

The Uniformly Gapped Substack: Existence and Stability of the Thermodynamic Spectral Gap

Part III of *Topological Phases of Matter in the Condensed-Mathematics
Paradigm*

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

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

14 July 2026

Abstract

Fix a spatial dimension d , a lattice L , local Hilbert spaces, a locality class, and a symmetry group G . The condensed moduli stack of G -symmetric quasi-local Hamiltonians $\mathfrak{Ham}_{d,G}$ carries a distinguished substack $\mathfrak{Gap}_{d,G}$, the *uniformly* gapped systems: over a profinite probe S , a family lies in $\mathfrak{Gap}_{d,G}(S)$ exactly when $\inf_{s \in S} \text{gap}(H_s) \geq \Delta$ for some $\Delta > 0$. We take the word *uniformly* seriously (a family whose pointwise gaps degenerate to zero is not a member) and study $\mathfrak{Gap}_{d,G}$ as the object on which the phase functor of the program rests. The governing constraint is a hard theorem: by Cubitt, Pérez-García and Wolf the spectral gap is undecidable, so there is no algorithm and no uniform criterion that decides membership in $\mathfrak{Gap}_{d,G}$ for general translation-invariant families. We therefore do not attempt to *decide* the gap. Existence of a thermodynamic gap is treated as a hypothesis that *defines* the substack, and the content of the paper is **stability and structure**, not decidability. We prove three results from cited present-day technology. First, at fixed finite volume the n -gap is Lipschitz in the interaction norm (with constant $2C_\Lambda$), so the gapped locus is open at every finite resolution and the gap is continuous along continuous families (Theorem III-A). Second, on the frustration-free / LTQO stratum the Bravyi–Hastings–Michalakis and Michalakis–Zwolak stability theorems, with the Nachtergaele–Sims–Young infinite-volume bulk-gap result, supply a volume-independent perturbation threshold, which we recast as: $\mathfrak{Gap}_{d,G}$ contains a condensed-open neighborhood of each stratum point along quasi-local perturbation directions (Theorem III-B). Third, a uniform gap forces uniform exponential clustering, with a correlation length bounded uniformly over the profinite base (Theorem III-C). Beyond the stratum we record what we cannot prove as three numbered conjectures: openness of $\mathfrak{Gap}_{d,G}$ in the full condensed topology,

detected by finite quotients (III-1); descent of uniform gappedness along profinite disorder hulls $\Omega = Q^{\mathbb{Z}^d}$ (III-2); and the identification of quasi-adiabatic continuation with the internal path components realizing $\mathcal{W}$, so that $\pi_0 \text{Shape}(\mathfrak{Gap}_{d,G}[\mathcal{W}^{-1}])$ equals the operator-algebraic set of gapped ground-state phases (III-3). A transverse-field Ising computation renders the gapless discriminant $\Sigma = \{g = 1\}$ as a numerical picture and exhibits, away from criticality, both gap persistence and the Lipschitz bound of Theorem III-A.

Contents

1	Introduction	3
1.1	The gapped substack and why uniformity is the point	3
1.2	The undecidability wall	4
1.3	Contributions	4
1.4	Relation to companion papers	5
2	Mathematical framework	6
2.1	Condensed probes and the Hamiltonian stack	6
2.2	Finite-volume gap and thermodynamic gap	6
2.3	The uniformly gapped substack	7
2.4	What “open substack” should mean	8
3	The undecidability wall and the definitional stance	8
3.1	The Cubitt–Pérez-García–Wolf theorem	8
3.2	Consequences for the substack	9
3.3	Why the program survives the wall	9
4	Finite-volume gap continuity	10
4.1	Weyl continuity of eigenvalues	10
4.2	Theorem III-A	10
4.3	Semicontinuity of the thermodynamic gap	11
5	Stability on the frustration-free / LTQO stratum	12
5.1	Frustration-freeness and local topological quantum order	12
5.2	The stability theorems, as cited	13
5.3	Theorem III-B: stability as condensed-openness	14
5.4	Spectral flow and automorphic equivalence	15
6	Uniform clustering from a uniform gap	16
6.1	Gap implies clustering	16
6.2	Theorem III-C: uniformity over the base	16
7	The structure of the gapped substack: conjectures	17
7.1	Openness in the full condensed topology	17
7.2	Uniform-gap descent along profinite covers	18

7.3	Quasi-adiabatic continuation realizes the equivalence class	18
8	Numerical illustration: the transverse-field Ising chain	19
8.1	The model	19
8.2	Exact diagonalization and the discriminant	20
8.3	Perturbative stability away from criticality	20
8.4	What the finite computation cannot do	21
8.5	The Haskell computation	21
9	Discussion	22
9.1	What the condensed paradigm contributes here	22
9.2	The two walls	22
9.3	The profinite-disorder thread	22
9.4	Limitations	23
10	Conclusion	23
	Code availability	24

1 Introduction

1.1 The gapped substack and why uniformity is the point

The program of which this is the third part organizes topological phases of matter around a single geometric object: the condensed moduli stack $\mathfrak{Ham}_{d,G}$ of G -symmetric local and quasi-local Hamiltonians on a lattice L in spatial dimension d . Part I [19] makes the interaction space into a Banach space $\mathcal{B}_F$ through an F -function locality norm and promotes the Heisenberg dynamics to a morphism of condensed objects; Part II [22] attaches the quasi-local C^* -algebra, its compact state space, and the solid K -theory invariant of [1]. The present paper concerns the substack

$$\mathfrak{Gap}_{d,G}(S) = \bigcup_{\Delta > 0} \{ H \in \mathfrak{Ham}_{d,G}(S) : \inf_{s \in S} \text{gap}(H_s) \geq \Delta \}, \quad S \text{ profinite}, \quad (1)$$

the *uniformly* gapped systems.

Everything downstream (the stabilized phase ∞ -groupoid $\mathfrak{Phase}_{d,G} = (\mathfrak{Gap}_{d,G}[\mathcal{W}^{-1}])^{\text{st}}$, the phase set $\text{Phases}_{d,G} = \pi_0 \text{Shape}(\mathfrak{Phase}_{d,G})$, and the invertible condensed phase spectrum $\mathbf{IP}_{d,G}^{\text{cond}}$ built in Part IV [20]) is a construction on top of (1). The phase functor is only as good as the substack it is built from.

The word *uniformly* is not decorative. A profinite family $s \mapsto H_s$ in which each H_s happens to be gapped, but with $\inf_s \text{gap}(H_s) = 0$, is exactly the kind of family that fails to define a robust invariant: as one moves through the base the gap can be driven arbitrarily small, correlations grow without bound, and the locally constant phase label ν_f of the program loses meaning at the degenerating points. The union over $\Delta > 0$ in (1) records that a member of $\mathfrak{Gap}_{d,G}(S)$ is a family for which *some* strictly positive Δ bounds the gap below

simultaneously over the whole probe. Keeping the $\inf_s$ explicit is the single most important discipline of this paper.

1.2 The undecidability wall

There is a hard theorem standing between us and any naive program to “compute the substack.” Cubitt, Pérez-García and Wolf [11] proved that the spectral gap is *undecidable*: there is a family of translation-invariant, nearest-neighbour Hamiltonians on a two-dimensional lattice, depending computably on a parameter, for which no algorithm decides whether the thermodynamic-limit system is gapped or gapless. The construction embeds the halting problem into the low-energy physics, so the question “is this system gapped?” is at least as hard as the halting problem, and in fact Π_1 -complete in the arithmetic hierarchy. The undecidability survives many natural restrictions.

We state the consequence for this paper in the strongest terms, because it shapes every claim we make.

*There is no general algorithm, decision procedure, or uniform criterion that determines membership in $\mathfrak{Gap}_{d,G}$ for arbitrary members of $\mathfrak{Ham}_{d,G}$. Existence of a thermodynamic gap is not, in general, a decidable property; it is an **assumption that defines** the substack $\mathfrak{Gap}_{d,G}$.*

This is why the paper is titled around *stability* rather than existence. We do not offer a test for gappedness: no such test can exist in general. Instead we take gappedness as a standing hypothesis and ask a question undecidability does not forbid: *given* a uniformly gapped family, what can be said about its neighborhood, its correlations, and its behavior along the profinite base? The answers are stability theorems, and they are the mathematically load-bearing content of Part III. Any statement in this paper that could be read as deciding, bounding from below, or algorithmically certifying the gap of a general member of $\mathfrak{Ham}_{d,G}$ would contradict [11] and is, by construction, absent.

1.3 Contributions

Within the boundary set by Section 1.2, we prove three theorems and formulate three conjectures.

- **Theorem III-A** (Section 4). At a fixed finite volume the n -gap of a Hamiltonian is $2C_\Lambda$ -Lipschitz in the interaction norm (equivalently 2-Lipschitz in the operator norm); consequently the gapped locus is open at every finite resolution, and along any S -continuous family the map $s \mapsto \text{gap}(H_s^\Lambda)$ is continuous. This is elementary (Weyl’s perturbation inequalities plus finite-volume norm equivalence), and it is the one place where openness of $\mathfrak{Gap}_{d,G}$ is unconditional. It is also the finite-volume statement that undecidability leaves untouched: a finite matrix has a computable gap; undecidability is a statement about the thermodynamic limit.
- **Theorem III-B** (Section 5). On the frustration-free / LTQO stratum, the stability theorems of Bravyi–Hastings–Michalakis [7, 6], Michalakis–Zwolak [24], and

Nachtergaele–Sims–Young [28] provide a perturbation threshold $\varepsilon_0 > 0$ independent of system size. We recast this as a statement about $\mathfrak{Gap}_{d,G}$: each point of the stratum has a condensed-open neighborhood, of interaction-norm radius ε_0 along quasi-local perturbation directions, lying inside $\mathfrak{Gap}_{d,G}$. Openness and stability of the substack hold *on the stratum*; we are careful to claim nothing beyond it.

- **Theorem III-C** (Section 6). A uniform gap over a family forces uniform exponential clustering. Combining Hastings–Koma [12] and Nachtergaele–Sims [25] with the uniform Lieb–Robinson velocity of Part I, a family with $\inf_s \text{gap}(H_s) \geq \Delta$ has a correlation length ξ bounded above uniformly over the profinite base. This is the structural feature separating uniformly gapped families from merely pointwise-gapped ones.
- **Conjectures III-1, III-2, III-3** (Section 7). Openness of $\mathfrak{Gap}_{d,G}$ in the full condensed topology on the physically relevant stratum, with the uniform-gap sheaf condition detected by finite quotients (III-1); descent of uniform gappedness along profinite covers, i.e. disorder hulls $\Omega = Q^{\mathbb{Z}^d}$ (III-2); and the identification of quasi-adiabatic continuation with the internal path components realizing $\mathcal{W}$, so that $\pi_0 \text{Shape}(\mathfrak{Gap}_{d,G}[\mathcal{W}^{-1}])$ equals the operator-algebraic set of gapped phases of Ogata [29] (III-3).

Section 8 illustrates the whole picture on the transverse-field Ising chain: exact diagonalization renders the gapless discriminant $\Sigma = \{g = 1\}$, the gap persists away from it in the manner Theorem III-A predicts, and, crucially, no finite computation decides the thermodynamic gap, exactly as Section 1.2 requires.

1.4 Relation to companion papers

This is Part III of a six-part series developing topological phases of matter in the condensed-mathematics paradigm. The parts are modular: each takes the previous as input and produces new structure.

Part I, Condensed Locality: Lieb–Robinson Estimates and Quasi-Local Dynamics on the Moduli Stack of Hamiltonians [19], supplies the F -function Banach space $\mathcal{B}_F$ and the uniform Lieb–Robinson velocity that Theorem III-C consumes; the quasi-adiabatic and automorphic-equivalence tools it formalizes are the ones Conjecture III-3 invokes.

Part II, Positivity, C^ -Norms, and Condensed State Spaces of Quasi-Local Algebras* [22], provides the infinite-volume GNS framework in which the thermodynamic gap of Section 2 is defined, and the crossed-product observable algebras through which the disorder hull of Conjecture III-2 acquires its K -theory.

Part IV, From Lattice Models to Effective Field Theories: Stabilization and the Invertible Condensed Phase Spectrum [20], is the immediate consumer of this paper: it group-completes the invertible sector of $\mathfrak{Phase}_{d,G} = (\mathfrak{Gap}_{d,G}[\mathcal{W}^{-1}])^{\text{st}}$ into $\mathbf{IP}_{d,G}^{\text{cond}}$. Its constructions presuppose that $\mathfrak{Gap}_{d,G}$ is a well-behaved, stable substack, which is what Theorems III-A–III-C, on their stratum, secure.

Part V, Physical Realizability of Bordism and Homotopy Classes by Gapped Lattice Systems [21], asks which abstract classes are realized by uniformly gapped lattice families; the “uniformly gapped” hypothesis it uses is the one defined here.

Part VI, *Topological Phases of Matter in the Condensed-Mathematics Paradigm: A Modular Research Program* [23], assembles the five modules and states the global Master Conjecture, of which Conjectures III-1–III-3 are the gap-theoretic components.

2 Mathematical framework

We recall only what is needed and fix notation consistent with the series. Details of the condensed formalism are in [31, 3] and Part I; the operator-algebraic framework is in [29] and Part II.

2.1 Condensed probes and the Hamiltonian stack

A *condensed set* is a sheaf on the site of profinite sets with finite, jointly-surjective covers [31]. For a topological space X its condensation $\underline{X}$ is defined on a profinite probe S by $\underline{X}(S) = \text{Cont}(S, X)$; on compactly generated spaces $X \mapsto \underline{X}$ is fully faithful, so ordinary parameter spaces, tori, Banach spaces such as $\mathcal{B}_F$, and compact disorder hulls all embed into a category with good algebraic and homological behavior. We use *light* condensed sets throughout to avoid set-theoretic size issues in the profinite inverse limits [10].

For a profinite probe S , an element of $\mathfrak{Ham}_{d,G}(S)$ is an S -continuous family $s \mapsto \Phi_s$ of G -symmetric quasi-local interactions with finite F -norm: equivalently, by Part I, a continuous map $S \rightarrow \mathcal{B}_F$ landing in the G -symmetric interactions. Each interaction Φ assigns to a finite $X \subset L$ a self-adjoint $\Phi(X)$ supported on X , and generates finite-volume Hamiltonians $H_\Lambda(\Phi) = \sum_{X \subseteq \Lambda} \Phi(X)$ for finite $\Lambda \subset L$ and a thermodynamic-limit dynamics and ground-state theory in the GNS sense of Part II. We write H for the interaction, H_s for the member at $s \in S$, and H_Λ for a finite-volume restriction when the interaction is fixed.

2.2 Finite-volume gap and thermodynamic gap

Two notions of gap appear, and their relationship is where the subtlety lives.

At *finite volume* Λ , H_Λ is a self-adjoint operator on the finite-dimensional Hilbert space $\mathcal{H}_\Lambda = \bigotimes_{x \in \Lambda} \mathcal{H}_x$. Let its eigenvalues, counted with multiplicity and in non-decreasing order, be $\lambda_0(H_\Lambda) \leq \lambda_1(H_\Lambda) \leq \dots$. For an integer $n \geq 1$ we define the *n-gap*

$$\text{gap}_{\Lambda,n}(H) = \lambda_n(H_\Lambda) - \lambda_0(H_\Lambda), \quad (2)$$

the spectral gap above the lowest n levels. When the ground state is non-degenerate one takes $n = 1$; when there is a ground-state *space* of dimension m (as in topological order or symmetry breaking) the physically relevant gap is the *m-gap*, the distance from the ground-state sector to the first genuinely excited level. The device of a fixed integer n sidesteps the discontinuity that a changing ground-state multiplicity would otherwise introduce; we return to this in Section 4.

At the *thermodynamic* level, following Part II and [29, 28], one works with an infinite-volume ground state ω , its GNS representation $(\mathcal{H}_\omega, \pi_\omega, \Omega_\omega)$, and the GNS Hamiltonian $H_\omega \geq 0$ implementing the dynamics with $H_\omega \Omega_\omega = 0$. The *thermodynamic gap* is

$$\text{gap}(H) = \sup\{\gamma \geq 0 : \text{spec}(H_\omega) \cap (0, \gamma) = \emptyset\}, \quad (3)$$

the width of the spectral gap of H_ω above 0 (with $\text{gap}(H) = 0$ if no such $\gamma > 0$ exists). This is the quantity appearing in (1). When a system has a uniform finite-volume gap— $\text{gap}_{\Lambda,m}(H) \geq \gamma$ for all Λ with γ independent of Λ —the thermodynamic gap satisfies $\text{gap}(H) \geq \gamma$; but the converse can fail, and a thermodynamic gap need not be accompanied by a uniform finite-volume lower bound taken as an axiom. The Nachtergaele–Sims–Young framework [28] is designed precisely to work with the infinite-volume gap directly rather than to assume uniform finite-size bounds, and we adopt its stance.

Remark 2.1. Undecidability [11] is a statement about (3), not (2). For any fixed finite Λ the number $\text{gap}_{\Lambda,n}(H)$ is an eigenvalue difference of an explicit finite matrix and is computable to any precision. What no algorithm can do is decide, from the local interaction data, whether lim-behavior produces $\text{gap}(H) > 0$ or $\text{gap}(H) = 0$. Every finite-volume statement in this paper is therefore safe from the wall; every thermodynamic statement is made *conditional* on a standing gap hypothesis.

2.3 The uniformly gapped substack

Definition 2.2 (Uniformly gapped substack). For a profinite probe S , an interaction family $H \in \mathfrak{Ham}_{d,G}(S)$ is *uniformly gapped* if there exists $\Delta > 0$ with $\inf_{s \in S} \text{gap}(H_s) \geq \Delta$. The uniformly gapped systems form the subpresheaf $\mathfrak{Gap}_{d,G} \subseteq \mathfrak{Ham}_{d,G}$ with

$$\mathfrak{Gap}_{d,G}(S) = \bigcup_{\Delta > 0} \mathfrak{Gap}_{d,G}^{\geq \Delta}(S), \quad \mathfrak{Gap}_{d,G}^{\geq \Delta}(S) = \{ H \in \mathfrak{Ham}_{d,G}(S) : \inf_{s \in S} \text{gap}(H_s) \geq \Delta \}.$$

The filtration by Δ is genuine: the strata $\mathfrak{Gap}_{d,G}^{\geq \Delta}$ are nested, $\mathfrak{Gap}_{d,G}^{\geq \Delta'} \subseteq \mathfrak{Gap}_{d,G}^{\geq \Delta}$ for $\Delta' \geq \Delta$, and a family can belong to the union without belonging to any single stratum uniformly as the probe varies. The following elementary observation records the sheaf-theoretic shape of the object and will be sharpened, conjecturally, in Section 7.

Proposition 2.3 (Restriction stability and the union structure). *Let $S' \rightarrow S$ be a map of profinite sets and $H \in \mathfrak{Gap}_{d,G}^{\geq \Delta}(S)$. Then the restriction $H|_{S'}$ lies in $\mathfrak{Gap}_{d,G}^{\geq \Delta}(S')$. Consequently each $\mathfrak{Gap}_{d,G}^{\geq \Delta}$ is a subpresheaf of $\mathfrak{Ham}_{d,G}$, and $\mathfrak{Gap}_{d,G}$ is the filtered union $\bigcup_{\Delta > 0} \mathfrak{Gap}_{d,G}^{\geq \Delta}$ of subpresheaves.*

Proof. A map $g : S' \rightarrow S$ of profinite sets sends the family $s \mapsto H_s$ to $s' \mapsto H_{g(s')}$. Since $g(S') \subseteq S$, $\inf_{s' \in S'} \text{gap}(H_{g(s')}) \geq \inf_{s \in S} \text{gap}(H_s) \geq \Delta$, so $H|_{S'} \in \mathfrak{Gap}_{d,G}^{\geq \Delta}(S')$. Presheaf functoriality is immediate from this together with the functoriality of $\mathfrak{Ham}_{d,G}$. The union statement is Definition 2.2. $\square$

Remark 2.4. Proposition 2.3 uses only inf monotonicity under restriction; it does *not* give the reverse implication, that a family gapped on every finite quotient of S is uniformly gapped on S . That reverse implication is a descent statement, and whether it holds is the content of Conjecture III-2. The gap between “gapped pointwise / at every finite resolution” and “uniformly gapped over the whole profinite probe” is exactly the phenomenon Definition 2.2 is built to track.

2.4 What “open substack” should mean

Because $\mathcal{B}_F$ is a Banach space, its condensation $\underline{\mathcal{B}}_F$ is a condensed $\mathbb{R}$ -vector space and $\mathfrak{Ham}_{d,G}$ sits over it. Openness of $\mathfrak{Gap}_{d,G}$ inside $\mathfrak{Ham}_{d,G}$ can be phrased in the condensed setting as follows. Call a subobject $U \subseteq \underline{\mathcal{B}}_F$ *condensed-open along a direction class* $\mathcal{D}$ if for every point $\Phi \in U$ there is $\varepsilon > 0$ such that every Ψ with $\Psi - \Phi \in \mathcal{D}$ and $\|\Psi - \Phi\|_F < \varepsilon$ again lies in U , and if this holds compatibly for probes: pulling back along any $g : S \rightarrow \underline{\mathcal{B}}_F$ whose image lies in the ε -ball around Φ lands in $U(S)$. For a full topology on $\mathfrak{Ham}_{d,G}$ one wants $\mathcal{D}$ to be all quasi-local directions and ε uniform; the honest situation, developed in Section 5, is that we can prove condensed-openness of $\mathfrak{Gap}_{d,G}$ only *on a stratum* and only *along the quasi-local perturbation class* for which the stability theorems apply. Recording this scope precisely is not a weakness of the result but the exact shape undecidability forces the result to take.

3 The undecidability wall and the definitional stance

3.1 The Cubitt–Pérez-García–Wolf theorem

We restate the constraint precisely, in the form we will use.

Theorem 3.1 (Cubitt–Pérez-García–Wolf [11]). *There is a fixed local Hilbert space dimension and an explicit, computable family $n \mapsto H(n)$ of translation-invariant nearest-neighbour Hamiltonians on $\mathbb{Z}^2$, indexed by a positive integer parameter n , with the following property. The map that sends n to the truth value of “the thermodynamic-limit system $H(n)$ is gapped” is undecidable: no Turing machine, given n , halts with the correct yes/no answer for all n . The problem is Π_1 -hard, and the two cases are sharply separated (a gap bounded below by a constant in the gapped case, dense spectrum above the ground state in the gapless case), so the undecidability is not an artifact of borderline or ambiguously-gapped systems.*

Two features of Theorem 3.1 matter for us. First, the hardness is already present for translation-invariant, nearest-neighbour, two-dimensional systems—the most classical and most physical corner of $\mathfrak{Ham}_{d,G}$ —so one cannot escape it by restricting to a tame-looking subclass defined by locality or symmetry alone. Second, the separation between the gapped and gapless cases is sharp, which forecloses the hope that some “soft” or approximate criterion could decide the gap up to an ε : the difficulty is not analytic delicacy near a threshold but genuine computational undecidability.

The undecidability persists in one spatial dimension [4] and under various symmetry restrictions, so it is not an artifact of high dimension. It is, however, a *fragile* phenomenon, in a way consonant with the viewpoint of this paper: the undecidable families are exactly the ones that are not stably gapped. Castilla-Castellano and Lucia [8] show that an arbitrarily small local perturbation of the one-dimensional construction restores decidability, by breaking the fine-tuned energy cancellation on which the halting encoding depends. Undecidability lives on the boundary between gapped and gapless—precisely the locus the stability theory of Section 5 stays away from, and precisely why that theory must assume, rather than certify, membership in the gapped substack.

3.2 Consequences for the substack

Theorem 3.1 has immediate and non-negotiable consequences for how $\mathfrak{Gap}_{d,G}$ may be studied.

Corollary 3.2 (No uniform gap criterion). *There is no algorithm that, given a finite description of a member $H \in \mathfrak{Ham}_{d,G}$ (its local interaction terms and symmetry data), decides whether $H \in \mathfrak{Gap}_{d,G}$ over the one-point probe. Equivalently, there is no computable function $c : \mathfrak{Ham}_{d,G} \rightarrow \mathbb{R}_{>0}$ such that $\text{gap}(H) \geq c(H)$ whenever H is gapped and $c(H)$ certifies gaplessness otherwise.*

Proof. Such an algorithm, restricted to the CPW family $n \mapsto H(n)$ of Theorem 3.1, would decide gappedness of $H(n)$ from the finite data determining $H(n)$, contradicting undecidability. $\square$

Consequently, “ $H \in \mathfrak{Gap}_{d,G}$ ” cannot be the output of a general computation; it is an input. We adopt the following stance, which we regard as forced rather than chosen.

Definitional stance. The uniformly gapped substack $\mathfrak{Gap}_{d,G}$ is defined by the gap hypothesis, not detected by a criterion. Membership is a hypothesis one places on a family. The theorems of this paper take the form “if a family (or a stratum point) is uniformly gapped and satisfies stated structural hypotheses, then its neighborhood, correlations, and descent behavior are controlled.” No theorem asserts, for a general family, that the hypothesis holds.

Remark 3.3. It is worth marking the boundary precisely, because it is easy to overstate the wall in either direction. *Undecidable:* the thermodynamic gap of a general translation-invariant family (Theorem 3.1). *Decidable / computable:* the finite-volume n -gap of any explicit finite H_Λ (Remark 2.1); the gap of any system for which one has a proof (frustration-free systems with a Knabe-type or martingale bound, exactly solvable models, etc.); and (this is the content of Section 5) the *persistence* of a gap under small perturbations once a gap is known to exist with the right structural hypotheses. Stability is a conditional, hypothesis-carrying statement, and conditional statements are exactly what undecidability permits.

3.3 Why the program survives the wall

It might seem that undecidability is fatal to a program that puts $\mathfrak{Gap}_{d,G}$ at its center. It is not, for the same reason that undecidability of the halting problem is not fatal to the theory of computation: one develops the theory of the objects *assuming* the property, and proves structural and stability results about the class so defined. The moduli-theoretic viewpoint is well suited to this. A stack does not need a decision procedure for membership to be a useful object; it needs good functorial behavior, a sensible topology, and stability of its distinguished subobjects under the maps one cares about. Those are precisely what Theorems III-A–III-C provide, on the stratum where they can be proved. The undecidability wall bounds the *global* reach of the theory (one cannot hope for a single computable classification of all gapped systems), but it leaves the local and stratified theory intact, and it is the local and stratified theory that feeds Parts IV and V.

4 Finite-volume gap continuity

We begin with the one unconditional result: at finite volume, the gap is a Lipschitz function of the interaction, and the gapped locus is open. This is the finite-volume shadow of the openness we would like for $\mathfrak{Gap}_{d,G}$, and it is the part of the picture undecidability leaves entirely alone.

4.1 Weyl continuity of eigenvalues

Lemma 4.1 (Weyl monotonicity). *Let A, B be self-adjoint operators on a finite-dimensional Hilbert space with eigenvalues $\lambda_0(A) \leq \dots \leq \lambda_{N-1}(A)$ and $\lambda_0(B) \leq \dots \leq \lambda_{N-1}(B)$ in non-decreasing order. Then for every k ,*

$$|\lambda_k(A) - \lambda_k(B)| \leq \|A - B\|,$$

where $\|\cdot\|$ is the operator norm.

Proof. This is Weyl's inequality. By the min-max theorem,

$$\lambda_k(A) = \min_{\dim V = k+1} \max_{0 \neq v \in V} \frac{\langle v, Av \rangle}{\langle v, v \rangle}.$$

For any subspace V and unit vector $v \in V$, $\langle v, Av \rangle = \langle v, Bv \rangle + \langle v, (A - B)v \rangle \leq \langle v, Bv \rangle + \|A - B\|$, so $\lambda_k(A) \leq \lambda_k(B) + \|A - B\|$. Symmetry in A, B gives the reverse, and the two together give the claim. $\square$

To pass from operator norm to the interaction F -norm we use finite-volume norm equivalence, which is the elementary end of the Lieb–Robinson machinery of Part I.

Lemma 4.2 (Finite-volume norm domination). *Fix a finite $\Lambda \subset L$. There is a constant $C_\Lambda < \infty$, depending on Λ and the F -function, such that for all interactions Φ, Ψ ,*

$$\|H_\Lambda(\Phi) - H_\Lambda(\Psi)\| \leq C_\Lambda \|\Phi - \Psi\|_F.$$

Proof. $H_\Lambda(\Phi) - H_\Lambda(\Psi) = \sum_{X \subseteq \Lambda} (\Phi(X) - \Psi(X))$, a finite sum of at most $2^{|\Lambda|}$ terms. The F -norm dominates the per-term operator norm up to the reweighting factor built into the norm [26]; summing the finitely many reweighting factors over $X \subseteq \Lambda$ yields a finite C_Λ with the stated bound. $\square$

4.2 Theorem III-A

Theorem III-A (Finite-volume gap continuity and finite-resolution openness). *Fix a finite $\Lambda \subset L$ and an integer $n \geq 1$.*

(i) *The n -gap is $2C_\Lambda$ -Lipschitz in the interaction norm:*

$$|\text{gap}_{\Lambda,n}(\Phi) - \text{gap}_{\Lambda,n}(\Psi)| \leq 2C_\Lambda \|\Phi - \Psi\|_F \quad \text{for all } \Phi, \Psi.$$

In particular $\Phi \mapsto \text{gap}_{\Lambda,n}(\Phi)$ is continuous on $\mathcal{B}_F$.

(ii) For every $\delta > 0$ the finite-volume gapped locus $\{\Phi \in \mathcal{B}_F : \text{gap}_{\Lambda,n}(\Phi) > \delta\}$ is open in $\mathcal{B}_F$.

(iii) Along any S -continuous family $H \in \mathfrak{Ham}_{d,G}(S)$ over a profinite probe S , the map $s \mapsto \text{gap}_{\Lambda,n}(H_s)$ is continuous, and for every $\delta > 0$ the set $\{s \in S : \text{gap}_{\Lambda,n}(H_s) > \delta\}$ is open in S .

Proof. (i) By (2), $\text{gap}_{\Lambda,n}(\Phi) - \text{gap}_{\Lambda,n}(\Psi) = (\lambda_n(H_\Lambda(\Phi)) - \lambda_n(H_\Lambda(\Psi))) - (\lambda_0(H_\Lambda(\Phi)) - \lambda_0(H_\Lambda(\Psi)))$. Lemma 4.1 bounds each parenthesized difference by $\|H_\Lambda(\Phi) - H_\Lambda(\Psi)\|$, and Lemma 4.2 bounds that by $C_\Lambda \|\Phi - \Psi\|_F$. The triangle inequality gives the factor 2.

(ii) Immediate from (i): a $2C_\Lambda$ -Lipschitz function has open strict super-level sets.

(iii) An S -continuous family is a continuous map $S \rightarrow \mathcal{B}_F$; composing with the continuous $\text{gap}_{\Lambda,n}$ of (i) gives a continuous $S \rightarrow \mathbb{R}$, and continuity gives openness of super-level sets. $\square$

Remark 4.3 (Why the fixed integer n). If one used the “physical” gap $\lambda_m - \lambda_0$ with m the ground-state multiplicity (itself a function of Φ), the map $\Phi \mapsto \text{gap}$ could jump where m changes, and Lipschitz continuity would fail there. Fixing n removes the discontinuity by decoupling the level index from the multiplicity. On any region where the ground-state multiplicity is locally constant and equal to m , the physical gap agrees with the m -gap and Theorem III-A applies verbatim. The transverse-field Ising computation of Section 8 exhibits exactly this: in the ferromagnetic phase the two lowest levels form a near-degenerate ground sector ($m = 2$ in the thermodynamic limit), and the physically meaningful gap is the 2-gap, not the 1-gap. As Λ grows, the ground-state multiplicity can change at isolated interaction values (level crossings or accidental degeneracies); the fixed- n device keeps $\text{gap}_{\Lambda,n}$ a well-defined, continuous ordered-eigenvalue difference across such events, and one reads off the physical gap by taking n equal to the ground-sector dimension on the region of interest, where it is locally constant. The continuity of *each* ordered eigenvalue (Lemma 4.1) never fails; only the identification of which level index is “the gap” can shift, and that shift happens on a measure-zero set of couplings.

4.3 Semicontinuity of the thermodynamic gap

At the thermodynamic level the gap is only *lower* semicontinuous in general, and even that requires care about the mode of convergence. We record the statement we can make and mark the pitfall.

Proposition 4.4 (Conditional persistence and one-sided continuity of the gap). *Suppose $\Phi_j \rightarrow \Phi$ in $\mathcal{B}_F$ and that all Φ_j and Φ admit infinite-volume ground states obtained as weak- $*$ limits of finite-volume ground states in the sense of Part II. If, for a subsequence, the GNS gaps satisfy $\text{gap}(\Phi_j) \geq \gamma$ for a fixed $\gamma > 0$, then $\text{gap}(\Phi) \geq \gamma$ provided the ground states converge weak- $*$ and the gap is realized by a convergent family of low-lying states. This γ -persistence is the conditional upper-semicontinuity direction. Unconditionally the thermodynamic gap is only lower semicontinuous, $\text{gap}(\Phi) \leq \liminf_j \text{gap}(\Phi_j)$, which supplies no lower bound (the collapse $\text{gap}(\Phi_j) = c > 0 \rightarrow \text{gap}(\Phi) = 0$ satisfies it); without the auxiliary hypotheses no lower bound on $\text{gap}(\Phi)$ survives at all.*

Proof sketch. The argument is the standard one for persistence of a spectral gap under strong resolvent convergence, adapted to the GNS setting of [28]: a spectral gap is the statement that the ground state minimizes energy with a margin γ against all orthogonal excitations, an inequality $\langle \psi, (H_\omega - \gamma)\psi \rangle \geq 0$ for $\psi \perp \Omega_\omega$; such inequalities pass to weak-* limits of states when the energy functionals converge, which they do along $\mathcal{B}_F$ -convergent interactions by the Lieb–Robinson continuity of the dynamics (Part I). Failure of the auxiliary hypotheses—non-convergence of ground states, or a low-lying spectrum that “escapes” in the limit—breaks the inequality; we do not claim a bound in that case, and by Corollary 3.2 no unconditional bound can exist. $\square$

Remark 4.5. Proposition 4.4 is deliberately hedged. It would be an error, and a violation of the undecidability wall, to state “the thermodynamic gap is continuous in the interaction” as an unconditional theorem: continuity would let one certify gaps by approximation, which Corollary 3.2 forbids. What fails is *upper* semicontinuity: along a sequence of gapped interactions a low-lying excited level can descend to the ground-state energy in the limit, collapsing the gap discontinuously, and this is precisely the mechanism the undecidable families of [11] exploit: the gap is present at every finite approximation and disappears only in the thermodynamic limit. Lower semicontinuity under convergence of the low-lying spectral data is therefore the honest statement, and it is exactly what the stability theory of Section 5 upgrades, on its stratum, to genuine openness with a uniform radius.

5 Stability on the frustration-free / LTQO stratum

We now reach the heart of the paper: the recasting of the Bravyi–Hastings–Michalakis, Michalakis–Zwolak and Nachtergaele–Sims–Young stability theorems as a statement that $\mathfrak{Gap}_{d,G}$ contains condensed-open neighborhoods on a well-defined stratum. Throughout, we are scrupulous about hypotheses; the stratum is exactly where the hypotheses hold with constants uniform in the system size.

5.1 Frustration-freeness and local topological quantum order

Definition 5.1 (Frustration-free interaction). An interaction Φ is *frustration-free* if there is a decomposition $H_\Lambda(\Phi) = \sum_{x \in \Lambda} h_x$ with each $h_x \geq 0$ a local term, such that the finite-volume ground-state space $\mathcal{G}_\Lambda = \ker H_\Lambda(\Phi)$ is the common kernel $\bigcap_x \ker h_x$; equivalently the finite-volume ground-state energy is 0 and every ground state is annihilated by every local term.

Frustration-freeness is a genuine restriction. Generic gapped systems are not frustration-free; the class includes the paradigmatic exactly-solvable models (Kitaev’s toric code and honeycomb model [18], AKLT-type chains, group-cohomology fixed-point models [9]) and the parent Hamiltonians of matrix-product and PEPS states, but it excludes, for instance, generic Chern insulators. We never extend a frustration-free theorem to a non-frustration-free system.

Definition 5.2 (Local topological quantum order, LTQO). A frustration-free family with ground-state projections P_Λ onto $\mathcal{G}_\Lambda$ satisfies *LTQO* with decay $\Omega_{\text{LTQO}}(\cdot)$ if for every ball $B_r(x)$ of radius r and every observable O supported on $B_r(x)$,

$$\|P_\Lambda O P_\Lambda - \frac{\text{Tr}(P_{B_{r'}(x)} O)}{\text{Tr } P_{B_{r'}(x)}} P_\Lambda\| \leq \|O\| \Omega_{\text{LTQO}}(r' - r)$$

for $r < r'$, where Ω_{LTQO} is a fixed rapidly-decaying function independent of Λ . Informally: ground states are locally indistinguishable, up to a boundary correction that decays in the distance to the region's edge.

The LTQO condition is the operator-algebraic expression of “topological” ground-state degeneracy: the degenerate ground states cannot be told apart by any local measurement. Michalakis and Zwolak [24] showed that, together with a uniform local gap, LTQO is essentially equivalent to stability of the gap and to an area law for the ground-state entanglement.

Definition 5.3 (The frustration-free / LTQO stratum). Fix $\gamma > 0$, a decay function Ω_{LTQO} , an F -function, and constants controlling the local-term norms and interaction range. The *frustration-free / LTQO stratum* $\mathfrak{S} = \mathfrak{S}(\gamma, \Omega_{\text{LTQO}}, F) \subseteq \mathfrak{Ham}_{d,G}$ is the collection of interactions Φ that are frustration-free (Definition 5.1), have finite-volume gap $\text{gap}_{\Lambda, m_\Lambda}(\Phi) \geq \gamma$ uniformly in Λ (with $m_\Lambda = \dim \mathcal{G}_\Lambda$), and satisfy LTQO with decay Ω_{LTQO} , all with the fixed constants. As a subpresheaf, $\mathfrak{S}(S)$ consists of S -families landing in $\mathfrak{S}$ for every $s \in S$ with the constants uniform over S .

By construction $\mathfrak{S} \subseteq \mathfrak{Gap}_{d,G}$: the uniform finite-volume gap γ forces $\text{gap}(\Phi) \geq \gamma$ (Section 2.2). The stratum is the domain of the stability theorems.

Remark 5.4 (The stratum is a proper subclass; what lies outside it). Frustration-freeness (Definition 5.1) is essential to the arguments of Theorems 5.5 to 5.7: the relative-bound and martingale/telescoping estimates that produce the size-independent threshold rest on the common-kernel structure and the local projectors P_Λ , and on LTQO stated through those projectors. Many physically important gapped systems are *not* of this form. Free-fermion topological insulators and Chern insulators, for instance, are gapped and stable, but they are not frustration-free commuting-projector systems (indeed the chiral ones cannot be, by the Kapustin–Fidkowski obstruction [15] discussed in Section 9.2), so they lie outside $\mathfrak{S}$, and Theorem III-B says nothing about them. Their stability is real but is established by different technology (single-particle / spectral-localizer and K -theoretic methods, [30]), not by the frustration-free stability theorems recast here. We do not attempt to widen $\mathfrak{S}$ to cover them; whether a single condensed-openness statement subsumes both classes is part of the open Conjecture III-1, not something the present theorem delivers.

5.2 The stability theorems, as cited

We quote the results we use in the form relevant here; the constants are those of the original sources.

Theorem 5.5 (Bravyi–Hastings–Michalakis [7]; short proof [6]). *Let Φ be a frustration-free interaction satisfying the topological-order conditions TQO-1 and TQO-2 (LTQO) with a*

uniform local gap $\gamma > 0$. There exists $\varepsilon_0 > 0$, depending only on γ , the LTQO decay, the local Hilbert space dimension, and the lattice geometry—and not on the system size—such that for any perturbation $V = \sum_x V_x$ with $\sup_x \|V_x\| \leq \varepsilon < \varepsilon_0$ and sufficiently fast decay, the perturbed finite-volume Hamiltonians $H_\Lambda(\Phi) + V_\Lambda$ have a spectral gap above their ground-state sector bounded below by $\gamma/2$ (say), uniformly in Λ , with the ground-state splitting within the sector decaying faster than any polynomial in the system size.

Theorem 5.6 (Michalakis–Zwolak [24]). *For frustration-free Hamiltonians, LTQO together with a uniform local gap implies stability of the gap under quasi-local perturbations, with a size-independent threshold; conversely, stability plus an area law essentially forces LTQO. The perturbed system retains a uniform gap and a ground-state space continuously connected to the unperturbed one by a spectral flow.*

Theorem 5.7 (Nachtergaele–Sims–Young [28]). *For frustration-free, topologically ordered quantum lattice systems, the bulk gap in the infinite-volume GNS representation is stable under quasi-local perturbations controlled by an F -function, with a threshold independent of the finite-volume approximations. The result is proved directly for the infinite-volume gap (3), without assuming a uniform finite-size lower bound as a separate axiom.*

The three results are complementary: [7, 6] establish the finite-volume, size-independent threshold; [24] identifies LTQO as the precise condition and connects it to the area law; [28] lifts the conclusion to the thermodynamic gap directly. What they share, and what we need, is a *single* quantity: a perturbation radius $\varepsilon_0 > 0$ that does not shrink with system size.

5.3 Theorem III-B: stability as condensed-openness

Theorem III-B (Conditional condensed-openness of $\mathfrak{Gap}_{d,G}$ on the stratum). *Let $\mathfrak{S} = \mathfrak{S}(\gamma, \Omega_{\text{LTQO}}, F)$ be a frustration-free / LTQO stratum (Definition 5.3), and let $\varepsilon_0 > 0$ be the size-independent perturbation threshold of Theorems 5.5 to 5.7, depending only on γ , the LTQO decay Ω_{LTQO} , the F -function, and d . Then for every $\Phi \in \mathfrak{S}$ and every quasi-local perturbation direction $V = \sum_x V_x$ with $\sup_x \|V_x\| \leq 1$ and F -controlled decay, the segment*

$$\{\Phi + tV : 0 \leq t < \varepsilon_0\} \subseteq \mathfrak{Gap}_{d,G},$$

with thermodynamic gap bounded below by $\gamma/2$ throughout.

Consequently $\mathfrak{Gap}_{d,G}$ is condensed-open along the quasi-local direction class at each point of $\mathfrak{S}$: for every profinite probe S and every S -family H landing in the ε_0 -ball of quasi-local perturbations around a point of $\mathfrak{S}$, one has $H \in \mathfrak{Gap}_{d,G}^{\geq \gamma/2}(S)$. In words: on the frustration-free / LTQO stratum, the uniformly gapped substack contains a condensed-open neighborhood of each point along quasi-local perturbation directions.

Proof. Fix $\Phi \in \mathfrak{S}$ and a quasi-local direction V normalized as stated. For $0 \leq t < \varepsilon_0$, the perturbation tV has per-site strength $\sup_x \|tV_x\| \leq t < \varepsilon_0$ and the F -controlled decay required by Theorems 5.5 to 5.7. Since Φ lies in the stratum, it satisfies the frustration-freeness, uniform local gap γ , and LTQO hypotheses of those theorems with the fixed constants; therefore each theorem applies to $\Phi + tV$ and yields a gap above the ground-state sector bounded below by $\gamma/2$, uniformly in the finite volume, and—by Theorem 5.7—a thermodynamic gap

$\text{gap}(\Phi + tV) \geq \gamma/2$. Since $\gamma/2 > 0$ is independent of t in the range, the whole segment lies in $\mathfrak{Gap}_{d,G}^{\geq \gamma/2} \subseteq \mathfrak{Gap}_{d,G}$.

For the condensed statement, let S be a profinite probe and $H \in \mathfrak{Ham}_{d,G}(S)$ a family with $H_s = \Phi + t(s)V(s)$ where, for each s , $\Phi \in \mathfrak{S}$, $V(s)$ is a normalized quasi-local direction, and $t(s) < \varepsilon_0$, with all constants uniform over S (this is what it means for the family to land in the ε_0 -ball along quasi-local directions with stratum-uniform data). The pointwise bound just proved gives $\text{gap}(H_s) \geq \gamma/2$ for every $s \in S$, hence $\inf_{s \in S} \text{gap}(H_s) \geq \gamma/2$ and $H \in \mathfrak{Gap}_{d,G}^{\geq \gamma/2}(S)$ by Definition 2.2. This is condensed-openness along the quasi-local class at Φ in the sense of Section 2.4. $\square$

Remark 5.8 (Exactly what is and is not claimed). Theorem III-B is a stability statement and nothing more. It does *not* decide whether any given Φ belongs to $\mathfrak{S}$: that requires knowing Φ is frustration-free with a uniform gap γ , which is an input, consistent with Corollary 3.2. It does *not* assert openness of $\mathfrak{Gap}_{d,G}$ at points outside $\mathfrak{S}$, nor along non-quasi-local directions, nor with a threshold uniform over all of $\mathfrak{Gap}_{d,G}$. The three “not”s are not timidity; each is a place where a stronger claim would either be false or would contradict undecidability. The uniform-in-size threshold ε_0 is the entire mechanism: it is what turns a family of finite-volume openness statements (Theorem III-A, which has $C_\Lambda \rightarrow \infty$ as Λ grows) into a single thermodynamic openness statement with a radius that does not collapse.

Corollary 5.9 (Openness of the stratum stratum-wise). *The intersection $\mathfrak{Gap}_{d,G} \cap (\Phi + \varepsilon_0 \mathcal{D}_{\text{ql}})$ contains the full ε_0 -ball along the quasi-local direction class $\mathcal{D}_{\text{ql}}$ at each $\Phi \in \mathfrak{S}$. In particular the assignment $\Phi \mapsto \gamma/2$ is a lower bound for the gap that is locally constant along quasi-local perturbations on $\mathfrak{S}$, which is the local constancy the phase label ν_f of the program requires.*

Proof. Immediate from Theorem III-B: every point of the ball has $\text{gap} \geq \gamma/2$, so the lower bound $\gamma/2$ is constant on the ball, and local constancy of a lower bound is what a locally constant phase label needs to be well defined on the gapped locus. $\square$

5.4 Spectral flow and automorphic equivalence

The stability of Section 5.3 comes with a dynamical companion: within a uniformly gapped neighborhood the ground-state spaces are related by a quasi-local automorphism, the spectral flow. This is the rigorous meaning of “the same phase” along a gapped path, and it is the interface with Part I and with the equivalence class $\mathcal{W}$.

Proposition 5.10 (Spectral flow inside a gapped segment). *Let $t \mapsto \Phi + tV$, $0 \leq t \leq t_1 < \varepsilon_0$, be a segment as in Theorem III-B, so $\text{gap}(\Phi + tV) \geq \gamma/2$ throughout. Then there is a quasi-local, F -norm continuous cocycle of $*$ -automorphisms $(\alpha_t)_{0 \leq t \leq t_1}$ of the quasi-local algebra, generated by a time-dependent quasi-local interaction, such that α_t maps the ground-state space of Φ to that of $\Phi + tV$, and $\alpha_t \rightarrow \text{id}$ as $\gamma \rightarrow \infty$ in the sense of the Lieb–Robinson tails. The automorphisms α_t belong to the class $\mathcal{W}$ of gapped quasi-local equivalences.*

Proof sketch. This is the quasi-adiabatic continuation / spectral-flow construction of Hastings and Wen [13] in the quasi-local formulation of Bachmann, Michalakis, Nachtergaele and

Sims [2] and Nachtergaele, Sims and Young [27]. Given a uniformly gapped path $\Phi + tV$, one forms the spectral-flow generator $D_t = \int F_\gamma(s) \tau_s^{(t)} \left(\frac{d}{dt} H(t) \right) ds$ with a filter F_γ whose Fourier support avoids the gap; the gap lower bound $\gamma/2$ makes D_t quasi-local with F -tails controlled by γ , and the resulting flow α_t intertwines the ground-state spaces. Continuity in the F -norm and membership in $\mathcal{W}$ are then as in Part I. The $\gamma \rightarrow \infty$ limit sharpens the filter to a delta and the flow to the identity on the ground-state sector. $\square$

Remark 5.11. Proposition 5.10 is what makes “ $\mathfrak{Gap}_{d,G}$ modulo $\mathcal{W}$ ” a sensible quotient along stratum segments: adiabatically connected uniformly gapped systems are $\mathcal{W}$ -equivalent by an explicit quasi-local automorphism. Conjecture III-3 below proposes that this is the whole story (that the internal path components of $\mathfrak{Gap}_{d,G}$ under quasi-adiabatic continuation are exactly the operator-algebraic gapped phases), but the general statement is beyond what Proposition 5.10 proves, because a general uniformly gapped path need not lie in a single stratum.

6 Uniform clustering from a uniform gap

The third theorem is structural: it isolates a property that *uniformly* gapped families have and pointwise-gapped families need not. A gap forces exponential decay of correlations; a *uniform* gap forces the decay rate to be uniform over the base. This is the sense in which $\mathfrak{Gap}_{d,G}$, and not the naive pointwise-gapped locus, is the correct object.

6.1 Gap implies clustering

Theorem 6.1 (Hastings–Koma [12]; Nachtergaele–Sims [25]). *Let H have a thermodynamic gap $\text{gap}(H) \geq \gamma > 0$ above a ground state ω , and let the dynamics obey a Lieb–Robinson bound with velocity v and F -tails. Then there are constants $c, C > 0$, depending only on γ , v and the F -function, such that for local observables A, B with disjoint supports,*

$$|\omega(AB) - \omega(A)\omega(B)| \leq C \|A\| \|B\| e^{-c \text{dist}(\text{supp } A, \text{supp } B)},$$

with $c \sim \gamma/v$ up to the geometric factors. The correlation length is $\xi = c^{-1}$.

The precise exponent and constants are as in [12]; the Lieb–Robinson input is the exponential-clustering theorem of [25]. We use them as cited.

6.2 Theorem III-C: uniformity over the base

Theorem III-C (Uniform clustering on a uniformly gapped family). *Let $H \in \mathfrak{Gap}_{d,G}^{\geq \Delta}(S)$ be a uniformly gapped family over a profinite probe S , so $\inf_{s \in S} \text{gap}(H_s) \geq \Delta > 0$, and suppose the family has a uniform Lieb–Robinson velocity v and common F -function (as provided by Part I for an S -continuous family in $\mathcal{B}_F$). Then there are constants $c, C > 0$, depending only on Δ , v and F —and not on s —such that for every $s \in S$ and all local A, B with disjoint supports,*

$$|\omega_s(AB) - \omega_s(A)\omega_s(B)| \leq C \|A\| \|B\| e^{-c \text{dist}(\text{supp } A, \text{supp } B)}.$$

In particular the correlation length function $s \mapsto \xi(s)$ is bounded above by c^{-1} uniformly over the profinite base.

Proof. Apply Theorem 6.1 at each $s \in S$. Its constants c, C depend on the input data (γ, v, F) only. By hypothesis the gap input can be taken to be the uniform value $\gamma = \Delta$ for every s (since $\text{gap}(H_s) \geq \Delta$), and the velocity and F -function are the common ones supplied by Part I for the S -continuous family. Thus the constants produced by Theorem 6.1 are the same for every s , and the clustering bound holds with s -independent c, C . The bound on $\xi(s) = c^{-1}$ is the final clause. $\square$

Corollary 6.2 (Failure for pointwise-gapped families). *If instead H_s is gapped for every $s \in S$ but $\inf_s \text{gap}(H_s) = 0$, then the per-point clustering rate $c(s) \sim \text{gap}(H_s)/v$ degenerates to 0 along a sequence s_j with $\text{gap}(H_{s_j}) \rightarrow 0$, and no uniform correlation-length bound exists. The uniform-gap hypothesis of Theorem III-C is therefore necessary for the conclusion, not merely sufficient.*

Proof. The Hastings–Koma rate is proportional to the gap; along s_j it tends to 0, so $\xi(s_j) \rightarrow \infty$ and $\sup_s \xi(s) = \infty$. $\square$

Remark 6.3. Corollary 6.2 is the technical vindication of Definition 2.2. It shows that the difference between “gapped at every probe point” and “uniformly gapped” is not a formality: only the latter guarantees a uniform correlation length, and a uniform correlation length is what one needs for the area laws, split property, and locally computable invariants that Parts II and IV build on. The $\inf_s$ in the definition of $\mathfrak{Gap}_{d,G}$ is precisely the hypothesis that rules out Corollary 6.2.

7 The structure of the gapped substack: conjectures

Beyond the stratum, we do not have theorems; we have conjectures, and we state them as such, with the reasons they are plausible and the obstacles to proving them. Each is scoped to respect the undecidability wall.

7.1 Openness in the full condensed topology

Conjecture III-1 (Openness on the physical stratum; finite-quotient detection). *On the physically relevant stratum—the closure, in a suitable sense, of the frustration-free / LTQO systems under quasi-local gapped equivalence, together with the systems for which a stability theory holds with uniform constants— $\mathfrak{Gap}_{d,G}$ is an open condensed substack of $\mathfrak{Ham}_{d,G}$. Moreover the uniform-gap sheaf condition of Definition 2.2 is detected by finite quotients of profinite probes: a family over $S = \varprojlim S_i$ lies in $\mathfrak{Gap}_{d,G}(S)$ if and only if its pushforwards to the finite quotients S_i lie in the corresponding finite-resolution gapped loci with a common bound Δ .*

Evidence and obstruction. Theorem III-B is exactly this statement restricted to $\mathfrak{S}$ and to quasi-local directions; the conjecture asserts that the phenomenon persists on the larger physical stratum and in all directions. The obstruction to a proof is precisely undecidability:

openness in the *full* $\mathfrak{Ham}_{d,G}$, with a uniform radius, would let one certify gaps of nearby systems, and near a CPW-type embedding [11] the gap can turn on and off in a way no uniform radius can survive. Restricting to the physical stratum is what excludes those pathological directions; making “physical stratum” precise, so that the restriction is both honest and non-vacuous, is the mathematical work the conjecture packages. The finite-quotient detection clause is a descent statement and is closely tied to Conjecture III-2.

7.2 Uniform-gap descent along profinite covers

The disorder-hull thread of the program ([5, 17] and Part II) makes the profinite structure of S physical: for a finite local-configuration set Q (the letter Q denotes the disorder alphabet, kept deliberately distinct from the locality F -function of Part I) the hull $\Omega = Q^{\mathbb{Z}^d}$ is compact and totally disconnected, hence profinite, and a disordered family $\omega \mapsto H_\omega$ is literally a Ω -point of $\mathfrak{Ham}_{d,G}$. Finite quotients of Ω are finite-resolution disorder data.

Conjecture III-2 (Uniform-gap descent). *Let $\Omega = Q^{\mathbb{Z}^d} = \varprojlim_i \Omega_i$ be a disorder hull with its profinite presentation, and let $H \in \mathfrak{Ham}_{d,G}(\Omega)$ be a disordered family. Then H is uniformly gapped, $H \in \mathfrak{Gap}_{d,G}(\Omega)$, if and only if each finite-resolution approximation $H^{(i)} \in \mathfrak{Ham}_{d,G}(\Omega_i)$ is gapped with a common bound Δ independent of i . Equivalently, uniform gappedness is a closed condition on the inverse system $(\Omega_i)_i$: the uniformly gapped families are the inverse limit of the finite-resolution uniformly-gapped ones at a fixed Δ .*

Evidence and obstruction. The “only if” direction is Proposition 2.3: a uniform gap over Ω restricts to a uniform gap over each Ω_i . The content is the “if” direction—that compatible finite-resolution gaps at a common Δ assemble to a uniform gap over the limit. This is a genuine descent question. The optimistic reason to expect it: continuity of the gap along $\mathcal{B}_F$ -convergent families (Theorem III-A at each finite volume) plus compactness of Ω suggests that a common finite-resolution bound should survive the limit. The obstruction: the gap is a *thermodynamic*, not finite-volume, quantity, and Proposition 4.4 shows only lower semicontinuity in general; the descent could fail if the low-lying spectrum reorganizes in the inverse limit. The conjecture asserts it does not, on the disorder hulls, at fixed Δ . A proof would presumably route through the crossed-product algebra $C(\Omega) \rtimes \mathbb{Z}^d$ of Part II and a gap-labelling argument in the manner of Bellissard [5].

7.3 Quasi-adiabatic continuation realizes the equivalence class

Conjecture III-3 (QAC path components are the operator-algebraic phases).

Quasi-adiabatic continuation defines, on $\mathfrak{Gap}_{d,G}$, the internal path components realizing the equivalence class $\mathcal{W}$: two uniformly gapped families are $\mathcal{W}$ -equivalent if and only if they are connected by a uniformly gapped path along which the spectral flow of Proposition 5.10 is quasi-local. Consequently

$$\pi_0 \text{Shape}(\mathfrak{Gap}_{d,G}[\mathcal{W}^{-1}]) \cong \{ \text{operator-algebraic gapped ground-state phases} \}$$

of Ogata [29], and the boxed slogan of the program—a topological phase is a component of the stabilized condensed stack of gapped systems—is, on this identification, a theorem rather than a definition.

Evidence and obstruction. Proposition 5.10 proves one direction locally: a uniformly gapped segment in a stratum gives a quasi-local $\mathcal{W}$ -equivalence. The theory of automorphic equivalence in [2] and the classification program of [29] make the target set precise in $d = 1$ and, for on-site finite symmetry, in $d = 2$. The obstruction to the full statement is twofold: a general uniformly gapped path need not remain in a single stratum, so Proposition 5.10 does not immediately apply along it; and the identification of π_0 Shape of the localized stack with the operator-algebraic set requires knowing that the condensed shape does not see more (or less) than the operator-algebraic equivalence, which is a comparison between two homotopy theories that has not been carried out. In low dimension, where [29] gives complete invariants, the conjecture is closest to reach; in higher dimension it is genuinely open, in step with the general state of the classification problem.

Remark 7.1 (The three conjectures as the gap-theoretic Master Conjecture). Conjectures III-1, III-2, III-3 are the gap-layer components of the program’s Master Conjecture (Part VI [23]). III-1 is the topology of the substack, III-2 its behavior under profinite disorder, III-3 the comparison of its localized homotopy type with the established operator-algebraic classification. None can be proved outright today; each is anchored to a theorem of this paper (III-B, III-A + Proposition 2.3, and Proposition 5.10 respectively) that establishes its local or one-directional content.

8 Numerical illustration: the transverse-field Ising chain

We close the mathematical development with a concrete computation that makes the abstract picture visible: the spectral gap of the transverse-field Ising chain, computed by exact diagonalization. It renders the gapless discriminant Σ as a one-dimensional picture and exhibits the stability that Theorem III-A quantifies. The accompanying Haskell program (Section 8.5) performs the computation and checks the properties as executable QuickCheck tests.

8.1 The model

On a chain of N spins with open boundary conditions, consider

$$H(g) = - \sum_{i=1}^{N-1} Z_i Z_{i+1} - g \sum_{i=1}^N X_i, \quad g \in [0, 2], \quad (4)$$

where X_i, Z_i are Pauli operators at site i . The model has a global $\mathbb{Z}_2$ symmetry $P = \prod_i X_i$, $[H(g), P] = 0$, and a quantum phase transition at $g = 1$ separating a ferromagnetic phase ($g < 1$, two-fold degenerate ground sector in the thermodynamic limit, the symmetry-broken states) from a paramagnetic phase ($g > 1$, unique ground state). By the Jordan–Wigner transformation the model is a free-fermion system with single-particle dispersion

$$\varepsilon_k(g) = 2\sqrt{1 + g^2 - 2g \cos k}, \quad (5)$$

whose minimum over k is $\min_k \varepsilon_k(g) = 2|1 - g|$, the thermodynamic single-particle gap. We use open boundary conditions: they keep the sparse structure and the $\mathbb{Z}_2$ symmetry of the matrix simple and avoid the fermion-parity (Neveu–Schwarz / Ramond) sector bookkeeping that periodic boundaries introduce through the Jordan–Wigner transformation. Periodic boundaries would reduce finite-size edge effects, but the qualitative gap picture—the discriminant at $g = 1$ and stability away from it—is independent of the boundary condition in the thermodynamic limit. Thus

$$\Sigma = \{ g : \text{gap}(H(g)) = 0 \} = \{ g = 1 \}, \quad (6)$$

a single point of the parameter interval: the gapless discriminant of this one-parameter family. Crossing $g = 1$ is the topological/quantum phase transition; away from it the system is gapped, with gap growing linearly, $\approx 2|1 - g|$. The exact gap also exhibits the Kramers–Wannier self-duality $\min_k \varepsilon_k(g) = g \min_k \varepsilon_k(1/g)$, immediate from (5) since $2|1 - g| = g \cdot 2|1 - 1/g|$; the Haskell suite of Section 8.5 checks this identity.

8.2 Exact diagonalization and the discriminant

In the computational (Z -) basis, (4) is a real symmetric $2^N \times 2^N$ matrix: the ZZ term is diagonal, the transverse field X_i is off-diagonal (a single spin flip). We diagonalize it exactly for small N by a cyclic Jacobi eigenvalue routine, sweeping g across $[0, 2]$ and recording the low-lying spectrum. The correctness of the eigensolver is not taken on faith: the QuickCheck properties of Section 8.5 check the exact trace identities

$$\sum_j \lambda_j(g) = \text{Tr } H(g) = 0, \quad \sum_j \lambda_j(g)^2 = \text{Tr } H(g)^2 = ((N - 1) + Ng^2) 2^N, \quad (7)$$

which follow from $\text{Tr}(Z_i Z_{i+1}) = \text{Tr}(X_i) = 0$ and the orthogonality of the Pauli strings, and which a wrong spectrum would violate. The cyclic Jacobi routine drives the off-diagonal Frobenius norm below 10^{-11} , and for the dimensions used ($N \leq 8$, matrices up to 256×256) the eigenvalues are accurate to near machine precision—the residuals in the identities (7) stay below 10^{-9} across this range, reaching $\sim 10^{-13}$ at $N = 4$. At substantially larger N both the accumulated floating-point error and the $O(2^{3N})$ cost of dense diagonalization grow, which is why we cap the dense computation at $N = 8$ and read the large- N trend from the analytic reference (5) instead.

The finite- N gap $\text{gap}_N(g) = \lambda_1(g) - \lambda_0(g)$ and the 2-gap $\lambda_2(g) - \lambda_0(g)$ together render (6). In the paramagnetic phase $g > 1$ the ground state is unique and $\text{gap}_N(g)$ is the physical gap, converging to $2(g - 1)$ as N grows. In the ferromagnetic phase $g < 1$ the two lowest levels form the near-degenerate ground sector (their splitting decays exponentially in N), so the physical gap is the 2-gap—precisely the m -gap of Remark 4.3 with $m = 2$. In both phases the relevant gap develops a sharp minimum near $g = 1$ that deepens with N : the finite-size fingerprint of the discriminant point (6). Figure 1 sketches the shape.

8.3 Perturbative stability away from criticality

To illustrate Theorem III-A and the stability picture, the program performs a perturbation experiment. Fix $g = 1.5$, deep in the gapped paramagnetic phase, and add a random local




Figure 1: Schematic of the transverse-field Ising gap across $g \in [0, 2]$. The exact thermodynamic gap (solid) is the “V” $2|1 - g|$ of (5), vanishing only at the discriminant $\Sigma = \{g = 1\}$ of (6). Finite- N gaps (dashed) round the minimum and deepen toward the “V” as N grows. Away from Σ the gap is bounded below and the family is uniformly gapped on any closed sub-interval avoiding $g = 1$.

field $V = \lambda \sum_i s_i Z_i$ with $s_i \in \{\pm 1\}$ and small λ . Exact diagonalization confirms

$$|\text{gap}_N(H(1.5) + V) - \text{gap}_N(H(1.5))| \leq 2 \|V\| \leq 2\lambda N, \quad (8)$$

the finite-volume Lipschitz bound of Theorem III-A (here C_Λ absorbed into the per-site normalization). The gap stays bounded away from 0 for all sampled perturbations: away from the discriminant the gapped locus is open and stability holds, exactly as Theorem III-A guarantees at finite volume. Near $g = 1$, by contrast, the same size of perturbation moves the small gap by a comparable amount and can push the finite system across the rounded minimum—the numerical trace of why openness of the substack is a statement scoped away from the discriminant.

8.4 What the finite computation cannot do

It is essential to state what Figure 1 does *not* establish, on pain of contradicting Section 3. No finite- N computation decides the thermodynamic gap. For the transverse-field Ising chain we happen to know the answer analytically—the model is exactly solvable, (5)—so the discriminant (6) is certain; but that certainty comes from the exact Jordan–Wigner solution, not from the diagonalization. For a general translation-invariant family the analogous finite- N data would be genuinely uninformative about the thermodynamic limit, because by Theorem 3.1 no finite computation, however large N , can decide whether the gap survives. The Ising chain is an illustration of the *structure* of the gapped substack—its discriminant, its stability away from criticality, the Lipschitz bound—not a method for detecting gaps. The distinction is the entire moral of Part III.

8.5 The Haskell computation

The computation is implemented in Haskell in `src/spectral-gap-stability/`. The module `TFIM` builds the Hamiltonian (4) as a real symmetric matrix and diagonalizes it by cyclic

Jacobi rotations; **Main** sweeps g over $[0, 2]$ for $N \in \{4, 6, 8\}$, prints the gap table, runs the perturbation experiment of Section 8.3, and exits with status 0 on success. The module **Properties** encodes as QuickCheck properties: the exact trace identities (7) (which double as a validation of the eigensolver); symmetry of the Hamiltonian matrix; positivity of the gap in the paramagnetic phase at fixed N (Section 8.2); the finite-volume Lipschitz bound (8) of Theorem III-A; and the monotone growth of the gap with g in the paramagnetic phase. The properties are stated one per claim and cite the corresponding equation or theorem, so that the executable tests track the paper’s assertions.

9 Discussion

9.1 What the condensed paradigm contributes here

Part III does not prove a new stability theorem; the analytic content is [7, 6, 24, 28, 12, 25]. Nor is the idea of studying the moduli space of gapped Hamiltonians and its topology new: it is developed, in ordinary topology and with effective-field-theory methods, by Hsin, Wang and collaborators [14]. What the condensed framing contributes is organizational, and it is not empty. First, it makes “uniformly gapped” a property of a *family* in a category where families, disorder hulls, and inverse limits are first-class objects: the inf_s of Definition 2.2 is the condition that a subobject of a condensed stack is cut out correctly, and Theorem III-C shows this is the condition with structural teeth. Second, it identifies exactly which classical theorems give *openness* of the distinguished subobject (Theorem III-B) and on which stratum, so that the downstream constructions of Parts IV and V know precisely what they may assume. Third, it places the undecidability wall where it belongs (as a statement about global decidability that leaves the local, stratified, functorial theory intact) rather than letting it either be ignored or overstated into paralysis.

9.2 The two walls

The program has two hard external constraints. The first is the undecidability of the gap [11], which is the governing constraint of this paper: it is why existence of a thermodynamic gap is a hypothesis and not a theorem, why openness is a stratum statement, and why Section 8.4 insists that no finite computation detects the gap. The second, the Kapustin–Fidkowski obstruction [15] to commuting-projector realizations of chiral phases, governs Part V [21] and does not bear directly here, except to note that the frustration-free / LTQO stratum of Definition 5.3 (on which we can prove stability) is close to the commuting-projector world where the second wall bites. Chiral phases (nonzero Hall conductance, [16]) are gapped but are not frustration-free commuting-projector systems, so they lie outside $\mathfrak{S}$; our stability theorem says nothing about them, and it should not, because their stability is a different and harder question.

9.3 The profinite-disorder thread

The disorder hull $\Omega = Q^{\mathbb{Z}^d}$ is where the condensed and profinite structure is not a repackaging of a smooth manifold but a match to the physical configuration space [5, 17, 30].

Conjecture III-2 is the gap-layer instance of the program’s recurring descent question: does a finite-resolution property, holding compatibly at every finite quotient of Ω with uniform constants, descend to the profinite limit? For gaps this is delicate precisely because the gap is thermodynamic. That the question is even well posed (that a disordered family is literally a Ω -point of $\mathfrak{Ham}_{d,G}$, and uniform gappedness a condition on that point) is a contribution of the condensed viewpoint, and settling it would connect the gap-labelling of [5] to the substack $\mathfrak{Gap}_{d,G}$ directly.

9.4 Limitations

We list the limitations plainly. Theorem III-B is conditional on membership in a frustration-free / LTQO stratum, which is undecidable to verify in general and is a proper subclass of gapped systems. Theorem III-C assumes the uniform Lieb–Robinson data of Part I, which restricts to the F -function locality class and excludes genuinely long-range interactions. Proposition 4.4 gives only lower semicontinuity, and deliberately so. The three conjectures are unproven, and Conjecture III-1 in particular cannot be strengthened to openness in the full $\mathfrak{Ham}_{d,G}$ without contradicting [11]. None of these limitations is incidental; each marks a boundary that the undecidability of the gap, or the restriction to a provable stability class, places on what can be claimed.

10 Conclusion

The uniformly gapped substack $\mathfrak{Gap}_{d,G}$ is the object on which the condensed theory of topological phases is built, and it is governed by a hard theorem: the spectral gap is undecidable, so membership cannot be the output of a computation and must be taken as a defining hypothesis. Within that constraint we proved what can be proved. At finite volume the gap is Lipschitz and the gapped locus open (Theorem III-A). On the frustration-free / LTQO stratum the classical stability theorems give a size-independent perturbation threshold, which we recast as condensed-openness of $\mathfrak{Gap}_{d,G}$ along quasi-local directions (Theorem III-B), with an explicit spectral flow realizing the equivalence class $\mathcal{W}$ (Proposition 5.10). A uniform gap forces a uniform correlation length (Theorem III-C), the structural property that distinguishes uniformly gapped families from pointwise-gapped ones and that makes $\inf_s$ the correct condition in the definition of the substack. Beyond the stratum we recorded three conjectures: openness on the physical stratum with finite-quotient detection (III-1), uniform-gap descent along profinite disorder hulls (III-2), and the identification of quasi-adiabatic path components with the operator-algebraic gapped phases (III-3), each anchored to a theorem of this paper and each scoped to respect the undecidability wall. The transverse-field Ising computation made the discriminant $\Sigma = \{g = 1\}$ and the stability away from it visible, while insisting that no finite computation detects the thermodynamic gap.

The substack is now a well-understood object on its provable stratum and a precisely delineated open problem beyond it. Part IV takes $\mathfrak{Gap}_{d,G}$ as given and group-completes its invertible sector into the condensed phase spectrum $\mathbf{IP}_{d,G}^{\text{cond}}$; Part V asks which classes that spectrum records are realized by uniformly gapped lattice families. Both rest on the stability established here, and both inherit its scope.

Code availability

The Haskell source for the transverse-field Ising computation of Section 8 (exact diagonalization, the base-only cyclic Jacobi eigensolver, and the QuickCheck property suite) is available at github.com/YonedaAI/topological-phases-of-matter, in `src/spectral-gap-stability/`.

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] Johannes Bausch, Toby S. Cubitt, Angelo Lucia, and David Pérez-García. Undecidability of the spectral gap in one dimension. *Physical Review X*, 10(3):031038, 2020.
- [5] Jean Bellissard, Andreas van Elst, and Hermann Schulz-Baldes. The noncommutative geometry of the quantum hall effect. *Journal of Mathematical Physics*, 35(10):5373–5451, 1994.
- [6] Sergey Bravyi and Matthew B. Hastings. A short proof of stability of topological order under local perturbations. *Communications in Mathematical Physics*, 307:609–627, 2011.
- [7] Sergey Bravyi, Matthew B. Hastings, and Spyridon Michalakis. Topological quantum order: stability under local perturbations. *Journal of Mathematical Physics*, 51(9):093512, 2010.
- [8] Laura Castilla-Castellano and Angelo Lucia. Instability of the undecidable behavior of the spectral gap in 1D, 2026. arXiv:2607.08686 [quant-ph].
- [9] Xie Chen, Zheng-Cheng Gu, Zheng-Xin Liu, and Xiao-Gang Wen. Symmetry protected topological orders and the group cohomology of their symmetry group. *Physical Review B*, 87:155114, 2013.
- [10] Dustin Clausen and Peter Scholze. Analytic stacks, 2023–2024. Lecture course, Max Planck Institute for Mathematics / IHES; recorded.
- [11] Toby S. Cubitt, David Pérez-García, and Michael M. Wolf. Undecidability of the spectral gap. *Nature*, 528:207–211, 2015.
- [12] Matthew B. Hastings and Tohru Koma. Spectral gap and exponential decay of correlations. *Communications in Mathematical Physics*, 265:781–804, 2006.

- [13] 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.
- [14] Po-Shen Hsin and Zhenghan Wang. On topology of the moduli space of gapped hamiltonians for topological phases. *Journal of Mathematical Physics*, 64(4):041901, 2023.
- [15] Anton Kapustin and Lukasz Fidkowski. Local commuting projector hamiltonians and the quantum hall effect. *Communications in Mathematical Physics*, 373(2):763–769, 2019.
- [16] Anton Kapustin and Nikita Sopenko. Local noether theorem for quantum lattice systems and topological invariants of gapped states. *Journal of Mathematical Physics*, 63(9):091903, 2022.
- [17] Johannes Kellendonk. Non commutative geometry of tilings and gap labelling. *Reviews in Mathematical Physics*, 7(7):1133–1180, 1995.
- [18] Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2–111, 2006.
- [19] Matthew Long and The YonedaAI Collaboration. Condensed locality: Lieb–robinson estimates and quasi-local dynamics on the moduli stack of hamiltonians, 2026. Part I of this series.
- [20] 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.
- [21] Matthew Long and The YonedaAI Collaboration. Physical realizability of bordism and homotopy classes by gapped lattice systems, 2026. Part V of this series.
- [22] Matthew Long and The YonedaAI Collaboration. Positivity, C^* -norms, and condensed state spaces of quasi-local algebras, 2026. Part II of this series.
- [23] 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.
- [24] Spyridon Michalakis and Justyna P. Zwolak. Stability of frustration-free hamiltonians. *Communications in Mathematical Physics*, 322:277–302, 2013.
- [25] Bruno Nachtergaele and Robert Sims. Lieb-robinson bounds and the exponential clustering theorem. *Communications in Mathematical Physics*, 265:119–130, 2006.
- [26] 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.

- [27] 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.
- [28] Bruno Nachtergaele, Robert Sims, and Amanda Young. Stability of the bulk gap for frustration-free topologically ordered quantum lattice systems. *Letters in Mathematical Physics*, 114:24, 2024.
- [29] Yoshiko Ogata. Classification of gapped ground state phases in quantum spin systems, 2021. Proceedings of the International Congress of Mathematicians (ICM 2022).
- [30] Emil Prodan and Hermann Schulz-Baldes. *Bulk and Boundary Invariants for Complex Topological Insulators: From K-Theory to Physics*. Mathematical Physics Studies. Springer, 2016.
- [31] Peter Scholze. Lectures on condensed mathematics, 2026. Joint work with Dustin Clausen; stable citable version of the 2019 Bonn course.