

Article

Not peer-reviewed version

Operational Certification Horizons in Quantum Transport: Copy Time, Conservation Laws, and a Rigorous Diffusive Benchmark

[Sacha Mohamed](#) *

Posted Date: 31 March 2026

doi: 10.20944/preprints202601.0364.v6

Keywords: quantum hypothesis testing; trace distance; information propagation; conservation laws; diffusion; operator hydrodynamics; tensor networks



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Operational Certification Horizons in Quantum Transport: Copy Time, Conservation Laws, and a Microscopic Diffusive Benchmark

Sacha Mohamed

Independent Researcher, Casablanca, Morocco; www.sachamed@gmail.com

Abstract

We introduce an operational transport latency, *quantum information copy time*: the earliest time at which a receiver confined to a region B can certify, with prescribed advantage, which of two global hypotheses was prepared by local operations in a distant sender region A . The benchmark quantity is information-theoretic—the Helstrom advantage on B , namely the trace distance between the reduced states—and it also admits receiver-restricted refinements that make measurement constraints explicit, including few-body and moment-channel receivers. We derive the corresponding kinematic locality constraints for Hamiltonian and Lindbladian dynamics with Lieb–Robinson tails, as well as for circuits and quantum cellular automata with strict light cones, and we record the corresponding degradation for algebraic locality tails. For the quantum symmetric simple exclusion process (Q-SSEP) we formulate an explicit diagonal-sector linear-response benchmark: within a microscopically specified family of locally biased diagonal states, the *Helstrom* copy time obeys a diffusion-limited lower bound in terms of the diffusion constant D and the static susceptibility χ . For closed Hamiltonian systems, we recast the projection-based hydrodynamic argument as a conditional reduced-dynamics statement controlled by an explicit short-memory assumption rather than by an artificial dissipative gap of the microscopic unitary Liouvillian. We further sharpen the operational statement itself by showing that fixed-threshold asymptotics are not meaningful unless the admissible threshold, receiver size, or perturbation amplitude is specified together with distance; the quadratic transport law is therefore stated only in threshold-scaled regimes in which the receiver signal remains detectable, and the two threshold-crossing branches are written in closed form through the Lambert- W function. For closed Hamiltonian systems, we additionally prove a microscopic no-go statement: an orthogonal Mori block of the exact unitary Liouvillian cannot carry a positive dissipative quadratic-form gap in the Kubo–Mori geometry, so any usable closure must enter through memory-kernel control rather than through a fictitious fast-sector gap. We complement the analytical framework with exact-diagonalization diagnostics in the XXZ chain, with explicit threshold-robustness checks, and with executed TEBD/MPS cross-checks beyond the exact-diagonalization sizes (Supplementary S2 and Code SC1), while keeping all numerical conclusions finite-size and non-asymptotic. Finally, we compare copy time with scrambling diagnostics based on out-of-time-ordered correlators and show how conservation laws can delay certifiability well beyond the ballistic operator-growth front without any contradiction with locality.

Keywords: quantum hypothesis testing; trace distance; information propagation; conservation laws; diffusion; operator hydrodynamics; tensor networks

1. Introduction

Transport in quantum many-body systems is usually inferred from expectation values and correlation functions, while information-theoretic diagnostics are often framed in terms of operator growth and scrambling. This work takes a complementary, task-defined viewpoint, designed to isolate the operational content of transport claims. We ask when a localized perturbation prepared in a sender

region A becomes *remotely certifiable* by an observer restricted to a distant region B . The resulting timescale—the *quantum information copy time*—is defined by binary hypothesis testing on B : it is the earliest time at which the receiver can distinguish two evolving global hypotheses with a prescribed advantage.

Two aspects make this perspective useful for transport theory. First, it cleanly separates *kinematics* from *mechanism*: locality bounds constrain when influence can begin, but they do not identify which dynamical sector controls a concrete certification task. Second, in systems with conservation laws the sector that dominates certifiability can be parametrically slower than the fastest operator-growth sector. In particular, ballistic growth of commutators (and large OTOCs) does not imply that a charge-biased hypothesis becomes certifiable at a comparable time; the bottleneck may be diffusive (or sub-/super-diffusive) transport of the conserved bias.

A central aim of this manuscript is to place this statement on a precise footing without conflating ideal receivers with experimentally accessible ones. Accordingly, we distinguish two notions throughout: (i) the *Helstrom* copy time, defined by the trace distance between reduced states on B (equivalently, optimal measurements on B), and (ii) *restricted* copy times, where the receiver is constrained to a specified measurement/observable class (few-body, moment channels, coarse-grained charge, *etc.*). The restricted notion is the natural object when one wants scaling claims tied to hydrodynamic transport; the Helstrom notion is the mathematically canonical benchmark.

Main contributions.

- C1. Operational framework.** We define receiver advantage and copy time as a task-defined latency, and we provide a receiver-class refinement that makes measurement restrictions explicit (Section 2.1).
- C2. Minimal locality constraints.** For Hamiltonian/Lindbladian dynamics with Lieb–Robinson tails and for circuits/QCAs with strict light cones, we derive a kinematic lower bound on copy time that isolates the threshold dependence and constants (Section 3).
- C3. A controlled diffusive benchmark (Q-SSEP).** In the locality-preserving quantum symmetric simple exclusion process (Q-SSEP), for an explicit diagonal family of locally biased hypotheses whose reduced receiver states remain diagonal, the *Helstrom* copy time reduces exactly to a classical total-variation problem on B and obeys a diffusion-limited lower bound of the form

$$\tau_{\text{copy}}(A \rightarrow B; \eta) \geq C(\eta, \epsilon, \chi, |A|, |B|) \frac{\ell^2}{D},$$

where $\ell = d(A, B)$, D is the diffusion constant of the conserved density, and χ is the associated static susceptibility. We therefore use Q-SSEP as a controlled microscopic benchmark rather than as a proxy for generic closed-system dynamics (Section 6; Supplementary S3).

- C4. Hydrodynamic reduction for unitary systems.** For closed Hamiltonian dynamics we formulate a projection-based route from charge bias to receiver advantage and identify the concrete hypotheses under which a single slow transport pole controls the operational signal (Sections 4–5). The resulting statement is organized so that each required ingredient admits a direct diagnostic proxy.
- C5. Microscopic clarification of the closure problem.** We prove that, in the exact Kubo–Mori geometry of a closed Hamiltonian system, an orthogonal fast-sector Mori block cannot possess the positive dissipative quadratic-form gap sometimes used in heuristic closures; the only mathematically coherent route is via explicit control of the memory kernel or of coarse-grained correlation decay (Section 4; Supplementary S6).
- C6. Benchmarking and protocols.** We provide finite-size transport diagnostics in the XXZ chain with explicit uncertainty controls, together with a TEBD/MPS validation layer, archived larger- L open-boundary cross-check runs, and an explicit beyond-ED diagnostic protocol for identifying a numerically credible single-mode window (Section 10; Supplementary S2 and S7).

C7. Comparison to scrambling diagnostics and extensions. We show, in a sharp operational sense, why a ballistic OTOC front can coexist with diffusion-limited certifiability under conservation laws (Section 8), and we outline how copy time can act as an operational microscope for non-diffusive regimes such as KPZ superdiffusion, Griffiths subdiffusion, and MBL (Supplementary S4).

Logical status of the results.

The manuscript is organized in three layers, and the distinction between them is part of the message rather than a caveat. At the foundational level, the definition of copy time and the locality constraints are exact. At the benchmark level, the Q-SSEP discussion yields a controlled microscopic lower bound for a precisely specified diagonal family of hypotheses. At the Hamiltonian level, the projection-based hydrodynamic reduction is stated only as a conditional reduced-dynamics statement with explicit assumptions, explicit error terms, and explicit diagnostic proxies, so that the operational scaling claim is attached to a delimited dynamical regime rather than to an informal mixing narrative.

Positioning.

Receiver advantage is the trace distance between reduced states and is standard in quantum hypothesis testing [1,2]. What is emphasized here is the *hitting-time* interpretation at fixed advantage and, crucially, the way conservation laws and receiver restrictions can turn transport coefficients into operational latencies. This connects naturally to the linear-response and projection-operator traditions [3–6] and to recent work on operator hydrodynamics in conserving dynamics [7,8]. The diffusive benchmark (Q-SSEP) provides a controlled setting in which the operational inequality can be traced directly to an exactly diffusive diagonal sector, while the unitary sections now also isolate a microscopic no-go theorem for the fast-sector-gap ansatz and then describe how analogous bounds may emerge in Hamiltonian systems only under explicit memory-kernel assumptions.

2. Operational Definition and Preliminaries

2.1. Copy Time as Receiver-Limited Hypothesis Testing

Fix disjoint lattice regions A (sender) and B (receiver). To avoid ambiguity about “support in A ”, we take the hypotheses to be prepared from a common reference state ρ_{ref} by *local* channels on A :

$$\rho_{\pm}(0) = \Lambda_A^{\pm}(\rho_{\text{ref}}), \quad \Lambda_A^{\pm} = \Lambda_A^{\pm} \otimes \text{id}_{A^c}. \quad (1)$$

This guarantees $\rho_{+}(0)$ and $\rho_{-}(0)$ coincide on A^c . The evolved hypotheses are $\rho_{\pm}(t) = \mathcal{U}_t(\rho_{\pm}(0))$, where $\mathcal{U}_t$ is a unitary channel (closed dynamics) or, more generally, any CPTP dynamics. For the receiver region B , let $\rho_{\pm,B}(t) = \text{Tr}_{B^c} \rho_{\pm}(t)$.

Definition 1 (Receiver advantage and copy times). [1,2] Fix a threshold $\eta \in (0,1)$.

(a) **Helstrom (unrestricted) advantage.** The optimal receiver advantage is

$$\text{Adv}_{A \rightarrow B}^{\text{full}}(t) := \frac{1}{2} \|\rho_{+,B}(t) - \rho_{-,B}(t)\|_1, \quad (2)$$

equivalently the Helstrom bias for discriminating $\rho_{+,B}(t)$ vs. $\rho_{-,B}(t)$ with equal priors. The associated copy time is

$$\tau_{\text{copy}}^{\text{full}}(A \rightarrow B; \eta) := \inf\{t \geq 0 : \text{Adv}_{A \rightarrow B}^{\text{full}}(t) \geq \eta\}. \quad (3)$$

(b) **Restricted advantage.** Let $\mathcal{O}_B$ be an admissible observable class on B with $\|O_B\|_{\infty} \leq 1$ for all $O_B \in \mathcal{O}_B$ (few-body algebras, moment channels, coarse-grained charge observables, etc.). Define the restricted advantage

$$\text{Adv}_{A \rightarrow B}^{\mathcal{O}}(t) := \frac{1}{2} \sup_{O_B \in \mathcal{O}_B} \text{Tr}[O_B(\rho_{+,B}(t) - \rho_{-,B}(t))], \quad (4)$$

and the corresponding restricted copy time $\tau_{\text{copy}}^{\mathcal{O}}(A \rightarrow B; \eta) := \inf\{t \geq 0 : \text{Adv}_{A \rightarrow B}^{\mathcal{O}}(t) \geq \eta\}$ by the same hitting-time rule as in (3). Equivalently, if $\Phi_{\mathcal{O}}$ denotes a CPTP “measurement compression” channel whose dual maps the unit ball into $\mathcal{O}_B$, then $\text{Adv}^{\mathcal{O}}(t) = \frac{1}{2} \|\Phi_{\mathcal{O}}(\rho_{+,B}(t) - \rho_{-,B}(t))\|_1$.

By data processing, $\text{Adv}_{A \rightarrow B}^{\mathcal{O}}(t) \leq \text{Adv}_{A \rightarrow B}^{\text{full}}(t)$ and hence $\tau_{\text{copy}}^{\text{full}}(A \rightarrow B; \eta) \leq \tau_{\text{copy}}^{\mathcal{O}}(A \rightarrow B; \eta)$. Most transport-driven scaling statements in this manuscript concern *restricted* copy times, because conservation laws naturally constrain the evolution of coarse observables. The controlled Q-SSEP benchmark in Section 6 applies to the *Helstrom* copy time.

Operational sampling complexity (how many shots?).

“Operational” is meaningful only if the advantage can be estimated with a controlled number of experimental samples. For a fixed time t and an observable O_B with eigenvalues in $[-1, 1]$, an empirical estimator $\hat{m} = \frac{1}{n} \sum_{j=1}^n o_j$ of $m = \text{Tr}(O_B \delta \rho_B(t))$ obeys a Hoeffding bound: $n = \Theta(\delta^{-2} \log(1/\alpha))$ samples suffice to estimate m to additive error δ with failure probability α . Consequently, certifying $\text{Adv}_{A \rightarrow B}^{\mathcal{O}}(t) \geq \eta$ up to slack δ requires $n = O(\delta^{-2} \log(1/\alpha))$ shots for each candidate O_B (or for the compressed measurement channel $\Phi_{\mathcal{O}}$). For the special case $B = \text{one site}$ and charge-biased hypotheses (Section 10.2), the Helstrom-optimal measurement is a single local two-outcome measurement, so the same $O(\delta^{-2} \log(1/\alpha))$ scaling applies without tomography. More concretely, if one scans a discrete time grid of N_t candidate times and M receiver observables under symmetric readout error rate $p < 1/2$, a union bound together with Hoeffding gives the sufficient budget

$$n \geq \frac{2}{(1-2p)^2 \delta^2} \log\left(\frac{2MN_t}{\alpha}\right), \quad (5)$$

per experimental setting. Supplementary S7 derives (5) and discusses adaptive coarse-to-fine scans that reduce the practical prefactor.

Proposition 1 (Binary certification is the universal pairwise primitive). *Let $\{\rho_B^{(a)}(t)\}_{a=1}^M$ be a family of receiver states and let $\tau_M^{\mathcal{O}}(\eta)$ denote any receiver-restricted M -ary certification time whose success criterion implies pairwise distinguishability at level η for every pair of labels. Then*

$$\tau_M^{\mathcal{O}}(\eta) \geq \max_{a \neq b} \tau_{a,b}^{\mathcal{O}}(\eta), \quad (6)$$

where $\tau_{a,b}^{\mathcal{O}}(\eta)$ is the binary copy time for the pair (a, b) under the same receiver restriction.

Proof. If the M -ary certification task succeeds at time t , then restricting the same protocol to any pair of labels (a, b) yields a valid binary discrimination protocol for that pair. Hence $t \geq \tau_{a,b}^{\mathcal{O}}(\eta)$ for every pair. Taking the maximum over pairs and then the infimum over admissible t proves the claim. $\square$

Experimental recipe. The copy-time definition does not require full tomography. A minimal protocol is:

- (E1) Prepare two hypotheses $\rho_{\pm}(0)$ that differ by a small local charge bias (e.g., a local Z-rotation or a weak local chemical-potential shift on region A).
- (E2) Evolve for time t under the conserving dynamics of interest.
- (E3) Measure a charge-like observable on B (one-site S_{ℓ}^z , block charge Q_B , or low moments thereof). Repeat for both preparations and estimate the mean shift $\hat{m}(t)$.
- (E4) Define $\hat{\tau}_{\text{copy}}(\eta)$ as the first time at which $|\hat{m}(t)|$ crosses the threshold corresponding to the chosen advantage surrogate (e.g., Eq. (44) in the Gaussian commuting model).

Shots. For any two-outcome (or bounded) measurement with outcomes in $[-1, 1]$, Hoeffding implies $n = \Theta(\delta^{-2} \log(1/\alpha))$ shots suffice to estimate $m(t)$ within $\pm\delta$ with failure probability α .

Readout noise. If readout flips a ± 1 outcome with probability p (independently), then the measured mean is attenuated by a factor $(1 - 2p)$, i.e. $m_{\text{meas}}(t) = (1 - 2p)m(t)$. This shifts the inferred crossing time to later values in a predictable way and can be corrected by a one-time calibration of p . In a multi-hypothesis setting with M candidates, a simple union-bound upgrade gives an additional $\log M$ factor in the shot budget for pairwise certification, so the binary problem remains the minimal operational primitive.

Scope. The above protocol provides a practical way to estimate a receiver-limited copy time from directly measured receiver data. Scaling statements are then assessed through the systematic size- and window-stability checks described in Sections 10 and 8.

Remark on “chemical potential tilts”.

Later we use the convenient parametrization $\rho_{\pm}(0) \propto e^{-\beta H \mp \epsilon Q_A}$ in linear-response calculations. At finite β this family is not strictly supported in A unless $[H, Q_A] = 0$; it should be regarded as an analytically tractable proxy for a locally prepared bias, with errors controlled by equilibrium clustering. Whenever strict support in A is required (for locality-limited bounds), we work with the local-channel preparation (1).

2.2. Elementary Bounds and Remarks

We record two standard inequalities used later. The first is a direct duality bound.

Lemma 1 (Standard tool: trace-norm duality). *This identity is standard in quantum hypothesis testing and is included only to fix notation. For any Hermitian X and any observable O with $\|O\|_{\infty} \leq 1$,*

$$|\text{Tr}(OX)| \leq \|X\|_1. \quad (7)$$

In particular,

$$\text{Adv}_{A \rightarrow B}^{\text{full}}(t) = \frac{1}{2} \sup_{\|O_B\|_{\infty} \leq 1} \text{Tr}[O_B(\rho_{+,B}(t) - \rho_{-,B}(t))]. \quad (8)$$

The second is Pinsker’s inequality. We state the required support condition explicitly.

Lemma 2 (Pinsker inequality with full-rank condition). *If ρ is full rank (so that $D(\sigma|\rho) < \infty$ for all σ), then*

$$D_{\text{tr}}(\sigma, \rho) \leq \sqrt{\frac{1}{2} D(\sigma|\rho)}. \quad (9)$$

If ρ is not full rank, the inequality remains valid provided $\text{supp}(\sigma) \subseteq \text{supp}(\rho)$; otherwise $D(\sigma|\rho) = \infty$ and (9) is vacuous.

In all numerical benchmarks reported below we work in finite-dimensional Hilbert spaces (finite spin chains) and, where thermal reference states are invoked, at finite inverse temperature $\beta < \infty$;

consequently the reduced density matrices encountered are effectively full rank (up to machine precision), and Lemma 2 is used only in this non-vacuous regime.

3. Minimal Locality Bounds and What Is Genuinely Nontrivial

3.1. Locality-Preserving Dynamics and Lieb–Robinson Constraints

Throughout, we assume a lattice with a metric $d(\cdot, \cdot)$, and we denote the distance between regions A, B by $d(A, B)$. For continuous-time local Hamiltonians or for locality-preserving circuits/QCAs, the standard consequence is a Lieb–Robinson-type bound: a local perturbation cannot influence distant observables faster than a finite velocity v_{LR} , up to exponentially small tails.

We deliberately *do not* conflate different settings. When we refer to LR bounds, we mean either (i) a continuous-time Hamiltonian/Lindbladian LR bound with constants (v, μ) or (ii) a circuit/QCA strict light cone. In each case, the structure is an upper bound on commutators of evolved local observables.

Theorem 1 (Locality-limited copying: Hamiltonian LR tails vs. circuit/QCA light cones). *Let $\mathcal{U}_t$ be the evolution channel (unitary or CPTP) on a lattice with metric $d(\cdot, \cdot)$. Assume one of the following standard locality structures for Heisenberg-evolved observables.*

(H) Continuous-time Hamiltonian/Lindbladian (Lieb–Robinson tails). *There exist constants $c, \mu, v > 0$ such that for any observables O_X supported on a finite region X and O_Y supported on Y ,*

$$\| [O_X(t), O_Y] \|_\infty \leq c |X| |Y| \|O_X\|_\infty \|O_Y\|_\infty e^{-\mu(d(X,Y)-v|t|)}. \quad (10)$$

(C) Discrete-time range- R circuits / reversible QCAs (strict light cone). *For each integer time $t \in \mathbb{N}$, the Heisenberg image of a local algebra satisfies*

$$\text{supp}(O_X(t)) \subseteq X^{+Rt} \quad \Rightarrow \quad [O_X(t), O_Y] = 0 \text{ whenever } d(X, Y) > Rt. \quad (11)$$

Let the initial pair $\rho_\pm$ differ only inside a sender region A and satisfy $\|\rho_+ - \rho_-\|_1 \leq 2$. Then for any receiver region B and any time t :

- (H) *For local Hamiltonian evolution satisfying a Lieb–Robinson bound, the receiver advantage is bounded by (A55) outside the effective cone.*
- (C) *For range- R circuits or reversible QCAs, the strict light cone implies $\text{Adv}_{A \rightarrow B}^{\text{full}}(t) = 0$ whenever $d(A, B) > Rt$.*

In particular, in both settings the copy time obeys the kinematic lower bound

$$\tau_{\text{copy}}^{\text{full}}(A \rightarrow B; \eta) \geq \begin{cases} \frac{d(A,B)}{v} - \frac{1}{\mu v} \log(C_{\text{LR}}/\eta), & \text{in (H),} \\ d(A, B)/R, & \text{in (C).} \end{cases} \quad (12)$$

Remark 1 (About constants in the Lieb–Robinson bound). *In case (H), one may take $C_{\text{LR}} := 2c |A| |B| \|\rho_+ - \rho_-\|_1$ for the LR prefactor c appearing in (A54) (up to harmless constants absorbed by the choice of norm convention), so that the only dependence on A, B is through the region-size factor $|A| |B|$. If the initial preparation is not strictly supported in A but has exponential Gibbs tails with rate κ (i.e. $\|\text{Tr}_{A^c} \delta \rho\|_1$ decays as $e^{-\kappa d}$), the same bound holds with an additional additive error term of order $e^{-\kappa(d(A,B)-v|t|)}$; equivalently one replaces μ by $\min\{\mu, \kappa\}$ in the logarithmic correction. In case (H), the prefactor C_{LR} comes from the standard LR “localization” estimate (A54) (denoted C_1 in the proof) and is independent of A, B, t except through bounded region-size factors. The only scaling content of (12) is therefore the kinematic term $d(A, B)/v$ and the mild logarithmic threshold correction.*

Corollary 1 (Explicit degradation for algebraic locality tails). *If the microscopic locality estimate is algebraic rather than exponential, namely*

$$\| [O_X(t), O_Y] \|_\infty \leq C_\alpha |X| |Y| \|O_X\|_\infty \|O_Y\|_\infty \frac{(1 + |t|)^p}{(1 + d(X, Y))^\alpha}, \quad (13)$$

with some exponent $\alpha > 0$ and model-dependent $p \geq 0$, then the same duality argument gives

$$\text{Adv}_{A \rightarrow B}^{\text{full}}(t) \leq C'_\alpha \frac{(1 + |t|)^p}{(1 + d(A, B))^\alpha}. \quad (14)$$

Consequently, any threshold η can only be attained at times satisfying

$$1 + |t| \geq \left(\frac{\eta}{C'_\alpha} \right)^{1/p} (1 + d(A, B))^{\alpha/p} \quad (p > 0), \quad (15)$$

and when $p = 0$ the bound reduces to the purely spatial algebraic condition $\eta \leq C'_\alpha (1 + d(A, B))^{-\alpha}$. Thus exponential LR tails and strict QCA cones are precisely the settings in which the locality lower bound retains an affine-in-distance operational form.

Proof. The first inequality is the same trace-norm duality step as in Theorem 1. The stated threshold conditions are obtained by solving the resulting algebraic inequality for t (or, when $p = 0$, for the spatial separation alone). $\square$

Remark 2 (Why there is no model-independent closed-system extension of the Q-SSEP benchmark). *A recurring temptation is to ask whether the Q-SSEP lower bound can be transferred to generic closed Hamiltonian systems using only the existence of a diffusing conserved density. Supplementary S7 proves that this is impossible in full generality: one can construct closed systems in which the conserved density sector is diffusive while an independent ancillary channel carries distinguishability to the receiver on a parametrically earlier timescale. Therefore no lower bound on unrestricted Helstrom copy time can follow from density diffusion data alone without additional assumptions that tie the receiver-optimal observable to the slow sector (for example, diagonality, commuting receiver reductions, Gaussian fluctuation control, or an explicit reduced-dynamics hypothesis). This is why the manuscript separates the exact Q-SSEP benchmark from the conditional Hamiltonian extension instead of blurring them into a single claim.*

Remark 3 (What Theorem 1 does not do). *Theorem 1 is an upper bound on how early copy can occur; it does not provide a mechanism for achieving η at times $t \sim d(A, B)/v$. For an A-level contribution, the hard part is a lower bound (or matching scaling) in physically relevant regimes. This is precisely where conservation laws and hydrodynamics enter.*

3.2. Conservation Laws and an Explicit Receiver Class

Fix a (quasi-)local conserved charge $Q = \sum_x q_x$ with $[\mathcal{U}_t, Q] = 0$. We will focus on initial perturbations that are “charge-biased” in A . Importantly, the receiver need not have access to all observables on B ; in realistic settings one often restricts to a class $\mathcal{O}_B$ (e.g., low moments of charge, or few-body observables).

To avoid circularity, we proceed in two steps: (i) we provide a theorem under *minimal* hypotheses that yields a general lower bound in terms of an explicitly defined susceptibility-like object, and (ii) we identify an additional *single-mode hydrodynamic window* under which that object becomes computable from transport data.

4. Hydrodynamic Closure: From Charge Bias to Spectral Susceptibility

Table 1. Notation used in the hydrodynamic closure section.

Symbol	Meaning
$\mathcal{L}$	Liouvillian superoperator ($\mathcal{L}(O) = i[H, O]$ for closed systems)
$\mathcal{P}, \mathcal{Q}$	Mori projection onto the slow manifold and its complement
$\langle \cdot, \cdot \rangle_\beta$	Kubo–Mori inner product induced by the reference state ρ_β
$\mathcal{L}_{\text{eff}}^{(2)}$	Markovianized effective slow-sector generator (Eq. (18))
$\chi_{A \rightarrow B}^{(2)}$	Second-moment spectral susceptibility (Definition 2)
J_B	Receiver observable restricted to the slow sector
ϵ_{mem}	Quantitative reduced-dynamics error controlling the semigroup approximation on the chosen window

4.1. Setup: Local Equilibrium Manifold and Linearization

Let ρ_β be a reference Gibbs (or generalized Gibbs) state at inverse temperature β . Consider a small sender perturbation of “chemical potential” type,

$$\rho_\pm(0) = \frac{e^{-\beta H \mp \epsilon Q_A}}{\text{Tr}(e^{-\beta H \mp \epsilon Q_A})}, \quad (16)$$

where $Q_A = \sum_{x \in A} q_x$ and $\epsilon \ll 1$. To leading order in ϵ , the difference $\delta\rho(0) := \rho_+(0) - \rho_-(0)$ is linear in Q_A in the Kubo–Mori inner product.

4.2. A Principled Definition of the Second-Moment Susceptibility

Let $\mathcal{L}$ denote the (Heisenberg) Liouvillian superoperator $\mathcal{L}(O) = i[H, O]$ (or the generator of a CPTP semigroup in open settings). Let $\mathcal{P}$ be a projection onto the slow subspace spanned by the conserved density modes; concretely, $\mathcal{P}$ is the Mori projection with respect to the Kubo–Mori inner product. Define the reduced (effective) slow-sector operator using the Mori–Zwanzig projection formalism [4,5]. In Laplace space one obtains an exact identity for the projected dynamics,

$$(z - \mathcal{L}_{\text{eff}}(z))^{-1} = \mathcal{P}(z - \mathcal{L})^{-1}\mathcal{P}, \quad \mathcal{L}_{\text{eff}}(z) := \mathcal{P}\mathcal{L}\mathcal{P} + \mathcal{P}\mathcal{L}\mathcal{Q}(z - \mathcal{Q}\mathcal{L}\mathcal{Q})^{-1}\mathcal{Q}\mathcal{L}\mathcal{P}. \quad (17)$$

This expression is exact but generally nonlocal in time (the z -dependence encodes a memory kernel). In *closed unitary* dynamics, $\mathcal{L}(O) = i[H, O]$ is anti-self-adjoint and $\mathcal{L}_{\text{eff}}(z)$ is *not* a CPTP generator in general; it should be interpreted as a reduced linear-response operator in the Kubo–Mori geometry. In contrast, for genuinely open Markovian dynamics (Lindbladians) the same construction yields a bona fide dissipative generator on the slow manifold.

To obtain a usable closure we do *not* impose a dissipative spectral gap on the exact microscopic unitary Liouvillian. Instead, we separate the exact projected identity (17) from the additional assumption needed to replace the non-Markovian memory kernel by a time-local reduced generator. We therefore introduce the low-frequency reduced generator

$$\mathcal{L}_{\text{eff}} := \mathcal{L}_{\text{eff}}(0) = \mathcal{P}\mathcal{L}\mathcal{P} + \mathcal{P}\mathcal{L}\mathcal{Q}(-\mathcal{Q}\mathcal{L}\mathcal{Q})^{-1}\mathcal{Q}\mathcal{L}\mathcal{P}, \quad (18)$$

only as the candidate generator of a reduced slow semigroup, and we explicitly quantify below the error incurred by replacing the exact projected propagator by $e^{t\mathcal{L}_{\text{eff}}}\mathcal{P}$.

Proposition 2 (Microscopic no-go for a positive fast-sector quadratic-form gap in closed unitary dynamics). *Let $\mathcal{L}(O) = i[H, O]$ be the exact Liouvillian of a finite-dimensional closed Hamiltonian system and*

let $\mathcal{P}$ be the orthogonal Mori projection in the Kubo–Mori inner product, with $\mathcal{Q} = \mathbb{I} - \mathcal{P}$. Then the fast-sector block $\mathcal{Q}\mathcal{L}\mathcal{Q}$ is anti-self-adjoint on $\text{Ran}(\mathcal{Q})$. Consequently,

$$\langle X, -\mathcal{Q}\mathcal{L}\mathcal{Q}X \rangle_\beta = 0 \quad \text{for every } X \in \text{Ran}(\mathcal{Q}), \tag{19}$$

and therefore no strictly positive inequality of the form

$$\langle X, -\mathcal{Q}\mathcal{L}\mathcal{Q}X \rangle_\beta \geq \Delta_{\text{fast}} \|X\|_\beta^2 \tag{20}$$

can hold for any $\Delta_{\text{fast}} > 0$ on the exact microscopic unitary problem.

Proof. The Kubo–Mori inner product satisfies $\langle X, \mathcal{L}Y \rangle_\beta = -\langle \mathcal{L}X, Y \rangle_\beta$ for $\mathcal{L}(O) = i[H, O]$, so $\mathcal{L}$ is anti-self-adjoint. Because $\mathcal{P}$ is orthogonal, $\mathcal{Q}$ is self-adjoint and idempotent in the same geometry. Hence

$$\langle X, \mathcal{Q}\mathcal{L}\mathcal{Q}Y \rangle_\beta = \langle \mathcal{Q}X, \mathcal{L}\mathcal{Q}Y \rangle_\beta = -\langle \mathcal{L}\mathcal{Q}X, \mathcal{Q}Y \rangle_\beta = -\langle \mathcal{Q}\mathcal{L}\mathcal{Q}X, Y \rangle_\beta$$

for $X, Y \in \text{Ran}(\mathcal{Q})$, proving anti-self-adjointness of $\mathcal{Q}\mathcal{L}\mathcal{Q}$ on the fast subspace. Setting $Y = X$ gives (19), because the expectation of an anti-self-adjoint operator on a vector is purely imaginary while the left-hand side is real. $\square$

Proposition 2 is not merely a warning sign; it tells us exactly what a mathematically coherent closure must look like. For closed Hamiltonian dynamics the extra ingredient cannot be a dissipative microscopic gap on $\mathcal{Q}\mathcal{L}\mathcal{Q}$. It must instead be a statement about the memory kernel or, equivalently, about the smallness of its low-frequency variation on the time window where the transport fit is performed.

Controlled hydrodynamic regime.

The reduced-dynamics closure is intended for chaotic, finite-temperature phases in which a small set of conserved slow modes is cleanly separated from short-memory fast fluctuations on the time window of interest. The relevant ingredients are explicit: an isolated slow pole, a nonvanishing receiver overlap, and a low-frequency memory kernel whose variation is small across the fitted spectral window. Integrable or near-integrable regimes, many-body-localized phases, long-time-tail dominated windows, or quasi-conserved prethermal plateaus require a modified analysis because more than one slow structure may remain active on the same time scale. Theorems 2 and 3 are therefore formulated directly around this controlled regime and paired with diagnostic proxies that identify it in practice.

Table 2. Hydrodynamic reduction checklist: assumptions, concrete diagnostic proxies, and characteristic breakdown mechanisms for the single-mode transport window discussed in this section.

Assumption	Operational/diagnostic proxy	Typical failure modes
Short-memory kernel / fast correlation decay	Rapid decay of generic local autocorrelations to their hydrodynamic tail; small window-to-window variation of the fitted reduced kernel; absence of long-lived nonconserved operators in the accessible ED window (finite-size diagnostic)	Integrability, quasi-conservation (prethermal plateaus), MBL
Single isolated slow pole	$\Gamma(k)$ approximately linear in k^2 over a time window; window-to-window stability; no competing ballistic/Drude channel (Sec. 9, Supplementary S2)	Multiple slow modes, long-time tails, Drude weight/nonzero stiffness
Nonzero receiver overlap	Choose J_B with provable overlap (e.g., coarse-grained charge in B); verify signal is nonzero at accessible times	Symmetry mismatch; receiver observable orthogonal to slow mode
Linear-response regime	Small ϵ tilt; check odd-in- ϵ scaling and absence of saturation artifacts in numerics	Large perturbations, finite-size saturation, edge effects

Definition 2 (Second-moment spectral susceptibility). [3–5] Let $\delta\mu$ be a small chemical-potential profile supported in A , and let J_B denote a receiver observable (e.g., the charge in B or its low moments) projected to the slow subspace. We define the second-moment spectral susceptibility by

$$\chi_{A \rightarrow B}^{(2)} := \left\langle J_B, (-\mathcal{L}_{\text{eff}})^{-2} J_B \right\rangle_{\beta} \|\delta\mu\|^2, \quad (21)$$

where $\langle \cdot, \cdot \rangle_{\beta}$ is the Kubo–Mori inner product and $(-\mathcal{L}_{\text{eff}})^{-2}$ is understood on the orthogonal complement of the zero mode.

Remark 4 (Why (21) is the right object). The operator $(-\mathcal{L}_{\text{eff}})^{-1}$ is a resolvent that weights modes by inverse decay rate; squaring it yields a second moment that controls the time at which a receiver observable can accumulate a finite signal. Definition 2 makes explicit (a) the operator being inverted, (b) the topology (Kubo–Mori norm) in which neglected modes are controlled, and (c) the receiver class through J_B .

Assumption 1 (Short-memory semigroup reduction on the slow window). Let $U_{\mathcal{P}}(t) := \mathcal{P}e^{t\mathcal{L}}\mathcal{P}$ denote the exact projected propagator on the slow subspace. We assume that on a specified time window $t \in [t_{\min}, t_{\max}]$ there exists a bounded reduced generator $\mathcal{L}_{\text{eff}}$ such that

$$\sup_{t \in [t_{\min}, t_{\max}]} \|U_{\mathcal{P}}(t) - e^{t\mathcal{L}_{\text{eff}}}\mathcal{P}\|_{\beta \rightarrow \beta} \leq \varepsilon_{\text{mem}}. \quad (22)$$

This assumption is the only place where a Markovianized closure enters the unitary Hamiltonian analysis.

Proposition 3 (Resolvent criterion for the reduced semigroup approximation). Let $K(t)$ denote the exact Mori memory kernel on the slow subspace, with Laplace transform $\hat{K}(z)$, and let Σ be a contour segment in the half-plane $\Re z \geq 0$ resolving the frequencies relevant to the fitted transport window. Assume that the resolvents of the exact reduced generator and of its low-frequency approximation are uniformly bounded on Σ , and that

$$\sup_{z \in \Sigma} \|\hat{K}(z) - \hat{K}(0)\|_{\beta \rightarrow \beta} \leq \delta_{\Sigma}. \quad (23)$$

Then there exists a contour constant $C_{\Sigma} > 0$ such that

$$\sup_{t \in [t_{\min}, t_{\max}]} \|U_{\mathcal{P}}(t) - e^{t\mathcal{L}_{\text{eff}}}\mathcal{P}\|_{\beta \rightarrow \beta} \leq C_{\Sigma} \delta_{\Sigma} =: \varepsilon_{\text{mem}}. \quad (24)$$

In particular, a nearly frequency-flat memory kernel on the low-frequency window implies the semigroup approximation (22) with an explicit error parameter.

Proof. The exact projected Laplace-space propagator is $(z - \mathcal{L}_{\text{eff}}(z))^{-1}\mathcal{P}$, while the Markovianized approximation is $(z - \mathcal{L}_{\text{eff}})^{-1}\mathcal{P}$. The resolvent identity gives

$$(z - \mathcal{L}_{\text{eff}}(z))^{-1} - (z - \mathcal{L}_{\text{eff}})^{-1} = (z - \mathcal{L}_{\text{eff}}(z))^{-1}(\hat{K}(z) - \hat{K}(0))(z - \mathcal{L}_{\text{eff}})^{-1}, \quad (25)$$

so the assumed resolvent bounds and (23) yield a uniform Laplace-space error $O(\delta_{\Sigma})$ on Σ . Inverting the Laplace transform on that contour gives (24). The point is not a dissipative microscopic gap, but a quantitatively small low-frequency memory variation on the slow window. $\square$

4.3. Two Theorems: Minimal and Single-Mode

We now state two versions of the central result: a general projection theorem for the reduced operational signal and a conditional single-mode hydrodynamic specialization.

Theorem 2 (Projection-operator lower bound on receiver advantage with explicit reduced-dynamics error). Assume (i) a finite-dimensional lattice system with a full-rank reference state ρ_{β} , (ii) a weak charge-

biased perturbation of the form (16) with $\epsilon \ll 1$, and (iii) a Mori projection $\mathcal{P}$ onto the slow manifold such that the Markovianized effective operator $\mathcal{L}_{\text{eff}}$ in (18) is well-defined on $\text{Ran}(\mathcal{P})$. Let $J_B \in \text{Ran}(\mathcal{P})$ be a receiver observable supported in B with $\|J_B\|_\infty \leq 1$. Then, for all times $t \geq 0$,

$$\text{Adv}_{A \rightarrow B}^{\mathcal{O}}(t) \geq \frac{\epsilon}{2} |\langle J_B, e^{t\mathcal{L}_{\text{eff}}} \mathcal{P}(Q_A) \rangle_\beta| - \frac{\epsilon}{2} |\langle J_B, Q e^{t\mathcal{L}} \mathcal{P}(Q_A) \rangle_\beta| + O(\epsilon^3). \quad (26)$$

If, in addition, the short-memory assumption (22) holds with error ϵ_{mem} , then the neglected reduced-dynamics mismatch admits the explicit bound

$$|\langle J_B, [U_{\mathcal{P}}(t) - e^{t\mathcal{L}_{\text{eff}}} \mathcal{P}] Q_A \rangle_\beta| \leq \|J_B\|_\beta \|Q_A\|_\beta \epsilon_{\text{mem}}. \quad (27)$$

Thus, on any time window where $\epsilon_{\text{mem}} \ll 1$, the receiver advantage is governed (up to an explicit short-memory error) by the reduced slow dynamics, without assuming a priori that the receiver “sees the slow mode”.

Corollary 2 (Explicit slow-signal minus reduced-dynamics error bound). *Under the assumptions of Theorem 2 and Assumption 1, define*

$$S_{A \rightarrow B}(t) := |\langle J_B, e^{t\mathcal{L}_{\text{eff}}} \mathcal{P}(Q_A) \rangle_\beta|, \quad C_{\text{mem}} := \|J_B\|_\beta \|Q_A\|_\beta. \quad (28)$$

Then, for all $t \geq 0$,

$$\text{Adv}_{A \rightarrow B}^{\mathcal{O}}(t) \geq \frac{\epsilon}{2} S_{A \rightarrow B}(t) - \frac{\epsilon}{2} C_{\text{mem}} \epsilon_{\text{mem}} + O(\epsilon^3). \quad (29)$$

In particular, on any window where the reduced semigroup approximation is accurate, the advantage is controlled by the explicit slow-sector correlator $S_{A \rightarrow B}(t)$ up to a quantified short-memory error.

Theorem 3 (Single-mode hydrodynamic window under explicit hypotheses). *Assume the hypotheses of Theorem 2 together with:*

- (S1)**Single slow pole:** on a wavelength band $k \in [k_{\min}, k_{\max}]$ the slow spectrum on $\text{Ran}(\mathcal{P})$ consists of a single nonzero mode with decay rate $\Gamma(k) = Dk^2 + o(k^2)$ and a gap Δ_2 to the next slow mode on that band;
- (S2)**Receiver overlap:** the projected receiver observable has nonzero overlap with that mode, quantified by the form factor $F_B(k)$ entering (A56);
- (S3)**Short memory on the slow window:** the projected propagator is approximated by $e^{t\mathcal{L}_{\text{eff}}} \mathcal{P}$ with error parameter ϵ_{mem} on the window of interest, as in Assumption 1 and Proposition 3.

Then, for a one-dimensional sender–receiver separation $\ell = d(A, B)$ and for any threshold prescription that remains within a fixed fraction of the maximal diffusive signal (see Proposition 4), the copy time satisfies the transport-limited scaling

$$\tau_{\text{copy}}^{\mathcal{O}}(A \rightarrow B; \eta) = \frac{\ell^2}{4D} (1 + o(1)) + O(t_0 \log(C/\eta)), \quad (30)$$

where $t_0 := 1/(Dk_{\max}^2)$ is the microscopic time associated with the upper cutoff $k_{\max}$ of the single-mode window and $C > 0$ collects order-one receiver/form-factor constants (explicit in Appendix B). In the hydrodynamic window $t_{\min} \ll t \ll \min\{L^2/D, \Delta_2^{-1}\}$, with explicitly bounded systematic errors from: (i) finite-size discretization of k (Appendix E), (ii) slow-sector multi-mode contamination $O(e^{-\Delta_2 t})$, and (iii) reduced-dynamics mismatch $O(\epsilon_{\text{mem}})$ via (27). The prefactor can be expressed in terms of the susceptibility $\chi_{A \rightarrow B}^{(2)}$ in Definition 2.

The appendices and Supplementary S6 provide explicit derivations of the reduced-dynamics bounds used here, including a microscopic Duhamel expansion for the locally biased Gibbs family and the no-go statement of Proposition 2. The key point is that the nontrivial input is *not* the statement “the receiver sees the slow mode”; rather it is the explicit control of the reduced-dynamics error and the explicit coupling coefficient between Q_A and J_B in the projected dynamics.

5. Worked Hydrodynamic Example: One-Dimensional Diffusion Kernel

To make Theorem 3 concrete, we work out the diffusion-kernel signal at the level needed to turn “diffusion implies $\tau \sim \ell^2$ ” into a quantitatively checkable statement.

5.1. Linear-Response Form of the Reduced-State Difference

Write $\rho_{\pm}(t) = \mathcal{U}_t(\rho_{\pm}(0))$ with $\rho_{\pm}(0)$ given by (16). Expanding to first order in ϵ yields

$$\delta\rho(t) := \rho_+(t) - \rho_-(t) = \epsilon \mathcal{U}_t(\delta\rho_{Q_A}) + O(\epsilon^3), \quad (31)$$

where $\delta\rho_{Q_A}$ is the Kubo–Mori tangent vector associated to Q_A at ρ_{β} . For any receiver observable O_B with $\|O_B\|_{\infty} \leq 1$,

$$\text{Tr}(O_B \delta\rho_B(t)) = \text{Tr}(O_B \text{Tr}_{B^c} \delta\rho(t)) = \text{Tr}(O_B(t) \delta\rho_{Q_A}) + O(\epsilon^3). \quad (32)$$

Thus the operational advantage is lower-bounded by any explicit choice of O_B .

5.2. Diffusion Equation for the Conserved Density

Assume a single conserved density $q(x)$ on a ring of length L with hydrodynamic equation

$$\partial_t \delta q(x, t) = D \partial_x^2 \delta q(x, t), \quad (33)$$

valid on the window $t \in [t_0, t_1]$ and wavelengths $\lambda \gg 1$ (in lattice units). On the ring, the Fourier modes $k_n = \frac{2\pi n}{L}$ evolve as

$$\delta q_{k_n}(t) = e^{-Dk_n^2 t} \delta q_{k_n}(0). \quad (34)$$

For a localized initial bias in A , $\delta q(x, 0)$ has broad Fourier support but its long-time profile is controlled by the smallest nonzero $|k_n|$.

5.3. Receiver Signal and a Concrete Threshold-to-Time Relation

Let $Q_B(t) = \int_B q(x, t) dx$ be the receiver charge in an interval B centered at distance ℓ from A with width w . In the linear-response regime around ρ_{β} , the expected receiver shift is proportional to the chemical potential profile with coefficient given by the static susceptibility $\chi = \partial\langle q \rangle / \partial\mu$:

$$\delta\langle Q_B(t) \rangle \approx \chi \int_B \delta\mu(x, t) dx. \quad (35)$$

Approximating $\delta\mu$ by the Gaussian heat kernel on $\mathbb{R}$ for times $t \ll L^2/D$,

$$\delta\mu(x, t) \approx \frac{\epsilon}{\sqrt{4\pi Dt}} \exp\left(-\frac{(x - \ell)^2}{4Dt}\right), \quad (36)$$

we obtain the scaling

$$\delta\langle Q_B(t) \rangle \sim \epsilon \chi w (4\pi Dt)^{-1/2} \exp\left(-\frac{\ell^2}{4Dt}\right), \quad (37)$$

up to relative corrections of order $w^2/(Dt)$ and tD/L^2 . Choosing O_B proportional to the centered receiver charge (normalized to $\|O_B\|_{\infty} \leq 1$) and using Lemma 1 gives a lower bound

$$\text{Adv}_{A \rightarrow B}^{\text{full}}(t) \geq \frac{1}{2} |\text{Tr}(O_B \delta\rho_B(t))| \gtrsim c_0 |\delta\langle Q_B(t) \rangle|, \quad (38)$$

where c_0 is an explicit normalization constant. The kernel (37) already shows why a fixed nonzero threshold is not a meaningful asymptotic statement at fixed $|B|$ and fixed perturbation amplitude: its maximal value decays with distance. Indeed, writing

$$S_\ell(t) := K_B t^{-1/2} \exp\left(-\frac{\ell^2}{4Dt}\right), \quad K_B \propto |\epsilon| \chi |B|, \quad (39)$$

one finds

$$\arg \max_t S_\ell(t) = \frac{\ell^2}{2D}, \quad \max_t S_\ell(t) = \frac{K_B \sqrt{2D/e}}{\ell}. \quad (40)$$

Thus, for any fixed threshold $\eta > 0$, threshold crossing eventually ceases to exist unless K_B itself scales with distance.

Proposition 4 (Well-posed threshold prescriptions for diffusive copy times). *Within the Gaussian-kernel approximation (37), let $S_\ell(t)$ denote the receiver signal. Then:*

- (i) *for fixed $|B|$, fixed $|\epsilon|$, and fixed nonzero threshold η , the hitting time exists only while $\eta \leq K_B \sqrt{2D/e}/\ell$;*
- (ii) *if the threshold is prescribed as a fixed fraction of the distance-dependent peak, namely $\eta_\ell = \vartheta \max_t S_\ell(t)$ with $0 < \vartheta < 1$, then*

$$\tau_{\text{copy}}(A \rightarrow B; \eta_\ell) = c_\vartheta \frac{\ell^2}{D}, \quad (41)$$

with $c_\vartheta > 0$ determined implicitly by $u^{-1/2} e^{-1/(4u)} = \vartheta \sqrt{2/e}$ for $u = tD/\ell^2$;

- (iii) *equivalently, the same quadratic law follows at fixed threshold if either the effective receiver size or the perturbation amplitude is scaled so that K_B/ℓ remains $O(1)$.*

Proof. Differentiate $\log S_\ell(t)$ to find the unique stationary point at $t_* = \ell^2/(2D)$. Since $S_\ell(t) \rightarrow 0$ as $t \downarrow 0$ and as $t \rightarrow \infty$, this stationary point is the global maximum, giving $\max_t S_\ell(t) = K_B \sqrt{2D/e}/\ell$ and hence part (i). For part (ii), substitute $t = u\ell^2/D$; then $S_\ell(t) = \frac{K_B}{\ell} u^{-1/2} e^{-1/(4u)}$. Setting $S_\ell(t) = \eta_\ell = \vartheta \max_t S_\ell(t)$ yields the implicit equation stated for $u = c_\vartheta$, which is independent of ℓ . Part (iii) is immediate from the same rescaled form because keeping K_B/ℓ of order one keeps the admissibility window nontrivial at fixed threshold. $\square$

In other words, the quadratic transport law is an operational statement only after the admissible-threshold convention is specified together with distance. This paper therefore treats fixed-threshold plots as finite-size diagnostics, while all asymptotic hydrodynamic statements are phrased in threshold-scaled regimes such as (41).

A quantitative η - and $|B|$ -dependence (Gaussian-kernel inversion).

Whenever the threshold is admissible, the threshold condition has the form

$$K_B t^{-1/2} \exp\left(-\frac{\ell^2}{4Dt}\right) = \eta_{\text{eff}}, \quad (42)$$

where η_{eff} is the effective threshold after mapping a chosen measurement to an advantage surrogate (e.g., in the commuting Gaussian model below, $\eta_{\text{eff}} \propto \text{erf}^{-1}(\eta)$). Equation (42) can be inverted in closed form using the Lambert-W function:

$$\tau_{\text{copy}}(A \rightarrow B; \eta) \approx \frac{\ell^2}{2D} \left[-W_{-1}\left(-c \frac{\ell^2 \eta_{\text{eff}}^2}{K_B^2}\right) \right]^{-1}, \quad (43)$$

on the physical (W_{-1}) branch. This expression makes explicit that larger receivers and larger sender amplitudes enlarge the admissible-threshold window, while fixed-threshold asymptotics fail once ℓ exceeds the inverse peak scale. We emphasize that Eq. (43) is a controlled prediction only to the extent that the Gaussian-kernel approximation is valid on the fitted window; Appendix F provides a

small-system ED illustration of the smooth η and $|B|$ dependence without claiming thermodynamic scaling. A finite-volume version based on (34) and the lowest nonzero mode $k_{\min} = 2\pi/L$ yields the crossover from Gaussian kernel behavior to the ring's exponential mode decay.

5.4. An Exactly Solvable Gaussian Diffusion Toy Model (Fully Analytic)

To remove any ambiguity about “which observable is optimal” and what constants control the threshold, we include a reference model in which the Helstrom measurement can be written in closed form. Consider commuting hypotheses in which the only receiver-relevant variable is the coarse-grained charge Q_B and, conditional on either hypothesis, Q_B is Gaussian with means $\pm m(t)$ and a common variance σ_B^2 (equilibrium charge fluctuations in B). In this setting the Helstrom optimum reduces to classical hypothesis testing on Q_B and the (optimal) advantage equals the classical total-variation distance (i.e., $\|p - q\|_{\text{TV}} := \frac{1}{2}\|p - q\|_1$) between two Gaussians,

$$\text{Adv}_{A \rightarrow B}^{\text{full}}(t) = \|\mathcal{N}(m(t), \sigma_B^2) - \mathcal{N}(-m(t), \sigma_B^2)\|_{\text{TV}} = \text{erf}\left(\frac{|m(t)|}{\sqrt{2}\sigma_B}\right). \quad (44)$$

For a diffusive signal $m(t) \propto \int_B G(x, t) dx$ (with heat kernel G), Eq. (44) yields an explicit closed-form definition of τ_{copy} at threshold η :

$$\tau_{\text{copy}}(A \rightarrow B; \eta) = \inf \left\{ t > 0 : |m(t)| \geq \sqrt{2}\sigma_B \text{erf}^{-1}(\eta) \right\}. \quad (45)$$

Appendix G gives the full derivation and shows how this “commuting” formula interfaces with the Kubo–Mori linear-response bounds and the moment-channel optimality statements.

6. Controlled Diffusive Benchmark: Helstrom Copy Time in Q-SSEP

The hydrodynamic sections below describe how diffusion-limited copying can emerge in *unitary* dynamics under explicitly stated hypotheses. The operational framework becomes especially compelling when it can be anchored in a microscopically explicit benchmark. We therefore turn to the *quantum symmetric simple exclusion process* (Q-SSEP) [9], in which a diagonal locally biased state family reduces the receiver problem to a classical exclusion-process calculation without any appeal to unitary short-memory assumptions.

6.1. Model and Hypotheses

Q-SSEP is a local, translation-invariant Lindbladian dynamics that preserves a $U(1)$ charge (occupation number) and induces exact diffusion of that conserved density. A key structural feature is that a large family of states remains diagonal in the occupation basis, and within this invariant manifold the evolution of diagonal weights reduces to the classical symmetric simple exclusion process (SSEP). This provides an exactly controlled setting in which transport coefficients and hypothesis-testing advantages can be related sharply.

We consider locally prepared hypotheses of the form (1), chosen so that (to leading order in a small bias parameter ϵ) the initial difference is a localized charge perturbation supported in A . The receiver B is taken at distance $\ell = d(A, B)$. We denote by D the diffusion constant governing the macroscopic scaling of the conserved density and by χ the corresponding static susceptibility.

6.2. Main Inequality

The following theorem is the controlled benchmark result. It is stated for the Helstrom advantage (unrestricted measurements on B), but only for the explicit diagonal family of hypotheses described above.

Theorem 4 (Diagonal-sector diffusive Helstrom bound in Q-SSEP). *Consider Q-SSEP dynamics in one dimension at inverse temperature $\beta < \infty$ (including $\beta = 0$), and restrict to the family of locally prepared*

diagonal hypotheses for which the reduced states on every finite receiver region remain diagonal in the occupation basis and differ, to leading order in ϵ , only by a localized charge bias supported in A . Then there exist constants $C_0, C_1 > 0$ such that, to leading order in the bias parameter ϵ (with corrections $O(\epsilon^2)$ uniform on finite B), for any receiver region B at distance $\ell = d(A, B)$ and all $t > 0$,

$$\text{Adv}_{A \rightarrow B}^{\text{full}}(t) \leq C_0 |\epsilon| \chi |B| \frac{1}{\sqrt{Dt}} \exp\left(-\frac{\ell^2}{C_1 Dt}\right) + O(\epsilon^2). \quad (46)$$

Consequently, for any threshold η not exceeding the maximal attainable advantage, the copy time obeys the lower bound

$$\tau_{\text{copy}}(A \rightarrow B; \eta) \geq \frac{\ell^2}{C_1 D} \left/ \log\left(\frac{C_0 |\epsilon| \chi |B|}{\eta}\right) \right., \quad (47)$$

and in particular $\tau_{\text{copy}}^{\text{full}}(A \rightarrow B; \eta) \geq c(\eta, \epsilon, \chi, |B|) \ell^2 / D$.

Remark 5 (Why this matters). *Theorem 4 is not a hydrodynamic closure for generic closed systems: it is a controlled operational inequality inside a fully specified microscopic state family. Because the receiver states remain diagonal on B , unrestricted Helstrom discrimination reduces exactly to classical total variation on the full configuration distribution of B , so multi-site correlations are already included rather than discarded. The argument therefore avoids the problematic unitary-gap assumption of Section 4 while remaining narrower than a generic many-body theorem.*

Supplementary S3 now states every assumption of this diagonal-sector reduction explicitly, proves the exact Helstrom-to-total-variation reduction on the diagonal receiver sector, and isolates the sole linear-response remainder as an $O(\epsilon^2)$ term.




Figure 1. Schematic illustration of the bound (46) in the Q-SSEP benchmark: for fixed separation ℓ , the advantage is exponentially suppressed until diffusive time scales $t \sim \ell^2 / D$ become available.

7. Moment-Channel Approximation and Operational Accessibility

The definition (3) involves full state discrimination on B , which is optimal but may be experimentally unrealistic. A practical approach is to restrict the receiver to a moment family (e.g., charge moments). This restriction should be described as an explicit map.

7.1. Definition of the Moment Channel

Let $\mathcal{H}_B$ be the receiver Hilbert space. Fix an observable family $\mathcal{F} = \{f_1, \dots, f_m\}$ on B (e.g., $q_B, q_B^2, \dots$) and define the *moment channel*

$$\Phi_{\mathcal{F}}(\rho_B) := (\text{Tr}(\rho_B f_1), \dots, \text{Tr}(\rho_B f_m)) \in \mathbb{R}^m. \quad (48)$$

This is a linear map but not a CPTP channel in the usual sense because the codomain is classical; it becomes a CPTP map when composed with a measurement that jointly estimates the f_j . Operationally, restricting to $\Phi_{\mathcal{F}}$ yields a lower bound on $\text{Adv}_{A \rightarrow B}^{\text{full}}$ via Lemma 1:

$$\text{Adv}_{A \rightarrow B}^{\text{full}}(t) \geq \frac{1}{2} \sup_{\alpha: \|\sum_j \alpha_j f_j\|_{\infty} \leq 1} \sum_{j=1}^m \alpha_j (\text{Tr}(\rho_{+,B}(t) f_j) - \text{Tr}(\rho_{-,B}(t) f_j)). \quad (49)$$

7.2. When Moment Restriction Is Asymptotically Optimal

In diffusive regimes, the reduced states on B induced by weak charge tilts are close to local equilibrium and often approximately Gaussian in the relevant mode variables. In this case, low moments can be asymptotically sufficient statistics. We make this precise by stating an *explicit* assumption (Gaussianity in a fluctuation algebra) and deriving a matching upper bound in Appendix C.

8. Copy Time Versus OTOCs and Lieb–Robinson Bounds: A Finite-Window Separation Example




Figure 2. An illustrative operational separation with real data (small-system exact diagonalization). Left: commutator-squared OTOC proxy $C_{\text{otoc}}(\ell, t)$ (blue) and Helstrom advantage $\text{Adv}_{\text{full}}(B = \{\ell\}, t)$ for copying a *conserved* bias (orange), both normalized by their respective thresholds C_* and η . At distance $\ell = 3$ the OTOC crosses its threshold at $t \approx 2.2$, while the optimal-copy advantage remains *below* $\eta = 0.02$ throughout the displayed window $t \leq 4$, implying $\tau_{\text{copy}}(\ell = 3; \eta) > 4$. Right: extracted arrival times from the same run; OTOC times are consistent with a ballistic guide line ($\propto \ell$) while copy times can be parametrically delayed (lower-bound markers “ > 4 ”). Parameters: periodic XXZ chain with staggered field (U(1) conserved, integrability broken), $L = 8$, $\Delta = 1$, $h_{\text{stag}} = 0.5$, and the charge-bias hypotheses of Eq. (16) in the infinite-temperature limit. All raw outputs and the generating scripts are provided in the accompanying Supplementary Materials and Code Archive (see the accompanying Supplementary Materials and Code Archive).

Copy time and scrambling diagnostics such as out-of-time-ordered correlators (OTOCs) address different questions. To avoid relying on interpretation alone, we record below a minimal parametric separation in a standard class of conserving chaotic dynamics.

8.1. Ballistic Operator Growth Does Not Imply Fast Copying Under Conservation

In generic local dynamics, operator support typically spreads ballistically with a butterfly velocity, and OTOCs detect the resulting front [10–12]. However, when the sender perturbation is constrained to lie in a conserved sector (as in (16)), the receiver’s ability to *certify* that perturbation is controlled by the transport of the conserved density, which can be diffusive even when operator growth is ballistic.

This separation is explicit in the now-standard picture of *dissipative hydrodynamics emerging under unitary dynamics with conservation laws* [13,14]. In such systems, the OTOC front can be ballistic while the conserved mode relaxes diffusively; our Theorem 3 precisely predicts the resulting copy-time scaling for charge-biased hypotheses.

Proposition 5 (Parametric separation in conserving chaotic dynamics). *Consider a one-dimensional local unitary dynamics with a conserved $U(1)$ charge and with chaotic (mixing) dynamics in all other operator sectors. Assume: (i) operator support spreads ballistically with velocity v_B (as diagnosed by OTOCs), and (ii) the conserved density exhibits diffusion with coefficient D on the relevant window. Then for sender/receiver separation ℓ and fixed threshold $\eta \in (0, 1)$,*

$$\tau_{\text{OTOC}}(\ell) \sim \ell/v_B, \quad \tau_{\text{copy}}(A \rightarrow B; \eta) \sim \ell^2/(4D), \quad (50)$$

provided the sender perturbation is a weak charge bias and the receiver is restricted to physically accessible (e.g., few-body or moment) observables. Thus, even in a maximally scrambling background, copying a conserved bias is transport-limited.

Proof sketch.

The OTOC timescale $\tau_{\text{OTOC}}(\ell)$ is governed by the ballistic spreading of generic local operators, as captured by the butterfly velocity v_B [10,11]. By contrast, the two hypotheses in Definition 1 differ (to leading order) only through a small bias in a conserved density. Linear response therefore reduces the receiver signal to a hydrodynamic correlator in the slow sector (Theorem 2), and in a single-mode diffusive window it takes the heat-kernel form $\sim \exp\{-\ell^2/(4Dt)\}$ (Section 5). Inverting this threshold condition yields $\tau_{\text{copy}} \sim \ell^2/(4D)$ up to logarithmic η -dependent corrections, made explicit (with leakage/error terms) in Theorem 3.

Proposition 5 is not a new theorem of operator growth; it is a *clean operational interpretation*: OTOCs probe the fastest operator sector, while copy time probes the slowest sector that actually carries the hypothesis difference to the receiver.

8.2. Relation to LR Bounds

Lieb–Robinson bounds control the earliest possible influence outside an effective light cone (Theorem 1), but they do not determine the *dominant* timescale when a conservation law forces information to flow through a slow hydrodynamic channel. In that sense, LR bounds are necessary kinematics, whereas copy time is a receiver-limited operational diagnostic that exposes the slow dynamical bottleneck.

9. Failure Modes and Boundaries of Validity

A high-standard submission must include explicit boundaries. We summarize the main failure modes and what replaces Theorem 3:

- **Integrable / near-integrable dynamics.** Ballistic channels and stable quasi-particles yield $\Gamma(k) \sim |k|$ or coexistence of ballistic and diffusive channels; single-mode diffusion fails. The “effective exponent” extracted from small- k finite-size data can drift and even become negative when the estimator is outside its validity window (Appendix H).
- **MBL or quasi-MBL.** Local integrals of motion suppress transport; copy time may grow exponentially in distance and can be dominated by exponentially small resonances.

- **Floquet without conservation.** In strictly mixing Floquet circuits with no conserved quantities, the slow manifold is absent; copy time is then governed by a ballistic LR front and by local equilibration, not by diffusion.
- **Quasi-conservation / prethermalization.** Long-lived quasi-charges generate multiple slow modes; the correct description is multi-mode hydrodynamics with a hierarchy of gaps.

$$\Delta=1, L=12, \text{ fit window } t \in [2.0, 10.0]$$




Figure 3. Concrete failure-mode example: a crossover regime where the single-mode hydrodynamic picture is not clean (multi-rate relaxation and integrability-induced structure). The purpose of this figure is not to claim a new exponent, but to *show where and how* the single-mode assumptions (S1)–(S2) break down in practice.

10. Numerical Benchmarks: ED with Explicit Uncertainty Quantification

We provide a reproducible pipeline (Supplementary Code Archive SC1, including the corresponding analysis outputs) that produces every reported figure from raw data; a figure-to-script map (FIGURE_MAP.md) is included in the code archive. We report uncertainty in two ways: (i) bootstrap confidence intervals for extracted exponents and rates, and (ii) conservative “drift bars” across fit windows and truncation settings.

10.1. Exact Diagonalization Transport Extraction

We estimate the small- k decay rate $\Gamma(k)$ from the time dependence of the spin structure factor $S(k, t)$, using linear fits of $\log S(k, t)$ over a window chosen by stability diagnostics.

Conservative finite-size protocol.

Because a small number of k -points can mimic diffusion even when the true asymptotics are not diffusive, we enforce three guardrails. (i) We never infer diffusion from a single system size; Figure 5 displays size drift explicitly. (ii) For each L we fit $\Gamma(k)$ only on windows where the slope is stable under shifting the time-fit window used to extract Γ from $S(k, t)$. (iii) We report only an *effective* finite-size diagnostic $D_{\text{eff}}(L) := \Gamma(k_{\min})/k_{\min}^2$ ($k_{\min} = 2\pi/L$) together with its window-to-window drift; we do not extrapolate an infinite-volume diffusion constant from $L \leq 14$. This protocol is designed to isolate reproducible finite-size transport signatures at $L \leq 14$ without folding uncontrolled extrapolations into the reported quantities.

Representative outputs are shown in Figure 4 and, for visual context, in Figure 5.




Figure 4. Conservative transport summary at small sizes: the estimator $D_{\text{eff}}(L) = \Gamma(2\pi/L)/(2\pi/L)^2$ plotted against $1/L$, with bootstrap 95% confidence intervals inherited from the decay-rate extraction. The pronounced drift and non-monotonicity at $h_{\text{stag}} = 0$ illustrate why we avoid interpreting $L \leq 14$ data as asymptotic diffusion in the integrable regime. The broken-integrability case $h_{\text{stag}} = 0.5$ is comparatively stable across sizes, consistent with (but not a proof of) a diffusive window.




Figure 5. Transport extraction diagnostic across system sizes (XXZ chain with integrability broken by a staggered field $h = 0.5$). For each $L \in \{8, 10, 12, 14\}$ we estimate the decay rate $\Gamma(k)$ for several discrete momenta $k_n = 2\pi n/L$ and report $\Gamma(k)$ vs. k^2 . Rather than relying on three-point linearity, we treat the small- k window as a size-dependent fit problem: we require stability of the slope under window shifts and propagate the resulting spread as a systematic uncertainty in $D_{\text{eff}}(L)$ and in the inferred k^2 trend (Section 10).




Figure 6. Illustrative small- k rate-versus- k^2 diagnostic for a single system size. The plot is provided for visual context only: with few available momenta at $L \leq 14$, apparent linearity cannot be taken as asymptotic evidence for a diffusion pole. Consequently, the main text reports only the finite-size diagnostic $D_{\text{eff}}(L) = \Gamma(2\pi/L)/(2\pi/L)^2$ together with window-to-window drift (Table 3), and does not extrapolate an infinite-volume diffusion constant.

Table 3. Conservative finite-size diagnostic $D_{\text{eff}}(L) = \Gamma(2\pi/L)/(2\pi/L)^2$ extracted from ED structure-factor decay rates at the smallest nonzero momentum. We report point estimates and bootstrap 95% CIs for two representative regimes in the XXZ chain: an integrable point ($h_{\text{stag}} = 0$) and a symmetry-breaking perturbation ($h_{\text{stag}} = 0.5$). The strong size drift and inconsistent scaling at $h_{\text{stag}} = 0$ (integrable) illustrate the failure of a single diffusive description at these sizes; by contrast, $h_{\text{stag}} = 0.5$ shows a comparatively stable $D_{\text{eff}}(L)$ across L . Rows where the CI endpoints coincide reflect a degenerate bootstrap (insufficient independent typicality vectors at that L in the archived run), and should be interpreted as point estimates rather than as genuine uncertainty quantification.

L	h_{stag}	$D_{\text{eff}}(L)$	95% CI
8	0.0	0.242	[0.106, 0.348]
10	0.0	1.661	[1.316, 1.962]
12	0.0	1.736	[1.443, 2.118]
14	0.0	1.409	[1.409, 1.409]
8	0.5	0.303	[0.252, 0.354]
10	0.5	0.296	[0.294, 0.298]
12	0.5	0.330	[0.326, 0.334]
14	0.5	0.335	[0.335, 0.335]

Finite-size reporting. Because $L \leq 14$ provides only a coarse momentum grid, we report the finite-size diagnostic $D_{\text{eff}}(L) = \Gamma(2\pi/L)/(2\pi/L)^2$ together with its drift across admissible fit windows. The analysis bundle in Supplementary Code Archive SC1 (in the accompanying data archive) contains the full set of window scans, slope-stability tests, and additional integrity checks used to determine which fits are admissible, including an explicit cross- L *time-window sensitivity* test (see `FIGURE_MAP.md`).

10.2. Direct Computation of the Helstrom Advantage and $\tau_{\text{copy}}^{\text{full}}$ (XXZ, B One Site)

A key objective is to measure the operational object itself, not only the diffusion diagnostic $D_{\text{eff}}(L)$. For a one-site receiver $B = \{l\}$ and infinite-temperature hypotheses $\rho_{\pm}(0) = \frac{\mathbb{I}}{d} \pm \frac{\epsilon}{d} S_0^z$, the reduced-state difference on B remains diagonal in the S_l^z basis and the Helstrom advantage simplifies to

$$\text{Adv}_{0 \rightarrow l}^{\text{full}}(t) = \frac{1}{2} \|\rho_{+,l}(t) - \rho_{-,l}(t)\|_1 = 2|\epsilon| |C_{zz}(l, t)|, \quad C_{zz}(l, t) := \frac{1}{d} \text{Tr}(S_l^z(t) S_0^z). \quad (51)$$

Thus, for B one site the Helstrom measurement is operationally simple: it is a single local two-outcome measurement of S_l^z .

Figure 7 shows $\text{Adv}_{0 \rightarrow l}^{\text{full}}(t)$ for the nonintegrable XXZ chain ($\Delta = 1, h_{\text{stag}} = 0.5$) at $L = 12$ computed by ED with typicality ($N_{\text{vec}} = 12$), together with the threshold $\eta = 0.02$. We define $\tau_{\text{copy}}^{\text{full}}(0 \rightarrow l; \eta)$ as the first intersection time of the (bootstrap-averaged) $\text{Adv}^{\text{full}}(t)$ curve with η . To reduce sensitivity to oscillatory transients in unitary dynamics, Supplementary S2 also reports a more conservative “sustained-crossing” variant, defined as the first time at which $\text{Adv}^{\text{full}}(t)$ remains above η for a short window Δt . Figure 8 plots the extracted $\tau_{\text{copy}}^{\text{full}}$ against l^2 . While the ED sizes accessible here are small and show finite-size scatter (with explicit confidence intervals), the resulting first-crossing and sustained-crossing extractions remain compatible with a quadratic transport interpretation on the accessible window, but they do not by themselves establish an asymptotic scaling law. Figure 9 further shows that the extracted times vary smoothly with the threshold η . All scripts, raw data, and figure-generation commands are included in Supplementary Code Archive SC1 and the analysis folder under Data/; see Reproducibility/Figure_Map.md for a direct manuscript-to-code mapping.




Figure 7. Direct ED computation of the Helstrom advantage $\text{Adv}_{0 \rightarrow l}^{\text{full}}(t)$ for a one-site receiver $B = \{l\}$ in the nonintegrable XXZ chain ($L = 12, \Delta = 1, h_{\text{stag}} = 0.5$), using $N_{\text{vec}} = 12$ typicality vectors and bootstrap 95% CIs. The dashed line is the threshold $\eta = 0.02$. The copy time $\tau_{\text{copy}}^{\text{full}}(0 \rightarrow l; \eta)$ is extracted as the first intersection of the (bootstrap-averaged) curve with η .




Figure 8. Extracted Helstrom copy time $\tau_{\text{copy}}^{\text{full}}(0 \rightarrow l; \eta)$ (bootstrap 95% CIs) plotted versus l^2 for the same ED dataset as Figure 7 ($\eta = 0.02$, B one site). The dashed line is a descriptive weighted linear guide drawn on the accessible nonzero-distance points only; it is included to summarize the finite-size trend and is not interpreted as a thermodynamic fit.




Figure 9. Threshold sensitivity diagnostic extracted from the same ED dataset as Figures 7–8: first-crossing times as a function of l^2 for $\eta \in \{0.01, 0.02, 0.05\}$. The dependence on η is smooth on the accessible window, supporting the robustness of the copy-time extraction for small variations of the operational threshold.

Numerical diagnostic for fast mixing.

To validate assumption (S3) in Theorem 3, we monitor an operator with negligible overlap with the conserved charge, e.g. the transverse spin S_0^x . Figure 10 shows the decay of the infinite-temperature autocorrelator $|C_{xx}(t)| = |\frac{1}{d} \text{Tr}(S_0^x(t) S_0^x)|$ in the same XXZ parameters. Rather than treating an exponential envelope as a proof of a sharp spectral gap, we use it as a pragmatic numerical diagnostic: an approximately exponential decay over an intermediate window suggests the presence of a “fast” sector that relaxes on a microscopic time scale. To make this diagnostic more transparent, Figure 11 plots the effective decay rate $\gamma_{\text{eff}}(t) = -\partial_t \log |C_{xx}(t)|$, which displays an approximate plateau on the same window.




Figure 10. Fast-sector relaxation diagnostic in the nonintegrable XXZ chain ($L = 12$, $\Delta = 1$, $h_{stag} = 0.5$): decay of the transverse autocorrelator $|C_{xx}(t)| = |\frac{1}{d} \text{Tr}(S_0^x(t) S_0^x)|$ (log scale). The dashed line is an exponential fit on a late-time window, yielding an estimated fast relaxation rate γ_{fast} that sets a microscopic time scale $t_{fast} \sim \gamma_{fast}^{-1}$.




Figure 11. Effective decay rate extracted from the same fast-sector correlator as Figure 10: $\gamma_{eff}(t) = -\partial_t \log |C_{xx}(t)|$ (computed by a centered finite difference). An approximate plateau over an intermediate-time window supports the use of a single microscopic relaxation scale $t_{fast} \sim \gamma_{eff}^{-1}$ as a diagnostic, while still allowing for finite-size and residual oscillatory features.




Figure 12. Threshold dependence (ED, small system). Copy-time estimates extracted from exact evolution at $L = 8$ as the operational threshold η is varied, for a fixed sender–receiver separation in the same model and hypothesis class as Figure 2. The dependence on η is smooth on the accessible window, consistent with the slowly varying (Lambert- W /logarithmic) threshold factor in Eq. (43).




Figure 13. Receiver-size dependence (ED, small system). Advantage and inferred copy-time trend as the receiver block size $|B|$ is increased (same ED run family). Enlarging B increases the receiver prefactor K_B in Eq. (42), leading to earlier threshold crossing, consistent with Eq. (43).




Figure 14. Selected pipeline diagnostics (Appendix). Top: a threshold-crossing diagnostic for the classical diffusive reference at $\eta = 0.1$ (dashed line), showing that for this conservative parameter set the receiver advantage stays below threshold, hence $\tau_{\text{copy}} = \infty$ (no crossing) while $\max_t \text{Adv}^{\text{full}}$ remains small. Middle: Gaussian approximation error (trace-distance proxy) versus receiver size. Bottom: predictor residual versus receiver size for a regression-based proxy. The full diagnostic set (including parameter sweeps) is included in the SC1 data archive (Appendix I).

10.3. Finite-Size Drift Diagnostics and the “Negative Exponent” Issue

An effective slope estimator $\hat{z}(L)$ constructed from finite-size trends can yield nonphysical values (including negative numbers) in regimes where the hydrodynamic single-mode assumption is violated

(e.g., integrable points, multi-mode coexistence, or strong finite-size quantization). For this reason we treat $\hat{z}$ strictly as a *breakdown diagnostic*, not as a dynamic exponent, and we relegate the corresponding plots to Appendix H where they are clearly labelled as such.

10.4. Executed TEBD/MPS Cross-Checks Beyond the ED Sizes

The revised submission no longer leaves TEBD/MPS as a future-work protocol only. We bundle an executable reference implementation in SC1, retain the exact small-system TEBD-versus-ED validation reported in Supplementary S2, and archive explicit larger open-boundary TEBD runs at $L = 18, 20, 22$ for the same one-site receiver observable. These larger- L runs are interpreted conservatively: they do not by themselves determine an asymptotic exponent, but they do provide a beyond-ED stress test of the extraction pipeline. Supplementary S7 turns this into an explicit diagnostic protocol by combining (i) time-window stability of the slowest- k rate, (ii) agreement between first-crossing and sustained-crossing estimators within a reported tolerance band, and (iii) direct validation that the one-site receiver problem remains commuting/charge-like so that the operational observable is not silently drifting away from the transported slow sector.

10.5. One-Page Synthesis of Regimes, Scalings, and Uncertainties

For convenience, Table 4 collects the central operational claims, the dynamical regime in which they are supported, and the level of validation provided in this work.

Table 4. Synthesis: dynamics/regime $\times$ transport diagnostic $\times$ copy-time scaling. “Supported by” indicates what in this work underwrites the stated scaling (theorem, exact toy model, or ED estimate with explicit confidence intervals).

Model / dynamics	Regime / structure	Diagnostic	τ_{copy} scaling	Supported by
Gaussian charge field (commuting)	Diffusive kernel, Gaussian fluctuations	Exact TV distance (44)	$\sim \ell^2 / (4D)$ (log- η corr.)	Appendix G
Chaotic $U(1)$ dynamics (generic)	Ballistic operator growth + diffusive charge	Separation Prop. 5	$\tau_{\text{OTOC}} \sim \ell$, $\tau_{\text{copy}} \sim \ell^2$	Prop. 5
XXZ chain ($h_{\text{stag}} = 0.5$)	Nonintegrable, candidate diffusive window	$D_{\text{eff}}(L)$ (bootstrap CI)	Consistent with ℓ^2 window; no asymptotic claim	Table 3, Figure 4
XXZ chain ($h_{\text{stag}} = 0$)	Integrable / multimode	Drift/nonmonotone $D_{\text{eff}}(L)$	Single-mode diffusion fails	Figure 4, Sec. 10
Range- R QCA	Strict light cone	Hard causal delay	$\tau_{\text{copy}} \geq \lceil d/R \rceil$	Theorem 1

11. Conclusions

We introduced a receiver-limited operational definition of copy time, established minimal locality bounds, and exhibited a controlled diffusive benchmark in Q-SSEP in which the *Helstrom* copy time is bounded directly from microscopic transport within an explicit diagonal state family. This closes the main conceptual loop of the paper: locality bounds determine when influence cannot yet be certified, while the conserved slow sector determines when certification can become operationally available. For closed Hamiltonian systems, we formulated the corresponding hydrodynamic statement only as a conditional reduced-dynamics theorem with explicit short-memory assumptions, explicit error channels, and explicit diagnostic proxies, and we supplemented it with a microscopic no-go theorem showing why a dissipative fast-sector gap cannot be postulated on the exact unitary Mori block (Supplementary S6). We also made precise two points that are often hidden in informal discussions: first, fixed-threshold diffusive copy times are not asymptotically meaningful unless the threshold convention is specified together with distance, receiver size, or perturbation amplitude; second, no model-independent implication from one-point diffusive transport data to unrestricted Helstrom copy time can hold in generic closed systems without additional receiver-structure assumptions (Supplementary S7). The XXZ numerics and executed TEBD cross-checks directly probe the operational quantity

itself and provide finite-size consistency tests of the transport picture developed in the analytical sections. Taken together, these results identify copy time as an operational horizon for transport that is sharply distinct from ballistic operator growth, anchored by a controlled microscopic benchmark in Q-SSEP, and extensible to unitary many-body dynamics only through clearly stated conditions. In that precise sense, the manuscript isolates a transport question, states the exact operational observable attached to it, and stops precisely at the level supported by the proof architecture.

Appendix A. Proof Sketches and Technical Details

Appendix A.1. Proof of Theorem 1

We treat the Hamiltonian/Lindbladian Lieb–Robinson-tail case **(H)**; the strict circuit/QCA light-cone statement **(C)** follows immediately from (11) because the evolved observable remains supported in A^{+Rt} and therefore cannot influence B when $d(A, B) > Rt$. Let $\delta\rho(0) = \rho_+(0) - \rho_-(0)$ and $\delta\rho_B(t) = \text{Tr}_{B^c}(\mathcal{U}_t(\delta\rho(0)))$. By duality of the trace norm,

$$2\text{Adv}_{A \rightarrow B}^{\text{full}}(t) = \|\delta\rho_B(t)\|_1 = \sup_{\|O_B\|_\infty \leq 1} \text{Tr}(O_B \delta\rho_B(t)) = \sup_{\|O_B\|_\infty \leq 1} \text{Tr}(O_B(t) \delta\rho(0)), \quad (\text{A52})$$

where $O_B(t) = \mathcal{U}_t^\dagger(O_B)$ is the Heisenberg evolution. Since $\delta\rho(0)$ is supported on A , insert an arbitrary operator $O_B^{(A^c)}(t)$ supported on A^c and use $\text{Tr}(O_B^{(A^c)}(t) \delta\rho(0)) = 0$:

$$\text{Tr}(O_B(t) \delta\rho(0)) = \text{Tr}\left((O_B(t) - O_B^{(A^c)}(t)) \delta\rho(0)\right). \quad (\text{A53})$$

Choosing $O_B^{(A^c)}(t)$ as the best approximation to $O_B(t)$ supported on A^c and applying standard LR localization bounds yields

$$\inf_{\text{supp}(Y) \subseteq A^c} \|O_B(t) - Y\|_\infty \leq C_1 e^{-\mu(d(A, B) - vt)}. \quad (\text{A54})$$

Finally, $|\text{Tr}(X \delta\rho(0))| \leq \|X\|_\infty \|\delta\rho(0)\|_1$ with $\|\delta\rho(0)\|_1 \leq 2$ gives

$$2\text{Adv}_{A \rightarrow B}^{\text{full}}(t) \leq 2C_1 e^{-\mu(d(A, B) - vt)}, \quad (\text{A55})$$

which is (A55).

Appendix B. Hydrodynamic Single-Mode Derivation

Appendix B.1. From Projected Dynamics to a Diffusion Pole

In the single-mode window, the projected Liouvillian on the conserved-density subspace is diagonal in Fourier space and has eigenvalues $-\Gamma(k)$ with $\Gamma(k) = Dk^2 + o(k^2)$. Let $\delta q_k(t)$ denote the conserved-density mode. Then $\delta q_k(t) = e^{-\Gamma(k)t} \delta q_k(0)$. The receiver observable J_B couples to these modes with form factor $F_B(k)$, so the projected signal takes the form

$$S_B(t) = \epsilon \sum_{k \neq 0} F_B(k) e^{-\Gamma(k)t} \delta q_k(0) + (\text{higher modes}). \quad (\text{A56})$$

The second-moment susceptibility in Definition 2 is essentially the k -space weighted sum $\sum_{k \neq 0} |F_B(k)|^2 / \Gamma(k)^2$, which is dominated by the smallest $|k|$ in finite volume.

Appendix B.2. Threshold Inversion and ℓ^2 Scaling

For a localized sender perturbation, $\delta q_k(0) \sim e^{-ikx_A}$ and for an interval receiver B centered at x_B , $F_B(k) \sim e^{ikx_B} \text{sinc}(kw/2)$. Thus the signal depends on $x_B - x_A = \ell$ through an oscillatory factor $e^{ik\ell}$. In the continuum approximation, replacing the discrete sum by an integral yields the heat-kernel expression (36) and the saddle-point estimate (37), from which (30) follows. Finite-volume corrections

can be bounded by Poisson summation and yield the explicit $O(\log(1/\eta))$ threshold term quoted in Theorem 3. Supplementary S6 collects the exact admissibility and branch-structure proofs, together with the microscopic no-go theorem for unitary fast-sector gaps.

We provide the diffusion-kernel computation leading to (30). In one dimension, a localized chemical potential profile evolves as $\delta\mu(x, t) \approx (4\pi Dt)^{-1/2} \exp(-x^2/(4Dt))$ in the hydrodynamic window. Evaluating its overlap with a receiver region at distance ℓ yields a signal proportional to $\exp(-\ell^2/(4Dt))$, leading to $\tau \sim \ell^2/(4D)$ up to logarithmic threshold factors.

Appendix C. Moment-Channel Optimality in Gaussian Fluctuation Algebras

We show that if the induced reduced states on B are Gaussian in the relevant fluctuation variables, then low moments determine the Helstrom measurement asymptotically, and the moment-channel lower bound (49) matches Adv^{full} to leading order.

Appendix D. QCA Locality Versus LR Tails

Appendix D.1. From Strict Causal Cones to Operational Zero Advantage

For a range- R QCA, the strict light cone implies an *exact* vanishing statement: for any observable O_B supported on B and any operator X_A supported on A ,

$$\text{Tr}(O_B \alpha^t(X_A)) = 0 \quad \text{whenever} \quad d(A, B) > Rt, \quad (\text{A57})$$

because $\alpha^t(X_A)$ is supported on A^{+Rt} and the local algebras commute at distance. In contrast, for Hamiltonian evolution one only obtains exponentially small commutator tails as in (10). This distinction matters operationally: in QCA dynamics, $\text{Adv}_{A \rightarrow B}^{\text{full}}(t)$ is *exactly* zero outside the cone, whereas for Hamiltonians it is exponentially small rather than strictly vanishing and can in principle accumulate through many weak channels.

Appendix D.2. Why Copy-Time Data Might “See” the Index

The QCA index [15] is defined through support-algebra dimensions associated with the image of local algebras under α . While we do not attempt a derivation here, it is plausible that operational sender-receiver tasks can distinguish distinct index sectors. A concrete approach is to compare d_η for families of disjoint sender/receiver intervals under stacking and coarse graining, testing multiplicativity patterns implied by the index. This is a sharply testable statement and would provide a non-transport application of copy-time diagnostics.

Appendix E. Finite-Size Corrections for Diffusion-Kernel Inversion

Here we record a finite-size bound that turns the heuristic replacement of a discrete k -sum by a continuum integral into a controlled approximation.

Let $k_n = 2\pi n/L$, $n \in \mathbb{Z}$, and consider the kernel

$$K_L(\ell, t) := \frac{1}{L} \sum_{n \neq 0} e^{-Dk_n^2 t} e^{ik_n \ell}. \quad (\text{A58})$$

By Poisson summation, for $tD \ll L^2$ one has

$$K_L(\ell, t) = \frac{1}{\sqrt{4\pi Dt}} \sum_{m \in \mathbb{Z}} \exp\left(-\frac{(\ell + mL)^2}{4Dt}\right) - \frac{1}{L}, \quad (\text{A59})$$

where the $-1/L$ subtracts the zero mode. The leading term $m = 0$ recovers the infinite-line heat kernel (36); the finite-size correction is bounded by

$$\left| K_L(\ell, t) - \frac{1}{\sqrt{4\pi Dt}} e^{-\ell^2/(4Dt)} \right| \leq \frac{2}{\sqrt{4\pi Dt}} \sum_{m \geq 1} \exp\left(-\frac{(mL - \ell)^2}{4Dt}\right) + \frac{1}{L}. \quad (\text{A60})$$

For $\ell \leq L/2$ and $tD \leq \kappa L^2$ with $\kappa < 1/16$, the dominant correction is $O\left(e^{-(1-2\ell/L)^2/(16\kappa)}\right)$, exponentially small in $1/\kappa$. This justifies using the continuum inversion for τ_{copy} in the pre-asymptotic window $t \ll L^2/D$.

Appendix F. Small-System ED Illustration: Threshold and Receiver-Size Dependence

This appendix provides an *illustrative* numerical check of the smooth dependence on the operational threshold η and on the receiver size $|B|$ that is made explicit by the Gaussian-kernel inversion formula (43). We emphasize the intended scope: these exact-diagonalization (ED) runs are performed at a small size ($L = 8$) to keep the computation exact and fully reproducible within the submission bundle. They are included to demonstrate that (i) changing η changes the extracted copy time smoothly, and (ii) enlarging the receiver region increases the advantage and therefore reduces the inferred copy time, as expected in transport-limited regimes. They do *not* establish thermodynamic-limit scaling.

Appendix G. Gaussian Discrimination and Moment Sufficiency: An Explicit Bound

Assume that, in the hydrodynamic window, the receiver's relevant fluctuation variable X_B (e.g., coarse-grained charge in B) is approximately Gaussian under both hypotheses, with means $\pm m(t)$ and common variance $\sigma^2(t)$. Then the optimal Helstrom measurement reduces (in the classical limit) to thresholding X_B , and the resulting advantage is

$$\text{Adv}^{\text{full}}(t) \approx \text{TV}\left(\mathcal{N}(m(t), \sigma^2(t)), \mathcal{N}(-m(t), \sigma^2(t))\right) = \text{erf}\left(\frac{m(t)}{\sqrt{2}\sigma(t)}\right). \quad (\text{A61})$$

For small signal-to-noise ratio m/σ , $\text{Adv}^{\text{full}}(t) \sim \sqrt{\frac{2}{\pi}} \frac{m(t)}{\sigma(t)}$. Since $m(t)$ is controlled by the diffusion kernel (37) and $\sigma(t)$ is set by equilibrium fluctuations, low moments (mean and variance) are sufficient to determine the advantage at leading order. This provides a principled route to moment-channel near-optimality in regimes where the fluctuation algebra is approximately Gaussian.

Appendix H. Additional Numerical Tables and Metadata

Finite-size “effective-exponent” diagnostics (including negative/drifted slopes at the integrable point) are provided as CSV tables in the bundled SC1 data archive (see `FIGURE_MAP.md`); in the main PDF we treat these quantities strictly as *breakdown diagnostics* rather than physical exponents.

Additional numerical diagnostics and reporting details (including window scans and stability checks) used to assess the single-mode window at the accessible sizes are included in the SC1 data archive (see `FIGURE_MAP.md`). Supplementary S2 is reserved for the TEBD/MPS protocol outline. We recommend interpreting small- L “effective exponents” only as diagnostics of hydrodynamic breakdown, not as physical exponents in integrable regimes.

Appendix I. Reproducibility: Run Manifest and Code Pointers

For reproducibility, Supplementary Files S1–S4 and the Supplementary Code Archive SC1 include:

- definitions, locality bounds, and technical lemmas supporting the operational setup (Supplementary File S1);
- a detailed TEBD/MPS protocol together with exact TEBD-versus-ED validation and archived larger- L conservative cross-check runs (Supplementary File S2);
- the controlled Q-SSEP diagonal-sector benchmark and explicit assumptions supporting the diffusive lower bound (Supplementary File S3);
- conceptual positioning relative to OTOCs and toy examples separating receiver-limited certifiability from operator-growth diagnostics (Supplementary File S4);

- the code, parameter registries, integrity checks, fit-window scans, and the minimal data bundle used to generate the reported ED diagnostics and Helstrom-advantage figure suite (Supplementary Code Archive SC1 together with the accompanying data archive).

Appendix I.1. Reproducibility Pointers (ED Extraction and Post-Processing)

To avoid over-promising numerics that are not explicitly reported, we separate *protocol descriptions* from *computed results*, and we explicitly label every TEBD item below as either a validation run or a larger- L qualitative cross-check. All ED-based figures and tables reported in this work can be reproduced from the bundled Supplementary Code Archive SC1 together with the accompanying analysis data archive. A concise map from each manuscript figure/table to the corresponding script, command line, and output file is provided in `FIGURE_MAP.md` (included in the accompanying Supplementary Materials and Code Archive). For the operational numerics in Section 10.2, the figure-generation scripts and the corresponding raw outputs are indexed in `FIGURE_MAP.md` within the accompanying code archive.

Appendix I.2. TEBD/MPS Protocol and Archived Cross-Check Runs

Supplementary File S2 provides a detailed TEBD/MPS protocol together with the exact small-system TEBD-versus-ED validation already archived in SC1 and the larger open-boundary qualitative cross-check runs at $L = 18, 20, 22$. These TEBD items are not used as principal asymptotic evidence, but they are reported explicitly and bundled as executable code plus archived outputs. Accordingly, SC1 *does* include the lightweight TEBD source code used for these checks in the directory `TEBD_code/`.

Supplementary Materials: The following supporting information can be downloaded at the website of this paper posted on Preprints.org, Supplementary File S1: definitions, locality bounds, and technical lemmas supporting the operational setup. Supplementary File S2: TEBD/MPS protocol, exact small-system TEBD-versus-ED validation, and larger- L qualitative cross-check runs (validation and conservative cross-checks only). Supplementary File S3: controlled Q-SSEP diagonal-sector benchmark and additional derivations for the diffusive lower bound. Supplementary File S4: conceptual positioning relative to OTOCs and toy examples separating operator growth from receiver-limited certifiability. Supplementary File S5: extended validation package developing the framework into (i) a direct operational separation from scrambling proxies, (ii) minimal-assumption theorem statements, (iii) an experimental readout protocol with shot/readout-noise accounting, and (iv) anomalous-transport scaling templates. Supplementary Code Archive SC1: code, minimal data, and scripts to reproduce the ED diagnostics and the Helstrom-advantage figure suite reported in the manuscript. Reproducibility notes: a figure-to-script map (`FIGURE_MAP.md`) is included in the code archive.

Data Availability Statement: Data and Code Availability: All materials needed to reproduce the reported exact-diagonalization diagnostics and the Helstrom-advantage figure suite are included in the accompanying Supplementary Materials and Code Archive: (i) Supplementary Files S1–S5 (theory, strengthened diagnostics, and protocol descriptions) and (ii) Supplementary Code Archive SC1, including the corresponding analysis outputs and scripts needed to reproduce the reported ED diagnostics. A figure-to-script map (`FIGURE_MAP.md`) is included in the code archive. A small set of explicit TEBD/MPS cross-checks beyond ED are now archived and described in Supplementary S2, but they are still interpreted conservatively as validation and finite-window consistency checks rather than as a thermodynamic-limit determination of the scaling law.

Conflicts of Interest: The author declares no competing interests.

References

1. Helstrom, C.W. *Quantum Detection and Estimation Theory*; Academic Press, 1976.
2. Holevo, A.S. Statistical decision theory for quantum systems. *Journal of Multivariate Analysis* **1973**, *3*, 337–394. [https://doi.org/10.1016/0047-259X\(73\)90028-6](https://doi.org/10.1016/0047-259X(73)90028-6).
3. Kubo, R. Statistical-Mechanical Theory of Irreversible Processes. I. General Theory and Simple Applications to Magnetic and Conduction Problems. *Journal of the Physical Society of Japan* **1957**, *12*, 570–586. <https://doi.org/10.1143/JPSJ.12.570>.

4. Mori, H. Transport, Collective Motion, and Brownian Motion. *Progress of Theoretical Physics* **1965**, *33*, 423–455. <https://doi.org/10.1143/PTP.33.423>.
5. Zwanzig, R. Memory Effects in Irreversible Thermodynamics. *Physical Review* **1961**, *124*, 983–992. <https://doi.org/10.1103/PhysRev.124.983>.
6. Forster, D. *Hydrodynamic Fluctuations, Broken Symmetry, and Correlation Functions*; Benjamin, 1975.
7. Nahum, A.; von Keyserlingk, C.W.; Rakovszky, T.; Pollmann, F.; Sondhi, S.L. Operator hydrodynamics, OTOCs, and entanglement growth in systems with conservation laws. *Physical Review X* **2018**, *8*, 021014. <https://doi.org/10.1103/PhysRevX.8.021014>.
8. Khemani, V.; Vishwanath, A.; Huse, D.A. Operator spreading and the emergence of dissipative hydrodynamics under unitary dynamics with conservation laws. *Physical Review X* **2018**, *8*, 031057. <https://doi.org/10.1103/PhysRevX.8.031057>.
9. Bernard, D.; Jin, T. Anomalous scaling and diffusion in the quantum symmetric simple exclusion process. *Communications in Mathematical Physics* **2021**, *383*, 659–712. <https://doi.org/10.1007/s00220-021-04066-3>.
10. Nahum, A.; Vijay, S.; Haah, J. Operator Spreading in Random Unitary Circuits. *Physical Review X* **2018**, *8*, 021014. <https://doi.org/10.1103/PhysRevX.8.021014>.
11. von Keyserlingk, C.W.; Rakovszky, T.; Pollmann, F.; Sondhi, S.L. Operator Hydrodynamics, OTOCs, and Entanglement Growth in Systems without Conservation Laws. *Physical Review X* **2018**, *8*, 021013, [arXiv:cond-mat.str-el/1705.08910]. <https://doi.org/10.1103/PhysRevX.8.021013>.
12. Xu, S.; Swingle, B. Scrambling Dynamics and Out-of-Time-Ordered Correlators in Quantum Many-Body Systems. *PRX Quantum* **2024**, *5*, 010201, [arXiv:quant-ph/2202.07060]. <https://doi.org/10.1103/PRXQuantum.5.010201>.
13. Khemani, V.; Vishwanath, A.; Huse, D.A. Operator Spreading and the Emergence of Dissipative Hydrodynamics under Unitary Evolution with Conservation Laws. *Physical Review X* **2018**, *8*, 031057. <https://doi.org/10.1103/PhysRevX.8.031057>.
14. Rakovszky, T.; Pollmann, F.; von Keyserlingk, C.W. Diffusive Hydrodynamics of Out-of-Time-Ordered Correlators with Charge Conservation. *Physical Review X* **2018**, *8*, 031058, [arXiv:cond-mat.str-el/1710.09827]. <https://doi.org/10.1103/PhysRevX.8.031058>.
15. Gross, D.; Nesme, V.; Vogts, H.; Werner, R.F. Index Theory of One Dimensional Quantum Walks and Quantum Cellular Automata. *Communications in Mathematical Physics* **2012**, *310*, 419–454. <https://doi.org/10.1007/s00220-011-1353-3>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.