

Glassy Word Problems: Ultraslow Relaxation, Hilbert Space Jamming, and Computational Complexity

Shankar Balasubramanian,^{1,2,3} Sarang Gopalakrishnan,⁴ Alexey Khudorozhkov⁵,⁶ and Ethan Lake^{1,6}

¹*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

³*Kavli Institute for Theoretical Physics, University of California, Santa Barbara, California 93106, USA*

⁴*Department of Electrical and Computer Engineering, Princeton University, Princeton, New Jersey 08544, USA*

⁵*Department of Physics, Boston University, Boston, Massachusetts 02215, USA*

⁶*Department of Physics, University of California Berkeley, Berkeley, California 94720, USA*



(Received 2 January 2024; revised 28 February 2024; accepted 1 April 2024; published 29 May 2024)

We introduce a family of local models of dynamics based on “word problems” from computer science and group theory, for which we can place rigorous lower bounds on relaxation timescales. These models can be regarded either as random circuit or local Hamiltonian dynamics and include many familiar examples of constrained dynamics as special cases. The configuration space of these models splits into dynamically disconnected sectors, and for initial states to relax, they must “work out” the other states in the sector to which they belong. When this problem has a high time complexity, relaxation is slow. In some of the cases we study, this problem also has high space complexity. When the space complexity is larger than the system size, an unconventional type of jamming transition can occur, whereby a system of a fixed size is not ergodic but can be made ergodic by appending a large reservoir of sites in a trivial product state. This finding manifests itself in a new type of Hilbert space fragmentation that we call fragile fragmentation. We present explicit examples where slow relaxation and jamming strongly modify the hydrodynamics of conserved densities. In one example, density modulations of wave vector q exhibit almost no relaxation until times $O(\exp(1/q))$, at which point they abruptly collapse. We also comment on extensions of our results to higher dimensions.

DOI: [10.1103/PhysRevX.14.021034](https://doi.org/10.1103/PhysRevX.14.021034)

Subject Areas: Quantum Physics, Statistical Physics

I. INTRODUCTION

A common paradigm in quantum dynamics is that isolated quantum systems usually thermalize if one waits long enough [1–5]. Indeed, assuming that interactions are spatially local, quantum systems tend to approach a form of local equilibrium on a timescale that is independent of system size, with the late-time dynamics governed by the hydrodynamic relaxation of a small number of conserved densities. The main possible exception to this rule is the many-body localized phase in strongly disordered systems [6,7] or in systems that effectively self-generate strong disorder [8–10]. The structures that can be rigorously shown to arrest thermalization in translation-invariant quantum systems—such as integrability [11], quantum scars [12–20], dynamical constraints [21–30], etc.—are

usually fine-tuned in some sense. However, they are still of experimental relevance since many realistic systems are near the fine-tuned points where these phenomena occur [31–36]. Systems near these points have long relaxation timescales and approximate conservation laws that are essentially exact on the timescale of realistic experiments on noisy quantum hardware and cold atom systems. Although the algebraic structure of integrable systems and systems with many-body scars has been well studied, a general understanding of the extent to which local Hilbert space constraints can arrest thermalization is still under development.

In this work, we introduce an alternative viewpoint for understanding thermalization in a large class of one-dimensional models with local Hilbert space constraints. We begin by developing a general framework for characterizing models with constrained dynamics in terms of semi-groups, algebraic structures that resemble groups but need not have inverses or an identity. This approach reproduces examples of constrained models known in the literature but also provides us with new examples with unusual properties. In particular, it enables us to leverage ideas from the

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

field of geometric group theory to construct (a) models with both exponentially long relaxation times and sharp thermalization transitions and (b) models where thermalization never occurs due to an unusual type of ergodicity breaking we call “Hilbert-space jamming” or “fragile fragmentation.”

The relation between constrained one-dimensional spin chains and semigroups can be summarized as follows (see Sec. II for a more formal discussion). Most examples of constrained systems in the literature—and all of the examples considered in this work—have constraints that can be formulated in a local product state basis. In these systems, the constraints place certain rules on the processes that the dynamics is allowed to implement, and it is these rules that endow the dynamics with the structure of a semigroup. This identification works by viewing each basis state of the on-site Hilbert space, in the computational basis, as a generator of a semigroup. Since any element of the semigroup can be written as a product of generators, a many-body computational-basis product state is naturally associated with an element of the semigroup, obtained by taking the product of generators from left to right along the chain. We call each computational-basis state a word, with each word being a presentation of a certain element in the semigroup. In this picture, the constraints are encoded by requiring that the dynamics preserve the semigroup element associated with each product state. In Sec. II, we show that all local dynamical constraints can be formulated in this way.

Of course, not all words represent distinct semigroup elements. For example, in the case where the semigroup is a group G with identity element e and elements g_1, g_2 , three distinct words of length 4 representing the same group element are $eeee$, $g_1 g_1^{-1} g_2^{-1} g_2$, and $g_1 g_2 g_2^{-1} g_1^{-1}$. Each distinct semigroup element is thus an equivalence class of words under the application of equivalence relations like $g_i g_i^{-1} = ee$. The most general local dynamics that preserves semigroup elements is precisely one that locally implements these equivalence relations.

The equivalence classes so defined produce multiple sectors of the Hilbert space (“Krylov sectors” or “fragments”) that the dynamical rules are unable to connect, breaking ergodicity and leading to Hilbert space fragmentation (HSF) [21–23]. Within a given fragment, thermalization of an initial basis state is a process by which the dynamics of the system “works out” which words represent the same semigroup element as the initial state. Crucially, when the problem of determining which words represent a given element is computationally hard, thermalization *within* each fragment is slow.

This general perspective is powerful because it maps the problem of thermalization in these models onto a well-known algorithmic problem, the word problem for semigroups (a perspective also adopted by Hastings and Freedman in Ref. [37] to provide examples of dynamics

exhibiting “topological obstructions” that provide a separation between the performances of QMC and quantum annealing). The word problem is the problem referenced above, namely, that of deciding whether two words represent the same semigroup element. This identification allows us to lift examples of computationally hard word problems from the mathematical literature to construct models with anomalously slow dynamics. In these models, the dynamics connects the basis states within each fragment in a manner that is much sparser than in generic systems (see Fig. 1 for an illustration), and it is this phenomenon that leads to long thermalization times.

The first part of this work focuses on models where the word problem takes an exponentially long time to solve. We place particular focus on the “Baumslag-Solitar model,” a spin-2 model with three-site interactions for which the relaxation time of a large class of initial states under any type of local dynamics (Hamiltonian, random unitary, etc.) is provably exponentially long in the system size. This model has a conserved charge, and this exponentially long timescale shows up as an exponentially slow hydrodynamic relaxation of density gradients. Not only is the relaxation timescale anomalous, but so is the functional form: A state with density gradients relaxes “gradually, then suddenly,” with an initially prepared density wave experiencing almost no relaxation for exponentially long times but then undergoing a sudden collapse at a sharply defined timescale. Despite this extremely slow hydrodynamics, the states involved are *not* dynamically frozen: Each configuration is rapidly locally fluctuating, and generic local autocorrelation functions decay rapidly.




FIG. 1. Schematic illustration of the difference between the Hilbert space connectivity of generic dynamics and semigroup dynamics, with each yellow dot representing a computational-basis product state in a single connected sector of the dynamics. In a 1D system of length L , typical dynamics (left) requires at most $O(L)$ steps of the dynamics (applications of a Hamiltonian or layers of a unitary circuit) to move between any two basis states. In semigroup dynamics (right), the number of steps needed scales as the Dehn function $\text{Dehn}(L)$ of the semigroup in question, which measures the word problem’s temporal complexity. When $\text{Dehn}(L)$ is large, the basis states in each sector are connected very sparsely, leading to long thermalization times.

In the second part of this work, we focus on word problems that have high *spatial* complexity: In other words, solving them requires not just many steps but also a large amount of additional space for scratch work. In other words, in these problems, deciding whether two words of length L are equivalent requires a derivation involving intermediate words much longer than L .

In the corresponding dynamical systems, one has (for any fixed L) pairs of basis states $|a\rangle, |b\rangle$ such that (i) $|a\rangle$ is not connected to $|b\rangle$ by the dynamics of a chain of length L , but (ii) $|a\rangle \otimes |c\rangle$ is connected to $|b\rangle \otimes |c\rangle$ by the dynamics of a longer chain, with $|c\rangle$ a fixed ancillary product state. This phenomenon can be viewed in two complementary ways: as a “fragile” form of Hilbert space fragmentation [38] as an unusual type of jamming that has no counterpart in known examples of jammed systems. Using constructions similar to the group model discussed above, we can construct examples where the amount of additional spatial resources grows extremely rapidly with L , not just exponentially but also as e^{e^L} , $e^{e^{e^L}}$, and so on.

This paper is organized as follows. In Sec. II, we introduce word problems for semigroups and groups, and relate them to fragmentation. In Sec. III, we use the complexity of the word problem to derive bounds on thermalization times, and in Sec. IV, we explore an explicit example of a high-complexity group word problem that yields dynamics with exponentially slow relaxation. We present numerical evidence and analytical estimates for the anomalously slow hydrodynamics of this model. In Sec. V, we introduce and analyze a family of group models with fragile fragmentation. Sections VI and VII, respectively, present examples based on semigroups and generalize our one-dimensional examples to two-dimensional loop models. Finally, we conclude with a discussion of future directions in Sec. VIII.

II. SEMIGROUP DYNAMICS AND CONSTRAINED 1D SYSTEMS

In this section, we introduce the general framework used to construct the models described above. We will refer to this framework as semigroup dynamics, which encompasses a general class of constrained dynamical systems whose constraints can be derived from the presentation of the underlying group or semigroup (to be defined below). These types of constraints are particularly appealing from a theoretical point of view: It turns out that we can rigorously characterize many properties—thermalization times, fragmentation, and so on—using tools from the field of geometric group theory. Broadly speaking, geometric group theory is concerned with characterizing the complexity and geometry of discrete groups (see Ref. [40] for an accessible introduction), ideas that will be made precise in the following.

As a starting point, we describe the necessary mathematical background needed to motivate group dynamics.

This discussion will center around the word problem, a century-old problem lying at the heart of results regarding the geometry and complexity of groups. We then see how algorithms solving the word problem can naturally be encoded into the dynamics of 1D spin chains, whose thermalization dynamics is controlled by the word problem’s complexity. Finally, we see how the structure of the word problem leads to Hilbert space fragmentation and discuss the properties of the group that control the severity of the fragmentation.

Throughout this paper, we mostly study constrained dynamics on 1D spin chains whose on-site Hilbert space is finite dimensional. [41] We only consider systems whose time evolution has constraints that can be specified in a local tensor product basis (referred to throughout as the computational basis), either directly or after the application of finite-depth local unitary circuits. In the latter case, we assume the unitary transformation has been done, to avoid loss of generality.

We let S denote the set of computational-basis state labels for the on-site Hilbert space, with individual basis states being written as letters in typewriter font ($|a\rangle, |b\rangle$, etc.):

$$\mathcal{H}_{\text{loc}} = \text{span}\{|a\rangle : a \in S\}. \quad (1)$$

Strings of letters are used as shorthand for tensor products, so, e.g., $|\text{word}\rangle = |w\rangle \otimes |o\rangle \otimes |r\rangle \otimes |d\rangle$. A ket with a single roman letter denotes a product state of arbitrary length, e.g., $|w\rangle = |\text{word}\rangle$.

A. Dynamical constraints and semigroups

We write $\text{Dyn}(t)$ to denote time evolution for time t under the dynamics in question, which may be performed using a set of unitary gates, a (possibly space- or time-dependent) Hamiltonian, or a bistochastic Markov chain. Having a dynamical constraint means that not all computational-basis product states $|w\rangle, |w'\rangle$ can be connected under the dynamics. We use $\varphi(|w\rangle)$ to denote the dynamical sector of the state $|w\rangle$, defined as the set of all computational-basis product states that $|w\rangle$ can evolve to; thus,

$$\langle w | \text{Dyn}(t) | w' \rangle \propto \delta_{\varphi(|w\rangle), \varphi(|w'\rangle)} \quad \forall t. \quad (2)$$

The tensor product of computational-basis states—that is, the stacking of one system onto the end of another—defines a binary operation on the dynamical sector labels, which we write as $\odot$:

$$\varphi(|w\rangle \otimes |w'\rangle) = \varphi(|w\rangle) \odot \varphi(|w'\rangle). \quad (3)$$

Since the tensor product is associative, so is $\odot$. Thus, the set of dynamical sectors is equipped with an associative binary operation, thereby endowing it with the structure of a semigroup, a generalization of a group that need not have

inverses or an identity element. [42] Since all $|w\rangle$ are formed as tensor products of the single-site basis states $|a\rangle$, the $\varphi(|a\rangle)$ —which we will write simply as a to save space—constitutes a generating set for this semigroup. The dynamical sector of a state $|w\rangle = |a_1\rangle \otimes \cdots \otimes |a_L\rangle$ is then determined simply by multiplying a_i along the length of the chain [43]:

$$\varphi(|w\rangle) = a_1 \odot \cdots \odot a_L. \quad (4)$$

Following usual group theory notation, we denote a semigroup with generating set S as

$$G = \text{semi}\langle S|R \rangle, \quad (5)$$

where R denotes a set of relations imposed on the product states that can be formed from elements of the on-site computational-basis states S . When writing presentations of groups we omit the inverse generators and the identity from S , and we likewise omit trivial relations like $aa^{-1} = e, ae = ea$ from R . For the remainder of the paper, presentations of semigroups that are not groups will always be denoted by $\text{semi}\langle S|R \rangle$, while presentations of groups will be denoted simply by $\langle S|R \rangle$.

The relations in R are determined by φ , namely, by which product states are related to one another under Dyn . Consider any two states $|w\rangle, |w'\rangle$ such that $\varphi(|w\rangle) = \varphi(|w'\rangle)$ define the same element of G . Since we are interested in dynamics that are geometrically local, it must be possible to relate $|w\rangle$ to $|w'\rangle$ using a series of *local* updates to $|w\rangle$. Thus, the set R must be expressible in terms of a set of equivalence relations that each involve only an $O(1)$ number of the elements of S —implying, in particular, that $|R|$ must be finite. A semigroup where both S and R are finite is said to be finitely presented, and all of the semigroups we consider are of this type.

Given a semigroup G , we use the notation Dyn_G to represent a general local dynamical process acting on $\mathcal{H} = \mathcal{H}_{\text{loc}}^{\otimes L}$ which satisfies the constraint (2) and hence preserves the dynamical sectors of all computational-basis states. The locality of the dynamics means that Dyn_G must be composed of elementary blocks (unitary gates or Hamiltonian terms) of constant length, which, when acting on computational-basis product states, implement the relations contained in R . Writing the relations in R as $r_l = r_r$ with $r_{l/r} = a_{l/r,1} \cdots a_{l/r,n}$, the locality of Dyn_G is determined by the maximal length of the $r_{l/r}$, which we denote as ℓ_R :

$$\ell_R = \max_{r_{l/r} \in R} |r_{l/r}|. \quad (6)$$

Here, ℓ_R will always be $O(1)$ (and when G is a group, one can show that there always exists a finite presentation of G such that $\ell_R \leq 3$; see Appendix A for the proof).

To be more explicit, first consider the case where Dyn_G corresponds to time evolution under a (geometrically) local Hamiltonian H . The semigroup constraint and locality of H show that $\langle w'|H|w\rangle$ can be nonzero only if the words w, w' differ by the *local* application of a relation in R . Consequently, H assumes the general form

$$H = \sum_i \sum_{r \in R} (\lambda_{i,r} |r_l\rangle \langle r_r|_i + \text{H.c.}), \quad (7)$$

where $|r_{l/r}\rangle_i = |a_{l/r,1}\rangle_i \otimes \cdots \otimes |a_{l/r,n}\rangle_{i+n}$ and $\lambda_{i,r}$ are arbitrary complex numbers. Note that the above Hamiltonian is only well defined if $|r_l| = |r_r|$ for all relations r . In cases where this does not hold, we will rectify this problem by adding a trivial character e to the on-site Hilbert space—with e defined to commute with all of the other generators a —which allows us to then “pad” the relations r in a way that ensures that $|r_l| = |r_r|$.

As a simple example, consider the group $\mathbb{Z}^2 = \langle x, y | xy = yx \rangle$, which, as we will see later, in some sense has trivial dynamics. Since the full generating set $S = \{x, x^{-1}, y, y^{-1}, e\}$ of this presentation has dimension $|S| = 5$, a Hamiltonian $H_{\mathbb{Z}^2}$ with $\mathbb{Z}^2$ -constrained dynamics thus acts most naturally on a spin-2 chain. The single nontrivial relation $xy = yx$ has length $\ell_R = 2$, and thus $H_{\mathbb{Z}^2}$ can be taken to be 2-local, assuming a form like

$$\begin{aligned} H_{\mathbb{Z}^2} = & \sum_i \sum_{a \in S} (\lambda_{1,i} |aa^{-1}\rangle \langle ee|_{i,i+1} + \lambda_{2,i} |a^{-1}a\rangle \langle ee|_{i,i+1} \\ & + \lambda_{3,i} |ae\rangle \langle ea| + \text{H.c.}) \\ & + \sum_i (\lambda_{4,i} |xy\rangle \langle yx| + \text{H.c.}) + \cdots, \end{aligned} \quad (8)$$

which describes two species of conserved particles, each of which defines a $U(1)$ conserved charge $n_a = \sum_i |a\rangle \langle a| - |a^{-1}\rangle \langle a^{-1}|$, where $a = x, y$. The explicit examples we consider in this work will not be more complicated than $H_{\mathbb{Z}^2}$ in terms of their degree of locality or the dimension of $\mathcal{H}_{\text{loc}}$, but their dynamical properties will be much richer.

The construction of group-constrained random unitary dynamics is similar to the Hamiltonian case. For random unitary dynamics, Dyn_G is constructed using ℓ_R -site unitary gates U whose matrix elements $\langle w'|U|w\rangle$ are nonzero only if the length- ℓ_R words w, w' satisfy $\varphi(|w\rangle) = \varphi(|w'\rangle)$. Such unitaries admit the decomposition

$$U = \bigoplus_{g \in G_{\ell_R}} U_g, \quad G_{\ell_R} = \{g | \exists |w\rangle : \varphi(|w\rangle) = g, |w| \leq \ell_R\}, \quad (9)$$

where G_{ℓ_R} denotes the set of elements of G expressible as products of precisely ℓ_R generators. A particularly natural realization of Dyn_G is when each U_g is drawn from an appropriate-dimensional Haar ensemble, although we do not need to specify to this case.

Existing examples of constrained 1D dynamics in the literature—from multipole conserving systems to models based on cellular automata and other constrained classical systems [21–28]—are all described by Dyn_G for an appropriate semigroup G and an appropriate kind of dynamics (random unitary, Hamiltonian, etc.) [44]. One of the main messages of this paper is that in addition to providing an organizing framework for discussing 1D constrained dynamics, approaching matters from this point of view enables a large arsenal of mathematical tools from the field of geometric group theory to be brought to bear, leading to general bounds on thermalization times, precise characterizations of ergodicity breaking, and explicit examples of models with extremely slow dynamics.

B. Word problem

We now formulate the semigroup word problem, a concept key for determining the thermalization behavior of models with semigroup dynamics. We refer to a computational-basis product state—defined by a string of generators in S , e.g., $w = a_1 \cdots a_L$ —as a *word*. In what follows, we will often slightly abuse notation by letting the symbol w stand for both an abstract string $a_1 \cdots a_L$ of generators in S and the associated computational-basis product state $|a_1 \cdots a_L\rangle$.

Words are naturally grouped into equivalence classes labeled by elements of G . Letting $W(S)$ denote the set of all words, we define these equivalence classes as

$$K_g \triangleq \{w \in W(S) | \varphi(|w\rangle) = g\}. \quad (10)$$

Any two words belonging to the same equivalence class can be deformed into one another by applying a sequence of relations in R . For any two $w, w' \in K_g$, we define a derivation from w to w' , written as $D(w \rightsquigarrow w')$, as the sequence of words appearing in this deformation:

$$D(w \rightsquigarrow w') = w \rightarrow u_1 \rightarrow u_2 \rightarrow \cdots \rightarrow u_n \rightarrow w', \quad (11)$$

where each arrow $\rightarrow$ indicates applying a *single* relation from R .

In Sec. III, we will see that the way in which Dyn_G thermalizes is determined by the complexity of the word problem for G , a fundamental problem in the fields of abstract algebra and computability theory. The word problem is defined by the following question:

Word problem: Given two words w, w' , does $\varphi(|w\rangle) = \varphi(|w'\rangle)$? In other words, is there a derivation $D(w \rightsquigarrow w')$?

A key result is that even in the case where both $|S|$ and $|R|$ are small, answering this question can be very difficult (even undecidably so; see Appendix C). For semigroups or groups that do not have an undecidable word problem, a

key problem is to determine the time and space complexity of algorithms that solve it. In what follows, we will introduce two functions characterizing the word problem's complexity: the Dehn function, which governs its time complexity, and the expansion length, which governs its space complexity.

1. Time complexity

Our operational definition of the time complexity of the word problem is the minimum length of a derivation linking w to w' , as illustrated in Fig. 2(a). [45] We denote this by $\text{Dehn}(w_1, w_2)$:

$$\text{Dehn}(w_1, w_2) \triangleq \min_D |D(w \rightsquigarrow w')|. \quad (12)$$

Note that $\text{Dehn}(w, w')$ measures the nondeterministic time complexity of the word problem since it is the maximum runtime of an algorithm that maps w to w' by blindly applying all possible relations in R to w in parallel, and it halts the first time w' appears in the resulting superposition of words. Time evolution under Dyn_G can be naturally regarded as a way of simulating this process, a connection enabling the derivation of the bounds found in Sec. III.

For any word w , we define the *length* $|w|$ of w as the number of generators appearing in w . To denote the subset of length- L words in K_g , we write

$$K_{g,L} \triangleq \{w \in K_g | |w| = L\} \quad (13)$$

as the set of length- L words in K_g (or equivalently, following our practice of letting w stand for both a word and a computational-basis product state, as the collection of product states associated with such words). The (worst-case) time complexity across all words in $K_{g,L}$ defines the function

$$\text{Dehn}_g(L) \triangleq \max_{|w|, |w'| \in K_{g,L}} \text{Dehn}(w, w'). \quad (14)$$

We are particularly interested in how $\text{Dehn}_g(L)$ scales asymptotically with L . In Appendix B, we show that this scaling is the same for any two finite presentations of the same semigroup. Thus, we may meaningfully speak about the Dehn function of a *semigroup* rather than a particular presentation thereof. This presentation independence imparts some degree of the robustness to the dynamical properties that we will derive.

For the case where G is a *group* (rather than just a semigroup), a fair amount more can be said. All groups have a distinguished identity element e , and in Appendix B, we show that [46]

$$\text{Dehn}_e(L) \gtrsim \text{Dehn}_g(L) \quad \forall g \in G. \quad (15)$$




FIG. 2. Semigroup dynamics and the word problem. (a) In the word problem, we are given a length- L word $|w\rangle$ and a series of rewriting rules that let us make local updates to the characters of the word. The word problem for $|w\rangle$ is the task of deciding whether or not a sequence of allowed updates can be found that transforms $|w\rangle$ into a particular reference word, here chosen to be $|e^L\rangle$. The Dehn function $\text{Dehn}(w)$ measures the time complexity of the word problem, namely, the minimal number of updates needed to connect $|w\rangle$ to $|e^L\rangle$. (b) In some situations, $|w\rangle$ can only be transformed into $|e^L\rangle$ by increasing the amount of available space, done here by appending $|ee\cdots\rangle$ onto the original word. The minimal amount of extra space needed defines the expansion length $\text{EL}(w)$, which captures the spatial complexity of the word problem. (c) For a semigroup $G = \langle S|R \rangle$ defined by generators $a \in S$ and relations between the generators $r_i \in R$, our construction defines a dynamics Dyn_G that acts on a 1D chain with an on-site Hilbert space $\mathcal{H}_{\text{loc}} = \{|a\rangle : a \in S\}$. The dynamics implemented by Dyn_G is restricted to local updates that preserve the semigroup element obtained by multiplying the generators along the length of the chain; the allowed updates are consequently fixed by the relations in R (with the figure drawn using the relations $r_1: ab = cd$, $r_2: ef = gh$, etc.).

We furthermore show that, as long as $|K_{g,L}|$ scales exponentially in L , then $\text{Dehn}_e(L) \sim \text{Dehn}_g(L)$. For groups, we thus define

$$\text{Dehn}(L) \triangleq \text{Dehn}_e(L) \quad (16)$$

as a simpler characterization of the word problem's time complexity. The calculation of $\text{Dehn}_e(L)$ also simplifies further for groups, as we may fix $|w'\rangle = |e^L\rangle$ in Eq. (14) without changing the asymptotic scaling of $\text{Dehn}(L)$. Thus, for groups, we mostly focus on computing

$$\text{Dehn}(L) = \max_{|w\rangle \in K_{e,L}} \text{Dehn}(w, e^L) \quad (G \text{ a group}). \quad (17)$$

Even for groups where $|S|, |R|$ are both $O(1)$, $\text{Dehn}(L)$ can scale in many different ways. To start, it is easy to verify that $\text{Dehn}(L) \sim L$ for all finite groups and that

$$G \text{ Abelian} \Rightarrow \text{Dehn}(L) \lesssim L^2, \quad (18)$$

with $\text{Dehn}(L) \sim L^2$ only when G is infinite. Indeed, for all such groups, such as the $G = \mathbb{Z}^2$ example above, $\text{Dehn}(L)$ is bounded from above by the time it takes to transport the conserved charges $n_a \triangleq \sum_i (|a\rangle\langle a|_i - |a^{-1}\rangle\langle a^{-1}|_i)$ across the system, which is about L^2 . For our purposes, we regard any G with $\text{Dehn}(L) \lesssim L^2$ as uninteresting, as for these groups Dyn_G thermalizes on timescales generically no slower than for conventional systems with conserved $U(1)$ charges.

A simple example of a group with an “interesting” Dehn function is the discrete Heisenberg group H_3 , which has the group presentation

$$H_3 = \langle x, y, z | xy = yxz, xz = zx, yz = zy \rangle \quad (19)$$

and possesses a Dehn function scaling as $\text{Dehn}(L) \sim L^3$ [47]. In Sec. IV, we give an example of a simple group where $\text{Dehn}(L) \sim 2^L$ and, in Sec. V, one with $\text{Dehn}(L) \sim 2^{2^L}$; examples of semigroups with similarly slow dynamics are given in Sec. VI. Going beyond these examples, Refs. [48,49] remarkably show that for any constant α , almost any function with growth $L^2 \leq f(L) < L^\alpha$ is the Dehn function of some finitely presented semigroup, including, for example, “unreasonable-looking” functions such as $L^{7\pi}$, $L^2 \log L$, and $L^e \log(L)^{\log(L)} \log \log L$.

In addition to the worst-case complexity of the word problem, we also consider its average-case complexity (both for groups and semigroups), a quantity that has received much less attention in the math literature (the only exception we are aware of is Ref. [50]). To this end, we define the typical Dehn function $\overline{\text{Dehn}}_g(L)$ as the number of steps needed to map a certain constant fraction of words in $K_{g,L}$ to one another:

$$\overline{\text{Dehn}}_g(L) \triangleq \sup\{t : \Pr_{w,w' \in K_{g,L}} [\text{Dehn}(w, w') \geq t] \geq 1/2\}. \quad (20)$$

Establishing rigorous results about $\overline{\text{Dehn}}$ is unfortunately much more difficult than for Dehn , and the landscape of

typical Dehn functions is comparatively less well explored. For the examples we focus on in this paper, however, a combination of physical arguments and numerics will nevertheless suffice to understand the rough asymptotic scaling of $\overline{\text{Dehn}}_g(L)$.

2. Space complexity

Another complexity measure is the (nondeterministic) space complexity of the word problem. [51] The space complexity is nontrivial in cases where, during the course of being deformed into w' , a word w must expand to a length $L' > L$. More formally, we define the expansion length [52] of a derivation $D(w \rightsquigarrow w') = w \rightarrow u_1 \rightarrow \dots \rightarrow u_n \rightarrow w'$ as the maximal length of the intermediate words u_i :

$$\text{EL}(D(w \rightsquigarrow w')) \triangleq \max_{u_i \in D(w \rightsquigarrow w')} |u_i|. \quad (21)$$

The relative expansion length $\text{EL}(w, w')$ between two words is then defined as the minimal expansion length of a derivation connecting them:

$$\text{EL}(w, w') \triangleq \min_{D(w \rightsquigarrow w')} \text{EL}(D(w \rightsquigarrow w')), \quad (22)$$

as illustrated in Fig. 2(b). The expansion length of a $K_{g,L}$ sector is likewise

$$\text{EL}_g(L) \triangleq \max_{w, w' \in K_{g,L}} \text{EL}(w, w'). \quad (23)$$

Similarly, with the Dehn function, we show in Appendix B that when G is a group, $\text{EL}_g(L) \lesssim \text{EL}_e(L)$ for all g , so for groups, we may use

$$\text{EL}(L) \triangleq \text{EL}_e(L) \quad (G \text{ a group}) \quad (24)$$

as a simple metric of the space complexity.

In Appendix B, we show that $\text{EL}(L) \lesssim L$ for all Abelian groups and that all finite groups have $\text{EL}(L) \leq L + C$ for some constant C ; such scalings are “uninteresting” from the perspective of space complexity. Just as with the Dehn function though, there exist simple semigroups for which $\text{EL}(L)$ grows extremely fast with L , the consequences of which will be explored in Sec. V.

C. Semigroup dynamics and Hilbert space fragmentation

The existence of multiple $K_{g,L}$ sectors means that Dyn_G is not ergodic as long as G is not a presentation of the trivial semigroup. The simplest case is when G is an Abelian group. In this case, the $K_{g,L}$ can be associated with the symmetry sectors of a global symmetry. This case is true simply because the sector that a given product state lives in can be determined by computing the expectation value of

the operators $n_g = \sum_i (|g\rangle\langle g|_i - |g^{-1}\rangle\langle g^{-1}|_i)$ for each generator g . Thus, Dyn_G for Abelian G is already very well understood, given the plethora of work on thermalization in the presence of global symmetries.

When G is not an Abelian group, global symmetries may still be present, but there inevitably exist other nonlocal conserved quantities that distinguish different dynamical sectors. Indeed, in Appendix D, we prove that the dynamical sectors of such models are *never* described by global symmetries alone. Since Dyn_G always has nonlocal conserved quantities if G is not an Abelian group, the lack of ergodicity due to these quantities leads to HSF [19,21–23,53], a phenomenon whereby the dynamics of initial states becomes trapped in disconnected subspaces of $\mathcal{H}$ —in our notation, simply the $K_{g,L}$ —whose existence cannot be attributed to the presence of global symmetries alone.

The original works on HSF [21,22] focused mainly on fragmentation in spin systems with conserved dipole moments. While these systems fall within our semigroup framework, [54] we instead find it more instructive to review the pair-flip model introduced in Ref. [55]. The pair-flip model is described by a spin- s Hamiltonian of the following form:

$$H_{\text{PF}} = \sum_i \sum_{a,b=-s}^s \lambda_{ab,i} (|aa\rangle\langle bb|_{i,i+1} + \text{H.c.}). \quad (25)$$

For generic choices of $\lambda_{ab,i}$, the model is nonintegrable. Note that the product states $|w\rangle$ where $w = a_1 \dots a_L$ are annihilated by H_{PF} if $a_i \neq a_{i+1}$ for all i . These product states alone provide $(2s+1)(2s)^{L-1}$ dynamically disconnected dimension-1 sectors not attributable to any local symmetry, meaning that H_{PF} exhibits HSF.

The dynamics generated by H_{PF} is in fact a special case of our construction applied to the group

$$\mathbb{Z}_2^{*(2s+1)} = \langle g_1, \dots, g_{2s+1} | g_i^2 = e \rangle, \quad (26)$$

where $*$ denotes the free product. Note that H_{PF} can be obtained from our general construction by considering a modified on-site Hilbert space $\mathcal{H}' = \text{span}\{|g\rangle, |g^{-1}\rangle : g \in S\}$, which contains no e generator. The relation $g_i^2 = e$ can be rewritten without e as $g_i^2 = g_j^2$ for all i, j , and a general local Hamiltonian acting on $(\mathcal{H}')^{\otimes L}$ that obeys the $\mathbb{Z}_2^{*(2s+1)}$ may accordingly be written as

$$H = \sum_i \sum_{g,h \in S} (\lambda_{gg,i} |gg\rangle\langle hh|_{i,i+1} + \text{H.c.}), \quad (27)$$

which matches the Hamiltonian in Eq. (25). Unfortunately, the Dehn function of $\mathbb{Z}_2^{*(2s+1)}$ is easily seen to scale as

$\text{Dehn}(L) \sim L^2$ and thus does not provide a complexity scaling that is unusual enough to warrant further study.

Having discussed the concept of HSF, let us further understand the structure and size of disconnected sectors under the dynamics Dyn_G . The number of such disconnected sectors is *at least* as large as the number of subspaces, each labeled by $K_{g,L}$ for some $g \in G$. Thus, the number of such subspaces—which we denote as $N_K(L)$ —depends on the number of semigroup elements expressible as words of length less than or equal to L . In particular, we define the geodesic length $|g|$ of an element $g \in G$ as the length of the shortest word representing g :

$$|g| = \min_{w \in K_g} |w|. \quad (28)$$

A word w with $\varphi(|w\rangle) = g$ satisfying $|w| = |g|$ is called a geodesic word. Then,

$$N_K(L) = |\{g \mid |g| \leq L\}|. \quad (29)$$

Semigroups with more elements “close” to the identity thus have dynamics with a larger number of Hilbert space fragments. As an example, one may readily verify that in the pair-flip model, $N_K(L) \sim (2s)^L$ grows exponentially as long as $s > 1/2$.

For many models with group constraints, including the pair-flip model, $N_K(L)$ exactly determines the number of Hilbert space fragments. However, for some semigroups, $N_K(L)$ is not the full story, with each $K_{g,L}$ further fragmenting into subspaces in a way controlled by the expansion-length function defined in Eq. (23). Understanding this phenomenon is the subject of Sec. V.

In the study of HSF, a distinction is often made between “strong” and “weak” HSF. This distinction was originally discussed in the context of Hamiltonian dynamics [21], where it was defined by violations of strong and weak ETH, respectively. Since we focus our discussion at a level where the nature of Dyn_G may or may not involve eigenstates, we instead adopt a slightly different definition in terms of the size of the largest $K_{g,L}$ sector, which we denote as $K_{\max,L}$ (in Appendix B, we prove that, for groups, the largest sector is in fact always the one associated with the identity, $K_{\max,L} = K_{e,L}$). We say that the dynamics is

- (1) weakly fragmented if $|K_{\max,L}|/|\mathcal{H}| \rightarrow O(1)$ as $L \rightarrow \infty$, and
 - (2) strongly fragmented if $|K_{\max,L}|/|\mathcal{H}| \rightarrow 0$ as $L \rightarrow \infty$.
- We find it useful to subdivide the strongly fragmented case into additional classes according to how quickly $|K_{\max,L}|/|\mathcal{H}|$ vanishes as $L \rightarrow \infty$. We say that the fragmentation in the case of strong HSF is
- (1) polynomially strong if $|K_{\max,L}|/|\mathcal{H}| \sim 1/\text{poly}(L)$, and
 - (2) exponentially strong if $|K_{\max,L}|/|\mathcal{H}| \sim \exp(-L)$.

Note that the above definitions are made without reference to any global symmetry sectors. Thus, if global symmetries happen to be present, the quantum numbers associated with them will constitute part of the elements g labeling the different dynamical sectors, and the above definition of strong or weak HSF will simply single out the largest, regardless of its symmetry quantum number(s). When symmetries are present, one could also quantify the degree of fragmentation by first fixing a quantum number, changing $K_{\max,L}$ to be the largest sector having that quantum number, and replacing $|\mathcal{H}|$ by the dimension of the chosen symmetry sector. However, since generic Dyn_G dynamics need not have any global symmetries, and since the result of the above procedure can depend sensitively on the chosen quantum number [56–58], we focus only on the above (simpler) definition, which maximizes over all symmetry sectors.

Our models provide a way of addressing two questions raised in Ref. [21]. The first question was whether or not 1D models exist with $0 < |K_{\max,L}|/|\mathcal{H}| < 1$ in the thermodynamic limit (known examples with weak fragmentation all have $|K_{\max,L}|/|\mathcal{H}| \rightarrow 1$ in the thermodynamic limit). Our construction answers this question in the affirmative, with examples provided by Dyn_G for any finite non-Abelian G (e.g., $G = S_3$).

The second question concerns the existence of models where $|K_{\max,L}|/|\mathcal{H}|$ vanishes more slowly than exponentially as $L \rightarrow \infty$ after specifying to a fixed symmetry sector (in fact, an affirmative answer to this question was already provided by the spin-1 Motzkin chain introduced in Ref. [59], where $|K_{\max,L}|/|\mathcal{H}| \sim L^{-3/2}$ [60]). Our models provide (many) more examples of this phenomenon, as one need only let G be a group with $L^2 \lesssim \text{Dehn}(L) \lesssim L^\infty$, a simple example being the discrete Heisenberg group H_3 [61]. One may further ask whether there are strongly fragmented systems where $|K_{\max,L}|/|\mathcal{H}|$ decays at a rate in between $1/\text{poly}(L)$ and $\exp(-L)$. The answer to this question is again affirmative, with one such example being the focus of Sec. IV.

III. SLOW THERMALIZATION AND THE TIME COMPLEXITY OF THE WORD PROBLEM

From the discussion of the previous section, it is natural to expect that systems with Dyn_G dynamics will have thermalization times controlled by the time complexity of the word problem, as diagnosed by the Dehn functions $\text{Dehn}_g(L)$ defined in Eq. (14). In this section, we make this relationship precise by using the functions $\text{Dehn}_g(L)$ to place lower bounds on various thermalization timescales. In the case of random unitary or classical stochastic dynamics, $\text{Dehn}_g(L)$ will be used to bound relaxation and mixing times; for Hamiltonian dynamics, it will appear in bounds for hitting times.

A common feature these timescales have is that they quantify when Dyn_G is able to spread initial product states across Hilbert space. This feature cannot usually be probed by looking at correlation functions of local operators, which would be the preferred method for thinking about thermalization. While the bounds derived in this section do not mandate that the relaxation of *local operators* also be controlled by $\text{Dehn}_g(L)$, in Sec. IV we will explore an explicit example in which they are. Until then, we focus on the more “global” diagnostics of relaxation and hitting times.

We note in passing that it is impossible to use $\text{Dehn}_g(L)$ —or any other quantity—to place *upper* bounds on thermalization timescales without making any additional assumptions about the details of Dyn_G (as without additional assumptions, we could always choose Dyn_G to be the evolution with a many-body localized Hamiltonian, and the relevant timescales would all diverge). Even in the case where Dyn_G is an appropriately constrained form of random unitary evolution, upper bounding the relaxation timescales requires techniques beyond those employed below and constitutes an interesting direction for future work.

A. Circuit dynamics

We first discuss the case where Dyn_G is generated by a G -constrained random unitary circuit, which can be mapped to the case where Dyn_G is a classical Markov process. We assume that Dyn_G is expressed as a brickwork circuit whose gates act on ℓ_R sites, with ℓ_R the maximum size of a relation in R . Note that $\text{Dyn}_G(t)$ will be used to denote $t \in \mathbb{N}$ time steps of this dynamics, with each time step consisting of ℓ_R staggered layers of gates.

Let us first look at how operators evolve under $\text{Dyn}_G(t)$. We use overbars to denote averages over circuit realizations, so acting on an operator $\mathcal{O}$ with one layer of the brickwork (corresponding to a time of $t = 1/\ell_R$) gives

$$\begin{aligned} \overline{\mathcal{O}(t = \ell_R^{-1})} &= \overline{\text{Dyn}_G^\dagger(\ell_R^{-1}) \mathcal{O} \text{Dyn}_G(\ell_R^{-1})} \\ &= \mathbb{E}_{\{U_{j,g}, U_{j',g'}\}} \bigoplus_j \bigoplus_{g \in G_{\ell_R}} U_{j,g}^\dagger \mathcal{O} \bigoplus_{j'} \bigoplus_{g' \in G_{\ell_R}} U_{j',g'}, \end{aligned} \quad (30)$$

where each $U_{j,g}$ acts on a length- ℓ_R block of sites and G_{ℓ_R} , as before, denotes the set of group elements expressible as words of length less than or equal to ℓ_R . Assuming the $U_{j,g}$ are drawn uniformly from the $|K_{g,\ell_R}|$ -dimensional Haar ensemble, [62] performing the average gives

$$\overline{\mathcal{O}(t = \ell_R^{-1})} = \bigoplus_j \left(\sum_{g \in G_{\ell_R}} \frac{\text{Tr}[\mathcal{O}_j \Pi_{K_{g,\ell_R}}]}{|K_{g,\ell_R}|} \Pi_{K_{g,\ell_R}} \right), \quad (31)$$

where we have defined

$$\Pi_{K_{g,\ell_R}} \triangleq \sum_{w \in K_{g,\ell_R}} |w\rangle\langle w| \quad (32)$$

as the projector onto the space of length- ℓ_R words w with $\varphi(|w\rangle) = g$, as well as—without loss of generality—taken $\mathcal{O}$ to factorize as $\mathcal{O} = \bigotimes_j \mathcal{O}_j$.

Thus, after a single step of the dynamics, all operators completely dephase and become diagonal in the computational basis, and operators violating the dynamical constraint evaluate to zero under Haar averaging. We may thus focus on diagonal operators without loss of generality, which we indicate using vector notation as $|\mathcal{O}(t)\rangle$. With this notation, Eq. (31) becomes

$$|\overline{\mathcal{O}(t = \ell_R^{-1})}\rangle = \mathcal{M}_1 |\mathcal{O}\rangle \quad (33)$$

with the matrix

$$\mathcal{M}_1 = \left(\sum_{g \in G_{\ell_R}} \tilde{\Pi}_{K_{g,\ell_R}} \right)^{\otimes L/\ell_R}, \quad (34)$$

where

$$\tilde{\Pi}_{K_{g,\ell_R}} \triangleq \frac{1}{|K_{g,\ell_R}|} \sum_{w,w' \in K_{g,\ell_R}} |w\rangle\langle w'| \quad (35)$$

projects onto the uniform superposition of states within K_{g,ℓ_R} .

Different layers of the brickwork likewise define matrices $\mathcal{M}_i$, with $i = 1, \dots, \ell_R$, where i denotes the staggering. Defining

$$\mathcal{M} \triangleq \prod_{i=1}^{\ell_R} \mathcal{M}_i, \quad (36)$$

diagonal operators evolve as

$$|\overline{\mathcal{O}(t)}\rangle = \mathcal{M}^t |\mathcal{O}\rangle. \quad (37)$$

Here, $\mathcal{M}$ is a symmetric [64] doubly stochastic matrix, with the smallest eigenvalue of 0 and the largest eigenvalue of 1. Because of the group constraint, $\mathcal{M}$ does not have a unique steady state, following from the fact that the Markovian dynamics is reducible as

$$\mathcal{M} = \bigoplus_{g \in G_L} \mathcal{M}_g. \quad (38)$$

However, each $\mathcal{M}_g$ defines an irreducible aperiodic Markov chain, whose unique steady state is the uniform distribution $|\pi_g\rangle \triangleq (1/|K_{g,L}|) \sum_{w \in K_{g,L}} |w\rangle$. Infinite-temperature circuit-averaged correlation functions are determined by the

mixing time and spectral gap of the $\mathcal{M}_g$, which we now relate to the Dehn function.

First, consider the mixing time of $\mathcal{M}_g$, which we may view as a characterization of the thermalization timescale within $K_{g,L}$:

$$t_{\text{mix}}(g) \triangleq \min\{t: \max_{|w\rangle \in K_{g,L}} \|\mathcal{M}_g^t |w\rangle - |\pi_g\rangle\|_1 \leq 1/2\}. \quad (39)$$

A basic result [65] about $t_{\text{mix}}(g)$ is that it is lower bounded by half the diameter of the state space that $\mathcal{M}_g$ acts on simply because, in order to mix, the system must, at the minimum, be able to traverse across most of state space. Thus, we have the bound

$$t_{\text{mix}}(g) \geq \frac{\text{Dehn}_g(L)}{2}. \quad (40)$$

However, this bound is, in fact, too loose. Intuitively, to saturate the bound, $\text{Dyn}_G(t)$ would need to immediately find the optimal path between any two nodes in configuration space. Since $\mathcal{M}$ generates a random walk on state space, the dynamics will instead diffusively explore state space in a less efficient manner. On general grounds, one might therefore expect a bound on $t_{\text{mix}}(g)$ that is the square of the rhs above. This guess, in fact, turns out to be essentially correct, with

$$t_{\text{mix}}(g) \geq \frac{\text{Dehn}_g^2(L)}{16 \ln |K_{g,L}|}, \quad (41)$$

which follows from Prop. 13.7 of Ref. [66] after using the fact that the equilibrium distribution of $\mathcal{M}_g$ is always uniform on K_g by virtue of $\mathcal{M}_g$ being doubly stochastic. Since $|K_{g,L}|$ is at most exponential in L , we thus have

$$t_{\text{mix}}(g) \geq C_g \frac{\text{Dehn}_g^2(L)}{L} \quad (42)$$

for some g -dependent $O(1)$ constant C_g . We generically expect the random walk generated by $\mathcal{M}$ to be the fastest mixing local Markov process; hence, we expect it to saturate the above bound on $t_{\text{mix}}(g)$, a prediction that we will confirm numerically for the example in Sec. IV.

Correlation functions for operators computed in states in $K_{g,L}$ are controlled by the relaxation time t_{rel} of $\mathcal{M}_g$, defined by the inverse gap of $\mathcal{M}_g$:

$$t_{\text{rel}}(g) \triangleq \frac{1}{1 - \lambda_2}, \quad (43)$$

where λ_2 is the second-largest eigenvalue of $\mathcal{M}_g$. This quantity admits a similar bound to $t_{\text{mix}}(g)$, as one can see that [65]

$$t_{\text{rel}}(g) \geq \frac{t_{\text{mix}}(g)}{\ln(2|K_{g,L}|)}, \quad (44)$$

and hence

$$t_{\text{rel}}(g) \geq C'_g \frac{\text{Dehn}_g^2(L)}{L^2}, \quad (45)$$

with C'_g another $O(1)$ constant. Thus, $\text{Dehn}_g(L)$ places lower bounds on both mixing and the decay of correlation functions. Note that the operators whose correlators decay as $t_{\text{rel}}(g)$ need not be local; indeed, the obvious ones to consider are projectors like $|w\rangle\langle w|$. In Sec. IV, we will, however, explore an example that possesses local operators that relax according to Eq. (45).

Note that the mixing and relaxation times are worst-case measures of the thermalization time. We can additionally define a “typical” mixing time $\overline{t}_{\text{mix}}(g)$ as

$$\overline{t}_{\text{mix}}(g) \triangleq \min\{t: \Pr_{|w\rangle}(\|\mathcal{M}_g^t |w\rangle - |\pi_g\rangle\|_1 \leq 1/2) \geq 1/2\}, \quad (46)$$

where $\Pr_{|w\rangle}$ indicates the probability over words uniformly drawn from $K_{g,L}$. By following similar logic as above, one can derive similar bounds on $\overline{t}_{\text{mix}}(g)$ in terms of the typical Dehn function.

B. Hamiltonian dynamics

We have thus far discussed mixing times of the Markov chains that arise from random unitary dynamics. However, purely Hamiltonian dynamics does not mimic that of a Markov chain, and being reversible, it possesses no direct analogues of mixing and relaxation times. Nevertheless, the Dehn function can still be used to place bounds on the timescales taken for time evolution to spread wave functions across Hilbert space. To be more quantitative, we focus on the hitting time $t_{\text{hit}}(g)$ of Dyn_G , which we define as the minimum time needed for product states in $K_{g,L}$ to “reach” all other product states in $K_{g,L}$. Since $\langle w' | e^{-iHt} | w \rangle$ will generically be nonzero for all $|w\rangle, |w'\rangle \in K_{g,L}$ as soon as $t > 0$, we need a slightly different definition of $t_{\text{hit}}(g)$ as compared with the Markov chain case.

To define $t_{\text{hit}}(g)$ more precisely, we first define the hitting amplitude between two computational-basis product states $|w\rangle, |w'\rangle \in K_{g,L}$ as

$$h_{ww'}(t; g) \triangleq |\langle w' | \text{Dyn}_G(t) | w \rangle|^2. \quad (47)$$

Note that $h_{w\cdot}(t; g)$ is a probability distribution on $K_{g,L}$, with $\sum_{w' \in K_{g,L}} h_{ww'}(t; g) = 1$ for all $|w\rangle \in K_{g,L}$. We define $t_{\text{hit}}(g, w, w')$ between two words w and w' as the minimum time for which $h_{ww'}(t; g)$ reaches a fixed fraction of its infinite-time average $\overline{h}_{ww'}(g)$, defined by

$$\overline{h_{ww'}}(g) \triangleq \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt h_{ww'}(t; g). \quad (48)$$

The hitting time is the maximum over all pairs of words $|w\rangle, |w'\rangle \in K_{g,L}$ of $t_{\text{hit}}(g, w, w')$, which is written as

$$t_{\text{hit}}(g) \triangleq \max_{|w\rangle, |w'\rangle \in K_{g,L}} \min \left\{ t : h_{ww'}(t; g) > \frac{1}{2} \overline{h_{ww'}}(g) \right\}. \quad (49)$$

To bound $t_{\text{hit}}(g)$, we first evaluate the hitting amplitude as

$$\begin{aligned} h_{ww'}(t; g) &= |\langle w' | e^{-iHt} | w \rangle|^2 \\ &= \left| \sum_{k=0}^{\infty} \frac{t^k}{k!} \langle w' | (-iH)^k | w \rangle \right|^2. \end{aligned} \quad (50)$$

We know that the minimum k such that $\langle w' | H^k | w \rangle \neq 0$ is given by $\text{Dehn}_g(w, w')$, which, for brevity, we denote by $d_{ww'}$ in the following. The first $d_{ww'}$ terms in the above sum will thus vanish; thus, truncating the above Taylor series at the leading nonzero term, we apply the remainder theorem to find

$$h_{ww'}(t; g) \leq \left(\frac{t^{d_{ww'}} |\langle w' | H^{d_{ww'}} | w \rangle|}{d_{ww'}!} \right)^2. \quad (51)$$

To diagnose the hitting time, we need to compare $h_{ww'}(t; g)$ with its long-time average $\overline{h_{ww'}}(g)$ by relating the above matrix element $|\langle w' | H^{d_{ww'}} | w \rangle|^2$ to $\overline{h_{ww'}}(g)$ as follows. Writing $\{|E\rangle\}$ for H 's eigenbasis, we find

$$\begin{aligned} |\langle w | H^{d_{ww'}} | w' \rangle|^2 &= \left| \sum_E \langle w | E \rangle \langle E | w' \rangle E^{d_{ww'}} \right|^2 \\ &\leq \left(\sum_E |\langle w | E \rangle|^2 |E|^{2d_{ww'}} \right) \left(\sum_E |\langle E | w' \rangle|^2 \right), \end{aligned} \quad (52)$$

where the second line follows from the Cauchy-Schwarz inequality. The first factor in parentheses is simply $\overline{h_{ww'}}(g)$, which can be readily verified:

$$\begin{aligned} \overline{h_{ww'}}(g) &\triangleq \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt h_{ww'}(t; g) \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{E, E'} \langle w | E \rangle \langle E | w' \rangle \langle w' | E' \rangle \langle E' | w \rangle \\ &\quad \times \int_0^T e^{-i(E-E')t} dt, \end{aligned} \quad (53)$$

which, assuming that the spectrum of $H|_{K_{g,L}}$ is nondegenerate (which we assume to be the case throughout this paper), gives

$$\overline{h_{ww'}}(g) = \sum_E |\langle w | E \rangle|^2 |\langle E | w' \rangle|^2. \quad (54)$$

Inserting this into Eq. (52), we obtain

$$|\langle w | H^{d_{ww'}} | w' \rangle|^2 \leq |\mathcal{H}| \overline{h_{ww'}}(g) \times \|H\|_{\infty}^{2d_{ww'}}. \quad (55)$$

Since H is local, $\|H\|_{\infty} = \lambda L$ for some $O(1)$ constant λ . We may thus write Eq. (51) as

$$h_{ww'}(t; g) \leq \left(\frac{(\lambda L t)^{d_{ww'}}}{d_{ww'}!} \right)^2 |\mathcal{H}| \overline{h_{ww'}}(g). \quad (56)$$

As long as $d_{ww'} = \Omega(L)$, $h_{ww'}(t; g)$ is much smaller than its equilibrium value when $t = \Omega(d_{ww'}/L)$. Indeed, from Stirling's approximation, we have

$$h_{ww'}(\eta d_{ww'}/L; g) \leq \frac{(\lambda \eta e)^{2d_{ww'}}}{2\pi d_{ww'}} |\mathcal{H}| \overline{h_{ww'}}(g), \quad (57)$$

which can always be made exponentially small in L by choosing η appropriately if $d_{ww'} = \Omega(L)$. We conclude that

$$t_{\text{hit}}(g) \geq \frac{\text{Dehn}_g(L)}{L}, \quad (58)$$

which is essentially the same bound as our naive result (40) for $t_{\text{rel}}(g)$ in the case of random circuit evolution. It would be interesting to see if this bound could be improved to the square of the rhs, as in Eq. (45).

Finally, we note that just as with the mixing time, a *typical* hitting time may also be defined as

$$\overline{t_{\text{hit}}}(g) \triangleq \sup \{ t : \Pr_{w, w' \in K_{g,L}} (t_{\text{hit}, ww'} \geq t) \geq 1/2 \}. \quad (59)$$

Arguments similar to the above then show that $\overline{t_{\text{hit}}}(g)$ admits a similar bound in terms of $\overline{\text{Dehn}}_g(L)$.

IV. BAUMSLAG-SOLITAR GROUP: EXPONENTIALLY SLOW HYDRODYNAMICS

The previous two sections have focused on developing parts of the general theory of semigroup dynamics. In this section, we take an in-depth look at a particular example that exhibits anomalously long relaxation times.

Our example will come from a family of groups whose Dehn functions scale exponentially with L , yielding word problems with large time complexity. [67] When $\text{Dehn}(L) \sim \exp(L)$, the hitting and mixing times discussed in the previous section are exponentially long. It is perhaps already rather surprising that a translation-invariant local Markov process or unitary circuit can have mixing times that are this long, but our example is most interesting for another reason: It *also* possesses a global symmetry whose conserved charge takes $\text{Dehn}(L)$ time to relax. This feature

allows the slow time complexity of the word problem to be manifested in the expectation values of simple local operators, rather than being hidden in hitting times between different product states.

Our example comes from a family of groups known in the math literature as the Baumslag-Solitar groups [68], which are parametrized by two integers $n, m \in \mathbb{N}$. Each group $\text{BS}(n, m)$ in this family is generated by two elements a and b obeying a single nontrivial relation, with the group presentation

$$\text{BS}(n, m) = \langle a, b | ba^m b^{-1} = a^n \rangle. \quad (60)$$

Models with $\text{Dyn}_{\text{BS}(n, m)}$ dynamics are thus most naturally realized in spin-2 systems with $\max(n, m)$ -local dynamics and on-site Hilbert space

$$\mathcal{H}_{\text{loc}} = \text{span}\{|a\rangle, |a^{-1}\rangle, |b\rangle, |b^{-1}\rangle, |e\rangle\}. \quad (61)$$

The simplest Baumslag-Solitar group that is interesting for our purposes is $\text{BS}(1, 2)$, and the rest of our discussion will focus on this case. For convenience, in what follows, we will write $\text{BS}(1, 2)$ simply as BS .

The nontrivial relation in BS reads

$$ab = ba^2, \quad (62)$$

so a generators duplicate when moved to the right of b generators. By taking inverses, the a generators are also seen to duplicate when moved to the left of b^{-1} generators:

$$b^{-1}a = a^2b^{-1}. \quad (63)$$

This duplication property means that an $O(n)$ number of b and b^{-1} generators can be used to grow an a generator by an amount of order $O(2^n)$, as

$$b^{-n}ab^n = a^{2^n}. \quad (64)$$

We will see momentarily how this exponential growth can be linked to the Dehn function of BS , which also scales exponentially.

Another key property of Eq. (62) is that it preserves the number of b 's. Thus, models with Dyn_G dynamics possess a $U(1)$ symmetry generated by the n_b , defined as

$$n_b \triangleq \sum_i n_{b,i} \triangleq \sum_i (|b\rangle\langle b|_i - |b^{-1}\rangle\langle b^{-1}|_i). \quad (65)$$

It is this conserved quantity that will display the exponentially slow hydrodynamics advertised above.

A. Geometric perspective on group complexity

To understand how dynamics in Dyn_{BS} works, we find it helpful to introduce a *geometric* way of thinking about the

group word problem. [69] Given a general discrete group G , we let CG_G denote the Cayley graph of G . Recall that CG_G is a graph with vertices labeled by elements of G and edges labeled by generators of G and their inverses, with two vertices h, k being connected by an edge g iff $h = gk$. As simple examples, $\text{CG}_{\mathbb{Z}_N}$ for $\mathbb{Z}_N = \langle g | g^N = e \rangle$ is a length- N closed cycle; $\text{CG}_{\mathbb{Z}^2}$ for $\mathbb{Z}^2 = \langle x, y | xy = yx \rangle$ is a 2D square lattice; and $\text{CG}_{\mathbb{Z} * \mathbb{Z}}$ for the free group $\mathbb{Z} * \mathbb{Z} = \langle x, y \rangle$ is a Bethe lattice with coordination number 4 [70].

It is useful to realize that from any Cayley graph CG_G , we can always construct a related 2-complex by associating oriented faces (or 2-cells) to each of CG_G 's elementary closed loops; the 2-cells have the property that the product of generators around their boundary is a relation in R . For the above-mentioned examples, $\mathbb{Z}_N$ would have one N -sided face, $\mathbb{Z}^2$ would have a face for each plaquette of the square lattice (with the generators around the plaquette boundaries reading $xyx^{-1}y^{-1}$), and the free product $\mathbb{Z} * \mathbb{Z}$ would have no faces at all (due to its lack of nontrivial relations). The 2-complex thus constructed is known as the Cayley 2-complex of G , and we abuse notation by also referring to it as CG_G .

Any group word $w \in W(G)$ naturally defines an oriented path in CG_G obtained by starting at an (arbitrarily chosen) origin of CG_G and moving along edges based on the characters in w . The end point of this path on the Cayley graph is thus the group element $\varphi(|w\rangle)$. Additionally, applying local relations in R to w deforms this path while keeping its end points fixed. This process gives a geometric interpretation of the subspaces $K_{g,L}$:

$$K_{g,L} = \{P | P = \text{length-}L \text{ path from origin to } g \text{ in } \text{CG}_G\}. \quad (66)$$

In particular, the (largest) subspace $K_{e,L}$ is identified with the set of all length- L closed loops in CG_G . The number $N_K(L)$ of Krylov sectors is the number of vertices of CG_G located a distance of less than or equal to L from the origin.

Having provided a relationship between Krylov subspaces and the Cayley 2-complex, what is the geometric interpretation of the Dehn function? Based on the definition (17), we can restrict our attention to deformations $D(w \rightsquigarrow e^L)$ for $w \in K_{e,L}$, which are simply homotopies that shrink the loop defined by w down to a point. The loop passes across one cell at each step, and the number of steps in $D(w \rightsquigarrow e^L)$ is the area enclosed by the surface swept out by the homotopy. In particular, the Dehn function of w is

$$\text{Dehn}(w, e^L) = \min_{S: \partial S = w} \text{area}(S), \quad (67)$$

where the minimum is over surfaces in CG_G with boundary w ; we thus refer to

$$\text{area}(w) \triangleq \text{Dehn}(w, e^L) \quad (68)$$

as the *area* of w . The Dehn function of the group is then the largest area of a word in $K_{e,L}$:

$$\text{Dehn}(L) = \max_{\text{loops } w \text{ of length } L} \text{Dehn}(w, e^L). \quad (69)$$

This perspective is important as it gives the algorithmic definition of $\text{Dehn}(L)$ a geometric meaning. The large-scale geometry of CG_G thus directly affects the complexity of the word problem and, consequently, the thermalization dynamics of Dyn_G .

B. Geometry of BS and the fragmentation of Dyn_{BS}

The simple examples (discrete groups, Abelian groups, and free groups) discussed above are all geometrically uninteresting, but the BS group is a notable exception. Note that CG_{BS} has the structure of an infinite branching tree of hyperbolic planes; it is illustrated in Fig. 3. To understand how this happens, recall that $b^{-n}ab^n = a^{2^n}$. Thus, for all n , $b^{-n}ab^na^{-2^n}$ forms a closed loop in CG_{BS} . This closed loop gives rise to a tiling of the hyperbolic plane, as shown in Fig. 3(a). Letting multiplication by a correspond to the motion along the $\hat{x}$ direction and multiplication by b to the motion along $\hat{y}$, the hyperbolic structure comes from the fact that moving by 2^n sites along

the $\hat{x}$ direction of the Cayley graph can either be accomplished by a direct path (multiplication by a^{2^n}) or a path that requires exponentially fewer steps, which first traverses n steps along the $\hat{y}$ direction before moving along $\hat{x}$ (multiplication by $b^{-n}ab^n$).

The full geometry of CG_{BS} is more complicated than a single hyperbolic plane, however. As shown in Fig. 3(a), this can be seen from the fact that the word bab cannot be embedded into the hyperbolic plane. Instead, this word must “branch out” into a new sheet, which also forms a hyperbolic plane. Figure 3(b) illustrates the consequences of this process for the Cayley graph, whose full structure consists of an infinite number of hyperbolic planes glued together in the fashion of a binary tree (formally, the presentation complex of BS is homeomorphic to $\mathbb{R} \times T_3$, where T_3 is a 3-regular tree).

The locally hyperbolic structure of CG_{BS} means that the number of vertices within a distance L of the origin—and hence the number of Krylov sectors $N_K(L)$ —grows exponentially with L . In Appendix E, we argue that the base of the exponent is very nearly $\sqrt{3}$:

$$N_K(L) \approx \sqrt{3}^L. \quad (70)$$




FIG. 3. Group geometry of $\text{BS}(1, 2)$. The Cayley graph of $\text{BS}(1, 2)$ is constructed from an infinite number of two-dimensional sheets glued together in a treelike fashion. (a) Single sheet of the Cayley graph, which resembles a hyperbolic plane. Orange arrows denote multiplication by a and purple arrows multiplication by b , with the boundary of each plaquette being the defining relation $aba^{-1}a^{-1}b^{-1} = e$. The diagonal purple lines denote edges that connect to different sheets; all diagonal edges connected to vertices at the same “height” connect to the same sheet. (b) How different sheets are glued together to form the full Cayley graph. The shaded yellow region denotes how a section of a single sheet is embedded in the full Cayley graph. The bold dashed path denotes the path traced out by a “large” word $w_{\text{large}}(n) = ab^{-n}a^{-1}b^na^{-1}b^{-n}ab^n$, which is homotopic to the identity ($\varphi(|w_{\text{large}}|) = e$) and possesses an area scaling exponentially with its length. In the figure, $n = 3$, and the path on the Cayley graph is ordered by points 1 to 8, according to the 8 different components of w_{large} (read from right to left). Thus, the point 2 corresponds to the word b^3 , the point 5 corresponds to $a^{-1}b^{-3}ab^3$, and so on.

Interestingly, despite having exponentially many Hilbert space fragments, it is known that the largest sector $K_{e,L}$ occupies a fraction of the full Hilbert space $\mathcal{H}$ that decreases subexponentially with L :

$$\frac{|K_e|}{|\mathcal{H}|} \sim e^{-\alpha L^{1/3}}, \quad (71)$$

where the numerical results in Fig. 4—obtained by sampling random words and postselecting on membership in $K_{e,L}$ —indicate that $\alpha \approx 1.84$. Models with Dyn_{BS} dynamics provide an example (indeed, the first that we are aware of) of a strongly fragmented model where the largest sector occupies neither an exponentially nor polynomially small fraction of Hilbert space.

C. Dehn function

We finally have the necessary tools to explain why the Dehn function of BS grows exponentially with L . Define the word

$$w_{\text{large}}(n) \triangleq \text{ab}^{-n} \text{a}^{-1} \text{b}^n \text{a}^{-1} \text{b}^{-n} \text{ab}^n, \quad (72)$$

which is of length $|w_{\text{large}}(n)| = 4n + 4$, belongs to $K_{e,L}$, and is shown for $n = 3$ in Fig. 3 as the thick dashed line in panel (b). As a path, this word can be broken into two legs. The first leg moves to the vertex labeled by a^{2^n} by first going “up” the hyperbolic plane of a given sheet, traversing one step along $\hat{x}$, and then coming back “down.” The second leg moves back to the origin but does so by passing along a *different* sheet of the tree. From Fig. 3, it is clear that the area of the minimal surface bounding $w_{\text{large}}(n)$ is exponentially large; more precisely, it is




FIG. 4. Probability p_e for a randomly chosen word to lie in the identity sector, determined according to the procedure described in Appendix E. The dashed line is a fit to $p_e \propto e^{-\alpha L^{1/3}}$ with $\alpha \approx 1.84$.

$$\text{area}(w_{\text{large}}(n)) = 2 \sum_{i=0}^{n-1} 2^i = 2^{n+1} - 2, \quad (73)$$

which scales exponentially with n . This construction thus shows that $\text{Dehn}(L) = \Omega(2^L)$ [71]. This bound is in fact tight (see Appendix E), so

$$\text{Dehn}(L) = \Theta(2^L). \quad (74)$$

From the results of Sec. III, this finding implies that Dyn_G has exponentially long mixing, relaxation, and hitting times.

A more sophisticated treatment needs to be given in order to understand the scaling of the *typical* Dehn function $\overline{\text{Dehn}}(L)$. To our knowledge, this question has not been answered in the math literature. While we will not provide a completely rigorous proof of $\overline{\text{Dehn}}(L)$ ’s scaling, a combination of physical arguments and numerics—which we relegate in their entirety to Appendix E—indicates that

$$\overline{\text{Dehn}}(L) \sim 2^{\sqrt{L}}. \quad (75)$$

The rough intuition behind this result is that a typical closed-loop path in $K_{e,L}$ will roughly execute a random walk along the tree part of CG_{BS} , reaching a “depth” of $\sqrt{L}$; from this point, it is then able to enclose an area exponential in this depth. Giving a rigorous proof of Eq. (75) could be an interesting direction for future research.

D. Exponentially slow hydrodynamics

An observation about the word $w_{\text{large}}(n)$ is that it contains a density wave of bs, with the profile of $n_{b,i}$ looking like two periods of a square wave as a function of i . Utilizing the fact that $\text{area}(w_{\text{large}}(n)) \sim 2^n$, one sees that this density wave takes an exponentially long time to relax, with a thermalization time $t_{\text{th}}(n) \sim [\text{area}(w_{\text{large}}(n))]^2 = \Omega(2^{2n})$ [where the square of the area comes from the square in Eq. (45)]. Furthermore, almost any word with an n_b density wave will admit a similar bound on t_{th} . Indeed, consider words $w_{\text{large}}(n, L)$ obtained by extending $w_{\text{large}}(n)$ to length $L > 4n + 4$ by inserting $L - (4n + 4)$ random characters at random points of $w_{\text{large}}(n)$ (we will assume, for simplicity, that in fact $w_{\text{large}}(n, L) \in K_{e,L}$, but this assumption is not crucial). The only way for $\text{area}(w_{\text{large}}(n, L))$ to be significantly smaller than $\text{area}(w_{\text{large}}(n))$ is if the added characters cancel out almost the entirety of the b density wave or cancel a large number of a’s located near the peaks and troughs of the density wave. For a random choice of $w_{\text{large}}(n, L)$, these situations are exponentially unlikely to occur, and we expect

$$\text{area}(w_{\text{large}}(n, L)) = \Omega(2^n) \quad (76)$$

with probability 1 in the large- n limit. The long relaxation time of the density wave is attributed to the long mixing time of BS because of its large Dehn function (see Sec. III). This feature is remarkable because probing the large-scale complexity of BS only requires studying the dynamics of *local* operators, namely, those that overlap with n_b . We henceforth will denote the relaxation and mixing times by t_{th} .

Of course, since 2^n is the *smallest* possible time needed to map $w_{\text{large}}(n)$ to e^{4n+4} , it gives only a lower bound on t_{th} . In the case where $\text{Dyn}_{\text{BS}}(t)$ is generated by a classical Markov process or random unitary circuit dynamics, $\text{Dyn}_G(t)$ effectively leads to words executing a random walk on the configuration space of loops. Due to the diffusive nature of this random walk, we expect that the true thermalization time instead scales as the square of its lower bound, namely,

$$t_{\text{th}}(w_{\text{large}}(n)) = O(2^{2n}), \quad (77)$$

an estimate that we will confirm shortly in Sec. IV E.

To examine the relaxation in more detail, consider a state $|w_{A,q}\rangle$ that contains a n_b density wave of momentum q and amplitude A but that is otherwise random. In other words, the expectation value of $n_{b,i}$ in $|w_{A,q}\rangle$ is (switching to schematic continuum notation)

$$n_b(x) = A \sin(qx), \quad (78)$$

and $|w_{A,q}\rangle$ is chosen randomly from the set of all states in $K_{e,L}$ that have this expectation value (with the restriction to $K_{e,L}$ done purely for notational simplicity).

The “depth” n_A that $w_{A,q}$ proceeds into CG_{BS} is equal to the contrast of the density wave, which we define as the integral of the b density over a quarter period of the density wave:

$$n_A \triangleq \int_0^{\pi/2q} n_b(x) dx = \frac{A}{q}. \quad (79)$$

We thus expect the density wave defined by $|w_{A,q}\rangle$ to have a thermalization timescale of

$$t_{\text{th}} = O(2^{2n_A}) = O(2^{2A/q}). \quad (80)$$

Note that this exponential timescale is not visible in the standard linear-response limit, in which one takes $A \rightarrow 0$ before taking $q \rightarrow 0$: The two limits lead to qualitatively different relaxation. In the standard linear-response limit, fluctuations at scale q will decorrelate on a *typical* relaxation timescale of about $2^{\sqrt{L}}$.

Beyond t_{th} , we would also like to know how the amplitude A —or, equivalently, the contrast n_A —behaves as a function of time. To estimate this, consider first how $w_{\text{large}}(n)$ relaxes. By the geometric considerations of

Sec. IV B, the shortest homotopy reducing $w_{\text{large}}(n)$ to the identity word must first map $w_{\text{large}}(n) \rightarrow w_{\text{large}}(n-1)$, then $w_{\text{large}}(n-1) \rightarrow w_{\text{large}}(n-2)$, and so on; see Fig. 9 for an illustration. The self-similarity of this process means that the time needed to map $w_{\text{large}}(m) \rightarrow w_{\text{large}}(m-1)$, namely, $t_{\text{th}}(w_{\text{large}}(m-1))$, is equal to the combined time needed to perform all of the maps $w_{\text{large}}(k) \rightarrow w_{\text{large}}(k-1)$, with $k < m$. Applying this observation to the density wave under consideration and using Eq. (77), we estimate

$$n_A(q, t_{\text{th}}(1-2^{-k})) \approx n_A(q, 0) - k \quad (81)$$

for $k \leq n_A(0)$. Writing k on the rhs in terms of the time $t = t_{\text{th}}(1-2^{-k})$, these arguments suggest that the density relaxes as

$$n_b(q, t) \approx \Theta(t_{\text{th}} - t)(n_{b,0} + \log_2(1 - t/t_{\text{th}} + 2^{-n_{b,0}}t/t_{\text{th}})), \quad (82)$$

where $t_{\text{th}} = C2^{2A/q}$, with C an $O(1)$ constant, $n_{b,0}$ is shorthand for $n_b(q, 0)$, and the $2^{-n_{b,0}}t/t_{\text{th}}$ inside the logarithm ensures that $n_b(q, t_{\text{th}}) = 0$.

An interesting consequence of Eq. (82) is that the relaxation of the density wave happens “all at once” in the large-contrast limit (e.g., small q at fixed A), meaning that the density profile remains almost completely unchanged until a time very close to t_{th} , at which point the density wave is abruptly destroyed. Indeed, we define the collapse timescale

$$t_{\text{col}}(\varepsilon) \triangleq \min\{t : n_A(q, t) < (1 - \varepsilon)n_A(q, 0)\} \quad (83)$$

as the time at which the collapse of the density wave becomes noticeable within a precision controlled by the constant $0 < \varepsilon < 1$. Then, the form (82) implies that

$$\lim_{\varepsilon \rightarrow 0} \lim_{A/q \rightarrow \infty} \frac{t_{\text{col}}(\varepsilon)}{t_{\text{th}}} = 1, \quad (84)$$

suggesting that the collapse is instantaneous in the large-contrast limit.

These arguments show that density waves collapse abruptly but do not tell us about the statistical distribution of collapse times obtained when considering an ensemble of realizations of the dynamics (specific disorder realizations of random unitary circuits, specific choices of Hamiltonian, etc.). In systems that relax suddenly, the width $\sigma_{t_{\text{th}}}$ of this distribution can be either parametrically smaller (in system size) than the mean collapse time t_{th} , or it can scale as $\Omega(t_{\text{th}})$. The former case occurs in Markov chains displaying a “cutoff” (see, e.g., Refs. [72,73]), which can occur when the chain describes random motion on a highly connected graph, or when it describes relaxation in a metastable system [where relaxation is

caused by rare events that yield a Poissonian scaling $\sigma_{t_{\text{th}}} \sim \sqrt{t_{\text{th}}} = o(t_{\text{th}})$. The latter case is typical in systems which relax diffusively, with an unbiased random walk on $\mathbb{Z}_d$ being a typical example. From the geometric considerations of this section, it is perhaps not clear which scenario should apply *a priori*. However, the numerics of the following section tentatively point to the latter scenario, with the distribution of thermalization times roughly obeying $\sigma_{t_{\text{th}}} \sim t_{\text{th}}$. Giving a more precise characterization of this scaling would be interesting to explore in future work.

E. Numerics

We now present the results of numerical simulations that let us take a more detailed look into the relaxation of n_b and confirm the predictions made in the previous subsection.

1. Stochastic circuits

Our simulations all treat Dyn_{BS} as time evolution under BS-constrained random unitary circuits. Because off-diagonal operators are rendered diagonal after a single step of random unitary dynamics [as was shown in Eq. (31)], without loss of generality, we can focus on the evolution of diagonal operators. In Sec. III, we showed that the product states associated with diagonal operators evolve in time according to the stochastic matrix $\mathcal{M}$ derived in Eq. (34). Explicitly, since the maximal size of an elementary relation in BS is $\ell_R = 3$ (e.g., $\text{abe} = \text{baa}$), $\mathcal{M}$ is most naturally constructed using three brickwork layers of three-site gates:

$$\mathcal{M} \triangleq \mathcal{M}_1 \mathcal{M}_2 \mathcal{M}_3, \quad \mathcal{M}_a \triangleq \bigotimes_{i=0+a}^{\lfloor L/3 \rfloor - 1} M_{3i+a}, \quad (85)$$

where, as in Eq. (34), the matrices M_i induce equal-weight transitions among all dynamically equivalent three-letter words:

$$M_i \triangleq \sum_{g_1, \dots, g_6 \in \mathcal{S} \cup \mathcal{S}^{-1} \cup \{e\}} \delta_{g_1 g_2 g_3, g_4 g_5 g_6} \times \frac{1}{|K_{g_1 g_2 g_3}(3)|} |g_1, g_2, g_3\rangle \langle g_4, g_5, g_6|_{i, i+1, i+2}, \quad (86)$$

where the $1/|K_{g_1 g_2 g_3}(3)|$ factor ensures that M_i is stochastic—in fact, doubly so on account of $M_i^T = M_i$. The steady-state distribution $|\pi_g\rangle$ of $\mathcal{M}$ within $K_{g,L}$ is accordingly given by the uniform distribution on $K_{g,L}$:

$$|\pi_g\rangle \triangleq \frac{1}{|K_{g,L}|} \sum_{w \in K_{g,L}} |w\rangle. \quad (87)$$

In practice, we do not actually diagonalize $\mathcal{M}$ but rather use the matrix elements of $\mathcal{M}$ to randomly sample updates

that may be applied to a computational-basis state $|w\rangle$, with the system thus remaining in a product state at all times. A single time step in our simulations corresponds to a single brickwork layer of $\mathcal{M}$, i.e., to the application of a single relation at each three-site block. Since $\lfloor L/3 \rfloor$ relations can be applied at each time step, in these units, it is $\text{Dehn}(L)/L$ —rather than $\text{Dehn}(L)$ —that lower bounds the mixing and relaxation times of $\mathcal{M}$.

2. Slow relaxation

We start by exploring how long it takes the state $|w_{\text{large}}(n, L)\rangle$ to relax under $\mathcal{M}$'s dynamics, where as above $w_{\text{large}}(n, L)$ is obtained from $|w_{\text{large}}(n)\rangle$ by padding it with $L - (4n + 4)$ identity characters inserted at random positions. This process is accomplished by initializing the system in $|w_{\text{large}}(n, L)\rangle$ and tracking the local $n_{b,i}$ density over time. We focus, in particular, on the moving average of the n_b density and the fluctuations thereof, defined as

$$\begin{aligned} \langle n_{b,i} \rangle(t) &\triangleq \frac{1}{T} \sum_{t'=t}^{t+T} n_{b,i}(t'), \\ \overline{\langle n_b \rangle^2}(t) &\triangleq \frac{1}{L} \sum_{i=1}^L \langle n_{b,i} \rangle^2, \end{aligned} \quad (88)$$

where $n_{b,i}(t) = \delta_{g_i(t), b} - \delta_{g_i(t), b^{-1}}$ is the b charge on site i at time step t [with $g_i(t)$ the i th entry of the state at time t] and T is a time window that is small relative to the thermalization timescales of interest but large enough to suppress short-time statistical fluctuations $n_{b,i}(t)$ (here, $\langle \cdot \rangle$ denotes averaging over this time window, while $\bar{\cdot}$ denotes averaging over space). The equilibrium distribution $|\pi_g\rangle \propto \sum_{w \in K_{g,L}} |w\rangle$ for $g = e$ satisfies $\langle \pi_e | n_{b,i} | \pi_e \rangle \approx 0$ for all i (see Appendix E), so for initial states $|w\rangle \in K_{e,L}$, any nonzero value of $\langle n_{b,i} \rangle(t)$ indicates a lack of equilibration (for $g \neq e$, the average $\langle \pi_g | n_{b,i} | \pi_g \rangle$ can be nonzero and must first be computed in order to diagnose equilibration).

The evolution of $\langle n_{b,i} \rangle(t)$ for a single realization of the dynamics initialized in $|w_{\text{large}}(n, L)\rangle$ is shown in Fig. 5(a) for $n = L/10$ and $L = 120$. From this evolution, we can clearly see the sudden collapse phenomena predicted above in Eq. (84): The n_b density wave hardly decays at all until very close to $t_{\text{th}} \sim 5 \times 10^6$, at which point $\langle n_{b,i} \rangle(t)$ rapidly becomes very close to zero. This finding is quantified in Fig. 5(b), which plots the spatially averaged n_b fluctuations $\overline{\langle n_b \rangle^2}$ for the same realization. A relatively good fit (dashed line) is obtained using the square of the “sudden collapse” function defined in Eq. (82).

We now investigate how the thermalization time of $|w_{\text{large}}(n, L)\rangle$ scales with n , continuing to fix $n/L = 10$ so as to keep both n, L extensive. Operationally, we define the thermalization time as the first time when the magnitude




FIG. 5. Stochastic circuit time evolution of observables associated with the b charge for the system initialized in a state $|w_{\text{large}}(n = L/10, L)\rangle$, which is a large-area word $w_{\text{large}}(n = L/10)$ with identities inserted at random places. (a) Time evolution of the spatial profile of b charge, $\langle n_{b,i} \rangle$, where averaging is performed over a time window of $T = 10^5$ brickwork layers. (b) Time evolution of observable $\langle n_b \rangle^2$ from Eq. (88). Thermalization occurs at time t_{th} , when this observable drops to zero, which corresponds to the b -charge density wave collapsing to a flat profile. The run corresponding to panel (a) is shown in gray, while the blue curve corresponds to an average over several independent runs, with the red shade showing the standard deviation. Time for each run has been rescaled by the respective thermalization time. The analytic formula from Eq. (82) is shown by the dashed green line (rescaled by the value at $t = 0$). (c) Thermalization time scaling exponentially with the system size L if the density of b 's in the initial word is kept fixed, $n = L/10$.

of fluctuations in n_b drops below a fixed fraction of their initial value:

$$t_{th} \triangleq \min\{t: \overline{\langle n_b \rangle^2}(t) < \frac{1}{10} \overline{\langle n_b \rangle^2}(0)\}. \quad (89)$$

Figure 5(c) shows the scaling of t_{th} with L , which is observed to admit an excellent fit to the predicted scaling of around 2^{2n} . As shown in Fig. 6, the *distribution* of thermalization times across different runs is additionally observed to be rather broad, with a standard deviation that scales approximately in the same way as t_{th} .

While this result confirms that the density wave present in $|w_{\text{large}}(n, L)\rangle$ relaxes on a timescale of $t_{th} \sim 2^{2n}$, it does




FIG. 6. Distribution of thermalization times t_{th} of the density wave defined by $|w_{\text{large}}(n, L)\rangle$ with $n = L/10$, shown for different system sizes with 300 circuit realizations. A rather broad distribution is observed.

not show that *all* states with b density waves of amplitude A and wavelength q relax on times of order $2^{2A/q}$. In fact, this cannot be true, and fast-relaxing density waves always exist. Indeed, consider the word $w_{\text{small}}(n, L)$ obtained from $w_{\text{large}}(n, L)$ by replacing all a s and a^{-1} s with e s. The Dehn time of this word is merely $\text{Dehn}(w_{\text{small}}(n, L)) \sim n^2$ since there are no a s present to slow down the dynamics of the b s (any aa^{-1} pairs created in between the segments of the density waves have a net zero number of a s and are thus ineffectual at providing a slowdown). Thus, the relaxation time of an initial state $|w\rangle$ carrying an n_b density wave cannot be predicted from knowledge of the conserved density alone—one must also have some knowledge about the distribution of a s in $|w\rangle$.

However, we expect that $t_{th} \sim 2^{2A/q}$ for *generic* states containing an amplitude- A , wavelength- q density wave. Indeed, as argued above near Eq. (76), this is simply because as long as the value of the a -charge $\sum_i (|a\rangle\langle a|_i - |a^{-1}\rangle\langle a^{-1}|_i)$ is nonzero in the regions between the segments of the density wave—which will almost always be true for a random density-wave state in the thermodynamic limit—the a s trapped “inside” the density wave will take an exponentially long time to escape. Note that this remark applies to a generic density-wave state $|w\rangle$, regardless of whether or not $|w\rangle \in K_{e,L}$.

In Fig. 7, we numerically investigate the relaxation of generic density waves by considering initial states that host density waves of momentum $q = 4\pi/3L$ and amplitude $A = nq$ (with n fixed at $n = L/10$) but that are otherwise random. Compared with $|w_{\text{large}}(n, L)\rangle$, these random density waves (which generically lie in different $K_{g,L}$ sectors) exhibit a much broader range of thermalization times; some thermalize very quickly, while some never thermalize




FIG. 7. Relaxation of n_b under stochastic circuit dynamics for initial states with random b density waves. The initial states are chosen to be words of the form $w_1 b^n w_2 b^{-n} w_3 b^n w_4 b^{-n} w_5$, where $n = L/10$ and the $w_{1,\dots,5}$ are random words containing only the characters $\{a, a^{-1}, e\}$. (a) Time evolution of observable $\overline{\langle n_b \rangle^2}$ from Eq. (88) for several independent runs with random initial states described above and the time averaging window $T = 100$. The time at which the density wave collapses and the final equilibrium value of $\overline{\langle n_b \rangle^2}$ are seen to change significantly for different choices of initial state, with some runs thermalizing faster than T (low-lying blue lines) and some runs not thermalizing within the whole displayed time window of $5000T$ (high-lying blue lines). (b) Thermalization times t_{th} of random density waves, postselected on initial states that exhibit a drop in $\overline{\langle n_b \rangle^2}$ of at least 75% during the displayed time window. Here, t_{th} is determined as the median across all runs (as opposed to the mean) due to the presence of a long tail in the distribution of t_{th} , the statistics of which are shown for $L = 80$ in the inset. Note that t_{th} determined in this way is seen to scale exponentially or faster with n .

over our chosen simulation time window [Fig. 7(a)]. Nevertheless, we still observe the average thermalization timescale to scale exponentially with n [Fig. 7(b)]. The large sample-to-sample fluctuations of t_{th} make it difficult to reliably extract the exact scaling behavior, but the above reasoning suggests that t_{th} continues to scale as 2^{2n} for typical initial states.

V. FRAGILE FRAGMENTATION AND THE SPACE COMPLEXITY OF THE WORD PROBLEM

Our discussion in the past few sections has focused on the way in which the time complexity of the word problem enters in the thermalization times of Dyn_G . In this section, we turn to the space complexity of the word problem. As discussed in Sec. II B, the space complexity of the word problem is determined by the maximal amount of space required to map between two words $w, w' \in K_{g,L}$, as diagnosed by the expansion-length function $\text{EL}_g(L)$ in Eq. (23). When the expansion length is large, transitioning from w to w' necessarily requires that w first grow to be much larger than its original size before shrinking down to w' . When $\text{EL}_g(L) > L$, the dynamics thus lacks the spatial resources needed to connect all states that describe the same group element. In this situation, Dyn_G cannot act ergodically in $K_{g,L}$, and thus $K_{g,L}$ become *further* fragmented. Each fragment now contains words that can be reached from some reference word w by derivations that do not involve intermediate words longer than L . We call this phenomenon fragile fragmentation in analogy to the notion of fragile topology in band theory [39]: Dyn_G is said to exhibit fragile fragmentation if there are pairs of words w, w' of length L such that Dyn_G on a system of length L does not connect $|w\rangle$ and $|w'\rangle$, but Dyn_G on a larger system of length $L' > L$ connects the “padded” words $|w\rangle \otimes |e^{L'-L}\rangle$ and $|w'\rangle \otimes |e^{L'-L}\rangle$ [74]. A schematic illustration of this definition is given in Fig. 8.

A simple example of this phenomenon that exists in higher dimensions is the jamming transition. In jammed systems, an ensemble of particles with hard-core repulsive interactions can exhibit a phase transition from a low-density mobile phase to a high-density jammed phase. The analog of fragmentation is the limited configuration space that particles can explore in the jammed phase. When the jammed particles are given more space, their density decreases, and when it drops below a critical value, the dynamics becomes ergodic. If the extra space is subsequently removed, ergodicity may again be broken, but the system may find itself in a previously inaccessible microstate. The models we discuss in this section are more drastic examples of this phenomenon: Unlike the examples with hard-core particle models, the models we study exhibit jamming even in one dimension, where the analog of the critical jamming density can be polynomially or exponentially small in system size.

To understand when Dyn_G exhibits fragile fragmentation, we need to compute the expansion length $\text{EL}(L)$. [75] Doing so for a general group can be rather difficult, as $\text{EL}(L)$ ’s definition involves a rather complicated minimization problem. However, if we already know the time complexity of the word problem—i.e., if we know the scaling of $\text{Dehn}(L)$ —it is possible to place a lower bound on $\text{EL}(L)$ [76]. Indeed, suppose a word $w \in K_{e,L}$ has an




FIG. 8. (a) Schematic of a jammed system, illustrated by densely packed, repulsively interacting particles. The interactions and high density prevent the system from rearranging itself with the space available to it. However, if the system size is increased for an intermediate period of time—allowing the particles to intermittently occupy a larger region of space before returning to their original volume—all different particle configurations can be reached. The amount by which the system size must be increased for ergodicity to be restored defines the expansion length EL . (b) Fragile fragmentation, which is the analogue of jamming in our dynamics. For a fixed system size, the dynamics is nonergodic, but in panel (c), (some degree of) ergodicity is restored when a reservoir of trivial ancillae $|0\rangle_A$ is appended to the system. The analogue of returning to the original system size in the jamming example is played by projecting the ancillae onto their original state $|0\rangle_A$ at the end of the time evolution.

expansion length $EL(w)$, implying that $|w|$ is at most $EL(w)$ during any homotopy from w to e^L . Then, the expansion length of any derivation $D(w \rightsquigarrow e)$ cannot exceed the total number of length- $EL(w)$ words; if it did, there would be at least one state that appears multiple times in $D(w \rightsquigarrow e)$ —which implies that such a derivation cannot be of minimal length. Since the number of length- $EL(w)$ words is $(2|S| + 1)^{EL(w)}$, we thus have $\text{Dehn}(w) \leq (2|S| + 1)^{EL(w)}$. By maximizing over all possible $w \in K_{e,L}$ and taking a logarithm, we obtain the general bound

$$EL(L) \geq \frac{\ln(\text{Dehn}(L))}{\ln(2|S| + 1)}. \quad (90)$$

This bound is interesting in that it connects the spatial and temporal complexities of the word problem. It also has the consequence that, to find examples with additional ergodicity breaking, we need only find a group with a super-exponential Dehn function, which we will do in Sec. VB. However, before doing so, we want to understand fragile fragmentation and the spatial complexity of the word problem in the simpler case of Dyn_{BS} dynamics. A general discussion of fragile fragmentation and its repercussions for thermalization will be given in Sec. VC.

A. Fragile fragmentation in BS dynamics

We saw in Sec. IV that the word $w_{\text{large}}(n)$ encloses an area of $O(2^n)$ on the BS Cayley graph, leading to a word problem with exponentially large time complexity. We now address the space complexity of the word problem for the BS group. At first glance, it may seem that the spatial complexity is also exponentially large. Indeed, the action of the naive homotopy mapping $w_{\text{large}}(n)$ to the trivial word is

to bring the two excursions that $w_{\text{large}}(n)$ makes along the b axis back “down” onto the a axis. This process would cause $w_{\text{large}}(n)$ to grow to a length of about 2^n over the course of the homotopy.

However, it turns out that $w_{\text{large}}(n)$ can be deformed in a way that does not require its size to significantly increase, via the process illustrated in Fig. 9. As shown in the figure, instead of collapsing the excursions “down,” we instead first make the loop “skinnier” by narrowing its extent along the a axis before bringing the excursions “down” after their area has become small enough.

It is then clear that the length of $w_{\text{large}}(n)$ does not grow by too much—at least, not by more than a factor linear in L —during this homotopy. In Appendix E, we prove that at large L ,

$$EL(L) \sim (1 + \alpha)L, \quad (91)$$

where $\alpha > 0$ is an $O(1)$ constant. Importantly, the fact that $\alpha > 0$ means that as L approaches $4(n + 1)$ from above, there will be an $L_* > 4(n + 1)$ at which $|w_{\text{large}}(n)| < L_*$ [so that $w_{\text{large}}(n)$ can still fit inside of the system], but $EL(w_{\text{large}}(n)) > L_*$; in this regime, the density wave defined by $w_{\text{large}}(n, L_*)$ cannot relax even at infinite times. Therefore, there is a transition at a finite density of es in the initial product state between a jammed regime (where homotopies are unable to contract) and an ergodic regime, leading to fragile fragmentation. Since the expansion length $EL(L) \propto L$, the severity of this jamming is comparable to that of conventional jammed systems in higher dimensions.

We can identify L_* numerically simply by decreasing L until $|w_{\text{large}}(n, L)|$ ceases to thermalize. The results are shown in Fig. 10. In this example, the length of the initial




FIG. 9. How a word w_{large} enclosing an exponentially large area can be deformed to the identity word without incurring an exponential amount of expansion. The deformation proceeds by making the loop defined by w_{large} narrower along the a direction before shrinking the loop in the b direction.

word (without identities) is $|w_{\text{large}}| = 4(n+1) = 44$. We observe that thermalization time diverges as L approaches $L_* \lesssim 50$ (for $L = 50$, we have observed only a single instance of thermalization at a very long time, $t_{\text{th}} \approx 5 \times 10^8$), confirming a nontrivial expansion length.

B. Exponential spatial complexity: Iterated Baumslag-Solitar group

We now present an example of group dynamics that exhibits fragile fragmentation with a zero density jamming transition: the iterated Baumslag-Solitar group [40]. This




FIG. 10. Thermalization time for stochastic BS dynamics as a function of L for a system initialized in a product state $|w_{\text{large}}(n, L)\rangle$ with $n = 10$ (so that the initial states for different L differ only by the density of es). The thermalization time diverges at some finite L , which defines the expansion length of the word $w_{\text{large}}(n)$. Each thermalization time is obtained by averaging over several independent runs with the same parameters. For $L = 50$, we observe a single thermalizing run, with thermalization time $t_{\text{th}} \approx 5 \times 10^8$ (not shown in the plot). No thermalization is observed for $L < 50$. Inset: expansion length of $w_{\text{large}}(n)$ for different n . For the BS group, EL is a linear function of n (and, therefore, of L).

group is (loosely speaking) constructed by embedding a BS group inside of itself. We refer to it as $\text{BS}^{(2)}$ and define it via the presentation

$$\text{BS}^{(2)} = \langle a, b, c \mid ab = ba^2, bc = cb^2 \rangle. \quad (92)$$

Thus, $\text{Dyn}_{\text{BS}^{(2)}}$ dynamics are most naturally realized in spin-3 chains with 3-local dynamics, whose local Hilbert space is obtained from that of the BS model by appending the two states $\{|c\rangle, |c^{-1}\rangle\}$. Note that like BS, $\text{BS}^{(2)}$ has a single conserved $U(1)$ charge given by the density of c generators, $n_c \triangleq \sum_i n_{c,i} = \sum_i (|c\rangle\langle c|_i - |c^{-1}\rangle\langle c^{-1}|_i)$.

Several facts about $\text{BS}^{(2)}$ (and related generalizations thereof) are proven in Appendix F. The most important result is that the Dehn function of $\text{BS}^{(2)}$ is a super-exponential function of L :

$$\text{Dehn}(L) \sim 2^{2^L}. \quad (93)$$

This finding can be intuited from the fact that the word $v(n) \equiv (c^{-n}b^{-1}c^n)a(c^{-n}bc^n)$ is equivalent to a double exponentially long string of as:

$$\begin{aligned} v(n) &= (c^{-n}b^{-1}c^n)a(c^{-n}bc^n) \\ &= b^{-2^n}ab^{2^n} \\ &= a^{2^{2^n}}, \end{aligned} \quad (94)$$

with the length of the rhs being double exponentially larger than that of $|v(n)|$ (and where “=” in the above denotes equality as elements of $\text{BS}^{(2)}$). By following the same strategy as in the construction of $w_{\text{large}}(n)$, we can construct a word $w_{\text{huge}}(n) \in K_{e,L}$ whose area grows double exponentially with its length, namely,

$$w_{\text{huge}}(n) = a^{-1}v(n)^{-1}av(n). \quad (95)$$

Like with BS, the slow dynamics of BS⁽²⁾ is manifest in the relaxation of the conserved charge n_c , which from the scaling of $\text{Dehn}(L)$, we expect to relax with an effective momentum-dependent diffusion constant that is double exponentially suppressed with q .

The general bound (90) implies that $\text{EL}(L) \gtrsim 2^L$. In Appendix F, we prove that this bound is in fact tight, so

$$\text{EL}(L) \sim 2^L. \quad (96)$$

Thus, unlike BS, there is no way to contract $w_{\text{huge}}(n)$ to the identity without it taking up an exponentially larger amount of space. [77] Therefore, the n_c density-wave pattern present in the state $|w_{\text{huge}}(n)\rangle \otimes |e^m\rangle$ will remain present even at *infinite* times unless $m \gtrsim 2^n$, i.e., unless the n_c density wave is exponentially dilute.

We now demonstrate this phenomenon numerically, using an extension of the analysis presented for BS. A direct implementation of the bistochastic circuit (86) is numerically rather expensive when investigating how n_c relaxes due to the requirement of needing simulations to be run for times double exponential in the characteristic scale of the n_c fluctuations under study. For this reason, we instead consider an *irreversible* modification of Eq. (86). We modify the dynamics so that $c c^{-1}$ and $c^{-1} c$ pairs can be annihilated but not created. Thus, our elementary stochastic gates M_i contain terms like $|e, e, e\rangle\langle c, c^{-1}, e|_{i,i+1,i+2}$ but *not* the transpose thereof, with the quantity

$$n_{|c|} \equiv \sum_i (|c\rangle\langle c|_i + |c^{-1}\rangle\langle c^{-1}|_i) \quad (97)$$

decreasing monotonically with time.

The merit of taking this approach is that the irreversible setting allows us to extract lower bounds on the relaxation time of n_c for the reversible setting. Our simulations are run by initializing the system in $|w_{\text{huge}}(n, L)\rangle$, a version of $w_{\text{huge}}(n)$ padded with $L - |w_{\text{huge}}(n)|$ e characters at random locations, so that $|w_{\text{huge}}(n, L)| = L$, and then tracking the time evolution of $n_{|c|}$. The results are shown in Fig. 11 for different values of n , which are necessarily very small on account of the double exponential growth of the Dehn function. In the main panel, we show the relaxation time of $|w_{\text{huge}}(3, L)\rangle$ as a function of L , defined by the time at which no c, c^{-1} characters remain in the evolved word. The inset shows $\text{EL}(n)$, defined as the minimal value of L for which relaxation (namely, reaching a state containing no c, c^{-1} characters) was observed to occur over 1000 runs of the dynamics. The extracted $\text{EL}(n)$ roughly conforms to our expectation of $\text{EL}(n) \sim 2^{\alpha n}$ for an $O(1)$ constant α ; however, the long timescales required to observe thermalization mean that statistical errors are rather large, and with our current data, we should not expect to obtain a perfectly exponential scaling.




FIG. 11. Thermalization time as a function of L for irreducible BS⁽²⁾ dynamics, where $c c^{-1}$ and $c^{-1} c$ pairs can be annihilated but not created. The dynamics is initialized in the state $|w_{\text{huge}}(n, L)\rangle$ with $n = 3$, and the system is considered to have thermalized when no c, c^{-1} characters remain. Inset: expansion length of $w_{\text{huge}}(n)$ for different n , defined as the minimal system size for which thermalization was observed to occur over 1000 runs of the dynamics. An approximately exponential dependence on n is observed.

C. Fragile fragmentation: Generalities

The BS⁽²⁾ example has a conserved density, so its failure to thermalize manifests itself as the freezing of a conserved density. In general, however, models exhibiting fragile fragmentation need not have conserved densities. Defining fragile fragmentation and finding reliable diagnostics for it in the general case are nontrivial tasks, which we address below. First, we provide a more precise definition of fragile fragmentation, to distinguish it from what we call intrinsic fragmentation. Second, we present a physical “decoupling” algorithm for detecting whether a system exhibits fragile fragmentation, given access to a large enough reservoir of ancillae. Third, we comment on the ways in which fragile fragmentation manifests itself in the dynamics of a thermalizing system.

1. Defining fragile fragmentation

We begin by precisely defining the notions of intrinsic and fragile fragmentation in a general context (i.e., without reference to the word problem). For concreteness, we specialize to quantum systems evolving under unitary dynamics, specified by a sequence of evolution operators $U_i(\mathcal{A})$, $i \in \mathbb{Z}_+$, acting on the system plus a collection of ancillae $\mathcal{A}$, with each ancilla assumed, for simplicity, to have the same on-site Hilbert space as the system itself. We assume that U_i forms a uniform family of time evolution operators that can be defined for any number of ancillae. Given any initial state $|\psi\rangle$ of the system, and a fixed reference state $|0\rangle_{\mathcal{A}}$ of the ancillae, we define an ensemble of states on the system as

$$\tau_{\mathcal{A}}(\psi) \equiv \{ \langle 0|_{\mathcal{A}}(U_i|\psi\rangle \otimes |0\rangle_{\mathcal{A}}) \}_{i=1}^{\infty}. \quad (98)$$

In other words, $\tau_{\mathcal{A}}$ is the ensemble of pure states one obtains by evolving the initial state $|\psi\rangle \otimes |0\rangle_{\mathcal{A}}$ for an arbitrary time and postselecting on the ancillae being in the final state $|0\rangle_{\mathcal{A}}$. Note that, in principle, we could make $\tau_{\mathcal{A}}$ depend explicitly on the state of the ancillae. However, for group dynamics, the most natural choice for this state is $|0\rangle_{\mathcal{A}} = |e\rangle_{\mathcal{A}}$, and for semigroups, an analogous state can be defined by augmenting the local Hilbert space with a character $|e\rangle$ that commutes with all characters. We therefore content ourselves with studying fragmentation for this particular choice of ancilla state. We define the Krylov sector of $|\psi\rangle$ extended to $\mathcal{A}$ as

$$K_{\psi}(\mathcal{A}) \equiv \text{span}(\tau_{\mathcal{A}}(\psi)) \subset \mathcal{H}_{\text{sys}}, \quad (99)$$

where $\mathcal{H}_{\text{sys}}$ is the Hilbert space of the system (without ancillae). We furthermore define the intrinsic Krylov sector associated with $|\psi\rangle$ as the limit

$$K_{\text{in},\psi} = \lim_{|\mathcal{A}| \rightarrow \infty} K_{\psi}(\mathcal{A}). \quad (100)$$

Under a generic thermalizing Hamiltonian or unitary dynamics, the intrinsic Krylov sector of any $|\psi\rangle$ will be the entire Hilbert space, $K_{\text{in},\psi} = \mathcal{H}_{\text{sys}}$. Intrinsic fragmentation occurs whenever this is not true, i.e., when there exist distinct initial states that do not mix under the dynamical rules even when the system is attached to an infinitely large bath (undergoing the same dynamics as the system).

When $|\mathcal{A}|$ is finite, each intrinsic Krylov sector may further shatter into many subsectors. This phenomenon (for $|\mathcal{A}|$ larger than the system) is what we have referred to above as fragile fragmentation. The expansion length associated with the dynamics is the minimal size of $\mathcal{A}$ below which additional subsectors form.

2. Probing fragile fragmentation

Our definition of fragile fragmentation above makes reference to postselection on the final state of the ancillae being $|0\rangle_{\mathcal{A}}$. The probability of postselection succeeding is clearly exponentially small in $|\mathcal{A}|$. We now present a more efficient algorithm for (i) identifying whether a given system exhibits fragile fragmentation and (ii) constructing the subspace $K_{\psi}(\mathcal{A})$ associated with an initial state $|\psi\rangle$ given a maximum expansion length $L + |\mathcal{A}|$. This procedure is more efficient than naively postselecting on the state of the ancillae in various regimes, which we discuss below.

The general algorithm proceeds as follows. We start with the state $|\psi\rangle \otimes |0\rangle_{\mathcal{A}}$ and evolve it under Dyn acting on the system plus ancillae for some time t_{th} . After time t_{th} , we repeat the following steps many times:

- (1) Measure the last site of the system plus ancillae in the computational basis.
- (2) If the outcome is e , decouple this site from the rest of the system. Otherwise, leave the site coupled.
- (3) Run the dynamics for a time t_{retherm} on the system plus remaining ancillae, and go to step 1.

The iteration stops when all ancilla sites have been decoupled: We know this is always possible since the initial state was originally decoupled from the ancillae. On physical grounds, we expect that the probability of getting outcome 0 in step 1 is $O(1)$ at all times, but for our purposes, it suffices for it to scale as $1/\text{poly}(L + |\mathcal{A}|)$.

We first discuss how this algorithm can be used to construct the subspace $K_{\psi}(\mathcal{A})$. To accomplish this goal, we set t_{th} to be the maximum possible thermalization time for the system plus ancillae, i.e., $t_{\text{th}} \sim \exp(L + |\mathcal{A}|)$. Any fragmentation that persists after t_{th} will persist to infinite time for the given spatial resources. One can take the rethermalization time after a measurement, t_{retherm} , to be much shorter (i.e., as a low-order polynomial in $L + |\mathcal{A}|$), as the measurement is a single-site perturbation to the equilibrated state, and it is not expected to take more than polynomial time to have a significant amplitude to yield $|e\rangle$. When the procedure terminates, it yields a state $|\psi'\rangle$ that (by hypothesis) is in $K_{\psi}(\mathcal{A})$. After many runs, the ensemble of generated states spans $K_{\psi}(\mathcal{A})$.

The procedure we described avoids the exponential overhead of postselection but still incurs the exponential overhead of mixing. If we want to reconstruct a state with overlap on all states in $K_{\text{in},\psi}$, this overhead cannot be avoided. Suppose, however, that we are not interested in full reconstruction of $K_{\text{in},\psi}$ but just in the simpler task of showing that adding ancillae and removing them (as above) partially lifts the fragmentation of the original system. More generally, suppose that we have constructed $K_{\psi}(\mathcal{A}_0)$ and want to know if enlarging $\mathcal{A}_0$ to $\mathcal{A}_1$ enlarges the sector. We can run the initial equilibration step to a much shorter time than the full equilibration time. We stop when we have compressed back down to $\mathcal{A}_0$, and we check if the resulting state is in $K_{\psi}(\mathcal{A}_0)$. [78] Finding a single state that lies outside $K_{\psi}(\mathcal{A}_0)$ suffices to establish fragile fragmentation. Thus, detecting fragile fragmentation can be accomplished without requiring a large bath or the system to fully thermalize.

Can this procedure be used to study other properties of the fragmentation? One additional quantity that can be computed using this method is the geodesic length of group elements. Recall that the geodesic length $|g|$ of an element $g \in G$ is the length of the shortest word that represents g , namely, $|g| = \min\{|w| : \varphi(|w\rangle) = g\}$. To compute this length, we repeat the sequential length-reducing procedure until the system freezes. Formally speaking, the system will freeze when $\langle e|\rho_{\text{end}}|e\rangle = 0$ for quantum dynamics (where ρ_{end} is the reduced density matrix of the spin at the end of the system) or $p(e_{\text{end}}) = 0$ for classical dynamics [where

$p(\cdot)$ denotes a marginal distribution for the last site of the chain]. When freezing occurs, the system size reaches the minimum length needed to support a word in the Krylov sector, therefore yielding the geodesic length of the word.

3. Fragile fragmentation and thermalization

We now discuss the unexpectedly subtle consequences of fragile fragmentation for the evolution of generic states under unitary dynamics that need not be time independent or have any local conserved densities. To keep our discussion concrete, we focus on Dyn_G dynamics for some group G , although the diagnostics we arrive at are much more general. As we will see, when G is a group (rather than just a semigroup), fragile fragmentation is particularly hard to detect locally; when G is instead a semigroup, simpler diagnostics exist (see Sec. VI for an example).

In a system exhibiting fragile fragmentation, a random word w of length L will contain many substrings s that have expansion length greater than L ; in particular, treating a specific substring s as subsystem A and $B = A^c$ as a bath of es, $|s\rangle_A \otimes |e\rangle_B$ exhibits fragile fragmentation so long as $|B| \ll \text{EL}(|s|)$. In reality, the substring is nested in the system as $|w\rangle = |w_L s w_R\rangle$, but we still claim that the action of Dyn_G can never map $|w\rangle$ to $|w'\rangle = |w_L s' w_R\rangle$, where s' is a word with expansion length asymptotically greater than L and in a different fragment to s which satisfies $\varphi(s) = \varphi(s')$. In particular, one might worry that the presence of the environment words $w_{L/R}$ can “catalyze” transitions of s , thereby sending $|w\rangle$ to $|w'\rangle$ despite s and s' living in different fragile sectors.

A simple argument shows that such catalysis cannot parametrically change the expansion length of a word. Indeed, suppose that, by contradiction, catalysis can occur. We can append w_L^{-1} to the left and w_R^{-1} to the right, increasing the length of the system by less than L . Then, the sequence

$$\begin{aligned} |e^{2|w_L|} s e^{2|w_R|}\rangle &\leftrightarrow |w_L^{-1} w_L s w_R w_R^{-1}\rangle \\ &\leftrightarrow |w_L^{-1} w_L s' w_R w_R^{-1}\rangle \\ &\leftrightarrow |e^{2|w_L|} s' e^{2|w_R|}\rangle \end{aligned} \quad (101)$$

is allowed by Dyn_G . Therefore, the space complexity of the transition $s \leftrightarrow s'$ is at most $2L$. For groups with asymptotically superlinear expansion lengths, this finding is a contradiction, and thus such a catalysis cannot occur.

To summarize, a random word contains large substrings that are frozen in some sense: If the initial state can be written as $|w_L s w_R\rangle$, time evolution under Dyn_G will never produce $|w_L s' w_R\rangle$. An immediate consequence of this fact is that the time-evolved reduced density matrix for a region A has $\langle s | \rho_A(t) | s' \rangle = 0$ at all times if s, s' are in distinct fragile fragments. Furthermore, since both s and s' can be locally generated, through the transition

$$|e^{2|s|}\rangle \leftrightarrow |s s^{-1}\rangle, \quad (102)$$

which only requires $2|s| \ll L$ of space, we expect, in general, $\langle s | \rho_A | s \rangle, \langle s' | \rho_A | s' \rangle$ to both be nonzero.

In conventional systems that exhibit the jamming transition, one can easily diagnose the jammed phase by computing local autocorrelation functions. However, in contrast, fragile fragmentation is hard to detect in this way because the frozen substrings can slide around in the system and locally change their configuration (while remaining in the same fragile sector). Thus, while one can write down a conserved quantity describing the frozen strings, such a quantity will generically be very nonlocal. Nevertheless, the observation that the dynamics does not connect pairs of states like $s w_R$ and $s' w_R$ regardless of w_R still allows one to construct a reasonable dynamical probe of fragile fragmentation. For any two words w, w' in the same Krylov sector, consider the two-point correlation function

$$\begin{aligned} C_{ww'}(t) &\equiv \text{Tr}(W X_{ww'}(t)) \\ &= 2 \cdot \text{Re} \sum_{\alpha, \beta} \langle w \alpha | \text{Dyn}_G^\dagger(t) | w \beta \rangle \langle w' \beta | \text{Dyn}_G(t) | w \alpha \rangle, \end{aligned} \quad (103)$$

where $W = |w\rangle\langle w|$ and $X_{ww'} = |w\rangle\langle w'| + \text{H.c.}$ are operators that can be fully supported in a subsystem of size $\max(|w|, |w'|)$. Suppose that $\text{EL}(w, w')$ is large so that all derivations between w and w' require large spatial resources. Then, if Dyn_G describes circuit dynamics, $C_{ww'}(t)$ is zero for small t since if both $\langle w \alpha | \text{Dyn}_G^\dagger(t) | w \beta \rangle$ and $\langle w' \beta | \text{Dyn}_G(t) | w \alpha \rangle$ are nonzero, one can transition from $|w \beta\rangle$ to $|w' \beta\rangle$ in time $2t$. Thus, $\text{Dehn}(w \beta, w' \beta)$ is expected to control when this quantity is nonzero. [79] Furthermore, for systems where the spacetime bound is saturated (like iterated BS) or close to being saturated, this timescale is dictated by $\text{EL}(w, w')$ and can be, at most, around $\exp(\text{EL}(w, w'))$.

Beyond this timescale, the system is able to undergo a large-scale rearrangement that connects w and w' (i.e., w and w' will no longer appear to live in disconnected fragile sectors). By measuring the onset timescale for $C_{ww'}(t)$ to become nonzero as a function of $|w|, |w'|$, one can diagnose the existence of and place bounds on large expansion lengths. Alternatively, computing $C_{w,w'}(t)$ can be thought of in the following way: Prepare an initial state ρ_0 where a subregion R is in the pure state $|w\rangle$, time evolve this state to get ρ , and measure the expectation value of $X_{w,w'}$ in the reduced density matrix ρ_R .

The above prescription suggests a heuristic way to determine the expansion length as follows (we leave various technical details of this proposal to future work). Define Φ_R to be a dephasing channel acting on region R of

the system (decomposing the system to the form ABR for convenience):

$$\Phi_R[\rho] \triangleq \sum_{\alpha_{A,B}, w_R, \beta_{A,B}} \rho_{\alpha_{A,B} w_R \alpha_B, \beta_{A,B} w_R \beta_B} |\alpha_{A,B} w_R \alpha_B\rangle \langle \beta_{A,B} w_R \beta_B|. \quad (104)$$

Starting in the state ρ_0 described above, one can alternate between time evolving under Dyn_G and applying Φ_R before measuring the expectation value of $X_{ww'}$. Call A the subsystem in which the state $|w\rangle$ is initially present. If $\text{dist}(A, R) \gg \text{EL}(w, w')$, then repeatedly dephasing the system should not change the value of $C_{ww'}(t)$ by much. However, if the dephasing channel is applied within a distance of $\text{EL}(w, w')$ from A , then we would expect a further suppression of $C_{ww'}(t)$, given that the dephasing eliminates many trajectories mapping w to w' . Finding the location of R where one crosses over between these two behaviors would thus provide an estimate of $\text{EL}(w, w')$.

The protocol discussed above is general but somewhat indirect. As we saw in Sec. VB, in specific examples, fragile fragmentation can have more direct and dramatic manifestations. In the next section, we show that when the group property is violated, fragile fragmentation generally has easier-to-detect physical consequences: In these cases, there can sometimes be a transition where, at small subsystem sizes, reduced density matrices are generically full rank, while above a threshold size, reduced density matrices have nontrivial kernels.

VI. SEMIGROUP EXAMPLES

In the explicit examples of Dyn_G dynamics studied above, G has been taken to have the structure of a group. However, the phenomena discussed so far are not limited to models with a group structure; indeed, all of the general results obtained in Secs. II and III are valid for any finitely presented semigroup. For semigroups, however, the geometric perspective adopted above in the discussion of Dyn_{BS} is less useful. [80] In this section, we introduce two new semigroup models that do not admit a group structure but that nevertheless have word problems exhibiting large time and space complexity, which we establish using combinatorial rather than geometric arguments. The first example has large time and small space complexity and shares similarity with the BS model. The second example has both large time and large space complexity, is qualitatively unique to non-group-based dynamics, and leads to a more direct characterization of fragile fragmentation than that provided by the general criterion of Sec. VC 3.

These models are inspired by the Motzkin spin chain and, more broadly, Motzkin dynamics (see Refs. [59, 81]). For the readers' convenience, we briefly summarize Motzkin dynamics. The local state space includes an identity character $|0\rangle$ (which plays the role of $|e\rangle$ in our group-based models) as well as left and right parentheses $|(\rangle, |)\rangle$. The dynamics is engineered so that the “nestedness” of the

parentheses remains preserved, where nestedness is defined by the number of left parentheses located to the left of matching right parentheses; for example, under the dynamics, the word “()” may evolve to “()()”, “((()”, or “0,” but not to “)(,” “((,” or “))”. These rules can be summarized formally by defining the Motzkin semigroup as

$$\text{Motz} = \text{semi}\langle 0, (,) | (0 = 0(,)0 = 0), () = 00 \rangle. \quad (105)$$

Alternatively, we may define a height field h_i that keeps track of the level of nestedness at site i via

$$h_i = \sum_{j < i} (|(\rangle \langle |)_j - |)\rangle \langle |)_j|. \quad (106)$$

The dynamics then preserves both h_L (the net difference between the number of left and right parentheses) and $\min_i h_i$ (which measures the extent of the nestedness).

A. Star-Motzkin model: Slow thermalization

Our first example has a local state space, which we label by $\{(\rangle, |)\rangle, 0, *\}$. As usual, all of what follows can be applied to Hamiltonian, random unitary, or classical stochastic dynamics. The purpose of the extra character $*$ is to slow down the dynamics of the parentheses, which is done by adding to the relations of Motz the relations

$$(*0 = ** (0*) =) ** \quad 0* = *0. \quad (107)$$

Thus, when a parenthesis moves past a $*$ character (in a certain direction), the $*$ character is duplicated. The combination of these sets of rules results in the dynamics that we refer to as Dyn_{*M} . One can readily see that Dyn_{*M} exhibits Hilbert space fragmentation. Indeed, if we ignore the $*$ character that was added, the Hamiltonian describes Motzkin dynamics and thus already possesses fragmentation, which cannot be described solely by a conserved parenthesis density. The sectors of these dynamics are labeled by a sequence of closed parentheses followed by a sequence of open parentheses taking the form $\rangle^m \langle^n$: This case corresponds to the total parentheses imbalance in the configuration. When the $*$ character is added, a label for the Krylov sectors becomes

$$K_{\ell, m, n} = \rangle^m *^\ell \langle^n, \quad (108)$$

where $\ell = 0, 1, \dots, O(L \cdot 2^{\max(m, n)})$. It is clear that there are exponentially more sectors with the addition of the $*$ character, and the structure of fragmentation is thus richer.

At some level, Dyn_{*M} resembles Dyn_{BS} , as the $*$ duplicates every time it is moved past a parenthesis, in a way similar to the duplication of as that occurs as they move past bs in BS. However, there are two differences. The first is that $*$ does not have a natural inverse. The second is that the underlying Motzkin dynamics does not satisfy properties of a group: If the characters (and) are identified with a generator and its inverse, then we

necessarily must also allow the rule $() \leftrightarrow ()$, which is absent in Dyn_{*M} . Nevertheless, in Appendix G, we show that the word problems for these models exhibit the same scaling of spatial and temporal complexities as in Dyn_{BS} :

- (1) Within the sector $K_{q,0,0}$, the expansion length is linear, $\text{EL}(w, w') = O(L)$.
- (2) Within the sector $K_{q,0,0}$, there exist words w, w' such that $\text{Dehn}(w, w') = O(q)$.

Since q can grow to be exponentially large in L , this last fact implies slow dynamics. In fact, one way to study this slow dynamics is to observe that atypical configurations of the $U(1)$ charge corresponding to the parentheses take a very long time to thermalize. Regarding the second result, we also provide a crisper characterization for the circumstances under which it takes a long time to transition between two words (see Appendix G).

B. Chiral star-Motzkin model: Fragile fragmentation

We now present an example where the expansion length is exponentially large, implying fragile fragmentation. In this example, we simply replace the $*$ character with a “chiral” version of the character, which we denote as $\triangleright$. The rules for these new characters are similar to those of $*$, except that the rules are only activated when a $\triangleright$ character is adjacent to). More specifically, we replace the rules in Eq. (107) as

$$0\triangleright) =)\triangleright\triangleright, \quad 0\triangleright = \triangleright 0. \quad (109)$$

Note that one could also add another chiral character $\triangleleft$ that only interacts with) and commutes with $\triangleright$, but the necessary physics is already illustrated for $\triangleright$.

We first discuss the intrinsic Krylov sectors of the dynamics. In particular, assuming the system is a large bath of 0s, then one can show that any configuration can be reduced to the canonical form:

$$R_{\vec{k}, \ell, m, n}^{\triangleright} =)^m \triangleright^\ell (\triangleright^{k_1} (\triangleright^{k_2} (\triangleright^{k_3} \dots (\triangleright^{k_n})). \quad (110)$$

Note that we can have a large number of $\triangleright$ characters locked between adjacent ‘(’ because $\triangleright$ characters cannot tunnel past ‘(’ characters. As a result, we label the Krylov sectors by the four indices $(\vec{k}, \ell, m, n)$, where $\dim \vec{k} = n$.

Chirality of the $\triangleright$ character plays a crucial role in the large space complexity. Define a *nest* as a collection of parentheses in the form $((\dots))$. In particular, $\triangleright$ characters embedded in a nest can exit the nest but cannot enter an adjacent one, due to the chirality constraint. As a result, the transition from $\triangleright^k (((\dots)))$ to $(((\dots))) \triangleright^k$ (which subsequently allows $\triangleright^k$ to enter the nest) is not possible unless $(((\dots)))$ is collapsed. Collapsing $(((\dots)))$ can, however, require exponentially large space if the nest $(((\dots)))$ contains a large number of $\triangleright$ characters. Thus, intuitively, an exponentially large bath is needed to unfreeze the system.

In Appendix G, we provide a more rigorous argument for when the expansion length connecting two words w and w'

can be exponentially large, and we also discuss an interesting consequence of the fragile fragmentation in this model, which does not have an analog in group dynamics. In particular, we argue that under unitary time evolution $\rho(t) = e^{-iHt} \rho(0) e^{iHt}$ where $\rho(0)$ is a product state, the subsystem density matrix $\rho_A(t) = \text{Tr}_{A^c}(\rho(t))$ exhibits a transition in its rank as $|A|$ is increased. When $|A| \ll \log L$, $\rho_A(t)$ is of full rank, and when $|A| \gg \log L$, $\rho_A(t)$ is no longer of full rank. Probing this property in a physically reasonable way is further discussed in Appendix G.

VII. GENERALIZATION TO TWO DIMENSIONS: GROUP LOOP MODELS

The discussion thus far has been restricted to 1D models of group dynamics. It is natural to wonder whether or not higher-dimensional models with similar behavior can be constructed, especially since the aforementioned phenomena are more prevalent in higher dimensions.

In this section, we discuss one 2D generalization of our group-based models that, in some sense, is the most faithful way of embedding the 1D group constraint in a two-dimensional system and that leads to a qualitatively new way of producing glassy dynamics and jamming in 2D. This process proceeds by fixing a group G [82] and considering loop models that possess one flavor of loop for each generator of G . Along any one-dimensional reference loop, one can associate a group element corresponding to the product of all of the generators corresponding to loops that the reference loop intersects. The dynamics is engineered so that this group element remains invariant under the dynamics.

This class of models can be viewed as a broad generalization of the construction in Refs. [83,84] (which studied dynamics) and Ref. [85] (which studied ground-state properties), and we expect similar robustness of the Hilbert space fragmentation in these models. For the sake of brevity, we discuss this construction at a high level—e.g., we largely use continuum language in order to avoid the notational burden incurred by an explicit lattice description—and defer a more comprehensive analysis to future work.

Given a discrete group G and a presentation thereof, the degrees of freedom in our 2D model are associated with directed loops labeled by generators of G . In the following, we describe how to place constraints on the dynamics of these loops to produce phenomenology similar to that present in our 1D models.

We start by observing that each directed reference path γ through space (microscopically, along the lattice) can be associated with a group word $w(\gamma)$. This word is determined by the ordered labels of the loops that γ intersects. Specifically, when proceeding along γ , each time γ intersects a loop labeled by the generator g_i , $w(\gamma)$ is multiplied by (i) g_i if the local tangent vector $\vec{r}$ of γ and the local tangent vector $\vec{\ell}$ of the loop can have a cross product $\vec{r} \times \vec{\ell}$ parallel to $\hat{z}$, and (ii) g_i^{-1} if $\vec{r} \times \vec{\ell}$ is antiparallel to $\hat{z}$.

This observation means that isolated closed loops can be regarded as implementing trivial relations in the words associated with the paths that pass through them, a fact we illustrate pictorially as follows.

$$\begin{array}{ccc}
 \gamma & \gamma & \gamma \\
 w(\gamma) = e & w(\gamma) = g_i^{-1} g_i & w(\gamma) = g_i g_i^{-1}
 \end{array} \quad (111)$$

Therefore, we associate processes nucleating a loop with a free expansion ($ee \rightarrow gg^{-1}$) and those annihilating a loop with a free reduction ($gg^{-1} \rightarrow ee$). With more loops, a more general situation might look like the following:

$$w(\gamma) = g_m^{-1} g_l g_l^{-1} g_m g_n \quad (112)$$

We now discuss how to implement relations of the group in terms of the loops. Suppose the group presentation is indicated by

$$G = \langle g_1, g_2, \dots, g_n \mid r_1, r_2, \dots, r_m \rangle, \quad (113)$$

where each of the r_i are words to be identified with the identity in G , and $|r_i| \leq 3$ for all i ; this restriction on the length of the relations follows from the fact that any group (but *not* any semigroup) exhibits a finite presentation satisfying $|r_i| \leq 3$ (see Appendix A for the proof). Writing $r_i = g_m g_n g_\ell$, this corresponds to an object that we refer to as a net:

$$r_i = \text{net}(g_m, g_n, g_\ell) \quad (114)$$

As another example, if $r_i = g_m g_n^{-1} g_\ell$ (i.e., one of the generators is replaced with its inverse), then the net looks like

$$r_i = \text{net}(g_m, g_n^{-1}, g_\ell) \quad (115)$$

and so on. Any loop configuration corresponding to one of the above nets can be created or destroyed without changing $\varphi(w(\gamma))$ [the group element associated with $w(\gamma)$] since, for any curve γ that cuts across the net, creating or destroying the net simply corresponds to applying the appropriate relation r_i at some point in the word $w(\gamma)$.

The dynamics we consider are generic dynamical processes that preserve $\varphi(w(\gamma))$ for all closed curves γ , which may be viewed as an unusual type of gauge constraint. Thus, the dynamics will include processes that nucleate and destroy nets associated with each r_i and will also include processes where lines are moved, stretched and contracted, and where intersections of lines are moved. It will also include processes that attach a loop with an intersection point of other loops, as shown below:

$$\text{net}(g_n, g_m, g_\ell) \leftrightarrow \text{net}(g_k, g_n, g_m, g_\ell) \quad (116)$$

where an analogous deformation occurs for g_k replaced with g_k^{-1} (in which case, the arrow is reversed). To avoid problems on the lattice where an unbounded number of joins can occur (requiring unphysically large degrees), we only allow for a join if the degree of the intersection point is below a certain threshold, set by $\max_i |r_i|$.

The dynamical processes described above are sufficient to simulate the full group dynamics. For example, suppose we want to determine whether the processes we wrote down suffice to simulate $g_m g_n = g_\ell^{-1}$ (assuming $r_i = g_m g_n g_l = e$ is a relation to the group). We can show that this is indeed the case by applying the following sequence of relations:

$$\text{Diagram 1} \rightarrow \text{Diagram 2} \rightarrow \text{Diagram 3} \rightarrow \text{Diagram 4} \rightarrow \text{Diagram 5} \quad (117)$$

where the first relation creates a net, the second relation corresponds to two joins, the third relation corresponds to two unjoins, and the last relation corresponds to a free reduction. A similar derivation shows that cyclic conjugates of relations (such as $g_l^{-1}r_i g_l = g_l^{-1}g_m g_n = e$) can similarly be simulated by the rules we have already discussed.

To summarize, our 2D dynamics contains the following processes:

- (1) Processes that deform loops in ways that do not create or destroy loop crossings.
- (2) Processes that nucleate and annihilate closed loops (performing free reductions and expansions).
- (3) Processes that join a free loop with an intersection of loops [86].
- (4) For each relation r_i , a process that nucleates or annihilates an r_i net.

The dynamical processes above were formulated for the case where G is a group, but it is also possible to generalize to the case when G is merely a semigroup, producing models reminiscent of the 2D generalizations of the Motzkin chain studied in Ref. [85]. Obtaining a more systematic understanding of the temporal and spatial complexity of the dynamics of these models would constitute an interesting avenue for future work.

One final observation we make is that there is a simple way to map group loop models to what we call tile models. Observe that a configuration of loops or nets splits the plane into a set of disjoint tiles (two adjacent tiles are joined by an edge that corresponds to a generator of the group). Then, it is possible to label tiles with an element $g \in G$ such that if two tiles have labels g and h and share an edge corresponding to generator k , then $gk = h$. For example, if $G = \mathbb{Z}$, each tile will be labeled by an integer, and two adjacent tiles have labels that differ by 1 (which is the generator of $\mathbb{Z}$). This case corresponds to a mapping to a height model, which the tile model is a generalization of in the case of arbitrary G . This mapping may play a role in understanding the nature of fragmentation in these models, which we leave to future work.

VIII. DISCUSSION

In this work, we have constructed a number of natural dynamical systems—with local few-body interactions—in which relaxation places anomalously expensive demands on a system's temporal and/or spatial resources. When the models have local conserved densities, the hydrodynamics of these densities is anomalous or frozen; even when conserved densities are absent, we have presented diagnostics for nonergodic behavior.

Our examples were all constructed in the context of models with intrinsic Hilbert space fragmentation. A natural question is whether the intrinsic fragmentation is essential to their physics. In our framework, dynamics without fragmentation is generated by finite presentations of the trivial group, which cannot have a superlinear Dehn

function (see Appendix B). Of course, this does not mean that the dynamics of models without fragmentation cannot be slow, but it does mean that any mechanism for slow thermalization must originate from something other than the Dehn function.

Our results lend themselves to several natural extensions. Most naturally, the anomalous hydrodynamic relaxation we saw in the BS model can be extended to other groups with presentations that manifest a conservation law. Whether these groups give rise to new classes of hydrodynamic relaxation is an interesting question for future work (a family of such examples will be presented in Ref. [87]). Another interesting direction is to investigate the ground states of Hamiltonians that implement group dynamics. A natural class of frustration-free Hamiltonians can be read off from the transfer matrices of bistochastic Markov processes [88]; their ground states are equal-weight superpositions of all the configurations in a sector, and their spectral gaps can be bounded by the Markov-chain gap. The tools developed here may be useful for undertaking a more detailed study of properties of these states.

The family of models we considered is restricted in the sense that the dynamical constraints can be expressed in the computational basis, so every computational-basis product state is in a definite dynamical sector of Hilbert space. More generally, one can consider constraints associated with a commuting set of projectors with entangled eigenstates. In one-dimensional spin chains, such commuting projectors can be deformed into unentangled projectors by a short-depth unitary circuit. However, this process just corresponds to conjugating the Hamiltonians or unitaries we have explored with short-depth unitary circuits, and it yields nothing qualitatively new. Obtaining something new in one dimension thus requires defining constraints using non-commuting projectors—which occurs in the (quantum-fragmented) Temperley-Lieb model [23,89,90]—and hence pursuing this direction requires systematic characterization of such constraints. In higher dimensions, however, there are sets of commuting projectors (like those of the toric code) whose simultaneous eigenstates are inherently long-range entangled. An interesting task for future work would be to extend our group-theoretic dynamical systems to these models and explore the resulting entanglement dynamics.

Finally, it would be interesting to explore the stability of our results with respect to weak violations of the constraints. One way to do this is by breaking the constraints in an isolated region of space or by bringing a thermal bath in contact with the system at its boundary. In the case of group dynamics, doing so destroys fragmentation, making the dynamics fully ergodic. Furthermore, it leads to all states being connected by at most $O(L^2)$ steps of the resulting dynamics since the $O(L)$ elements of a given product state's geodesic word need to be moved at most an $O(L)$ distance to the constraint-free region, at which point they can be changed into the elements of any other geodesic

word. However, it is still possible for the dynamics in this case to have long (even exponentially long in L) thermalization times due to bottlenecks in Fock space that arise from the finite spatial extent of the constraint-free region. A specific example of this case is presented in Ref. [91].

ACKNOWLEDGMENTS

S. B. and E. L. thank Soonwon Choi and Tim Riley for useful discussions. S. B. thanks Aram Harrow for helpful discussions on the cutoff phenomenon. A. K. is grateful to Claudio Chamon and Andrei Ruckenstein for related discussions. S. B. was supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. E. L. was supported by a Miller Research Fellowship. A. K. was supported in part by DOE Award No. DE-FG02-06ER46316 and the Quantum Convergence Focused Research Program, funded by the Rafik B. Hariri Institute at Boston University. S. B., S. G., and E. L. performed some of this work at the Boulder Summer School, supported by Grant No. NSF DMR-1560837. This research was also supported in part by the National Science Foundation under Grants No. NSF PHY-1748958 and No. NSF PHY-2309135, the Heising-Simons Foundation, and the Simons Foundation (216179, LB).

APPENDIX A: SEMIGROUP PRESENTATIONS

In this appendix, we review some basic notions about discrete semigroup presentations and prove a few small results mentioned in the main text.

Formally, a discrete semigroup G is a set equipped with a binary associative operation; when G is a group, it additionally has a distinguished element that acts as the identity, and each member of the set has a corresponding inverse. It is common to discuss a semigroup G in terms of a specific set of generators S and relations R between them, writing

$$G = \text{semi}\langle S|R \rangle \quad (\text{A1})$$

to signify this relationship (in the main text, when G is a group, we will simply write $G = \langle S|R \rangle$ and omit inverse generators and the identity from S , as well as trivial relations involving the identity from R). For example, one might think of the group $\mathbb{Z}$ as being defined by a single generator $S = \{x\}$ that obeys no nontrivial relations. However, this viewpoint is too narrow since it is possible for different choices of S and R to produce the same semigroup. As an example, consider the groups

$$\begin{aligned} G_1 &= \langle x \rangle, \\ G_2 &= \langle x, y | x^m = y^n, xy = yx \rangle, \end{aligned} \quad (\text{A2})$$

where m, n are relatively prime. These two groups are isomorphic, $G_1 \cong G_2 \cong \mathbb{Z}$, with the isomorphism associating an element $x^a y^b$ with the integer $am + bn$; note that

this finding is true despite the fact that G_1 and G_2 have a different number of generators and relations.

A given semigroup, in general, admits an infinite number of different presentations, but below, we will prove that the group-theoretic functions defined in the main text—the Dehn function, expansion length, and so on—have asymptotic scaling behaviors that are presentation independent.

We can exploit this fact to choose presentations satisfying some particular desired property. For example, we may be concerned with choosing a model of dynamics where the Hamiltonian or unitary gates are as local as possible. Since the locality of the dynamics is limited by the maximum size of the relations in R , we would thus like to minimize the size of the relations. To this end, we have the following proposition:

Proposition 1. Every finitely generated group has a presentation $\langle S|R \rangle$ where all relations $r_i \in R$ have length $|r_i| \leq 3$.

Proof. Consider the Cayley 2-complex CG_G of a finitely presentable group G (see Sec. IV for a brief discussion of its definition) obtained from a finite presentation P . While the exact structure of CG_G depends on P , CG_G can always be subdivided to obtain a simplicial complex where each 2-cell has three edges. Since each 2-cell in the complex corresponds to a relation, $|r_i| \leq 3$ for all $r_i \in R$ in the subdivided complex, thereby defining a presentation P' whose relations all have length less than or equal to 3. Since P is finite, this subdivision is completed after only a finite number of steps, and P' is consequently also finite. ■

Note that while the Cayley 2-complex of a semigroup can also be subdivided, the lack of translation invariance in a semigroup's Cayley complex means that the resulting subdivision may yield a presentation with infinitely many generators (an illustrative example is to compare the semigroup $\mathbb{N} \times \mathbb{N}$ with the group $\mathbb{Z} \times \mathbb{Z}$).

As mentioned above, the group-theoretic properties we are interested in from the point of view of group dynamics is largely insensitive to the choice of presentation. For example, recall the growth function $N_K(L) \triangleq |\{g, |g| \leq L\}|$ defined in Eq. (29) of the main text, which measures the number of dynamical sectors as a function of system size. The scaling of $N_K(L)$ with L is independent of the choice of presentation, allowing us to meaningfully talk about the growth function of a semigroup, rather than of a presentation:

Proposition 2. Let $N_{K,P}$ denote the growth function for a particular presentation $P = \text{semi}\langle S|R \rangle$ of G . Then,

$$N_{K,P} \sim N_{K,P'} \quad (\text{A3})$$

for all finite presentations P, P' of G .

Proof. Let $P = \text{semi}\langle a_1, \dots, a_{|S|} | R \rangle$ and $P' = \text{semi}\langle a'_1, \dots, a'_{|S'|} | R' \rangle$. Then, since both P, P' present G , each a'_i can be expressed as a product of a finite number of a_j . Let $n_{PP'}$ denote the maximal number of generators of P that appear

when writing the a'_i in terms of these generators. Let $|g|_P$ also denote the geodesic distance of $g \in G$ with respect to the presentation P . Then, $|g|_{P'} \leq n_{PP'}|g|_P$. Thus,

$$N_{K,S'}(L) \leq N_{K,S}(n_{PP'}L). \quad (\text{A4})$$

We may also perform a similar rewriting of generators of P in terms of those of P' . After running the same argument, we find that there exist $O(1)$ constants $n_{PP'}, n_{P'P}$ such that

$$N_{K,P'}(L/n_{P'P}) \leq N_{K,P}(L) \leq N_{K,P'}(n_{PP'}L), \quad (\text{A5})$$

and hence $N_{K,P} \sim N_{K,P'}$. ■

Similar reasoning can be applied to show that the Dehn function and expansion length (see Sec. II or the following appendix for definitions) of a semigroup are presentation independent:

Proposition 3. Let $\text{Dehn}_P(L)$ and $\text{EL}_P(L)$ be the Dehn function and expansion length of a semigroup with a finite presentation $P = \langle S | R \rangle$. Then, if P, P' are any two such finite presentations,

$$\begin{aligned} \text{Dehn}_P(L) &\sim \text{Dehn}_{P'}(L), \\ \text{EL}_P(L) &\sim \text{EL}_{P'}(L). \end{aligned} \quad (\text{A6})$$

APPENDIX B: A PRIMER IN GEOMETRIC GROUP THEORY

In this appendix, we state and prove some useful facts about the geometry and complexity of finitely presentable discrete groups. Many of the statements derived below are well-known results in the math literature, and we have tried to provide citations to the original works when appropriate. A particularly accessible review of background material relevant to the discussion to follow can be found in Ref. [40]; a more advanced reference is Ref. [92]. As a small notational convenience, in the following, the notation $w \sim g$ will be used as shorthand to denote that the word w evaluates to g ; in the main text, this was written as $\varphi(|w\rangle) = g$:

$$w \sim g \leftrightarrow \varphi(|w\rangle) = g. \quad (\text{B1})$$

We are mostly interested in infinite groups since finite ones have trivial large-scale geometry (in a sense soon to be made precise). Finitely presentable infinite semigroups are, of course, very easily constructed; indeed, it is easily verified that any group presentation where the number of generators exceeds the number of nontrivial relations [93] will generate an infinite group. Free groups (on $n > 1$ generators) and Abelian groups, in some sense, define opposite extremes since the free group has a Cayley graph that is embeddable in the hyperbolic space $\mathbb{H}^n$, while Abelian groups have Cayley graphs that are embeddable in $\mathbb{R}^n$. Most of the interesting cases for us correspond to

when an intermediate amount of “Abelian-ness” is introduced to the non-Abelian free group.

1. Time complexity: The Dehn function

The definition of the Dehn function (14) in the main text relates only to the (worst-case) complexity of deforming a given word $w \in K_e$ into the identity word [recall that $K_{g,L}$ is the set of length- L words that represent the element g , namely, $K_g = \{w | |w| = L, g(w) = g\}$]. In the main text, we claimed that focusing on the complexity of words in K_e —as opposed to K_g for $g \neq e$ —was done without loss of generality. We now prove that studying the complexity of the word problem in $K_{g \neq e}$ indeed does not produce anything that is not already captured by $\text{Dehn}(L)$:

Proposition 4. For a given element g of geodesic distance $|g| \leq L$, define the g -sector Dehn function as

$$\text{Dehn}_g(L) \triangleq \max_{w, w' \in K_{g,L}} \text{Dehn}(w, w'), \quad (\text{B2})$$

where $\text{Dehn}(w, w')$ is the minimum number of applications of relations needed to transform w into w' . Then,

$$\text{Dehn}_g(L) \sim \text{Dehn}_e(L) \triangleq \text{Dehn}(L) \quad (\text{B3})$$

for all g .

Proof. For two words $w_{1,2}$ in the same K_g sector, any deformation (a.k.a. based homotopy) of w_1 to w_2 gives a deformation between the length- $2L$ word $w_1 w_2^{-1} \in K_e(2L)$ and e . Thus, the minimal number of steps needed to relate w_1 to w_2 cannot be asymptotically smaller than the minimal number of steps needed to deform $w(w')^{-1}$ to e . This finding implies

$$\text{Dehn}(w(w')^{-1}, e) \lesssim \text{Dehn}(w, w'), \quad (\text{B4})$$

where $\lesssim$ denotes the equivalence of additional contributions linear in L . [94] Thus, $\text{Dehn}(L) \lesssim \text{Dehn}_g(L)$. Conversely, since we can deform e to $w_1^{-1} w_2 \sim e$ in time around $\text{Dehn}(w_1 w_2^{-1})$, w_1 can be deformed into $w_1(w_1^{-1} w_2) = w_2$ in time less than or around $\text{Dehn}(L)$. Thus, we also have $\text{Dehn}(w, w') \lesssim \text{Dehn}(w_1 w_2^{-1})$, so up to factors of order L , we have

$$\text{Dehn}(w_1 w_2^{-1}) \sim \text{Dehn}(w, w') \Rightarrow \text{Dehn}_g(L) \sim \text{Dehn}(L). \quad (\text{B5})$$

■
The above argument means that—if the amount of space available to us is not restricted—all sectors asymptotically have the same worst-case complexity. However, if we require that all words involved have length less than or equal to L , Proposition 4 is modified to read

$$\text{Dehn}_g(L) \lesssim \text{Dehn}(L). \quad (\text{B6})$$

In this case, all sectors $K_{g,L}$ for which the geodesic length [95] of g satisfies $|g|/L < 1$ in the $L \rightarrow \infty$ limit will still have $\text{Dehn}_g(L) \sim \text{Dehn}(L)$. On the other hand, sectors where $|g|/L = 1$ as $L \rightarrow \infty$ will not have enough “free space” for the above argument to work and will thus have $\text{Dehn}_g(L) < \text{Dehn}(L)$ [in the case where fragile fragmentation occurs (see Sec. V), this case remains true provided we define $\text{Dehn}(w, w') = 0$ if w_1, w_2 are in different subsectors of $K_{g,L}$]. Sectors where $|g|/L = 1$ as $L \rightarrow \infty$ are, however, necessarily exponentially smaller than those with $|g|/L < 1$ (see Proposition 10), and hence, a random state will be in a sector with $\text{Dehn}_g(L) \sim \text{Dehn}(L)$ with high probability.

We now provide examples of simple groups and their Dehn functions and discuss which sorts of groups we can expect to have large Dehn functions. A basic result is that only infinite non-Abelian groups can have interesting Dehn functions due to the following easily verified statements:

Fact 1. All finite groups have $\text{Dehn}(L) \lesssim CL$ for some (presentation-dependent) constant C related to the diameter of the group’s Cayley graph [96].

Fact 2. $\text{Dehn}(L) \lesssim L^2$ for any Abelian group [97].

Groups that are *too* non-Abelian also have quadratic Dehn functions:

Proposition 5. Free groups have $\text{Dehn}(L) \sim L^2$.

Proof. This finding follows from the fact that the Cayley graphs of free groups are trees, and thus any word $w \in K_e$ is a (potentially backtracking) path on the tree that contains no nontrivial loops. Such paths can be contracted to the trivial path by applying $O(L)$ relations $xx^{-1} = ee$ and $O(L^2)$ relations $xe = ex$ [98]. ■

To look for interesting Dehn functions, we thus need infinite groups with a moderate amount of Abelian-ness, which (roughly speaking) possess nontrivial loops at all length scales. One can start by finding examples of groups where $\text{Dehn}(L)$ scales as a higher-order polynomial. It was shown in Ref. [47] that these turn out to be virtually nilpotent groups. The simplest example is as follows:

Example 1. The discrete Heisenberg group

$$H_3 = \langle x, y, z | [x, y]z^{-1}, [x, z], [y, z] \rangle \quad (\text{B7})$$

has $\text{Dehn}(L) \sim L^3$ [47].

The group $\text{BS}(1, 2)$ [99] studied in detail in the main text is the simplest group known to us with an exponential Dehn function, whose properties we discuss in detail in Appendix E.

2. Space complexity: The expansion length

As discussed in the main text, the *space* complexity of the word problem is given by the maximal size of words that one must encounter when reducing a $w \in K_e$ to the identity, as measured by the expansion length $\text{EL}(w)$, a quantity originally introduced by Gromov in Ref. [100].

If $\text{EL}(w) > |w|$, then w must expand by a nontrivial amount while being reduced to the identity.

As we did with the area, we may use the expansion length to define a distance metric between any two words w_1, w_2 that represent the same group element. Instead of using $\text{EL}(w_1 w_2^{-1})$ (which we do not do because the homotopy relating w_1 to w_2 needs to be properly based), we define the relative expansion length $\text{EL}(w, w')$ between two words $w_1 \sim w_2$ by replacing e by w_2 in the above definition:

$$\text{EL}(w, w') \triangleq \min_{\{\Delta_{w_1 \rightarrow w_2}\}} \max_t |\Delta_{w_1 \rightarrow w_2}(t)|, \quad (\text{B8})$$

where $\Delta_{w_1 \rightarrow w_2}(t)$ are all paths in the Cayley graph with end points fixed at e and $g(w_1) = g(w_2)$. When the homotopy is trivial, i.e., when $w_1 = w_2$, the relative expansion length vanishes.

The expansion length of a *group* can be defined as the worst-case spatial complexity of words in K_e , as seen in our definition of the Dehn function:

$$\text{EL}(L) \triangleq \max_{w \in K_{e,L}} \text{EL}(w). \quad (\text{B9})$$

A comprehensive survey of geometric properties of EL can be found in Ref. [101].

Similar to the Dehn function, the asymptotic scaling of $\text{EL}(L)$ is independent of the choice of (finite) presentation. Also like the Dehn function, considering expansion in other K_g sectors does not yield anything new. Just as with Proposition 4, one can similarly show the following:

Proposition 6. For a given g of geodesic distance $|g| \leq L$, define the expansion lengths

$$\text{EL}_g(L) \triangleq \max_{w_1, w_2 \in K_{g,L}} \text{EL}(w, w'). \quad (\text{B10})$$

Then,

$$\text{EL}_g(L) \sim \text{EL}_e(L) \triangleq \text{EL}(L) \quad (\text{B11})$$

for all g .

Most groups we are familiar with have $\text{EL}(L) \lesssim L$. For example, it is easy to see that all Abelian groups have $\text{EL}(L) \lesssim CL$ for some constant C . One can also show the following:

Proposition 7. All finite groups have $\text{EL}(L) < L + C$ for some constant C .

Proof. The proof proceeds according to the same one that would be used in demonstrating the correctness of Example 1. Let $w \in K_e$ be a length- L closed loop in the Cayley graph of a finite group, and let N be the number of vertices in the Cayley graph. Then, since N is finite, we may write

$$w = \prod_j w_j, \quad w_j \sim e, |w_j| \leq N, \forall j, \quad (\text{B12})$$

since a path in the Cayley graph can only reach at most N different vertices before returning to its starting point. Let $M = \text{EL}(N)$. If we append M identity characters to the end of w , we can use them to turn any one of the w_j into e without increasing the length of the appended word. Since we can follow this procedure for all of the w_j , we thus have $\text{EL}(L) = L + M$ as claimed. ■

Furthermore, we will see later that, despite having exponential time complexity [$\text{Dehn}(L) \sim 2^L$], $\text{BS}(1, 2)$ only has linear spatial complexity [$\text{EL}(L) \sim L$]. This finding is part of a more general result that asynchronously combable groups [which $\text{BS}(1, 2)$ is an example of] have $\text{EL}(L) \lesssim L$ [102].

However, it is relatively straightforward to construct examples where $\text{EL}(L)$ grows faster than linearly. This process is done simply by finding groups with Dehn functions that scale as $\text{Dehn}(L) = \omega(2^L)$ due to the “spacetime” bound mentioned in the main text:

Proposition 8. (Ref. [76]) For a finitely presentable group generated by n_g generators,

$$\text{Dehn}(L) \lesssim (2n_g + 1)^{\text{EL}(L)}. \quad (\text{B13})$$

This formula can be rewritten as $\text{EL}(L) \gtrsim \log_{2n_g+1} \text{Dehn}(L)$, which grows superlinearly if $\text{Dehn}(L)$ grows superexponentially.

Proof. For $w \in K_e$ with expansion length $\text{EL}(w)$, the number of words that w can possibly visit as it is reduced to the identity is $|\{a_i, a_i^{-1}, e\}|^{\text{EL}(w)} = (2n_g + 1)^{\text{EL}(w)}$. Since the shortest reduction of w to e cannot visit a given word w' more than once, the number of steps in the reduction is (often very loosely) upper bounded by $(2n_g + 1)^{\text{EL}(w)}$. ■

3. Growth rates of groups

A simple measure of a group’s geometry is how fast the group grows, namely, how the number of elements within a distance L of the origin of the Cayley graph grows as L is increased. To this end, let $\mathcal{B}(L)$ denote the radius- L ball centered at the origin of the Cayley graph, namely,

$$\mathcal{B}(L) \triangleq \{g \mid |g| \leq L\}, \quad (\text{B14})$$

and define the growth function (or group volume) as $N_K(L) \triangleq |\mathcal{B}(L)|$. Physically, $N_K(L)$ places a lower bound on the number of Krylov sectors that models with Dyn_G dynamics fragment into [$N_K(L)$ equaling the number of sectors in the absence of fragile fragmentation].

A basic property of a group is the asymptotic scaling of N_K with L . It is straightforward to show that $N_K \sim O(1)$ for finite groups and $N_K \sim \text{poly}(L)$ for Abelian groups. Free groups provide the simplest examples where $N_K \sim \exp(L)$, which is the maximum possible growth rate [103].

We now ask what implications group expansion has for the time and space complexity measures introduced above.

It turns out that exponential growth is needed in order to have $\text{Dehn}(L) > \text{poly}(L)$:

Proposition 9. If a finitely presented group G possesses a superpolynomial Dehn function, then G has exponential growth.

Proof. The contrapositive of this proposition follows from Gromov’s theorem that all finitely generated groups with polynomial growth are virtually nilpotent (namely, have a nilpotent subgroup of finite index) [105]. Since virtually nilpotent groups have the same Dehn functions as nilpotent ones, we may combine Gromov’s theorem with the fact that all nilpotent groups have $\text{Dehn}(L) \sim L^d$ for some d [47] to arrive at the result. ■

Note that the converse to Proposition 9 is obviously false, as free groups on more than one generator provide examples of groups with exponential growth but with polynomial [in fact, $\text{Dehn}(L) \sim L$] Dehn functions.

We would also like to find the sizes of the different sectors K_g . One result along these lines is that the sector sizes must get small as $|g|$ gets large:

Proposition 10. Define

$$\mathcal{D} \triangleq \sum_{g \mid |g| \leq L} |K_{g,L}| = (2n_g + 1)^L \quad (\text{B15})$$

as the total number of words of length L . Then, the size of $K_{g,L}$ is upper bounded as

$$\frac{|K_{g,L}|}{\mathcal{D}} \leq C \exp\left(-c \frac{|g|^2}{L}\right) \quad (\text{B16})$$

for some g, L -independent constants C, c .

In the context of groups (like BS) with exponential growth, this finding reveals that almost all Krylov sectors are exponentially smaller than the largest ones.

Proof. The proof follows from connecting the counting of walks in $K_{g,L}$ with the heat kernel on the Cayley graph. Consider the symmetric simple lazy walk on the Cayley graph, and let $p_L(g, h)$ be the probability that a length- L walk starting at h ends at g . Then,

$$\begin{aligned} p_L(g, h) &= \frac{1}{2n_g + 1} \left(p_{L-1}(g, h) + \sum_{k \in \partial g} p(k, h) \right) \\ &= p_{L-1}(g, h) + \frac{1}{2n_g + 1} \sum_{k \in \partial g} (p(k, h) - p(g, h)), \end{aligned} \quad (\text{B17})$$

where the sum over k runs over the neighbors of g in the Cayley graph. By translation invariance of the Cayley graph, we can fix $h = e$ without loss of generality and simply write $p_L(g)$ for $p_L(g, e)$.

The previous equation can be written more succinctly as

$$\delta_L p = \Delta p, \quad (\text{B18})$$

where Δ is the normalized graph Laplacian and δ_L denotes the discrete derivative along the “time” direction determined by L . Thus, the different sizes of the $K_{g,L}$ are determined by using the heat equation to evolve a delta function concentrated on e for a total time of L . The claim we are trying to prove then follows from estimates of the discrete heat kernel Greens function developed in the graph theory literature; see, e.g., Ref. [106] for a review. ■

Various other facts follow from the observation that $p_L(x)$ obeys the heat equation. For example, it implies that $p_L(x)$ obeys strong maximum and minimum principles, which guarantees that all local maxima and minima of $p_L(x)$ occur on the boundaries of its domain of definition. It also means that since in groups of exponential growth almost all group elements in $\mathcal{B}(L)$ have geodesic distance close to L , almost all sectors $K_{g,L}$ will contain a number of elements exponentially smaller than $\mathcal{D}$.

In addition to the above asymptotic bound, we can also prove that K_e is always the largest sector:

Proposition 11. For all L and all g , $|g| \leq L$, we have [107]

$$|K_{e,L}| \geq |K_{g,L}|. \quad (\text{B19})$$

Proof. We aim to show that $p_L(e, e) \geq p_L(e, g)$ for all g , L . For simplicity of notation, let $L \in 2\mathbb{N}$. Then, using $p_L(g, h) = p_L(gk, hk)$ for all g, h, k on account of translation invariance of the Cayley graph, we have

$$\begin{aligned} p_L(e, e) &= \sum_g p_{L/2}(e, g)^2 \\ &= \sqrt{\sum_h p_{L/2}(e, h)^2} \sqrt{\sum_{h'} p_{L/2}(g, h')^2} \\ &\geq \sum_h p_{L/2}(e, h) p_{L/2}(h, g) \\ &= p_L(e, g), \end{aligned} \quad (\text{B20})$$

where, in the third line, we used the Cauchy-Schwarz inequality. ■

It is also possible to make statements about the absolute size of K_e . In particular, $|K_e|$ admits different bounds depending on the growth rate of the group:

Proposition 12. (Ref. [108]) For a group with growth function $N_K(L)$ scaling as a polynomial of L ,

$$\begin{aligned} L^{d_1} &\lesssim V(L) \lesssim L^{d_2} \\ \Rightarrow (L \log L)^{-d_2/2} &\lesssim \frac{|K_e|}{\mathcal{D}} \lesssim L^{-d_1/2}. \end{aligned} \quad (\text{B21})$$

For a group with exponential growth,

$$\frac{|K_e|}{\mathcal{D}} \lesssim e^{-L^{1/3}}. \quad (\text{B22})$$

In the main text, we saw numerically that this bound is saturated for BS (see Fig. 4). Thus, Dyn_{BS} is as close to being weakly fragmented as a group dynamics model with exponentially slow relaxation can be.

We can also connect the above bound with a previous result to derive the following:

Corollary 1. Groups with superpolynomial Dehn functions have exponentially many group sectors for words of a fixed length, and they have an identity sector K_e that contains a fraction of all length- L words that scales at most as $e^{-L^{1/3}}$.

Proof. This follows by combining Proposition 12 with Proposition 9. ■

APPENDIX C: SEMIGROUP DYNAMICS AND UNDECIDABILITY

While not relevant to the examples studied in the main text, it is amusing to note that the word problem for semigroups is not just computationally hard but can even be undecidable. For example, a foundational result in computability theory is the undecidability of the semigroup word problem, as first proven by Markov [109]. This result implies the existence of finitely presented semigroups for which $\text{Dehn}(w, w')$ grows faster than any recursive function. For such groups, one can never be sure whether or not two words $|w\rangle, |w'\rangle$ are in the same K_g , showing that establishing the number of dynamical sectors is impossible, in general.

Remarkably, rather simple examples of semigroups with undecidable word problems are known, with the required $|\mathcal{H}_{\text{loc}}|$ being as low as five. An explicit example from Ref. [110] serves as an illustration of how simple the examples can be: The semigroup in question has a five-element generating set $S = \{v, w, x, y, z\}$, and the seven relations

$$\begin{aligned} R = \{ &vx = xv, vy = yv, wx = xw, wy = yw, \\ &xz = zxv, yz = zyw, x^2v = x^2vz \}. \end{aligned} \quad (\text{C1})$$

The word problem is also undecidable when one specifies further to the setting of a finitely presented group [111,112]. Going further still, the Adyan-Rabin theorem [113,114] states that nearly all “reasonable” properties of finitely presented groups—their Dehn times, whether or not they are finite or Abelian, even whether or not they are a presentation of the trivial group—are undecidable.

As a corollary, the existence of Hilbert space fragmentation itself is therefore undecidable: Even if ergodicity seems to be broken for all system sizes below L , for some

group presentations one can, in general, never be sure that it will not be restored at system size $L + 1$. This result complements recent work on the undecidability of physical problems relating to local Hamiltonians, e.g., the determination of spectral gaps, ground-state phase diagrams, and so on [115–118].

APPENDIX D: NO COMPLETE SYMMETRY LABELS FOR NON-ABELIAN G

In this appendix, we will prove that the Krylov sectors of Dyn_G cannot be associated with the quantum numbers of any global symmetry if G is non-Abelian (as is the case for all of the examples of interest). In fact, we will prove a more general result. To state the result, we will define a locality-preserving unitary as any unitary operator $\mathcal{U}$ such that, for all local operators $\mathcal{O}$, the conjugated operator

$$\mathcal{O}^{\mathcal{U}} \triangleq \mathcal{U}^\dagger \mathcal{O} \mathcal{U} \quad (\text{D1})$$

is also local. The generators of any global symmetry are locality-preserving unitaries, but we will not need to assume that $\mathcal{U}$ commutes with Dyn_G .

We will prove the following result:

Proposition 1. When G is non-Abelian, the Krylov sectors of Dyn_G cannot be fully distinguished by the eigenvalues of any set of locality-preserving unitaries.

Proof. We argue by contradiction. Assume that there exists a set of unitaries $\mathcal{U}_a$ whose expectation values in a given computational-basis product state $|w\rangle$ allow one to determine the group element $\varphi(|w\rangle)$ associated with w (and, therefore, the $K_{g,L}$ that $|w\rangle$ belongs to). Consider the states

$$|w_{gh}\rangle = |e^l g e^m h e^n\rangle, \quad (\text{D2})$$

with l, m, n all proportional to the system size L . We can write

$$|w_{gh}\rangle = \mathcal{O}_{g,l} \mathcal{O}_{h,l+m+1} |e^{l+m+n+2}\rangle, \quad (\text{D3})$$

where $\mathcal{O}_{g,i} = |g\rangle\langle i| + \text{H.c.}$ Our goal is to show that $|w_{gh}\rangle$ cannot be an eigenstate of any of the $\{\mathcal{U}_a\}$ if $gh \neq hg$.

Denoting $|e^{l+m+n+2}\rangle$ as $|e\rangle$ for simplicity and defining $|w_{gh;a}\rangle = \mathcal{O}_{g,l} \mathcal{O}_{h,l+m+1} \mathcal{U}_a |e^{l+m+n+2}\rangle$, we have

$$\begin{aligned} \langle w_{gh;a} | \mathcal{U}_a | w_{gh} \rangle &= \langle e | \mathcal{U}_a^\dagger \mathcal{O}_{g,l+1} \mathcal{O}_{h,l+m+2} \mathcal{U}_a \mathcal{O}_{g,l+1} \mathcal{O}_{h,l+m+2} | e \rangle \\ &= \langle e | \mathcal{O}_{g,l+1}^\dagger \mathcal{O}_{h,l+m+2}^\dagger \mathcal{O}_{g,l+1} \mathcal{O}_{h,l+m+2} | e \rangle. \end{aligned} \quad (\text{D4})$$

Since we have assumed m to be extensive and $\mathcal{U}_a$ to be locality preserving, we always have $m + 1 > 2\text{supp}(\mathcal{O}_{g,i}^\dagger)$ in the thermodynamic limit, so the above expectation value splits as

$$\begin{aligned} \langle w_{gh;a} | \mathcal{U}_a | w_{gh} \rangle &= \langle e | \mathcal{O}_{g,l+1}^\dagger \mathcal{O}_{g,l+1} | e \rangle \\ &\times \langle e | \mathcal{O}_{h,l+m+2}^\dagger \mathcal{O}_{h,l+m+2} | e \rangle. \end{aligned} \quad (\text{D5})$$

In particular, in the translation-invariant case where $\langle e | \mathcal{O}_{g,l+1}^\dagger \mathcal{O}_{g,l+1} | e \rangle$ is independent of l , we have

$$\langle w_{gh;a} | \mathcal{U}_a | w_{gh} \rangle = \langle w_{hg;a} | \mathcal{U}_a | w_{hg} \rangle. \quad (\text{D6})$$

Taking g, h to be any two generators such that $gh \neq hg$ as group elements, the above then shows that states in different $K_{g,L}$ cannot be distinguished by eigenvalues of the $\{\mathcal{U}_a\}$, provided that (i) the above assumption about translation invariance can be removed and (ii) $|w_{gh}\rangle$ is not orthogonal to $|w_{gh;a}\rangle$ for all choices of l, m , and all choices of g, h such that $gh \neq hg$. We deal with condition (i) by noting that when the $\mathcal{U}_a$ are not translation invariant (as is the case for, e.g., modulated symmetries), we may exploit the fact that the sectors to which $|w_{gh}\rangle$ and $|w_{hg}\rangle$ belong are independent of l, m , and hence l, m can be varied without changing the eigenvalues of the respective states under $\{\mathcal{U}_a\}$. For condition (ii), we need only note that

$$\begin{aligned} \langle w_{gh;a} | w_{gh} \rangle &= \langle e | \mathcal{U}_a^\dagger \mathcal{O}_{g,l}^\dagger \mathcal{O}_{h,l+m+1}^\dagger | e \rangle \\ &= \langle e | \mathcal{U}_a^\dagger (|e\rangle\langle l| + |g\rangle\langle g|_l) \\ &\times (|e\rangle\langle l+m+1| + |h\rangle\langle h|_{l+m+1}) | e \rangle \\ &= \langle e | \mathcal{U}_a^\dagger | e \rangle, \end{aligned} \quad (\text{D7})$$

which is always nonzero by the unitarity of $\mathcal{U}_a$. ■

APPENDIX E: GEOMETRY AND COMPLEXITY OF BAUMSLAG-SOLITAR GROUPS

In this section, we state and prove some facts about the group geometry of the Baumslag-Solitar group $\text{BS}(1, 2)$ and some of its simple generalizations. The original work in which these groups were defined is Ref. [99]. We always work with the presentation

$$\text{BS}(1, 2) = \langle a, b | ab = baa \rangle. \quad (\text{E1})$$

Any discussion below can in fact be modified only slightly in order to work for the groups $\text{BS}(1, q) = \langle a, b | ab = ba^q \rangle$; however, for concreteness, we specify to $q = 2$ throughout, and as in the main text, we write BS for $\text{BS}(1, 2)$. A useful fact about our chosen presentation for BS is that it admits the matrix representation [40]

$$a = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad b = \begin{pmatrix} 1/2 & 0 \\ 0 & 1 \end{pmatrix}, \quad (\text{E2})$$

which allows the word problem to be solved in linear time and is helpful in numerical studies of BS 's group geometry [119].

Our notation carries over from the previous section on general aspects of group geometry. We also introduce the notation $n_b(w)$ to denote the net number of b s that appear in a word w , namely,

$$n_b(w) \triangleq \sum_{i=1}^{|w|} (\delta_{w_i, b} - \delta_{w_i, b^{-1}}). \quad (\text{E3})$$

1. Worst-case complexity

We start by analyzing the worst-case time and space complexity of words in BS by computing the Dehn and expansion-length functions.

To orient ourselves, it is helpful to recall that the Cayley graph of BS is homeomorphic to the product of the real line with a 3-regular tree [for BS(1, q), it is a $(q + 1)$ tree], with the depth of a given word w along the tree being controlled by the relative number of b s and b^{-1} s that w contains (see Fig. 3 of the main text). We refer to this tree as the b tree and use terminology whereby multiplying by b (b^{-1}) moves one “up” (“down”) on the sheet of the b tree, while multiplying by a (a^{-1}) moves one “right” (“left”) within the sheet (along the a axis). Because of the hierarchical nature of the graph, multiplying by b —namely, moving deeper into the b tree—moves one to “larger scales” while multiplying by b^{-1} does the opposite.

Because of the tree structure, it is clear that the group volume

$$N_K(L) \sim \lambda^L \quad (\text{E4})$$

grows exponentially with L for some constant $\lambda > 1$. A naive guess for the value of λ is as follows. First, we realize that the exponential growth comes from the tree structure and that motion by one node on this tree is always possible through the use of at most two group generators (since to perform an arbitrary move on the tree, one must multiply by either b , b^{-1} , or ab). Since the number of points at depths $d \leq L$ of a 3-regular tree goes like 3^L , we estimate $N_K(L) \sim 3^{L/2}$, giving $\lambda \approx \sqrt{3}$. This estimate is actually extremely close to the numerically computed scaling, as we show in Fig. 12.

a. Dehn function

We now turn to computing the Dehn function and expansion length of BS. The Dehn function must be large since it is easy to construct words with $\text{Dehn}(w) \sim 2^{|w|}$ [71]:

Proposition 13. Define $w_n \triangleq b^{-n} a b^n$ so that $w_n \sim a^{2^n}$ [120]. Then, the word

$$w_{\text{big}} = w_n a w_n^{-1} a^{-1} \quad (\text{E5})$$




FIG. 12. Number $N_K(L)$ of Krylov sectors, equal to the number of distinct BS group elements whose geodesic length is less than or equal to L .

has area

$$\text{Dehn}(w_{\text{big}}) = 2^{n+1} - 2 \sim 2^{|w_{\text{big}}|}. \quad (\text{E6})$$

A visual illustration of this statement is given in Fig. 3. The proof, which is a condensed version of the original proof in Ref. [71], is as follows:

Proof. To determine $\text{Dehn}(w_{\text{big}})$, we need to find the minimal number of relations needed to turn w_{big} into the identity word. Geometrically, this value corresponds [121] to the number of 2-cells in the Cayley complex CG_{BS} that form a minimal spanning surface $S_{w_{\text{big}}}$ with w_{big} as its boundary. Note that the loop defined by w_{big} is entirely contained within two sheets of CG_{BS} . It is furthermore clear that, if we only consider bounding surfaces $S_{w_{\text{big}}}$ contained within these two sheets, the minimal bounding surface has area

$$2 \sum_{k=0}^n 2^k = 2^{n+1} - 2. \quad (\text{E7})$$

Therefore, we need only show that this area cannot be reduced by considering surfaces that extend into other sheets. However, this case is clearly true, as a 2-cell of $S_{w_{\text{big}}}$ that lives on any other sheet will necessarily make an unwanted contribution to $\partial S_{w_{\text{big}}}$. More formally, we can recognize that, being homeomorphic to the product of $\mathbb{R}$ with a 3-tree, the Cayley 2-complex is contractible, and thus the S_{big} found above is the unique bounding surface. ■

Note that by considering a homotopy that shrinks w_{big} down to the identity by first making the loop narrower along the a axis before shrinking it along the b axis, the length of the word does not parametrically increase. This explanation is intuitive for the following fact:

Fact 3. [Ref. [102]] BS has linear expansion length $\text{EL}(L) \sim L$.

The previous proposition immediately implies the existence of exponentially many (in L) elements of $K_{e,L}$ with an exponentially large (also in L) area. The Dehn function thus grows at least as $\text{Dehn}(L) \gtrsim 2^L$. We now show that this bound is in fact tight. The existing proof of this fact in the mathematical literature appears to be to note that BS is an “asynchronously automatic” group [122] and that any such group has $\text{Dehn}(L) \lesssim 2^L$ [123]. In the following, we give a more elementary proof:

Proposition 14. BS has an exponential Dehn function:

$$\text{Dehn}(L) \sim 2^L. \quad (\text{E8})$$

Proof. The construction above tells us that $\text{Dehn}(L) \gtrsim 2^L$. A matching upper bound can be proven using an algorithm that converts an input word to a particular standard form.

We prove this case by simply moving all occurrences of b in w to the left and all occurrences of b^{-1} to the right, duplicating a and a^{-1} s along the way as needed and (optionally, for our present purposes) eliminating bb^{-1} pairs as they are encountered.

It is clear that at most an exponential in L number of additional a s are generated during this process of shuffling the b s and b^{-1} s around (as usual, this means exponential up to a polynomial factor, here linear in L), so the resulting word is

$$w' = b^k w_a b^{-l}, \quad (\text{E9})$$

where w_a is a word containing only e, a, a^{-1} and with length $|w_a| \lesssim 2^L$. Since $w \sim e$, we know that $k = l$ and that $w_a \sim e$. Thus, an additional 2^L relations suffice to reduce w_a to e and hence w' to e ; thus, 2^L is also an upper bound on $\text{Dehn}(L)$. ■

In Fact 12, we saw that BS has linear expansion length, meaning that there exists constants C, D such that $\text{EL}(w) \leq CL + D$ for all $w \in K_{e,L}$. The argument in Ref. [102], however, does not tell us whether $C = 1$ or $C > 1$, which is needed for determining whether or not BS exhibits fragile fragmentation. The following proposition answers this question in the affirmative:

Proposition 15. There exist constants α, D with $\alpha > 0$ such that

$$\text{EL}(L) \geq (1 + \alpha)L + D. \quad (\text{E10})$$

Proof. It is enough to show that $\text{EL}(w) \geq (1 + \alpha)L + D$ for a particular length- L word w , which we choose to be w_{big} in Eq. (E5) with n such that $L = 4(n + 1)$. By the contractibility of the Cayley 2-complex CG_{BS} , for all $0 \leq m \leq 2^n$, any homotopy from w_{big} to the identity must pass through the vertex a^m of CG_{BS} . Therefore,

$$\text{EL}(w_{\text{big}}) \geq \max_{0 \leq m \leq 2^n} 2|a^m|, \quad (\text{E11})$$

where $|a^m|$ is the geodesic distance (not word length) of a^m .

We now argue that there exists an $2^{n-1} < m < 2^n$, an $\alpha > 0$, and a constant c_1 such that $|a^m| > (2 + 2\alpha)n + c_1$. By the contractibility of CG_{BS} , any geodesic of a^m will be contained within a single sheet. Furthermore, since m is exponentially large in n in the worst case, the geodesic should reach a height of at least $n - c_2$ on the sheet, where c_2 is a sufficiently large $O(1)$ constant. If the geodesic does not reach such a height, then it must make a much larger number of steps in the a direction resulting in a larger perimeter. Letting $n_{|b|}$ ($n_{|a|}$) be the number of b, b^{-1} (a, a^{-1}) characters that appear in the geodesic, we see that $n_{|b|} \geq 2(n - c_2)$, so $\alpha \geq 0$.

To compute the perimeter, we have to determine the number of steps the smallest length word makes in the a direction in order to reach a^m . Since at its highest point $n_{|b|}/2$ steps down along the b direction are needed, we can intersperse any of these $n_{|b|}/2 > n - c_2$ steps with at least one step along the a direction. Note that if two consecutive steps along the a direction are taken, these can be pulled past the previous b step to form a single a step, thereby reducing the total length of the loop. Therefore, the minimal length loop has b steps interspersed with at most a single a step. Since there are exponentially many (in n) choices of m , and since each a^m has a unique geodesic $\alpha = 0.49$. ■

While this demonstrates that words in $K_{e,L}$ must expand by an amount proportional to L before being mapped to the identity word, this statement is only meaningful in the large- L limit, and in practice, $\text{EL}(L)/L$ can be very close to 1 for modest choices of L (as was seen in the numerics of Sec. IV E).

2. Average-case complexity

We now examine the average-case complexity of words in BS (or more precisely, the complexity of typical words in BS). Unfortunately, only a few results of average-case Dehn functions are known (one example for nilpotent groups is Ref. [50]), and average-case complexity has not been studied for groups with superpolynomial Dehn functions. We emphasize that our analysis is heuristic, relying on several reasonable claims backed up by numerics in some places. Letting $\overline{\text{Dehn}}(L)$ denote the Dehn function of a typical word in $K_{e,L}$, we make the following claim:

Claim 1. With probability $1 - \epsilon$, a randomly chosen word in $K_{e,L}$ for BS has a Dehn function

$$\overline{\text{Dehn}}(L) \gtrsim 2^{f(\epsilon)\sqrt{L}}, \quad (\text{E12})$$

for $f(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$.

Proving this claim—which we believe is within reach—constitutes an interesting direction for future research.

a. Distribution of geodesic lengths for random words

Note that $\overline{\text{Dehn}}(L)$ measures the typical area of words in $K_{e,L}$. The difficulty in computing $\overline{\text{Dehn}}(L)$ lies in the fact

that it is hard to sample words from $K_{e,L}$ due to the constraint that such words form closed loops in the Cayley 2-complex $\mathbf{CG}_{\text{BS}}$. A much easier task is to determine the geometry of typical length- L paths in $\mathbf{CG}_{\text{BS}}$, regardless of whether these paths form closed loops. Understanding this easier problem will allow us to build intuition for computing the scaling of $\overline{\text{Dehn}}(L)$.

The precise question we address is the following: Given a random length- L word w , what is the geodesic distance $|g(w)|$ of w ? We can efficiently obtain an answer numerically by randomly sampling words w and approximately computing their geodesic distances. In order for this procedure to be efficient, it is helpful to realize that any word may always be brought into the following canonical form:

Proposition 16. (Ref. [124]) Any word can always be mapped to a unique standard form

$$w_{knl} = b^k a^n b^{-l}, \quad k, n, l \in \mathbb{Z}, n, l \geq 0, \quad (\text{E13})$$

where n can be even only if at least one of k, l is zero [125].

Proof. This finding follows from the arguments around Eq. (E9) or the use of the matrix representation (E2). ■

Since k, n, l are unique, the w_{knl} serve as a set of canonical representatives for each g . To obtain the geodesic of an arbitrary word, we first reduce it to this form and then make use of the following:

Proposition 17. (Ref. [124]) The geodesic distance of a word w is determined by the exponents k, n, l appearing in its canonical form w_{knl} as

$$\frac{1}{2}(k + l + \log_2 |n|) \leq |g(w_{knl})| \leq 4(k + l + \log_2 |n| + 1) \quad (\text{E14})$$

when $n \neq 0$, and $|g(w_{knl})| = |k - l|$ if $n = 0$.

All that remains is to find a way of determining k, n, l given an arbitrary word w . For this, we use the matrix representation (E2), in which w_{knl} becomes

$$w_{knl} = \begin{pmatrix} 2^{l-k} & n2^{-k} \\ 0 & 1 \end{pmatrix}. \quad (\text{E15})$$

Therefore, to find k, n, l for an input string w , we proceed as follows. We first find the matrix corresponding to w by explicit matrix multiplication, yielding a result of the form $\begin{pmatrix} A & B \\ 0 & 1 \end{pmatrix}$. We see immediately that

$$\log_2 A = n_b(w) = k - l, \quad (\text{E16})$$

the number of b s contained in w minus the number of b^{-1} s. Then, we note the following:

- (i) If $B \in \mathbb{Z}$, either $k = 0$ or $l = 0$. Which one of these scenarios holds depends on $\text{sgn}(n_b)$: If $n_b < 0$, then $k = 0, l = -n_b$ and $n = B$, while if $n_b > 0$, then $l = 0, k = n_b$ and $n = 2^{n_b} B$.

- (ii) If $B \notin \mathbb{Z}$, then $k > 0$, and k is determined by the number of significant figures after the decimal point when B is represented in binary, [126] after which both l and n are determined.

Numerically implementing the above procedure for several values of L gives the histogram shown in Fig. 13, where, for simplicity, we plot an estimate of $|g(w)|$ as $k + l + \log_2(|n| + 1)$, which has the same asymptotic scaling as the true value of the geodesic. Thus, we see that the probability of obtaining $|g(w)| = 0$ (namely, $w \in K_e$) is suppressed [we have already seen that $\Pr(w \in K_e) \sim e^{-L^{1/3}}$] and that the typical geodesic distance goes as $|g(w)| \sim \sqrt{L}$.

Since we are sampling random words, $n_b(w) = k - l$ will converge to a Gaussian of width $\sqrt{L}$; thus, a typical word will reach a depth of $\sqrt{L}$ on the b -tree part of the Cayley graph. We also expect the $k + l$ contribution to the geodesic estimate (E14) to scale as $\sqrt{L}$; indeed, this expectation can be numerically verified to be the case. The fact that a random word typically has $|g(w)| \sim \sqrt{L}$ then means that the $\log |n|$ contribution to the geodesic estimate scales as $\log |n| \lesssim \sqrt{L}$. Thus, words that traverse *exponentially* far along the a axis of the Cayley graph (namely, those with $\log |n| \sim L$) are rare, meaning that we should not expect the Dehn time of typical words to be close to the worst-case result. To determine whether $\log |n| < \sqrt{L}$ or $\log |n| \sim \sqrt{L}$, we simply make a histogram of $(\log |n|)/\sqrt{L}$, with the result shown in Fig. 14. We see that the $\log |n|$ part makes an $O(1)$ contribution to the expected geodesic distance, and we learn the following:

Observation 1. A typical randomly chosen word travels to a depth of about $\sqrt{L}$ along the b tree of $\mathbf{CG}_{\text{BS}}$ and reaches a distance of about $2\sqrt{L}$ along the a axis.




FIG. 13. Histogram of the geodesic length estimate $|g(w)| = k + l + \log_2(|n| + 1)$ of randomly chosen length- L words in BS. The scaling collapse indicates that typical words have a geodesic length scaling as $\sqrt{L}$.




FIG. 14. Histogram of the distance $|n|$ that a randomly-chosen length- L word proceeds along the a axis of BS's Cayley graph. The scaling collapse indicates that typical words have $|n| \sim 2\sqrt{L}$.

Before moving on, we note as an aside that a randomly chosen word in $K_{e,L}$ is numerically observed to be essentially homogeneous as far as the conserved n_b density is concerned, meaning that $(1/|K_{e,L}|) \sum_{w \in K_{e,L}} \langle w | n_{b,i} | w \rangle \approx 0$ for all i (although, as far as we can tell, there is no symmetry that enforces this expectation value to vanish identically). This case is demonstrated in Fig. 15, and it is important for numerically determining thermalization times in the manner of Sec. IV E.

b. Dehn times of typical words

The rough intuition leading to Claim 1 is as follows. Let $L \in 2\mathbb{N}$ for simplicity, and consider a random word $w = w_L w_R$ in $K_{e,L}$, with $|w_L| = |w_R| = L/2$. Then, to the extent that $w_{L,R}$ behave like random length- $L/2$ words, the midpoint of the loop defined by w will be at a distance of about $2\sqrt{L}$ along the a axis from the origin, as follows from Observation 1. Thus, we may expect that the area of w




FIG. 15. The $\langle w | n_{b,i} | w \rangle$ averaged over 10^5 random words w of length $L = 100$ in the identity sector $K_{e,L}$ of BS.

also scales as $2\sqrt{L}$. The only exception is if w_R is very close to w_L^{-1} , which implies that w subtends little area, but we will argue that such an event is unlikely.

We begin with an argument about certain types of words that we will argue are likely to occur as subwords of typical elements in $K_{e,L}$:

Claim 2. Consider a length- L word w chosen randomly, subject to the following constraints. First, $n_b(w) = 0$. Second, the returning walk induced on $\mathbb{Z}$ by restricting to the b, b^{-1} characters of w is constrained to $\mathbb{Z}^{\geq 0}$ (i.e., it is a *Dyck walk*), meaning that the cumulative sums $n_b(x) = \sum_{j=1}^x (\delta_{w_j, b} - \delta_{[w_j], b})$ are positive for all x .

Let $S_{\text{Dyck}}(w)$ be the set of all length- L words $w' \sim w$ obeying the above constraint. For w' drawn randomly from $S_{\text{Dyck}}(w)$, w and w' satisfy $d(w, w') \sim 2\sqrt{L}$ with high probability.

Argument: The Dyck walk property means that all $w' \in S_{\text{Dyck}}(w)$ are reducible to a^{n_w} for some n_w , with n_w shared by all words in $S_{\text{Dyck}}(w)$. This case can be seen by an inspection of the Cayley graph or by recalling the canonical form (E13) [the walks corresponding to a group element are read right to left, so $b^k a^n b^{-l} \in S_{\text{Dyck}}(w)$ only if $l = k = 0$].

Since w was chosen randomly from the set of words obeying the Dyck walk condition, the b walk defined by w is exponentially likely to reach a height of $O(\sqrt{L})$, namely, to have $\max_x n_b(x) \sim \sqrt{L}$. A random w fulfilling this condition can be seen to be exponentially likely to reduce to a word a^{n_w} of length $n_w \sim 2\sqrt{L}$, as can be seen by, e.g., following the procedure that brings w into canonical form (E13).

Suppose now that n_w scales as $2\sqrt{L}$, and consider a random element $w' \in S_{\text{Dyck}}(w)$. We claim that the distance between w and w' is exponentially likely to scale as $d(w, w') \sim 2\sqrt{L}$ because w, w' are exponentially likely to travel on different sheets of the Cayley graph for nearly all of the length of their walks. More precisely, consider the projection of w, w' onto the b tree, and define the branch point $\text{Br}(w, w')$ as the largest depth of a vertex in the tree visited by both w and w' . We claim that $\text{Br}(w, w')$ is exponentially likely to be $O(1)$, so w, w' indeed travel on separate sheets for nearly the entirety of their trajectories.

To demonstrate this case more carefully, consider how returning random walks on CG_{BS} behave when projected onto the b tree. When the walk on the tree moves to larger scales of CG_{BS} , we will refer to it as moving “upwards,” and when it moves to smaller scales, we will say that it moves “downwards.” At a given vertex in CG_{BS} , moving upwards while staying on the same sheet can be done by moving directly upwards, or by moving to the left or right by an even number of steps and then moving up. Moving upwards onto a different sheet, on the other hand, is done by moving left or right by an odd number of steps before

moving up. Finally, moving downwards (on the same sheet) can be done by moving an arbitrary amount either left or right, and then moving downwards. Thus, the probability of moving downwards is equal to the probability of moving upwards, despite the fact that moving upwards can be done on either of the two sheets. More precisely, let $p^{\nwarrow}$, $p^{\nearrow}$, and $p^{\downarrow}$ be the probabilities of moving up on the same sheet, up on a different sheet, and down, respectively. Then,

$$\begin{aligned} p^{\nwarrow} &= \frac{4}{5} \frac{1}{4} \sum_{k \in \mathbb{Z}} \sum_{l=|k|}^{\infty} \binom{2l}{l+|k|} (1/4)^{2l} = \frac{1}{3}, \\ p^{\nearrow} &= \frac{1}{2} - p^{\nwarrow} = \frac{1}{6}, \\ p^{\downarrow} &= \frac{1}{2}. \end{aligned} \quad (\text{E17})$$

The fact that $p^{\downarrow} = 1/2$ means that, as far as motion in the tree is concerned, the walk will *not* move ballistically upwards or downwards, but it will instead move diffusively. Using the above transition probabilities, the expected behavior of $\text{Br}(w, w')$ can be calculated analytically using generating functions. The details are rather messy, however, and since we regard the claim of $\text{Br}(w, w')$ as being intuitively rather clear, we will be content with a numerical demonstration of this result, which we provide in Fig. 16.

Now, we return to our discussion of the distance between w and w' . As just argued, $\text{Br}(w, w')$ is exponentially likely to be $O(1)$. The contractibility of CG_{BS} means that the minimal bounding surface linking w to w' must consist of all cells bounded by w and the a axis that lie at a depth greater than $\text{Br}(w, w')$, together with the analogous set of cells for w' [a similar argument arose in the proof of Proposition 13,




FIG. 16. Probability for two returning random walks w, w' on the 3-tree—with transition probabilities as in Eq. (E17)—to have a given branch point $\text{Br}(w, w')$, confirming that $\text{Br}(w, w') = O(1)$ with constant probability.

where we effectively had $\text{Br}(w, w') = 0$]. Since each of these contributions to the bounding surface consists of around $2^{\sqrt{L}}$ cells, we indeed have $d(w, w') \sim 2^{\sqrt{L}}$. $\square$

We are now in a position to argue for the correctness of Claim 1. While Dyck walks do not constitute a constant fraction of all returning walks, [127] the basic idea is to realize that a generic word in K_e is likely to contain at least two subwords of size around $\sqrt{L}$ obeying the Dyck property, allowing an application of the above argument.

First, when sampling a random word in K_e , we can first sample uniformly from all $5^{L/2}$ words w_L of length $L/2$ and then sample from all words w_R of length $L/2$ such that $w_L w_R \sim e$. Since w_L is chosen randomly, we expect that, with unit probability in the $L \rightarrow \infty$ limit, the walk defined by w_L reaches a distance of order around $2^{\sqrt{L}}$ along the a axis of the Cayley graph (this was numerically demonstrated to be the case in Fig. 14). Since $|w_L| \sim L$, reaching this distance is only possible if w_L contains an excursion along a particular sheet of the b tree which reaches a maximal depth of about $\sqrt{L}$. Therefore, w_L must contain at least one subword $w_{L,D}$ that performs a Dyck walk of height around $\sqrt{L}$.

Given w_L , we now consider sampling words w_R such that $w_L w_R \sim e$. Since w_R must also traverse a distance of around $2^{\sqrt{L}}$ along the a axis, it must contain a subword $w_{R,D}$ that performs a Dyck walk of height around $\sqrt{L}$. By the contractibility of the Cayley 2-complex, the areas that $w_{L,D}$ and $w_{R,D}$ define with respect to the a axis can cancel out only if the excursions that $w_{L,D}$ and $w_{R,D}$ perform occur along the *same* sheets of the b tree for a fraction of their respective walks exponentially close to 1. The chance for this to occur is, however, exponentially small in the depth of the walk (since the number of such walks on the b tree grows exponentially in their length), which goes as $\sqrt{L}$. Therefore, the areas contributed by $w_{L,D}$ and $w_{R,D}$ are exponentially unlikely to cancel, and thus the area of $w = w_L w_R$ will scale as $2^{\sqrt{L}}$ with probability approaching 1 as $L \rightarrow \infty$.

We end our argument for Claim 1 by providing a degree of evidence from numerics which supports the conclusions of the above arguments. We consider the distance that randomly chosen words in $K_{e,L}$ travel from the origin. More precisely, for each randomly chosen $w \in K_{e,L}$, we subdivide w into two equal-length halves as $w = w_L w_R$ and then compute the typical amount that w_L extends along the a axis of CG_{BS} . Sampling uniformly from $K_{e,L}$ is accomplished by sampling random words and then postselecting on them being in $K_{e,L}$, which, for large L , is numerically costly on account of the postselection succeeding with probability around $e^{-\alpha L^{1/3}}$. This case limits our numerics to relatively modest system sizes, for which finite-size effects are rather strong. Nevertheless, the results in Fig. 17 show that typical words in $K_{e,L}$ have w_L , which indeed extends a




FIG. 17. Statistical properties of words in $K_{e,L}$. For w chosen randomly from $K_{e,L}$ (with L even), we write $w = w_L w_R$ with $|w_L| = |w_R| = L/2$ and then bring w_L into canonical form as $w_L = b^{k_L} a^{n_L} b^{-l_L}$. Left panel: histograms of $\log(|n_L| + 1)/\sqrt{L}$. We see that typical words in $K_{e,L}$ reach a distance of around $2\sqrt{L}$ along the a axis of the Cayley graph at their midpoints. Right panel: histograms of $n_{b,L} = k_L - l_L$, showing that typical words in $K_{e,L}$ travel a depth of around $\sqrt{L}$ into the b tree at their midpoints.

distance of around $2\sqrt{L}$ along the a axis of CG_{BS} . As argued above, this finding implies the correctness of Claim 1 (as the only way for this to be false is if w_L, w_R almost always traverse the same path of the b tree, which we have already demonstrated is exponentially unlikely to be the case). $\square$

Finally, we note that by following similar logic as in the proof of Proposition 4, the same scaling as in Claim 1 can be argued to occur for randomly chosen words and not just those in K_e . More precisely, from Claim 1, we expect that if w is a random length- L word and w' is subsequently randomly chosen from $K_{g(w)}(L)$, then $d(w, w') \sim 2\sqrt{L}$ with constant probability.

APPENDIX F: ITERATED BAUMSLAG-SOLITAR GROUPS

In this appendix, we discuss a construction for modifying the Baumslag-Solitar group $\text{BS}(n, m)$ in a way that lets us achieve Dehn functions scaling superexponentially in L . From the bound (B13), this immediately allows us to construct examples of models with nontrivial fragility lengths.

Variants of $\text{BS}(n, m)$ can be constructed whose Dehn functions grow as fast as $\text{Dehn}(L) \sim \exp^{[\log_2 L]}(1)$ [68], an extremely quickly growing function of L . However, we will content ourselves with studying the simplest variants, for which the Dehn function scales double exponentially, $\text{Dehn}(L) \sim 2^{2^L}$. A family of groups $\text{BS}^{(2)}(n, m, o, p)$ with this scaling are defined in their simplest presentations by three generators, a, b, c , where a, b satisfy the relations of $\text{BS}(n, m)$ and b, c satisfy those of $\text{BS}(o, p)$:

$$\begin{aligned} \text{BS}^{(2)}(n, m, o, p) \\ = \langle a, b, c | a^m b = b a^n, b^o c = c b^p \rangle. \end{aligned} \quad (\text{F1})$$

We focus on the case $n = o = 1, m = p = 2$, which is the simplest nontrivial example. For notational brevity, we write $\text{BS}^{(2)}(1, 2, 1, 2)$ simply as $\text{BS}^{(2)}$.

The Cayley graph of $\text{BS}^{(2)}$ consists of an infinite hierarchy of BS Cayley graphs. Multiplying by c increases the “scale” of the BS graph generated by a, b , while multiplying by c^{-1} decreases the scale. Basic facts about the geometry of $\text{CG}_{\text{BS}^{(2)}}$ such as its growth rate and the size of the identity sector are relatively difficult to address numerically, partly because, unlike BS, $\text{BS}^{(2)}$ does not admit a simple linear representation, making the evaluation $g(w)$ for arbitrary words w rather involved [128].

It is perhaps intuitively reasonable that $\text{BS}^{(2)}$ should have a superexponential Dehn function since the c characters can be used to create exponential expansion of the b characters, which in turn can create superexponential expansion of the a characters. The following result finetunes this intuition:

Proposition 18. $\text{BS}^{(2)}$ has Dehn function

$$\text{Dehn}(L) \sim 2^{2^L}. \quad (\text{F2})$$

This result appears to be well known and was stated in Ref. [40] without an explicit proof, which we provide below for completeness. The key result we need to complete the proof is known as Britton’s lemma [129], [130] which is stated as follows:

Lemma 1. Let G be a group with presentation S . Furthermore, let G contain two isomorphic subgroups $H, K \subset G$, with $\phi: H \rightarrow K$ the isomorphism between them. Define the group

$$G_\phi \equiv \langle S, t | t^{-1} H t = \phi(H) \rangle. \quad (\text{F3})$$

Now, any word w on $\{S, t\}^*$ can be written in the form

$$w = g_0 t^{\epsilon_1} g_1 t^{\epsilon_2} g_2 \cdots t^{\epsilon_{n-1}} g_{n-1} t^{\epsilon_n} g_n, \quad g_i \in G, \epsilon_i = \pm 1. \quad (\text{F4})$$

Britton’s lemma states that if $w \sim e$ represents the identity in G_ϕ , then there must be some i such that either

- (1) $n = 0$ and $g_0 = e$,
- (2) $\epsilon_i = -1, \epsilon_{i+1} = 1$, and $g_i \in H$, or
- (3) $\epsilon_i = 1, \epsilon_{i+1} = -1$, and $g_i \in K$.

Thus, we can simplify any word $w \sim e$ representing the identity in G_ϕ . Case 1 above is trivial. In case 2, we can replace the occurrence of $t^{-1} h t$ with $\phi(h)$, while in case 3, we may replace $t k t^{-1}$ with $\phi^{-1}(k)$. After this reduction, we are left with a new word w' that still represents e , and we can apply the lemma again. This process guarantees that we will always be able to apply a successive series of reductions to eliminate all t s from and $w \sim e$ in $\{S, t\}^*$

to obtain a word $w_G \in \{S\}^*$, $w_G \sim e$. The application to $\text{BS}(1,2)$ is clear, where we take $G = \mathbb{Z}$, $H = G$, and $K = 2\mathbb{Z}$.

We now return to a proof of the Dehn function scaling:

Proof. We first construct a lower bound. Define the word $w_n = c^{-n}bc^n$ so that $w_n \sim b^{2^n}$. Then, we feed this word into the construction of the large-area word w_{big} for BS , by defining $w'_n = w_n a w_n^{-1}$. Next, we claim that the word

$$w_{\text{huge}} = (w'_n)^{-1} a w'_n a^{-1} \quad (\text{F5})$$

has area

$$\text{Dehn}(w_{\text{huge}}) \sim 2^{2^n}, \quad (\text{F6})$$

which is double exponential in $|w_{\text{huge}}|$. This finding follows from an argument similar to the one we gave for the area of w_{big} in BS . The treelike structures of the sheets of the BS Cayley graph give treelike structures both for words built from b, c and those built from a, b . Letting $\tilde{w}_{\text{huge}} = b^{-2^n} a b^{2^n}$, we find

$$\text{Dehn}(w_{\text{huge}}) = \text{Dehn}(\tilde{w}_{\text{huge}} a \tilde{w}_{\text{huge}}^{-1} a^{-1}) + O(2^n). \quad (\text{F7})$$

However, using the results from our study of BS , we know that $\text{Dehn}(\tilde{w}_{\text{huge}} a \tilde{w}_{\text{huge}}^{-1} a^{-1}) \sim 2^{2^n}$. Thus, $\text{Dehn}(L) \geq 2^{2^L}$ asymptotically.

We now need to provide a matching upper bound, which we can do by combining Britton's lemma with our earlier result about BS . Note that $\text{BS}^{(2)}$ can be obtained from BS using just the type of extension that appears in Britton's lemma, where $G = \text{BS}$, $H = \langle b \rangle$, $K = \langle b^2 \rangle$. Then, we know that if we are given $w \sim e$, $|w| = L$ in $\text{BS}^{(2)}$, we can obtain a word $w' \sim e$ in BS after at most $O(L)$ applications of $c^{-1}bc = b^2$. The maximum amount that $|w|$ can grow under these substitutions is $O(2^L)$. Thus, an upper bound on $\text{Dehn}(L)$ in $\text{BS}^{(2)}$ can be obtained by an upper bound on $\text{Dehn}(2^L)$ in BS . Using our previous result on the latter, we conclude that

$$\text{Dehn}(L) \lesssim 2^{2^L}, \quad (\text{F8})$$

and thus when combined with the lower bound, we also have that $\text{Dehn}(L) \lesssim 2^{2^L}$.

Note that $\text{BS}^{(2)}$ also provides an example with a super-linear expansion length:

Corollary 2. $\text{BS}^{(2)}$ has exponential expansion length,

$$\text{EL}(L) \sim 2^L. \quad (\text{F9})$$

Proof. From the general bound (B13) and our above result about the Dehn function of $\text{BS}^{(2)}$, we know that $\text{EL}(L) \geq 2^L$. The upper bound follows from the above

application of Britton's lemma and the fact that the expansion length of BS is only $O(L)$. ■

It is easy to generalize the above example to construct groups with even-faster-growing space and time complexity:

Corollary 3. Define the group $\text{BS}^{(l)}$ through the presentation [40]

$$\text{BS}^{(l)} = \langle a_0, \dots, a_l | a_{i-1} a_i = a_i a_{i-1}^2, i = 1, \dots, n \rangle. \quad (\text{F10})$$

This group has the Dehn function and expansion length

$$\text{Dehn}(L) \sim \exp^{(l)}(L), \quad \text{EL}(L) \sim \exp^{(l-1)}(L). \quad (\text{F11})$$

APPENDIX G: DETAILED ANALYSIS OF NONGROUP EXAMPLES

In this appendix, we provide some detailed analysis studying both the time and space complexity of the non-group examples presented in the main text.

1. Star-Motzkin model

Recall the Star-Motzkin model from Sec. VI A. There are two sources of fragmentation: The first originates from the parentheses, and the second from the interaction of the parentheses with the $*$ character.

In the following analysis, we rely on the fact that we can write down a nonlocal conserved quantity under the dynamics that necessarily reflects the interaction between the Motzkin degrees of freedom with the $*$ degrees of freedom. This nonlocal conserved quantity is

$$Q = \sum_i 2 \sum_{j < i} n_{(j-n)_j} n_{*,i}, \quad (\text{G1})$$

where $n_{c,i} = |c\rangle\langle c|_i$. The interpretation of this operator can be understood pictorially. A configuration of parentheses (ignoring the $*$ character) can be mapped onto a height configuration. The height h_i is written as $h_i = \sum_{j < i} n_{(j-n)_j}$. Bringing a $*$ character at height h_i down to height $h_j < h_i$ by a sequence of local updates creates $2^{h_i-h_j}$ such $*$ characters at height h_j . The charge operator simply counts the total number of $*$ characters if *all* $*$ characters are brought down to zero height. In the situation where the height becomes negative, we can rescale $2^{\sum_{j < i} \hat{n}_{(j-\hat{n})_j}}$ to $2^{\sum_{j < i} \hat{n}_{(j-\hat{n})_j} - \bar{h}}$, where $\bar{h}$ is the value of the lowest height. The interpretation of the exponential charge is then equivalent to the previously introduced definition if $\bar{h}$ is redefined to be at zero height.

We now proceed to understanding how to label Krylov sectors of the dynamics. For simplicity, we start with analyzing the Krylov sector K_Q corresponding to the balanced sector for the parentheses, with total value Q for the nonlocal conserved quantity. We expect this model

to exhibit fragile fragmentation with a linear expansion length (much like Dyn_{BS}), and as such, we discuss connectivity of configurations assuming they are appended to a reservoir of 0 characters of length αL .

We claim that the dynamics is ergodic within the Krylov sector K_Q so long as $\alpha = O(1)$ is sufficiently large. The proof follows from finding a path between any two states in K_Q given the desired size of the reservoir of $O(L)$. For this purpose, we construct a reference state R , which we show can be reached from all states given the provided space:

$$R = ((\dots (*^{m_1}) \dots *^{m_{k-2}}) *^{m_{k-1}}) *^{m_k}, \quad (\text{G2})$$

where the configuration $((\dots (*^{m_1}) \dots *^{m_{k-2}}) *^{m_{k-1}}) *^{m_k}$ holds all of the nonlocal conserved quantity, and $m_i \in \{0, 1\}$ are the numbers appearing in the binary representation of Q :

$$Q = \sum_{i=1}^k m_i 2^{k-i}. \quad (\text{G3})$$

Without any $*$ characters ($Q = 0$), the dynamics is entirely ergodic within K_0 . Therefore, the goal is to show that when $Q \neq 0$, all of the $*$ s can be isolated to a reference configuration R on one side of the system.

To this end, given a configuration C , we first isolate $O(\log Q)$ number of 0 characters on one side of the system and use these to create a nest $((\dots () \dots))$ of parentheses. Next, we construct an algorithm for localizing the entirety of the nonlocal conserved quantity into the nest, thereby reducing C to the reference configuration R . Consider the $*$ character nearest to the nest. Labeling the nest by “ q ” (the current amount of the conserved quantity inside of it), this character is positioned as such:

$$q(((\dots (* \dots, \quad (\text{G4})$$

where the number of open parentheses is p . Next, we move the $*$ character via the following sequence:

$$(\dots ((((* \rightarrow (\dots ((** \rightarrow (\dots (**(*) \rightarrow (\dots (**(*) \rightarrow \dots. \quad (\text{G5})$$

Repeating, we obtain the canonical form

$$** (* (* \dots (* (* \dots. \quad (\text{G6})$$

The nest may then absorb the two $*$ characters adjacent to it, forming $q'(* \dots (* \dots$, where $q' = q + 1$. The next step is to collapse any paired parentheses: In other words, $()$ is sent to 00. For instance, this kind of collapse will occur for the configuration

$$q(((\dots (*) \dots \quad (\text{G7})$$

since the canonical form is $q'(* \dots (* \dots$.

After absorbing one unit of the conserved quantity in the nest and performing the collapse process, we iterate these two steps. By construction, this algorithm will eventually localize the conserved quantity in the nest, forming the reference configuration R . So long as the reservoir is large enough, it is possible to transition from any configuration $C \in K$ to R , therefore proving ergodicity.

This result indicates that the dynamics is ergodic within the K sector [up to a mild form of fragile fragmentation that exists due to $\text{EL}(L) > L$]. It also indicates that the dynamics may be slow since transporting the nonlocal conserved quantity out of a region appears to take a very large number of steps. Thus, the relaxation times of $\text{Dyn}_{*,M}$ mimic that of Dyn_{BS} . To formalize this finding, we define the notion of an h restriction of C . First, we perform a preprocessing step where we eliminate as many $()$ pairs as possible in C . The new configuration $\tilde{C}$ is what we call a clean version of C . An h restriction of $\tilde{C}$, which we denote by $S_h(\tilde{C})$ is formed by first drawing a reference line at a height h (note that the height profile of the parentheses is shifted, so the minimum height is at 0). Here, $S_h(\tilde{C})$ denotes a set of contiguous configurations above height h —an example is denoted in Fig. 18. We also label the total nonlocal conserved quantity in a contiguous configuration c by $q_c = \sum_{i \in c} n_{i,*} 2^{h_i}$. Though this quantity should not be associated with a “charge” corresponding to a symmetry, in an abuse of notation, we henceforth call it a charge.

Suppose two contiguous configurations c and c' have charges q_c and $q_{c'}$. If we want to transport an amount of charge Δq between these two configurations, how much time will it take? Note that because c and c' are supported




FIG. 18. Illustration of the process used to define h restrictions. We first clean the configuration, as shown in the upper panel. Then, we draw a line at height h and consider all contiguous regions above this line, which are shaded in red in the bottom panel. We label the nonlocal charge in each of these regions by q_i .

above height h , charge must be pumped in or out of them at increments of 2^h . Therefore, the amount of time required is at least $O(\Delta q/2^h)$. Since the value of the charge supported in c or c' can be exponentially large, transporting some fraction of the charge in c to c' can take an exponentially long time so long as h is not too large.

With this in mind, we can discuss how long it takes to transition between two configurations. Denote the h restriction of C to be $S_h(\tilde{C})$ and that of C' to be $S_h(\tilde{C}')$. Label the charges of the h restriction of C to be (in decreasing order) $q_1 \geq q_2 \geq \dots \geq q_{n_C}$ and for C' to be $q'_1 \geq q'_2 \geq \dots \geq q'_{n_{C'}}$. Here, n_C denotes the number of contiguous regions in the h restriction of C . If $n_C \neq n_{C'}$, then some number of contiguous regions need to be created. [131] In this case, assuming that $n_C \geq n_{C'}$ without loss of generality, we construct the two vectors

$$\begin{aligned}\vec{q} &= \langle q_1, q_2, \dots, q_{n_C} \rangle, \\ \vec{q}' &= \langle q'_1, q'_2, \dots, q'_{n_{C'}}, 0, \dots, 0 \rangle,\end{aligned}\quad (\text{G8})$$

where the number of 0s in $\vec{q}'$ is $n_C - n_{C'}$, indicating a number of yet-to-be created contiguous configurations. The number of charge that needs to be transferred in and out of these contiguous configurations is at least $\Delta q = \|\vec{q} - \vec{q}'\|_1$. The amount of time required for this process is therefore

$$t_{\text{hit}} \geq 2^{-h} \|\vec{q} - \vec{q}'\|_1. \quad (\text{G9})$$

As a result of this bound, there is a simple method for checking whether the time needed to proceed between two configurations is very long. Given two configurations C and C' , we first construct clean versions and successively raise the value of h until their h -restricted charge vectors are significantly different in 1-norms. At this point, so long as h is not too large, we know that it will take a long time to traverse between these configurations.

Note that the hitting times strongly depend on the total value of Q . In particular, we have the obvious bound $Q \leq L \cdot 2^{\max_i h_i}$. For a randomly chosen height profile, $\max_i h_i = O(\sqrt{L})$; thus, we expect $Q \sim \exp(\sqrt{L})$. As a result, we may expect that, for generic sectors, the hitting time could scale around $\exp(\sqrt{L})$, the same scaling as the one argued for in the typical Dehn function of the Baumslag-Solitar group [see Eq. (75) and Appendix E 2].

2. Chiral star-Motzkin

We now provide a more in-depth analysis of the chiral star-Motzkin model discussed in Sec. VI B. Note that, like in the star-Motzkin dynamics, the chiral star-Motzkin dynamics features a nonlocal conserved quantity:

$$Q_R = \sum_i 2^{\sum_{j<i} n_{(j-n)j}} n_{\triangleright, i}. \quad (\text{G10})$$

To understand why large spatial resources are needed, we first consider the following warm-up example. Suppose we have the configuration

$$C = ((\dots (\triangleright) \dots) \triangleright) ((\dots (\triangleright) \dots)) \quad (\text{G11})$$

and we want to convert it to the configuration

$$C' = ((\dots (\triangleright) \dots)) ((\dots (\triangleright) \dots) \triangleright). \quad (\text{G12})$$

In essence, we want to move a single unit of charge from one of the nests to the other one. We can move the $\triangleright$ out of the first nest, yielding

$$((\dots (\triangleright) \dots)) \triangleright \triangleright ((\dots (\triangleright) \dots)), \quad (\text{G13})$$

but unlike in the star-Motzkin model, we cannot move it into the other cluster. The only way to proceed is to collapse the entire cluster, giving

$$(\dots (\triangleright) \dots) \triangleright^{2^h+2} (\dots () \dots) \rightarrow (\dots (\triangleright) \dots) \triangleright^{2^h+2}, \quad (\text{G14})$$

where h is the height of the nest that was collapsed. Next, we may reinsert an empty nest of parentheses forming

$$(\dots (\triangleright) \dots) (\dots () \dots) \triangleright^{2^h+2}, \quad (\text{G15})$$

which, when 2^h of the $\triangleright$ s are used to populate the center of the right nest with an $\triangleright$, gives C' .

Let us more rigorously discuss when a large amount of spatial resources is needed. We work in an intrinsic Krylov sector with fully matched parentheses ($m = n = 0$) and with a fixed value of the nonlocal conserved quantity $Q_R = Q$. Consider an h restriction of configurations C and C' . Label the charges of the contiguous regions (in decreasing order) for C and C' with $\vec{q}$ and $\vec{q}'$; this notation was introduced in the discussion of the star-Motzkin model. We need to transfer an amount of charge to convert between charge configurations $\vec{q}$ and $\vec{q}'$; we define

$$\Delta q_{\max} = \max_i |q_i - q'_i|. \quad (\text{G16})$$

This value is (a lower bound on) the maximum amount of charge that has to be transferred out of a single contiguous region. Some of this charge can be deposited below height h . However, the maximum amount of charge below height h is $L2^h$. Therefore, if $\Delta q_{\max} - L2^h$ is large, then this remaining amount of charge must be transferred to a different contiguous region. As there are at most L contiguous regions, at some point in time during the charge transfer process, an amount of charge has to be inserted in a contiguous region of charge at least $\Delta q_{\max}/L - 2^h$. This process requires an amount of space $O(\Delta q_{\max} 2^{-h}/L)$ since,

at minimum, the entire region needs to be collapsed down to height h before inserting the charge. Therefore, if $\Delta q_{\max} = \epsilon Q$, then this space can be very large. This bound is, in fact, extremely loose but sufficient for our purpose of showing large space complexity.

Therefore, as in the quasifragmented example, given two configurations C and C' , one first constructs clean versions of these configurations and then selects a height h such that the h restriction of C and C' has large $\Delta q_{\max}$. If this is the case, the space complexity scales linearly in $\Delta q_{\max} 2^{-h}$ up to polynomial factors in L .

3. Nongroups and thermalization

Since we could also construct group-based examples with large time and space complexities, it is an interesting question to ask whether there are qualitatively new features that nongroup-based constraints provide to the dynamics.

We answer this question by examining the structure of reduced density matrices of subsystems under the dynamics. Recall that under group dynamics Dyn_G , reduced density matrices of subsystems, defined as $\rho_A = \text{Tr}_{A^c}(\text{Dyn}_G^\dagger \rho_0 \text{Dyn}_G)$, have nonzero values along their diagonals. To explain why, consider a decomposition of the system $S = AA^c$. Start with an initial product state

$$|\psi_0\rangle = |u\rangle_A \otimes |v\rangle_{A^c}, \quad (\text{G17})$$

and suppose v has m zeros (and can be converted under the dynamics to some canonical form $0^m v'$). Under the dynamics, we can perform a sequence of transitions that converts $|u\rangle_A \otimes |0^m v'\rangle_{A^c}$ to $|0^{|A|}\rangle_A \otimes |0^{m-|A|}uv'\rangle_{A^c}$ and, subsequently,

$$|0^{|A|}\rangle_A \otimes |0^{m-|A|}uv'\rangle_{A^c} \rightarrow |w\rangle_A \otimes |w^{-1}0^{m-|A|-|w|}uv'\rangle_{A^c}. \quad (\text{G18})$$

If the number of zeros in the initial state is large enough, then any word w can be produced in A , therefore implying that all diagonal elements of ρ_A will be nonzero. Note that this crucially relies on the existence of inverses, hence the reason why it is special to a group structure.

However, this property *no longer* applies for dynamics that are not based on groups. Instead, we argue that for certain nongroup examples, it is not possible to attain all words u in subsystem A . To see this case, let us consider the chiral star-Motzkin model from the previous subsection. Consider the initial word in the subsystem to be

$$u = (((\cdots (\triangleright) \cdots))) \quad (\text{G19})$$

and the entire initial state $|u\rangle \otimes |v\rangle$ to be in K_Q for some large value of Q . Let us assume that v is generic enough that a constant fraction of it is filled with “0” characters. We define

$$n_{*,A} = \sum_{i \in A} |*\rangle \langle *|_i \quad (\text{G20})$$

and track the probability distribution $p(n_{*,A})$ over time, with

$$p(n, t) = \langle \psi(t) | \mathcal{P}_n | \psi(t) \rangle, \quad (\text{G21})$$

where $\mathcal{P}_n$ projects onto configurations with $n_{*,A} = n$. Let us first suppose that $|A| \ll \log L$. Then, we contend that $p(n, t) > 0$ for all $0 \leq n \leq |A|$. This case can be simply shown as follows: In order to allow for $0 \leq n \leq |A|$, we need to be able to annihilate all of the parentheses in A . Since this process can be completed in roughly $2^{|A|/2}$ space and $|A| \ll \log L$, configurations with all possibly densities $n_{*,A}$ can be reached.

However, when $|A| \gg \log L$, it is no longer the case that all configurations with densities $0 \leq n_{*,A} \leq |A|$ are achievable. To see this result, consider sectors with $Q \gg \exp(|A|)$. Note that in order to have configurations with $n_{*,A} = |A|/2 + q$, we need to annihilate at least q pairs of parentheses in A . If we set the height of the h restriction to be $h = |A|/2 - q$ and set $Q \gg \exp(|A|)$, there will be at least two contiguous regions. To transport an amount of charge 2^q from the contiguous region in A to a contiguous region outside A requires at least $2^q/\text{poly}(L)$ space. If $q \gg \log L$, then this process is not possible. Therefore, we find that for subsystems of size $\gg \log L$, no configurations with charge $n_{*,A} = |A|/2 + q$ are reachable with $q \gg \log L$. This finding implies an unusual property that the reduced density matrix ρ_A will not have full rank unless the size of the subsystem is smaller than $\log L$. In this example, the consequences are easily observable since the value of $n_{*,A}$ is a local operator.

-
- [1] J. M. Deutsch, *Quantum statistical mechanics in a closed system*, *Phys. Rev. A* **43**, 2046 (1991).
 - [2] M. Srednicki, *Chaos and quantum thermalization*, *Phys. Rev. E* **50**, 888 (1994).
 - [3] M. Rigol, V. Dunjko, and M. Olshanii, *Thermalization and its mechanism for generic isolated quantum systems*, *Nature (London)* **452**, 854 (2008).
 - [4] R. Nandkishore and D. A. Huse, *Many-body localization and thermalization in quantum statistical mechanics*, *Annu. Rev. Condens. Matter Phys.* **6**, 15 (2015).
 - [5] L. D'Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, *From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics*, *Adv. Phys.* **65**, 239 (2016).
 - [6] F. Alet and N. Laflorencie, *Many-body localization: An introduction and selected topics*, *Comp. Rend. Phys.* **19**, 498 (2018).
 - [7] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, *Colloquium: Many-body localization, thermalization, and entanglement*, *Rev. Mod. Phys.* **91**, 021001 (2019).

- [8] M. Schiulaz and M. Müller, *Ideal Quantum Glass Transitions: Many-Body Localization without Quenched Disorder*, in *AIP Conference Proceedings*, Vol. 1610 (American Institute of Physics, New York, 2014), pp. 11–23.
- [9] T. Grover and M. P. Fisher, *Quantum disentangled liquids*, *J. Stat. Mech.* (2014) P10010.
- [10] N. Yao, C. Laumann, J. I. Cirac, M. D. Lukin, and J. Moore, *Quasi-many-body localization in translation-invariant systems*, *Phys. Rev. Lett.* **117**, 240601 (2016).
- [11] A. L. Retore, *Introduction to classical and quantum integrability*, *J. Phys. A* **55**, 173001 (2022).
- [12] N. Shiraishi and T. Mori, *Systematic construction of counterexamples to the eigenstate thermalization hypothesis*, *Phys. Rev. Lett.* **119**, 030601 (2017).
- [13] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner *et al.*, *Probing many-body dynamics on a 51-atom quantum simulator*, *Nature (London)* **551**, 579 (2017).
- [14] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, *Quantum scarred eigenstates in a Rydberg atom chain: Entanglement, breakdown of thermalization, and stability to perturbations*, *Phys. Rev. B* **98**, 155134 (2018).
- [15] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, *Weak ergodicity breaking from quantum many-body scars*, *Nat. Phys.* **14**, 745 (2018).
- [16] S. Moudgalya, S. Rachel, B. A. Bernevig, and N. Regnault, *Exact excited states of nonintegrable models*, *Phys. Rev. B* **98**, 235155 (2018).
- [17] C.-J. Lin and O. I. Motrunich, *Exact quantum many-body scar states in the Rydberg-blockaded atom chain*, *Phys. Rev. Lett.* **122**, 173401 (2019).
- [18] M. Schecter and T. Iadecola, *Weak ergodicity breaking and quantum many-body scars in spin-1xy magnets*, *Phys. Rev. Lett.* **123**, 147201 (2019).
- [19] S. Moudgalya, B. A. Bernevig, and N. Regnault, *Quantum many-body scars and Hilbert space fragmentation: A review of exact results*, *Rep. Prog. Phys.* **85**, 086501 (2022).
- [20] A. Chandran, T. Iadecola, V. Khemani, and R. Moessner, *Quantum many-body scars: A quasiparticle perspective*, *Annu. Rev. Condens. Matter Phys.* **14**, 443 (2023).
- [21] P. Sala, T. Rakovszky, R. Verresen, M. Knap, and F. Pollmann, *Ergodicity breaking arising from Hilbert space fragmentation in dipole-conserving Hamiltonians*, *Phys. Rev. X* **10**, 011047 (2020).
- [22] V. Khemani, M. Hermele, and R. Nandkishore, *Localization from Hilbert space shattering: From theory to physical realizations*, *Phys. Rev. B* **101**, 174204 (2020).
- [23] S. Moudgalya and O. I. Motrunich, *Hilbert space fragmentation and commutant algebras*, *Phys. Rev. X* **12**, 011050 (2022).
- [24] N. Pancotti, G. Giudice, J. I. Cirac, J. P. Garrahan, and M. C. Bañuls, *Quantum East model: Localization, non-thermal eigenstates, and slow dynamics*, *Phys. Rev. X* **10**, 021051 (2020).
- [25] M. van Horssen, E. Levi, and J. P. Garrahan, *Dynamics of many-body localization in a translation-invariant quantum glass model*, *Phys. Rev. B* **92**, 100305(R) (2015).
- [26] P. Brighi, M. Ljubotina, and M. Serbyn, *Hilbert space fragmentation and slow dynamics in particle-conserving quantum East models*, *SciPost Phys.* **15**, 093 (2023).
- [27] Z. Lan, M. van Horssen, S. Powell, and J. P. Garrahan, *Quantum slow relaxation and metastability due to dynamical constraints*, *Phys. Rev. Lett.* **121**, 040603 (2018).
- [28] J. P. Garrahan, *Aspects of non-equilibrium in classical and quantum systems: Slow relaxation and glasses, dynamical large deviations, quantum non-ergodicity, and open quantum dynamics*, *Physica (Amsterdam)* **504A**, 130 (2018).
- [29] G. De Tomasi, D. Hetterich, P. Sala, and F. Pollmann, *Dynamics of strongly interacting systems: From Fock-space fragmentation to many-body localization*, *Phys. Rev. B* **100**, 214313 (2019).
- [30] S. Moudgalya, A. Prem, R. Nandkishore, N. Regnault, and B. A. Bernevig, *Thermalization and Its Absence within Krylov Subspaces of a Constrained Hamiltonian*, in *Memorial Volume for Shoucheng Zhang* (World Scientific, Singapore, 2022), pp. 147–209.
- [31] T. Kohlert, S. Scherg, P. Sala, F. Pollmann, B. Hebbe Madhusudhana, I. Bloch, and M. Aidelsburger, *Exploring the regime of fragmentation in strongly tilted Fermi-Hubbard chains*, *Phys. Rev. Lett.* **130**, 010201 (2023).
- [32] T. Kinoshita, T. Wenger, and D. S. Weiss, *A quantum Newton's cradle*, *Nature (London)* **440**, 900 (2006).
- [33] A. Scheie, N. Sherman, M. Dupont, S. Nagler, M. Stone, G. Granroth, J. Moore, and D. Tennant, *Detection of Kardar-Parisi-Zhang hydrodynamics in a quantum Heisenberg spin-1/2 chain*, *Nat. Phys.* **17**, 726 (2021).
- [34] N. Malvania, Y. Zhang, Y. Le, J. Dubail, M. Rigol, and D. S. Weiss, *Generalized hydrodynamics in strongly interacting 1D Bose gases*, *Science* **373**, 1129 (2021).
- [35] D. Wei, A. Rubio-Abadal, B. Ye, F. Machado, J. Kemp, K. Srakaew, S. Hollerith, J. Rui, S. Gopalakrishnan, N. Y. Yao *et al.*, *Quantum gas microscopy of Kardar-Parisi-Zhang superdiffusion*, *Science* **376**, 716 (2022).
- [36] S. Scherg, T. Kohlert, P. Sala, F. Pollmann, B. Hebbe Madhusudhana, I. Bloch, and M. Aidelsburger, *Observing non-ergodicity due to kinetic constraints in tilted Fermi-Hubbard chains*, *Nat. Commun.* **12**, 4490 (2021).
- [37] M. B. Hastings and M. H. Freedman, *Obstructions to classically simulating the quantum adiabatic algorithm*, *Quantum Inf. Comput.* **13**, 1038 (2013).
- [38] In analogy to fragile topology [39], as the ergodicity of the dynamics is modified by the addition of trivial ancillae.
- [39] H. C. Po, H. Watanabe, and A. Vishwanath, *Fragile topology and Wannier obstructions*, *Phys. Rev. Lett.* **121**, 126402 (2018).
- [40] M. Clay and D. Margalit, *Office Hours with a Geometric Group Theorist* (Princeton University Press, Princeton, NJ, 2017).
- [41] Generalizations to 2D are briefly discussed in Sec. VII.
- [42] We focus on the more general structure of a semigroup (as opposed to a group) since it provides the most general framework for discussing 1D constrained dynamics and does not complicate the bounds we obtain on thermalization dynamics; in addition, most of the existing models studied in the literature are described by semigroups rather than groups.

- [43] We are mostly interested in chains with open boundary conditions; for PBC, one may arbitrarily choose a site to serve as the “start” of the word.
- [44] This case is not true for models with *quantum* Hilbert space fragmentation [23,26], violating the above assumptions by virtue of having constraints that cannot be formulated in a local product basis.
- [45] An important caveat here is that proving $\varphi(|w\rangle) = \varphi(|w'\rangle)$ does not need to require outputting an explicit derivation $D(w_1 \rightsquigarrow w_2)$, as for a *particular* G there may exist other (faster) algorithms for constructing a proof. We therefore must stress that when we speak of the time complexity of the word problem, we are referring to its time complexity within the so-called Dehn proof system, a logical system where the *only* manipulations one is allowed to perform are the implementations of local relations in R . If one wants an algorithm that works for all choices of G , operating within this system is generically the best one can do.
- [46] Throughout this work, $\sim, \gtrsim, \lesssim$ will be used to denote (in) equality in the scaling sense. For example, when we write $f(L) \sim g(L)$, we mean that there exists a constant $C > 0$ such that $f(L/C) \leq g(L) \leq f(LC)$.
- [47] S. M. Gersten, D. F. Holt, and T. R. Riley, *Isoperimetric inequalities for nilpotent groups*, *Geom. Funct. Anal.* **13**, 795 (2003).
- [48] N. Brady and M. R. Bridson, *There is only one gap in the isoperimetric spectrum*, *Geom. Funct. Anal.* **10**, 1053 (2000).
- [49] M. V. Sapir, J.-C. Birget, and E. Rips, *Isoperimetric and isodiametric functions of groups*, *Ann. Math.* **156**, 345 (2002).
- [50] R. Young, *Averaged Dehn functions for nilpotent groups*, *Topology* **47**, 351 (2008).
- [51] The same caveat we made for the time complexity also applies here—we are explicitly referring to the space complexity within the “Dehn proof system.”
- [52] The expansion-length function is called the filling length function in the math literature.
- [53] S. Moudgalya and O. I. Motrunich, *From symmetries to commutant algebras in standard Hamiltonians*, *Ann. Phys. (N.Y.)* **455**, 169384 (2023).
- [54] Concretely, the strongly fragmented models of Refs. [21,22] may be phrased as semigroup dynamics for the “dipole semigroup” $\text{Dip} = \text{semi}(0, 1, 2|R)$, where $R = \{020 = 101, 121 = 202, 120 = 201, 021 = 102\}$.
- [55] L. Caha and D. Nagaj, *The pair-flip model: A very entangled translationally invariant spin chain*, *arXiv*: 1805.07168.
- [56] A. Morningstar, V. Khemani, and D. A. Huse, *Kinetically constrained freezing transition in a dipole-conserving system*, *Phys. Rev. B* **101**, 214205 (2020).
- [57] C. Pozderac, S. Speck, X. Feng, D. A. Huse, and B. Skinner, *Exact solution for the filling-induced thermalization transition in a one-dimensional fracton system*, *Phys. Rev. B* **107**, 045137 (2023).
- [58] C. Wang and Z.-C. Yang, *Freezing transition in particle-conserving East model*, *Phys. Rev. B* **108**, 144308 (2023).
- [59] S. Bravyi, L. Caha, R. Movassagh, D. Nagaj, and P. W. Shor, *Criticality without frustration for quantum spin-1 chains*, *Phys. Rev. Lett.* **109**, 207202 (2012).
- [60] This result follows from the fact that the probability of a length- L random walk on $\mathbb{Z}$ being a returning Motzkin walk scales as about $L^{-3/2}$.
- [61] If Dyn_G possesses global symmetries and one modifies the above definitions to operate within a fixed symmetry sector, then various models are known where a *specific* symmetry sector with polynomially strong fragmentation exists [56–58].
- [62] Note that unlike in Ref. [63], here we take the $U_{j,g}$ to be nontrivial in *every* K_{g,ℓ_R} sector, even the one-dimensional ones (where $U_{j,g}$ acts as multiplication by a random phase). This process leads to a comparatively simpler expression for the Markov generator to appear below.
- [63] H. Singh, B. A. Ware, R. Vasseur, and A. J. Friedman, *Subdiffusion and many-body quantum chaos with kinetic constraints*, *Phys. Rev. Lett.* **127**, 230602 (2021).
- [64] Strictly speaking, $\mathcal{M} = \mathcal{M}^T$ only if we double each step of the evolution by including a reflection-related pattern of gates, namely, $\mathcal{M} \rightarrow \prod_{i=1}^{\ell_R} \mathcal{M}_i \prod_{j=1}^{\ell_R} \mathcal{M}_{\ell_R+1-j}$; we will tacitly assume this is always the case.
- [65] D. A. Levin and Y. Peres, *Markov Chains and Mixing Times*, Vol. 107 (American Mathematical Society, Providence, 2017).
- [66] R. Lyons and Y. Peres, *Probability on Trees and Networks*, Vol. 42 (Cambridge University Press, Cambridge, England, 2017).
- [67] However, the *spatial* complexity of these examples—which will be discussed in Sec. VA—is trivial.
- [68] G. Baumslag, *A non-cyclic one-relator group all of whose finite quotients are cyclic: To Bernhard Hermann Neumann on his 60th birthday*, *J. Aust. Math. Soc.* **10**, 497 (1969).
- [69] This picture is most useful for groups, rather than semigroups; see Sec. VI and Appendix A for discussions on this point.
- [70] While the exact structure of CG_G is different for different presentations of G , the large-scale structure of CG_G (scalings of geodesic lengths, graph expansion, and so on) is presentation independent (see, e.g., Ref. [40] for background information). Therefore, we will mostly refer to the Cayley graph of a *group* rather than a particular presentation thereof.
- [71] S. M. Gersten, *Dehn functions and l 1-norms of finite presentations*, in *Algorithms and Classification in Combinatorial Group Theory* (Springer, New York, 1992), pp. 195–224.
- [72] D. Aldous and P. Diaconis, *Shuffling cards and stopping times*, *Am. Math. Mon.* **93**, 333 (1986).
- [73] P. Diaconis, *The cutoff phenomenon in finite Markov chains*, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 1659 (1996).
- [74] For semigroups, we would replace e with a character 0 , which freely moves past other characters.
- [75] Our notation in this discussion will assume that G is a group, for which $\text{EL}(L) \triangleq \text{EL}_e(L) \gtrsim \text{EL}_g(L)$ for all g . For general semigroup dynamics, one would need to consider the $\text{EL}_g(L)$ s separately.

- [76] S. M. Gersten and T. R. Riley, *Filling length in finitely presentable groups*, *Geometriae Dedicata* **92**, 41 (2002).
- [77] Essentially, the best one can do is to first eliminate the cs and expand out $w_{\text{huge}}(n)$ into $w_{\text{large}}(2^n)$, at which point the word lies entirely in the BS subgroup generated by a, b and can be contracted without a further large expansion in length.
- [78] For word problems, this can be done via a computational-basis measurement; for systems with fragmentation in a nonproduct state basis, this measurement requires more care.
- [79] For Hamiltonian dynamics, this quantity is never strictly zero, but using similar arguments from Sec. III, we expect this correlator to take a long time to reach its equilibrium value.
- [80] To understand why, one may envision some notion of a Cayley 2-complex that labels elements of the semigroup. However, since there is no concept of an identity element, loops on the graph do not carry the same meaning that they do for groups. Furthermore, paths or loops on the Cayley 2-complex are no longer freely deformable because the semigroup relations are not invariant under cyclic conjugation, due to the lack of inverses.
- [81] R. Movassagh and P. W. Shor, *Supercritical entanglement in local systems: Counterexample to the area law for quantum matter*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 13278 (2016).
- [82] As we will see shortly, our 2D models are simplest to write down when G is a group rather than a semigroup, although this is not a fundamental restriction.
- [83] D. T. Stephen, O. Hart, and R. M. Nandkishore, *Ergodicity breaking provably robust to arbitrary perturbations*, *Phys. Rev. Lett.* **132**, 040401 (2024).
- [84] C. Stahl, R. Nandkishore, and O. Hart, *Topologically stable ergodicity breaking from emergent higher-form symmetries in generalized quantum loop models*, *SciPost Phys.* **16**, 068 (2024).
- [85] S. Balasubramanian, E. Lake, and S. Choi, *2D Hamiltonians with exotic bipartite and topological entanglement*, [arXiv:2305.07028](https://arxiv.org/abs/2305.07028).
- [86] For dynamics suitably defined on the lattice, processes 1 and 3 are governed by essentially the same process.
- [87] S. Balasubramanian, E. Lake, and A. Mahadevan (to be published).
- [88] D. S. Rokhsar and S. A. Kivelson, *Superconductivity and the quantum hard-core dimer gas*, *Phys. Rev. Lett.* **61**, 2376 (1988).
- [89] M. T. Batchelor and A. Kuniba, *Temperley-Lieb lattice models arising from quantum groups*, *J. Phys. A* **24**, 2599 (1991).
- [90] B. Aufgebauer and A. Klümper, *Quantum spin chains of Temperley-Lieb type: Periodic boundary conditions, spectral multiplicities and finite temperature*, *J. Stat. Mech.* (2010) P05018.
- [91] Y. Han, X. Chen, and E. Lake, *Exponentially slow thermalization and the robustness of hilbert space fragmentation*, [arXiv:2401.11294](https://arxiv.org/abs/2401.11294).
- [92] M. Sapir, *Asymptotic invariants, complexity of groups and related problems*, *Bull. Math. Sci.* **1**, 277 (2011).
- [93] By a “trivial relation,” we mean a free reduction or expansion, namely, a relation of the form $aa^{-1} = e$.
- [94] These additional contributions occur because we need $O(L)$ steps just to eliminate trivial words like ww^{-1} . However, extra L steps can be important when we do not allow words of length greater than L to appear at any stage of the deformation (which we very well may want to do on physical grounds).
- [95] Recall that the geodesic length of a group element is defined by the length of the shortest word representing g , namely, $|g| = \min\{w: \varphi(|w\rangle) = g\}$.
- [96] Since finite groups have trivial geometry (in mathematical language, they are quasi-isometric to the trivial group), this provides another sense in which $O(L)$ contributions to Dehn functions are trivial.
- [97] The L^2 comes from applying relations to move around generators past other ones and next to their inverses.
- [98] In the mathematical literature, a convention is often used where $\text{Dehn}(w)$ is the number of relations needed to reduce w to the identity e , allowing the length of the word to change in the process (so that, e.g., ee may be replaced by e). For physical models of group dynamics, where the total length of the word is always fixed, $\text{Dehn}(w)$ will generically be longer than the value obtained with this convention by a factor of L , corresponding to moves of the form $ea = ae$ that need to be performed during the reduction of w to e^L .
- [99] G. Baumslag and D. Solitar, *Some two-generator one-relator non-Hopfian groups* (1962).
- [100] M. Gromov, *Asymptotic invariants of infinite groups*, Technical Report No. P00001028, 1992.
- [101] T. Riley, *Filling functions*, [arXiv:math/0603059](https://arxiv.org/abs/math/0603059).
- [102] S. M. Gersten, *Isoperimetric and isodiametric functions of finite presentations*, *Geom. Group Theory* **1**, 79 (1993).
- [103] Groups with growth rates intermediate between polynomial and exponential exist—the “simplest” examples are groups arising from automorphisms of trees [104]—but it is conjectured that all such groups have only infinite presentations.
- [104] R. Grigorchuk and I. Pak, *Groups of intermediate growth: An introduction*, *Enseign. Math.* **54**, 251 (2008).
- [105] M. Gromov, *Groups of polynomial growth and expanding maps*, *Publ. Math. l’IHÉS* **53**, 53 (1981).
- [106] E. B. Davies, *Heat kernel bounds, conservation of probability and the feller property*, *J. Anal. Math.* **58**, 99 (1992).
- [107] Note that in order for this proposition to be true, we need the words in K_g to be drawn from the alphabet $S \cup S^{-1} \cup \{e\}$, where S is the generating set. If we were to just use $S \cup S^{-1}$, then it obviously cannot be strictly true, as in this case one could, e.g., have a bipartite Cayley graph such that $|K_{e,L}| = 0$ for odd L .
- [108] W. Woess, *Random Walks on Infinite Graphs and Groups* (Cambridge University Press, Cambridge, England, 2000), p. 138.
- [109] A. Markov, *Impossibility of certain algorithms in the theory of associative systems* (1948).
- [110] G. S. Tseitin, *An associative calculus with an insoluble problem of equivalence*, *Trudy Mat. Inst. Steklova* **52**, 172 (1958).

- [111] P. S. Novikov, *On the algorithmic unsolvability of the word problem in group theory*, Trudy Mat. Inst. Steklova **44**, 3 (1955).
- [112] W. W. Boone, *The word problem*, Ann. Math. **70**, 207 (1959).
- [113] M. O. Rabin, *Recursive unsolvability of group theoretic problems*, Ann. Math. **67**, 172 (1958).
- [114] S. I. Adyan, *Algorithmic unsolvability of problems of recognition of certain properties of groups*, Dokl. Akad. Nauk SSSR, **103**, 48 (1955).
- [115] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, *Undecidability of the spectral gap*, Nature (London) **528**, 207 (2015).
- [116] J. Bausch, T. S. Cubitt, A. Lucia, and D. Perez-Garcia, *Undecidability of the spectral gap in one dimension*, Phys. Rev. X **10**, 031038 (2020).
- [117] N. Shiraishi and K. Matsumoto, *Undecidability in quantum thermalization*, Nat. Commun. **12**, 5084 (2021).
- [118] J. Bausch, T. S. Cubitt, and J. D. Watson, *Uncomputability of phase diagrams*, Nat. Commun. **12**, 452 (2021).
- [119] For $\text{BS}(1, q)$, the $1/2$ in $\mathfrak{b}$ is replaced by $1/q$.
- [120] Note that w_n is nearly a geodesic for a^{2^n} ; the true geodesic with our presentation is instead $b^{-(n-1)}a^2b^{n-1}$, which is one character shorter than w_n .
- [121] Up to the aforementioned $O(|w_{\text{big}}|^2)$ contributions.
- [122] D. B. Epstein, *Word Processing in Groups* (CRC Press, 1992).
- [123] G. Baumslag, S. M. Gersten, M. Shapiro, and H. Short, *Automatic groups and amalgams*, J. Pure Appl. Algebra **76**, 229 (1991).
- [124] J. Burillo and M. Elder, *Metric properties of Baumslag–Solitar groups*, Int. J. Algebra Comput. **25**, 799 (2015).
- [125] If both are nonzero, then we can simplify to $b^k a^{2^n} b^{-l} \sim b^{k-1} a^n b^{-(l-1)}$.
- [126] There is a small subtlety here since this approach is incorrect if $l=0$ and $n=2^m n_o$ with $m > 0, n_o \in 2\mathbb{Z} + 1$. However, in this case, it is easy to show that the naive values of (k, n, l) one obtains from this procedure are $(k, n, l)_{\text{naive}} = (k - m, n_o, -m)$. Since $l_{\text{naive}} < 0$, this mistake is easy to identify and correct.
- [127] The number of length- L Dyck walks scales around $L^{-3/2}2^L$, while the number of returning walks scales around $L^{-1/2}2^L$; thus, the fraction of returning walks that are also Dyck walks scales around $1/L$.
- [128] It is in fact an outstanding open question if there exist groups with $\text{Dehn}(L) > 2^L$ for which the word problem can be solved in linear time (a matrix representation for $\text{BS}^{(2)}$ would answer this question in the affirmative).
- [129] R. C. Lyndon, P. E. Schupp, R. Lyndon, and P. Schupp, *Combinatorial Group Theory*, Vol. 188 (Springer, New York, 1977).
- [130] We thank Tim Riley for suggesting this proof strategy.
- [131] These new regions can either be created from scratch or can be created by borrowing part of an existing contiguous region. However, one can show that in both cases, the amount of time required to pump charge q into this new region will be at least $q/2^h$. Thus, in what follows, we can assume, without loss of generality, that new contiguous configurations are constructed from scratch.