

Condensed Locality: Lieb–Robinson Estimates and Quasi-Local Dynamics on the Moduli Stack of Hamiltonians

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

14 July 2026

Part I of a six-part series on topological phases of matter in the condensed-mathematics paradigm

Abstract

The condensed-mathematics program for topological phases replaces the topological space of admissible interactions by a condensed object probed by profinite sets, so that continuous families, profinite disorder hulls, and inverse limits are handled inside a single category with exact homological algebra. This opening paper supplies the analytic ground floor. Fixing a Nachtergaele–Sims–Young F -function on a lattice, we recall that the G -symmetric interactions of finite F -norm form a Banach space $\mathcal{B}_F$, and we condense it: $\underline{\mathcal{B}}_F(S) = \text{Cont}(S, \mathcal{B}_F)$ on profinite probes. We prove three theorems, each with a complete proof assembled from present-day results. First, $\mathcal{B}_F$ is a real Banach space and the Heisenberg dynamics is a strongly continuous one-parameter group of $*$ -automorphisms of the quasi-local algebra obeying a Lieb–Robinson bound. Second, $\underline{\mathcal{B}}_F$ is a condensed $\mathbb{R}$ -vector space and the passage to dynamics is a morphism of condensed sets $\mathbb{R} \times \underline{\mathcal{B}}_F \rightarrow \underline{\text{Aut}}(\mathcal{A})$, by full faithfulness of condensation on compactly generated spaces; for separable systems the morphism is one of light condensed sets. Third, an S -continuous family of interactions over a profinite base is exactly a norm-convergent system of locally constant families over finite quotients, and the load-bearing point is that a uniform F -norm bound forces a Lieb–Robinson velocity and constants that are uniform over the base. Quasi-adiabatic continuation and Bachmann–Michalak–Nachtergaele–Sims automorphic equivalence are then the morphism-level tools that generate the equivalence class $\mathcal{W}$ inverted by the companion papers. We close with three numbered conjectures: that $\mathfrak{Ham}_{d,G}$ is a condensed higher stack, that time-dependent F -function interactions generate a condensed group with quasi-adiabatic continuation as internal path-lifting, and that Lieb–Robinson estimates satisfy descent along finite

quotients of a disorder hull, together with a Haskell simulation of the transverse-field Ising chain that exhibits the light cone and fits the velocity.

Contents

1	Introduction	2
1.1	The locality problem in the condensed program	2
1.2	What this paper proves, and what it only conjectures	4
1.3	Relation to companion papers	4
2	Quantum lattice systems and interaction spaces	5
2.1	Lattice, local Hilbert spaces, and the quasi-local algebra	5
2.2	F -functions and the interaction Banach space	6
3	The Heisenberg dynamics and the Lieb–Robinson bound	7
3.1	Finite-volume dynamics and the infinite-volume limit	7
3.2	The generator as a symmetric derivation	9
4	Condensation of the interaction space	10
4.1	Condensed sets and profinite probes	10
4.2	The condensed interaction space and the condensed dynamics	11
5	Uniform Lieb–Robinson bounds over a profinite base	13
5.1	Uniformity over the base	13
5.2	Uniform clustering: the bridge to the gapped substack	15
6	Quasi-adiabatic continuation and the class $\mathcal{W}$	15
7	Conjectures: toward the condensed stack	16
8	The light cone in the transverse-field Ising chain	17
9	Discussion	18
10	Conclusion	19
A	Constants in the Lieb–Robinson bound	19
	References	20

1 Introduction

1.1 The locality problem in the condensed program

A gapped quantum lattice system is specified by a lattice L of sites in dimension d , a finite-dimensional Hilbert space at each site, and a rule that assigns to each finite cluster of sites

a Hermitian operator supported there, an *interaction*. The set of admissible G -symmetric interactions, together with a locality (decay) condition, is the topological space $\mathcal{I}_{d,G}$ of the program this series develops. The organizing proposal is to stop treating $\mathcal{I}_{d,G}$ as a bare topological space and instead pass to its condensation,

$$X \longmapsto \underline{X}, \quad \underline{X}(S) = \text{Cont}(S, X),$$

with S ranging over profinite sets, and then to build over it the moduli stack $\mathfrak{Ham}_{d,G}$ of Hamiltonians, the uniformly gapped substack $\mathfrak{Gap}_{d,G}$, the stabilized phase ∞ -groupoid $\mathfrak{Phase}_{d,G}$, and finally the invertible condensed phase spectrum $\mathbf{IP}_{d,G}^{\text{cond}}$. In one sentence, the thesis of the program is that

a topological phase is a component of the stabilized condensed stack of gapped systems.

The advantage of this move is not a new Chern number; it is that condensed abelian groups form a Grothendieck abelian category with exact derived functors, so that continuous families, profinite disorder, analytic completions, symmetry, stacking, defects, and transition loci can be manipulated together, where topological groups behave badly.

None of this is available until one controls the analysis at the bottom. Before $\mathfrak{Ham}_{d,G}$ can be a condensed object at all, three things must be true and must be *uniform in the probe*: the space of interactions must be a Banach space so that its condensation is a condensed $\mathbb{R}$ -vector space; the Heisenberg dynamics must depend continuously on the interaction and the time in the Banach topology, so that a family of dynamics is a morphism rather than a set-theoretic assignment; and the Lieb–Robinson estimates that make the dynamics quasi-local must survive the passage to profinite limits with constants that do not degenerate. This paper establishes exactly these facts, and no more. The word “uniform” is doing real work throughout: a family of quasi-local dynamics whose Lieb–Robinson velocity is finite pointwise but unbounded over the base is not a morphism of condensed sets, just as a family whose gap is positive pointwise but has infimum zero is not a member of $\mathfrak{Gap}_{d,G}$.

The tools are classical. Lieb and Robinson proved a finite group velocity for quantum spin systems in 1972 [11]; the modern reformulation through F -functions, due to Nachtergaele, Sims, and Young [18, 19] and surveyed by Hastings [8], gives exactly the Banach-space packaging we need. The condensed formalism is that of Clausen and Scholze [21, 20, 6], with the pyknotic variant of Barwick and Haine [3]. The closest precedent on the physics side is the homotopical, sheaf-theoretic study of parametrized families of spin systems by Beaudry and collaborators [4], which organizes families of gapped systems with ordinary, not condensed, topology; the present paper is the condensed refinement of that viewpoint at the locality layer. Our contribution is to show that these two mature bodies of work fit together on the nose: the analytic estimate that Nachtergaele–Sims–Young prove is precisely the continuity statement that condensation requires, and the profinite probes of condensed mathematics are precisely matched to the inverse limits of finite-volume and disorder data that lattice physics produces.

1.2 What this paper proves, and what it only conjectures

We adopt the editorial stance of the whole series: a statement is labelled Theorem or Proposition only when it has a complete proof from cited present-day results, and everything past that boundary is a numbered Conjecture. Under that discipline the results are as follows. Theorem 3.3 (labelled I-A) records that the F -normed interactions form a real Banach space $\mathcal{B}_F$ and that the Heisenberg dynamics is a strongly continuous one-parameter group of $*$ -automorphisms with a Lieb–Robinson bound; this is a repackaging of [18]. Theorem 4.4 (I-B) upgrades this to condensed mathematics: $\underline{\mathcal{B}}_F$ is a condensed $\mathbb{R}$ -vector space, and the assignment $(t, \Phi) \mapsto \tau_t^\Phi$ is a morphism of condensed sets, because it is continuous in the Banach topology and condensation is fully faithful on compactly generated spaces. Theorem 5.1 (I-C) is the analytic heart: over a profinite base, an S -continuous family of interactions is the same datum as a norm-convergent compatible system over the finite quotients, and a uniform bound on the F -norm produces a Lieb–Robinson light cone whose velocity and prefactors are uniform in $s \in S$. Theorem 5.4 draws the uniform exponential clustering consequence that Part III will need, and Theorem 6.2 records that quasi-adiabatic continuation and automorphic equivalence [10, 2] furnish the morphism-level generators of the equivalence class $\mathcal{W}$.

The conjectures mark the edge of what today’s estimates deliver. Theorem 7.1 (I-1), that $\mathfrak{Ham}_{d,G}$ is a condensed *higher* stack (the groupoid of gauge and quasi-local automorphisms glues along finite covers of profinite probes), is not proved here and is not asserted anywhere in the series as a theorem. Theorem 7.2 (I-2) proposes that the quasi-local automorphisms generated by time-dependent F -function interactions form a condensed group for which quasi-adiabatic continuation is an internal path-lifting. Theorem 7.3 (I-3) proposes that Lieb–Robinson estimates satisfy descent along finite quotients of a disorder hull $\Omega = Q^{\mathbb{Z}^d}$, so that quasi-local invariants are determined by finite-resolution data. These three are the locality-layer instances of the program’s Master Conjecture.

1.3 Relation to companion papers

This is Part I of six. The series is modular: each part takes the previous ones as input and produces new structure, rather than assembling into a single monolith.

Part II, *Positivity, C^* -Norms, and Condensed State Spaces of Quasi-Local Algebras* [14], takes the quasi-local algebra $\mathcal{A}$ and the automorphism morphism built here and attaches the condensed state space and the solid- K -theory invariant of Aoki [1]. The Lieb–Robinson bound of Theorem 3.3 is what makes the Heisenberg dynamics act on $\mathcal{A}$ by genuine $*$ -automorphisms, which Part II needs before it can speak of states and positivity. Part III, *The Uniformly Gapped Substack* [16], cuts out $\mathfrak{Gap}_{d,G} \subset \mathfrak{Ham}_{d,G}$ and studies stability of the gap; its exponential-clustering input is our Theorem 5.4, and its automorphic equivalences are the $\mathcal{W}$ generated by our Theorem 6.2. Part IV, *From Lattice Models to Effective Field Theories* [12], group-completes the invertible sector into $\mathbf{IP}_{d,G}^{\text{cond}}$; the continuity of dynamics in the interaction (Theorem 4.4) is what lets its parametrized families be morphisms of condensed objects. Part V, *Physical Realizability of Bordism and Homotopy Classes* [13], asks which abstract classes are realized by uniformly gapped lattice families, and inherits the uniform-over-the-base discipline established here. Part VI [15] composes the five modules

and states the Master Conjecture. Where those papers speak of “an S -continuous family of Hamiltonians”, the precise meaning is the one fixed by Theorem 5.1.

The rest of the paper is organized as follows. Section 2 fixes the lattice, the quasi-local algebra, the F -functions, and the Banach space $\mathcal{B}_F$. Section 3 recalls the Heisenberg dynamics and the Lieb–Robinson bound and proves I-A. Section 4 recalls exactly as much condensed mathematics as we use and proves the condensed vector-space and dynamics statements (I-B). Section 5 proves the uniform Lieb–Robinson bound over a profinite base and the compatible-finite-data description (I-C), and draws the clustering corollary. Section 6 treats quasi-adiabatic continuation and the generation of $\mathcal{W}$. Section 7 states the three conjectures. Section 8 reports the transverse-field Ising simulation. Sections 9 and 10 discuss limitations and conclude.

2 Quantum lattice systems and interaction spaces

2.1 Lattice, local Hilbert spaces, and the quasi-local algebra

We work in the standard operator-algebraic setting for quantum spin systems; the reader who wants proofs of the folklore statements below will find them in [18] and the references there.

Definition 2.1 (Lattice and local structure). A *lattice* is a countable set L equipped with a metric dist such that the balls $b_x(r) = \{y \in L : \text{dist}(x, y) \leq r\}$ are finite and their cardinality grows at most polynomially: there are $\kappa, \nu > 0$ with $|b_x(r)| \leq \kappa(1+r)^\nu$ for all $x \in L$ and $r \geq 0$. The integer lattice $\mathbb{Z}^d$ with the ℓ^1 metric is the running example, with $\nu = d$. To each site $x \in L$ we attach a finite-dimensional Hilbert space $\mathcal{H}_x$ of dimension n_x , uniformly bounded: $\sup_x n_x < \infty$.

For a finite $X \subset L$ write $\mathcal{H}_X = \bigotimes_{x \in X} \mathcal{H}_x$ and let $\mathcal{A}_X = \mathcal{B}(\mathcal{H}_X)$ be the matrix algebra of observables supported on X . If $X \subset Y$ then $\mathcal{A}_X \hookrightarrow \mathcal{A}_Y$ by $A \mapsto A \otimes \mathbf{1}_{Y \setminus X}$, and the *local algebra* is the union

$$\mathcal{A}_{\text{loc}} = \bigcup_{X \in L} \mathcal{A}_X,$$

the colimit over finite subsets $X \in L$. Its completion in the operator norm is the *quasi-local algebra*

$$\mathcal{A} = \overline{\mathcal{A}_{\text{loc}}}^{\|\cdot\|},$$

a unital C^* -algebra; this is the standard C^* -inductive-limit (quasi-local) construction of Bratteli–Robinson [5]. Because L is countable and each $\mathcal{A}_X$ is finite-dimensional, $\mathcal{A}$ is separable. An element $A \in \mathcal{A}$ is *local* if $A \in \mathcal{A}_X$ for some finite X , and its *support* $\text{supp } A$ is the smallest such X . We write $\mathcal{A}$ for what the notation table of the series calls the observable algebra $\mathcal{A}$; the calligraphic letter is the standard one in the Lieb–Robinson literature and we keep it to avoid clashing with generic operators.

A symmetry is a group G acting on $\mathcal{A}$ by $*$ -automorphisms, on-site in the internal case (a representation $u_x : G \rightarrow \mathcal{U}(\mathcal{H}_x)$ inducing $\alpha_g = \bigotimes_x \text{Ad}(u_x(g))$) and by isometries of (L, dist) in the spatial case. Everything below is compatible with a fixed such G ; to keep the analysis in the foreground we suppress G from the notation and reinstate it only where symmetry changes a statement.

2.2 F -functions and the interaction Banach space

The decay of interactions is measured against a weight on the metric, an F -function.

Definition 2.2 (F -function). An F -function on (L, dist) is a non-increasing function $F : [0, \infty) \rightarrow (0, \infty)$ satisfying

$$(F1) \text{ Uniform summability: } \|F\| := \sup_{x \in L} \sum_{y \in L} F(\text{dist}(x, y)) < \infty;$$

$$(F2) \text{ Convolution bound: } C_F := \sup_{x, y \in L} \sum_{z \in L} \frac{F(\text{dist}(x, z)) F(\text{dist}(z, y))}{F(\text{dist}(x, y))} < \infty.$$

On $\mathbb{Z}^d$ the functions $F_\varepsilon(r) = (1 + r)^{-(d+\varepsilon)}$ with $\varepsilon > 0$ are F -functions, as are the exponentially weighted $F(r) = e^{-\theta r}(1 + r)^{-(d+\varepsilon)}$. The next lemma, which we use constantly, says that one may always sharpen the decay of a given F -function by an exponential factor at no cost in the two constants; it is the mechanism by which a finite velocity appears.

Lemma 2.3 (Exponential reweighting). *Let F be an F -function and $a \geq 0$. Then $F_a(r) := e^{-ar}F(r)$ is an F -function with*

$$\|F_a\| \leq \|F\|, \quad C_{F_a} \leq C_F.$$

Proof. F_a is positive and non-increasing. For (F1), $e^{-ar} \leq 1$ gives $F_a(r) \leq F(r)$, so $\|F_a\| \leq \|F\|$. For (F2), the triangle inequality $\text{dist}(x, y) \leq \text{dist}(x, z) + \text{dist}(z, y)$ and monotonicity of $t \mapsto e^{-at}$ give $e^{-a \text{dist}(x, z)} e^{-a \text{dist}(z, y)} = e^{-a(\text{dist}(x, z) + \text{dist}(z, y))} \leq e^{-a \text{dist}(x, y)}$, hence

$$\begin{aligned} \frac{F_a(\text{dist}(x, z)) F_a(\text{dist}(z, y))}{F_a(\text{dist}(x, y))} &= e^{-a(\text{dist}(x, z) + \text{dist}(z, y) - \text{dist}(x, y))} \frac{F(\text{dist}(x, z)) F(\text{dist}(z, y))}{F(\text{dist}(x, y))} \\ &\leq \frac{F(\text{dist}(x, z)) F(\text{dist}(z, y))}{F(\text{dist}(x, y))}. \end{aligned}$$

Summing over z and taking the supremum over x, y yields $C_{F_a} \leq C_F$. $\square$

To keep the constants concrete—the knowledge base for this program insists that the F -norm never be left implicit—we verify the axioms for the standard polynomial weights.

Lemma 2.4 (Polynomial F -functions on $\mathbb{Z}^d$). *On $(\mathbb{Z}^d, \ell^1)$ the function $F_\varepsilon(r) = (1 + r)^{-(d+\varepsilon)}$ is an F -function for every $\varepsilon > 0$, with $\|F_\varepsilon\| < \infty$ and $C_{F_\varepsilon} \leq 2^{d+\varepsilon+1} \|F_\varepsilon\|$, both independent of the base point by translation invariance.*

Proof. Throughout, $|x - y|_1 = \text{dist}(x, y)$ denotes the ℓ^1 distance and $|y|_1 = \text{dist}(0, y)$ the distance to the origin. The number of sites at ℓ^1 -distance n from a fixed point is $O(n^{d-1})$, so

$$\|F_\varepsilon\| = \sum_{y \in \mathbb{Z}^d} (1 + |y|_1)^{-(d+\varepsilon)} \leq c_d \sum_{n \geq 0} (1 + n)^{d-1} (1 + n)^{-(d+\varepsilon)} = c_d \sum_{n \geq 0} (1 + n)^{-(1+\varepsilon)} < \infty.$$

For (F2), split the sum over z according to whether $|x - z|_1 \geq \frac{1}{2}|x - y|_1$ or $|z - y|_1 \geq \frac{1}{2}|x - y|_1$; at least one holds by the triangle inequality. In the first case $(1 + |x - y|_1)/(1 + |x - z|_1) \leq 2$, so $F_\varepsilon(|x - z|_1)/F_\varepsilon(|x - y|_1) \leq 2^{d+\varepsilon}$ and the summand $F_\varepsilon(|x - z|_1)F_\varepsilon(|z - y|_1)/F_\varepsilon(|x - y|_1)$ is at most $2^{d+\varepsilon}F_\varepsilon(|z - y|_1)$, whose sum over z is $2^{d+\varepsilon}\|F_\varepsilon\|$. The second case is symmetric, giving $C_{F_\varepsilon} \leq 2^{d+\varepsilon+1}\|F_\varepsilon\| < \infty$. $\square$

Definition 2.5 (Interaction and F -norm). An *interaction* is a map Φ that assigns to each finite $X \subseteq L$ a self-adjoint element $\Phi(X) = \Phi(X)^* \in \mathcal{A}_X$, and in the G -symmetric case commutes with the symmetry, $\alpha_g(\Phi(X)) = \Phi(gX)$ for all $g \in G$. Its F -norm is

$$\|\Phi\|_F := \sup_{x,y \in L} \frac{1}{F(\text{dist}(x,y))} \sum_{X \ni x,y} \|\Phi(X)\|,$$

the inner sum ranging over finite X containing both x and y . The *interaction Banach space* is

$$\mathcal{B}_F := \{ \Phi \text{ an interaction} : \|\Phi\|_F < \infty \}.$$

The F -norm controls, uniformly in the site, the total weight of interaction terms that link any two sites. A finite-range interaction of bounded strength has finite F -norm for every F -function; a two-body interaction $\Phi(\{x,y\})$ decaying like $F(\text{dist}(x,y))$ does too. The point of Theorem 2.5 is that everything downstream depends on Φ only through the single number $\|\Phi\|_F$ (and the constants of the chosen F), which is what will let us pass to families.

Proposition 2.6 ($\mathcal{B}_F$ is a Banach space). $(\mathcal{B}_F, \|\cdot\|_F)$ is a real Banach space, and the map $\Phi \mapsto \|\Phi\|_F$ is a genuine norm on it.

Proof. Absolute homogeneity and the triangle inequality are inherited from the operator norm inside the supremum and the sum, and $\|\Phi\|_F = 0$ forces $\Phi(X) = 0$ for every X (take $x, y \in X$). Only completeness needs an argument. Let $(\Phi^{(n)})_n$ be Cauchy in $\|\cdot\|_F$. For each fixed finite X and any two sites $x, y \in X$,

$$\|\Phi^{(n)}(X) - \Phi^{(m)}(X)\| \leq F(\text{dist}(x,y)) \|\Phi^{(n)} - \Phi^{(m)}\|_F,$$

so $(\Phi^{(n)}(X))_n$ is Cauchy in the finite-dimensional space $\mathcal{A}_X$ and converges to some self-adjoint $\Phi(X) \in \mathcal{A}_X$; the limit is symmetric if each $\Phi^{(n)}$ is. Fix $\varepsilon > 0$ and N with $\|\Phi^{(n)} - \Phi^{(m)}\|_F \leq \varepsilon$ for $n, m \geq N$. For every pair x, y and every finite truncation of the sum $\sum_{X \ni x,y}$, letting $m \rightarrow \infty$ gives $F(\text{dist}(x,y))^{-1} \sum_{X \ni x,y} \|\Phi^{(n)}(X) - \Phi(X)\| \leq \varepsilon$ for $n \geq N$, uniformly in x, y ; hence $\|\Phi^{(n)} - \Phi\|_F \leq \varepsilon$ and $\Phi \in \mathcal{B}_F$ with $\Phi^{(n)} \rightarrow \Phi$. Thus $\mathcal{B}_F$ is complete. $\square$

We record the two elementary structures that condensation will act on. The self-adjointness constraint makes $\mathcal{B}_F$ a real (not complex) Banach space; scaling an interaction by a real number and adding interactions termwise are continuous, so $\mathcal{B}_F$ is in particular a topological real vector space and an abelian topological group under addition.

3 The Heisenberg dynamics and the Lieb–Robinson bound

3.1 Finite-volume dynamics and the infinite-volume limit

For a finite $\Lambda \subseteq L$ the *local Hamiltonian* of $\Phi \in \mathcal{B}_F$ is the self-adjoint operator

$$H_\Lambda^\Phi = \sum_{X \subseteq \Lambda} \Phi(X) \in \mathcal{A}_\Lambda,$$

a finite sum, and the finite-volume *Heisenberg dynamics* is the one-parameter group of $*$ -automorphisms

$$\tau_t^{\Phi, \Lambda}(A) = e^{itH_\Lambda^\Phi} A e^{-itH_\Lambda^\Phi}, \quad A \in \mathcal{A}.$$

The Lieb–Robinson bound controls the spatial spread of $\tau_t^{\Phi, \Lambda}(A)$ and is the key to removing the cutoff Λ .

Theorem 3.1 (Lieb–Robinson bound, F -function form [18]). *Let $\Phi \in \mathcal{B}_F$ and let $X, Y \subseteq L$ be disjoint. For $A \in \mathcal{A}_X$, $B \in \mathcal{A}_Y$, and every finite $\Lambda \supseteq X \cup Y$,*

$$\|[\tau_t^{\Phi, \Lambda}(A), B]\| \leq \frac{2\|A\| \|B\|}{C_F} \left(e^{2\|\Phi\|_F C_F |t|} - 1 \right) \sum_{x \in X} \sum_{y \in Y} F(\text{dist}(x, y)),$$

with $\|F\|$ and C_F as in Theorem 2.2. The bound is uniform in Λ .

This is Theorem 3.4 of [18] in the present notation; we take it as given. The velocity is made explicit by combining it with reweighting.

Proposition 3.2 (Light cone and velocity). *Let $\Phi \in \mathcal{B}_{F_a}$ for the reweighted function $F_a(r) = e^{-ar} F(r)$ of Theorem 2.3, with $a > 0$. Then for disjoint X, Y and $A \in \mathcal{A}_X$, $B \in \mathcal{A}_Y$,*

$$\|[\tau_t^{\Phi, \Lambda}(A), B]\| \leq \frac{2\|A\| \|B\| \|F\|}{C_{F_a}} \min\{|X|, |Y|\} e^{-a(\text{dist}(X, Y) - v|t|)}, \quad v = \frac{2\|\Phi\|_{F_a} C_{F_a}}{a},$$

where $\text{dist}(X, Y) = \min_{x \in X, y \in Y} \text{dist}(x, y)$. In particular the commutator is exponentially small outside the light cone $\text{dist}(X, Y) \geq v|t|$.

Proof. Apply Theorem 3.1 with F_a in place of F . For the geometric sum, each $x \in X$ has $\sum_{y \in Y} F_a(\text{dist}(x, y)) \leq e^{-a \text{dist}(X, Y)} \sum_{y \in L} F(\text{dist}(x, y)) \leq \|F\| e^{-a \text{dist}(X, Y)}$ using $e^{-a \text{dist}(x, y)} \leq e^{-a \text{dist}(X, Y)}$; summing over the smaller of X, Y gives the factor $\min\{|X|, |Y|\}$. Finally $e^{2\|\Phi\|_{F_a} C_{F_a} |t|} - 1 \leq e^{2\|\Phi\|_{F_a} C_{F_a} |t|} = e^{av|t|}$, and the two exponentials combine. $\square$

Theorem 3.3 (Banach interaction space and strongly continuous LR dynamics, I-A). *Fix an F -function F and $\Phi \in \mathcal{B}_F$.*

1. $\mathcal{B}_F$ is a real Banach space (Theorem 2.6).
2. For each $A \in \mathcal{A}$ the limit $\tau_t^\Phi(A) = \lim_{\Lambda \uparrow L} \tau_t^{\Phi, \Lambda}(A)$ exists in operator norm, uniformly for t in compact sets, and defines a strongly continuous one-parameter group $\{\tau_t^\Phi\}_{t \in \mathbb{R}}$ of $*$ -automorphisms of $\mathcal{A}$.
3. The infinite-volume dynamics satisfies the Lieb–Robinson bound of Theorem 3.1 and the light-cone estimate of Theorem 3.2 with the same constants (the finite-volume bounds being uniform in Λ).

Proof. Part (1) is Theorem 2.6. For part (2): for local A and finite volumes $\Lambda \subseteq \Lambda'$, the difference $\tau_t^{\Phi, \Lambda'}(A) - \tau_t^{\Phi, \Lambda}(A)$ is expressed by Duhamel's formula as a time integral of a commutator of $H_{\Lambda'}^\Phi - H_\Lambda^\Phi$ with $\tau_s^{\Phi, \Lambda}(A)$; the Lieb–Robinson bound of Theorem 3.1 bounds that commutator by the tail $\sum_{X \ni x, X \not\subseteq \Lambda} \|\Phi(X)\|$ against F , which is a convergent tail of a

$\|\cdot\|_F$ -finite sum and hence tends to 0 as $\Lambda \uparrow L$, uniformly for t in compacts. This is the Cauchy criterion, giving a limit $\tau_t^\Phi(A)$ for local A ; the estimate is uniform in A on norm-balls, so the limit extends to all of $\mathcal{A} = \overline{\mathcal{A}_{\text{loc}}}$ by density and is an isometric $*$ -homomorphism. The group law and strong continuity pass to the limit. This is the content of [18, §3–4]; see also [8] for the argument in the exponential case and [11] for the original velocity. Part (3) is the uniformity in Λ of Theorems 3.1 and 3.2, which survives the limit. $\square$

The proof of Theorem 3.3(2) already contains the estimate we need for the condensed statement, because it is really a statement about how τ depends on Φ . We isolate it.

Lemma 3.4 (Continuity of the dynamics in the interaction). *Fix a local observable $A \in \mathcal{A}_X$ and a time horizon $T > 0$. On the closed ball $\mathcal{B}_F(B) = \{\Phi : \|\Phi\|_F \leq B\}$ there is a constant $K = K(A, T, B, F) < \infty$ such that*

$$\sup_{|t| \leq T} \|\tau_t^\Phi(A) - \tau_t^\Psi(A)\| \leq K \|\Phi - \Psi\|_F \quad \text{for all } \Phi, \Psi \in \mathcal{B}_F(B).$$

Moreover $t \mapsto \tau_t^\Phi(A)$ is norm-continuous, so $(t, \Phi) \mapsto \tau_t^\Phi(A)$ is jointly continuous $\mathbb{R} \times \mathcal{B}_F(B) \rightarrow \mathcal{A}$.

Proof. Work first in a finite volume $\Lambda \supseteq X$ and interpolate along the straight line $\Phi_r = \Psi + r(\Phi - \Psi)$, $r \in [0, 1]$. Duhamel gives

$$\tau_t^{\Phi, \Lambda}(A) - \tau_t^{\Psi, \Lambda}(A) = i \int_0^t \tau_s^{\Phi, \Lambda}([H_\Lambda^\Phi - H_\Lambda^\Psi, \tau_{t-s}^{\Psi, \Lambda}(A)]) ds.$$

The generator difference is $H_\Lambda^\Phi - H_\Lambda^\Psi = \sum_{Z \subseteq \Lambda} (\Phi - \Psi)(Z)$. Bounding the commutator of each term $(\Phi - \Psi)(Z)$ with the Lieb–Robinson-evolved $\tau_{t-s}^{\Psi, \Lambda}(A)$ by Theorem 3.1, and summing the resulting weights against F , produces

$$\|[H_\Lambda^\Phi - H_\Lambda^\Psi, \tau_{t-s}^{\Psi, \Lambda}(A)]\| \leq 2\|A\| \|\Phi - \Psi\|_F |X| \|F\| e^{2BC_F|t-s|},$$

where the exponential comes from the Lieb–Robinson factor at $\|\Psi\|_F \leq B$ and the factor $|X|\|F\|$ from summing F over the support of A against all sites. Integrating over $|s| \leq |t| \leq T$ gives the claim in finite volume with $K = 2\|A\|\|X\|\|F\|T e^{2BC_FT}$, uniform in Λ ; letting $\Lambda \uparrow L$ using Theorem 3.3(2) gives it for the infinite-volume dynamics. Norm-continuity in t is the strong continuity of Theorem 3.3(2). Joint continuity follows from the two one-variable estimates and the triangle inequality. $\square$

3.2 The generator as a symmetric derivation

The strongly continuous group τ^Φ has an infinitesimal generator, and the F -norm controls it on local observables. This is the infinitesimal counterpart of the whole picture: the same single number $\|\Phi\|_F$ that bounds the dynamics also bounds its generator, term by term.

Proposition 3.5 (Symmetric derivation generating the dynamics). *Fix $\Phi \in \mathcal{B}_F$. For local $A \in \mathcal{A}_X$ the sum*

$$\delta_\Phi(A) = i \sum_{Z: Z \cap X \neq \emptyset} [\Phi(Z), A]$$

converges in operator norm, with $\|\delta_\Phi(A)\| \leq 2|X|F(0)\|\Phi\|_F\|A\|$, and defines a symmetric derivation δ_Φ on the dense $*$ -subalgebra $\mathcal{A}_{\text{loc}}$. Its closure generates τ^Φ , i.e. $\tau_t^\Phi = \exp(t\overline{\delta_\Phi})$ as strongly continuous groups, and $\frac{d}{dt}\big|_{t=0}\tau_t^\Phi(A) = \delta_\Phi(A)$ for local A .

Proof. For $A \in \mathcal{A}_X$ only clusters Z meeting X contribute. Bounding each commutator by $2\|\Phi(Z)\|\|A\|$ and using $\sum_{Z \ni x}\|\Phi(Z)\| \leq F(0)\|\Phi\|_F$ for each $x \in X$ —this is the F -norm evaluated at the diagonal pair $y = x$, since $F(0)^{-1}\sum_{Z \ni x}\|\Phi(Z)\| \leq \|\Phi\|_F$ —gives

$$\|\delta_\Phi(A)\| \leq 2\|A\| \sum_{x \in X} \sum_{Z \ni x} \|\Phi(Z)\| \leq 2|X|F(0)\|\Phi\|_F\|A\|,$$

so the sum converges absolutely. The Leibniz rule $\delta_\Phi(AB) = \delta_\Phi(A)B + A\delta_\Phi(B)$ and the symmetry $\delta_\Phi(A^*) = \delta_\Phi(A)^*$ are inherited from the commutator, so δ_Φ is a symmetric derivation on $\mathcal{A}_{\text{loc}}$. That its closure generates τ^Φ and that the group differentiates to δ_Φ on local elements is the standard generation theorem for quasi-local dynamics [18, 8]; the Lieb–Robinson bound of Theorem 3.3 is exactly what makes the finite-volume generators $\delta_\Phi^\Lambda(A) = i[H_\Lambda^\Phi, A]$ converge to $\delta_\Phi(A)$ as $\Lambda \uparrow L$. $\square$

The assignment $\Phi \mapsto \delta_\Phi$ is real-linear and, on each $\mathcal{A}_X$, bounded in the F -norm: $\|\delta_\Phi(A) - \delta_\Psi(A)\| \leq 2|X|F(0)\|\Phi - \Psi\|_F\|A\|$. This is the infinitesimal shadow of Theorem 3.4; once condensed it says that the passage to generators is itself a morphism $\underline{\mathcal{B}}_F \rightarrow \underline{\text{Der}(\mathcal{A}_{\text{loc}})}$, continuous in the F -norm, of which Theorem 4.4 is the exponentiated form.

Example 3.6 (Transverse-field Ising chain). Take $L = \mathbb{Z}$, $\mathcal{H}_x = \mathbb{C}^2$, and

$$\Phi_{\text{TFIM}}(\{x, x+1\}) = -J\sigma_x^z\sigma_{x+1}^z, \quad \Phi_{\text{TFIM}}(\{x\}) = -h\sigma_x^x,$$

all other $\Phi(X) = 0$. This is finite range, so with any F -function—for instance F_ε of Theorem 2.4—the F -norm is finite. The heaviest diagonal pair $x = y$ collects the on-site field together with the two incident bonds, contributing $(|h| + 2|J|)/F(0)$; a nearest-neighbour pair collects the single bond between them, contributing $|J|/F(1)$; and no pair at distance ≥ 2 contributes. Since F is non-increasing, $F(1) \leq F(0)$, so

$$\|\Phi_{\text{TFIM}}\|_F = \max\left\{\frac{|h|+2|J|}{F(0)}, \frac{|J|}{F(1)}\right\} \leq \frac{2|J|+|h|}{F(1)} < \infty,$$

and Theorem 3.3 applies for all J, h . The velocity bound of Theorem 3.2 is $v = 2\|\Phi\|_{F_a}C_{F_a}/a$; its optimization over the reweighting parameter a produces a finite group velocity, which we exhibit numerically in Section 8. The chain is gapped except on the critical line $|J| = |h|$; in the language of the program the critical line is a slice of the gapless discriminant Σ_f , and the two gapped phases sit in different components of $\mathfrak{Phase}_{d,G}$. The present paper says nothing about the gap—that is Part III—but it does guarantee that the dynamics on both sides has the same uniform light cone as long as J, h stay bounded.

4 Condensation of the interaction space

4.1 Condensed sets and profinite probes

We recall only what we use. A *profinite set* is a cofiltered limit $S = \varprojlim_i S_i$ of finite sets, equivalently a compact, Hausdorff, totally disconnected topological space. The *site of profinite*

sets has covers the finite jointly surjective families; a *condensed set* is a sheaf on this site (equivalently, on the subcategory of extremally disconnected profinite sets, where sheaf and product-preserving presheaf coincide). Condensed abelian groups form a Grothendieck abelian category, so kernels, cokernels, and derived functors behave as in ordinary homological algebra [21]; the pyknotic formulation is [3]. Write **Cond** for the category of condensed sets.

Definition 4.1 (Condensation). For a topological space X , its *condensation* $\underline{X} \in \mathbf{Cond}$ is

$$\underline{X}(S) = \text{Cont}(S, X), \quad S \text{ profinite},$$

the set of continuous maps $S \rightarrow X$, with the evident restriction along maps of profinite sets. For a finite set S this is just X^S , so $\underline{X}(*) = X$ recovers the points.

The one structural fact we lean on is full faithfulness. A topological space is *compactly generated* if its topology is final for the maps out of compact Hausdorff spaces into it; metrizable spaces and, more generally, first-countable spaces are compactly generated.

Theorem 4.2 (Full faithfulness of condensation [21]). *The functor $X \mapsto \underline{X}$ from compactly generated Hausdorff spaces to condensed sets is fully faithful and preserves finite products. In particular, for compactly generated Hausdorff X, Y ,*

$$\text{Cont}(X, Y) \xrightarrow{\sim} \text{Hom}_{\mathbf{Cond}}(\underline{X}, \underline{Y}),$$

so a continuous map is the same datum as a morphism of condensed sets, and $\underline{X \times Y} \cong \underline{X} \times \underline{Y}$.

Two comments fix the size of the formalism we need. First, a Banach space V is metrizable, hence compactly generated, so Theorem 4.2 applies to it; moreover $\underline{V}$ is a module over the condensed ring $\mathbb{R}$, i.e. a *condensed $\mathbb{R}$ -vector space*, because addition $V \times V \rightarrow V$ and scaling $\mathbb{R} \times V \rightarrow V$ are continuous and condensation preserves finite products. Second, for separable target spaces one may replace profinite probes by their *light* (metrizable, equivalently second-countable) counterparts and work in the category $\mathbf{Cond}^{\text{light}}$ of light condensed sets [6]; this avoids the set-theoretic subtleties of the full theory (no inaccessible cardinal is needed) and is the correct home for separable spin systems, whose observable algebra $\mathcal{A}$ is separable by Theorem 2.1.

4.2 The condensed interaction space and the condensed dynamics

Proposition 4.3 (Condensed structure of $\mathcal{B}_F$). *$\underline{\mathcal{B}_F}$ is a condensed real vector space and a condensed abelian group. For a profinite set S ,*

$$\underline{\mathcal{B}_F}(S) = \text{Cont}(S, \mathcal{B}_F)$$

is the real Banach space of continuous $\mathcal{B}_F$ -valued functions on S with the supremum norm $\|\Phi_\bullet\|_{\infty, F} = \sup_{s \in S} \|\Phi_s\|_F$, and the vector-space operations are pointwise.

Proof. $\mathcal{B}_F$ is a real Banach space (Theorem 2.6), so addition and real scaling are continuous; condensation preserves finite products (Theorem 4.2), so it carries these to morphisms $\underline{\mathcal{B}_F} \times \underline{\mathcal{B}_F} \rightarrow \underline{\mathcal{B}_F}$ and $\mathbb{R} \times \underline{\mathcal{B}_F} \rightarrow \underline{\mathcal{B}_F}$ satisfying the vector-space axioms internally. The description of $\underline{\mathcal{B}_F}(S)$ is Theorem 4.1; $\text{Cont}(S, \mathcal{B}_F)$ with S compact and $\mathcal{B}_F$ Banach is complete in the sup-norm, hence itself Banach. $\square$

We now assemble the dynamics. Let $\text{Aut}(\mathcal{A})$ carry the topology of strong convergence: a net $\beta_i \rightarrow \beta$ iff $\beta_i(A) \rightarrow \beta(A)$ in norm for every $A \in \mathcal{A}$. Because $\mathcal{A}$ is separable, this topology is metrizable and $\text{Aut}(\mathcal{A})$ is a Polish group, in particular compactly generated, so Theorem 4.2 applies to it.

Theorem 4.4 (Condensed dynamics morphism, I-B). *Let F be an F -function, $\mathcal{A}$ the associated quasi-local algebra, and $\text{Aut}(\mathcal{A})$ its automorphism group in the strong topology.*

1. $\underline{\mathcal{B}}_F$ is a condensed $\mathbb{R}$ -vector space (Theorem 4.3).
2. The assignment $D : (t, \Phi) \mapsto \tau_t^\Phi$ is a continuous map $\mathbb{R} \times \mathcal{B}_F \rightarrow \text{Aut}(\mathcal{A})$.
3. Consequently D condenses to a morphism of condensed sets

$$\underline{D} : \mathbb{R} \times \underline{\mathcal{B}}_F \longrightarrow \underline{\text{Aut}(\mathcal{A})}, \quad \mathbb{R} \times \underline{\mathcal{B}}_F \cong \underline{\mathbb{R} \times \mathcal{B}_F},$$

and $\underline{D}$ is a morphism of light condensed sets whenever the interactions are drawn from a separable closed subspace of $\mathcal{B}_F$ (for instance the finite-range interactions, or those with a fixed countable generating family of clusters).

Proof. Part (1) is Theorem 4.3. For part (2), fix $A \in \mathcal{A}$ and $\varepsilon > 0$; we show $(t, \Phi) \mapsto \tau_t^\Phi(A)$ is continuous, which is the definition of strong continuity of D . Choose a local A' with $\|A - A'\| < \varepsilon/3$ (possible since $\mathcal{A}_{\text{loc}}$ is dense). Automorphisms are isometric, so $\|\tau_t^\Phi(A) - \tau_{t'}^{\Phi'}(A)\| \leq 2\|A - A'\| + \|\tau_t^\Phi(A') - \tau_{t'}^{\Phi'}(A')\| < \frac{2\varepsilon}{3} + \|\tau_t^\Phi(A') - \tau_{t'}^{\Phi'}(A')\|$. By Theorem 3.4 the second term is small when $|t - t'|$ and $\|\Phi - \Phi'\|_F$ are small, on any ball $\|\Phi\|_F, \|\Phi'\|_F \leq B$. Hence D is continuous at (t, Φ) . Part (3): $\mathbb{R} \times \mathcal{B}_F$ is metrizable, hence compactly generated, and so is $\text{Aut}(\mathcal{A})$ by separability of $\mathcal{A}$; Theorem 4.2 turns the continuous D into a morphism $\underline{D}$ and identifies $\mathbb{R} \times \underline{\mathcal{B}}_F$ with $\underline{\mathbb{R} \times \mathcal{B}_F}$. When the interactions lie in a separable closed subspace $V_0 \subseteq \mathcal{B}_F$, both $\mathbb{R} \times V_0$ and $\text{Aut}(\mathcal{A})$ are separable metrizable, so the restricted morphism is one of light condensed sets [6]. $\square$

Theorem 4.4 is the sense in which “the dynamics is a morphism, not just an assignment.” The diagram

$$\begin{array}{ccc} \mathbb{R} \times \underline{\mathcal{B}}_F & \xrightarrow{\underline{D}} & \underline{\text{Aut}(\mathcal{A})} \\ \text{pr}_2 \downarrow & \nearrow \Phi \mapsto \tau_\bullet^\Phi & \\ \underline{\mathcal{B}}_F & & \end{array}$$

commutes internally in **Cond**: freezing a profinite family of interactions and letting time vary is a $\mathbb{R}$ -action on $\underline{\text{Aut}(\mathcal{A})}$ over $\underline{\mathcal{B}}_F$. This action is the object that Part IV group-completes on the invertible gapped sector.

Remark 4.5 (Why the Banach topology is not negotiable). The continuity in Theorem 4.4(2) is stated for the $\|\cdot\|_F$ -topology on $\mathcal{B}_F$, and Theorem 3.4 genuinely uses that norm: the Lieb–Robinson factor $e^{2BC_F|t|}$ and the weight $\|F\|$ are finite only because $\|\Phi - \Psi\|_F$ measures interaction differences against the same F -function. A weaker topology (pointwise convergence of $\Phi(X)$ for each cluster X , say) does not control the tail sums and does not make D continuous. This is why the ground floor of the program is $\underline{\mathcal{B}}_F$ for a fixed F and not the condensation of some coarser space of formal interactions; long-range interactions outside every F -function class can violate the finite-velocity bound and fall outside the theorem.

5 Uniform Lieb–Robinson bounds over a profinite base

The previous section made the dynamics a morphism out of $\mathcal{B}_F$. We now show that when one plugs in an actual profinite family, the Lieb–Robinson estimate holds simultaneously and uniformly, and we identify what an S -continuous family of interactions is in concrete terms.

5.1 Uniformity over the base

Theorem 5.1 (Profinite families and the uniform light cone, I-C). *Let $S = \varprojlim_i S_i$ be a profinite set and $\Phi_\bullet \in \text{Cont}(S, \mathcal{B}_F) = \underline{\mathcal{B}}_F(S)$ a continuous family of interactions.*

1. (Compatible finite data.) $\Phi_\bullet$ is a norm limit, in the sup-norm $\|\cdot\|_{\infty, F}$ of Theorem 4.3, of locally constant families: there are indices $i_1 \leq i_2 \leq \dots$ and interactions $\Phi_\bullet^{(k)}$ each factoring through the finite quotient $S \rightarrow S_{i_k}$ with $\|\Phi_\bullet - \Phi_\bullet^{(k)}\|_{\infty, F} \rightarrow 0$. Equivalently, $\underline{\mathcal{B}}_F(S)$ is the completion of $\varinjlim_i \text{Cont}(S_i, \mathcal{B}_F)$ in the sup-norm.
2. (Uniform bound.) $B := \sup_{s \in S} \|\Phi_s\|_F < \infty$, attained because S is compact and $s \mapsto \|\Phi_s\|_F$ is continuous.
3. (Uniform light cone.) If in addition $\Phi_s \in \mathcal{B}_{F_a}$ with $B_a := \sup_{s \in S} \|\Phi_s\|_{F_a} < \infty$ for the reweighting F_a , then the Lieb–Robinson bound of Theorem 3.2 holds for every $s \in S$ simultaneously with one velocity

$$v = \frac{2 B_a C_{F_a}}{a}$$

and prefactors independent of s . Consequently the map $s \mapsto \tau_t^{\Phi_s}$ is a morphism $\underline{S} \rightarrow \underline{\text{Aut}}(\mathcal{A})$ obtained by restricting $\underline{D}$ of Theorem 4.4 along $\Phi_\bullet$, and it factors through the sub-condensed set of automorphisms with light cone of slope at most v .

Proof. (1) A continuous map from a profinite (compact, totally disconnected) space to a metric space is a uniform limit of locally constant maps. Concretely, fix k and use uniform continuity of $\Phi_\bullet$ on the compact S : there is a finite clopen partition of S on which $\Phi_\bullet$ varies by at most $1/k$ in $\|\cdot\|_F$; refining, the partition is pulled back from some finite quotient S_{i_k} , and choosing one value of $\Phi_\bullet$ per block defines $\Phi_\bullet^{(k)}$ factoring through S_{i_k} with $\|\Phi_\bullet - \Phi_\bullet^{(k)}\|_{\infty, F} \leq 1/k$. Locally constant families are exactly the elements of $\varinjlim_i \text{Cont}(S_i, \mathcal{B}_F)$ (a locally constant map factors through a finite quotient), and their sup-norm closure is all of $\text{Cont}(S, \mathcal{B}_F)$ by the density just shown; completeness of $\text{Cont}(S, \mathcal{B}_F)$ (Theorem 4.3) identifies it with the completion of the colimit.

(2) $s \mapsto \|\Phi_s\|_F$ is continuous because $\Phi_\bullet$ is continuous and the norm is 1-Lipschitz; a continuous real function on a compact set is bounded and attains its supremum.

(3) The estimate of Theorem 3.2 depends on the interaction *only* through $\|\Phi\|_{F_a}$, entering the velocity $v = 2\|\Phi\|_{F_a} C_{F_a}/a$ and the prefactor not at all. Replacing $\|\Phi_s\|_{F_a}$ by its uniform bound B_a gives a single velocity $v = 2B_a C_{F_a}/a$ valid for all s ; the prefactor $2\|A\| \|B\| \|F\| C_{F_a}^{-1} \min\{|X|, |Y|\}$ is already s -independent. Thus every $\tau_t^{\Phi_s}$ obeys the same light-cone estimate. That $s \mapsto \tau_t^{\Phi_s}$ is a morphism of condensed sets is Theorem 4.4(3) precomposed with the point $\underline{S} \rightarrow \underline{\mathcal{B}}_F$ named by $\Phi_\bullet$; the light-cone slope bound is preserved because it holds pointwise with the uniform v . $\square$

The three parts say, in order: an S -continuous family is a compatible tower of finite-resolution families; such a family is automatically norm-bounded; and a norm bound is exactly what makes the whole family share one light cone. The last point is the reason the condensed reformulation is not merely cosmetic. In the topological category one would have a map $S \rightarrow \mathcal{I}_{d,G}$ and a velocity function $s \mapsto v(s)$ with no guarantee of a common bound; the Banach structure of $\mathcal{B}_F$ upgrades “finite velocity at each s ” to “one finite velocity for all s ”, which is what a morphism into a condensed set of quasi-local dynamics requires.

The uniformity is compatible with change of probe, which is what makes it useful for descent.

Proposition 5.2 (Naturality of the dynamics over the base). *For a continuous map $\varphi : S' \rightarrow S$ of profinite sets, pullback of families $\varphi^* : \underline{\mathcal{B}}_F(S) \rightarrow \underline{\mathcal{B}}_F(S')$, $\Phi_\bullet \mapsto \Phi_{\varphi(\cdot)}$, commutes with the passage to dynamics: for each fixed t the square*

$$\begin{array}{ccc} \underline{\mathcal{B}}_F(S) & \xrightarrow{\tau_t} & \underline{\text{Aut}}(\mathcal{A})(S) \\ \varphi^* \downarrow & & \downarrow \varphi^* \\ \underline{\mathcal{B}}_F(S') & \xrightarrow{\tau_t} & \underline{\text{Aut}}(\mathcal{A})(S') \end{array}$$

commutes, and the uniform bound transfers, $\sup_{s' \in S'} \|\Phi_{\varphi(s')}\|_F \leq \sup_{s \in S} \|\Phi_s\|_F$, so the pulled-back family has light cone of slope at most that of the original. In particular the finite-quotient approximations of Theorem 5.1(1) form a compatible system of quasi-local dynamics with a common velocity.

Proof. Commutativity is naturality of the morphism τ_t of condensed sets (Theorem 4.4) evaluated on φ : a morphism of condensed sets is a natural transformation of the underlying functors, and its naturality square on the map $\varphi : S' \rightarrow S$ is exactly the diagram. The bound transfers because $\{\Phi_{\varphi(s')} : s' \in S'\} \subseteq \{\Phi_s : s \in S\}$ as interactions, so its supremum F -norm cannot exceed that of the original family; the light-cone slope $v = 2B_a C_{F_a}/a$ is monotone in the uniform bound B_a by Theorem 5.1(3). $\square$

Corollary 5.3 (Disorder families). *Let $\Omega = Q^{\mathbb{Z}^d}$ be a profinite disorder hull for a finite local-configuration alphabet Q (a compact, totally disconnected space); here Q denotes the configuration alphabet, and F is reserved for the F -function. A disordered interaction $\omega \mapsto \Phi_\omega$ with $\sup_{\omega \in \Omega} \|\Phi_\omega\|_{F_a} < \infty$ is an element $\Phi_\bullet \in \underline{\mathcal{B}}_F(\Omega)$ and hence, by Theorem 5.1, has a Ω -uniform light cone. In particular the disorder-averaged and almost-sure Lieb–Robinson velocities coincide with the deterministic bound $v = 2(\sup_{\omega} \|\Phi_\omega\|_{F_a})C_{F_a}/a$.*

Proof. The map $\omega \mapsto \Phi_\omega$ is continuous into $\mathcal{B}_F$ by hypothesis (uniform-norm continuity), so it is an element of $\text{Cont}(\Omega, \mathcal{B}_F) = \underline{\mathcal{B}}_F(\Omega)$; apply Theorem 5.1(3) with $S = \Omega$. The averaged and almost-sure velocities are bounded by the uniform one because the estimate is deterministic and holds for every ω . $\square$

This corollary is the locality-layer appearance of the disorder thread that runs through the series: the profinite probe Ω is the physical configuration space, not a repackaging of a smooth manifold, and Theorem 5.3 is what Theorem 7.3 below proposes to sharpen into a descent statement.

5.2 Uniform clustering: the bridge to the gapped substack

Part III needs not just the light cone but its standard spectral consequence: a uniform gap forces uniform exponential decay of ground-state correlations. We state the version that our uniformity provides; the single-Hamiltonian statement is due to Hastings–Koma [9] and Nachtergaele–Sims [17].

Proposition 5.4 (Uniform exponential clustering, input to Part III). *Let $\Phi_\bullet \in \mathcal{B}_{F_a}(S)$ over a profinite S with uniform bound $B_a = \sup_s \|\Phi_s\|_{F_a} < \infty$, and suppose the family is uniformly gapped: there is $\Delta > 0$ with $\text{gap}(H^{\Phi_s}) \geq \Delta$ for all $s \in S$, where $\text{gap}(\cdot)$ is the spectral gap above the ground state of the infinite-volume GNS Hamiltonian. Then there are constants $C, \mu > 0$, depending only on B_a, C_{F_a}, a , and Δ —and in particular not on s —such that for all local A, B with disjoint supports and all $s \in S$,*

$$|\langle \Omega_s | AB | \Omega_s \rangle - \langle \Omega_s | A | \Omega_s \rangle \langle \Omega_s | B | \Omega_s \rangle| \leq C \|A\| \|B\| e^{-\mu \text{dist}(\text{supp } A, \text{supp } B)},$$

where Ω_s is the ground state of H^{Φ_s} .

Proof. For a fixed s this is the Hastings–Koma / Nachtergaele–Sims exponential clustering theorem [9, 17]: a spectral gap Δ together with a Lieb–Robinson velocity v yields correlation decay with rate $\mu = \text{const} \cdot \Delta / (v + \Delta)$ and constant C depending on the same data. The only point to check is uniformity. By Theorem 5.1(3) the velocity $v = 2B_a C_{F_a} / a$ and the Lieb–Robinson prefactors are the same for every s ; by hypothesis the gap lower bound Δ is the same for every s . The clustering rate and constant are explicit functions of (v, Δ, C_{F_a}, a) only, so they are s -independent. $\square$

Theorem 5.4 is stated conditionally on a uniform gap because deciding the gap is not a locality question and, by the Cubitt–Pérez-García–Wolf undecidability theorem [7]¹, cannot be. Our contribution is exactly the clause “and in particular not on s ”: the locality estimates are uniform, so whatever gap hypothesis Part III imposes, its clustering consequence is automatically uniform over the base.

6 Quasi-adiabatic continuation and the class $\mathcal{W}$

The moduli-stack picture inverts a class $\mathcal{W}$ of gapped adiabatic, quasi-local equivalences to define phases. At the level of a single gapped path the rigorous content of $\mathcal{W}$ is quasi-adiabatic continuation (Hastings–Wen [10]) and the automorphic equivalence of Bachmann–Michalakis–Nachtergaele–Sims [2]. We recall these as the morphism-level tools they are, staying within what is proved. Throughout we keep the Heisenberg picture fixed in Section 3.1, where automorphisms act by $A \mapsto U^* A U$, so that $\tau_t^\Phi(A) = e^{itH^\Phi} A e^{-itH^\Phi}$; this convention is what puts the positive sign in the commutator flows below.

¹Cited for scope, not used: there is no algorithm deciding whether a translation-invariant nearest-neighbour family is gapped. Part III therefore treats “uniformly gapped” as the hypothesis defining $\mathfrak{Gap}_{d,G}$, never as a conclusion.

Definition 6.1 (Quasi-adiabatic generator). Let $[0, 1] \ni r \mapsto \Phi(r) \in \mathcal{B}_{F_a}$ be a norm- C^1 path of interactions, uniformly gapped with gap $\geq \Delta > 0$, and let P_r be the ground-state projection of $H^{\Phi(r)}$. The *quasi-adiabatic generator* is

$$\mathcal{D}(r) = \int_{-\infty}^{\infty} W_{\gamma}(t) \tau_t^{\Phi(r)} \left(\frac{d}{dr} H^{\Phi(r)} \right) dt,$$

where W_{γ} is a Hastings weight function whose Fourier transform is supported outside $(-\Delta, \Delta)$ and decays faster than any polynomial (indeed almost exponentially).

Proposition 6.2 (Quasi-adiabatic continuation generates $\mathcal{W}$ [10, 2]). *With the hypotheses of Theorem 6.1:*

1. $\mathcal{D}(r)$ is a quasi-local self-adjoint operator: its finite-volume approximations obey a Lieb–Robinson bound for a reweighted F -function $F_{a'}$ with $a' < a$, and $\|\mathcal{D}(r)\|_{F_a}$ is bounded uniformly in r in terms of Δ and $\sup_r \|\frac{d}{dr} \Phi(r)\|_{F_a}$.
2. The flow α_r generated by $\{\mathcal{D}(r)\}$ (the solution of $\frac{d}{dr} \alpha_r = i \alpha_r \circ [\mathcal{D}(r), \cdot]$, $\alpha_0 = \text{id}$) is a strongly continuous family of $*$ -automorphisms of $\mathcal{A}$ with $\alpha_r(P_0) = P_r$ for all r ; equivalently α_r intertwines the ground states along the path.
3. Two interactions connected by such a uniformly gapped path are automorphically equivalent: there is a quasi-local automorphism α_1 , obeying a Lieb–Robinson bound, with α_1 mapping the ground state of $H^{\Phi(0)}$ to that of $H^{\Phi(1)}$. This automorphism is the generator of $\mathcal{W}$ for the pair.

Proof (citation). Parts (1)–(2) are the construction of Hastings–Wen [10], made quantitative in the F -function calculus of Nachtergaele–Sims–Young [18]: the quasi-locality of $\mathcal{D}(r)$ follows from the Lieb–Robinson bound of Theorem 3.3 for $\tau_t^{\Phi(r)}$ and the fast decay of W_{γ} , and the intertwining property is the defining feature of the Hastings weight. Part (3) is the automorphic equivalence theorem of Bachmann–Michalakis–Nachtergaele–Sims [2]: integrating the flow of part (2) over $r \in [0, 1]$ yields the quasi-local α_1 . We claim nothing beyond these results; in particular we do not assert that α_1 assembles over a profinite base into a condensed automorphism, which is the content of Theorem 7.2. $\square$

Remark 6.3 (What $\mathcal{W}$ is, and is not, here). $\mathcal{W}$ is part of the *definition* of the phase problem, not a theorem: isomorphism, gapped homotopy, stable equivalence, and K -theoretic equivalence need not coincide, and the program must say which one it inverts. Theorem 6.2 pins down the adiabatic, quasi-local generators of $\mathcal{W}$ for a single gapped path. Turning these generators into an honest morphism structure over $\underline{\mathcal{B}}_F$ (a condensed group acting on $\mathfrak{Ham}_{d,G}$ with quasi-adiabatic continuation as internal path-lifting) is beyond the present estimates and is posed as Theorem 7.2.

7 Conjectures: toward the condensed stack

The theorems above make $\underline{\mathcal{B}}_F$ a condensed $\mathbb{R}$ -vector space and the dynamics a morphism out of it, uniformly over profinite probes. They do *not* establish that the whole moduli problem

$\mathfrak{Ham}_{d,G}$ (which remembers gauge and quasi-local automorphisms, not just interactions) is a condensed stack. We isolate the three statements that the locality layer would need, and label them as conjectures with stable numbers.

Conjecture 7.1 (Condensed higher stack, I-1). The moduli problem $\mathfrak{Ham}_{d,G}$ of G -symmetric quasi-local Hamiltonians is a condensed higher stack, not merely a condensed set: the assignment sending a profinite probe S to the groupoid (or ∞ -groupoid) of S -families of Hamiltonians and their gauge and quasi-local automorphisms satisfies descent along finite jointly surjective covers of profinite sets. Its π_0 over the point recovers the condensed set $\underline{\mathcal{B}}_F$ modulo relabelling, and its automorphism sheaf is generated by the quasi-local flows of Theorem 6.2.

What is proved here is the underlying-condensed-set statement (Theorem 4.3) and the morphism property of dynamics (Theorem 4.4); the gluing of the automorphism groupoid along covers is open. We flag, as the knowledge base insists, that this stack structure must not be asserted as a theorem anywhere in the series.

Conjecture 7.2 (Condensed group of quasi-local automorphisms, I-2). The quasi-local automorphisms generated by time-dependent interactions in a fixed $\mathcal{B}_{F_a}$ -ball, with the Lieb–Robinson bound of Theorem 6.2(1), assemble into a condensed group $\underline{\mathcal{G}}_{F_a}$ acting on $\underline{\mathcal{B}}_F$ (and on $\mathfrak{Ham}_{d,G}$). Under this action, quasi-adiabatic continuation is an internal path-lifting: a uniformly gapped S -family of paths lifts to an S -family of automorphisms in $\underline{\mathcal{G}}_{F_a}$, continuously in the probe.

The obstruction to a proof is exactly the uniformity of the quasi-adiabatic construction over a base: Theorem 6.2 produces α_1 for each path, and Theorem 5.1 makes the underlying Lieb–Robinson data uniform, but assembling the α_1 into a continuous S -family of automorphisms with a common quasi-locality modulus is not contained in the cited theorems.

Conjecture 7.3 (Descent of Lieb–Robinson estimates along a disorder hull, I-3). Let $\Omega = Q^{\mathbb{Z}^d}$ be a profinite disorder hull with its $\mathbb{Z}^d$ -action, and $\Phi_\bullet \in \underline{\mathcal{B}}_F(\Omega)$ a disordered family. The Lieb–Robinson estimates of Theorem 5.1 satisfy descent along the finite quotients $\Omega \rightarrow \Omega_i$: the quasi-local dynamics over Ω is the limit of the quasi-local dynamics over the finite quotients in a way that determines all quasi-local invariants from finite-resolution data. Equivalently, the sheaf $S \mapsto \{\text{quasi-local dynamics over } S\}$ is the right Kan extension of its restriction to finite quotients of Ω .

Theorem 5.3 is the unconditional shadow of Theorem 7.3: it gives the uniform velocity over Ω , which is the π_0 -level statement. The descent claim is the assertion that the entire quasi-local structure, not just the velocity, is finite-resolution data; this is what Part II would use to make the crossed-product observable algebra $C(\Omega) \rtimes \mathbb{Z}^d$ functorial in the finite quotients.

8 The light cone in the transverse-field Ising chain

The accompanying Haskell package `src/lieb-robinson-locality/` makes the light cone of Theorem 3.6 visible and fits a velocity. We summarize what it computes; the code is

the authoritative specification. It is available at github.com/YonedaAI/topological-phases-of-matter under `src/lieb-robinson-locality/`, so the numbers below reproduce from a single `ghc` (or `cabal`) build.

The `FFunction` module implements F -functions as first-class values with combinators for power-law decay $F_\varepsilon(r) = (1 + r)^{-(d+\varepsilon)}$, exponential decay, and the reweighting $F \mapsto F_a$ of Theorem 2.3; it exposes numerical estimates of $\|F\|$ and C_F on a truncated $\mathbb{Z}$. The `Ising` module builds the finite-volume transverse-field Ising chain, evolves a local operator $A = \sigma_0^x$ by the exact matrix exponential $\tau_t(A) = e^{itH} A e^{-itH}$, and computes the commutator norm

$$c(x, t) = \left\| [\tau_t(\sigma_0^x), \sigma_x^x] \right\|$$

on a grid of sites x and times t . Plotting the level sets of $c(x, t)$ shows the characteristic cone $|x| \approx v_{\text{fit}} t$ outside of which c is exponentially small; a least-squares fit through the origin of the front position—the largest site x whose commutator norm reaches 10% of the row maximum at that time—against t recovers a group velocity v_{fit} consistent with the analytic bound of Theorem 3.2. For the isotropic point $J = h = 1$ the fitted velocity is of order the known $v_{\text{LR}} = 2 \max(J, h)$ scale, and it stays finite and bounded as (J, h) range over a compact box, illustrating the uniform-over-the-base content of Theorem 5.1: the velocity does not blow up as the parameters vary within bounds.

The `Properties` module states QuickCheck properties that formalize the paper’s elementary claims: monotonicity of F -functions (power-law and exponential), the two reweighting inequalities $\|F_a\| \leq \|F\|$ and $C_{F_a} \leq C_F$ of Theorem 2.3 (both tested on the truncated constants), the subadditivity $\|\Phi + \Psi\|_F \leq \|\Phi\|_F + \|\Psi\|_F$ of the interaction norm, and the positivity and monotonicity in the norm bound of the analytic velocity of Theorem 3.2. `Main.hs` runs the simulation, prints the fitted velocity alongside the analytic Lieb–Robinson upper bound, checks deterministically that the former does not exceed the latter, and exits with status zero.

Remark 8.1. The simulation is finite-volume and therefore does not, and cannot, verify the infinite-volume statements of Theorems 3.3 and 4.4; it verifies the finite-volume Lieb–Robinson estimate of Theorem 3.1, which is the uniform-in- Λ input those theorems pass to the limit. That is the honest scope of a numerical check: it exhibits the mechanism (a finite velocity, stable under parameter variation) whose analytic control is the content of the theorems.

9 Discussion

What the condensed packaging buys. Nothing in Theorem 3.3 is new as analysis; it is Nachtergaele–Sims–Young in a fixed notation. The step that is genuinely a step is Theorem 4.4 together with Theorem 5.1: the Lieb–Robinson estimate is not merely a bound but a *continuity* statement in the Banach norm, and continuity is exactly what condensation requires. Once one sees this, the dynamics is a morphism $\mathbb{R} \times \mathcal{B}_F \rightarrow \text{Aut}(\mathcal{A})$ for free, and profinite families (finite-volume towers, disorder hulls, inverse limits of parameter spaces) inherit a uniform light cone with no extra work. The payoff is organizational, as the program concedes throughout: we do not compute a new invariant, we make the invariants of later parts live over a base that includes disorder and continuous families on the same footing as single Hamiltonians.

Limitations. Three boundaries are worth stating plainly. First, everything is tied to a fixed F -function; interactions with slower-than- F tails, and genuinely long-range models, are outside the theorems, and for some of them the finite-velocity picture is known to fail. Second, the condensed *set* structure is all that is proved: the higher-stack structure of $\mathfrak{Ham}_{d,G}$ (Theorem 7.1) and the condensed group of automorphisms (Theorem 7.2) are conjectural, and we have been careful never to use them as if proved. Third, the light-condensed reduction requires separability; for a non-separable interaction space one is in the full condensed formalism, with its attendant set-theoretic care, and the statements about $\text{Aut}(\mathcal{A})$ being Polish use separability of $\mathcal{A}$.

Where the thread continues. The automorphic equivalence of Theorem 6.2 is the generator of the class $\mathcal{W}$ that Part III inverts to form the phase groupoid; the uniform clustering of Theorem 5.4 is the correlation-decay input to the stability theory there; and the condensed morphism of Theorem 4.4 is what lets Part IV’s parametrized families be morphisms of condensed spectra. The disorder corollary (Theorem 5.3) and its conjectural sharpening (Theorem 7.3) feed the crossed-product picture of Part II. In each case the locality layer supplies uniformity, and uniformity is what makes the later constructions functorial in the probe rather than merely pointwise.

10 Conclusion

We have built the analytic ground floor of the condensed-mathematics program for topological phases. The F -normed G -symmetric interactions form a real Banach space $\mathcal{B}_F$ (Theorem 3.3), its condensation $\underline{\mathcal{B}}_F$ is a condensed $\mathbb{R}$ -vector space, and the Heisenberg dynamics is a genuine morphism $\mathbb{R} \times \underline{\mathcal{B}}_F \rightarrow \text{Aut}(\mathcal{A})$ of condensed sets (Theorem 4.4) because the Lieb–Robinson bound is a continuity statement in the Banach norm. Over a profinite base an S -continuous family of interactions is a compatible tower of finite-resolution families, and a uniform F -norm bound forces a single Lieb–Robinson velocity and one light cone across the whole base (Theorem 5.1), with uniform exponential clustering as a corollary under a uniform gap (Theorem 5.4) and quasi-adiabatic continuation supplying the generators of the equivalence class $\mathcal{W}$ (Theorem 6.2). The three conjectures I-1, I-2, I-3 mark exactly where today’s estimates stop: the higher-stack structure, the condensed group of automorphisms, and descent along a disorder hull are the locality-layer instances of the program’s Master Conjecture, and they are stated as conjectures precisely because the cited technology does not yet reach them. The subsequent parts build positivity, the gap, effective field theories, and realizability on top of this floor; each of them takes “ S -continuous family with a uniform light cone” as a primitive, and that primitive is what this paper makes precise.

A Constants in the Lieb–Robinson bound

For reference we collect the constants used above, all for a fixed F -function F on (L, dist) and its reweighting F_a .

- $\|F\|$ (uniform summability) and C_F (convolution) are the two constants of Theorem 2.2; reweighting gives $\|F_a\| \leq \|F\|$ and $C_{F_a} \leq C_F$ (Theorem 2.3).
- The commutator prefactor in Theorem 3.1 is $2\|A\|\|B\|/C_F$ and the time factor is $e^{2\|\Phi\|_F C_F |t|} - 1$.
- The velocity in Theorem 3.2 is $v = 2\|\Phi\|_{F_a} C_{F_a}/a$; minimizing the bound over $a > 0$ (subject to $\Phi \in \mathcal{B}_{F_a}$) gives the sharpest cone the method provides.
- The Lipschitz constant of the dynamics in the interaction (Theorem 3.4) is $K = 2\|A\| \|\text{supp } A\| \|F\| T e^{2B C_F T}$ on the ball $\|\cdot\|_F \leq B$ and time horizon T .
- The uniform velocity over a profinite base (Theorem 5.1) is $v = 2B_a C_{F_a}/a$ with $B_a = \sup_{s \in S} \|\Phi_s\|_{F_a}$.

References

- [1] Ko Aoki. (semi)topological K -theory via solidification, 2024.
- [2] Sven Bachmann, Spyridon Michalakis, Bruno Nachtergaele, and Robert Sims. Automorphic equivalence within gapped phases of quantum lattice systems. *Communications in Mathematical Physics*, 309:835–871, 2012.
- [3] Clark Barwick and Peter Haine. Pyknotic objects, I. Basic notions, 2019.
- [4] Agnès Beaudry, Michael Hermele, Juan Moreno, Markus J. Pflaum, Marvin Qi, and Daniel Spiegel. Homotopical foundations of parametrized quantum spin systems. *Reviews in Mathematical Physics*, 36(9):2460003, 2024.
- [5] Ola Bratteli and Derek W. Robinson. *Operator Algebras and Quantum Statistical Mechanics 1: C^* - and W^* -Algebras, Symmetry Groups, Decomposition of States*. Texts and Monographs in Physics. Springer, Berlin, 2 edition, 1987.
- [6] Dustin Clausen and Peter Scholze. Analytic stacks, 2023–2024. Lecture course, Max Planck Institute for Mathematics / IHES; recorded.
- [7] Toby S. Cubitt, David Pérez-García, and Michael M. Wolf. Undecidability of the spectral gap. *Nature*, 528:207–211, 2015.
- [8] Matthew B. Hastings. Locality in quantum systems. In *Quantum Theory from Small to Large Scales (Les Houches Summer School, Vol. 95)*, pages 171–212. Oxford University Press, 2010.
- [9] Matthew B. Hastings and Tohru Koma. Spectral gap and exponential decay of correlations. *Communications in Mathematical Physics*, 265:781–804, 2006.

- [10] Matthew B. Hastings and Xiao-Gang Wen. Quasi-adiabatic continuation of quantum states: The stability of topological ground state degeneracy and emergent gauge invariance. *Physical Review B*, 72:045141, 2005.
- [11] Elliott H. Lieb and Derek W. Robinson. The finite group velocity of quantum spin systems. *Communications in Mathematical Physics*, 28(3):251–257, 1972.
- [12] Matthew Long and The YonedaAI Collaboration. From lattice models to effective field theories: Stabilization and the invertible condensed phase spectrum, 2026. Part IV of this series.
- [13] Matthew Long and The YonedaAI Collaboration. Physical realizability of bordism and homotopy classes by gapped lattice systems, 2026. Part V of this series.
- [14] Matthew Long and The YonedaAI Collaboration. Positivity, C^* -norms, and condensed state spaces of quasi-local algebras, 2026. Part II of this series.
- [15] Matthew Long and The YonedaAI Collaboration. Topological phases of matter in the condensed-mathematics paradigm: A modular research program, 2026. Part VI of this series.
- [16] Matthew Long and The YonedaAI Collaboration. The uniformly gapped substack: Existence and stability of the thermodynamic spectral gap, 2026. Part III of this series.
- [17] Bruno Nachtergaele and Robert Sims. Lieb-robinson bounds and the exponential clustering theorem. *Communications in Mathematical Physics*, 265:119–130, 2006.
- [18] Bruno Nachtergaele, Robert Sims, and Amanda Young. Quasi-locality bounds for quantum lattice systems. I. Lieb-Robinson bounds, quasi-local maps, and spectral flow automorphisms. *Journal of Mathematical Physics*, 60(6):061101, 2019.
- [19] Bruno Nachtergaele, Robert Sims, and Amanda Young. Quasi-locality bounds for quantum lattice systems. Part II. Perturbations of frustration-free spin models with gapped ground states. *Annales Henri Poincaré*, 23:393–511, 2022.
- [20] Peter Scholze. Lectures on analytic geometry, 2026. Joint work with Dustin Clausen; stable citable version of the 2019/20 Bonn course.
- [21] Peter Scholze. Lectures on condensed mathematics, 2026. Joint work with Dustin Clausen; stable citable version of the 2019 Bonn course.