

Hidden Manifolds

Optimal Transport, Accessibility Landscapes, and the Reconstruction of Constraint

Flyxion

Independent Researcher

Abstract

Observable flows are rarely fundamental. What we observe are movements—of goods, people, information, or probability mass—while what we seek are the hidden constraint fields that make certain trajectories admissible and others impossible. This essay develops that thesis through a detailed examination of two papers on optimal transport (OT) and its inverse: a 2026 Nature Communications study by Gaskin, Demirel, Wolfram, and Duncan applying neural inverse OT to global agricultural trade, and a 2019 paper by Stuart and Wolfram establishing the Bayesian foundations of inverse OT with applications to migration flows. Both papers demonstrate that OT constitutes a genuine substrate theory for flow phenomena, outperforming covariate-based gravity models by treating cost as something inferred from structure rather than assumed from covariates. The gravity model tradition, we argue, exemplifies a grammar approach: it parameterizes the cost landscape using observable proxies—distance, tariffs, language, colonial ties—and declares the cost derived. OT is a substrate: it recovers cost from the flows themselves. But OT as deployed in both papers is a static substrate. It identifies the cost matrix at a moment but provides no equations governing its temporal evolution. We develop a field-theoretic completion using the Relativistic Scalar-Vector Plenum (RSVP) theory, in which the cost field $C_{ij}(t)$ is the coarse-grained scalar constraint field Φ , the transport plan T_{ij} is the discretized vector flow \mathbf{v} , and the entropy regularization $-\varepsilon H(T)$ is the accessibility field S performing genuine

thermodynamic work. The RSVP field equations provide the missing dynamics; the Sinkhorn iteration is reinterpreted as RSVP gradient flow; and the dual OT admissibility condition $f_i + g_j \leq C_{ij}$ corresponds to the RSVP entropy production inequality $\sigma_{\text{prod}} \geq 0$. The empirical findings of Gaskin et al.—the disproportionate cost burden on African wheat importers, the Brexit cost asymmetry, the Asia-Pacific accessibility dynamics—are reinterpreted as deformations of the RSVP cost field under geopolitical shocks. The framework is then extended beyond trade: migration, cognition, and physical inverse problems all belong to the same scientific activity—reconstructing hidden manifolds of constraint from visible trajectories. We close with the philosophical implications of this shift: from a science of objects and volumes to a science of accessibility, constraint, and hidden geometry.

Introduction: The Problem of Hidden Structure

Trajectories Are Not Fundamental

Across the sciences, the objects we observe most readily are not the objects that explain. We observe flows of goods between nations but not the cost landscape generating them. We observe migration between countries but not the friction field that makes certain crossings costly and others easy. We observe neural firing patterns but not the constraint geometry that selects which patterns are admissible. We observe the paths of particles through a detector but not the fields that curved those paths. In each case, the visible trajectory is a projection of an invisible structure, and the scientific work is to invert the projection—to reconstruct the hidden manifold from the motion it produces.

This is not a new observation. The history of physics is largely a history of inferring fields from trajectories: Newton inferred the inverse-square gravitational field from planetary orbits; Maxwell inferred the electromagnetic field from moving charges and currents; Einstein inferred the curvature of spacetime from the geodesics of light and matter. The same logic applies, with appropriate modifications, to economics, demography, biology, and cognition. In every case, the trajectory is data; the field is the explanation. The formal theory of inverse problems

across these domains has been developed systematically [Kaipio and Somersalo, 2006, Tarantola, 2005].

What is new, in the work examined in this essay, is the systematic mathematical machinery for performing this inversion in non-physical domains. Optimal transport theory provides exactly such machinery for flow phenomena: given observed flows between sources and destinations, it recovers the latent cost field that would make those flows optimal. The result is not a curve-fit to observed data but a reconstruction of the hidden geometry—the accessibility landscape—that the flows are navigating.

From Trade Volumes to Trade Geometry

Global trade is a paradigmatic example of the general pattern. The observable quantity is trade volume: country i exports T_{ij} units of commodity k to country j in year t . These numbers are what FAO compiles, what governments report, what economists analyze. But trade volumes are not the fundamental object. They are the visible outcome of a deeper structure: the cost landscape $C_{ij}(t)$ that determines how difficult it is to move goods from i to j at time t .

The cost landscape is invisible in the sense that it is not directly reported by any authority. It encodes the aggregate effect of tariffs, logistics costs, customs delays, political relations, institutional trust, sanctions, currency mismatches, and the thousand informal barriers and facilitators that shape whether a trade route is cheap or expensive. A tariff can be looked up in a treaty database; political hostility cannot. An institutional trust relationship built over decades of commercial exchange leaves no single document. These hidden quantities are what the cost landscape encodes, and recovering them from observed flows is the inverse problem that the papers under review address.

Two recent contributions tackle this problem from different angles. Stuart and Wolfram [2019] establish the Bayesian foundations of inverse optimal transport (IOT): given noisy observations of a transport plan T , infer the cost matrix C that generated those flows. They demonstrate the approach on European migration data and establish the mathematical conditions under which the inverse problem is well-posed. Gaskin et al. [2026] apply a neural network implementation of in-

verse OT to global agricultural trade from 2000 to 2022, demonstrating that the recovered cost field outperforms traditional gravity models by an order of magnitude in prediction accuracy and reveals patterns—in particular the disproportionate cost burden placed on African wheat importers by the Ukraine war—that trade volume data alone cannot show.

Both papers represent genuine scientific advances. But they also leave open the question that this essay addresses: why does the cost landscape have the structure it does, and what equations govern its evolution? We develop a field-theoretic answer using the Relativistic Scalar-Vector Plenum (RSVP) framework, in which the cost field is a dynamical object governed by nonlinear PDEs coupled to the flow field and an accessibility field. The OT framework recovers what the cost field is; RSVP proposes how it got there and where it is going.

Three Levels of Description

Before proceeding, it is important to make explicit a hierarchy that runs throughout this essay, because the argument operates simultaneously at three levels that must be kept distinct.

Grammar is the symbolic system used to describe observable phenomena by assigning them to pre-defined categories. The gravity model is a grammar: it parameterizes the cost landscape using measurable covariates—distance, tariffs, language, colonial ties—and fits parameters to reproduce observed trade volumes. Grammar is useful and often predictively adequate, but it does not generate the cost from first principles. It assumes the form of the cost and fits it to data.

Semantics is the operational procedure for recovering cost from observed flows without presupposing its functional form. Inverse optimal transport is semantics: given observed trade flows, it infers the cost matrix that would make those flows optimal, using the OT forward equations as structural constraints. Semantics is more powerful than grammar because it recovers rather than assumes the cost structure. But semantics is still static: it identifies the cost at a moment without explaining its temporal dynamics.

Substrate is the dynamical field theory from which cost structure emerges and through which it evolves. RSVP is proposed here as a candidate substrate: a field

triple (Φ, \mathbf{v}, S) governed by coupled nonlinear PDEs whose coarse-grained projections onto trade networks produce the cost matrices that inverse OT recovers. Substrate is the most powerful level because it explains cost structure and predicts its evolution—but it also requires the strongest assumptions and is the most speculative.

The progression

Grammar \longrightarrow Semantics \longrightarrow Substrate

is not merely a stylistic hierarchy. It corresponds to a progression in what is assumed versus what is derived. Grammar assumes the most (the functional form of cost) and derives the least (parameter values from data). Semantics assumes less (only the OT structure) and derives more (the cost matrix itself). Substrate assumes the least at the level of cost (it derives cost from field dynamics) but requires the most at the level of the field equations (the RSVP Lagrangian and admissibility constraints).

Readers should note that Sections 2–6 of this essay operate primarily at the established mathematics level, discussing OT theory, Bayesian inference, and the two papers directly. Beginning with Section 7, the essay transitions to speculative theory construction. The RSVP field equations are not presented as implications of the OT papers but as a proposed field-theoretic extension. The identifications between OT and RSVP variables are presented as ansätze and projection relationships, not equalities, except where explicitly noted otherwise.

On the Status of the Substrate Hypothesis

The progression from Grammar to Semantics is empirical and established. The progression from Semantics to Substrate is conjectural. The RSVP framework is not presented as a consequence of inverse optimal transport, nor as a uniquely determined completion of it. Rather, RSVP is proposed as a candidate dynamical completion possessing three desirable properties: temporal continuity, spatial locality, and admissibility constraints that correspond to the irreversibility principle.

The argument of this essay therefore operates in two independent layers. The first layer concerns inverse optimal transport and relies only upon established mathe-

matics. The second layer concerns the possibility that the inferred cost landscape may be generated by an underlying field theory. The first layer stands independently of the second. The success or failure of RSVP does not affect the validity of the inverse OT results discussed herein. RSVP is offered as a research program motivated by those results, not implied by them.

Hidden Manifolds as a General Scientific Pattern

The deepest claim of this essay can be stated as a single sentence, which we offer as an opening thesis rather than a conclusion:

Observable flows are projections of hidden manifolds of constraint.

This thesis unifies the examples that follow. Newton’s planetary trajectories are projections of the gravitational field. Scattering cross-sections are projections of the potential energy landscape. Migration flows are projections of the friction field that makes certain border crossings costly. Trade volumes are projections of the cost landscape. Neural outputs are projections of the constraint geometry of cognition. In each case, recovering the hidden manifold from its visible projection is the central scientific activity—and the activity that OT formalizes for flow phenomena.

The structure is always:

$$\text{Observed Flow} \xrightarrow{\text{inverse problem}} \text{Hidden Geometry}$$

with the hidden geometry being the object that explains, predicts, and generalizes. The difference between approaches lies in how much structure they impose on the hidden geometry: gravity models impose a parametric form, inverse OT imposes the OT structure, and RSVP imposes field-theoretic dynamics.

Inverse Problems as Scientific Method

The Forward Problem

Every domain discussed in this essay shares a common mathematical structure. There exists a hidden object x (a field, a potential, a cost landscape, a constraint

geometry) and an observable output y (trajectories, flows, measurements, trade volumes) related by a forward map \mathcal{F} :

$$y = \mathcal{F}(x). \quad (1)$$

Examples across domains:

Domain	Hidden structure x	Observable y
Newtonian gravity	Gravitational potential	Planetary orbits
Electromagnetism	Electromagnetic field	Charged particle paths
Medical tomography	Tissue density distribution	X-ray attenuation
Inverse scattering	Potential energy landscape	Scattering cross-sections
Migration modeling	Friction field	Migration flows
Trade economics	Cost landscape	Trade volumes

In every case, the forward map \mathcal{F} is known (or assumed). Given x , computing y is the forward problem. It is typically well-posed: y is uniquely determined by x , and small changes in x produce small changes in y .

The Inverse Problem

The inverse problem asks: given y , recover x . This is the direction of scientific discovery. It is typically *ill-posed* in the sense of Hadamard: solutions may not exist (if y was not generated by any x in the assumed class), may not be unique (multiple x values produce the same y), or may not depend continuously on y (small noise in y can produce large errors in the recovered x).

Optimal transport addresses the inverse problem for flow phenomena. Given observed trade flows $y = T_{ij}(t)$, it recovers the cost matrix $x = C_{ij}(t)$ using the OT forward equations as constraints. The recovery is possible because the OT forward map $\mathcal{F} : C \mapsto T$ has a well-characterized structure (the Sinkhorn form (8)) that constrains which cost matrices are consistent with which flow matrices.

Regularization, Ill-Posedness, and Identifiability

Because inverse problems are generically ill-posed, recovery requires regularization: additional assumptions that restrict the space of admissible solutions. The universal structure of regularized inversion is:

$$\hat{x} = \arg \min_x [\|y - \mathcal{F}(x)\|^2 + \lambda R(x)], \quad (2)$$

where $\|y - \mathcal{F}(x)\|^2$ measures the misfit between the predicted and observed output and $R(x)$ is a regularization functional that encodes prior knowledge about admissible solutions. The parameter λ controls the trade-off between data fidelity and regularity.

This structure unifies the approaches discussed in this essay. The Bayesian inverse OT of Stuart and Wolfram [2019] uses the Gaussian likelihood as the misfit term and the uniform prior over $[0, 1]^{m+n+mn}$ as the regularizer. The neural inverse OT of Gaskin et al. [2026] uses the loss (18) as the misfit term and the implicit regularization of the neural network architecture. The RSVP inverse problem uses the field equations (26)–(29) as regularity constraints on the temporal evolution of the cost field.

Weak and Strong Reconstruction

Two notions of reconstruction should be distinguished throughout this essay.

Weak reconstruction seeks any hidden structure capable of reproducing the observed flows. Multiple cost landscapes may generate transport plans that are observationally equivalent; weak reconstruction recovers one element of this equivalence class. Inverse OT generally provides weak reconstruction.

Strong reconstruction seeks the unique hidden structure responsible for the observed flows. This requires additional constraints—field equations, symmetry conditions, boundary conditions—sufficient to select a unique element from the equivalence class. The RSVP program aspires to strong reconstruction through the introduction of dynamical constraints and admissibility conditions. Whether these constraints are sufficient for uniqueness remains an open mathematical question addressed in Section 17.

The distinction matters because much of the criticism of field-theoretic approaches to economics amounts to observing that weak reconstruction is all that has been established. The response is not to claim strong reconstruction prematurely but to identify the additional constraints that would make it possible.

The Newtonian Analogy and Its Appeal

The gravity model of trade, first proposed by Tinbergen [1962] and since developed into a sophisticated econometric tradition [Anderson and van Wincoop, 2003, Galichon, 2016, Yotov, 2022], begins from an analogy with Newton’s law of universal gravitation. Just as the gravitational force between two masses is proportional to their product and inversely proportional to the square of the distance between them, bilateral trade between two countries is assumed to be proportional to their economic masses—proxied by GDP or total output and expenditure—and inversely proportional to some measure of the distance or cost between them:

$$T_{ij}(t) \approx \frac{O_i(t) E_j(t)}{C_{ij}(t)}. \quad (3)$$

The analogy has obvious attractions. It is mathematically simple, has clear economic motivation (large economies trade more; distant economies trade less), and has been empirically successful across a wide range of commodity types, time periods, and country pairs. The gravity model has become the workhorse of international trade econometrics.

The Covariate Representation of Cost

The crucial operational step in any gravity model is the specification of C_{ij} . Since the cost is not directly observable, it must be represented through observable proxies. The standard specification is

$$\log C_{ij}(t) = \sum_k \alpha_k \pi_{i,k}(t) + \sum_l \beta_l \chi_{j,l}(t) + \sum_m \gamma_m \rho_{ij,m}(t), \quad (4)$$

where $\pi_{i,k}$ are exporter-side regressors (GDP, institutional quality, export infrastructure), $\chi_{j,l}$ are importer-side regressors (GDP, import tariffs, regulatory envi-

ronment), and $\rho_{ij,m}$ are bilateral covariates: geographic distance between capitals, a dummy for shared land border, a dummy for shared language, a dummy for colonial relationship, a dummy for membership in a common trade agreement, and measures of applied tariff rates.

The parameters α , β , γ are estimated by regression—typically Poisson Pseudo Maximum Likelihood [PPML; Silva and Tenreyro, 2006] to handle the zero-flow problem and the multiplicative structure of trade data. The result is an estimated cost function that can be used to predict bilateral trade flows and to conduct counterfactual analysis: what would trade have looked like if this tariff had not existed, or if this trade agreement had been in force?

The Problem of Unobservables

The gravity model tradition has identified, with considerable precision, the limitations of the covariate approach. Multilateral resistance terms—the costs of each country’s trade with all other countries, which affect its bilateral trade with any given partner—are theoretically required by the structural gravity model but are not directly observable and must be estimated, typically through fixed effects. This introduces identification challenges when fixed effects are correlated with the regressors of interest.

More fundamentally, many of the most important determinants of trade costs are not quantifiable at all. Political tension between governments leaves traces in diplomatic cables and news reporting but not in any standard dataset. Consumer preferences against products from specific countries—the informal boycotts and substitution effects that can reshape trade patterns without any formal policy change—enter only as unexplained residuals. Institutional trust between trading communities, built through decades of successful exchange and mutual investment, is entirely absent from any covariate. Strategic realignment—China’s decision in 2020 to impose anti-dumping tariffs on Australian barley as a political signal—involves motivations that no economic covariate captures.

These unobservables are absorbed into the error term, which is exactly where the most interesting trade dynamics live. A gravity model fitted on pre-2022 data will assign the Ukraine wheat shock to its error term, not to its cost specification. The

model sees the volume change but not the accessibility change.

Gravity as Grammar

The gravity model is a grammar in the precise sense developed elsewhere in this research program. A grammar is a symbolic system that covers the observable phenomena by assigning them to pre-defined categories—here, the covariate categories of distance, tariff, language, and so on. The grammar is useful: it organizes the data, enables comparison across contexts, and provides a vocabulary for policy analysis. But it is not a substrate. It does not generate the cost from first principles; it describes the cost using symbols whose meaning is assumed rather than derived.

The map/territory distinction is operative here. The gravity model's covariate set is a map of trade costs: a compressed symbolic representation that captures some of the territory's structure. But the territory—the actual cost landscape, with all its unobservable determinants—is always richer than any map of it. When the territory changes in ways the map's categories do not cover (a pandemic, a war, a political rupture, an informal norm shift), the map fails. The model's residuals grow; its predictions degrade; its fixed effects absorb the change without explaining it.

Optimal transport does not solve the map/territory problem by itself. But it makes a different move: instead of mapping the cost through covariates, it reconstructs the cost from the flows. It lets the territory speak for itself.

Why Reconstruction Matters: Prediction versus Explanation

The Distinction

A model may predict well without identifying underlying structure. A regression fitted to historical data can achieve high R^2 while remaining agnostic about mechanisms. The gravity model, in its most successful specifications, predicts bilateral trade volumes with surprising accuracy—but its predictions are hostage to the stability of its covariate structure. When that structure changes (a war, a sanctions regime, a pandemic), the model's predictions degrade precisely because the model describes correlates of cost rather than cost itself.

Reconstruction attempts something different. It does not fit a curve through data; it recovers a latent structure that is hypothesized to causally generate the data. The recovered structure can then be interrogated: where is cost high? Why? What mechanism produced this deformation of the cost landscape? These questions are not answerable from a gravity model regression, but they are answerable from the cost matrix recovered by inverse OT.

This distinction matters philosophically. In the gravity model tradition, the cost C_{ij} is a dependent variable: it is explained by the covariates. In the inverse OT tradition, the cost C_{ij} is the object of inference: it is recovered from the flows. The former assumes what it wants to explain; the latter recovers what it wants to understand.

The Limits of Covariate Thinking

The fragility of covariate-based models in complex adaptive systems is well documented across many fields. In economics, Goodhart's Law states that once a measure becomes a target, it ceases to be a good measure—because agents respond to the measure itself, not to the underlying structure it proxies. In statistics, omitted variable bias corrupts coefficient estimates whenever the included covariates are correlated with excluded ones. In philosophy of science, the underdetermination of theory by data means that many different covariate specifications can fit the same data equally well, making it impossible to choose between them on empirical grounds alone.

These problems are acute in trade modeling. The set of possible covariates is large and only partially observed. The relationships between covariates change over time as geopolitical alignments shift. The error term absorbs not just measurement noise but systematically excluded variables—precisely the variables that are hardest to quantify but most important for understanding what actually drives trade. Inverse OT replaces the covariate approach with a structural assumption (that trade flows minimize cost subject to supply-demand constraints) and lets the data identify the cost directly.

Geometry Emerges from Flow

The Inversion

The conceptual move at the heart of the OT approach to trade is a reversal of the usual direction of inference. Traditional models assume geometry (the cost structure) and predict flow (trade volumes). Inverse OT infers geometry from flow. This is the same inversion that characterizes field discovery in physics: we do not observe the gravitational field directly, we infer it from the trajectories of masses that move under its influence.

This inversion has a consequence that is easy to overlook: the geometry that emerges from flow is not the geometry that was assumed in advance. It is not constrained to take the form of a linear combination of covariates, or a symmetric matrix, or a function of geographic distance. It can be asymmetric, non-smooth, time-varying, and correlated with political variables that no dataset directly measures. The cost matrix recovered by inverse OT is, in this sense, a more honest representation of the actual trade landscape than any gravity model can provide.

Metric Geometry versus Accessibility Geometry

The cost matrix $C_{ij}(t)$ recovered by inverse OT is not a metric in the mathematical sense. A metric satisfies three conditions: non-negativity, symmetry ($d(i, j) = d(j, i)$), and the triangle inequality ($d(i, k) \leq d(i, j) + d(j, k)$). The OT cost matrix satisfies none of these in general: it can be asymmetric ($C_{ij} \neq C_{ji}$), the triangle inequality can fail (it may be cheaper to route through a third country than to trade directly), and its values depend on political conditions that have no geometric interpretation.

This motivates the introduction of *accessibility geometry* as an alternative to metric geometry. Define the accessibility from i to j at time t as the negative of the cost: $A_{ij}(t) = -C_{ij}(t)$ (so that high accessibility means low cost). The accessibility landscape has the following properties that distinguish it from a metric:

Asymmetry: $A_{ij} \neq A_{ji}$ in general. Exporting from Australia to China and importing from China to Australia face different regulatory environments, infrastructure, and institutional contexts. The Brexit cost increase was experienced asymmetri-

cally by UK importers of European goods versus European importers of UK goods.

Time dependence: $A_{ij}(t)$ evolves continuously in response to geopolitical events, infrastructure investment, institutional development, and shocks. It is not a static property of the pair (i, j) but a state of the system at time t .

Non-Euclidean structure: The accessibility landscape is not embeddable in Euclidean space. There is no assignment of coordinates to countries such that Euclidean distances equal trade costs. The Ukraine wheat shock created a discontinuous deformation of the accessibility landscape that is incompatible with any smooth Euclidean embedding.

Trade therefore lives on an *accessibility manifold*—a directed, time-varying, non-Euclidean geometric object—rather than on geographic space. The gravity model mistake is to confuse accessibility geometry with metric geometry: to assume that geographic distance is a good proxy for trade cost. It sometimes is, but the cases where it fails—wars, sanctions, trust relationships, institutional differences—are precisely the cases of greatest policy interest.

Cost as Emergent Geometry

The cost matrix $C_{ij}(t)$ recovered by inverse OT can be interpreted geometrically without yet invoking RSVP. Consider the set of countries as nodes in a weighted directed graph, where the weight on the edge from i to j at time t is $C_{ij}(t)$. This weighted graph defines a geometry on the set of countries: a notion of “distance” (in the sense of cost) that is not necessarily symmetric, not necessarily triangle-inequality-satisfying, and not necessarily correlated with physical geography.

In this geometric interpretation, the Ukraine war is a deformation of the graph’s edge weights: certain edges (Ukraine-to-Africa, Russia-to-Africa) become heavier while others (alternative supplier-to-Africa) may or may not compensate. Brexit is a persistent asymmetric weight change on edges between the UK and EU countries, different from the corresponding changes for Ireland. A free-trade agreement is a systematic reduction in edge weights for a set of country pairs.

This geometric language is precise enough to be useful without requiring field theory. The cost matrix is a finite-dimensional geometric object—a weighted directed graph—and its evolution over time is a sequence of graph deformations.

The question that RSVP then addresses is: what continuous field over the underlying geographic manifold would produce this sequence of graph deformations when projected to the node level? That is the reconstruction problem, and it is harder than the OT problem.

The Optimal Transport Framework as Substrate

The Forward Problem

Let $\mu \in \mathbb{R}^m$ be a supply vector (total exports of each country) and $\nu \in \mathbb{R}^n$ a demand vector (total imports of each country). A transport plan $T \in \mathbb{R}_+^{m \times n}$ is a non-negative matrix satisfying the marginal constraints

$$\sum_i T_{ij} = \nu_j \quad \forall j, \quad \sum_j T_{ij} = \mu_i \quad \forall i. \quad (5)$$

The optimal transport problem, originating with Monge [1781] and formalized in the modern relaxed form by Kantorovich [1942], finds the plan that minimizes total cost:

$$T^* = \arg \min_{T: \text{satisfies (5)}} \sum_{i,j} T_{ij} C_{ij}. \quad (6)$$

In practice the entropy-regularized formulation is used, adding the negative entropy $H(T) = -\sum_{ij} T_{ij} (\log T_{ij} - 1)$ as a regularization term:

$$T_\varepsilon^* = \arg \min_{T: \text{satisfies (5)}} \left[\sum_{i,j} T_{ij} C_{ij} - \varepsilon H(T) \right]. \quad (7)$$

The unique solution to (7) has the Sinkhorn form

$$T_\varepsilon^* = \Pi e^{-C/\varepsilon} \Omega, \quad (8)$$

where $\Pi = \text{diag}(e^{\lambda_1/\varepsilon}, \dots, e^{\lambda_m/\varepsilon})$ and $\Omega = \text{diag}(e^{\eta_1/\varepsilon}, \dots, e^{\eta_n/\varepsilon})$ are diagonal scaling matrices of Lagrange multipliers enforcing the marginal constraints. The scaling

matrices are found iteratively via the Sinkhorn-Knopp algorithm [Sinkhorn, 1964, Cuturi, 2013]:

$$\Omega^{(l+1)} = \frac{\nu}{M^\top \Pi^{(l)}}, \quad \Pi^{(l+1)} = \frac{\mu}{M \Omega^{(l+1)}}, \quad (9)$$

where $M = e^{-C/\varepsilon}$ and division is element-wise.

The Dual Problem and Admissibility

The dual of the regularized OT problem (7) is [see Villani, 2003, Galichon, 2016, for the full duality theory]

$$\max_{f \in \mathbb{R}^m, g \in \mathbb{R}^n} \langle f, \mu \rangle + \langle g, \nu \rangle - \varepsilon \sum_{i,j} e^{(f_i + g_j - C_{ij})/\varepsilon}. \quad (10)$$

In the limit $\varepsilon \rightarrow 0$, the last term enforces the *dual admissibility condition*:

$$f_i + g_j \leq C_{ij} \quad \forall i, j. \quad (11)$$

This condition has a natural economic interpretation: f_i is the cost of collecting a unit of good at source i , and g_j is the minimum profit achievable by delivering it at destination j . The admissibility condition states that the combined cost of collection and delivery must not exceed the transport cost. If $f_i + g_j > C_{ij}$, the route is dominated and should not be used. The OT solution selects exactly those routes that satisfy the admissibility condition with equality.

The Gravity Model as a Restricted Grammar

The gravity model imposes a specific functional form on C :

$$\log C_{ij}(t) = \sum_k \alpha_k \pi_{i,k}(t) + \sum_l \beta_l \chi_{j,l}(t) + \sum_m \gamma_m \rho_{ij,m}(t), \quad (12)$$

where $\pi_{i,k}$ are exporter-side regressors, $\chi_{j,l}$ importer-side regressors, and $\rho_{ij,m}$ bilateral covariates. The parameters α, β, γ are estimated by regression. This approach

has several structural limitations identified by Gaskin et al.:

Trade costs are generally asymmetric ($C_{ij} \neq C_{ji}$), but most standard covariates (geographic distance, shared language) are symmetric. The error term absorbs unobservable influences—changing political relations, consumer preferences, informal barriers—that are precisely what one most wants to understand. The functional form is assumed rather than derived, so misspecification directly biases inference.

The OT approach replaces (12) with no functional form assumption. The cost $C_{ij}(t)$ is inferred from the observed flows $T_{ij}(t)$ by solving the inverse problem. This is a genuine improvement: the substrate (cost minimization) is specified, and the grammar (functional form of C) is not imposed.

The Inverse Problem: Learning Cost from Flow

Stuart-Wolfram: Bayesian Inverse OT

Stuart and Wolfram [2019] formulate the inverse problem in a Bayesian framework [building on the general theory of statistical inverse problems in Kaipio and Somersalo, 2006]. Introduce latent variables $u \in \mathbb{R}_+^m$, $v \in \mathbb{R}_+^n$, $W \in \mathbb{R}_+^{m \times n}$ mapping to normalized supply $p = \mathcal{M}_n(u)$, demand $q = \mathcal{M}_n(v)$, and cost $C = \mathcal{M}_{n \times n}(W)$ via

$$\mathcal{M}_n(u)_j = \frac{u_j}{\sum_{\ell} u_{\ell}}, \quad \mathcal{M}_{n \times n}(W)_{ij} = \frac{W_{ij}}{\sum_{k,\ell} W_{k\ell}}. \quad (13)$$

The forward map is $T^* = \mathcal{G}(u, v, W) = \mathcal{F}(\mathcal{M}_n(u), \mathcal{M}_n(v), \mathcal{M}_{n \times n}(W))$. Assuming observed flows are corrupted by Gaussian noise $\eta \sim \mathcal{N}(0, \sigma^2 I)$:

$$T_{\text{obs}} = \mathcal{G}(u, v, W) + \eta, \quad (14)$$

the posterior distribution is

$$P(u, v, W \mid T_{\text{obs}}) \propto \exp\left(-\frac{1}{2\sigma^2} \|T_{\text{obs}} - \mathcal{G}(u, v, W)\|^2\right) \cdot \mathbf{1}_{\mathcal{U}}(u, v, W), \quad (15)$$

where $\mathcal{U} = [0, 1]^{m+n+mn}$ is the prior support. Sampling from (15) via the Random

walk Metropolis within Gibbs algorithm provides both the MAP estimate of C and full uncertainty quantification.

For graph-based costs—where C_{ij} equals the shortest path distance between nodes i and j in a graph with edge weights f —the inverse problem reduces to learning the edge weights from observed flows. Stuart and Wolfram demonstrate this on European migration data: the graph has countries as nodes and edges between bordering countries, and the algorithm recovers the cost of crossing each border from observed migration flows.

Identifiability

A critical technical observation concerns identifiability. The cost matrix C is determined only up to additive row and column shifts: for any vectors $\alpha \in \mathbb{R}^m, \beta \in \mathbb{R}^n$,

$$C_{ij} \mapsto C_{ij} + \alpha_i + \beta_j \quad (16)$$

leaves the transport plan T invariant, since the shift is absorbed by the scaling matrices Π and Ω . This is a gauge freedom in the inverse problem. Gaskin et al. resolve it by imposing $u(0) = C_{\max}$ —zero-flow edges have maximal cost—which forces $\alpha_i = \beta_j = 0$ under the maximum constraint. Stuart and Wolfram address it through the prior structure and the graph-based cost parameterization, which has fewer degrees of freedom than a general cost matrix.

This gauge freedom is not a mere mathematical nuisance. It signals that the cost matrix C is not directly observable: only its effect on the transport plan is. The gauge-invariant content of C is precisely what drives flows, which is the RSVP scalar field's role: Φ is a potential, defined up to additive constants, and only its gradient (or more precisely, its variation across the network) has physical content.

Gauge Freedom and Physical Observables

The cost matrix C_{ij} is determined only up to additive row and column shifts: for any vectors $\alpha \in \mathbb{R}^m, \beta \in \mathbb{R}^n$,

$$C_{ij} \mapsto C_{ij} + \alpha_i + \beta_j \quad (17)$$

leaves the transport plan T invariant, since the shift is absorbed by the scaling matrices Π and Ω . This is a gauge freedom in the inverse problem.

The physical content of the gauge freedom is this: the *absolute level* of costs is not observable through flow patterns. Only the *differences and variations* in cost across the network drive observable flows. This is analogous to electrostatics, where the electric potential ϕ is defined only up to an additive constant, while the electric field $\mathbf{E} = -\nabla\phi$ is the gauge-invariant quantity with direct physical significance.

For trade networks, the gauge-invariant content of C_{ij} is the set of cost *differentials*: how much more expensive is route (i, j) compared to route (i, k) ? These differentials determine which routes are used and which are abandoned. The absolute level of costs—whether trade is in general cheap or expensive—is not recoverable from flow data alone.

In the RSVP framework, the same gauge freedom appears as an additive shift in the scalar field: $\Phi \mapsto \Phi + c$ for a constant c . Since Φ enters the cost via the line integral (22), a constant shift in Φ produces a constant shift in all path integrals, which corresponds exactly to the OT gauge transformation $C_{ij} \mapsto C_{ij} + \text{const}$. The physical quantity in RSVP is therefore $\nabla\Phi$ (the cost gradient), not Φ itself. This is the field-theoretic content of the OT gauge freedom, and it means that the RSVP scalar field is itself defined only up to an additive constant—a property it shares with all gauge potentials in physics.

Recovery Does Not Imply Interpretation

Inverse optimal transport recovers latent costs but does not automatically explain their origin. A recovered increase in transport cost between two countries may arise from sanctions, infrastructure failure, political hostility, elevated insurance premiums, institutional distrust, informational barriers, or some combination thereof. The inverse problem identifies the existence and magnitude of hidden friction; it does not decompose that friction into causal components.

This limitation is not a weakness of the OT framework but a feature of any reconstruction approach. Medical tomography recovers the density distribution inside a body but does not identify which specific pathology causes an anomaly; additional clinical knowledge is required for diagnosis. Similarly, inverse OT recovers the cost

landscape but does not identify which geopolitical factor caused a particular cost deformation; additional domain knowledge is required for causal attribution.

Accordingly, inverse OT should be viewed as a reconstruction tool that provides the material for causal analysis, not as a complete explanatory theory in itself. The RSVP framework is one attempt to provide the governing equations that would explain the dynamics of the recovered cost field; additional empirical work is required to test whether those equations are correct.

Gaskin et al. [2026] train a neural network $u : T \mapsto C$ to solve the inverse problem by minimizing the loss

$$J = \left\| \hat{T}(\hat{C}) - T \right\|_2^2 + \sum_{(i,j) \in \mathcal{S}} (C_{ij} - C_{\max})^2, \quad (18)$$

where $\hat{T}(\hat{C})$ is the transport plan produced by Sinkhorn’s algorithm applied to the estimated cost $\hat{C} = u(T)$, and $\mathcal{S} = \{(i, j) : T_{ij} = 0\}$ is the set of zero-flow edges. The loss is differentiable with respect to the network parameters because the Sinkhorn algorithm is differentiable with respect to C .

Uncertainty quantification is achieved by training an ensemble of 10 networks and sampling from the FAO data uncertainty (each bilateral flow has two reports, from exporter and importer, which frequently disagree). The pushforward measure $\rho_C = u_{\#} \rho_T$ gives a distribution over cost matrices whose spread reflects both data uncertainty and the degree to which different cost matrices are consistent with the observed flows.

The key empirical findings deserve careful attention because they reveal what the cost matrix carries that trade volumes do not. For the Ukraine wheat shock: African importers saw large volume drops and large cost increases simultaneously, while European importers saw volume drops but smaller cost increases or even cost decreases. Canada saw a 91% volume drop with nearly zero cost change; Zambia saw a 97% volume drop with a 0.48 unit cost increase. This decoupling of volume and cost is precisely what the gravity model cannot capture—gravity sees a uniform “distance” between countries that does not change when geopolitics shifts—and what the OT cost matrix reveals directly.

The Physical Interpretation of Entropy Regularization

Regularization as Computation versus Regularization as Physics

The entropy term $-\varepsilon H(T)$ in the regularized OT problem (7) is typically presented as a computational convenience: it makes the problem strictly convex, guarantees a unique solution, and enables the Sinkhorn iteration which is dramatically more efficient than general linear programming. But entropy regularization also has a physical interpretation that is worth taking seriously.

The regularized objective

$$\min_T \left[\sum_{ij} T_{ij} C_{ij} - \varepsilon H(T) \right] = \min_T \left[\sum_{ij} T_{ij} C_{ij} + \varepsilon \sum_{ij} T_{ij} (\log T_{ij} - 1) \right] \quad (19)$$

is the free energy minimization of a system with energy $\sum_{ij} T_{ij} C_{ij}$ and entropy $H(T)$, at temperature ε [cf. Jaynes, 2003, Shannon, 1948, for the information-theoretic perspective on maximum entropy inference]. The Sinkhorn solution (8) then has the form of a Gibbs distribution:

$$T_{ij} \propto e^{-C_{ij}/\varepsilon}, \quad (20)$$

modulated by the scaling matrices that enforce the marginal constraints. This is the standard statistical mechanical result: at temperature ε , the probability of using route (i, j) is exponentially suppressed by the cost of that route relative to the temperature. At $\varepsilon \rightarrow 0$, only the minimum-cost routes are used; at large ε , all routes are equally accessible.

The parameter ε therefore plays the role of a thermodynamic temperature for the trade network: it measures the degree to which traders are willing to use suboptimal routes. Low ε corresponds to a highly optimizing, cost-sensitive trade system; high ε corresponds to a diffuse, poorly optimized system in which geographic and political noise dominates. The empirical choice of ε (which affects the balance between fidelity to data and computational stability) is thus not merely technical but physically meaningful—it should be estimated from the data, not tuned by hand.

Entropy Regularization and the Accessibility Field

In the RSVP framework, the entropy-accessibility field $S : M \rightarrow \mathbb{R}_{>0}$ encodes the local thermodynamic accessibility of configurations. High S means many configurations are locally accessible; low S means the system is in a highly constrained, low-entropy region.

The correspondence with OT entropy regularization is direct. Define $S_{ij} = \varepsilon$ at each network edge (i, j) . Then the regularized OT objective becomes

$$\min_T \sum_{ij} [T_{ij} C_{ij} + S_{ij} T_{ij} (\log T_{ij} - 1)], \quad (21)$$

which is the free energy minimization on a network where each edge has its own temperature S_{ij} . In the spatially uniform case $S_{ij} = \varepsilon$ for all (i, j) , this reduces to the standard entropy-regularized OT. In the spatially varying case, high- S edges (accessible, low-constraint regions) attract more flow than their costs alone would suggest; low- S edges (inaccessible, high-constraint regions) repel flow even if their costs are moderate.

This reinterpretation is not merely terminological. It suggests that the OT entropy regularization parameter ε should vary across network edges and over time, and that its variation encodes physical information about accessibility constraints—sanctions, infrastructure quality, political relations, logistical capacity—that the cost matrix C alone cannot capture. The two papers under review treat ε as a global constant; the RSVP generalization would allow $\varepsilon_{ij}(t)$ to be a field quantity with its own dynamics.

From Cost Matrices to Fields: The Reconstruction Gap

Why a Cost Matrix Is Not Yet a Field

The cost matrix $C_{ij}(t)$ is a finite-dimensional discrete object: an $m \times n$ array of non-negative numbers indexed by country pairs, defined at a sequence of time points. It has no notion of spatial continuity (the cost between Paris and Berlin is not constrained to be consistent with any notion of smooth variation across France and Germany), no intrinsic dynamics (the cost in year $t + 1$ is inferred independently of

the cost in year t), and no locality (a shock in one edge weight propagates to other edges only through the flow constraints, not through any field equation governing C itself).

A field $\Phi(x, t)$ on a manifold M is fundamentally different. It is defined at every point of M , not just at a finite set of nodes. It satisfies differential equations that constrain its behavior at nearby points. It has a dynamics that connects its values at different times. And it has a locality property: the field value at (x, t) is determined by local conditions—the field and its derivatives at nearby points and recent times.

The transition from $C_{ij}(t)$ to $\Phi(x, t)$ therefore requires:

1. **Continuity:** an embedding of the discrete node set into a continuous manifold M .
2. **Locality:** an assumption that $C_{ij}(t)$ is determined by the values of Φ along the route from i to j , not by global field properties.
3. **Dynamics:** a PDE governing the evolution of Φ that reduces to the observed $C_{ij}(t)$ sequences when projected to node pairs.
4. **Admissibility:** a constraint on Φ -trajectories that excludes physically inadmissible configurations.

None of these are provided by inverse OT. They are additional assumptions that constitute the speculative field-theoretic extension developed in the following sections.

The Reconstruction Gap and Its Central Question

The most important theoretical question in this essay is not the RSVP field equations. It is:

Under what conditions does a cost matrix $C_{ij}(t)$ uniquely determine an RSVP field configuration (Φ, \mathbf{v}, S) ?

This is a reconstruction theorem that does not yet exist. Its absence is the central gap between the OT and RSVP levels of description.

To see why the question is non-trivial, note that many different scalar fields Φ on M can produce the same line integrals $C_{ij} = \int_{x_i}^{x_j} \Phi d\ell$ if the paths of integration

are not fixed and the field has sufficient freedom. Even fixing the paths, the map from Φ to C_{ij} is a projection from an infinite-dimensional function space to a finite-dimensional matrix, and such projections are generically non-injective. Additional constraints—the RSVP field equations, the admissibility condition, boundary conditions, or symmetry requirements—are needed to make the inverse map well-defined.

The analogous theorem in physics is the reconstruction of a spacetime metric from observed geodesics. In general relativity, the metric $g_{\mu\nu}$ is not uniquely determined by a finite set of geodesic trajectories; additional assumptions (symmetry, asymptotic flatness, equations of state) are required. The reconstruction problem for RSVP fields from cost matrices is structurally similar, and similarly hard.

We flag this gap explicitly rather than concealing it, because it is precisely where further mathematical work is required. The RSVP framework is a candidate answer to the reconstruction problem, not a solution to it.

Effective Theories and Coarse-Graining

One way to approach the reconstruction gap is to treat the relationship between OT and RSVP as one of effective theories at different scales of description, rather than as a direct identification.

Consider the following hierarchy:

Scale	Object	Analog
Microscopic	Individual decisions, contracts, logistics	Particles
Meso	Trade flows $T_{ij}(t)$	Kinetic theory
Coarse	Cost matrix $C_{ij}(t)$	Thermodynamics
Field	Scalar field $\Phi(x, t)$	Hydrodynamics

In statistical mechanics, the transition from particles to thermodynamics does not require that each thermodynamic quantity correspond to a specific particle; it requires that thermodynamic quantities be recovered as statistical averages over particle configurations. Similarly, the transition from cost matrices to RSVP fields does

not require that C_{ij} be literally a line integral of Φ ; it requires that cost matrices be recoverable as coarse-grained projections of field configurations.

Under this reading, the identifications $C_{ij} = \mathcal{P}[\Phi]$ and $T_{ij} = \mathcal{D}[\mathbf{v}]$ (where \mathcal{P} and \mathcal{D} are projection and discretization functionals respectively) are effective-theory relationships, valid at the cost-matrix scale but not asserting microscopic identity. This framing is easier to defend because it does not require RSVP to be the unique completion of OT—only a consistent one.

RSVP as Field-Theoretic Completion of OT

Identifying the Field Variables

The remainder of this essay develops a speculative field-theoretic extension of the OT framework. The identifications introduced below are ansätze and coarse-grained projection relationships, not mathematical equalities. They constitute research hypotheses, not established results. The reader should distinguish the following claim levels throughout:

Level	Examples in this essay
Established mathematics	OT duality, Sinkhorn convergence, Bayesian inverse OT
Interpretive analogy	Cost landscape as accessibility field, entropy regularization as temperature
Research hypothesis	RSVP field equations governing $C(t)$, $\varepsilon \leftrightarrow S$ identification, admissibility correspondence

We now introduce the RSVP field variables as a candidate completion of the OT framework, with the explicit understanding that the projection relationships stated below require further mathematical development to be placed on a rigorous footing.

The network setting has nodes (countries) at positions x_i and edges (i, j) connecting them. The continuous limit of the network is a manifold M on which the RSVP fields (Φ, \mathbf{v}, S) are defined.

Definition 9.1 (Cost Field Projection — Research Hypothesis). *The RSVP scalar field*

$\Phi : M \rightarrow \mathbb{R}$ is hypothesized to project onto the cost matrix via

$$C_{ij}(t) = \mathcal{P}_{ij}[\Phi(\cdot, t)] \approx \int_{x_i}^{x_j} \Phi(x, t) d\ell, \quad (22)$$

where the integral is along the dominant trade route connecting i to j and \mathcal{P}_{ij} denotes this projection functional. The cost is therefore not a function of the nodes alone but of the field Φ along the path between them. This identification requires a choice of route and is non-unique without additional constraints on Φ .

This definition reproduces the graph-based cost structure of Stuart and Wolfram: the cost of crossing an edge is determined by the field value along that edge, and the cost of a multi-hop path is the sum of edge costs (in the linear approximation). It also immediately explains the gauge freedom (17): adding a constant to Φ at all points along all paths connecting country i shifts C_{ij} by the same amount for all j , which is exactly the row-shift gauge transformation.

Definition 9.2 (Flow Field Discretization — Research Hypothesis). *The transport plan $T_{ij}(t)$ is hypothesized to arise as the discretization of the RSVP vector flow \mathbf{v} via*

$$T_{ij}(t) = \mathcal{D}_{ij}[\mathbf{v}(\cdot, t)] \approx \int_{\partial V_{ij}} \mathbf{v}(x, t) \cdot \hat{n}_{ij} dA, \quad (23)$$

where $V_{ij} = \partial V_i \cap \partial V_j$ is the shared boundary between the Voronoi cells of nodes i and j , and \hat{n}_{ij} is the outward normal from V_i . This is a discretization relationship, not an identity.

Definition 9.3 (Accessibility Field Ansatz). *Ansatz: the entropy regularization parameter at edge (i, j) may be interpreted as a coarse-grained accessibility field:*

$$\varepsilon_{ij}(t) \approx \frac{1}{|V_{ij}|} \int_{V_{ij}} S(x, t) dA, \quad (24)$$

the average of S over the shared boundary. This is an interpretive identification, not a derivation. Its empirical content is that $\varepsilon_{ij}(t)$ should correlate with observable proxies for trade route accessibility—infrastructure quality, institutional development, logistical performance indices. Testing this correlation against the cost matrices of Gaskin et al. [2026] would constitute a direct empirical examination of the ansatz.

Under these identifications, the entropy-regularized OT problem (7) becomes a discretization of the RSVP free energy functional on the trade network, with each

term in the objective corresponding to a term in the RSVP Lagrangian.

RSVP Field Equations as Cost Dynamics

Derivation Principles

Before stating the field equations, we explain why the RSVP Lagrangian takes the form it does rather than some other form. Field theories are normally motivated by a set of construction principles, and the RSVP Lagrangian is no exception.

Locality: The Lagrangian density at a point x depends only on the fields and their first derivatives at x , not on field values at distant points. This is the standard locality assumption of continuum field theory and corresponds to the physical assumption that trade costs at a location respond to local accessibility conditions, not to global network properties.

Lowest-order derivative expansion: Among all local Lagrangians, the simplest are those involving at most first derivatives of the fields. The RSVP Lagrangian includes $(\partial_\mu \Phi)^2$ (the scalar kinetic term), $F_{\mu\nu} F^{\mu\nu}$ (the vector kinetic term, which involves first derivatives of \mathbf{v}), and coupling terms involving Φ and $\nabla \cdot \mathbf{v}$. Higher-derivative terms are excluded on grounds of simplicity; they would appear as corrections in a systematic expansion.

Symmetry: The Lagrangian should respect the symmetries of the underlying system. The trade network has no preferred spatial direction (at the global scale, no direction is privileged), suggesting that the scalar field should enter through its magnitude and that the vector field should enter through the rotationally invariant combination $\nabla \cdot \mathbf{v}$.

Admissibility: The entropy production inequality $\sigma_{\text{prod}} \geq 0$ is imposed as a constraint on trajectories rather than derived from the Lagrangian. This corresponds to imposing the Second Law as an external constraint on the variational problem.

Under these principles, the leading-order RSVP Lagrangian is essentially unique up to the choice of coupling constants a, b, c, d, λ, μ . Specifically, the terms $-\frac{1}{2}(\partial\Phi)^2$, $-\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$, $-\Phi \nabla \cdot \mathbf{v}$, $-\lambda(\mathbf{v} \cdot \nabla S)$, and $-\mu S(\nabla \cdot \mathbf{v})^2$ are the lowest-order local invariants involving the three fields and their first derivatives. The potential $V(\Phi, S)$ encodes self-interactions and equilibrium structure, with the Φ^4 term allowing spontaneous

symmetry breaking.

The RSVP Lagrangian density is

$$\mathcal{L} = -\frac{1}{2}(\partial_\mu \Phi)^2 - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \Phi \nabla \cdot \mathbf{v} - \lambda(\mathbf{v} \cdot \nabla S) - \mu S(\nabla \cdot \mathbf{v})^2 + V(\Phi, S), \quad (25)$$

where $F_{\mu\nu} = \partial_\mu v_\nu - \partial_\nu v_\mu$ is the vector field strength and $V(\Phi, S) = \frac{a}{2}\Phi^2 + \frac{b}{4}\Phi^4 + \frac{c}{2}(S - S_0)^2 + d\Phi^2 S$ is the interaction potential. The Euler-Lagrange equations give the dynamics for each field.

For the cost field Φ :

$$\square\Phi - \nabla \cdot \mathbf{v} + a\Phi + b\Phi^3 + 2d\Phi S = 0. \quad (26)$$

This is the equation of motion for the cost field. It says that the cost field at any point in the trade network is not arbitrary: it evolves according to a nonlinear wave equation driven by the divergence of trade flows and coupled to the accessibility field. A region where trade flows are converging ($\nabla \cdot \mathbf{v} < 0$) is one where cost pressure is building; a region where flows are diverging ($\nabla \cdot \mathbf{v} > 0$) is one where cost pressure is being released.

In the static equilibrium limit ($\square\Phi = 0, \partial_t \mathbf{v} = 0$):

$$\nabla^2\Phi = \nabla \cdot \mathbf{v} - a\Phi - b\Phi^3 - 2d\Phi S, \quad (27)$$

which is a nonlinear Poisson equation for the cost field sourced by the flow divergence. In regions with large inflows ($\nabla \cdot \mathbf{v} \ll 0$), Φ accumulates, driving costs up. In regions with large outflows ($\nabla \cdot \mathbf{v} \gg 0$), Φ relaxes, driving costs down. This is qualitatively consistent with the empirical observation that countries that absorbed diverted trade flows (e.g., Saudi Arabia absorbing Australian barley exports following the China ban) saw cost decreases, while countries that lost access to sources (African wheat importers following the Ukraine war) saw cost increases.

For the flow field \mathbf{v} :

$$\partial_\mu F^{\mu\nu} = \partial^\nu \Phi - \lambda \partial^\nu S - 2\mu S \partial^\nu (\nabla \cdot \mathbf{v}), \quad (28)$$

which has the form of a Maxwell equation with a composite source current. The trade flow is driven by the cost gradient $\nabla \Phi$ (flows move from high-cost to low-cost regions), the accessibility gradient $\lambda \nabla S$ (flows prefer accessible routes), and the divergence feedback $2\mu S \nabla (\nabla \cdot \mathbf{v})$ (regions of high trade activity generate self-organizing pressure).

For the accessibility field S :

$$S = S_0 - \frac{1}{c} [\lambda \nabla \cdot \mathbf{v} - \mu (\nabla \cdot \mathbf{v})^2 + d\Phi^2], \quad (29)$$

an algebraic equation relating accessibility to cost and flow divergence. High constraint load (Φ^2 large, i.e., high costs) decreases accessibility; high inflow ($\nabla \cdot \mathbf{v} < 0$ in our sign convention) increases it through the linear term. This captures the fact that heavily used trade routes become more accessible over time through infrastructure development, regulatory harmonization, and institutional trust-building—while high-cost routes remain underused and inaccessible.

The Sinkhorn Iteration as RSVP Gradient Descent

The Sinkhorn iteration (9) can be interpreted as gradient descent on the RSVP effective action at the network scale. Define the discrete RSVP free energy at scale Λ (the network resolution) as

$$F_\Lambda[T, C, S] = \sum_{ij} T_{ij} C_{ij} - \sum_{ij} S_{ij} H_{ij}(T), \quad (30)$$

where $H_{ij}(T) = -T_{ij}(\log T_{ij} - 1)$ is the per-edge entropy. The minimum of F_Λ with respect to T subject to the marginal constraints (5) is exactly the Sinkhorn solution (8) with $\varepsilon_{ij} = S_{ij}$.

The Sinkhorn iteration alternates between updating Π (enforcing the row marginals, i.e., supply constraints) and Ω (enforcing the column marginals, i.e., demand constraints). In RSVP terms, this alternation corresponds to iterating between the Φ -

equation (26) and the v-equation (28): the cost field adjusts to the flow divergence, then the flow adjusts to the updated cost field, and so on until convergence. The convergence of the Sinkhorn algorithm corresponds to the RSVP system reaching a static admissible equilibrium.

Proposition 9.1 (Sinkhorn as RSVP Relaxation). *The Sinkhorn iteration (9) is the discretization of the RSVP gradient flow*

$$\partial_t \Pi = -\nabla_{\Pi} F_{\Lambda}, \quad \partial_t \Omega = -\nabla_{\Omega} F_{\Lambda} \quad (31)$$

projected onto the constraint manifold defined by the marginal conditions (5).

The proposition follows from the observation that the Lagrange multiplier updates in Sinkhorn's algorithm are exactly the gradient descent steps on F_{Λ} with respect to the dual variables λ_i (encoding supply) and η_j (encoding demand), projected to maintain non-negativity.

This interpretation gives the Sinkhorn algorithm a physical meaning: it is not merely an efficient numerical method but the actual relaxation dynamics of the cost-flow system toward its equilibrium configuration. The number of Sinkhorn iterations required for convergence is therefore related to the relaxation time of the RSVP system at the network scale—a quantity that itself varies across the network and over time.

The Admissibility Correspondence

The dual OT admissibility condition (11) states that the sum of the collection cost f_i and the delivery profit g_j must not exceed the transport cost C_{ij} . In RSVP terms, the dual variables f_i and g_j are the values of the RSVP potential Φ at the node locations: $f_i = \Phi(x_i)$ and $g_j = -\Phi(x_j)$ (with sign conventions from the direction of trade). The admissibility condition becomes

$$\Phi(x_i) - \Phi(x_j) \leq \int_{x_i}^{x_j} \Phi \, dl, \quad (32)$$

which is a non-negativity condition on the line integral of Φ along any path between nodes. This is satisfied automatically when $\Phi \geq 0$ everywhere along the

path—which is the RSVP condition that the constraint field has no negative regions.

The RSVP admissibility constraint $\sigma_{\text{prod}} \geq 0$ (entropy production non-negative) then corresponds to the statement that the OT dual solution (f, g) cannot be improved by any perturbation that violates the dual admissibility condition: the trade system is at the boundary of admissible configurations where no route can be used more efficiently without violating the cost constraint. This is the economic content of the RSVP Second Law: trade routes that are used are used as efficiently as constraints permit; unused routes are inadmissible given current cost and accessibility conditions.

Reinterpreting the Empirical Findings

Retrodictive and Predictive Uses of the Framework

Any theoretical framework can be used retrospectively to reinterpret known events. The reinterpretations below should therefore be understood primarily as demonstrations of consistency between the RSVP framework and the Gaskin et al. findings, not as confirmations of the framework. The empirical case studies show that the RSVP language is adequate for describing the observed cost field deformations; they do not show that the RSVP field equations *predicted* those deformations in advance.

The framework gains empirical credibility only to the extent that it successfully predicts future deformations of the cost landscape before they occur. Specific falsifiable predictions are stated in Section 17. Until those predictions are tested, the reinterpretations below should be read as illustrations of the framework’s conceptual reach rather than as evidence for its correctness.

The Ukraine war shock, as documented by Gaskin et al., is best understood as an abrupt deformation of the RSVP cost field Φ in the wheat trade network. Before the shock, the cost field had established a stable configuration with low Φ along the Black Sea corridors (high accessibility, well-established trade routes) and higher Φ in alternative corridors. The shock imposed a sudden increase in Φ along the Ukraine-export corridors (port closures, blockades, insurance cost spikes) and si-

multaneously altered the accessibility field S by reducing it in affected routes.

The RSVP cost field equation (26) predicts how the shock propagates. An increase in Φ along Ukraine's export corridors raises C_{ij} for Ukraine- j routes, which reduces T_{ij} (fewer goods flow along these routes). The reduced flow means the divergence term in (26) changes, and the cost field adjusts throughout the network. Specifically, countries that previously sourced wheat from Ukraine now face higher costs; if alternative suppliers exist (Russia, Canada, Australia), the flow v redirects, and Φ decreases along the alternative corridors as those routes become better-used and more accessible.

The empirical finding that African importers experienced both volume drops and cost increases, while European importers experienced volume drops but smaller cost increases, is explained by the accessibility field S . European countries have higher baseline accessibility (better infrastructure, established alternative suppliers, lower logistical friction) so their S_{ij} values for alternative routes are higher—meaning they can absorb the shock by switching to alternative suppliers without a large cost penalty. African importers have lower baseline accessibility for alternative routes, so the same volume drop is accompanied by a larger cost increase because the alternative routes have lower S and therefore higher effective costs.

This is not merely a post-hoc redescription. The RSVP framework predicts that the cost-volume relationship should be heterogeneous across the network in a way that correlates with the accessibility field, and that the accessibility field should correlate with infrastructure quality, institutional development, and historical trade relationships. These predictions can be tested against the Gaskin et al. cost estimates.

Brexit as Persistent Cost Field Asymmetry

The Brexit finding—that UK trade costs with Europe increased more than Irish trade costs with Europe—is a signature of a persistent asymmetry in the cost field Φ introduced by a regulatory boundary. The UK exit from the common market and customs union increased Φ along UK-EU routes by introducing non-tariff barriers, customs delays, and regulatory divergence. Ireland (ROI), remaining in the common market, did not experience this increase.

The asymmetry $C_{\text{UK-EU}} > C_{\text{ROI-EU}}$ is a stable deformation of the cost field, not a transient one. The RSVP field equation (26) predicts that this asymmetry will persist as long as the regulatory boundary remains—it is a stable attractor of the cost field dynamics, maintained by the ongoing divergence between UK and EU regulatory frameworks (which enters through the potential $V(\Phi, S)$ as a boundary condition).

Interestingly, the Gaskin et al. finding that UK vegetable imports from Morocco increased (with a precipitous drop in trade costs) while Irish imports from Morocco did not follow the same pattern is interpretable as the UK seeking new admissible routes under the changed cost field. In RSVP terms: the UK's original routes to European suppliers became less admissible (higher Φ), so the flow \mathbf{v} redirected toward alternative admissible routes where Φ was lower—Morocco in this case. Ireland, whose European routes remained admissible, had no such pressure to redirect.

Free Trade Agreements as Accessibility Field Manipulation

The Asia-Pacific findings reveal an important distinction between two mechanisms for cost reduction: genuine trade liberalization (which reduces Φ) and accessibility enhancement (which increases S). The Gaskin et al. finding that ChAFTA (2015) had little measured effect on Australian trade costs with China, because costs had already fallen substantially between 2000 and 2015, suggests that the formal agreement ratified and institutionalized accessibility gains that had already occurred through the deepening of trade relationships.

In RSVP terms: the accessibility field S along Australia-China routes had been increasing since China's WTO accession in 2001, reflecting growing familiarity, better logistics, and more established relationships. By 2015, the routes were already highly accessible (S near its maximum), and the formal agreement added little to accessibility. The agreement did lower Φ slightly for some commodities (wine, dairy), but the dominant effect had already occurred.

By contrast, the Modi visit in 2015 and the subsequent rapid increase in India-China trade represents a sudden jump in S for India-China routes: a policy event that increased the accessibility of routes that were previously formally open but

informally constrained. The cost drop of 33% in sugar trade from India to China between 2015 and 2022 corresponds to the accessibility field S rapidly approaching its equilibrium value for a trade relationship whose formal barriers had been low but whose informal barriers were high.

The Inverse Problem as RSVP Field Inference

Learning $C(t)$ versus Inferring $\Phi(x, t)$

Both the Stuart-Wolfram and Gaskin et al. approaches treat each time step as an independent inference problem: they learn $C(t)$ from $T(t)$ without using the dynamics of C between time steps. This is a significant limitation: it means the learned cost matrix at time t does not inform the estimate at time $t + 1$, and the temporal evolution of costs is observed only retrospectively rather than predicted prospectively.

The RSVP framework resolves this limitation by providing a dynamics for $\Phi(x, t)$. Given observations of $T(t)$ —which correspond to observations of \mathbf{v} at the network nodes—the RSVP inverse problem is to infer $\Phi(x, t)$ from the observed flow field, using the field equations (26)–(29) as constraints. This is a dynamical inverse problem: the cost field is inferred continuously, and its evolution is governed by the PDEs, not by an independent neural network estimate at each time step.

Formally, the RSVP inverse problem for trade networks is:

Definition 11.1 (RSVP Trade Inverse Problem). *Given observations $\{T_{ij}(t_k)\}_{k=1}^K$ of trade flows at times $t_1 < \dots < t_K$, find the RSVP field triple (Φ, \mathbf{v}, S) on $M \times [t_1, t_K]$ such that:*

1. *The fields satisfy the RSVP field equations (26)–(29).*
2. *The discrete approximations (22)–(24) reproduce the observed transport plans $T(t_k)$ via the Sinkhorn forward map.*
3. *The admissibility constraint $\sigma_{\text{prod}} \geq 0$ is satisfied throughout.*

This inverse problem is more constrained than the neural IOT approach of Gaskin et al. because the RSVP field equations connect the cost estimates at different times. It is also more informative: the inferred $\Phi(x, t)$ is a continuous field over the trade

network, not a matrix of node-to-node costs, and it can be used to predict cost evolution between observation times and to attribute cost changes to specific physical mechanisms (flow divergence, accessibility changes, potential landscape shifts).

The Bayesian-RSVP Synthesis

The Stuart-Wolfram Bayesian framework is compatible with and complementary to the RSVP dynamics. The posterior distribution (15) over cost matrices can be replaced by a posterior distribution over RSVP field configurations, with the RSVP field equations imposed as structural constraints (analogous to the graph-based cost structure in Stuart and Wolfram):

$$P(\Phi, \mathbf{v}, S \mid \{T(t_k)\}) \propto \exp\left(-\sum_k \frac{1}{2\sigma^2} \left|T(t_k) - \hat{T}[\Phi, \mathbf{v}, S](t_k)\right|^2\right) \cdot \mathbf{1}_{\mathcal{A}}(\Phi, \mathbf{v}, S), \quad (33)$$

where $\hat{T}[\Phi, \mathbf{v}, S](t_k)$ is the Sinkhorn transport plan computed from the RSVP fields at time t_k , and \mathcal{A} is the admissible region (configurations satisfying the RSVP field equations and the entropy production inequality). The indicator $\mathbf{1}_{\mathcal{A}}$ plays the role of the prior in the Stuart-Wolfram framework: it encodes the structural constraints on the cost field.

This Bayesian-RSVP synthesis has several advantages over either approach alone. The RSVP field equations constrain the temporal evolution of costs, reducing the effective dimensionality of the inference problem. The Bayesian framework provides uncertainty quantification for the inferred fields, propagating data uncertainty to the cost field estimate. And the admissibility constraint $\sigma_{\text{prod}} \geq 0$ provides a physical prior that excludes cost field configurations whose evolution would violate thermodynamic irreversibility—a constraint with genuine predictive content.

What OT Cannot Yet Explain and What RSVP Provides

The Missing Dynamics

OT in both papers is inherently static: it finds the optimal transport plan given a cost matrix, or infers the cost matrix from an observed plan. It does not explain

why costs change. The temporal dimension in Gaskin et al. is handled by training the network on time-indexed data and inferring $C(t)$ at each t independently; there is no model of $C(t + 1)$ given $C(t)$.

RSVP provides the missing dynamics via the field equations (26)–(29). These equations predict how the cost field evolves given the current flow configuration and accessibility structure. They allow, in principle, the prediction of future cost changes from current field states—a capability that neither the neural IOT approach nor the Bayesian IOT approach possesses.

The Missing Spatial Structure

The OT framework treats the cost matrix C_{ij} as a matrix of numbers indexed by node pairs, with no spatial structure beyond what is encoded in the graph topology. The Stuart-Wolfram graph-based cost (shortest paths on a border-adjacency graph) introduces some spatial structure, but it is still determined by a finite-dimensional vector of edge weights rather than a continuous field.

RSVP provides continuous spatial structure via the field $\Phi(x, t)$. The cost between any two countries is not a primitive but is derived from the line integral of Φ along the trade route. This allows the cost field to be spatially smooth (nearby countries have similar costs), to respond to spatial events (a blockade raises Φ in a specific region), and to be resolved at finer spatial scales than the country-level network.

The Missing Physical Interpretation of ε

As noted in Section 4, the entropy regularization parameter ε has a physical interpretation as the thermodynamic temperature of the trade network—the degree to which traders tolerate suboptimal routes. In Gaskin et al., ε is treated as a fixed hyperparameter (the paper notes a hyperparameter sweep was performed). In Stuart and Wolfram, ε appears as the regularization parameter chosen relative to the noise level σ .

RSVP provides a dynamical interpretation: $\varepsilon_{ij}(t) = S_{ij}(t)$ is the accessibility field, which evolves according to the equation of state (29). This means ε is not a global constant but a field quantity that varies across the network and over time, responding to changes in cost and flow. A trade route that becomes highly trafficked be-

comes more accessible (S_{ij} increases), effectively lowering the temperature for that route and making the flow more deterministic. A route that is disrupted or abandoned becomes less accessible (S_{ij} decreases), raising the temperature and making the flow more diffuse.

This interpretation has concrete predictive content: $\varepsilon_{ij}(t)$ should correlate with observable proxies for trade route accessibility—infrastructure quality indices, bilateral institutional quality, logistical performance indices. Testing whether the learned ε in Gaskin et al. correlates with such proxies would be a direct empirical test of the RSVP identification.

Toward a General Theory of Hidden Fields

Trade, Migration, and the Universality of the Inverse Problem

Gaskin et al. note explicitly that their methodology is general: it applies to any flow system, not only agricultural trade. Stuart and Wolfram illustrate this immediately with European migration flows [Raymer et al., 2013, Abel and Sander, 2014], treating the number of people moving between countries as a transport plan whose underlying cost encodes the difficulty of migration along each route. The same Bayesian inverse OT framework recovers the cost of crossing each border from observed migration statistics, providing estimates that are sensitive to the distinction between formal policy barriers and informal cultural or economic ones.

Migration as a flow phenomenon has exactly the structure of trade as a flow phenomenon. There is a supply of would-be migrants at each origin, a demand (or absorptive capacity) at each destination, and a cost landscape that determines which movements actually occur. The cost landscape for migration includes visa requirements and enforcement stringency, economic wage differentials, cultural distance (language, religion, shared diaspora networks), physical geography (distance, border infrastructure), and political climate at the destination. Most of these are not directly measurable, and the gravity model analogs for migration have the same limitations as for trade: they parameterize the cost through proxies and absorb the unobservables into residuals.

Inverse OT applied to migration is not merely an analogy with trade; it is the

same mathematical structure applied to a different domain. The hidden manifold is the migration accessibility landscape—the field that determines, at each moment, which origin-destination pairs are accessible, at what cost, and with what uncertainty. Recovering this field from observed migration stocks and flows is the inverse problem. The RSVP interpretation is immediate: the cost field Φ encodes migration resistance, the flow field \mathbf{v} encodes the actual migration trajectories, and the accessibility field S encodes the degree to which each route is operationally open to would-be migrants.

Cognition as Trajectory on a Constraint Manifold

The extension to cognition is structurally analogous, though phenomenologically different. Cognitive processes can be understood as trajectories on a constraint manifold: sequences of mental states that move through a configuration space governed by constraints of consistency, attention, memory, and relevance. The observed trajectory is behavior—the verbal output, the decision made, the inference drawn—and the hidden manifold is the constraint geometry that made some trajectories admissible and others impossible.

This interpretation connects to predictive processing frameworks [Friston, 2010] and to the broader project of understanding inference as constrained optimization under uncertainty [Jaynes, 2003, MacKay, 2003]. But the inverse problem formulation sharpens it: given the observable trajectory (the behavior), what does the constraint manifold look like? What is the hidden accessibility geometry that the cognitive system is navigating?

The analogy with inverse OT is not perfect, because cognitive trajectories are not obviously solving a cost-minimization problem in the same sense as trade flows. But the structural parallel holds: both involve a hidden field (the constraint geometry in cognition, the cost landscape in trade) that is not directly observable, and both require an inverse inference from observable motion to recover the hidden structure. The compression in cognitive observation is if anything more severe than in trade: we observe behavioral outputs that are projections of neural activity that is itself a projection of a high-dimensional constraint space.

What RSVP adds to cognitive modeling is the field-theoretic perspective: the con-

straint manifold is not static but evolves dynamically under the influence of attention (the flow field \mathbf{v}), accumulated knowledge and priors (the scalar constraint field Φ), and the accessibility of mental states given current cognitive resources (the entropy-accessibility field S). A thought that is accessible in a high- S cognitive state (relaxed, well-rested, unconstrained attention) may be inadmissible in a low- S state (fatigued, anxious, overloaded). The inverse problem for cognition is then: given observable cognitive outputs, recover the accessibility landscape that produced them.

Physics: Inferring Fields from Trajectories

The deep structure of the inverse problem is most explicit in physics, where the entire enterprise of field theory can be understood as recovering hidden manifolds from visible trajectories. Newton inferred the gravitational field from the elliptical orbits of planets; the field was not observed, only its effect on trajectories. Maxwell inferred the electromagnetic field from the trajectories of charged particles and the patterns of light. Einstein inferred spacetime curvature from the geodesics of light and matter, treating the curvature as the hidden manifold whose structure determines which trajectories are admissible.

In each case the logical structure is: trajectories are observed, the field is inferred, and the field equations provide the dynamics that explain why the observed trajectories occurred and predict future ones. The inverse problem is not a secondary activity but the primary mode of field discovery. We never observe the gravitational field directly; we infer it from the motion it produces.

Modern experimental physics has made this structure explicit through inverse scattering problems, tomographic reconstruction, and lattice gauge theory [Taranola, 2005, Kaipio and Somersalo, 2006]. In inverse scattering, the hidden manifold is the potential energy landscape, and the observable trajectories are scattering cross-sections. The potential is inferred by inverting the scattering matrix. In tomography, the hidden manifold is the density distribution inside a body, and the observable trajectories are X-ray attenuation measurements; the reconstruction algorithm inverts the Radon transform. In lattice gauge theory, the hidden manifold is the gauge field configuration, and the observable quantities are correlation functions of gauge-invariant operators inferred from Monte Carlo sampling of the path

integral.

The inverse OT approach to trade is a direct descendant of this physical tradition [Villani, 2003, Peyré and Cuturi, 2019], translated into an economic context. The hidden manifold is the cost field; the observable trajectories are trade flows; and the neural network or Bayesian algorithm performs the inversion. What RSVP provides is the field equations for the hidden manifold—the analog of Maxwell’s equations or the Einstein field equations for the trade cost landscape.

Compressed Causality and the Stability of Reconstruction

Observed Flows as Compressed Information

The trade matrix $T_{ij}(t)$ is a compressed observation of the full causal history of the trade relationship between countries i and j . It is compressed in multiple senses. Temporally: it aggregates all transactions over a year into a single number. Causally: it reflects the combined effect of thousands of individual decisions, each responding to a complex mix of prices, regulations, relationships, and expectations. Informationally: it loses the intermediate states—the negotiation, the contract, the logistics arrangement—that connect the geopolitical cause to the trade effect.

This compression raises a fundamental question: how much of the hidden manifold can be reconstructed from the compressed observation [cf. the general theory of causal inference from compressed data in Pearl, 2009, Cover and Thomas, 2006]? The Stuart-Wolfram analysis addresses this through identifiability: the inverse problem is overdetermined when the cost structure is low-dimensional (Toeplitz, graph-based) and underdetermined when it is high-dimensional (general non-symmetric). As dimension increases, the posterior distribution over cost matrices becomes diffuse, reflecting the fact that many cost matrices are consistent with the observed flows.

This is the compressed causality problem: not all causal information survives the compression from the full causal trajectory to the observed trade matrix. The information that survives is precisely the gauge-invariant content of the cost field—the part that determines flows and cannot be absorbed into row-column shifts. What is

lost is the absolute level of costs, information about unused routes, and the within-year dynamics of trade adjustment.

Conditions for Successful Reconstruction

Stuart and Wolfram identify several conditions under which reconstruction succeeds. The problem is overdetermined (more constraints than unknowns) when the cost structure is sufficiently constrained—when the cost is Toeplitz, or when it is determined by a graph with far fewer edges than country pairs. In these cases, the posterior over costs is well-concentrated and the reconstruction is reliable. When the cost structure is fully general (every country-pair cost is independent), the problem is underdetermined and the posterior is diffuse.

Gaskin et al. address this through the neural network architecture: by training a network that maps trade matrices to cost matrices, they implicitly impose a low-dimensional structure on the cost landscape—the structure that the network’s weights can represent. The network’s capacity acts as a regularizer, preventing the trivial solution where each time step’s cost matrix is fitted independently with no shared structure.

From the RSVP perspective, the RSVP field equations themselves impose the most principled regularization. They say that the cost field at time $t + 1$ is not independent of the cost field at time t —it is related through the field equations (26). This temporal regularization uses physical structure rather than the implicit regularization of a neural network architecture, and it yields a reconstruction that is interpretable in terms of identified mechanisms (flow divergence, accessibility change, potential landscape shift) rather than latent network representations.

Trade as a Reconstruction Problem

The unified picture is: trade flows are compressed observations of an underlying cost field, and both the OT forward problem and the OT inverse problem are instances of the general reconstruction problem of recovering a hidden manifold from its observable projections. The gravity model is a prior over the shape of the hidden manifold that privileges a particular low-dimensional parameterization. Inverse OT is a data-driven reconstruction that imposes the OT structure but not

the functional form. RSVP inverse OT is a dynamically constrained reconstruction that uses the field equations as additional prior structure.

Each level of the hierarchy adds structure. The gravity model adds the most interpretable structure but is most vulnerable to misspecification. Inverse OT adds less structure but is more robust. RSVP inverse OT adds more structure through the dynamics but requires the RSVP field equations to be approximately correct—a stronger assumption that yields stronger predictions.

Philosophical Implications

The Territory Behind the Data

The shift from gravity models to inverse OT to RSVP field theory represents, at the philosophical level, a shift in what we take to be the fundamental object of study. In the gravity model tradition, trade is fundamentally a collection of bilateral transactions, each explained by the characteristics of the two parties and the distance between them. The data are the primary reality; the model is a summary of the data.

In the inverse OT tradition, trade is fundamentally an equilibrium flow on a hidden accessibility landscape. The data (trade volumes) are secondary—they are projections of the primary reality (the cost landscape). The model is not a summary of the data but a reconstruction of the structure that generated the data. This is the epistemic priority of the hidden manifold over the visible trajectory.

In the RSVP tradition, the hidden manifold is itself a dynamical object, governed by field equations that connect different moments in time and different locations in space. The cost landscape at any moment is not an independent object but a state of a dynamical system whose history and future are constrained by the field equations and the admissibility condition. Understanding trade means understanding the dynamics of the cost field, not merely its current state.

Constraint Before Content

The philosophical shift implicit in this progression can be stated as a principle: constraint precedes content. The accessible configurations of a system—the ad-

missible trajectories, the states reachable from a given initial condition subject to the field equations and the entropy production inequality—are prior to the specific configurations observed at any given time. What we observe is always a sample from what is admissible; the task of science is to recover the admissibility structure.

This is the reverse of the standard empiricist picture, in which observations are primary and theories are summaries of observations. In the constraint-before-content picture, the admissibility structure is primary and observations are samples from it. The scientific work is not to summarize the samples but to infer the structure that generated them. This is inverse reasoning: from effects to causes, from trajectories to fields, from visible motion to hidden geometry.

The OT papers make this structure vivid in the economic context. The disproportionate cost increase faced by African wheat importers after the Ukraine war is not explained by anything visible in the trade volume data alone. It is explained by the deformation of the hidden cost field—a change in the accessibility landscape that volume data can only partially reveal. Recovering the cost field from the volume data is the scientific act; without it, the observation is brute fact rather than explained phenomenon.

From Objects to Accessibility

The deepest implication is ontological. The gravity model treats countries as objects with economic mass and distance-like separation—a point-particle ontology applied to trade. Optimal transport treats trade as a flow on a cost landscape—a field ontology applied to economics. RSVP generalizes this: the cost field, the flow field, and the accessibility field are the fundamental objects, and the countries are nodes in a network that the fields govern.

This shift from object ontology to field ontology is the same shift that physics underwent in the nineteenth and twentieth centuries. Newton’s mechanics is an ontology of particles with properties (mass, position, velocity) interacting through forces. Maxwell’s electrodynamics is an ontology of fields that permeate space and time, and particles that respond to those fields. General relativity makes the field primary in the strongest sense: spacetime itself is a field, and particles are excitations or trajectories in that field.

The OT approach to trade and migration represents a similar conceptual shift in economics and demography. Countries are not the fundamental objects; the cost field is. Trade volumes are not the fundamental data; the accessibility landscape is. The scientific explananda are not why some countries trade more than others but what the cost field looks like and why it has the shape it does.

RSVP makes this shift explicit through its field-theoretic formulation. The field triple (Φ, \mathbf{v}, S) is the fundamental ontology; the node-level trade statistics are its compressed projections. Recovering the field from its projections—the inverse OT problem extended to RSVP dynamics—is the project of understanding not just what happens in trade networks but why.

The Hierarchy of Scientific Description

The grammar/semantics/substrate hierarchy introduced in Section 1.3 can now be made precise in light of the full development.

Grammar: Parameterized Description

A grammar assigns observable phenomena to pre-defined categories and fits parameters to reproduce observed outputs. Gravity models are the canonical example in trade economics: the cost C_{ij} is parameterized as a log-linear combination of measurable covariates, and regression estimates the parameters. The grammar is useful when the categories are stable and the parameters are interpretable, but it fails when the underlying structure changes in ways the categories do not cover.

The formal structure of grammar is: assume a parameterized cost function $C_{ij}(\theta)$, fit θ to data T , report $\hat{\theta}$. The hidden manifold is identified with the parameterized family; the fitted parameters are the model's answer to what the hidden manifold looks like.

Semantics: Structural Recovery

Semantics recovers the structure of the hidden object from data without presupposing its functional form. Inverse OT is the canonical example: it recovers C_{ij} directly from T_{ij} , using only the OT structural constraint that trade flows minimize cost. The recovered cost matrix is the model's answer to what the hidden

manifold looks like, and it is more faithful than the grammar answer because it does not constrain the manifold's shape in advance.

The formal structure of semantics is: given data y and a forward map \mathcal{F} , solve $\hat{x} = \arg \min_x \|y - \mathcal{F}(x)\|^2 + \lambda R(x)$. The regularization $R(x)$ encodes whatever prior knowledge is available about admissible solutions.

Substrate: Dynamic Generation

A substrate provides the field equations from which the hidden object evolves. RSVP is the candidate substrate proposed here: the scalar field Φ , vector field \mathbf{v} , and accessibility field S are governed by coupled nonlinear PDEs whose solutions, projected to the trade network, produce the cost matrices recovered by inverse OT.

The formal structure of substrate is: specify a Lagrangian $\mathcal{L}(\Phi, \mathbf{v}, S, \partial\Phi, \partial\mathbf{v}, \partial S)$, derive the Euler-Lagrange equations, impose the admissibility constraint $\sigma_{\text{prod}} \geq 0$, and read off the cost field as the projection $C_{ij} = \mathcal{P}_{ij}[\Phi]$.

Hidden Manifold as Invariant Object

The hidden manifold is the object that all three levels describe, each with different fidelity and different assumptions. The grammar describes it through a low-dimensional parameterization. Semantics describes it through data-driven recovery. The substrate describes it as a dynamical field with equations of motion.

The hidden manifold is not any one of these descriptions but the invariant object that underlies all of them. It is what would remain if we could remove all modeling assumptions and observe the cost landscape directly. Since direct observation is impossible—the cost landscape is by definition not observable—all we can do is approach it from different directions with different tools and check whether the descriptions are mutually consistent.

The strongest evidence for the reality of the hidden manifold would be the convergence of all three approaches: that gravity model parameters, inverse OT cost matrices, and RSVP field solutions all point to the same underlying accessibility structure. That convergence has not yet been demonstrated, but demonstrating it is the empirical program that the hidden manifold thesis proposes.

Open Mathematical Problems

The argument of this essay rests on several mathematical claims that remain unproved. We state them explicitly as open problems, both to be honest about the current state of the framework and to identify concrete directions for further work.

Uniqueness of Field Reconstruction

The central open problem is: under what conditions does a sequence of cost matrices $\{C_{ij}(t_k)\}_{k=1}^K$ uniquely determine an RSVP field configuration (Φ, \mathbf{v}, S) on $M \times [t_1, t_K]$?

As discussed in Section 7.2, the projection $\mathcal{P}_{ij} : \Phi \mapsto C_{ij}$ is generically non-injective. Additional constraints are needed. The question is whether the RSVP field equations, the admissibility condition $\sigma_{\text{prod}} \geq 0$, and appropriate boundary conditions are sufficient to make the inverse map well-defined. The analogous result in physics—the reconstruction of a metric from geodesic data—is highly non-trivial and requires strong symmetry assumptions. A reconstruction theorem for RSVP fields from cost matrices would be the foundational result justifying the entire framework.

Network-to-Manifold Convergence

The OT framework operates on a finite graph (countries as nodes, bilateral flows as edges). The RSVP framework operates on a continuous manifold. The transition between them requires a continuum limit: as the number of nodes increases and their spacing decreases, does the discrete cost matrix converge to the restriction of a continuous field Φ to node pairs?

This is a question about the consistency of the coarse-graining procedure. In statistical mechanics, the analogous question is whether the discrete lattice model converges to a continuum field theory in the thermodynamic limit. For trade networks, the question is complicated by the irregular spacing of countries, the asymmetry of flows, and the time-varying nature of the graph structure. A precise convergence theorem would require specifying the topology of the manifold M , the embedding of countries into M , and the regularity conditions on Φ .

Accessibility Field Estimation

The accessibility field ansatz $\varepsilon_{ij}(t) \approx \langle S \rangle_{\partial V_{ij}}$ has an empirical implication: the entropy regularization parameter at each edge should correlate with observable measures of trade route accessibility. Can $S_{ij}(t)$ be estimated directly from data, independently of the cost matrix, using logistics performance indices, bilateral institutional quality measures, or infrastructure quality data?

If yes, this would provide an independent check on the ansatz: the estimated $S_{ij}(t)$ should predict the optimal $\varepsilon_{ij}(t)$ in the Gaskin et al. framework better than a constant global ε . If no correlation is found, the ansatz is disconfirmed.

Dynamic Inverse Optimal Transport

Both the Stuart-Wolfram and Gaskin et al. approaches treat successive time steps as independent inference problems. A dynamic extension would incorporate the RSVP field equations as temporal constraints: the inferred $C(t+1)$ is constrained by the inferred $C(t)$ and the RSVP dynamics. This amounts to replacing the static posterior (33) with a filtering distribution over field trajectories.

The mathematical challenge is that the RSVP field equations are nonlinear PDEs whose solution is computationally demanding. Integrating them into a Bayesian or neural inference loop requires either efficient numerical solvers or a surrogate model that approximates the field dynamics. This is an active area of research in scientific machine learning and data-driven PDE discovery.

Empirical Tests of the RSVP Identification

The framework makes several predictions that could in principle be tested against the cost matrices inferred by Gaskin et al. [2026]:

The spatial correlation structure of $C_{ij}(t)$ should match the correlation structure predicted by the RSVP field equations—in particular, geographically proximate country pairs should have more strongly correlated cost changes than distant pairs.

The temporal autocorrelation of $C_{ij}(t)$ should follow the relaxation timescale of the RSVP field, which is determined by the coupling constants λ and μ .

The relationship between cost changes and flow divergence should satisfy the

static equilibrium condition (27): regions of net inflow should see cost increases and regions of net outflow should see cost decreases.

The entropy regularization parameter $\varepsilon_{ij}(t)$ inferred by Gaskin et al. (if estimated per-edge rather than globally) should correlate with the World Bank Logistics Performance Index and similar accessibility measures.

None of these tests have yet been performed. Performing them would constitute the first empirical examination of the RSVP framework in an economic context.

A Minimal Reconstruction Conjecture

The strongest version of the RSVP program requires proving that cost matrices uniquely determine field configurations—a theorem that does not yet exist and is likely to require strong additional assumptions. A weaker and more attainable standard is sufficient for many practical purposes.

Minimal Reconstruction Conjecture: If a family of RSVP field configurations exists whose coarse-grained projections reproduce the observed cost matrices and whose dynamics correctly predict future cost observations, then the RSVP framework possesses explanatory utility even in the absence of a uniqueness theorem.

Under this weaker criterion, predictive adequacy precedes ontological uniqueness. The framework is validated not by proving that the hidden field is unique but by demonstrating that the field equations make correct predictions about how the cost landscape will evolve. This is the same standard by which quantum electrodynamics is validated: not by proving that the electromagnetic field is the unique field consistent with all historical observations, but by showing that its equations predict future measurements with extraordinary accuracy.

Conclusion: Reconstructing Hidden Worlds

The significance of the optimal transport framework for trade modeling is not merely that it predicts flows better than the gravity model, though it does—often by an order of magnitude in RMSE. Its deeper contribution is methodological and ontological. It demonstrates that observable flows can be used to reconstruct hid-

den accessibility landscapes, transforming trade from a problem of bilateral exchange into a problem of latent geometry. The same conceptual move appears in inverse problems across economics, physics, cognition, and migration, and the RSVP framework provides the field-theoretic substrate that explains why the hidden landscape has the structure it does and how it evolves.

The three-layer architecture clarifies the relationships among the approaches surveyed: OT is the substrate (cost minimization as the generative principle of trade flows), Bayesian or neural inverse OT is the semantics (the operational procedure for inferring cost from observed flows), and gravity models are the grammar (a symbolic compression of cost structure in terms of observable covariates). The grammar is useful when precise inference is unavailable; the semantics is necessary for rigorous inference; but only the substrate plus its field-theoretic dynamics—the RSVP completion—can explain why costs have the structure they do and predict how they will evolve.

The empirical findings of Gaskin et al. become, under this framework, instances of a general pattern: geopolitical shocks, regulatory changes, and accessibility dynamics produce deformations of the hidden cost field that propagate through the trade network in ways governed by the field equations. The Ukraine wheat shock is a sudden increase in Φ along Black Sea corridors that propagates through the network's vector flow field, with disproportionate cost impact in regions where the accessibility field S is low. Brexit is a persistent cost field asymmetry maintained by a regulatory boundary condition. The pre-agreement cost drops in Asia-Pacific are increases in the accessibility field S along established trade corridors that the formal agreements ratified but did not originate.

The general principle is: observable flows are projections of hidden manifolds of constraint. Science proceeds by inverting those projections—by recovering the hidden structure from visible motion. This is what Newton did with gravity, Maxwell with electromagnetism, Einstein with spacetime, and what Gaskin, Stuart, Wolfram, and their collaborators are doing with trade and migration. The hidden manifolds are always there; the challenge is to find the right inverse procedure to recover them.

The work of developing that procedure—field-theoretically, computationally, and

empirically—continues.

Scope of the Claim

This essay does not establish that trade networks are governed by RSVP dynamics. Nor does it establish a unique reconstruction of hidden fields from observed flows. Its more modest claim is that inverse optimal transport demonstrates the recoverability of latent accessibility structure, and that field-theoretic completions such as RSVP provide one possible framework for understanding the dynamics of that structure. The validity of this broader program ultimately depends upon future mathematical reconstruction theorems and empirical tests of the kind identified in Section 17.

Remark 18.1 (Hidden Manifold Principle). *Observable trajectories do not constitute the primary object of scientific explanation. Rather, trajectories are projections of admissible regions within a hidden constraint manifold. The central scientific task is therefore the inverse reconstruction of that manifold from its observable projections. This principle is instantiated by Newtonian gravitation (orbits project the gravitational field), by electromagnetism (particle paths project the electromagnetic field), by inverse scattering (cross-sections project the potential landscape), by inverse optimal transport (trade flows project the cost landscape), and by the RSVP program (cost matrices project the accessibility field). The unifying structure is:*

$$\text{Observed Flow} \xrightarrow{\text{inverse problem}} \text{Hidden Geometry} \xrightarrow{\text{field equations}} \text{Dynamics and Prediction}.$$

References

- Guy J. Abel and Nikola Sander. Quantifying global international migration flows. *Science*, 343(6178):1520–1522, 2014.
- James E. Anderson and Eric van Wincoop. Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192, 2003.

- Thomas Cover and Joy Thomas. *Elements of Information Theory*. Wiley, 2nd edition, 2006.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. 26, 2013.
- Karl Friston. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- Alfred Galichon. *Optimal Transport Methods in Economics*. Princeton University Press, 2016.
- Thomas Gaskin, Guven Demirel, Marie-Therese Wolfram, and Andrew Duncan. Modelling global trade with optimal transport. *Nature Communications*, 17:2947, 2026. doi: 10.1038/s41467-026-69694-5.
- Edwin T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- Jari Kaipio and Erkki Somersalo. *Statistical and Computational Inverse Problems*. Springer, 2006.
- Leonid V. Kantorovich. On the translocation of masses. *Doklady Akademii Nauk USSR*, 37:199–201, 1942.
- David J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
- Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences*, 1781.
- Judea Pearl. *Causality*. Cambridge University Press, 2nd edition, 2009.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5–6):355–607, 2019.
- James Raymer et al. Integrated modeling of european migration. *Journal of the American Statistical Association*, 108(503):801–819, 2013.
- Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

João Santos Silva and Silvana Tenreiro. The log of gravity. *Review of Economics and Statistics*, 88(4):641–658, 2006.

Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Annals of Mathematical Statistics*, 35:876–879, 1964.

Andrew M. Stuart and Marie-Therese Wolfram. Inverse optimal transport. *arXiv preprint*, arXiv:1905.03950, 2019. URL <https://arxiv.org/abs/1905.03950>.

Albert Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 2005.

Jan Tinbergen. Shaping the world economy. *The Twentieth Century Fund*, 1962.

Cédric Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2003.

Yoto V. Yotov. *Gravity at Sixty: The Workhorse Model of Trade*. CESifo Working Paper 9584, 2022.

This essay is a theoretical commentary on two papers: Gaskin et al. [2026] and Stuart and Wolfram [2019]. All RSVP field-theoretic constructions are original to the present author and form part of the ongoing RSVP research program documented at github.com/standardgalactic. The author holds no institutional affiliation.