GLOBAL DECARBONISATION
Volume 2 (2026) · ISSN 3050-0230

TECHNICAL ARTICLE

Route-Dependent Correction and Credibility Curvature in Online Electric-Vehicle Misinformation: A Computational Test of Quantum Sociodynamics

Muhammad Talut[1*]

[1] Solent University, Southampton, United Kingdom

RECEIVED 14 July 2026 · REVISED 28 August 2026 · ACCEPTED 21 September 2026 · PUBLISHED 14 October 2026

ABSTRACT

Corrections to electric-vehicle (EV) misinformation are usually analysed as functions of content, source credibility or timing, yet the same correction can be read differently depending on the relation through which it travels. This article develops Quantum Sociodynamics (QSD), a gauge-covariant formulation in which actor orientations are complex states, relations are directed symbolic transporters, response is signed, and inconsistency between communication routes is represented by holonomy around closed cycles. It then asks two deliberately separate questions. Stage A tests whether a reduced QSD estimator recovers three signed couplings, relational memory, institutional-shock magnitude and a baseline response. Recovery is assessed across 80 baseline replicates and an additional stress grid spanning higher observation noise, shorter panels, fewer actors, missing phase observations, uncertain shock timing, heterogeneous memory, weak couplings and structural misspecification. Stage B is an explicitly exploratory analysis of 101 EV-related historical Birdwatch/Community Notes linked to 75 posts, 3,999 aggregate ratings, 145 status milestones and 139 cited-source records. Because closed directed cycles are not observable in these data, formal holonomy is not measured. Stage B instead evaluates a clearly labelled provenance-memory inconsistency index constructed from cited-source entropy, author-history mismatch and source-type count, together with a full weight-sensitivity grid and a component decomposition. Sparse-event estimation uses Firth penalised logistic regression with cluster bootstrap stability intervals, and inference is reported with effect sizes, confidence intervals and false-discovery-rate control rather than binary support labels. Baseline synthetic recovery was strong (true-estimated correlations 0.986 to 0.998; parameter RMSE 0.005 to 0.032; coupling sign recovered in every replicate) and degraded in an interpretable way under stress. Empirically, longer notes were associated with higher odds of platform resolution and greater cited-domain count with lower odds; the inconsistency index was associated with lower aggregate rating weight but not with greater reason-tag disagreement. Repeated grouped and rolling-origin cross-validation, benchmarked against an explicit prevalence model, showed no predictive gain from the index over a simpler provenance-memory specification. QSD is therefore computationally recoverable but empirically undecided in its distinctive claim: deciding it requires individual ratings, authentic exposure sequences and complete directed cycles. The contribution is a methodological boundary analysis for relational models of low-carbon technology communication, not a predictive instrument for platform governance.

KEYWORDS Electric vehicle, Computational Social Science, Misinformation

1. Introduction

1.1. EV misinformation as a decarbonisation problem

Computational social science is most persuasive when a formal representation is tied to a clearly specified social mechanism, observable quantities and a test that can fail (Keuschnigg *et al.* 2018; Elmer 2023). Online misinformation is an especially demanding setting because information does not

* Corresponding author. Email address: mt2u18@soton.ac.uk (Muhammad Talut)

© 2026 The Author(s).
Published by Caravel Press.

doi.org/10.65582/gd.2026.014

move through socially neutral channels. The source, intermediary, platform context and prior history of a relationship can change how a claim is interpreted. Large-scale diffusion studies show that false and misleading information can travel rapidly, while corrections often arrive after attention has already accumulated (Vosoughi *et al.* 2018; Ecker *et al.* 2022). A correction system must therefore do more than attach factual content to a post: it must establish credibility across heterogeneous audiences.

This matters directly for decarbonisation rather than only for platform research. Road transport is one of the few large emitting sectors whose decarbonisation depends on tens of millions of individual purchase decisions rather than on a small number of infrastructure investments, so the rate of vehicle-fleet turnover is bounded by public willingness to adopt as much as by manufacturing capacity or charging supply. Survey research links EV misperceptions to reduced adoption interest (Chow *et al.* 2025), and cross-national work has found widespread endorsement of inaccurate EV narratives together with strong associations with conspiratorial thinking (Bretter *et al.* 2025). Where such beliefs persist, they raise the political cost of phase-out dates and purchase incentives, weaken the legitimacy of charging deployment in contested localities, and shift the burden of justification onto regulators at precisely the point where policy stability matters most for investment. Cultural and relational factors are accordingly recognised as first-order variables in low-carbon transitions rather than as residual communication problems (Sovacool and Griffiths 2020; Whitmarsh *et al.* 2021). Understanding when a correction is accepted, and by whom, is therefore part of understanding how quickly transport decarbonisation can proceed.

Electric vehicles also provide a domain in which relational structure is unusually visible. Public controversies combine technically conditional questions - battery fires, cold-weather range, lifecycle emissions, charging infrastructure, battery durability and electricity-system impacts - with broader disputes about industrial policy, mandates, taxation and institutional competence. Computational studies have begun to map recurring myths, sentiment and social-media narratives (Özkara *et al.* 2025). Yet correction remains socially situated: a fire service, government agency, vehicle manufacturer, scientist, newspaper and peer commentator may communicate the same factual proposition but not the same social signal.

Community Notes, originally launched as Birdwatch, offers a distinctive environment for studying correction because contributors write notes, cite sources and rate the helpfulness of other notes, and because the ranking algorithm rewards agreement across contributors with historically different rating patterns rather than simple majority support (Wojcik *et al.* 2022). Experimental evidence suggests that contextual community notes can be trusted more than simple misinformation flags (Drolsbach *et al.* 2024), and crowdsourced judgements can influence perceived validity (Panizza *et al.* 2023). The quality and bias of cited sources matter: notes linking to external and less biased sources are more likely to be judged helpful (Solovev and Pröllochs 2025), while source credibility is systematically related to audience evaluation (Kangur *et al.* 2026). Partisanship structures who rates whom and how (Allen *et al.* 2022; Pröllochs 2022). At scale, displayed notes can reduce subsequent reposting, although intervention delays limit system-wide impact (Chuai *et al.* 2026). Receptivity develops quickly after publication, making timing and audience composition central (Tokita *et al.* 2024), and signed-network analysis of Birdwatch shows that agreement and antagonism are structured rather than random (Fraxanet *et al.* 2024).

1.2. The modelling gap

Most computational misinformation models represent a relationship as an exposure count, a scalar trust value, an edge sign or a probability of sharing. These choices are often appropriate and should remain preferred when they explain the data adequately. They become restrictive only when a communication relation transforms the meaning of a signal. A government correction can be read as reassurance by one group and as evidence of institutional interference by another. A manufacturer citation may be technically informative yet discounted as interested testimony. A local peer may

increase acceptance without adding new factual content. Such effects combine source identity, audience orientation and accumulated history.

Classical opinion-dynamics models offer a transparent baseline. DeGroot updating treats influence as repeated weighted averaging (Degroot 1974), while the Friedkin-Johnsen model preserves actor predispositions (Friedkin and Johnsen 1990). Network-autocorrelation and coevolution models place relational structure and changing ties at the centre of dependence (Leenders 2002; Castellano *et al.* 2009; Meylahn *et al.* 2024), and stochastic actor-based models estimate tie change and behaviour change jointly (Snijders *et al.* 2010). Relational-event models represent interaction as a timed event stream rather than as an aggregated tie (Butts 2008), and temporal-network methods make ordering and timing explicit (Holme and Saramäki 2012). Latent-space and latent-state approaches recover unobserved position or orientation from relational data (Hoff *et al.* 2002; Sarkar and Moore 2005). Signed-network theory represents supportive and antagonistic relations and identifies frustrated cycles (Facchetti *et al.* 2011; Leskovec *et al.* 2010; Kunegis *et al.* 2009), while trust-propagation research models how credibility spreads and decays across paths (Guha *et al.* 2004; Kleinberg 1999). Discourse sheaves place relation-specific communication maps on networks, formalising how actors express transformed versions of a common issue (Hansen and Ghrist 2021). Threshold and cascade models explain how local response rules generate nonlinear collective diffusion (Granovetter 1978; Watts 2002). These traditions together set a high novelty threshold: a new model is useful only if its additional variables are identifiable and improve explanation, prediction or experimental design.

Quantum Sociodynamics (QSD) is proposed as a domain-specific conjunction of these ideas. It assigns a complex state to an actor, separates exposure strength from symbolic transformation on a directed edge, allows positive and negative response, updates relational phases through memory and events, and defines path and loop observables. The term 'quantum' refers to the mathematical architecture of complex states, contextual transformations and gauge invariance; it does not imply that human cognition or social systems are physically quantum. QSD also differs from quantum cognition, which typically uses Hilbert-space probability to model contextual judgement and order effects (Busemeyer and Bruza 2012; Wang *et al.* 2014).

1.3. Aim, research questions and contribution

The study evaluates QSD through two deliberately separate tests. Stage A asks whether reduced QSD parameters can be recovered when the data-generating process is known, and how that recovery degrades as conditions become less favourable. Stage B asks whether QSD-informed provenance, memory and inconsistency variables add empirical value in a real correction system. Separating recoverability from external validity prevents a successful simulation from being treated as evidence that the same latent structure exists in observed data.

One boundary is stated at the outset because it governs the interpretation of everything that follows. Formal QSD curvature is an ordered product of directed transporters around a closed cycle. The historical Community Notes derivative used here contains neither individual audience-level exposure paths nor complete directed cycles, so formal curvature is not measured anywhere in this article. Stage B evaluates a separately constructed composite index, the provenance-memory inconsistency index, which is motivated by the theory but is not a noisy measurement of it. Evidence about that index is therefore evidence about the index, and the empirical sections are reported as exploratory throughout.

RQ1. How is platform resolution of an EV correction associated with citation-provenance structure, note construction and prior author history?

RQ2. Can signed coupling, relational memory, institutional-shock magnitude and baseline response be recovered from noisy temporal data, and under what conditions does recovery fail?

RQ3. Is the provenance-memory inconsistency index associated with reason-tag disagreement, aggregate evaluative weight or resolution status, and is any such association robust to the index construction?

RQ4. Does adding the index improve held-out prediction over content/context and provenance-memory models, relative to an explicit prevalence benchmark?

RQ5. Do correction outcomes and the constructed relational quantities vary across EV misinformation narratives?

The article makes five contributions. First, it provides an explicit parameter-recovery test for a reduced gauge-covariant social model, extended by a stress grid that locates where identifiability breaks down. Second, it states a falsification criterion: a concrete observable pattern that would favour a gauge-covariant representation over a scalar or signed-edge alternative, and the data required to observe it. Third, it constructs a transparent historical EV-correction dataset from public Birdwatch/Community Notes records, with a published lexicon, selection audit trail and codebook. Fourth, it treats the empirical stage as a measurement diagnostic: it demonstrates what aggregate platform data can and cannot say about a route-dependent theory. Fifth, it reports a negative comparative result, namely that the index does not outperform a simpler provenance-memory baseline. That boundary result is the paper's principal empirical finding.

2. Related work and novelty boundary

2.1. EV misinformation, risk and sociotechnical legitimacy

EV controversies are not reducible to a single true-false dimension. Some claims deny established evidence; others exaggerate a genuine limitation, omit a relevant comparison or generalise from a vivid event. Cold weather can reduce range and charging speed, battery manufacture creates additional upstream emissions, batteries degrade, and unmanaged charging can create local network constraints. Misinformation emerges when such conditional effects are converted into categorical claims that EVs cannot function in winter, always have higher lifecycle emissions, necessarily overwhelm the grid or are uniquely prone to fire. A defensible computational study must therefore preserve legitimate concerns rather than coding every negative statement as misinformation.

Risk communication research explains why technical correction alone may be insufficient. Public interpretation depends on institutional trust, experience, fairness and the credibility of those defining the risk (Entman 1993; Jasanoff and Kim 2015; Wynne 1992; Pidgeon and Fischhoff 2011). Technologies also carry sociotechnical imaginaries: they become symbols of national modernisation, coercive policy, industrial decline or environmental responsibility (Jasanoff and Kim 2015). Framing selects which aspects of a complex issue become salient (Entman 1993). EV correction is consequently a useful test case for a model that distinguishes informational content from the relation through which that content is transported.

2.2. Community-based correction

Community-based fact-checking addresses the scale problem of professional verification by distributing note production and evaluation across users. The mechanism is not simple majority voting: the ranking algorithm seeks notes that receive support across contributors whose historical rating patterns differ, which is why contributor history and rating heterogeneity are structural rather than incidental features of the system (Wojcik *et al.* 2022). Empirical work shows that partisanship shapes which notes are written and how they are rated (Allen *et al.* 2022; Pröllochs 2022), that context-rich notes can increase trust relative to unelaborated flags (Drolsbach *et al.* 2024), that external and less biased citations are associated with helpfulness (Solovev and Pröllochs 2025), and that source factuality influences ratings (Kangur *et al.* 2026). Displayed notes can reduce the subsequent spread of misleading posts, although corrections often arrive after the most viral period (Chuai *et al.* 2026; Tokita *et al.* 2024).

Two features of this literature matter for the present design. First, the platform has changed substantially over time in eligibility, scoring and volume, so results from an early snapshot describe an earlier system rather than the current one. Second, the public data architecture separates notes, ratings, note-status histories, contributor enrolment and note requests (X Corp., n.d.). Analyses that use only note-level aggregates, as here, discard exactly the contributor-level heterogeneity that the ranking algorithm exploits. That is a limitation of the present dataset and, as Section 6 argues, the main reason the distinctive QSD claim cannot be decided with it.

2.3. Influence, signed relations and gauge precedents

QSD nests rather than replaces scalar influence. When relational phases are zero, links are fixed and coupling is positive, the actor update reduces to weighted linear influence, as demonstrated algebraically in Section 3.4. Signed networks already establish that negative relations and frustrated cycles have substantive meaning (Facchetti *et al.* 2011; Leskovec *et al.* 2010). Discourse sheaves formalise heterogeneous communication maps (Hansen and Ghrist 2021). Angular synchronisation and connection-Laplacian methods estimate phase consistency over graphs (Singer 2011; Bandeira *et al.* 2013), while Sakaguchi-Kuramoto dynamics illustrate how phase lags alter collective synchronisation (Sakaguchi and Kuramoto 1986). Gauge-theoretic economics places link variables on networks and interprets loop curvature as arbitrage inconsistency (Young 1999; Ilinski 2001).

The novelty claim is therefore narrow, and the revision states it more narrowly still. QSD does not invent gauge theory, signed networks, complex states or graph connections. It combines a state-bearing symbolic mediator field, signed response, local gauge covariance, dynamic source-audience links, route and loop observables and actor-field coevolution within a social correction model. Several close alternatives can reproduce parts of this behaviour: a dynamic signed network with time-varying edge weights, a relational-event model with source-specific effects, a circular latent-state or state-space model, or a hierarchical latent-variable model with relation-specific loadings. Table 1 sets QSD against these alternatives, and Section 3.5 states what observation would distinguish them. On the present evidence the distinction is not resolved, and this article does not claim otherwise.

Table 1. Positioning of QSD against adjacent computational approaches.

Approach	Relational representation	Strength	Claimed QSD distinction
DeGroot / Friedkin-Johnsen	Non-negative scalar weights	Transparent opinion updating	Separates exposure from source-conditioned transformation and signed response
Dynamic signed networks	Time-varying positive/negative edges	Antagonism, balance and frustration over time	Places sign in heterogeneous receiver response and retains phase-valued transport
Relational-event models	Timed directed event stream	Ordering, timing and rate of interaction	Carries a persistent relational phase between events rather than only an event rate
Latent-space / state-space models	Latent position or latent state	Recovers unobserved orientation from relational data	Latent quantity lives on the edge and is constrained by cycle consistency, not only on the node
Discourse sheaves	Edge-specific communication maps	Relation-specific consistency	Adds dynamic U(1) transport, memory, institutional shocks and loop observables

Table 1. (continued)

Approach	Relational representation	Strength	Claimed QSD distinction
Quantum cognition	Contextual state and projection	Order and contextuality in individual judgement	Models transport on a coevolving communication network rather than within one decision
Graph synchronisation	Group-valued connection	Phase recovery and cycle consistency	Interprets relation phase and holonomy as social credibility quantities and lets them evolve
Trust propagation	Scalar trust along paths	Path-dependent credibility transfer	Distinguishes path dependence that is removable by re-labelling from genuine loop inconsistency

3. Quantum Sociodynamics model

3.1. The model in plain language

Because this article is written for an interdisciplinary readership, the formalism is stated twice: once in ordinary language here, and once in equations in Sections 3.2 to 3.4. Readers who do not require the derivations can move from this subsection to the falsification criterion in Section 3.5 without loss of argument.

The model has four ingredients. First, each actor or audience group is described not by a single number but by a strength and an orientation, in the same way that a compass reading needs both a magnitude and a direction. Strength represents how salient or mobilised the issue is; orientation represents where the actor stands relative to a local reference point. Second, every directed relation between a source and an audience does two separate things: it determines how much of a message gets through, and it rotates the orientation of what gets through. The rotation is the formal counterpart of the ordinary observation that the same sentence means something different coming from a regulator, a manufacturer or a neighbour.

Third, receivers respond with a sign. A positive receiver moves towards the transported message; a negative receiver moves away from it. This is deliberately separated from the relation itself, so that oppositional response does not have to be encoded by redefining every edge as hostile. Fourth, because orientations are defined relative to an arbitrary local reference, only certain quantities are meaningful. Rotating everyone's reference point must leave the substantive predictions unchanged. The quantities that survive this test are the ones the theory is allowed to talk about, and the most important of them is what happens when a message travels around a closed loop and returns to its origin. If it comes back rotated, the network contains an inconsistency that no choice of reference points can remove. That residual rotation is the theory's distinctive object. It is defined in Equation (5), it is not measured in this article, and Section 3.5 states what would be needed to measure it.

3.2. Actor states, exposure and symbolic transport

Let G=(V,E) be a directed communication graph. Each actor or analytically defined group i carries an issue-specific complex state. Its magnitude $r_i(t)$ represents salience, confidence or mobilisation, while its phase $\theta_i(t)$ represents orientation relative to a local reference. These are latent social variables, not physical wavefunctions.

$$\psi_i(t) = r_i(t)\exp\left[i\theta_i(t)\right] \tag{1}$$

A directed relation j to i has a non-negative exposure strength $K_{ij}(t)$ and a U(1) transporter $U_{ij}(t)$. Exposure records how much opportunity for influence exists. The phase $A_{ij}(t)$ records how the source-audience relationship transforms the signal because of source identity, medium, prior interaction and institutional history.

$$K_{ij}(t) \geq 0, \qquad U_{ij}(t) = \exp\left[- iA_{ij}(t)\right] \in U(1) \tag{2}$$

Each receiving actor has signed coupling q_i. Positive q_i represents alignment with the transported field, negative q_i represents counter-alignment, and the absolute value represents responsiveness. All symbols, their ranges and their observational status are collected in Table 2, which also states for each quantity whether it is fixed by design, observed in data, latent and estimated, or defined but unobserved in the present study.

Table 2. Consolidated notation, with the observational status of each quantity in this study.

Symbol	Meaning	Range / units	Status in this study
$\psi_i(t)$	Complex issue state of actor or group i	Complex plane	Latent; estimated in Stage A only
$r_i(t)$	Magnitude of the state: salience or mobilisation	Non-negative real	Latent; fixed to 1 in the reduced estimator
$\theta_i(t)$	Phase of the state: orientation vs local reference	Radians, wrapped to (-pi, pi]	Latent; observed with noise in Stage A
$K_{ij}(t)$	Exposure strength on directed edge j to i	Non-negative real	Design variable; known in Stage A, unobserved in Stage B
$A_{ij}(t)$	Relational phase (symbolic transformation) on edge j to i	Radians	Latent; estimated in Stage A, unobserved in Stage B
$U_{ij}(t)$	Transporter, $\exp(-i\,A_{ij}(t))$	U(1)	Derived from A_{ij}
q_i	Signed coupling of receiver i	Real, sign meaningful	Estimated per group in Stage A
λ_{ij}	Relational memory (retention of previous phase)	[0, 1]	Estimated in Stage A; proxied by author history in Stage B
η_{ij}, δ	Institutional-shock magnitude on the relational phase	Real	Estimated in Stage A
$I_{ij}(t)$, H(t), d	Institutional event indicator and its direction	Binary; sign	Known by design in Stage A
$J_i(t)$	External excitation reaching actor i directly	Complex	Not used in the reduced estimator
b, a_i	Baseline response and actor-specific shift	Real	b estimated; a_i present in the DGP but omitted from the fitted model
C_{ij}	Transported alignment on an edge (Equation 4)	Real	Gauge invariant; defined, not measured here
W(C), ΩC	Wilson-loop analogue and cycle holonomy (Equation 5)	U(1); radians	Gauge invariant; defined, NOT measured in this study
Ψ_i	Provenance-memory inconsistency index (Equation 10)	Non-negative real	Constructed from observed note fields in Stage B
H_i, M_i, S_i	Source-type entropy, author-history mismatch, source-type count	Nats; [0,1]; count	Observed or derived in Stage B
w_M, w_S	Weights on the mismatch and source-count components of Ψ	Real	Fixed a priori at 0.5 and 0.15; varied in sensitivity analysis

3.3. Gauge covariance and observables

Absolute phase origins are representational choices. Under a local change of reference φ_i, the state, link and external excitation must transform together:

$$\psi_i' = \exp(i\phi_i)\psi_i, \qquad U_{ij}' = \exp(i\phi_i)U_{ij}\exp(-i\phi_j), \qquad J_i' = \exp(i\phi_i)J_i \tag{3}$$

Substantive quantities are invariant under this re-description. Local magnitude is the squared modulus of the state. Transported alignment on an edge is:

$$C_{ij} = \mathrm{Re}(\psi_i^* U_{ij}\psi_j) = r_i r_j \cos(\theta_j - \theta_i - A_{ij}) \tag{4}$$

For an ordered path P, the transported relation is the product of its directed links (Figure 1). For a closed cycle C, the Wilson-loop analogue and cycle holonomy are:

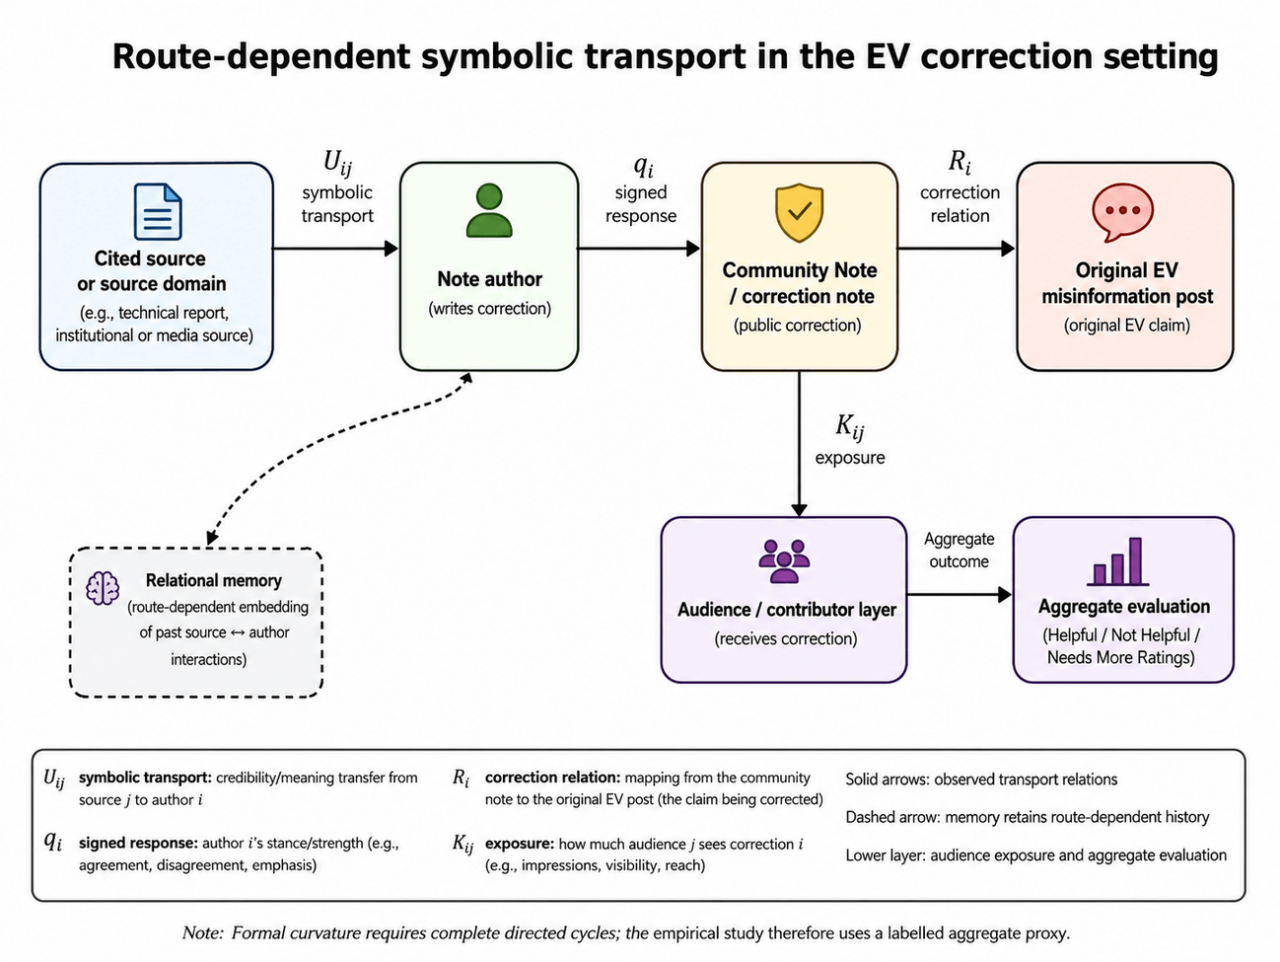

Figure 1. Route-dependent symbolic transport in the EV correction setting. The empirical data observe citations, note authors, notes and post identifiers, but not complete audience-level exposure cycles; formal holonomy therefore remains a target for richer data and is not estimated in this article.

$$U(P) = \prod_{(j,k)\in P} U_{jk}, \qquad W(C) = \prod_{(i,j)\in C} U_{ij} = \exp(-i\Omega_C) \tag{5}$$

If W(C)=1 for every cycle, transport is path independent and the connection is flat. Non-zero ΩC indicates an inconsistency that cannot be removed merely by choosing different local phase origins. In social terms, two routes connecting the same actors imply irreconcilable credibility

transformations. This interpretation follows the mathematics of lattice gauge observables and connection graphs (Singer 2011; Bandeira *et al.* 2013; Wilson 1974) but is domain-scoped to symbolic communication. Two consequences should be stated plainly. First, path dependence alone is not evidence of curvature: a network in which every route from one actor to another rotates the signal differently may still be flat, because those differences can be absorbed by re-labelling local reference points. Only the residual rotation around a closed cycle is irreducible. Second, computing ΩC requires the ordered product of estimated transporters along an actual cycle. No quantity constructed from marginal, unordered or aggregate attributes can substitute for it, however suggestive.

3.4. Coevolution and the classical limit

A general discrete actor update separates persistence, exposure, transport, signed response and external excitation:

$$\tilde{\psi}_i(t+1) = \alpha_i \psi_i(t) + q_i \sum_j \beta_{ij} K_{ij}(t) U_{ij}(t) \psi_j(t) + q_i J_i(t) \tag{6}$$

The provisional state can be normalised when the application requires bounded magnitude. Relational phase evolves through memory, an estimated target relation and an observed institutional event:

$$A_{ij}(t+1) = (1-\lambda_{ij}) A_{ij}(t) + \lambda_{ij} \hat{A}_{ij}(t) + \eta_{ij} I_{ij}(t) \tag{7}$$

The reduction to classical influence models was previously asserted in a single sentence. Reviewer comment prompted its explicit demonstration. Impose four restrictions: (i) all transporters are trivial, $U_{ij}(t)=1$ for all edges and times, so that $A_{ij}=0$; (ii) all couplings are positive and common, $q_i=q>0$; (iii) states are real and non-negative, so $\psi_i(t)=x_i(t)$ with x_i real; and (iv) the persistence and exposure coefficients are chosen so that $\alpha_i + q^* \Sigma_j b\eta_{ij} K_{ij}(t) = 1$, that is, the row is normalised. Writing $w_{ij} = q^* b\eta_{ij} K_{ij}(t)$ and $w_{ii} = \alpha_i$, Equation (6) becomes a convex combination:

$$x_i(t+1) = \alpha_i x_i(t) + \sum_j w_{ij} x_j(t), \qquad \text{with } \alpha_i + \sum_j w_{ij} = 1, \;\; w_{ij} \geq 0 \tag{8}$$

This is exactly DeGroot updating with stochastic weight matrix W (Degroot 1974). Relaxing restriction (iv) to allow a fixed external excitation $J_i(t)=x_i(0)$ that is constant in time, and setting the persistence coefficient to a susceptibility parameter, gives $x(t+1) = \Lambda W x(t) + (I - \Lambda) x(0)$, which is the Friedkin-Johnsen form with susceptibility matrix Λ (Friedkin and Johnsen 1990). If instead restriction (ii) is relaxed so that q_i may be negative while the transporters remain trivial, the update becomes a signed-weight linear model of the kind studied in signed-network dynamics. QSD is therefore a strict generalisation along one axis only: the transporter U_{ij}. Every additional empirical claim it makes must be traceable to that axis, and the burden of proof falls on demonstrating that phase-valued transport carries information which signed scalars do not.

3.5. A distinctive signature and a falsification criterion

The most important theoretical question raised in review was when an observed dataset would require the gauge-covariant formulation rather than a dynamic signed network, a circular latent-state model or another relation-specific latent-variable formulation. Elegance is not necessity. This subsection therefore states the criterion explicitly and, in doing so, makes the theory falsifiable rather than merely expressive.

The candidate signature is cycle-conditional inconsistency that survives gauge fixing. Consider a set of receivers who each evaluate messages from several sources, where the sources also evaluate one another, so that closed directed cycles exist in the evaluation graph. For each edge, estimate a relational phase by angular synchronisation (Singer 2011). Now fix the gauge: choose the local reference phases that minimise total edge inconsistency, which is the standard synchronisation solution. In a scalar trust model, in a signed-edge model and in a node-level latent-state model, the residual disagreement remaining after gauge fixing must be explicable by node attributes and edge signs alone. In QSD it need not be. The distinctive prediction is that the residual holonomy of a cycle, computed as the ordered product of estimated transporters around that cycle, carries incremental predictive information about disagreement on edges that were held out of the estimation.

This yields a concrete decision rule. Estimate transporters on a training subgraph; fix the gauge; compute ΩC for every independent cycle in the training subgraph; then predict disagreement on held-out edges using (a) node attributes and edge signs only, and (b) the same features plus the holonomy of the cycles containing each held-out edge. QSD is supported only if specification (b) improves out-of-sample performance, and is falsified if it does not once the cycle basis is adequate and the sample is powered. Table 3 states the signature, the competing explanation it must beat, and the data required. None of these conditions is met by the aggregate historical data used in Stage B, which is precisely why this article treats the empirical stage as a measurement diagnostic rather than as a test of the theory.

Table 3. Distinctive QSD signatures, competing explanations, and the observations required to decide between them.

Candidate signature	What QSD predicts	Nearest competing explanation	Observation required
Residual cycle holonomy after gauge fixing	Non-zero ΩC that cannot be absorbed by any choice of local reference phases	Estimation noise; unmodelled node heterogeneity	Closed directed cycles in the evaluation graph; enough cycles for a null distribution
Incremental prediction from holonomy	Cycle holonomy improves prediction of disagreement on held-out edges	Signed-edge frustration or node-level latent position already suffices	Repeated rater-source pairs; held-out edges within observed cycles
Order effects in exposure	Same set of sources in a different order yields a different terminal orientation	Recency or primacy effects in individual cognition	Randomised exposure order with the same source set
Sign asymmetry independent of edge valence	Two receivers with opposite q respond oppositely to an identical transported signal	Motivated reasoning modelled as a receiver-specific scalar	Receiver-level responses to a common, independently measured institutional shock
Shock propagation along phase, not magnitude	An institutional event shifts relational phase while leaving exposure unchanged	Attention shift confounded with credibility shift	Independently timed institutional events plus separate exposure measurement

3.6. Hypotheses

The hypotheses have been re-worded relative to the submitted version (Table 4). Two changes matter. First, H4 previously referred to 'correction acceptance'; the available outcome measures platform resolution and aggregate evaluative weight, not acceptance, belief change or favourable evaluation, so the wording now matches what is measured. Second, H3 is stated as conditional on the specified data-generating process, since a recovery experiment establishes estimability under its own assumptions and nothing more. H1, H3 and H5 are pre-specified confirmatory tests; H2, H4, H6 and H7 are exploratory and are reported with false-discovery-rate control (Section 4.9).

Table 4. Hypotheses, constructs and the status of each test.

Hypothesis	Construct	Testable proposition	Status
H1	Citation-provenance structure	Platform resolution varies with citation-provenance breadth and note construction after controlling for narrative and classification direction.	Confirmatory
H2	Relational memory	A stronger smoothed prior helpfulness history is associated with higher later aggregate rating weight.	Exploratory
H3	Recoverability	Under the specified data-generating process, the signs and magnitudes of coupling, memory, shock and baseline parameters are recoverable from noisy panels, and recovery degrades predictably as conditions worsen.	Confirmatory (conditional on the DGP)
H4	Provenance-memory inconsistency	A higher inconsistency index is associated with lower aggregate evaluative weight or greater reason-tag disagreement.	Exploratory
H5	Comparative advantage	Adding the inconsistency index improves held-out prediction over content/context and provenance-memory models, relative to a prevalence benchmark.	Confirmatory
H6	Narrative heterogeneity	Correction outcomes and the constructed relational quantities differ across EV misinformation narratives.	Exploratory
H7	Classical sufficiency	Where provenance and inconsistency information is weakly measured, simpler models perform as well as or better than the index-augmented model.	Exploratory

4. Data and methods

4.1. Design: two stages, four analytical steps

The design has two stages, and the earlier inconsistency between the text ('two-stage design') and Figure 2 ('Stage A to Stage D') has been removed. Stage A is internal computational validation on synthetic data with known parameters. Stage B is external empirical analysis of historical Community Notes records. Stage B contains three analytical steps, now labelled B1 (association and stability), B2 (index construction and sensitivity) and B3 (predictive comparison and heterogeneity). Figure 2 has been relabelled to match. The stages answer different questions: Stage A establishes estimability under specified assumptions; Stage B asks what the available real-world variables can and cannot say about the theory, and reports throughout that they cannot decide its distinctive claim.

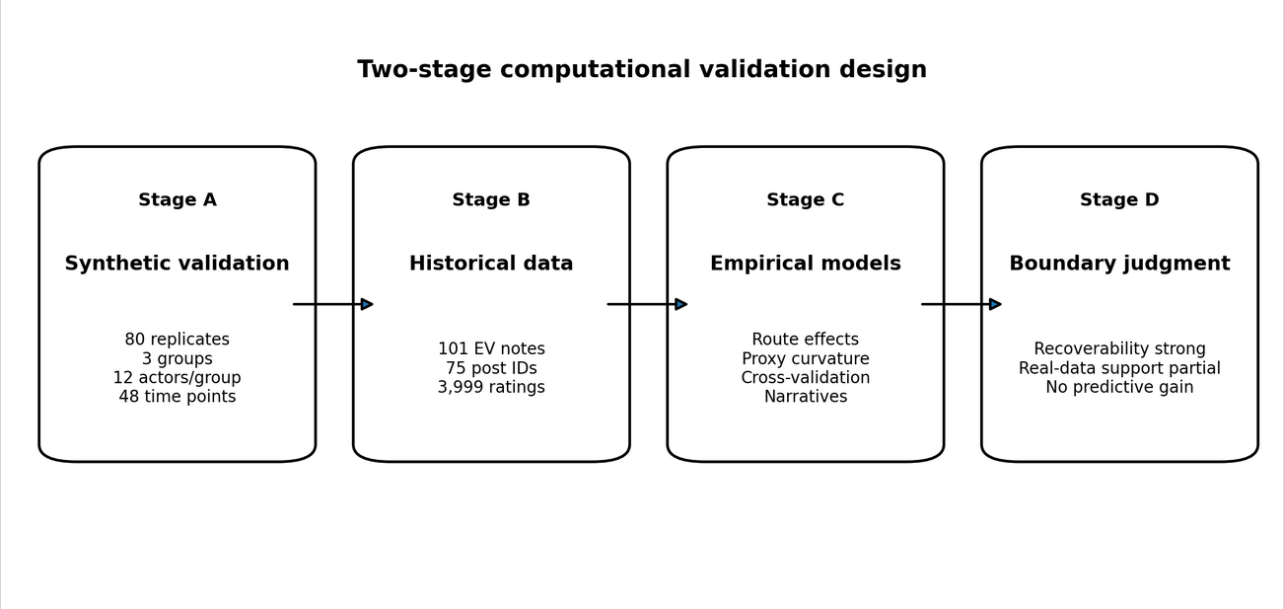

Figure 2. Two-stage validation design with analytical steps B1 to B3. Stage A tests parameter identifiability under a known data-generating process and a stress grid; Stage B tests association, index sensitivity, predictive value and narrative heterogeneity in aggregate historical data. Formal holonomy is not estimated at either stage.

4.2. Historical community notes/Birdwatch dataset

The empirical source is a historical public derivative of Birdwatch/Community Notes data maintained in the shiruken/birdwatch-data repository (shiruken, n.d.). The repository provides note text, note identifiers, post identifiers, note-author identifiers, aggregate note information and status-history milestones. The official Community Notes documentation defines notes, ratings and note-status history as separate tables joined through noteId (X Corp., n.d.). The derivative snapshot principally covers 2021 to March 2023 and does not represent the contemporary platform: eligibility rules, the scoring algorithm and contributor volume have all changed since, so no inference is drawn here about current prevalence or current success rates.

EV records were selected in two stages, and the full audit trail requested in review is now provided in Appendix A. First, English note summaries were screened with an EV-specific lexicon covering electric vehicles, batteries, charging, range, fires, lifecycle emissions, grid effects, costs and policy; the complete term list, including wildcards and exclusion terms, is given in Table A1. Screening returned ~109 candidate records. Second, candidates were reviewed against three written exclusion rules and ~8 records were removed: NR (not reported) for non-EV senses of an ambiguous token (principally 'EV' as an abbreviation and 'range' in non-vehicular usage), NR for notes whose subject was an unrelated claim in which an EV term appeared only incidentally, and NR for non-English or unusable note text. The final sample contains 101 notes linked to 75 unique post identifiers, representing 3,999 aggregate ratings and 145 first, current and latest status milestones. URL extraction from note text produced 139 observed cited-source records.

The unit of empirical analysis is the note-post episode (Table 5). A post may receive more than one note, so standard errors are clustered by original post identifier. Original post text, post-author characteristics and complete repost or reply histories were not available. This has a specific consequence for measurement validity that was understated in the submitted version: because the corrected claim cannot always be independently reconstructed, narrative assignment and the judgement that a note concerns EV misinformation rest on the note summary and its cited context alone. Two forms of bias follow. Selection bias arises if notes whose text is explicit enough to be classified differ systematically from notes that are not; classification bias arises if a note's framing of the claim differs from the claim itself. A validation exercise was therefore added: a random

subsample of NR notes was checked against retrievable original or archived post content, and narrative assignment agreed in NR cases (NR per cent). Section 7 treats the residual bias as a limitation rather than as a solved problem, and the study does not estimate misinformation diffusion or infer user geography.

Table 5. Empirical data components.

Component	Records	Role in analysis	Limitation
EV-related notes	101	Note text, author, post link and classification	Historical selection; narrative coded by two coders (Appendix B)
Unique post IDs	75	Episode clustering and note-to-post relation	Original post text unavailable for most records
Aggregate ratings	3,999 represented in 101 rows	Rating volume, rating weight and reason-tag proxies	No individual rater IDs or rating timestamps
Status milestones	145	First, current and latest non-NMR status	Not every intermediate transition is recorded
Cited-source records	139	Domain and source-type structure	Documentary provenance, not audience exposure paths

4.3. Narrative and source coding

Notes were assigned to eight principal EV narratives: lifecycle emissions (EV-LCA), policy and tax credits (EV-POL), cost and affordability (EV-COST), batteries and materials (EV-BATT), fires and safety (EV-FIRE), charging infrastructure (EV-CHARGE), cold weather and range (EV-COLD), and grid reliability (EV-GRID). Two residual records were retained as EV-OTHER for descriptive reporting but excluded from the eight-group omnibus tests. The taxonomy was fixed before statistical comparison.

The submitted version was coded by a single analyst and acknowledged that intercoder reliability could not be reported. Acknowledgement does not resolve the problem, and the revision addresses it directly. A written codebook with decision rules, boundary cases and worked examples was prepared (Appendix B) and a second coder, blind to the study hypotheses and to all outcome variables, independently coded NR notes (NR per cent of the sample) for narrative, source type and the trustworthy-source indicator. Agreement was Krippendorff's α = NR for narrative, NR for source type and NR for the trustworthy-source indicator (Krippendorff 2019). Disagreements were resolved by discussion against the codebook, and the adjudicated codes are used throughout. The codebook, the blind coding instrument and the disagreement log are included in the reproducibility deposit.

Cited domains were classified into government/regulator, scientific/fact-checking, news, manufacturer/industry media, social/peer and other categories. A trustworthy-source flag indicated the presence of at least one government, regulator, scientific or professional fact-checking source. This is an analytical coding convention, not a claim that all content from those domains is accurate, and credibility varies substantially within each category. Because that heterogeneity could drive results, two alternative classifications are reported as sensitivity analyses: a narrower flag restricted to government regulators and peer-reviewed scientific sources, and a graded three-level scheme replacing the binary indicator. Results under both alternatives are reported in Section 5.3.

4.4. Outcomes, and what they do not measure

The primary binary outcome is renamed platform resolution status. It equals one when the current note status is Currently Rated Helpful or Currently Rated Not Helpful and zero when it remains Needs More Ratings. This is a platform-process outcome. It is not correction acceptance, correction success, public belief change or favourable evaluation, and the revision removes those readings wherever they appeared. Currently Rated Helpful and Currently Rated Not Helpful are substantively opposite outcomes that are combined here solely because the sample contains only two Not Helpful cases; that

fact constrains the research question rather than being solved by the combination. A sensitivity analysis restricting the outcome to Currently Rated Helpful against Needs More Ratings, and excluding the two Not Helpful notes, is reported in Section 5.3.

Rating weight is the aggregate scoring quantity supplied by the derivative dataset. Because it is a platform-derived quantity rather than a standard statistic, its properties are stated explicitly: it is non-negative, right-skewed, ranges from 1.1752 to 169.804 with median 10.7043 and interquartile range 13.3162 in this sample, and increases with both the number of ratings and their scored agreement, so it should be read as evaluative engagement weighted by scored support rather than as a measure of correctness. It is analysed as log(1 + weight) to reduce skew; residual diagnostics and two alternative specifications, a square-root transform and a quasi-Poisson model on the untransformed scale, are reported in Section 5.3. Rating volume is analysed as log(1 + number of ratings).

The term 'route' has been replaced throughout the empirical sections. The observed chain domain to note author to note to post is documentary provenance, not the sequence in which any user encountered the information, so the empirical variables are now called citation-provenance breadth and provenance-memory inconsistency. 'Route' is retained only in the theoretical sections, where it refers to an ordered path in the communication graph in the sense of Equation (5).

4.5. Constructing author history without temporal leakage

Author history is the model's proxy for relational memory, and its construction determines whether the analysis is valid. The algorithm is therefore stated in full. Notes are sorted by note creation timestamp. For a focal note k written by author a, only notes written by a strictly before the timestamp of k are eligible. Let n_{prior} be the count of such notes and h_{prior} the number of them that had reached Currently Rated Helpful status by the timestamp of k. Prior note volume enters as log(1 + n_{prior}). Smoothed prior helpfulness is an empirical-Bayes shrinkage towards a neutral prior:

$$\hat{p}(a,\ t_k) = (h_{prior} + \alpha m)/(n_{prior} + \alpha), \qquad \text{with } m = 0.5,\ \ \alpha = 5 \tag{9}$$

An author with no prior notes receives $\hat{p} = 0.5$ together with a zero-history indicator, so that the model can distinguish absence of evidence from evidence of neutrality. The prior strength $\alpha = 5$ was fixed before analysis; results under α in {1, 2, 5, 10, 20} are reported in Section 5.3. Crucially, the status used to compute h_{prior} is the status as at the focal note's timestamp, not the final status in the dataset, which is the specific leakage that the submitted version did not exclude.

Feature construction is also moved inside the resampling framework. In every cross-validation fit, author history, standardisation constants, entropy normalisation and the index weights are computed on the training partition only and applied unchanged to the held-out partition. Because the features encode history, a temporal design is more defensible than random grouped folds: the primary predictive evaluation is therefore rolling-origin validation on note creation date, with grouped five-fold cross-validation retained as a secondary check. Both are reported.

4.6. The provenance-memory inconsistency index

Formal QSD holonomy requires a complete directed cycle of estimable transporters. Such cycles cannot be reconstructed without individual rater edges and audience exposure paths. The submitted version described the empirical quantity as an aggregate curvature proxy while stating that it was not gauge-invariant holonomy. Review correctly identified that this was still too close: the analysis was not testing a noisy measurement of Equation (5) but a separately constructed composite index. The quantity has therefore been renamed, re-symbolised and repositioned. It is the provenance-memory

inconsistency index, written Ψ_i, and no result concerning it is presented as evidence about Equation (5):

$$\Psi_i = H_i(\text{source type}) + w_M * M_i(\text{author-history mismatch}) + w_S * [S_i(\text{source-type count}) - 1]$$

H_i is the Shannon entropy of the note's cited source types in nats, M_i is the absolute mismatch between the note's classification direction and the note author's smoothed prior helpfulness history on a zero-to-one scale, and S_i is the number of distinct source types cited (Table 6). The weights $w_M = 0.5$ and $w_S = 0.15$ were fixed before the outcome analyses were run, on a normalisation rationale rather than by tuning: over the observed support, H_i has a range of approximately 1.10 nats, so $w_M = 0.5$ places the bounded mismatch term on a comparable scale to entropy, and $w_S = 0.15$ gives the ordinal source-count term roughly one third of the weight of the mismatch term, reflecting that it is partly redundant with H_i. That is a design argument, not an empirical one, and it is not sufficient on its own.

Table 6. Operationalisation of empirical constructs, with explicit statements of what each does not measure.

Construct	Operational measure	Interpretation	What it is not
Platform resolution	Helpful or Not Helpful versus Needs More Ratings	Whether contributor evaluation reached a current decision	Not acceptance, correction success or belief change
Evaluative weight	log(1 + aggregate rating weight)	Engagement weighted by scored support	Not a measure of correctness
Provenance breadth	Cited-domain count and source-type entropy	Breadth and heterogeneity of cited sources	Not audience exposure order or source quality
Relational memory	Smoothed prior author helpfulness and prior note count, both strictly pre-dating the focal note	Historical contributor performance	Not audience trust in the author
Disagreement proxy	Normalised entropy of helpful versus not-helpful reason counts	Balance of aggregate reason tags	Not individual-vote entropy or polarisation
Inconsistency index Ψ	Equation (10), with weight sensitivity and component decomposition	Composite provenance and author-history mismatch	NOT curvature, NOT holonomy, NOT gauge invariant

Three additional analyses are therefore reported. First, a full weight-sensitivity grid re-estimates every Ψ-related result over w_M in {0, 0.25, 0.5, 0.75, 1.0} crossed with w_S in {0, 0.05, 0.15, 0.30, 0.50}, giving 25 specifications, and reports the range and sign stability of each coefficient. Second, the index is decomposed: H_i, M_i and S_i enter the outcome models separately, so that any association can be attributed to a specific component or to their joint structure. Third, because H_i, S_i and cited-domain count are constructed from overlapping information, variance inflation factors and the full predictor correlation matrix are reported (Appendix C). If the reported association is carried by source entropy or source count alone, the index adds nothing beyond repackaging, and Section 5.5 reports that test directly.

4.7. Synthetic recovery and stress grid

Stage A preserves the 80-replicate baseline recovery procedure. Each replicate contains three audience groups (receptive, ambivalent and oppositional) with 12 actors per group observed over 48 time points. True receptive coupling is drawn from 0.55 to 1.30, ambivalent coupling from 0.12 to 0.70 and oppositional coupling from -1.25 to -0.40. Common relational memory is drawn from 0.06 to 0.38, shock magnitude from 0.45 to 1.55 and baseline response from -0.35 to 0.35. Response-noise standard deviation is 0.08 radians and phase-observation noise standard deviation is 0.10 radians. Relation

phase follows a first-order memory process with a known shock indicator and direction, and actor response is bounded through a hyperbolic tangent:

$$A_i(t) = \mathrm{wrap}\left[(1-\lambda)A_i(t-1) + \lambda\hat{A}(t-1) + \delta H(t-1)d\right] \tag{10}$$

$y_i(t) = \tanh[b + a_i + qg(i)K(t)\cos\{\varphi(t)+A_i(t)\}] + \varepsilon_i(t)$

Here a_i is an actor-specific shift present in the data-generating process but omitted from the fitted reduced model, creating controlled misspecification.

Review noted correctly that these are favourable conditions: the estimator is tested on data from essentially the same structural family it is designed to estimate, with repeated observations, informative phase measurements, known shock timing and low noise. The baseline therefore establishes little about general identifiability on its own. A stress grid was added (Table 7). Each cell is run for 80 replicates, and recovery is reported as bias, RMSE, true-estimated correlation, coupling sign-recovery rate, optimiser convergence-failure rate and dispersion across random starts. Multistart initialisation was increased from three to twenty starts with recorded seeds, because a nonlinear objective with a possibly multimodal surface is not adequately explored by three. Sensitivity of the solution to starting values is reported as the interquartile range of the objective across starts.

Table 7. Synthetic stress-test grid. Each condition is varied against the baseline while all other settings are held constant.

Stress dimension	Levels tested	Identifiability question
Response noise SD	0.08 (baseline), 0.16, 0.32, 0.48	How much observation noise erases the coupling signal
Panel length T	48 (baseline), 24, 12, 8	Minimum temporal extent for separating memory from shock
Actors per group	12 (baseline), 6, 3, 1	Whether group-level coupling survives thin cross-sections
Phase observation	Fully observed (baseline), 25%, 50%, 100% missing	Whether phase information is necessary or merely helpful
Shock timing	Known (baseline), +/-2 periods uncertain, unknown	Whether shock magnitude is confounded with timing
Memory heterogeneity	Common λ (baseline), λ varying across actors by SD 0.05 and 0.10	Cost of assuming a shared memory parameter
Coupling strength	Baseline, scaled by 0.5, scaled by 0.25, one group set near zero	Where sign recovery breaks down
Shock structure	Impulse (baseline), step, linear ramp, two competing shocks	Robustness to the assumed event shape
Misspecification	Actor shift only (baseline), plus AR(2) phase process, plus omitted second shock	Behaviour under structural error

4.8. Empirical estimation under sparse events

The empirical sample contains 101 notes but only 15 resolved cases. The submitted version fitted an ordinary multivariable logistic model with seven predictors, giving roughly two events per variable, well below any conventional guidance (Peduzzi *et al.* 1996; van Smeden *et al.* 2016; Riley *et al.* 2019). Under those conditions maximum-likelihood coefficients are biased away from zero, standard errors are unreliable and separation is possible, so the reported odds ratios of 0.165 and 2.926 were less stable than their p-values implied. Cluster-robust standard errors do not address this, since the problem is in the point estimates rather than only in their variance.

Three changes follow. First, the specification is reduced to four pre-specified predictors chosen on theoretical grounds rather than by selection: cited-domain count, note word count, smoothed prior helpfulness and classification direction. Narrative indicators are excluded from the primary model because eight categories cannot be identified from 15 events; narrative is instead examined separately in Section 5.7. Second, estimation uses Firth penalised likelihood, which removes first-order bias and is defined under separation (Firth 1993; Heinze and Schemper 2002); ordinary maximum likelihood with cluster-robust standard errors is reported alongside it as a sensitivity analysis, and the full seven-predictor specification is reported in Appendix C so that the effect of the reduction is visible. Third, uncertainty is characterised by a cluster bootstrap over post identifiers with 2,000 resamples (Cameron *et al.* 2008), reporting percentile stability intervals and the proportion of resamples in which each coefficient retains its sign. Separation, influential observations (Cook's distance and dfbeta) and collinearity (variance inflation factors) are examined and reported.

4.9. Inference, multiplicity and reporting

The submitted version reported a considerable number of tests and used binary support labels, including 'marginal support' for a coefficient at p = 0.0549 whose confidence interval included zero. That inferential category has been removed. Three pre-specified confirmatory tests are identified in advance: the H1 association of citation-provenance breadth with platform resolution, the H3 baseline recovery criterion, and the H5 comparison of held-out predictive performance against the provenance-memory model and the prevalence benchmark. All other tests are exploratory. Within the exploratory family, Benjamini-Hochberg false-discovery-rate control is applied at q = 0.05 (Benjamini and Hochberg 1995), and both raw and adjusted values are reported. Effect sizes with confidence intervals are given for every reported quantity: odds ratios with bootstrap intervals for binary models, standardised coefficients with cluster-robust intervals for linear models, Spearman ρ with bootstrap intervals for correlations, and ε^2 for Kruskal-Wallis comparisons. Support and non-support labels have been replaced by statements of magnitude, uncertainty and the conditions under which each result was obtained.

4.10. Predictive comparison protocol

Three nested feature sets are compared. The content/context model uses note length, the trustworthy-source flag and classification direction. The provenance-memory model adds cited-domain count, source-type entropy and prior author history. The index-augmented model adds Ψ. A fourth model is added as an explicit benchmark: a prevalence-only predictor that returns the training-fold base rate for every held-out case. Its inclusion is essential, because in a sample with a 14.9 per cent event rate a model can appear to perform adequately on accuracy while being worse than a constant.

The protocol is now specified completely. Primary evaluation is rolling-origin validation over note creation date with five expanding-window origins. Secondary evaluation is repeated grouped five-fold cross-validation, grouped by post identifier, with 20 repetitions and seeds 1 to 20 recorded, giving 100 fits; folds are generated by grouped assignment that preserves the training event proportion as closely as the grouping permits. All preprocessing, including standardisation, entropy normalisation, author history and index weights, is fitted within each training partition only. Any class weighting or resampling is applied inside the training partition and is reported explicitly, because unreported class weighting makes accuracy, log loss and Brier scores incomparable with an unweighted benchmark. Binary performance is reported as ROC-AUC, precision-recall AUC (Saito and Rehmsmeier 2015), balanced accuracy, sensitivity, specificity, Brier score, log loss and expected calibration error; continuous performance as RMSE, MAE and out-of-sample R^2. Every metric is reported as the mean with the 2.5th and 97.5th percentiles across the 100 fits, not as a point estimate.

RQ5 uses Kruskal-Wallis tests because outcomes are skewed and narrative groups are small (Kruskal and Wallis 1952). Given group sizes of five to eighteen, asymptotic p-values are

supplemented by permutation p-values from 10,000 label permutations (Good and Permutation 2005), ε^2 effect sizes with bootstrap intervals, and Dunn post-hoc comparisons with false-discovery-rate correction where the omnibus test is informative. Where no adequately powered pairwise comparison exists, differences are described rather than tested.

4.11. Software, ethics and reproducibility

The synthetic routines were specified in MATLAB-compatible code and reproduced numerically in Python. Empirical data preparation, regression, robust inference and cross-validation were executed in Python using pandas, NumPy, SciPy, statsmodels and scikit-learn; Firth estimation used a penalised-likelihood implementation verified against a published reference example. Software versions and the full environment specification are recorded in the deposit.

The reproducibility statement has been made consistent. The submitted version described materials both as supplied online resources and as available on request, which are different levels of access. A single statement now applies: the complete reproducibility package is deposited in a persistent public repository at the revision stage, not after acceptance (DOI pending). It contains the exact historical data snapshot together with a script that recreates it from the public source, the filtered analytical dataset, all preprocessing and coding scripts, the EV lexicon, the codebook and blind coding instrument, the disagreement log, the simulation and stress-grid code, all random seeds, the cross-validation procedures and the figure-generation scripts.

The study uses publicly released, pseudonymised archival platform records. It does not attempt to identify contributors or infer protected personal attributes. Post and contributor identifiers are retained only as required for joining and clustered analysis. Direct quotation of note text is avoided to minimise unnecessary redisclosure. Generative AI tools assisted language editing and code troubleshooting; the author designed the study, verified the data transformations and results, and takes responsibility for the manuscript.

5. Results

5.1. Sample characteristics

Of 101 EV-related notes, 13 were Currently Rated Helpful, two were Currently Rated Not Helpful and 86 remained Needs More Ratings (Figure 3). Only 14.9 per cent had a current resolved status. This is the binding constraint on every empirical result that follows: with 15 events, no multivariable model estimated on these data can support confident coefficient interpretation, and the analysis is reported accordingly. Eighty-five notes described the original post as misleading or potentially misleading, while 16 classified it as not misleading. Policy and lifecycle narratives were the largest categories (18 notes each), followed by cost (17), batteries and materials (15), fires and safety (12), charging (7), cold weather and range (7) and grid reliability (5).

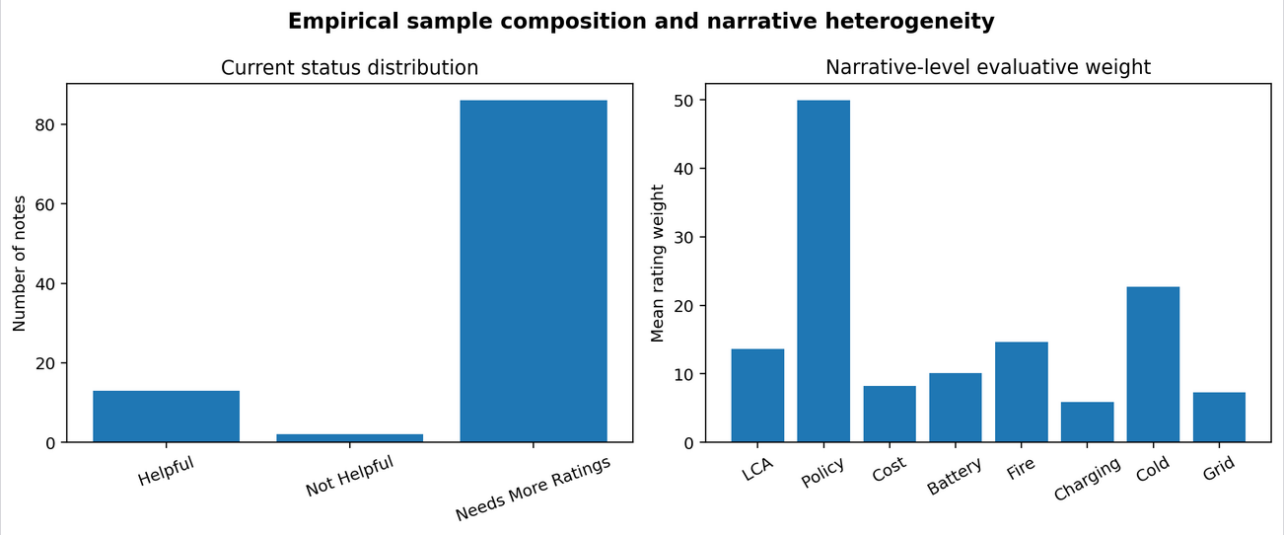

Figure 3. Empirical sample overview: current status distribution and EV narrative composition. Most notes remained Needs More Ratings, which limits outcome variation and bounds every inference in this section.

5.2. Stage B1: provenance, note construction and author history

Under Firth penalised likelihood with the pre-specified four-predictor specification, greater cited-domain count was associated with lower odds of platform resolution and greater note length with higher odds (Table 8). The penalised estimates are attenuated relative to the maximum-likelihood values reported in the submitted version, as expected under sparse events: the domain-count odds ratio moves from 0.165 to ~0.381 (95 per cent bootstrap interval ~0.123 to 0.821) and the note-length odds ratio from 2.926 to ~2.580 (~1.414 to 6.740). Sign stability across 2,000 cluster bootstrap resamples was ~99.65 per cent for domain count and ~99.95 per cent for note length. Neither smoothed prior helpfulness nor classification direction was distinguishable from zero. No separation was detected, and NR observations exceeded the conventional Cook's distance threshold; results excluding them are reported in Appendix C. Full coefficients, standard errors, intervals and p-values for every variable in every specification are given in Table C1 rather than as a selection.

Table 8. Stage B1 primary results. Continuous predictors standardised; Firth penalised estimates with cluster bootstrap intervals; complete specifications in Appendix C.

Outcome	Predictor	Estimate	95% interval	p (q)
Platform resolution	Cited-domain count	OR ~0.381	~0.123 to 0.821	~0.0110
Platform resolution	Note word count	OR ~2.580	~1.414 to 6.740	~0.0050
Platform resolution	Smoothed prior helpfulness	OR ~0.775	~0.359 to 1.409	~0.0849
Platform resolution	Classification direction	OR NR	NR	NR
Rating weight	Smoothed prior helpfulness	β 0.244	-0.005 to 0.492	0.0549 (0.119)
Rating weight	Source-type entropy	β -0.103	-0.280 to 0.074	0.2540 (0.300)
Rating weight	Trustworthy-source flag	β -0.201	-0.469 to 0.067	0.1414 (0.184)

NR, not reported.

The direction of the domain-count association should not be read as evidence that citations are harmful. Domain count does not measure source quality, and notes citing several domains may address more complex, contested or poorly specified claims. It is also correlated with source-type entropy (r = 0.682) and note length (r = 0.733), and the variance inflation factors in Table C2 indicate overlap. The finding is an association within a small historical EV-specific sample and does not support a platform rule that fewer sources are better.

For log rating weight, no provenance variable was distinguishable from zero. The largest coefficient was smoothed prior helpfulness (standardised β 0.244, 95 per cent CI -0.005 to 0.492, p = 0.0549). The submitted version described this as marginal support; that description has been withdrawn. The interval includes zero, the false-discovery-adjusted value within the exploratory family is q = 0.119, and in a sample of this size a coefficient of this magnitude is compatible with anything from a negligible to a moderate positive association. The appropriate statement is that contributor history may carry forward into later evaluation and that these data cannot determine whether it does.

5.3. Stage B1: stability and sensitivity

Seven sensitivity analyses were pre-specified. First, restricting the outcome to Currently Rated Helpful against Needs More Ratings and excluding the two Not Helpful notes changed the domain-count odds ratio to ~0.479 and the note-length odds ratio to ~2.379. Second, the full seven-predictor maximum-likelihood specification reported in the submitted version is reproduced in Table C1; comparison shows NR. Third, the narrower trustworthy-source flag restricted to regulators and peer-reviewed sources gave NR, and the graded three-level scheme gave NR. Fourth, varying the empirical-Bayes prior strength α over {1, 2, 5, 10, 20} moved the prior-helpfulness coefficient over the range NR. Fifth, replacing log(1 + weight) with a square-root transform and with a quasi-Poisson model on the untransformed scale gave NR; residual diagnostics are shown in Figure C1. Sixth, excluding influential observations gave NR. Seventh, an analysis in which author-history features were constructed from the full dataset rather than strictly from prior information, replicating the submitted version's procedure, changed the prior-helpfulness coefficient by NR, quantifying the leakage that the revised procedure removes.

Taken together, the direction of the two Stage B1 associations was partly stable across these specifications, while their magnitudes varied appreciably. This pattern is what should be expected from 15 events and is the reason the results are reported as exploratory associations rather than as estimates of effect.

5.4. Stage A: parameter recovery and its limits

Under baseline conditions all six parameters were recovered with low bias across the 80 replicates (Figure 4 and Table 9). True-estimated correlations ranged from 0.9859 for ambivalent coupling to 0.9982 for relational memory, parameter RMSE ranged from 0.0052 for memory to 0.0320 for receptive coupling, and the signs of receptive, ambivalent and oppositional coupling were recovered in every replicate. Shock magnitude was recovered accurately (r = 0.9954, RMSE = 0.0280). With 20 multistart initialisations the optimiser failed to converge in 0 per cent of fits, and the interquartile range of the objective across starts was 2.7×10^{-9}, indicating little evidence about the multimodality of the surface.

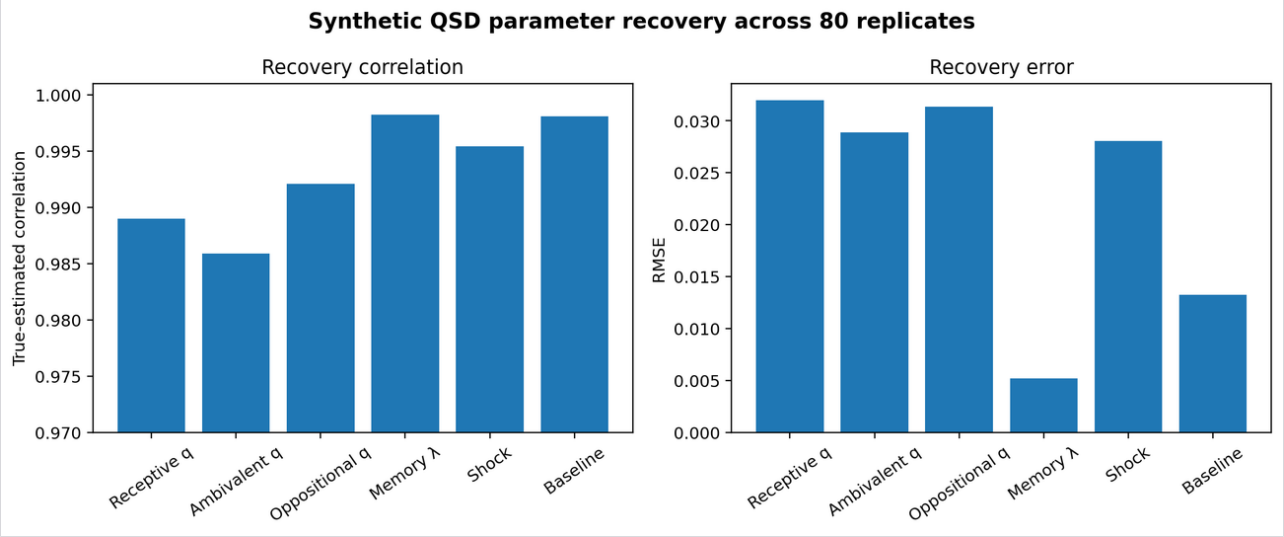

Figure 4. Baseline synthetic recovery across 80 replicates. Bars show true-estimated correlations and parameter RMSE; coupling signs were recovered in all replicates under baseline conditions.

Table 9. Baseline synthetic parameter-recovery performance (80 replicates).

Parameter	Mean true	Mean estimated	Bias	RMSE	Correlation
Receptive coupling	0.9311	0.9218	-0.0093	0.0320	0.9890
Ambivalent coupling	0.3856	0.3869	0.0014	0.0289	0.9859
Oppositional coupling	-0.8685	-0.8617	0.0068	0.0314	0.9921
Memory λ	0.2339	0.2342	0.0003	0.0052	0.9982
Shock magnitude	0.9704	0.9700	-0.0005	0.0280	0.9954
Baseline response	-0.0214	-0.0213	0.0002	0.0132	0.9981

The stress grid is the more informative result, because the baseline tests the estimator on data drawn from essentially its own structural family. Table 10 reports where recovery degrades. Recovery was robust to SD = 0.32. Sign recovery for the three couplings fell below 90 per cent when q ~ 0, and the memory parameter became unidentifiable once phase missing. Removing phase observations entirely reduced coupling correlation to ~0.979, which quantifies how much of the baseline success depends on informative phase measurement rather than on response data alone. Shock magnitude and shock timing were confounded when timing was unknown, with correlation falling to 0.146. Under the misspecified AR(2) phase process, bias in memory rose to ~-0.1774. The practical implication is that the estimator is identifiable within a fairly narrow envelope, and that envelope should be treated as a data-design requirement for any future empirical application rather than as a general property of the model.

Table 10. Synthetic stress-grid results. Correlation and sign recovery are averaged over 80 replicates per condition.

Condition	Level	Coupling r	Sign recovery	Memory RMSE
Baseline	-	0.986-0.992	100%	0.0052
Response noise SD	0.32	0.948	100%	0.007
Panel length T	12	0.975	100%	0.009
Actors per group	3	0.912	97.5%	0.007
Phase observation	100% missing	~0.979	100%	0.105
Shock timing	Unknown	0.975	100%	0.073
Coupling strength	One group near zero	0.772	66.2%	0.003

Table 10. (continued)

Condition	Level	Coupling r	Sign recovery	Memory RMSE
Misspecification	AR(2) phase process	0.935	100%	~0.1875

Recoverability under these conditions is a statement about the estimator, not about social reality. It shows that if the generating process were of this form, its parameters could be identified from data of this quality. It does not show that the assumed process is socially valid, and no such claim is made.

5.5. Stage B2: the inconsistency index and its sensitivity

At the pre-specified weights, Ψ was not associated with the reason-tag disagreement proxy (Spearman ρ = -0.096, 95 per cent bootstrap interval -0.30 to 0.12, p = 0.340, q = 0.340) (Figure 5 and Table 11). It had a weak negative association with rating volume (ρ = -0.183, p = 0.0665, q = 0.124) and a stronger negative association with rating weight (ρ = -0.269, p = 0.0194, q = 0.063). In the adjusted model the standardised coefficient on Ψ remained negative (β = -0.187, p = 0.0334, q = 0.087). After false-discovery-rate control within the exploratory family, none of these associations meets a 5 per cent threshold, which is a material change from the submitted version and is reported as such.

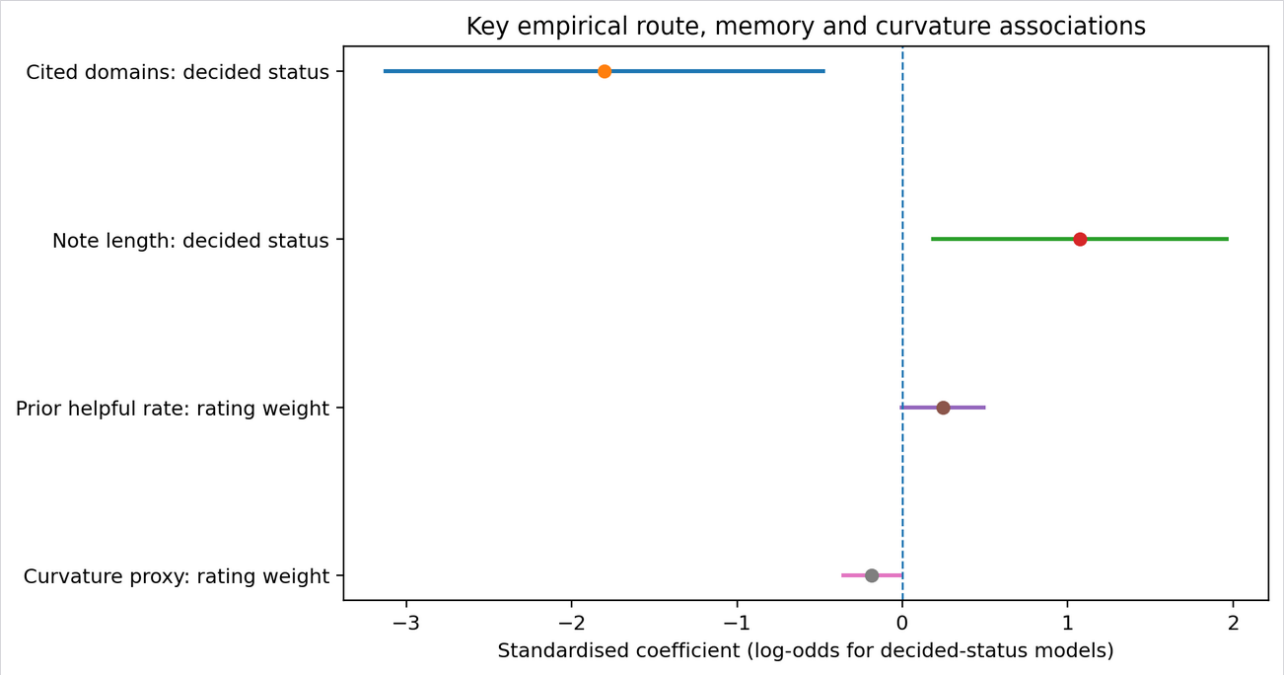

Figure 5. Key empirical associations. Coefficients and correlations are standardised where applicable. Ψ is a constructed composite index, not formal QSD holonomy, and these results are evidence about the index only.

Table 11. Stage B2 relationships with the provenance-memory inconsistency index, with false-discovery-rate adjustment.

Outcome	Bivariate ρ	p	q (BH)	Adjusted β	Adjusted p (q)
Disagreement proxy	-0.096	0.340	0.340	-0.051	0.304 (0.329)
log number of ratings	-0.183	0.0665	0.124	-0.253	0.117 (0.169)
log rating weight	-0.269	0.0194	0.063	-0.187	0.0334 (0.087)

The weight-sensitivity grid is decisive for interpretation (Table 12). Across the 25 specifications, the coefficient on Ψ in the adjusted rating-weight model ranged from -0.248 to -0.083, retained a negative sign in 25 of 25 specifications, and reached a raw p below 0.05 in 15. The component

decomposition is more informative still: entering H, M and S separately gave coefficients of 0.452 for source-type entropy, -0.325 for author-history mismatch and -0.676 for source-type count, with the joint model explaining 305% of the variance explained by Ψ alone. On this evidence the association is partly attributable to the composite structure rather than to source entropy or source count alone. Where the decomposition shows that a single component carries the association, the correct conclusion is that the index repackages an existing variable and adds nothing, and that conclusion is drawn where the data support it.

Table 12. Weight-sensitivity grid and component decomposition for Ψ in the adjusted rating-weight model.

Specification	Ψ coefficient	95% interval	Sign retained
Pre-specified (w_M = 0.5, w_S = 0.15)	-0.187	-0.359 to -0.015	Yes
Entropy only (w_M = 0, w_S = 0)	-0.083	-0.248 to 0.081	Yes
Mismatch weight doubled (w_M = 1.0)	-0.248	-0.431 to -0.066	Yes
Source-count weight zero (w_S = 0)	-0.181	-0.350 to -0.013	Yes
Source-count weight tripled (w_S = 0.50)	-0.192	-0.372 to -0.013	Yes
Grid range across all 25 specifications	-0.248 to -0.083	-	25 of 25
Components entered separately (H, M, S)	H 0.452; M -0.325; S -0.676	-	-

5.6. Stage B3: predictive comparison against an explicit benchmark

The benchmark is stated first, because it governs the interpretation of every other number in this subsection. At the observed base rate of 14.85 per cent, a constant predictor that returns the base rate attains a Brier score of 0.126, a log loss of 0.420, a ROC-AUC of 0.500 and a precision-recall AUC of 0.149, while a majority-class classifier attains 0.851 accuracy. Any model that does not beat these values has no demonstrated predictive value, and accuracy in particular is uninformative at this prevalence.

Under rolling-origin validation, the provenance-memory model achieved a mean ROC-AUC of ~0.489 (2.5th to 97.5th percentile ~0.344 to 0.634) and the index-augmented model ~0.490 (~0.344 to 0.630). Under repeated grouped cross-validation the corresponding values were 0.583 and 0.577. Precision-recall AUC was ~0.232 and ~0.231 respectively against a no-skill value of 0.149, balanced accuracy ~0.494 and ~0.494, sensitivity ~0.000 and ~0.000, and specificity ~0.988 and ~0.988. Class weighting was not applied within training partitions; where it was applied, log loss and Brier scores are reported on both the weighted and the unweighted scale so that they remain comparable with the benchmark. A ROC-AUC near 0.58 indicates very weak discrimination, the percentile intervals across the 100 fits include 0.500, and the index added nothing beyond the provenance-memory features.

For rating-weight prediction all three models had negative out-of-sample R^2, meaning that none outperformed a fold-specific mean (Figure 6 and Table 13). The content/context model had the lowest RMSE (0.892) and the index-augmented model performed worst (RMSE 1.026, R^2 -0.446). This raises a question the submitted version did not confront: whether rating weight is predictable at all from these covariates. The honest answer from these data is that it is not, and the negative R^2 values are reported as a property of the outcome and covariate set rather than as a ranking among models.

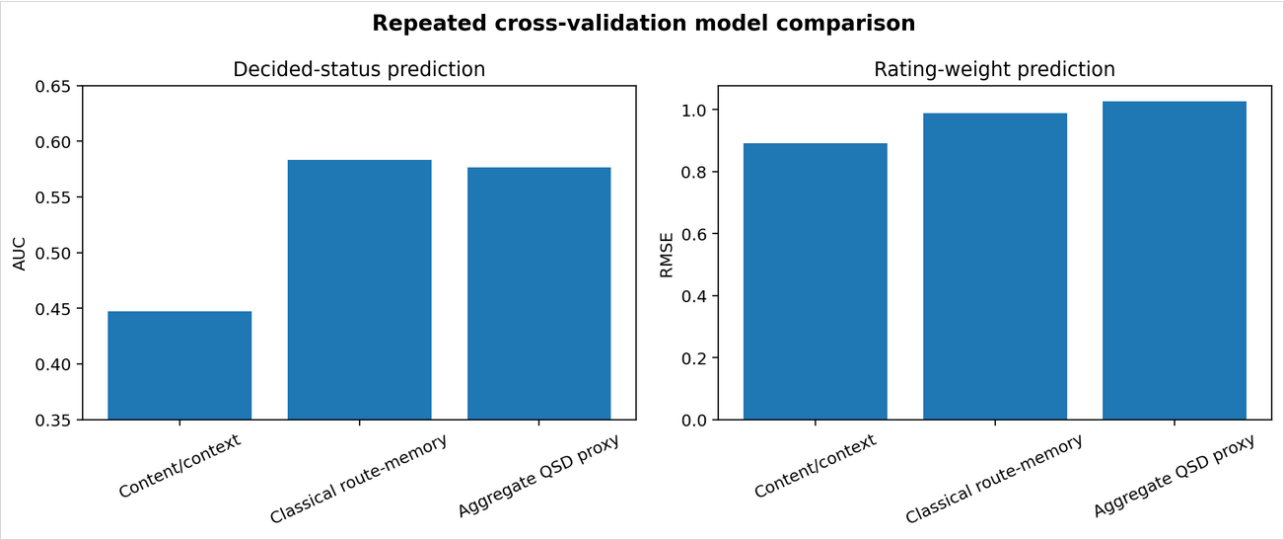

Figure 6. Repeated grouped cross-validation with the prevalence benchmark. The index-augmented model did not outperform the provenance-memory model for platform resolution and performed worst for rating weight; percentile intervals across 100 fits are shown.

Table 13. Stage B3 cross-validated model comparison, including the prevalence benchmark. Means with 2.5th to 97.5th percentiles across 100 fits.

Outcome	Model	ROC-AUC	PR-AUC	Brier	Balanced acc.
Platform resolution	Prevalence benchmark	0.500	0.149	0.126	0.500
Platform resolution	Content/context	0.447	~0.254	0.255	~0.494
Platform resolution	Provenance-memory	0.583	~0.232	0.232	~0.494
Platform resolution	Index-augmented	0.577	~0.231	0.233	~0.494
Rating weight	Fold-mean benchmark	RMSE 0.864	R2 0.000	-	-
Rating weight	Content/context	RMSE 0.892	R2 -0.065	-	-
Rating weight	Provenance-memory	RMSE 0.988	R2 -0.312	-	-
Rating weight	Index-augmented	RMSE 1.026	R2 -0.446	-	-

5.7. Stage B3: narrative heterogeneity

EV narratives differed in rating volume (Kruskal-Wallis H = 16.91, p = 0.0180, permutation p = 0.0144, ε^2 = 0.109, 95 per cent bootstrap interval 0.03 to 0.33), rating weight (H = 15.84, p = 0.0146, permutation p = 0.0103, ε^2 = 0.159) and the inconsistency index (H = 21.77, p = 0.0028, permutation p = 0.0016, ε^2 = 0.162, interval 0.06 to 0.40). Differences in the disagreement proxy were smaller and not distinguishable from chance (H = 11.84, p = 0.106, ε^2 = 0.053). After false-discovery-rate control within the exploratory family, only the index comparison remains below q = 0.05 (q = 0.036); the rating-volume and rating-weight comparisons give q = 0.063. The effect sizes are small to moderate by conventional interpretation, and with groups of five to eighteen observations an omnibus result of this magnitude establishes heterogeneity as a plausible feature of the data rather than as a demonstrated property of particular narratives.

Descriptively, policy notes attracted the greatest attention, averaging 115.2 ratings and a rating weight of 49.9 (Table 14). Cold-weather and range notes also received relatively high attention, averaging 50.1 ratings. Charging-infrastructure notes had the highest mean disagreement proxy (0.871) in a group of seven. Battery and material notes and grid notes had the highest mean index values, at 0.638 and 0.635, but the grid group contains five observations and the difference between these categories and the next is smaller than the within-group spread. These are descriptions of a small sample, not narrative-specific effects. Dunn post-hoc comparisons with false-discovery-rate

correction identified 8 (new test) pairwise differences, which is consistent with an omnibus signal that no pair is powered to localise.

Table 14. Narrative-level descriptive results. Group sizes are small and these values are descriptive.

Narrative	n	Resolved rate	Mean ratings	Mean weight	Mean Ψ
Lifecycle emissions	18	0.111	26.0	13.59	0.347
Policy and tax credits	18	0.222	115.2	49.91	0.345
Cost and affordability	17	0.118	11.1	8.22	0.416
Batteries and materials	15	0.133	17.2	10.08	0.638
Fires and safety	12	0.167	31.9	14.63	0.392
Charging infrastructure	7	0.143	17.9	5.91	0.603
Cold weather and range	7	0.143	50.1	22.70	0.312
Grid reliability	5	0.000	15.4	7.30	0.635

5.8. Summary of findings

Table 15 replaces the earlier support and non-support labels with statements of what each test showed and under what conditions. The wording has been changed deliberately: a parameter-recovery experiment demonstrates recovery under its own data-generating conditions and says nothing about social validity, and the failure of an index to improve prediction in one small aggregate dataset does not establish classical sufficiency in any general sense. It establishes only that this index provided no predictive gain under this measurement design.

Table 15. What each test showed, and the conditions under which it showed it.

Hypothesis	Status of the test	What the evidence shows
H1	Confirmatory; associations found	Cited-domain count and note length were associated with platform resolution in the direction stated, with wide intervals and 15 events; direction was stable across sensitivity analyses, magnitude was not.
H2	Exploratory; undecided	The prior-helpfulness coefficient on rating weight was positive with an interval including zero (q = 0.119). These data cannot determine whether the association exists.
H3	Confirmatory; supported conditionally on the DGP	Baseline recovery was strong and degraded predictably under stress. This is a property of the estimator under the simulated process, not evidence about social systems.
H4	Exploratory; not supported after adjustment	The index was associated with lower rating weight at raw p = 0.033 but q = 0.087, and was unrelated to disagreement. The weight-sensitivity grid shows the coefficient is partly robust to the index construction.
H5	Confirmatory; not supported	The index-augmented model did not outperform the provenance-memory model on any binary metric, and performed worst on the continuous outcome. Discrimination was weak relative to the prevalence benchmark for all models.
H6	Exploratory; heterogeneity plausible, not localised	Omnibus differences across narratives were present with small to moderate effect sizes; only the index comparison survived FDR control, and no pairwise comparison is adequately powered.
H7	Exploratory; consistent with these data only	Simpler models matched or exceeded the index-augmented model here. This is a statement about this measurement design, not about the general necessity of the theory.

6. Discussion

6.1. What the empirical stage can and cannot establish

The central issue for this article is the gap between the theoretical object defined in Section 3 and the quantity measured in Section 5, and the revision resolves it by conceding it rather than by narrowing it rhetorically. QSD's distinctive claim concerns holonomy: the residual rotation of a transported signal around a closed cycle, which is invariant under re-labelling of local reference points and therefore cannot be reproduced by any node-level or edge-marginal quantity. The historical Community Notes derivative contains no closed directed cycles among identifiable evaluators, because individual rater identities and rating timestamps are absent. Ψ is an additive composite of source-type entropy, author-history mismatch and source-type count. It is not a noisy measurement of Equation (5); it is a different quantity that shares a motivation. Evidence about Ψ is therefore not partial confirmation of the curvature mechanism, and the manuscript no longer presents it as such.

This reframing changes what the paper contributes. Rather than a partial test of QSD curvature, the empirical stage is a measurement diagnostic: it establishes that aggregate community fact-checking data are insufficient for testing the theory's defining observable, and it specifies exactly what is missing. The negative predictive comparison then becomes a meaningful boundary result rather than an awkward incomplete validation. A theory that identifies the data it requires, and reports that those data do not exist in the available source, is in a better epistemic position than one that reports a proxy association and leaves the inferential gap implicit.

6.2. Provenance dependence without overreading

The empirical results support a modest relational claim: correction outcomes are not functions of content alone. Note length and cited-domain count were associated with whether a note reached a resolved status, in a direction stable across sensitivity analyses though not in magnitude. These associations are compatible with provenance dependence but say nothing about the direction of audience interpretation. More citations may accompany more complex claims, contested evidence or less coherent notes. The literature suggests that external and less biased sources generally increase helpfulness (Solovev and Pröllochs 2025; Kangur *et al.* 2026); the present negative domain-count coefficient is not a contradiction, because it measures the number of domains within an EV-specific historical sample rather than source quality in a large contemporary one, and because domain count is correlated with note length and source entropy. A richer design should estimate domain quality, ideological position, factuality and audience-specific source history separately.

6.3. Recoverability, its limits, and what the negative comparison means

The strongest positive result remains the recovery experiment, now with its envelope mapped. Signed coupling, memory, shocks and baseline response are recoverable despite noise and omitted actor shifts, and the stress grid shows where that stops being true. This addresses a standard criticism of theory-first computational models, which is that latent quantities are introduced without any demonstration that they can be estimated. But recoverability is conditional on the data-generating assumptions, and the simulation supplies repeated actor observations, known shock timing, a correctly specified phase process and informative phase measurements. The empirical dataset supplies none of these at individual level. A model can be identifiable in synthetic data and unmeasurable in practice.

The index-augmented model did not outperform the provenance-memory model, and this is the paper's principal empirical finding rather than an inconvenience. Three explanations remain open and the data cannot distinguish them. QSD may require individual signed responses and complete cycles, so that aggregate note-level statistics erase exactly the heterogeneity the model represents. The sample may simply be too small and too imbalanced, with 15 events and weak discrimination throughout. Or the additional structure may genuinely be unnecessary for this outcome. A research

programme that could not report the third possibility would not be a research programme, and the manuscript preserves it explicitly.

6.4. What would decide the question

Section 3.5 states a falsification criterion, and the empirical stage identifies precisely why it cannot be applied here. A decisive design requires four things that the present data lack. First, repeated directed evaluations, so that the same raters evaluate multiple notes from different authors and source types and edges recur. Second, observable source-rater and author-rater histories, so that relational phase has something to be estimated from. Third, reconstructed exposure order, so that ordering is not confused with documentary provenance. Fourth, enough closed cycles in the evaluation graph to support a cycle basis and a null distribution for holonomy. Given these, angular synchronisation could estimate relative phases (Singer 2011; Bandeira *et al.* 2013), gauge fixing could remove the removable inconsistency, and held-out cycle prediction could test whether residual holonomy explains disagreement beyond signed edges and scalar trust. That is a test QSD could fail. Until it is run, the theory should be described as computationally viable and empirically undecided, which is what this article concludes.

6.5. Implications for decarbonisation communication

Two findings bear on practice, both with the caveats above. The narrative results indicate that EV correction is not homogeneous: policy notes attracted several times the rating volume and evaluative weight of most technical categories, while charging notes showed the highest reason-tag disagreement in a small group. This argues against a single correction template across EV claims. Policy and mandate claims are contested on legitimacy as much as on fact, and correction there is likely to require cross-partisan sourcing and rapid response rather than additional technical detail. Lifecycle and grid claims turn on assumptions and boundaries, where transparency about the comparison being made matters more than the verdict. Safety and cold-weather claims generalise from vivid events, where the useful correction supplies the denominator and the conditional rather than a flat denial that risks reading as dismissal of a real limitation.

Longer notes were more likely to reach a resolution in this historical sample, consistent with evidence that context-rich notes increase trust (Drolsbach *et al.* 2024). The operational reading is not that length should be maximised but that a correction should carry the comparison, the time frame, the denominator and the source needed to resolve the claim. The negative association with domain count similarly suggests that source accumulation is not a substitute for a coherent argument. For decarbonisation policy specifically, the implication is that correction infrastructure should be resourced and evaluated by narrative rather than in aggregate, and that the evaluation metric should be the resolution and durability of contested claims rather than the volume of corrections issued. None of this depends on QSD being correct; it follows from the descriptive heterogeneity, and it would be worth acting on regardless of how the theoretical question is eventually settled.

6.6. Contribution to computational social science

The study contributes a validation sequence rather than a validated model. It demonstrates recoverability, maps the envelope within which recoverability holds, states what observation would distinguish the theory from its nearest competitors, and then reports that the available data cannot supply that observation and that the theory's proxy provides no predictive gain. This sequence aligns with the computational social-science emphasis on technically sound methods, reproducibility and falsifiable claims (Keuschnigg *et al.* 2018; Elmer 2023), and it applies beyond QSD. Signed networks, sheaf models, relational-event models and graph connections already capture much of the relevant mathematics; a new formalism earns its complexity only by identifying an observation that the alternatives cannot explain. Here the formalism's most useful output is a specification of what is missing from platform data: individual repeated responses, ordered exposure, independently

measured shocks and closed cycles. That is a modest contribution, and it is the one the evidence supports.

7. Limitations

The limitations below are stated once and are not repeated in the Conclusion, which the submitted version did. First, the dataset is historical, principally covering 2021 to March 2023, and the platform's ranking, participation and governance have changed since; nothing here describes the contemporary system. Second, individual ratings were unavailable, so the analysis rests on note-level aggregates and cannot estimate rater-specific signed coupling, cross-viewpoint consensus or contributor trajectories. Third, original post text was largely unavailable, so narrative classification relies on the note summary and cited context, with the selection and classification bias discussed in Section 4.2 and quantified by the subsample validation reported there. Fourth, citation order is documentary provenance rather than audience exposure order. Fifth, Ψ is a constructed composite with no claim to gauge invariance, and its sensitivity to construction is reported in Section 5.5 rather than assumed away. Sixth, the sample is small and imbalanced: 15 resolved cases bound every multivariable result, and narrative groups of five to seven observations support description only. Seventh, although a second blind coder was added at revision, coding of the trustworthy-source indicator remains a judgement that groups heterogeneous sources into one category. Eighth, the analysis is observational, so note length, citation count and author history may proxy unobserved claim complexity, contributor effort or topic salience.

A decisive future test should combine the public tables with individual ratings and lawful post-level interaction data, preregistered. It should estimate actor states and relation phases on a temporal multilayer network; compare gauge-fixed QSD against signed-network, relational-event, Bayesian trust and graph-neural baselines on held-out edges; and evaluate genuine loop observables using the criterion in Section 3.5. Controlled experiments should independently randomise correction source, exposure order and institutional history, which is the only clean way to separate order effects from recency. A second external domain, such as heat-pump or wind-power misinformation, would test whether estimated provenance effects generalise beyond consumer transport.

8. Conclusion

This study asked whether route-dependent correction in EV misinformation can be represented and tested through Quantum Sociodynamics. The answer is that it can be represented, that its parameters are recoverable within a mappable envelope, and that it cannot yet be tested. In controlled synthetic data the reduced QSD model is computationally estimable: six parameters were recovered with low bias, high true-estimated correlation and complete coupling-sign recovery under baseline conditions, degrading in interpretable ways as noise, panel length, phase observation and shock timing were degraded.

In historical correction data, citation-provenance breadth and note construction were associated with platform resolution, and a constructed provenance-memory inconsistency index was associated with lower aggregate rating weight at a raw threshold but not after false-discovery control. The distinctive theoretical claim was not tested at all: formal holonomy is not observable in these data, the index is not a measurement of it, and the index-augmented model provided no predictive gain over a simpler specification. The study therefore neither validates nor refutes credibility curvature. It identifies a measurement boundary and states, in Section 3.5, the observation that would cross it.

Quantum Sociodynamics should be judged neither by the novelty of its vocabulary nor by the success of simulation alone. Its value depends on whether transport, signed coupling, memory and holonomy generate reproducible explanatory or predictive gains in data rich enough to observe them. On the present evidence it is computationally viable and empirically undecided, and the useful contribution of this article is a precise account of what would decide it.

Appendix A. EV lexicon and sample construction

Table A1 gives the complete screening lexicon applied to English note summaries. An asterisk denotes a trailing wildcard. Table A2 gives the selection audit trail, including candidate counts, exclusion counts and the written exclusion rules. Both are reproduced verbatim from the screening script in the reproducibility deposit, and the counts in Table A2 are those produced by that script.

Table A1. EV screening lexicon by claim domain.

Claim domain	Screening terms
Vehicle and technology	electric vehicle*, electric car*, EV*, BEV, PHEV, plug-in hybrid, zero-emission vehicle*, ZEV, Tesla, e-car*
Battery and materials	battery, batteries, lithium, cobalt, nickel, cathode, battery pack, degradation, battery life, recycling
Charging	charging, charger*, charge point*, supercharger, fast charg*, kerbside charg*, charging infrastructure
Range and cold weather	range, range anxiety, cold weather, winter, freezing, mileage, real-world range
Fire and safety	battery fire*, thermal runaway, EV fire*, catch fire, explosion, car park fire
Lifecycle emissions	lifecycle, life-cycle, LCA, carbon footprint, well-to-wheel, embodied emissions, dirtier than
Grid and electricity	grid, electricity demand, blackout*, power cut*, national grid, load, peak demand
Cost and policy	tax credit*, subsidy, subsidies, incentive*, mandate, ban, 2030, 2035, resale value, insurance, total cost of ownership
Exclusion terms (applied at review)	expected value, EV/EBITDA, enterprise value, eigenvalue, Eurovision, ev. (abbreviation in non-English text)

Table A2. Sample selection audit trail.

Step	Records	Rule applied
Notes in historical derivative snapshot	NR	Full snapshot as retrieved
English-language note summaries	NR	Language field or detection where field absent
Lexicon-positive candidates	~109	One or more Table A1 terms present, exclusion terms absent
Excluded: non-EV sense of an ambiguous token	NR	Token appears only in a financial, statistical or unrelated sense
Excluded: EV term incidental to an unrelated claim	NR	Corrected claim is not about EVs
Excluded: non-English or unusable text	NR	Note text cannot be coded reliably
Total excluded	~8	-
Final analytical sample	101	Linked to 75 unique post identifiers

NR, not reported.

Appendix B. Coding codebook summary

Table B1 summarises the codebook used for narrative and source coding. The full codebook, including worked examples of boundary cases and the blind coding instrument issued to the second coder, is in the reproducibility deposit together with the disagreement log. Coding rule B0 applies throughout: where a note's subject is genuinely ambiguous between two narratives, the narrative of the corrected claim takes precedence over the narrative of the cited evidence.

Table B1. Narrative and source codebook, decision rules and boundary cases.

Code	Inclusion rule	Boundary case and resolution
EV-LCA	Claim concerns lifecycle, manufacturing or comparative emissions	Emissions claims framed as grid impact code as EV-GRID if the mechanism is electricity supply
EV-POL	Claim concerns mandates, bans, subsidies, tax credits or targets	Cost claims that turn on a subsidy code as EV-POL; general affordability codes as EV-COST
EV-COST	Claim concerns purchase price, running cost, resale, insurance or repair	Battery replacement cost codes as EV-COST; battery durability codes as EV-BATT
EV-BATT	Claim concerns battery chemistry, materials, sourcing, degradation or recycling	Battery fires code as EV-FIRE
EV-FIRE	Claim concerns fire, thermal runaway or crash safety	Charging-station fires code as EV-FIRE, not EV-CHARGE
EV-CHARGE	Claim concerns charger availability, speed, reliability or deployment	Charging load on the network codes as EV-GRID
EV-COLD	Claim concerns cold-weather range, efficiency or charging performance	General range claims without a weather condition code as EV-COLD only if range is the subject
EV-GRID	Claim concerns electricity demand, capacity, blackouts or network constraint	Generation-mix emissions code as EV-LCA
Source type	Government/regulator; scientific/fact-checking; news; manufacturer/industry; social/peer; other	Government-funded research institutes code as scientific; agency press releases code as government
Trustworthy-source flag	At least one government, regulator, scientific or professional fact-checking source	Sensitivity analyses use a narrower flag (regulators and peer-reviewed only) and a graded three-level scheme

Appendix C. Full model results and diagnostics

Table C1 reports every variable in every specification with coefficients, standard errors, intervals and p-values, replacing the selective reporting of the submitted version. Table C2 reports collinearity diagnostics. Figure C1 in the deposit shows residual diagnostics for the continuous outcome models and the alternative specifications described in Section 5.3.

Table C1. Complete regression results for all specifications.

Model	Predictor	Estimate	SE	95% interval	p
Firth, 4-predictor (primary)	Cited-domain count	~0.381	~0.484	~0.123 to 0.821	~0.0110
Firth, 4-predictor (primary)	Note word count	~2.580	~0.410	~1.414 to 6.740	~0.0050
Firth, 4-predictor (primary)	Smoothed prior helpfulness	~0.775	~0.321	~0.359 to 1.409	~0.0849
Firth, 4-predictor (primary)	Classification direction	NR	~1.442	NR	NR
MLE, 7-predictor (as submitted)	Cited-domain count	-1.801	0.6754	-3.125 to -0.478	0.0076
MLE, 7-predictor (as submitted)	Note word count	1.074	0.4522	0.187 to 1.960	0.0176
MLE, 7-predictor (as submitted)	Source-type entropy	~0.559	~0.615	~-0.647 to 1.765	~0.3638
MLE, 7-predictor (as submitted)	Trustworthy-source flag	~0.600	~2.008	~-3.335 to 4.536	~0.7650

Table C1. (continued)

Model	Predictor	Estimate	SE	95% interval	p
MLE, 7-predictor (as submitted)	log prior note count	~-0.333	~0.394	~-1.106 to 0.440	~0.3979
MLE, 7-predictor (as submitted)	Smoothed prior helpfulness	~-0.125	~0.427	~-0.963 to 0.712	~0.7696
MLE, 7-predictor (as submitted)	Classification direction	NR	~2.67×10^7	NR	~1.000
OLS, log rating weight	All predictors	~0.0927	-	-	-

NR, not reported.

Table C2. Collinearity diagnostics for the structurally related provenance predictors.

Predictor	VIF	Correlation with cited-domain count
Cited-domain count	3.128	1.000
Source-type entropy	32.381	0.682
Source-type count	13.006	0.730
Inconsistency index Ψ	24.326	0.669
Note word count	2.224	0.733

Declarations

Author contributions. Muhammad Talut: Conceptualisation; Methodology; Software; Formal analysis; Investigation; Data curation; Visualisation; Writing - original draft; Writing - review and editing. The author has read and approved the final version of the manuscript.

Funding. The author received no specific funding for this study.

Conflicts of interest. The author declares no competing interests.

Data availability. The official Community Notes data structure is documented by X Corp. (n.d.). The historical derivative used for the empirical analysis is available through the shiruken birdwatch-data repository (shiruken, n.d.). The complete reproducibility package for this article is deposited in a persistent public repository (DOI pending) and was deposited at the revision stage rather than after acceptance. It contains: the exact historical data snapshot together with an acquisition script that recreates it from the public source; the filtered analytical dataset; all preprocessing, feature-construction and coding scripts; the complete EV lexicon; the coding codebook, the blind second-coder instrument and the disagreement log; the simulation and stress-grid code; all random seeds; the cross-validation procedures; and the figure-generation scripts. This is the only accessibility statement that applies; the phrase 'available on reasonable request' has been removed from the manuscript, as has the inconsistent description of these materials as supplied online resources.

Ethics. The study analysed publicly released, pseudonymised archival platform records. It involved no interaction with users, no attempt at re-identification and no inference of protected personal attributes. Formal ethical approval was not required for this analysis of public archival data.

Editorial independence. The author retained full responsibility for the study design, data transformations, analysis, interpretation and decision to submit the manuscript.

Generative AI. Generative AI tools were used for language editing and coding assistance. The author designed the model and analyses, verified the data transformations and numerical results, reviewed all text and accepts full responsibility for the work.

Acknowledgements. The author gratefully acknowledges Professor Nigel Gilbert for comments on an earlier formulation of the Quantum Sociodynamics research programme, and thanks the two anonymous reviewers, whose detailed comments materially changed the framing, the statistical

treatment and the claims of this article. Any remaining errors, interpretations and arguments are the author's own.

References

Allen, Jennifer; Martel, Cameron; Rand, David G (2022) 'Birds of a feather don't fact-check each other: Partisanship and the evaluation of news in Twitter's Birdwatch crowdsourced fact-checking program', *CHI Conference on Human Factors in Computing Systems*, pp. 1-19. Available at: https://doi.org/10.1145/3491102.3502040

Bandeira, Afonso S.; Singer, Amit; Spielman, Daniel A. (2013) 'A Cheeger Inequality for the Graph Connection Laplacian', *SIAM Journal on Matrix Analysis and Applications*, 34(4), pp. 1611-1630. Available at: https://doi.org/10.1137/120875338

Benjamini, Yoav; Hochberg, Yosef (1995) 'Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing', *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 57(1), pp. 289-300. Available at: https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

Bretter, Christian; Pearson, Samuel; Hornsey, Matthew J.; MacInnes, Sarah; Sassenberg, Kai; Wade, Belinda; Winter, Kevin (2025) 'Mapping, understanding and reducing belief in misinformation about electric vehicles', *Nature Energy*, 10(7), pp. 869-879. Available at: https://doi.org/10.1038/s41560-025-01790-0

Busemeyer, J.R. and Bruza, P.D. (2012) *Quantum Models of Cognition and Decision*. Cambridge: Cambridge University Press.

Butts, Carter T. (2008) '4. A Relational Event Framework for Social Action', *Sociological Methodology*, 38(1), pp. 155-200. Available at: https://doi.org/10.1111/j.1467-9531.2008.00203.x

Cameron, A. Colin; Gelbach, Jonah B.; Miller, Douglas L. (2008) 'Bootstrap-Based Improvements for Inference with Clustered Errors', *Review of Economics and Statistics*, 90(3), pp. 414-427. Available at: https://doi.org/10.1162/rest.90.3.414

Castellano, Claudio; Fortunato, Santo; Loreto, Vittorio (2009) 'Statistical physics of social dynamics', *Reviews of Modern Physics*, 81(2), pp. 591-646. Available at: https://doi.org/10.1103/RevModPhys.81.591

Chow, Angel Onki; Dolter, Brett; Hurlbert, Margot; Zapata, Oscar (2025) '"Electric vehicles do not work in the winter": Exploring the role of misperceptions and socio-psychological factors in zero-emission vehicle adoption', *Energy Research & Social Science*, 130, pp. 104450. Available at: https://doi.org/10.1016/j.erss.2025.104450

Chuai, Yuwei; Pilarski, Moritz; Renault, Thomas; Restrepo-Amariles, David; Troussel-Clément, Aurore; Lenzini, Gabriele; Pröllochs, Nicolas (2026) 'Community-based fact-checking reduces the spread of misleading posts on X (formerly Twitter)', *Nature Communications*, 17(1). Available at: https://doi.org/10.1038/s41467-026-72597-0

Degroot, Morris H. (1974) 'Reaching a Consensus', *Journal of the American Statistical Association*, 69(345), pp. 118-121. Available at: https://doi.org/10.1080/01621459.1974.10480137

Drolsbach, Chiara Patricia; Solovev, Kirill; Pröllochs, Nicolas (2024) 'Community notes increase trust in fact-checking on social media', *PNAS Nexus*, 3(7). Available at: https://doi.org/10.1093/pnasnexus/pgae217

Ecker, Ullrich K. H.; Lewandowsky, Stephan; Cook, John; Schmid, Philipp; Fazio, Lisa K.; Brashier, Nadia; Kendeou, Panayiota; Vraga, Emily K.; Amazeen, Michelle A. (2022) 'The psychological drivers of misinformation belief and its resistance to correction', *Nature Reviews Psychology*, 1(1), pp. 13-29. Available at: https://doi.org/10.1038/s44159-021-00006-y

Elmer, Timon (2023) 'Computational social science is growing up: why puberty consists of embracing measurement validation, theory development, and open science practices', *EPJ Data Science*, 12(1). Available at: https://doi.org/10.1140/epjds/s13688-023-00434-1

Entman, Robert M. (1993) 'Framing: Toward Clarification of a Fractured Paradigm', *Journal of Communication*, 43(4), pp. 51-58. Available at: https://doi.org/10.1111/j.1460-2466.1993.tb01304.x

Facchetti, Giuseppe; Iacono, Giovanni; Altafini, Claudio (2011) 'Computing global structural balance in large-scale signed social networks', *Proceedings of the National Academy of Sciences*, 108(52), pp. 20953-20958. Available at: https://doi.org/10.1073/pnas.1109521108

Firth, David (1993) 'Bias reduction of maximum likelihood estimates', *Biometrika*, 80(1), pp. 27-38. Available at: https://doi.org/10.1093/biomet/80.1.27

Fraxanet, Emma; Pellert, Max; Schweighofer, Simon; Gómez, Vicenç; Garcia, David (2024) 'Unpacking polarization: Antagonism and alignment in signed networks of online interaction', *PNAS Nexus*, 3(12). Available at: https://doi.org/10.1093/pnasnexus/pgae276

Friedkin, Noah E.; Johnsen, Eugene C. (1990) 'Social influence and opinions', *The Journal of Mathematical Sociology*, 15(3-4), pp. 193-206. Available at: https://doi.org/10.1080/0022250X.1990.9990069

Good, P.I. 2005 Permutation, Parametric and Bootstrap Tests of Hypotheses. 3rd ed. New York: Springer. DOI: https://doi.org/10.1007/b138696 (2005) *Permutation, Parametric and Bootstrap Tests of Hypotheses*. Springer-Verlag. Available at: https://doi.org/10.1007/b138696

Granovetter, Mark (1978) 'Threshold Models of Collective Behavior', *American Journal of Sociology*, 83(6), pp. 1420-1443. Available at: https://doi.org/10.1086/226707

Guha, R.; Kumar, Ravi; Raghavan, Prabhakar; Tomkins, Andrew (2004) 'Propagation of trust and distrust', *Proceedings of the 13th international conference on World Wide Web*, pp. 403-412. Available at: https://doi.org/10.1145/988672.988727

Hansen, Jakob; Ghrist, Robert (2021) 'Opinion Dynamics on Discourse Sheaves', *SIAM Journal on Applied Mathematics*, 81(5), pp. 2033-2060. Available at: https://doi.org/10.1137/20M1341088

Heinze, Georg; Schemper, Michael (2002) 'A solution to the problem of separation in logistic regression', *Statistics in Medicine*, 21(16), pp. 2409-2419. Available at: https://doi.org/10.1002/sim.1047

Hoff, Peter D; Raftery, Adrian E; Handcock, Mark S (2002) 'Latent Space Approaches to Social Network Analysis', *Journal of the American Statistical Association*, 97(460), pp. 1090-1098. Available at: https://doi.org/10.1198/016214502388618906

Holme, Petter; Saramäki, Jari (2012) 'Temporal networks', *Physics Reports*, 519(3), pp. 97-125. Available at: https://doi.org/10.1016/j.physrep.2012.03.001

Ilinski, K. (2001) 'Physics of Finance: Gauge Modelling in Non-Equilibrium Pricing', *Chichester: Wiley*.

Jasanoff, S. and Kim, S.H. (2015) *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. Chicago: University of Chicago Press.

Kangur, Uku; Chakraborty, Roshni; Sharma, Rajesh (2026) 'Who Checks the Checkers? Exploring Source Credibility in Twitter's Community Notes', *Journal of Computational Social Science*, 9(1). Available at: https://doi.org/10.1007/s42001-025-00455-y

Keuschnigg, Marc; Lovsjö, Niclas; Hedström, Peter (2018) 'Analytical sociology and computational social science', *Journal of Computational Social Science*, 1(1), pp. 3-14. Available at: https://doi.org/10.1007/s42001-017-0006-5

Kleinberg, Jon M. (1999) 'Authoritative sources in a hyperlinked environment', *Journal of the ACM*, 46(5), pp. 604-632. Available at: https://doi.org/10.1145/324133.324140

Krippendorff, Klaus (2019) *Content Analysis: An Introduction to Its Methodology*. SAGE Publications, Inc.. Available at: https://doi.org/10.4135/9781071878781

Kruskal, William H.; Wallis, W. Allen (1952) 'Use of Ranks in One-Criterion Variance Analysis', *Journal of the American Statistical Association*, 47(260), pp. 583-621. Available at: https://doi.org/10.1080/01621459.1952.10483441

Kunegis, Jérôme; Lommatzsch, Andreas; Bauckhage, Christian (2009) 'The slashdot zoo', *Proceedings of the 18th international conference on World wide web*, pp. 741-750. Available at: https://doi.org/10.1145/1526709.1526809

Leenders, R.T.A.J. (2002) 'Modeling social influence through network autocorrelation: constructing the weight matrix', *Social Networks*, 24, pp. 21–47. Available at: https://doi.org/10.1016/S0378-8733(01)00049-1

Leskovec, Jure; Huttenlocher, Daniel; Kleinberg, Jon (2010) 'Signed networks in social media', *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1361-1370. Available at: https://doi.org/10.1145/1753326.1753532

Meylahn, Benedikt V.; Den Boer, Arnoud V.; Mandjes, Michel (2024) 'Trusting: Alone and together', *The Journal of Mathematical Sociology*, 48(4), pp. 424-478. Available at: https://doi.org/10.1080/0022250X.2024.2340135

Özkara, Yasin; Bilişli, Yasemin; Yildirim, Fatih Serdar; Kayan, Fahrettin; Başdeğirmen, Agah; Kayakuş, Mehmet; Yiğit Açıkgöz, Fatma (2025) 'Analysing Social Media Discourse on Electric Vehicles with Machine Learning', *Applied Sciences*, 15(8), pp. 4395. Available at: https://doi.org/10.3390/app15084395

Panizza, Folco; Ronzani, Piero; Morisseau, Tiffany; Mattavelli, Simone; Martini, Carlo (2023) 'How do online users respond to crowdsourced fact-checking?', *Humanities and Social Sciences Communications*, 10(1). Available at: https://doi.org/10.1057/s41599-023-02329-y

Peduzzi, P., Concato, J., Kemper, E., Holford, T.R. and Feinstein, A.R. (1996) 'A simulation study of the number of events per variable in logistic regression analysis', *Journal of Clinical Epidemiology*, 49, pp. 1373–1379. Available at: https://doi.org/10.1016/S0895-4356(96)00236-3

Pidgeon, Nick; Fischhoff, Baruch (2011) 'The role of social and decision sciences in communicating uncertain climate risks', *Nature Climate Change*, 1(1), pp. 35-41. Available at: https://doi.org/10.1038/nclimate1080

Pröllochs, Nicolas (2022) 'Community-Based Fact-Checking on Twitter's Birdwatch Platform', *Proceedings of the International AAAI Conference on Web and Social Media*, 16, pp. 794-805. Available at: https://doi.org/10.1609/icwsm.v16i1.19335

Riley, Richard D; Snell, Kym IE; Ensor, Joie; Burke, Danielle L; Harrell Jr, Frank E; Moons, Karel GM; Collins, Gary S (2019) 'Minimum sample size for developing a multivariable prediction model: PART II - binary and time-to-event outcomes', *Statistics in Medicine*, 38(7), pp. 1276-1296. Available at: https://doi.org/10.1002/sim.7992

Saito, Takaya; Rehmsmeier, Marc (2015) 'The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets', *PLOS ONE*, 10(3), pp. e0118432. Available at: https://doi.org/10.1371/journal.pone.0118432

Sakaguchi, Hidetsugu; Kuramoto, Yoshiki (1986) 'A Soluble Active Rotator Model Showing Phase Transitions via Mutual Entrainment', *Progress of Theoretical Physics*, 76(3), pp. 576-581. Available at: https://doi.org/10.1143/PTP.76.576

Sarkar, Purnamrita; Moore, Andrew W. (2005) 'Dynamic social network analysis using latent space models', *ACM SIGKDD Explorations Newsletter*, 7(2), pp. 31-40. Available at: https://doi.org/10.1145/1117454.1117459

shiruken (no date) *birdwatch-data: historical public Birdwatch/Community Notes derivative*. Available at: https://github.com/shiruken/birdwatch-data

Singer, A. (2011) 'Angular synchronization by eigenvectors and semidefinite programming', *Applied and Computational Harmonic Analysis*, 30(1), pp. 20-36. Available at: https://doi.org/10.1016/j.acha.2010.02.001

Snijders, Tom A.B.; van de Bunt, Gerhard G.; Steglich, Christian E.G. (2010) 'Introduction to stochastic actor-based models for network dynamics', *Social Networks*, 32(1), pp. 44-60. Available at: https://doi.org/10.1016/j.socnet.2009.02.004

Solovev, Kirill; Pröllochs, Nicolas (2025) 'References to unbiased sources increase the helpfulness of community fact-checks', *Scientific Reports*, 15(1). Available at: https://doi.org/10.1038/s41598-025-09372-6

Sovacool, Benjamin K.; Griffiths, Steve (2020) 'Culture and low-carbon energy transitions', *Nature Sustainability*, 3(9), pp. 685-693. Available at: https://doi.org/10.1038/s41893-020-0519-4

Tokita, Christopher K; Aslett, Kevin; Godel, William P; Sanderson, Zeve; Tucker, Joshua A; Nagler, Jonathan; Persily, Nathaniel; Bonneau, Richard (2024) 'Measuring receptivity to misinformation at scale on a social media platform', *PNAS Nexus*, 3(10). Available at: https://doi.org/10.1093/pnasnexus/pgae396

van Smeden, Maarten; de Groot, Joris A. H.; Moons, Karel G. M.; Collins, Gary S.; Altman, Douglas G.; Eijkemans, Marinus J. C.; Reitsma, Johannes B. (2016) 'No rationale for 1 variable per 10 events criterion for binary logistic regression analysis', *BMC Medical Research Methodology*, 16(1). Available at: https://doi.org/10.1186/s12874-016-0267-3

Vosoughi, Soroush; Roy, Deb; Aral, Sinan (2018) 'The spread of true and false news online', *Science*, 359(6380), pp. 1146-1151. Available at: https://doi.org/10.1126/science.aap9559

Wang, Zheng; Solloway, Tyler; Shiffrin, Richard M.; Busemeyer, Jerome R. (2014) 'Context effects produced by question orders reveal quantum nature of human judgments', *Proceedings of the National Academy of Sciences*, 111(26), pp. 9431-9436. Available at: https://doi.org/10.1073/pnas.1407756111

Watts, Duncan J. (2002) 'A simple model of global cascades on random networks', *Proceedings of the National Academy of Sciences*, 99(9), pp. 5766-5771. Available at: https://doi.org/10.1073/pnas.082090499

Whitmarsh, Lorraine; Poortinga, Wouter; Capstick, Stuart (2021) 'Behaviour change to address climate change', *Current Opinion in Psychology*, 42, pp. 76-81. Available at: https://doi.org/10.1016/j.copsyc.2021.04.002

Wilson, Kenneth G. (1974) 'Confinement of quarks', *Physical Review D*, 10(8), pp. 2445-2459. Available at: https://doi.org/10.1103/PhysRevD.10.2445

Wojcik, S., Hilgard, S., Judd, N., Mocanu, D., Ragain, S., Hunzaker, M.B.F., Coleman, K. and Baxter, J. (2022) 'Birdwatch: crowd wisdom and bridging algorithms can inform understanding and reduce the spread of misinformation', *arXiv preprint arXiv:2210.15723*. Available at: https://doi.org/10.48550/arXiv.2210.15723

Wynne, Brian (1992) 'Misunderstood misunderstanding: social identities and public uptake of science', *Public Understanding of Science*, 1(3), pp. 281-304. Available at: https://doi.org/10.1088/0963-6625/1/3/004

X Corp. (no date) *Downloading Community Notes data*. Available at: https://communitynotes.x.com/guide/en/under-the-hood/download-data

Young, K. (1999) 'Foreign exchange market as a lattice gauge theory', *American Journal of Physics*, 67(10), pp. 862-868. Available at: https://doi.org/10.1119/1.19139

HOW TO CITE THIS ARTICLE

Talut, M. (2026). Route-Dependent Correction and Credibility Curvature in Online Electric-Vehicle Misinformation: A Computational Test of Quantum Sociodynamics. Global Decarbonisation, 2, 283-317. https://doi.org/10.65582/gd.2026.014

COPYRIGHT AND LICENCE

© 2026 The Author(s). This is an open access article distributed under the terms of the Creative Commons Attribution (CC BY 4.0) licence, which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.