Liouville and no-work certification on Fisher scalar sector

Open Technical Note

J. Dunkley

19th December 2025

Contents

1 Purpose and scope

This note records a small but operational check of the reversible sector described in the later part of the "Answer" paper, in the setting where the Fisher scalar sector is coupled to an antisymmetric transport term. The aim is not to introduce new theory, but to make several structural claims concrete: we write the relevant operators explicitly, state the minimal identities we rely on, and report numerical observations showing that these identities survive discretisation in a clean, high resolution testbed.

We work on a two dimensional periodic domain, so the antisymmetric part J has one independent component and can be represented by a single amplitude field $a(x)$. This is the simplest nontrivial case in which the Liouville compatibility constraint $\nabla \cdot (\rho J) = 0$ can be enforced by a canonical projection, and in which the "no-work" statement can be checked in two distinct forms. The restriction to two dimensions is deliberate: it keeps the construction transparent and avoids introducing gauge choices for a general antisymmetric tensor. A generalisation to higher dimensions is conceptually straightforward but is not attempted here.

> *Remark.* Everything in this note should be read as an operational sanity check of ongoing research, not a formal stance. The continuum identities are known from the programme; what we add is an explicit discretisation in which adjointness, projection, and the no-work certificate can be measured directly. We make no claim that these tests, by themselves, settle questions of universality, continuum limits, or model selection in applications.

2 Setting and notation

Let $\mathbb{T}^2$ denote the flat two torus. We write $x = (x^1, x^2)$ and $\nabla = (\partial_1, \partial_2)$. The density $\rho(x)$ is assumed smooth and strictly positive, and we treat the Fisher scalar chemical potential in the canonical form

$$\mu := \log \rho. \tag{2.1}$$

The antisymmetric operator J is represented by a scalar amplitude $a(x)$ through the canonical area form ε:

$$J_{ij}(x) := a(x)\, \varepsilon_{ij}, \qquad \varepsilon_{12} = +1,\ \varepsilon_{21} = -1,\ \varepsilon_{11} = \varepsilon_{22} = 0. \tag{2.2}$$

We distinguish three inner products, chosen to match the "Answer" conventions in the scalar Fisher sector:

$$\langle f, g \rangle_{L^2} := \int_{\mathbb{T}^2} f(x)\, g(x)\, \mathrm{d}x, \tag{2.3}$$

$$\langle f, g \rangle_{A} := \int_{\mathbb{T}^2} f(x)\, g(x)\, \mathrm{d}x, \qquad \text{(amplitude space, flat)} \tag{2.4}$$

$$\langle s, v \rangle_{V} := \int_{\mathbb{T}^2} s(x) \cdot v(x)\, \mathrm{d}x, \qquad \text{(vector space, flat).} \tag{2.5}$$

For this technical note we deliberately keep these flat, since the Liouville defect operator already carries the density weight explicitly. The weighted pairing enters through the definition of the defect itself, as in Section 3.

3 Liouville defect operator and its adjoint

Define the *Liouville defect* of the amplitude a to be the divergence of the weighted antisymmetric current:

$$s(a) := \nabla\cdot(\rho J) \qquad \text{with} \qquad (\rho J)_i := \rho\, J_{ij}\, \mathbf{e}_j. \tag{3.1}$$

In two dimensions, with $J_{ij} = a\,\varepsilon_{ij}$, this can be written componentwise as

$$Da := s(a) = \begin{pmatrix} -\partial_2(\rho a) \\ \partial_1(\rho a) \end{pmatrix}. \tag{3.2}$$

We call a (or equivalently J) *Liouville-compatible* if

$$Da = 0. \tag{3.3}$$

This is precisely the local compatibility condition used in the reversible sector of the programme: it is stronger than antisymmetry alone and is the condition that makes the reversible current compatible with probability conservation in the Fisher scalar sector.

The formal adjoint D^* with respect to the flat pairings $\langle\cdot,\cdot\rangle_A$ and $\langle\cdot,\cdot\rangle_V$ is the unique operator satisfying

$$\langle Da, v\rangle_V = \langle a, D^* v\rangle_A \qquad \text{for smooth periodic fields } a \text{ and } v. \tag{3.4}$$

A short integration by parts yields

$$D^* \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \rho\big(\partial_2 v_1 - \partial_1 v_2\big). \tag{3.5}$$

3.1 Canonical repair as an orthogonal projection

Given a general amplitude a, the "Answer" construction suggests a canonical way to remove its Liouville defect by orthogonally projecting onto $\ker D$. In the present setting this is naturally expressed through the normal operator D^*D. Formally, the excess component δa is obtained by solving

$$D^*D\,\delta a = D^*D\,a, \tag{3.6}$$

and then defining the repaired amplitude

$$a_{\text{cons}} := a - \delta a. \tag{3.7}$$

If δa solves Eq. (3.6), then $a_{\text{cons}} \in \ker D$ in the sense that $Da_{\text{cons}} = 0$ in the range of the discretised operator. In finite precision numerics one expects Da_{cons} to be small at the level of solver tolerance and roundoff.

> *Remark* (A concrete kernel element). A particularly transparent Liouville-compatible amplitude is obtained by fixing $\rho J = \varepsilon$ pointwise. In two dimensions this is simply
>
> $$a_{\text{ex}}(x) = \frac{1}{\rho(x)}, \tag{3.8}$$
>
> for which $Da_{\text{ex}} = 0$ holds exactly at the continuum level. This element is useful as

a reference direction when testing repair numerically.

4 Reversible drift and the no-work certificate

The reversible contribution to the density dynamics driven by J and the Fisher scalar potential $\mu = \log \rho$ takes the local form

$$v_{\mathrm{rev}} := \nabla \cdot \big(\rho \, J \, \nabla \mu\big). \tag{4.1}$$

There are two distinct "no-work" statements one can write down.

The first is an algebraic identity that holds for any antisymmetric J:

$$P_1 := \int_{\mathbb{T}^2} \rho \, \nabla \mu \cdot (J \nabla \mu) \, \mathrm{d}x = 0, \tag{4.2}$$

since $u \cdot Ju = 0$ pointwise for antisymmetric J. This identity is true but weak: it does not constrain the divergence form $\nabla \cdot (\rho J \nabla \mu)$.

The second is the continuity-based pairing

$$P_2 := \int_{\mathbb{T}^2} \mu \, v_{\mathrm{rev}} \, \mathrm{d}x, \tag{4.3}$$

which is sensitive to Liouville compatibility. In the "Answer" framework, the reversible term is required to do no work in this stronger sense when the Liouville constraint is enforced. Operationally, the distinction between Eq. (4.2) and Eq. (4.3) is part of what makes the Liouville condition nontrivial.

5 What is tested in this note

The numerical experiments reported later are designed to check the following minimal structural points in a single consistent discretisation.

5.1 Adjointness of D and D^*

We discretise Eqs. (3.2) and (3.5) on a periodic grid and check the adjoint identity Eq. (3.4) on random fields, measuring the absolute discrepancy.

5.2 Repair as a projection onto $\ker D$

We implement the repair map $P : a \mapsto a_{\mathrm{cons}}$ defined by Eqs. (3.6) and (3.7) using a matrix free conjugate gradient solve, and measure three properties: $\|D(Pa)\|$ is driven to numerical roundoff, $P(Pa) \approx Pa$ (idempotency), and when $a = a_{\mathrm{ex}} + \delta$ with $a_{\mathrm{ex}} = 1/\rho$, the repaired field aligns strongly with a_{ex}.

5.3 No-work certificate in two forms

We compute v_{rev} from Eq. (4.1) and evaluate P_1 and P_2 from Eqs. (4.2) and (4.3) for two cases: a Liouville-compatible amplitude $a_{\text{ex}} = 1/\rho$, and a deliberately distorted amplitude $a_{\text{viol}} = a_{\text{ex}} + \delta$. The purpose is to exhibit that P_1 is blind to Liouville breaking, while P_2 and $\|v_{\text{rev}}\|$ respond at order one.

5.4 Holonomy mechanism on an analytic loop

As a separate mechanism check, we compute the winding number of a simple analytic loop $Z(\theta) = \mathrm{e}^{\mathrm{i}\theta}$ and show that a mild smoothing leaves the winding unchanged at the observed discretisation error. This is included only to calibrate the winding reader used elsewhere in the programme.

5.5 Kramers–Kronig note

We include a minimal Kramers–Kronig residual calculation, but only as a bookkeeping placeholder. As discussed later, a naive FFT-based Hilbert transform on a finite symmetric band is not yet accurate enough to serve as quantitative evidence for admissibility of inferred generators.

6 Numerical observations

This section records the numerical observations corresponding to Section 5. The emphasis is on checking that the operator identities used in the reversible sector admit a clean and stable discrete realisation, rather than on performance or on any particular application. All tests below use a periodic grid on $\mathbb{T}^2$ and centred finite differences, so that discrete integration by parts is as close as possible to the continuum calculation.

6.1 Discretisation used for the checks

Let the grid be $n_x \times n_y$ with spacings $\Delta x, \Delta y$ and periodic boundary conditions. We discretise ∂_1, ∂_2 by centred differences. The discrete inner products are taken as the obvious Riemann sums, for example

$$\langle f, g\rangle_{L^2,h} \coloneqq \sum_p f_p g_p \,\Delta x\,\Delta y, \qquad \langle s, v\rangle_{V,h} \coloneqq \sum_p s_p \cdot v_p \,\Delta x\,\Delta y. \tag{6.1}$$

The discrete versions D_h and D_h^* are obtained by applying these stencils to Eqs. (3.2) and (3.5) without additional filtering.

The density ρ is a smooth strictly positive field on the grid, constructed from a small number of random Fourier modes and then shifted and rescaled so that $\rho > 0$ pointwise and $\langle\rho\rangle = 1$. For each test we form the reference Liouville-compatible amplitude $a_{\mathrm{ex}} = 1/\rho$ from Eq. (3.8), and then distort it by an independent random perturbation δ:

$$a_{\mathrm{raw}} \coloneqq a_{\mathrm{ex}} + \delta, \qquad \delta \sim \text{mean-zero, smooth random field on the grid.} \tag{6.2}$$

The repair map P is implemented by solving the discrete normal equations Eq. (3.6) using a matrix free conjugate gradient method with a tolerance of order 10^{-10} in the Euclidean norm of the residual.

6.2 Adjointness of D_h and D_h^*

The first check is the discrete analogue of Eq. (3.4). For independent random fields a and v we measure

$$|\langle D_h a, v\rangle_{V,h} - \langle a, D_h^* v\rangle_{A,h}|. \tag{6.3}$$

At 256×256 resolution, the mean discrepancy over 44 trials is at roundoff level, indicating that the discretisation respects the intended integration by parts structure.

6.3 Repair as a projection onto $\ker D_h$

For each trial we compute $a_{\mathrm{cons}} \coloneqq P(a_{\mathrm{raw}})$, and measure the Liouville defect norms $\|D_h a_{\mathrm{raw}}\|_{V,h}$ and $\|D_h a_{\mathrm{cons}}\|_{V,h}$, together with the idempotency defect $\|P(P(a_{\mathrm{raw}})) - P(a_{\mathrm{raw}})\|_{A,h}$. We also record the alignment of the repaired amplitude with the explicit kernel element $a_{\mathrm{ex}} = 1/\rho$ using the flat amplitude pairing:

$$\langle a_{\mathrm{cons}}, a_{\mathrm{ex}}\rangle_{A,h}. \tag{6.4}$$

This quantity is not an orthogonality diagnostic; it is included here only to confirm that, when the raw field is constructed as $a_{\mathrm{ex}} + \delta$, the repair map returns a field strongly aligned with the reference kernel direction.

A representative run (the one discussed in this note) produced the following aggregate results:

Quantity (mean over 44 trials, grid 256×256)	Value
Adjoint discrepancy $\lvert\langle D_h a, v\rangle_{V,h} - \langle a, D_h^* v\rangle_{A,h}\rvert$	2.092×10^{-16}
Raw defect norm $\|D_h a_{\mathrm{raw}}\|_{V,h}$	7.676345×10^{1}
Repaired defect norm $\|D_h a_{\mathrm{cons}}\|_{V,h}$	8.999801×10^{-14}
Alignment $\lvert\langle a_{\mathrm{cons}}, a_{\mathrm{ex}}\rangle_{A,h}\rvert$	9.999583×10^{-1}
Idempotency $\|P(P(a)) - P(a)\|_{A,h}$	5.822×10^{-16}
Mean CG iterations (worst of two passes)	5.358×10^{2}
Mean final CG residual (worst of two passes)	9.830×10^{-11}

Two qualitative points are worth extracting.

First, the repaired field is numerically Liouville-compatible: the defect norm collapses from order 10^{1} to order 10^{-13} at this resolution. Second, the repair map is numerically a projector: $P(P(a))$ returns $P(a)$ to roundoff. These are the two operational properties needed for the "repair" interpretation in Section 3.1.

Remark (On the diagnostic $|\|Da_{\mathrm{raw}}\| - \|Da_{\mathrm{ex}}\||$). In the present setup $a_{\mathrm{ex}} = 1/\rho$ is chosen precisely because it lies in $\ker D$, so $Da_{\mathrm{ex}} = 0$ at the continuum level. At finite resolution the measured $\|D_h a_{\mathrm{ex}}\|_{V,h}$ is at numerical floor. Consequently $|\|D_h a_{\mathrm{raw}}\|_{V,h} - \|D_h a_{\mathrm{ex}}\|_{V,h}|$ is essentially $\|D_h a_{\mathrm{raw}}\|_{V,h}$. This is not a defect decomposition and should not be interpreted as such.

6.4 No-work certificate: algebraic versus continuity-based

We now evaluate the reversible drift v_{rev} from Eq. (4.1) and the two "power" functionals Eqs. (4.2) and (4.3). Recall that P_1 vanishes for any antisymmetric J, whereas P_2 is sensitive to whether the reversible term is compatible with the Liouville constraint.

We report the three quantities P_1, P_2, and $\|v_{\mathrm{rev}}\|_{L^2,h}$, for a Liouville-compatible amplitude $a_{\mathrm{ex}} = 1/\rho$ and for a distorted amplitude $a_{\mathrm{viol}} = a_{\mathrm{ex}} + \delta$.

Quantity	$a_{\mathrm{ex}} = 1/\rho$	$a_{\mathrm{viol}} = a_{\mathrm{ex}} + \delta$
$P_1 = \int \rho\, \nabla\mu \cdot (J\nabla\mu)\, \mathrm{d}x$	-5.434×10^{-25}	4.432×10^{-25}
$P_2 = \int \mu\, v_{\mathrm{rev}}\, \mathrm{d}x$	2.527×10^{-25}	8.272×10^{-25}
$\|v_{\mathrm{rev}}\|_{L^2,h}$	2.653×10^{-17}	9.687×10^{-2}

The behaviour matches intended separation of roles.

For $a_{\mathrm{ex}} = 1/\rho$, the reversible drift is numerically zero and both P_1 and P_2 are at floating point floor. For the distorted amplitude a_{viol}, the algebraic quantity P_1 remains at numerical floor (as it must, by antisymmetry), while the actual reversible drift becomes order 10^{-1} in L^2. This is the operational content of the "no-work certificate" in the scalar Fisher sector: antisymmetry is insufficient, and Liouville compatibility is the extra constraint that collapses the divergence form $v_{\mathrm{rev}} = \nabla \cdot (\rho J \nabla\mu)$.

Remark (What this does and does not show). These checks support a specific local statement: the Liouville constraint can be enforced by a canonical projection and, when enforced, it makes the reversible drift vanish (in this specific Fisher-scalar choice $\mu = \log\rho$). This note does not address the more delicate questions of how J should be inferred from data, how stable the projection is under coarse graining in applications, or how the reversible sector interacts with the full metriplectic structure once G is included. Those belong in separate analyses.

7 Holonomy mechanism and Kramers–Kronig caveat

This section records two auxiliary points. The first is a minimal holonomy or winding computation on an analytic loop, included only to calibrate the winding reader used elsewhere in the programme. The second is a caveat on a naive Kramers–Kronig residual computed with an FFT-based Hilbert transform on a finite band.

7.1 Winding number on a simple analytic loop

Consider the loop $Z(\theta) = e^{i\theta}$, $\theta \in [0, 2\pi)$. Its winding number about the origin is exactly 1. A standard numerical winding reader computes

$$n(Z) := \frac{1}{2\pi}\Big(\arg Z(2\pi^-) - \arg Z(0)\Big), \tag{7.1}$$

where arg is unwrapped along the sampled points. Using 4096 points and comparing a "microscopic" reader with a mildly blurred one (local three point averaging), we observed

$$n_{\text{micro}} \approx 0.999756, \qquad n_{\text{blur}} \approx 0.999756, \tag{7.2}$$

with no discrepancy at the level of the printed precision. The undercount relative to 1 is a discretisation artefact of the particular unwrapping and endpoint convention, and provides a convenient scale for what "integer to within numerical tolerance" means when the same reader is applied to more structured loops.

> *Remark* (Relation to the programme). The programme uses holonomy or winding diagnostics in settings where a complex reader $Z(\lambda)$ is built from physical data or from a control family ρ_λ. The calculation above does not address those constructions. It only shows that, for the simplest analytic loop, the winding reader is stable under mild smoothing, as a sanity check for the topological mechanism.

7.2 Kramers–Kronig residual: what was (and was not) checked

A Kramers–Kronig (KK) consistency dial is useful when a susceptibility $\chi(\omega)$ is known to be the boundary value of a function analytic in the upper half plane, with suitable decay. In that setting the real and imaginary parts are Hilbert transforms of each other, and KK provides a sharp admissibility constraint.

In this note we implemented only a minimal placeholder version of a KK residual, using an FFT-based Hilbert transform on a finite symmetric frequency band. Even for the simplest analytic one-pole susceptibility

$$\chi(\omega) = \frac{1}{\gamma - i\omega}, \tag{7.3}$$

the resulting baseline residual was large (of order unity in the chosen normalisation). This does *not* indicate a failure of analyticity of χ. It indicates that the naive finite-band periodic Hilbert transform is not a faithful numerical proxy for the principal value integral on $\mathbb{R}$. A quantitative KK dial requires a more careful discretisation (for example, a semi-infinite quadrature, explicit principal value handling, or a windowed and de-aliased transform).

> *Remark* (Why this is mentioned at all). KK is conceptually important in the unified-operator story, where one wants admissibility conditions on inferred generators. However, the version included in the present diagnostic script should be treated as

> a bookkeeping hook only. A proper KK dial is a separate numerical task.

8 Integration within the UIH programme

The “Answer” paper develops the irreversible sector as a Fisher-metric gradient flow and then sketches how reversible and irreversible parts can be combined into a unified operator $\mathcal{K} = G + J$. Within that story, the reversible component J is not an arbitrary antisymmetric perturbation: it is constrained by probability conservation in the weighted geometry, and it is precisely this constraint that permits the “no-work” interpretation.

The present technical note supports that part of the story in an explicitly operational way, in the simplest nontrivial setting.

8.1 What is now operational rather than schematic

The core operator claims checked here are:

Adjointness: the discrete operators corresponding to Eqs. (3.2) and (3.5) satisfy the adjoint identity Eq. (3.4) to roundoff on a high resolution periodic grid. This is the minimal finite-dimensional analogue of the continuum integration by parts used repeatedly in the programme.

Canonical repair: there is a constructive and numerically stable map P that takes a general antisymmetric amplitude a and returns a Liouville-compatible amplitude $a_{\text{cons}} \in \ker D$ by solving the normal equation Eq. (3.6). In our tests, the repaired defect $\|D_h a_{\text{cons}}\|$ collapses to roundoff and P is idempotent to roundoff. This is an operational form of the “repair” idea in the later part of the Answer paper.

No-work separation: antisymmetry alone enforces the algebraic cancellation $P_1 = 0$, but does not guarantee that the reversible divergence term $v_{\text{rev}} = \nabla \cdot (\rho J \nabla \mu)$ vanishes. When Liouville compatibility is enforced by P, the reversible drift collapses to numerical floor. When Liouville compatibility is violated at order one, the reversible drift is order one while P_1 remains essentially zero. This sharp separation is precisely the content of the no-work certificate in the Fisher scalar setting.

8.2 How it connects to other parts of the work

These observations fit naturally into the broader programme in three ways.

First, they provide a concrete procedure for maintaining structural constraints in numerical solvers that use $\mathcal{K} = G + J$. In particular, any inferred or proposed antisymmetric component can be passed through the projection P to obtain a Liouville-compatible reversible component. This is closely analogous in spirit to gauge fixing or Hodge-type projections, but it is tied directly to the Fisher scalar geometry used in the programme.

Second, they clarify the interpretation of the reversible sector as a form of “protection” that is operational rather than microscopic. The constraint $\nabla \cdot (\rho J) = 0$ is checkable and enforceable at the level of the effective operator, and the no-work property can be verified directly. In this sense the reversible sector can be stabilised without appeal to additional microscopic symmetries.

Third, they support the internal consistency of the unified-operator decomposition. The programme's irreversible sector is already tightly constrained by convex regularity and the Fisher functional. The present check shows that the reversible sector can be constrained comparably tightly by Liouville compatibility, and that the distinction matters operationally.

9 Limitations and remaining questions

This note is intentionally narrow. The following items remain open, and should be treated explicitly if these diagnostics are used in broader arguments.

9.1 Dimensional generality

All calculations here are in 2D, where an antisymmetric tensor has one degree of freedom and can be encoded by a scalar amplitude. In 3D, J has three independent components, and the Liouville defect operator D acts on a vector-valued amplitude. One expects an analogous projection onto $\ker D$ to exist, but the details (and any gauge-type redundancies) need to be written down and tested. This is a natural next step if one wants the same operational checks for realistic three dimensional applications.

9.2 Continuum limit and discretisation dependence

We observed roundoff-level adjointness and projection properties on a 256×256 periodic grid with centred differences. This is strong evidence that the discrete operators reflect the continuum structure in that setting, but it is not a theorem. A more complete analysis would show how the measured errors scale with grid refinement and how robust the projection is under alternative discretisations (finite volume, spectral, staggered grids).

9.3 Beyond the scalar Fisher choice $\mu = \log \rho$

The no-work check was performed for the canonical scalar Fisher choice $\mu = \log \rho$. In the full programme one often considers a Fisher term plus additional potentials, and the generator G is present as well. The interaction of the projected J with these additional structures, and with the full metriplectic evolution, is not analysed here.

9.4 Holonomy and KK as programme-level dials

The holonomy mechanism check is performed only on an analytic toy loop and does not yet validate holonomy dials built from UIH readers. The KK residual in the present diagnostic script is not yet numerically faithful. Both are conceptually important within the broader programme, but both require separate dedicated numerical work before they can be used as quantitative evidence.

10 Practical implications and cautions

Operationally, the main implication is that Liouville repair can be treated as a first-class numerical primitive: one can enforce $\nabla \cdot (\rho J) = 0$ by a canonical projection that is stable and idempotent in practice, and doing so has a visible effect on the reversible drift. This supports the view that the reversible sector is a constrained geometric object rather than an arbitrary antisymmetric add-on.

The main caution is rhetorical rather than technical: the results here are clean but local. They certify that a particular set of identities in the "Answer" framework can be realised numerically on a periodic grid at high resolution. They do not by themselves establish universality claims, do not validate KK dials, and do not replace the need for application-specific tests in gravity, spectroscopy, or other domains.

A Reproducibility notes

nomogenetics.com/python/uih_diag.py

The diagnostic script used for these checks implements D_h and D_h^* using centred differences on a periodic grid, constructs a smooth positive ρ from random Fourier modes, and computes the projection P by a matrix free conjugate gradient solve of the normal equations Eq. (3.6). The reported run used a 256×256 grid and 44 independent trials for the anomaly projection suite. The no-work check used $\mu = \log \rho$ and compared $a_{\text{ex}} = 1/\rho$ to a distorted amplitude $a_{\text{viol}} = a_{\text{ex}} + \delta$.

We emphasise again that the KK residual reported by the script is not a quantitative KK test; it is included only to keep the place of such a dial visible in the diagnostic workflow.