a declared representation-uncertainty screen lowered that value to approximately 0.88. These probabilities characterize the simulated design under the declared data-generating model. The finite sweep discriminated prespecified generators, whereas nonparametric recovery required the complete response map. We also recorded the associated convex gradient-flow and Langevin dynamics.

Keywords: reciprocal graph functional, structural identifiability, equilibrium map, inverse problem, gradient flow, stochastic dynamics

Submitted to: *Journal of Physics A: Mathematical and Theoretical*

‡ ORCID: 0009-0001-8868-7497
§ Corresponding author. ORCID: 0000-0001-9457-7019

1. Introduction

Positive ratio fields on weighted graphs arise whenever local quantities are measured relative to site-specific reference values. Given only controlled boundary conditions and the locations of the resulting constrained equilibria, can the scalar function that generates the graph potential be recovered? That is the inverse problem answered here. The forward problem is standard: fixing a generator, convexity yields a unique equilibrium for each boundary condition. The inverse problem is more subtle, because equilibrium locations are highly compressed summaries of the generator, and because every candidate with the same curvature at the reference produces the same leading near-equilibrium behavior.

Ratio data of this kind are common in biology: concentrations relative to a homeostatic baseline, expression levels relative to a control, compartment abundances normalized by capacity, firing rates relative to spontaneous activity. Such ratios are dimensionless, orientation-sensitive under naive linear scores, and naturally paired with their inverses, which is what motivates a reciprocal generator. These biological examples motivate the choice of functional form; the question settled here is mathematical.

Reciprocal and ratio-based potentials sit alongside a broader literature on accounting principles relating order and dissipation, from Schrödinger's negative entropy [1] to dissipative adaptation [2, 3], with open-reaction-network thermodynamics supplying signed work and free-energy balances for driven chemical systems [4, 5] and gradient-flow and large-deviation identities supplying a mathematical language for dissipative models [6, 7, 8]. Our concern is not a new dissipation identity but the weighted-graph use of a particular reciprocal state functional and, principally, the solution of its inverse problem.

The scalar characterization $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ among continuous reciprocal normalized costs satisfying a multiplicative d'Alembert (ratio-composition) identity and unit log-curvature was established previously by Washburn and Zlatanović [9]. We proved it here in the present formulation, so that the paper is self-contained, and use it as an input. The principal new result is Proposition 3: on a two-site, one-edge system, recovery of a reciprocal generator, uniquely after unit-curvature calibration, from the complete minimizer-location map, together with a finite-grid design for discriminating nearby members of the calibrated reciprocal family.

In the vocabulary of identifiability analysis, Proposition 3 is a structural identifiability statement in the classical sense of Bellman and Åström [10]: it asserts that the noiseless, population-level map from generator to observable equilibrium locations is injective on a specified class, which is a property of the model and the experiment design rather than of any dataset. Practical identifiability is a different question, concerning how sharply a generator can be pinned down from finitely many noisy observations, and for nonlinear models it is typically assessed by profile likelihood or related methods [11, 12]. Section 3 settles the structural question; Section 4 reports a finite-grid parametric feasibility calculation bearing on the practical one, with a full profile-likelihood treatment left to future work. Read as an inverse problem, the statement is closest to constitutive-function recovery from equilibrium response data; it is adjacent to inverse optimization on networks [13, 14] and to data-driven recovery of governing equations [15], but to our knowledge neither framework supplies the specific injectivity argument established here. Douglas's classical inverse calculus of variations [16] provides historical background. Graph propagation and the gradient-flow and Langevin identities are standard once the potential is fixed.

Seen as mathematical physics, the object of study is an inverse problem for a constitutive generator on a weighted network: the forward map is a controlled equilibrium response, and the question is whether that response determines the generator that produced it. Candidate realizations include equilibrium response in networked transport and in biochemical exchange models, where the state variable is naturally a ratio to a local reference and the control variables are boundary values and conductances. The identifiability theorem holds independently of any particular realization, and the physical assumptions a realization would require are separated out in Section 6.

Because the article combines theorems, modeling choices, and synthetic calculations, Table 1 separates the claims by category. Every probability quoted below is a simulated selection probability under a declared data-generating model; operating characteristics of a calibrated experiment await a physical realization. Section 6 collects the scope conditions of the construction in one place.

Table 1. Scope of claims. Rows separate results that follow from the stated definitions, choices adopted before any result applies, calculations performed on synthetic data, and questions deferred to future work.

Status	Content
Mathematical results	Scalar uniqueness of J (Theorem 1, proved here in the present formulation and previously established in [9]); termwise propagation to a weighted graph (Corollary 2); convexity and unique constrained minimizers (Proposition 1); local three-site separation (Proposition 2); structural identifiability from the complete two-site minimizer map (Proposition 3); the dissipation identity, exponential relaxation, Gibbs reversibility, and zero-noise limit of the specified overdamped model (Propositions 4–6).
Modeling assumptions	Locality and pairwise additivity; one common local shape ϕ for every site and edge term; graph topology, reference state, and site/edge assignment; the coupling λ; a constant positive-definite friction matrix.
Synthetic calculations	Every selection probability, Fisher and Cramér–Rao figure, and $\hat{p}$ simulation below: conditional design calculations at declared replicate count, noise, grid, and nuisance width, with calibration to a named platform deferred.
Deferred to future work	Measured data and a named analyte, membrane, or transporter; nonparametric recovery of g from a finite grid; a continuum limit of the finite-graph construction; a physical or thermodynamic interpretation of the functional, which awaits the calibration program sketched in Section 6.

Section 2 recalls the scalar cost and the graph functional; Section 3 establishes structural identifiability from the complete minimizer-location map; Section 4 quantifies what finitely many controlled conditions discriminate; Section 5 records the gradient-flow and Langevin realization; Section 6 states the limitations and what a physical reading would require; Section 7 concludes.

2. Reciprocal graph potential

Throughout, $x > 0$ denotes a multiplicative deviation: the ratio of an observed quantity to its reference value. Thus $x = 1$ is the absence of deviation, while x and x^{-1} are opposite deviations of equal multiplicative magnitude. We seek a continuous scalar $F : \mathbb{R}_{>0} \to \mathbb{R}$ that measures the cost of such a deviation, and we impose on it four requirements: reciprocity, normalization, composition, and calibration.

- **Reciprocity.** $F(x) = F(x^{-1})$ for all $x > 0$.
- **Normalization.** $F(1) = 0$.
- **Ratio-composition law (RCL).** Writing $\Phi := F + 1$,

  $$\Phi(xy) + \Phi(x/y) = 2\,\Phi(x)\,\Phi(y) \qquad (x, y > 0),$$

  equivalently $F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y)$. Near the reference the resulting cost is quadratic in $\log x$; away from it, RCL constrains all scales at once. That exact scale-consistency, not mere second-order agreement, is what will later distinguish J from the quadratic-log comparator Q.
- **Calibration.** Unit log-curvature at the reference, stated in the derivative-free form $\lim_{t\to 0} 2F(e^t)/t^2 = 1$.

Of the four requirements, composition carries the substantive structural content; the others fix symmetry, origin, and units. Under additional side conditions, reciprocity and normalization can be recovered from RCL rather than postulated independently [9], but the hypotheses needed for that recovery vary with the reduction, so we list them separately.

Multiplicative data alone do not imply RCL; we adopt it as an additional all-scale functional-form assumption. It is the assumption that distinguishes J from the nearby unit-curvature rivals introduced below (the quadratic-log cost Q, the relative-entropy cost D, and the reciprocal family J_p). The primary finite boundary/weight grid tests prespecified members of the J_p family; only the complete continuous map in Proposition 3 identifies an unrestricted generator.

Theorem 1 (Uniqueness of the reciprocal cost)**.** *Let $F : \mathbb{R}_{>0} \to \mathbb{R}$ be continuous and satisfy reciprocity, normalization, RCL, and calibration. Then $F = J$, where*

$$J(x) = \tfrac{1}{2}\big(x + x^{-1}\big) - 1 = \cosh(\log x) - 1. \tag{1}$$

Proof. Put $\mathcal{H}(t) = F(e^t) + 1$. Normalization gives $\mathcal{H}(0) = 1$, reciprocity makes $\mathcal{H}$ even, and the transform $x = e^t$, $y = e^u$ converts RCL into d'Alembert's functional equation $\mathcal{H}(t+u) + \mathcal{H}(t-u) = 2\mathcal{H}(t)\mathcal{H}(u)$ for this continuous even function. Its continuous solutions are classical [17]: $\mathcal{H}(t) = \cosh(kt)$, $\mathcal{H}(t) = \cos(kt)$, or $\mathcal{H} \equiv 1$, with second derivatives at the origin equal to k^2, $-k^2$, and 0 respectively. Since each branch is smooth and even with $\mathcal{H}(0) = 1$, Taylor's theorem turns the derivative-free calibration $\lim_{t\to 0} 2F(e^t)/t^2 = 1$ into $\mathcal{H}''(0) = 1$. That excludes the oscillatory cosine branch (nonpositive curvature) and the constant branch (vanishing curvature), and fixes $k = 1$ on the hyperbolic branch. Hence $F(e^t) = \cosh t - 1$, which is (1). □

Supplemental S1 records in full the regularity step behind the classification, that a continuous solution of d'Alembert's equation with $\mathcal{H}(0) = 1$ is automatically smooth; [9] is the prior source of this scalar characterization. We use J from here on. Equivalently, $J(x) = (x-1)^2/(2x)$. The function is nonnegative, vanishes only at $x = 1$, is strictly convex ($J''(x) = x^{-3} > 0$), and diverges as $x \to 0^+$ and as $x \to \infty$. Algebraically, $2J(x/y) = x/y + y/x - 2$ is the additive symmetrization of the scalar Itakura–Saito divergence, sometimes called the COSH divergence [18]. The Itakura–Saito divergence is the Bregman divergence generated by $-\log$ and has established information-geometric and clustering uses [19, 20].

2.1. Graph functional

Fix a finite set of sites X and a symmetric edge set $\mathcal{E}$ of unordered adjacent pairs. Equip each site s with a weight $w_s > 0$ and each edge $\{s, s'\}$ with a weight $w_{ss'} > 0$. A positive multiplier field is a map $m : X \to \mathbb{R}_{>0}$; a boundary condition pins a subset $X_\partial \subseteq X$ to prescribed positive values b. The uniform field assigns 1 to every site and is the unique absolute zero of the functional constructed below. Concretely, $m(s)$ may represent a local concentration, expression level, or abundance divided by its reference; the choice of reference is part of the model specification and must be fixed before the outcome is inspected. Collectively, the graph, weights, references, and measurement transforms constitute the representation $\mathcal{M}$.

Because site and edge attributes generally have different physical dimensions, their relative strength is controlled by an explicit site-to-edge coupling $\lambda > 0$ fixed from geometry, transport calibration, or separate training data. Let a_{ref} and κ_{ref} be positive reference scales fixed once for the entire protocol, and write $\bar{a}_s = a_s/a_{\mathrm{ref}}$ and $\bar{\kappa}_{ss'} = \kappa_{ss'}/\kappa_{\mathrm{ref}}$. The reference scales must not be renormalized separately in each sweep condition; otherwise a one-edge normalization would erase the intended conductance variation. Then

$$w_s = \bar{a}_s, \qquad w_{ss'} = \lambda\,\bar{\kappa}_{ss'}.$$

Equivalently,

$$C_\lambda[m] = \sum_s \bar{a}_s\, J(m_s) + \lambda \sum_{\{s,s'\}} \bar{\kappa}_{ss'}\, J(m_s/m_{s'}).$$

A common overall multiplicative scale of C_λ leaves every constrained minimizer unchanged and is therefore not identifiable from equilibrium locations. The relative coupling λ, by contrast, is not absorbed by that common rescaling, and is identifiable only relative to the fixed $(\bar{a}, \bar{\kappa})$ normalization and the experimental design; uncertainty in it propagates into every predicted minimizer and selection probability. Supplemental S4 collects the resulting non-identifiabilities: which quantities equilibrium locations and relaxation data each determine, and the fixed normalization convention they require.

Candidate physical attributes for the dimensionless measures $(\bar{a}, \bar{\kappa})$ include normalized compartment volume or molecule-count capacity for sites, and contact area, channel density, or measured transport conductance for edges [21]. Two objects are kept distinct from these physical graph weights and do not enter $C_{\mathcal{M}}$: the statistical inverse-variance weights of the observation model, which belong in the likelihood, and any global scale constant relating the functional to a physical energy (Section 6). The representation is frozen before any outcome is scored, a convention taken up again there.

Definition 1 (Reference-relative deviation functional)**.** For a positive multiplier field m on X,

$$C_{\mathcal{M}}[m] \;=\; \sum_{s\in X} w_s\, J\big(m(s)\big) \;+\; \sum_{\{s,s'\}\in\mathcal{E}} w_{ss'}\, J\!\left(\frac{m(s)}{m(s')}\right). \tag{2}$$

The edge term is well defined because $J(m(s)/m(s')) = J(m(s')/m(s))$ by reciprocity. Site terms penalize absolute deviation of each field value from its reference; edge terms penalize relative deviation between adjacent sites, so that $C_{\mathcal{M}} = 0$ means agreement with the reference. Throughout, $C_{\mathcal{M}}$ is a dimensionless bookkeeping functional rather than a thermodynamic free energy (Section 6).

Pinned-site constants and the reduced objective. The sum in (2) runs over all sites, so pinned boundary sites contribute constant terms $w_s J(b(s))$ to $C_{\mathcal{M}}$. Throughout we write $y_s = \log m(s)$ for the free (unpinned) log-coordinates, so that $J(e^t) = \cosh t - 1$ and each free site contributes a cosh term. To keep the two objects apart we write $C^{\mathrm{full}}_{\mathcal{M}}[m]$ for the complete sum (2) and, for a fixed boundary assignment b,

$$C^{\mathrm{free}}_{\mathcal{M},b}[y] := C^{\mathrm{full}}_{\mathcal{M}}[e^y] - \sum_{s\in X_\partial} w_s\, J\big(b(s)\big)$$

for the boundary-reduced objective on the free log-coordinates, so that $C^{\mathrm{full}}_{\mathcal{M}}$ and $C^{\mathrm{free}}_{\mathcal{M},b}$ differ by a constant depending on b alone. At fixed b we abbreviate $\widetilde{C} := C^{\mathrm{free}}_{\mathcal{M},b}$, and every free-coordinate objective and every reported $\Delta C_{\mathcal{M}}$ in Sections 3–4 refers to $C^{\mathrm{free}}_{\mathcal{M},b}$ at fixed b. Where no ambiguity arises we retain the plain symbol $C_{\mathcal{M}}$. Because the pinned constants cancel in every cost difference taken under the same boundary assignment and in the free-coordinate gradient, dropping them affects no minimizer and no reported $\Delta C_{\mathcal{M}}$. Throughout, $\Delta C_{\mathcal{M}}$ refers to changes in the free-coordinate part alone under a fixed boundary; if a boundary value varies during an interval, its cost change must be retained explicitly.

Pairwise additivity as a modeling ansatz. Equation (2) is a pairwise-additive graph ansatz. Theorem 1 selects the scalar cost J for a single positive ratio; it does not force locality, pairwise additivity, or the absence of higher-order (three-site or plaquette) interactions. Those are separate modeling assumptions, adopted here for simplicity (Table 1).

Granting the pairwise-additive ansatz, the scalar theorem propagates to the graph once a common local shape is assumed. We stated this as a corollary rather than a theorem: it applies Theorem 1 termwise, and its content resides in the hypotheses (L), (G), and (H) below.

Corollary 2 (Propagation of scalar uniqueness under a pairwise homogeneous ansatz)**.** *Let $C : (\mathbb{R}_{>0})^X \to \mathbb{R}_{\geq 0}$ be a deviation functional on a finite weighted graph satisfying:*

(L) Locality. *There are continuous functions $f_s \geq 0$ and $h_{ss'} \geq 0$ with $C[m] = \sum_s f_s(m_s) + \sum_{\{s,s'\}} h_{ss'}(m_s, m_{s'})$.*

(G) Ratio-invariant, reciprocal edges. *Each edge term depends on its endpoints only through their ratio and is inversion-invariant: $h_{ss'}(x, y) = \theta_{ss'}(x/y)$ with $\theta_{ss'}(r) = \theta_{ss'}(1/r)$. (We write "ratio-invariant" rather than "gauge-invariant" to avoid suggesting a stronger gauge symmetry than dependence on the endpoint ratio alone.)*

(N) Normalization. *$C[m] = 0$ if and only if $m \equiv 1$ (listed for clarity; it follows from (L), (G), (H), and (A) once all weights are positive).*

(H) Homogeneity. *A single continuous scalar shape ϕ governs every local term: $f_s = \omega_s \phi$ and $\theta_{ss'} = \eta_{ss'}\phi$ for constants $\omega_s, \eta_{ss'} > 0$.*

(A) Scalar axioms. *ϕ satisfies the scalar hypotheses of Theorem 1 (reciprocity, normalization, RCL, and unit-curvature calibration).*

Then $\phi = J$, and writing $w_s = \omega_s$, $w_{ss'} = \eta_{ss'}$, one has $C = C_{\mathcal{M}}$. Within the class fixed by (L), (G), and (H), the scalar axioms leave only the weights as free parameters of the graph functional. That conclusion is conditional on the class: the graph topology, the choice of reference state, the assignment of which variables are sites and which contacts are edges, and the homogeneous-shape ansatz itself are modeling choices made before (L)–(H) apply, not consequences of the scalar axioms.

Proof. By locality (L) and homogeneity (H),

$$C[m] = \sum_s \omega_s\, \phi(m_s) + \sum_{\{s,s'\}} \eta_{ss'}\, \phi\left(\frac{m_s}{m_{s'}}\right),$$

where the edge form follows from (G): $h_{ss'}(x, y) = \theta_{ss'}(x/y)$ with $\theta_{ss'} = \eta_{ss'}\phi$ and $\phi(r) = \phi(1/r)$. Hypothesis (A) places ϕ under Theorem 1, so $\phi = J$. Renaming $w_s = \omega_s$ and $w_{ss'} = \eta_{ss'}$ yields $C = C_{\mathcal{M}}$. Moreover (N) is then automatic: each summand is a nonnegative multiple of J and vanishes only at argument one, so $C[m] = 0$ if and only if $m \equiv 1$. □

Proposition 1 (Static minimization)**.** *Let $\mathcal{M}$ have weights $w_s, w_{ss'} > 0$, and let (X_∂, b) be a boundary condition with nonempty free set. Then:*

(i) $C_\mathcal{M}[m] \geq 0$ *for every field* m*;*

(ii) over unrestricted fields, the uniform field is the unique absolute minimizer and has $C_\mathcal{M} = 0$*;*

(iii) for every fixed boundary assignment there is a unique constrained minimizer on the feasible set;

(iv) if at least one pinned value differs from one, the constrained minimum of $C_\mathcal{M}^{\mathrm{full}}$ *is strictly positive, since that pinned site contributes a positive constant.*

Proof. Every summand is a nonnegative multiple of J and vanishes only when its argument equals one, which gives (i), the unique absolute zero at $m \equiv 1$ in (ii), and (iv), since a pinned value $b(s) \neq 1$ contributes the positive constant $w_s J(b(s))$. For (iii), in log-coordinates the free-site Hessian is $\operatorname{diag}(w_s \cosh y_s) \succ 0$ and the edge terms are convex, so $\widetilde{C}$ is strictly convex on every affine feasible set; the same cosh terms make it coercive, so the constrained minimizer exists and is unique. Supplemental S2 gives the argument in full. □

Thus the functional (2) is nonnegative and strictly convex in log coordinates, with a unique constrained minimizer for each fixed boundary. That constrained minimizer is the object the rest of the paper works with: Section 3 asks what a sweep of boundaries and weights reveals about the generator, Section 4 quantifies what finitely many such conditions discriminate, and Section 5 records the dynamics for which the cost is the potential.

3. Structural identifiability from controlled boundaries

The uniqueness of J among scalar costs satisfying these assumptions does not by itself identify J from data. Near the reference every continuous unit-curvature cost agrees with J through quadratic order in $\log x$, so a near-equilibrium design confined to small deviations separates J weakly from the quadratic-log comparator Q and from nearby members of the reciprocal family J_p. Identifiability therefore requires a controlled sweep containing asymmetric weight configurations, which break the relevant reflection symmetry and expose higher-order structure in the stationarity map; boundary magnitude alone does not, since in the equal-weight two-site geometry with boundary log-ratio u the stationarity condition is symmetric about $v = u/2$ for every reciprocal generator. We therefore consider first a local three-site separation at one fixed operating point, then a structural recovery statement from a boundary/weight sweep, and finally a brief analysis explaining why the equal-weight symmetric geometry cannot do the job by location alone.

The two unit-curvature reciprocal comparators used below are

$$Q(x) := \tfrac{1}{2}(\log x)^2, \qquad J_p(x) := \frac{\cosh(p \log x) - 1}{p^2} \quad (p > 0).$$

Thus $J_1 = J$, while J_p converges pointwise to Q as $p \to 0$. Figure 1 summarizes the two controlled geometries and shows why large, asymmetric drives are informative.

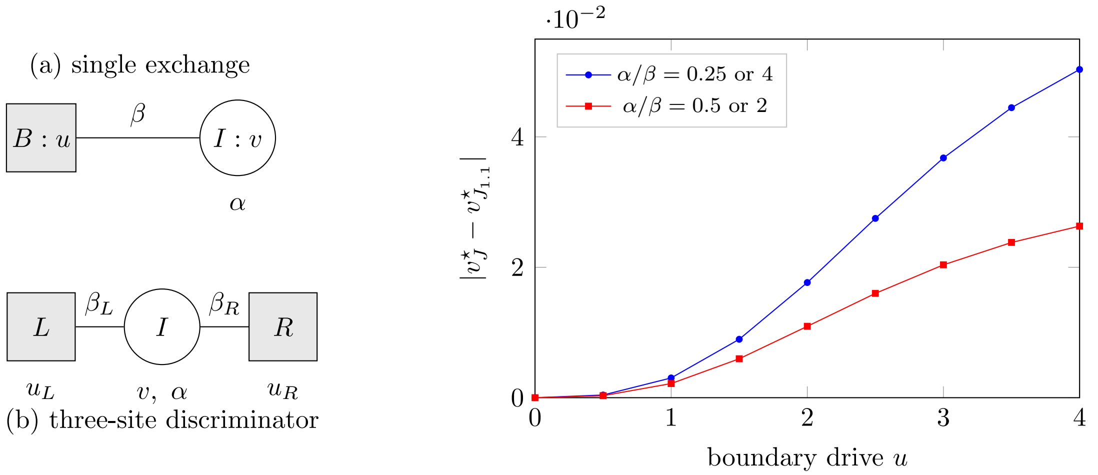

Figure 1. Controlled geometries and the source of finite-sweep discrimination. Left: pinned sites are shaded and the free site is open; site and edge weights are labeled. Panel (b) is the chain L–I–R (site names only; the friction matrix R is unrelated). Right: predicted equilibrium-location separation between J and $J_{1.1}$ grows beyond the shared near-reference quadratic regime. Reciprocal weight ratios give equal absolute separation; $\alpha/\beta = 1$ gives exactly zero separation at every drive and is retained as a quality control.

3.1. Asymmetric three-site objective

Consider the chain L–I–R with interior site I free and ends pinned. Let $v = \log m_I$, $u_L = \log m_L$, $u_R = \log m_R$, with free-site weight $\alpha > 0$ and edge weights $\beta_L, \beta_R > 0$. Unequal pinned boundaries and unequal edge weights break the reflection symmetry that, in the equal-weight two-site problem with boundary log-ratio u, takes the form $v \mapsto u - v$ and forces every reciprocal cost to share the geometric-mean critical point.

Proposition 2 (Local three-site separation of J from Q)**.** *The one-free-site objective for a scalar cost with $\Phi_0(t) := F(e^t)$ is*

$$C_I(v) = \alpha\,\Phi_0(v) + \beta_L\,\Phi_0(u_L - v) + \beta_R\,\Phi_0(v - u_R). \tag{3}$$

For J ($\Phi_0 = \cosh - 1$) the minimizer v_J solves

$$\alpha \sinh v - \beta_L \sinh(u_L - v) + \beta_R \sinh(v - u_R) = 0, \tag{4}$$

whereas for the quadratic-log comparator Q ($\Phi_0 = \frac{1}{2}t^2$) it solves

$$\alpha v - \beta_L(u_L - v) + \beta_R(v - u_R) = 0, \qquad v_Q = \frac{\beta_L u_L + \beta_R u_R}{\alpha + \beta_L + \beta_R}. \tag{5}$$

The two objectives share the same quadratic Taylor polynomial at the reference, while their minimizer maps first differ at cubic order in the boundary amplitudes; $v_J \neq v_Q$ for all parameters outside a measure-zero set.

Proof. Both $\cosh -1$ and $\frac{1}{2}t^2$ are strictly convex and the arguments in (3) are affine in v, so each C_I is strictly convex with a unique critical point, giving (4)–(5). By strict convexity of the J-objective, $v_J = v_Q$ holds if and only if v_Q is also its critical point, i.e. if and only if

$$\mathcal{A}(u_L, u_R, \alpha, \beta_L, \beta_R) := \alpha \sinh v_Q - \beta_L \sinh(u_L - v_Q) + \beta_R \sinh(v_Q - u_R) = 0,$$

with v_Q given by (5). The vanishing of the $O(\|u\|^3)$ term alone does not establish this; we argue instead from analyticity. The map $\mathcal{A}$ is real-analytic in $(u_L, u_R, \alpha, \beta_L, \beta_R)$ on $\alpha, \beta_L, \beta_R > 0$ and is not identically zero: at $\alpha = \beta_L = \beta_R = 1$, $u_R = 0$, $u_L = 3s$ one has $v_Q = s$ and

$$\mathcal{A} = 2\sinh s - \sinh 2s = 2\sinh s\,(1 - \cosh s),$$

which is nonzero for every $s \neq 0$. The zero set of a real-analytic function that is not identically zero has empty interior and Lebesgue measure zero, so $\mathcal{A} \neq 0$, and hence $v_J \neq v_Q$, for all $(u_L, u_R, \alpha, \beta_L, \beta_R)$ outside a measure-zero set. To locate the leading separation, substitute $\sinh t = t + t^3/6 + \cdots$ into $\mathcal{A}$; the linear part vanishes by the definition of v_Q in (5), and with $C_I''(v_Q) = \alpha + \beta_L + \beta_R + O(\|u\|^2)$ and $\|u\| := \max\{|u_L|, |u_R|\}$ one obtains

$$v_J - v_Q = -\frac{\alpha v_Q^3 - \beta_L(u_L - v_Q)^3 + \beta_R(v_Q - u_R)^3}{6(\alpha + \beta_L + \beta_R)} + O(\|u\|^5),$$

so the minimizer maps agree through quadratic order in $\|u\|$ and first differ at cubic order. The objectives share their quadratic Taylor polynomial and first differ at quartic order, since $J(e^t) = \cosh t - 1 = t^2/2 + t^4/24 + O(t^6)$ and $Q(e^t) = t^2/2$ both have vanishing cubic term; the induced minimizer maps nonetheless first differ at cubic order in the boundary amplitudes, because stationarity differentiates the objective. □

The cubic near-reference separation accounts for the local proximity of J and Q: a design confined to small deviations sees a shared parabolic well. The numerical design of Section 4 therefore works away from that regime, obtaining its comparator predictions by solving the stated stationarity equations at large fixed boundary values and weight asymmetry. Stationarity conditions for the additional comparators S, D, and H_{el}, defined in Section 4.2, are recorded in Supplemental Material S6.

3.2. Structural recovery of the generator

Proposition 2 separates J from one comparator at one configuration, and only at cubic order near the reference. A stronger, non-perturbative statement holds on a two-site, one-edge system: sweeping the boundary and weight ratio over their entire continuous range recovers the whole reciprocal generator from minimizer locations alone. This is a structural (population-level, noiseless) identifiability statement; it concerns the complete minimizer-location map, not recovery from a finite noisy sweep, which is quantified separately in Section 4.

Proposition 3 (Structural identifiability from the complete minimizer map)**.** *Consider a free site I (log-coordinate v, weight $\alpha > 0$) joined by a single edge (weight $\beta > 0$) to a neighbor pinned at u. For a reciprocal cost with $\Phi_0 \in C^2$, even, $\Phi_0(0) = 0$, and $\Phi_0''(t) > 0$ for all t (so $g := \Phi_0'$ is odd, strictly increasing, with $g(0) = 0$), the objective $\alpha\Phi_0(v) + \beta\Phi_0(u - v)$ is strictly convex, with unique minimizer $v^\star(u, \alpha/\beta)$ characterized by*

$$\alpha\, g(v^\star) = \beta\, g(u - v^\star). \tag{6}$$

As $u > 0$ and $\alpha/\beta > 0$ vary, the complete minimizer map $(u, \alpha/\beta) \mapsto v^\star$ determines $g(s)/g(w)$ for every $s, w > 0$. Without a scale calibration, this determines g only up to a positive multiplicative constant. Within the calibrated class $g'(0) = 1$, it determines g uniquely. Consequently no comparator distinct from J in that calibrated class reproduces the minimizer locations of J across all $(u, \alpha/\beta)$.

The proposition is stated for a whole class of reciprocal generators, not for J alone: its proof uses only the listed hypotheses on Φ_0, and J is the special case $g = \sinh$.

Proof of Proposition 3. The hypothesis $\Phi_0''(t) > 0$ for all t implies strict convexity of $v \mapsto \alpha\Phi_0(v) + \beta\Phi_0(u - v)$ and $g'(t) > 0$ everywhere, giving a unique minimizer solving the stationarity condition (6). Fix $u > 0$. Since g is odd and strictly increasing with $g(0) = 0$, $g \le 0$ on $v \le 0$ and $g(u - v) \ge g(u) > 0$ there, so the objective's derivative $\alpha g(v) - \beta g(u - v)$ is strictly negative on $v \le 0$; symmetrically it is strictly positive on $v \ge u$. The unique minimizer $v^\star$ therefore lies in the open interval $(0, u)$, where $g > 0$ throughout. Write $r = \alpha/\beta$ and $G(v, r) = r\, g(v) - g(u - v)$. Then $\partial_v G = r\, g'(v) + g'(u - v) > 0$ and $\partial_r G = g(v)$, so the implicit function theorem gives

$$\frac{dv^\star}{dr} = -\frac{g(v^\star)}{r\, g'(v^\star) + g'(u - v^\star)} < 0,$$

since $g(v^\star) > 0$ on $(0, u)$. Thus $v^\star$ decreases strictly in r. The endpoint limits are $v^\star \to u$ as $r \to 0^+$ and $v^\star \to 0$ as $r \to \infty$: if a sequence $r_n \to 0^+$ had $v^\star(r_n) \to v_\infty \le u$ with $v_\infty < u$, then (6) would force $g(u - v_\infty) = 0$, hence $v_\infty = u$, a contradiction; likewise $r_n \to \infty$ with $v^\star(r_n) \to v_\infty \ge 0$ and $v_\infty > 0$ would force $g(v_\infty) = 0$. Hence $v^\star$ sweeps $(0, u)$ bijectively. For any target $w \in (0, u)$ a unique ratio places the minimizer at $v^\star = w$, and (6) reads $r = g(u - w)/g(w)$. Given arbitrary $s, w > 0$, choose $u = s + w$: the ratio that places the minimizer at w is $g(s)/g(w)$. The minimizer-location map therefore yields every positive ratio of g. Fix any reference $w_0 > 0$ and define $h(s) := g(s)/g(w_0)$ for $s > 0$ from those ratios; then h is known, C^1, and odd after the extension $h(-s) = -h(s)$, $h(0) = 0$. Without calibration this determines g only up to the positive factor $g(w_0)$. Under $g'(0) = 1$ one has $h'(0) = 1/g(w_0)$, so

$$g(w_0) = \frac{1}{h'(0)}, \qquad g(s) = g(w_0)\, h(s) = \frac{h(s)}{h'(0)} \qquad (s > 0),$$

or equivalently $g(s) = \lim_{w \to 0^+} w \cdot g(s)/g(w)$. The odd extension then fixes g on $(-\infty, 0)$. With $\Phi_0(0) = 0$ and $\Phi_0(t) = \int_0^t g(s)\, ds$, the generator (and hence $F(x) = \Phi_0(\log x)$)

is determined. For J one has $g = \sinh$; any comparator matching J on all minimizer locations has $g = \sinh$ and equals J. □

Operationally, the controlled weight ratio acts as a query variable for g: the ratio that places the observed minimizer at w directly returns $g(u-w)/g(w)$. Therefore, within the calibrated class of Proposition 3, a complete minimizer map consistent with J reconstructs $g = \sinh$. Conversely, reconstruction of a generator inconsistent with sinh under unit curvature would rule out J within the calibrated class.

Propositions 2 and 3 are complementary: a single operating point tests local separation at large drive, while the complete continuous map identifies the generator. The finite grid of Section 4 discriminates prespecified parametric generators and estimates p; the three-site discriminator (Supplemental S10) is a single-configuration contrast, distinct from structural recovery.

3.3. Equal-weight symmetric case

On the equal-weight two-site chain with a single pinned boundary $a = e^u$ and equal site/edge weights, every strictly convex even log-coordinate cost Ψ used with the same shape in the site and edge terms yields an objective of the form $\Psi(v)+\Psi(u-v)$ up to an additive constant. Differentiating gives $\Psi'(v)-\Psi'(u-v) = 0$; since such a Ψ has Ψ' odd and strictly increasing, the unique solution is $v^\star = u/2$, i.e. $z^\star = \sqrt{a}$, independently of the particular shape of Ψ. Location alone therefore cannot identify J under equal weights and reflection symmetry: every such cost, including every J_p and the quadratic-log limit Q, shares the same geometric-mean critical point. Static magnitude comparisons between nearby reciprocal unit-curvature costs, such as J, Q, and nearby J_p, are also weak near moderate ratios because they agree through second order. Details, including the two-site closed form under unequal weights and a static score table, are collected in Supplemental Material S5. Location-based identifiability requires breaking the relevant reflection symmetry; the three-site design does so with both unequal pinned boundaries and unequal edge weights, although other asymmetric configurations are possible.

The two designs differ in how readily they could be instantiated. The primary two-site sweep needs only one controlled exchange with a maintained reservoir. The three-site geometry does not reduce to a single exchange current: the asymmetry that separates generators generically drives a through-current, so any realization would require sequential edge gating or explicit edge-current coordinates (Section 6). We therefore treat the three-site design as a mathematical discriminator rather than a candidate experiment.

4. Finite-sweep discrimination of generators

The primary calculation is a two-site boundary/weight sweep; a secondary asymmetric three-site discriminator is summarized at the end, with full detail in Supplemental S10.

4.1. Primary two-site boundary/weight sweep

The design is the minimal asymmetric configuration: one free site joined by a single edge to a neighbor pinned at boundary multiplier $a = e^u$, with no other edge incident on the free site. Under a maintained boundary the free log-coordinate obeys the undriven scalar flow $\rho\,\dot v = -C'(v)$, where $\rho > 0$ is the one-dimensional friction (the scalar case of the matrix R), with $C(v) = \alpha(\cosh v - 1) + \beta(\cosh(u - v) - 1)$. Since C is strictly convex and coercive, the flow relaxes monotonically to the unique passive equilibrium $v^\star(u, \alpha/\beta)$ of (6), at the exponential rate recorded later in Proposition 5; the comparisons below use only the stationarity condition (6) and are therefore independent of the relaxation model. Because each candidate generator predicts a full vector of equilibrium locations across a prespecified grid, the natural analyses are (i) pairwise model selection between fully specified generators by aggregate residual, and (ii) least-squares estimation of the reciprocal-family index p (with $p = 1$ giving J).

For the whole reciprocal family this equilibrium is available in closed form. Writing $r = \alpha/\beta$, the stationarity condition (6) for J_p reads $\alpha \sinh(pv) = \beta \sinh\big(p(u - v)\big)$; the substitution $X = e^{pv}$ makes it linear in X^2 and gives

$$v_p^\star(u, r) = \frac{1}{2p} \log\left(\frac{r + e^{pu}}{r + e^{-pu}}\right), \tag{7}$$

which reduces at $p = 1$ to the two-site formula of Supplemental S5. Predicted locations, their derivatives in p, and hence the Fisher information below are therefore available in closed form rather than only by root solving. Equation (7) also fixes the informative part of the grid: at $r = 1$ the ratio collapses to e^{pu}, so $v_p^\star = u/2$ for every p, and the equal-weight conditions predict a common location across the entire family.

Observation-noise-only pairwise selection Fix in advance the grid $u \in \{2, 3, 4\}$ crossed with $\alpha/\beta \in \{0.25, 0.5, 1, 2, 4\}$ (15 points), with $n = 7$ replicates per condition at $\sigma = 0.05$. Each replicate is one independent simulated measurement at its operating condition. By (7) the conditions with $\alpha/\beta = 1$ predict the same location for every member of the family, so they carry zero local Fisher information about p and contribute identical residuals to both candidates: they cancel exactly from any pairwise residual difference. They are therefore mathematically non-informative for this statistic and are retained as symmetry controls. Selecting J over a nearby prespecified rival J_p by smaller aggregate residual has the observation-noise-only simulated selection probabilities in Table 2: 0.997 at $p = 1.1$, 0.921 at $p = 1.05$, and 0.720 at $p = 1.02$. These figures assume every boundary u and weight ratio α/β is known exactly, which is what "observation-noise-only" means here and for the three-site 0.828 figure.

Representation-uncertainty screening for the primary sweep The three-site representation-uncertainty analysis does not transfer to this geometry. We therefore re-simulate the sweep under a screening errors-in-variables model with shared and condition-specific errors of common width w on the boundaries u and log weight

Table 2. Observation-noise-only simulated selection probabilities for pairwise comparison of J with a fully specified rival J_p in the 15-point single-exchange sweep ($n = 7$ independent replicates per condition, $\sigma = 0.05$, $N = 2 \times 10^4$, seed 20260713). The first row reports the simulated probability of selecting J when J generates the data; the second reports selection of J when J_p generates the data. This is simple-versus-simple model selection, not a composite test of $\{p \neq 1\}$. The last row before the separations gives the binomial Monte Carlo standard error, common to both probability rows.

p	1.02	1.05	1.10	1.15	1.20
Pr(select J | J true)	0.720	0.921	0.997	1.000	1.000
Pr(select J | J_p true)	0.282	0.078	0.003	0.000	0.000
Monte Carlo s.e.	0.003	0.002	0.0004	< 0.001	< 0.001
$\max_k \lvert v^\star_J - v^\star_{J_p} \rvert$	0.011	0.026	0.050	0.073	0.095

ratios, the same realized nuisance entering both candidate generators' data-generating maps while the analyst scores against nominal predictions (Supplemental S9.3). At the nominal width $w = 0.10$, the simulated probability of selecting J against $J_{1.1}$ falls from 0.997 to ≈ 0.88; under the same screen, the model-specific probability of selecting J when $J_{1.1}$ generates the data is ≈ 0.10. At $w = 0.15$ it is ≈ 0.79. A reduced 12-point grid that drops the $\alpha/\beta = 1$ controls gives 0.877 at $w = 0.10$; because those conditions cancel exactly from the pairwise statistic, the gap from 0.882 is Monte Carlo variation at $N = 4000$ per width, not a loss of information. Shared-only errors are milder than condition-specific ones of the same width (0.97 versus 0.77 at $w = 0.10$); adding a batch-level random effect of size 0.03 on the replicates does not materially change the $w = 0.10$ figure. Against nearer rivals the collapse is sharper: at $w = 0.10$, the screening selection probability is ≈ 0.72 at $p = 1.05$ and ≈ 0.59 at $p = 1.02$. Figure 2 summarizes the primary sweep. Headline probabilities use $N = 2 \times 10^4$ Monte Carlo datasets (binomial Monte Carlo standard error at most 0.004); width scans use $N = 4000$ per point (standard error at most 0.008).

Estimation of p, not only pairwise selection Table 2 reports simple-versus-simple selection probabilities against fixed J_p alternatives. Separately, under the exact-parameter Gaussian observation model the local Fisher information at $p = 1$ yields the Cramér–Rao lower bound $\mathrm{sd}(\hat{p}) \geq 0.017$ for the 15-point sweep (versus 0.13 at the single three-site point). Direct Monte Carlo of the least-squares estimator minimizes the sum of squared deviations between the 15 condition means and predicted locations on the prespecified grid $p \in [0.85, 1.35]$ with spacing 0.005; ties and boundary solutions are retained at the corresponding grid value. Its idealized Wald standard error is the inverse square root of the Fisher information evaluated at $p = 1$. The estimator attains that scale when parameters are known exactly (bias ≈ 0, empirical sd 0.017, Wald 95% coverage 0.95, type-I error of $H_0\colon p = 1$ near 0.05), and the simulated rejection proportion for H_0 at $p_{\mathrm{true}} = 1.1$ was 1.000 over $N = 2000$ datasets. Under the same $w = 0.10$ nuisance model the empirical sd inflates to ≈ 0.05 and naive exact-parameter Wald intervals undercover. The 0.017 figure is a lower bound under exact parameters; achievable precision falls once calibration errors are admitted, and a protocol claiming interval coverage would profile, integrate, or bootstrap those errors.

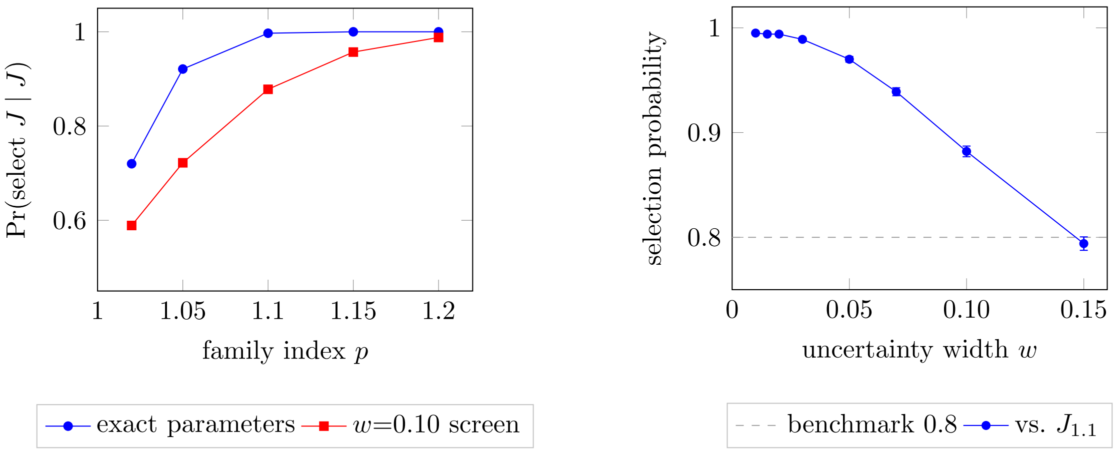

Figure 2. Primary single-exchange sweep. Left: probability of selecting J over a fixed rival J_p when J generates the data, with all boundaries and weight ratios known exactly and under the $w = 0.10$ representation-uncertainty screen. Right: the corresponding screening probability against $J_{1.1}$ versus nuisance width w (Supplemental S9.3). Error bars are ± 1 binomial Monte Carlo standard error ($N = 2 \times 10^4$ left, $N = 4000$ right) and are smaller than the markers at several points. The 0.8 line is a design benchmark for comparing conditions. The plotted quantity is the simulated probability of selecting the data-generating model, which is distinct from the rejection rate of a level-controlled test.

Feasibility and scope The required boundary ratios $a = e^u \approx 7.4$, 20, and 54.6 are analyte- and platform-dependent; whether ratios of order 10–50 are attainable under single-current closure is a question for a named physical realization (Section 6). A 12-point grid (dropping $\alpha/\beta = 1$) reduces the total number of observations from 105 to 84 without discarding any pairwise-selection information. The sweep is a within-family discriminator: it separates prespecified members of the J_p family and estimates p. Nonparametric reconstruction of g uses the complete continuous map of Proposition 3. Each two-site condition separates the non-reciprocal comparators Q, S, and D weakly; the large single-configuration separations for those comparators come from the asymmetric three-site discriminator (summary below; Supplemental S10).

4.2. Secondary discriminator (summary)

As a secondary mathematical discriminator, an asymmetric three-site system separates J from four fixed comparators calibrated to unit curvature at the reference:

$$Q(x) = \tfrac{1}{2}(\log x)^2, \qquad S(x) = \tfrac{1}{2}(x-1)^2, \qquad D(x) = x \log x - (x-1),$$

together with a linear-edge elastic reference H_{el}, whose full objective is given in Supplemental S10. At the declared operating point $(\alpha, \beta_L, \beta_R, a_L, a_R) = (1, 2, 1, 10, 0.2)$ the nearest comparator is D, with log-location separation $|v_J - v_D| = 0.056$. The observation-noise-only simulated selection probability is 0.828; under the stated representation-uncertainty screening model it falls to ≈ 0.33. Full locations, figure, decision rule, and sensitivity analysis are in Supplemental S10.

5. Gradient-flow and Langevin realization

Sections 3–4 used only the stationarity condition, so the results above hold independently of how a system reaches its constrained minimizer. We recorded the dynamics for which $\widetilde{C}$ is the potential for two reasons: the equilibrium analysis above implicitly assumes that constrained minimizers are the states a relaxing system reaches, and the Gibbs realization identifies $\widetilde{C}$ as the potential of a reversible diffusion. The results are standard overdamped gradient-flow theory specialized to $\widetilde{C}$, stated as consequences of convexity.

In the log-coordinates y of Section 2, boundary coordinates are held fixed, so the gradient below is taken over free coordinates only.

Proposition 4 (Interval dissipation identity for the driven flow)**.** *Suppose the free coordinates are absolutely continuous and obey the overdamped driven flow*

$$R\,\dot{y}(t) = -\nabla\widetilde{C}\big(y(t)\big) + \xi(t), \tag{8}$$

where R is a constant symmetric positive-definite friction (inverse-mobility) matrix, $\xi(t)$ is a forcing term, and $P(t) = \xi(t)\cdot\dot{y}(t) \in L^1$. Set $E(t) = E(t_0) + \int_{t_0}^{t} P(s)\,ds$. Then

$$\big(\widetilde{C}(y(t_2)) - \widetilde{C}(y(t_1))\big) + \int_{t_1}^{t_2} \dot{y}^\top R\,\dot{y}\,ds \;=\; E(t_2) - E(t_1), \tag{9}$$

and, discarding the nonnegative quadratic integral, $\Delta\widetilde{C} \le \int P \le \int \max\{0, P\}$. At fixed boundary this difference equals the corresponding change in $C_{\mathcal{M}}$, by the convention of Section 2.

Proof. From (8), $\nabla\widetilde{C} = \xi - R\dot{y}$, so $d\widetilde{C}/dt = P - \dot{y}^\top R\dot{y}$ along the trajectory; integrating over $[t_1, t_2]$ gives (9). Since $R \succ 0$, discarding $\dot{y}^\top R\dot{y} \ge 0$ gives the stated inequality, and $\int P \le \int \max\{0, P\}$ is immediate. □

Identity (9) is the exact interval balance of an overdamped gradient system [6, 7, 8], specialized to the reciprocal functional $\widetilde{C}$. What it would take to read ξ, $\xi\cdot\dot{y}$, and $\dot{y}^\top R\dot{y}$ as a physical force, power, and dissipation is set out in Section 6.

Units and notation. The multiplier field m, the log-coordinates y, the cost $C_{\mathcal{M}}$, and its reduction $\widetilde{C}$ are dimensionless, hence so are the gradient $\nabla\widetilde{C}$, the generalized force ξ, and the noise level ε in (10), an effective temperature entering only as $\widetilde{C}/\varepsilon$ in the invariant density. The friction matrix R carries units of time and the mobility $M := R^{-1}$ units of time^{-1}. Only the product of mobility and elapsed time is meaningful without an independent calibration; Supplemental S4 records the remaining unit conventions and the scale non-identifiabilities they leave. Calligraphic $\mathcal{M}$ always denotes the graph representation (as in $C_{\mathcal{M}}$), italic M the unrelated mobility matrix.

With the forcing switched off, the same flow relaxes to the static optimum of Proposition 1 at a rate set by the weights and the friction.

Proposition 5 (Exponential relaxation of the undriven flow)**.** *Suppose $\xi \equiv 0$ on the free coordinates, the free set is nonempty, and $R \succ 0$ is constant, so* (8) *reads $R\dot{y} = -\nabla\widetilde{C}(y)$. Let $y^\star$ be the unique constrained minimizer of Proposition 1, and set $\mu := \min_{s \text{ free}} w_s > 0$, $m_{\text{mob}} := \lambda_{\min}(R^{-1}) = 1/\lambda_{\max}(R) > 0$ (smallest eigenvalue of the mobility $M = R^{-1}$). Then for all $t \geq t_0$,*

$$\widetilde{C}\big(y(t)\big) - \widetilde{C}(y^\star) \leq e^{-2\mu m_{\text{mob}}(t-t_0)}\,\Delta_0, \qquad \|y(t) - y^\star\| \leq \sqrt{2\Delta_0/\mu}\, e^{-\mu m_{\text{mob}}(t-t_0)},$$

where $\Delta_0 := \widetilde{C}(y(t_0)) - \widetilde{C}(y^\star)$. In particular, for the multiplier field $m(t) = e^{y(t)}$ recovered from the log-coordinates, $C_{\mathcal{M}}[m(t)] \to C_{\mathcal{M}}[m^\star]$ and $m(t) \to m^\star := e^{y^\star}$ exponentially at a rate of at least μm_{mob}.

The proof is the standard strong-convexity and Polyak–Łojasiewicz argument [22, 23] and is recorded in Supplemental S2.

The flow (8) arises as the zero-noise deterministic limit of a coarse-coordinate Langevin model on the selected log-multiplier coordinates:

$$dy_t^\varepsilon = M\big(-\nabla\widetilde{C}(y_t^\varepsilon) + \xi(t)\big)\,dt + \sqrt{2\varepsilon}\,M^{1/2}\,dW_t, \qquad M := R^{-1} \succ 0, \tag{10}$$

with W a standard vector Brownian motion and $\varepsilon > 0$ a noise level.

Proposition 6 (Finite-horizon zero-noise limit and unforced reversibility)**.** *By construction $\widetilde{C} \in C^\infty$ and $M = R^{-1}$ is a constant symmetric positive-definite matrix, so the drift $b(t, y) = M(-\nabla\widetilde{C}(y) + \xi(t))$ is locally Lipschitz in y; this standing regularity, not the hypotheses below, supplies the drift condition. Assume in addition that: (L1) ξ is continuous (or piecewise C^1) on a finite interval $[t_0, T]$; (L2) the initial condition $y^\varepsilon(t_0) = y_0$ is fixed independently of ε; and (L3) every free-site weight is strictly positive. Under (L3) the* cosh *site terms make $\widetilde{C}$ strongly convex with modulus $\mu := \min_{s \text{ free}} w_s > 0$, whence the dissipativity bound $y \cdot \nabla\widetilde{C}(y) \geq \mu\|y\|^2 - \|\nabla\widetilde{C}(0)\|\,\|y\|$. Then the diffusion* (10) *is nonexploding and satisfies the finite-horizon quadratic Lyapunov bound* $\sup_{0<\varepsilon\leq\varepsilon_0} \mathbb{E}[\sup_{t\in[t_0,T]} V(y_t^\varepsilon)] < \infty$ *with $V(y) = \frac{1}{2}y^\top M^{-1}y$, together with finite-horizon compact containment; no finite moment of the exponentially growing cost $\widetilde{C}$ is claimed. Consequently:* zero-noise convergence, *as $\varepsilon \to 0$ the process y^ε converges in probability, uniformly on $[t_0, T]$, to the solution of* (8)*; and* unforced reversibility, *if $\xi \equiv 0$ the diffusion is reversible with invariant Gibbs measure $\pi_\varepsilon \propto e^{-\widetilde{C}/\varepsilon}\,dy$, so $\widetilde{C}$ is its potential. For constant ξ the drift remains a gradient of the tilted potential $\widetilde{C}_\xi(y) = \widetilde{C}(y) - \xi \cdot y$, which is coercive under (L3) with integrable Gibbs weight $\pi_{\varepsilon,\xi} \propto \exp(-\widetilde{C}_\xi/\varepsilon)$; a constant* *vector tilt is therefore conservative in the represented coordinates, producing a force-shifted equilibrium rather than a maintained nonequilibrium steady state with sustained current.*

The Lyapunov and containment bounds follow from the standing regularity of $\widetilde{C}$ and M together with (L1)–(L3); they, the reversibility argument, and the finite-horizon localization are proved in Supplemental S3 via Itô's formula and the Burkholder–Davis–Gundy and Grönwall inequalities [24, 25], with constant-tilt numerics in Supplemental S7. The matched mobility and noise covariance $2\varepsilon M$ in (10) encode the detailed-balance modeling assumption; with that choice the unforced diffusion is reversible with respect to the stated Gibbs measure. A different covariance would break the correspondence between $\widetilde{C}$ and the invariant density.

6. Limitations and physical outlook

The results above concern a convex functional on a weighted graph and the inverse problem of recovering its generator. A physical or thermodynamic interpretation of $C_{\mathcal{M}}$ calls for further ingredients, set out here together with the scope conditions of the construction.

Relation to conventional free energies For ideal-mixture concentration transport, standard free-energy expressions are relative-entropy-like rather than of the form J [4, 5]; the comparator D of Section 4 is representative of that family. Any energy-like reading of $C_{\mathcal{M}}$ would therefore be an empirical calibration rather than a derivation: one would introduce a conversion constant with units of energy per dimensionless cost unit, estimate it on a quasistatic dataset held separate from any test interval, and treat the nonlinear form of the functional, not its scale, as the testable content. As Supplemental S4 records, such a constant is identified only relative to a fixed normalization of the graph weights, so the normalization convention would have to be declared in advance and never refitted.

Constitutive requirements for a dynamical interpretation The identity of Proposition 4 becomes a physical work balance only after ξ is identified with a measurable generalized force, $\xi \cdot \dot{y}$ with physical power, and $\dot{y}^\top R \dot{y}$ with physical dissipation. Each identification is an assumption about a particular system, to be supplied by that realization. In particular the boundary log-ratio u and any external tilt ξ play distinct roles and must not be double-counted: a maintained reservoir enters through the edge term and shifts the passive equilibrium, which is a boundary condition rather than an applied force.

Open modeling questions Four modeling commitments would need independent justification in any physical instantiation. First, the pairwise homogeneous ansatz of Corollary 2 imposes one common local shape on every site and edge term; a system in which site and edge deviations are penalized differently falls outside it. Second, the free-site term $\alpha\Phi_0(v)$ requires a restoring mechanism independent of the reservoir edge, since passive exchange with a single reservoir alone would equilibrate at the boundary; folding a binding, elastic, or partitioning contribution into the one-coordinate potential is legitimate only given a measured equilibrium relation and demonstrated timescale separation. Third, the one-coordinate reduction assumes a single open exchange current. With two simultaneous edge currents the physical power cannot be reconstructed from the rate of change of a single amount and one effective affinity, and two maintained unequal reservoirs generically sustain a through-current at steady concentration; such configurations require explicit edge-current coordinates and lie outside the present model. Fourth, fixing the representation before any outcome is scored rules out one form of post hoc adjustment; the physical meaningfulness of the chosen graph, references, and measurement transforms, and their invariance under equivalent descriptions, are separate questions, so a representation-sensitivity analysis belongs in any protocol built on this model.

Requirements for a physical realization A physical program built on these results would name an analyte and platform, measure the constitutive relation for the free-site term, calibrate the conversion constant on independent quasistatic data, and only then test the work bound $\Delta\widetilde{C} \leq \int \max\{0, P\}$ of Proposition 4 as a necessary consistency screen applied after the generator and representation have been fixed. That program is separate from the identifiability question settled here and rests on measurements a physical realization would supply; the finite-grid calculations of Section 4 accordingly characterize the design rather than projected experimental performance.

7. Discussion and conclusions

Taken together, the results separate two questions that are easily conflated: selection of a scalar cost by functional assumptions, and structural identification of that cost from controlled equilibrium locations. Both are mathematical, and the second is the contribution of this paper. Within the prespecified reciprocal family, the numerical sweep provides a finite-sample parametric counterpart to the structural-identifiability result: it discriminates prespecified members and estimates p, while recovery of an unrestricted generator draws on the complete map. It also indicates that calibration and representation uncertainty, rather than observation noise, would dominate the error budget of an experimental implementation.

Two directions follow. Mathematically, the open question is how far Proposition 3 extends beyond one free coordinate: whether the complete minimizer map on a larger graph still determines the generator, and what minimal family of boundary and weight configurations suffices. Physically, the open items are those of Section 6, chief among them an independent free-site restoring mechanism with the required timescale separation and a derivation of J as a microscopic potential rather than an axiomatic choice. Until then, the finite-grid results here quantify how much a controlled sweep could distinguish within a given model.

Data and code availability

No measured datasets were created or analyzed in this study; every numerical value in Section 4 and Supplemental S9–S10 follows from the fixed equations and observation models stated there. The Python code, pinned dependencies and tested interpreter version, expected outputs, independent second implementation, and regression tests required to reproduce all synthetic results are provided as supplementary research data with this article (Monte Carlo seed `20260713`); Supplemental S11 lists the entry points. Deferred proofs and numerical details are in Supplemental S1–S11.

Author contributions

Jonathan Washburn: Conceptualization, Investigation, Methodology, Resources, Software, Supervision, Validation, Writing – original draft, Writing – review & editing. **Megan Simons**: Formal analysis, Investigation, Project administration, Validation, Visualization, Writing – review & editing.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Conflict of interest

The authors declare that they have no competing interests.

Acknowledgments

We thank colleagues at the Recognition Physics Institute (Elshad Allahyarov, Sebastian Pardo-Guerra, Anil Thapa, Dylan Funk, Philip Beltracchi, and Emma Tully) for discussion and review, in particular on the separation between the mathematical results and their possible physical readings, on the three-site design, and on the through-flux limitation.

[1] E. Schrödinger. *What is Life? The Physical Aspect of the Living Cell.* Cambridge University Press, 1944.

[2] J. L. England. Statistical physics of self-replication. *Journal of Chemical Physics* **139**, 121923 (2013). doi:10.1063/1.4818538.

[3] J. L. England. Dissipative adaptation in driven self-assembly. *Nature Nanotechnology* **10**, 919–923 (2015). doi:10.1038/nnano.2015.250.

[4] R. Rao and M. Esposito. Nonequilibrium thermodynamics of chemical reaction networks: wisdom from stochastic thermodynamics. *Physical Review X* **6**, 041064 (2016). doi:10.1103/PhysRevX.6.041064.

[5] H. Qian. Phosphorylation energy hypothesis: open chemical systems and their biological functions. *Annual Review of Physical Chemistry* **58**, 113–142 (2007). doi:10.1146/annurev.physchem.58.032806.104550.

[6] A. Mielke. A gradient structure for reaction–diffusion systems and for energy-drift-diffusion systems. *Nonlinearity* **24**, 1329–1346 (2011). doi:10.1088/0951-7715/24/4/016.

[7] S. Adams, N. Dirr, M. A. Peletier, and J. Zimmer. From a large-deviations principle to the Wasserstein gradient flow: a new micro–macro passage. *Communications in Mathematical Physics* **307**, 791–815 (2011). doi:10.1007/s00220-011-1328-4.

[8] A. Mielke, M. A. Peletier, and D. R. M. Renger. On the relation between gradient flows and the large-deviation principle, with applications to Markov chains and diffusion. *Potential Analysis* **41**, 1293–1327 (2014). doi:10.1007/s11118-014-9418-5.

[9] J. Washburn and M. Zlatanović. Uniqueness of the canonical reciprocal cost. *Mathematics* **14**(6), 935 (2026). doi:10.3390/math14060935.

[10] R. Bellman and K. J. Åström. On structural identifiability. *Mathematical Biosciences* **7**(3–4), 329–339 (1970). doi:10.1016/0025-5564(70)90132-X.

[11] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* **25**(15), 1923–1929 (2009). doi:10.1093/bioinformatics/btp358.

[12] A. F. Villaverde, A. Barreiro, and A. Papachristodoulou. Structural identifiability of dynamic systems biology models. *PLoS Computational Biology* **12**(10), e1005153 (2016). doi:10.1371/journal.pcbi.1005153.

[13] R. K. Ahuja and J. B. Orlin. Inverse optimization. *Operations Research* **49**(5), 771–783 (2001). doi:10.1287/opre.49.5.771.10607.

[14] T. C. Y. Chan, R. Mahmood, and I. Y. Zhu. Inverse optimization: theory and applications. *Operations Research* **73**(2), 1046–1074 (2025). doi:10.1287/opre.2022.0382.

[15] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences* **113**(15), 3932–3937 (2016). doi:10.1073/pnas.1517384113.

[16] J. Douglas. Solution of the inverse problem of the calculus of variations. *Transactions of the American Mathematical Society* **50**(1), 71–128 (1941). doi:10.1090/S0002-9947-1941-0004740-5.

[17] J. Aczél. *Lectures on Functional Equations and Their Applications.* Academic Press, New York, 1966.

[18] A. Cichocki and S. Amari. Families of alpha-, beta- and gamma-divergences: flexible and robust measures of similarities. *Entropy* **12**, 1532–1568 (2010). doi:10.3390/e12061532.

[19] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh. Clustering with Bregman divergences. *Journal of Machine Learning Research* **6**, 1705–1749 (2005).

[20] S. Amari. *Information Geometry and Its Applications.* Springer, Tokyo, 2016. doi:10.1007/978-4-431-55978-8.

[21] T. F. Weiss. *Cellular Biophysics, Volume 1: Transport.* MIT Press, Cambridge, MA, 1996.

[22] B. T. Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics* **3**(4), 864–878 (1963). doi:10.1016/0041-5553(63)90382-3.

[23] H. Karimi, J. Nutini, and M. Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak–Łojasiewicz condition. In P. Frasconi, N. Landwehr, G. Manco, and J. Vreeken (eds.), *Machine Learning and Knowledge Discovery in Databases*, pp. 795–811, Springer, Cham (2016). doi:10.1007/978-3-319-46128-1_50.

[24] R. Khasminskii. *Stochastic Stability of Differential Equations*. 2nd ed., Springer, Berlin, 2012. doi:10.1007/978-3-642-23280-0.

[25] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*. 3rd ed., Springer, Berlin, 2012. doi:10.1007/978-3-642-25847-3.