

Complexity as information in spin-glass Gibbs states and metastates: Upper bounds at nonzero temperature and long-range models

N. Read

Department of Physics, Yale University, P.O. Box 208120, New Haven, Connecticut 06520-8120, USA
and Department of Applied Physics, Yale University, P.O. Box 208284, New Haven, Connecticut 06520-8284, USA



(Received 24 January 2022; accepted 24 March 2022; published 23 May 2022)

In classical finite-range spin systems, especially those with disorder such as spin glasses, a low-temperature Gibbs state may be a mixture of a number of pure or ordered states; the complexity of the Gibbs state has been defined in the past roughly as the logarithm of this number, assuming the question is meaningful in a finite system. As nontrivial pure-state structure is lost in finite size, in a recent paper [Phys. Rev. E **101**, 042114 (2020)] Höller and the author introduced a definition of the complexity of an infinite-size Gibbs state as the mutual information between the pure state and the spin configuration in a finite region, and applied this also within a metastate construction. (A metastate is a probability distribution on Gibbs states.) They found an upper bound on the complexity for models of Ising spins in which each spin interacts with only a *finite* number of others, in terms of the surface area of the region, for all $T \geq 0$. In the present paper, the complexity of a metastate is defined likewise in terms of the mutual information between the Gibbs state and the spin configuration. Upper bounds are found for each of these complexities for general finite-range (i.e., short- or long-range, in a sense we define) mixed p -spin interactions of discrete or continuous spins (such as m -vector models), but only for $T > 0$. For short-range models, the bound reduces to the surface area. For long-range interactions, the definition of a Gibbs state has to be modified, and for these models we also prove that the states obtained within the metastate constructions are Gibbs states under the modified definition. All results are valid for a large class of disorder distributions.

DOI: [10.1103/PhysRevE.105.054134](https://doi.org/10.1103/PhysRevE.105.054134)

I. INTRODUCTION

One of several uses of the term “complexity” in the physical and mathematical sciences refers to the decomposition of an equilibrium (Gibbs) state in a short-range classical statistical mechanics system as a mixture of ordered or “pure” Gibbs states [1–3], or to the similar decomposition of stationary (in time) statistical states of a corresponding dynamical system (i.e., with the same Hamiltonian) into ergodic components [4]. (Indeed, the Gibbs states are stationary, and the pure states are ergodic.) Each weight in the mixture can be viewed as the probability for the system to be in the corresponding pure state, and Palmer [5] suggested the name complexity for the entropy of the set of weights, and thus for (roughly speaking) the logarithm of the number of pure or ergodic components in such systems. He also suggested that, in some systems such as short-range spin glasses (SGs), the complexity could be extensive, that is, proportional to system size.

There are two issues with the last proposal. One is that it had to refer to a system of finite size (i.e., having a finite number of degrees of freedom), and typical finite-size models have only a unique pure or ergodic state; a nontrivial pure- or ergodic-state decomposition can arise only in an infinite-size version [1–4]. Glossing over this fact for a moment, and treating the pure states as somehow defined approximately for finite size, van Enter and van Hemmen [6] showed that the complexity of pure states in short-range models cannot

be extensive. (In contrast, ordinary entropy differs on both counts: it is well defined in the finite-size case, and extensive.)

In their deep studies of short-range finite-dimensional classical SGs which used rigorous notions of Gibbs and pure states in strictly infinite size [7,8], Newman and Stein (NS) frequently emphasized the idea of restricting the Gibbs probability distribution for the spins to only the spins in a finite-size region or “window,” such as a hypercube Λ_W of side W . The entropy of such a distribution is finite even when the total system size is infinite. Further, if, for example, one considers different finite sizes of the system, with the interactions among the spins inside and near the window fixed, then at zero temperature different ground states may be seen in the window. It is clear, for example, in a model with interactions between only nearest-neighbor Ising spins that, except in certain degenerate models, the logarithm of the number of such ground states cannot be greater than a constant times the surface area $\sim W^{d-1}$ of the window. This is more restrictive than simply being subextensive, even if extensive is taken to mean $\sim W^d$. Because, in such systems, ground states are the zero-temperature version of pure states, this raises the expectation that the same bound should also hold for some notion of the complexity of a Gibbs state as seen in a window at nonzero temperature. We are not aware of a precise formulation and proof of such a result in print until recently.

In previous work, Höller and the author [9] (we refer to this paper as HR), building on the notions of complexity from

Refs. [5,6] but using the spin distribution restricted to a finite window Λ_W , arrived at a definition of complexity, relative to the window, of an infinite-size Gibbs state (i.e., of its pure-state decomposition), as the *mutual information* between the spins in the window and the pure states. Here, the mutual information represents the average amount of information about which pure state the system is in that is obtained by an observation of the spins in the window. This definition has several desirable properties, and in some situations it reduces to the entropy of the set of weights of pure states as $W \rightarrow \infty$. For Ising spins and nearest-neighbor interactions it was easy to show [9] that it is less than a constant times the surface area $\sim W^{d-1}$ of the window, for any temperature $T \geq 0$.

While the definition of complexity as mutual information in HR is very general, the upper bound on it obtained there (essentially by counting distinct spin configurations on the boundary of the window) is limited in scope in two ways, as just mentioned: (1) the bound does not cover the case of models in which each spin interacts directly with infinitely many others, but which, because the interactions fall off sufficiently rapidly with distance, nonetheless behaves like a short-range model, with well-defined Gibbs states (we say models in this larger class are of finite range); (2) while the bound extends trivially to other models of spins that take a finite number of distinct values, it diverges for continuous spins, such as vector spins with rotation-invariant interactions.

In Sec. III of this paper, we prove upper bounds on the disorder-average of the complexity that overcome both of these limitations for a broad class of models, but only for the case of strictly positive temperature. For interactions involving mean-zero bonds J_{ij} between the pair of spins at i, j , we prove that the disorder average of the complexity (mutual information) of a Gibbs state in a SG is bounded at any $T > 0$ by

$$\frac{1}{T^2} \sum_{\{i, j\}: i \in \Lambda_W, j \notin \Lambda_W} \text{Var } J_{ij}, \quad (1.1)$$

where $\text{Var } J_{ij}$ is the variance of J_{ij} . The bound extends naturally to models with interactions among sets of p spins, instead of $p = 2$. [For these, the precise statements are in Sec. III C 1, inequality (3.67), and in most general form (3.70), which includes non-SGs and nondisordered models.] It applies to continuous spins as well as to discrete, provided the spins are vectors of magnitude ≤ 1 , and it reduces to the surface area $\sim W^{d-1}/T^2$ for nearest-neighbor interactions. For the one-dimensional power-law model [10], in which $i, j \in \mathbf{Z}$ and $\text{Var } J_{ij} = |i - j|^{-2\sigma}$ ($\sigma > 1/2$), it gives a bound $\propto W^{2-2\sigma}/T^2$ for $1/2 < \sigma < 1$, and constant/ T^2 for $\sigma > 1$. While this bound may not be useful if one wishes to take $T \rightarrow 0$ at fixed W , nonetheless for any fixed $T > 0$ it gives a bound on the asymptotic growth of the average complexity as $W \rightarrow \infty$.

The infinite-size Gibbs states that we consider in these results are obtained from taking a thermodynamic limit. As the limit may not exist directly, it is necessary to utilize a metastate, a probability distribution on infinite-size Gibbs states for given bonds [7,8,11]; a metastate carries information about finite-size systems. The same bound applies to the disorder-average both of the complexity of a Gibbs state

sampled from the metastate, and also of the complexity of the metastate-average state (MAS), which is itself a Gibbs state. We further define a concept of the *complexity of the metastate* itself, again relative to the window. This does not refer to a pure-state decomposition, but describes the logarithm of the number of Gibbs states over which the metastate ranges; it vanishes for a trivial metastate (i.e., one supported on a single Gibbs state for the given bonds). Its disorder average is subject to the same bound too. The three complexities are related by a simple formula: the complexity of the MAS is the sum of the other two.

In long-range models, such as cases $\sigma < 1$ in the one-dimensional power-law models, from a rigorous point of view, questions arise even about the definition of Gibbs states [12]. As those models are of interest in the current work, we address those problems as well. It turns out that the concept of relative entropy, to which mutual information is closely related (and which also appears in the proofs of the bounds already discussed), is very useful in proving the existence of metastates and the nature of the Gibbs states in the long-range models, and the methods used are needed for proving the bounds on complexity in these models as well. For these reasons we also include these results here, in the form of Appendix A. We also use similar methods in a proof in Appendix B that there is a unique Gibbs state at $T > 0$ in the short-range case $\sigma > 1$ of the one-dimensional power-law models.

In Sec. III C 5 we give a final discussion of the relation of the results to the problems of SGs.

II. SPIN-GLASS MODELS

We begin by detailing the SG models we have in mind; this section can be skipped or skimmed by knowledgeable readers. (For a general reference, see, e.g., Ref. [13].) The notation Λ , Λ' , ... will stand for sets with elements i . The basic general form of a SG Hamiltonian, due to Edwards and Anderson (EA) [14], is

$$H = - \sum_{\{i, j\} \in \mathcal{E}} J_{ij} s_i s_j - h \sum_i s_i, \quad (2.1)$$

where at present $s_i = \pm 1$ for all sites (vertices) i in some (possibly infinite) index set Λ are Ising spins, and \mathcal{E} is a set of edges, that is, unordered pairs $\{i, j\}$ of sites i, j in Λ ; the real numbers J_{ij} are random variables called bonds, and h is a magnetic field. (To be clear, the symbol Λ for a set of i will not always denote the system as a whole, but here it does.) At present, this expression is formal, that is, we do not yet concern ourselves with convergence of the sum. Unless stated otherwise, we assume the bonds (and their generalizations J_X below) are independent random variables with mean zero and finite variance, and at first we can assume they are Gaussian. (A random variable with mean zero will be termed “centered.”) We write s for $(s_i)_{i \in \Lambda}$ and $J = (J_{ij})_{\{i, j\} \in \mathcal{E}}$. We also write $s|_{\Lambda'}$ for the restriction $(s_i)_{i \in \Lambda'}$ of s to its values on a subset $\Lambda' \subseteq \Lambda$, and similarly for J . (Similar notation will be used for other indexed sets and their restrictions.) We can view a configuration s as a function from Λ into $\{-1, +1\}$, so $s \in S_\Lambda = \{-1, +1\}^\Lambda$, the set of all such functions; for $\Lambda = \mathbf{Z}^d$, we write S for S_Λ . The probability distribution on J

is written $v(J)$, and the operation of expectation with respect to v is denoted by \mathbf{E} or $[\dots]_v$.

While some central calculations can be carried out more generally, we mainly assume that the joint distribution of all the bonds is homogeneous, that is, invariant under some symmetry group that permutes the sites at which spins are located. For the finite-range cases, we assume that the sites i label positions \mathbf{x}_i in the lattice \mathbf{Z}^d , embedded in d -dimensional Euclidean space, and use the induced Euclidean metric to define distances on \mathbf{Z}^d , with the distance between nearest neighbors being 1. Then we assume the distributions of the bonds are invariant under the group \mathbf{Z}^d of translations. We can identify subsets Λ with sets of lattice sites, and use finite such portions Λ of this setup to define finite-size systems, that is, with free boundary conditions; periodic boundary conditions, which ensure translation-invariance in a finite-size system, or other boundary conditions, can be handled with only minor modifications. All the general statements and bounds obtained below are independent of the boundary conditions used. In the basic EA model, the set \mathcal{E} is the set of nearest-neighbor pairs, with the same value of $\text{Var } J_{ij} = 1$, say, for each such pair. Another interesting model is the power-law model [10,15] (mentioned in the one-dimensional case in the introduction), in which \mathcal{E} is the set of all pairs, and $\text{Var } J_{ij} = |\mathbf{x}_i - \mathbf{x}_j|^{-2d\sigma}$.

For infinite-range models, the basic example is the Sherrington-Kirkpatrick (SK) model [16], in which for $i \in \Lambda = \{0, \dots, N-1\}$, the joint distribution of bonds is invariant under the symmetric group S_N that acts by permuting the sites, and the variance is $\text{Var } J_{ij} = 1/N$ for all $\{i, j\}$.

Both types of model can be extended to p -spin interactions for $p \geq 1$. For these, we use notation X, Y, \dots for arbitrary finite subsets of Λ , and also define, for Ising spins, $s_X = \prod_{i \in X} s_i$. Then a p -spin Hamiltonian in system Λ has the form [17]

$$H = H(s) = - \sum_{X \subseteq \Lambda: |X|=p} J_X s_X, \quad (2.2)$$

where the sum is over distinct subsets containing p sites; this is called a p -spin model (for the specified p). Thus the EA Hamiltonian with $h = 0$ is a $p = 2$ model, and $p = 1$ terms are random magnetic fields. Models that contain terms for more than one, or possibly all finite, p ,

$$H = - \sum_{X \in \mathcal{X}(\Lambda)} J_X s_X, \quad (2.3)$$

where $\mathcal{X}(\Lambda)$ is the set of all finite subsets of Λ [we set $\mathcal{X}(\mathbf{Z}^d) = \mathcal{X}$], are called mixed p -spin models. The $X = \emptyset$ term cancels in the Gibbs weights, and so it can be set to zero; also, we write $J = (J_X)_{X \in \mathcal{X}(\Lambda)}$ and $J|_{\mathcal{X}(\Lambda')} = (J_X)_{X \in \mathcal{X}(\Lambda')}$ for $\Lambda' \subseteq \Lambda$. For Λ finite, the mixed p -spin Hamiltonians include all possible Ising Hamiltonians, because $\{s_X : X \subseteq \Lambda\}$ forms a complete set of functions of the spins, which is orthonormal with respect to the inner product defined by summation over all $s \in S_\Lambda$ with weight $2^{-|\Lambda|}$; from an arbitrary function $H(s)$, its expression as $-\sum_{X \subseteq \Lambda} J_X s_X$ (the Fourier-Walsh expansion) can be obtained by using orthonormality. For mixed p -spin SG models, we again require the joint distributions of J to be homogeneous, and there are both finite- and infinite-range versions; in the infinite-range SK-type versions, for $|X| = p$,

the variance scales as $\text{Var } J_X \propto N^{-(p-1)}$ [17]. If desired, here (and also for the m -vector models that follow) we can allow the $p = 1$ term to have nonzero mean (a uniform magnetic field), with no effect on any of the later results.

For vector spins, we consider m -component unit vectors s_i (which we also call m -vector spins) for each $i \in \Lambda$, and the standard inner product on \mathbf{R}^m , with 2-spin interaction terms $-J_{ij} s_i \cdot s_j$ that possess $O(m)$ symmetry. The space of spin configurations is now $s \in S_\Lambda = (\mathbf{S}^{m-1})^\Lambda$, where \mathbf{S}^{m-1} is the unit sphere in \mathbf{R}^m . Summation over the Ising variables is now replaced by integration over S_Λ , using as measure the product of uniform [i.e., $O(m)$ -invariant] measures which are the same on each \mathbf{S}^{m-1} factor. More generally, we could consider m -vector mixed p -site models, and require the interactions to be $O(m)$ invariant, at least for $p > 1$; to be $O(m)$ invariant, a term for a set X of p sites must be a product of 2-spin factors $s_i \cdot s_j$ ($i \neq j$), where for $m > 1$, sites i, j are allowed to appear more than once in the product. Thus for given $p > 1$ and $m > 1$ there is more than one way (in fact, a countable infinity of ways) to construct distinct such terms for each X (with $|X| = p$), and so the terms in H would not be indexed solely by X . $m > 1$ -vector spin models with anisotropic interactions would be similar. For all these cases, in general the $p \geq 1$ terms would have to be indexed using pairs (X, x) in place of X , where X is again a set of sites, and x labels distinct interaction types for each X ; then the random bonds would become $J_{(X,x)}$, the products s_X become $s_{(X,x)}$ [constructed to obey $|s_{(X,x)}| \leq 1$ for all (X, x)] and sums would be taken over (X, x) (we leave the trivial details of such changes to the following results to the reader). [We note that for the Fourier-Walsh expansion, for general $m > 1$, for each i the two basis functions, 1 and s_i , in the Ising case ($m = 1$) must be replaced by an infinite orthonormal set, which could be taken to be constructed from products of components of s_i , i.e., the traceless symmetric tensors on \mathbf{S}^{m-1} , yielding infinite sums even for finite Λ .]

The case of general discrete-spin models, such as Potts spins, is similar to the m -vector models. Such a spin variable can be represented as a discrete set of vectors in \mathbf{S}^{m-1} for some m , and so will not be discussed further here. In every case, we assume that the underlying measure for sums or integrals over s_i for each i is uniform, meaning that it is invariant under a transitive action of a compact group of (measurable) symmetries, namely, the group $O(m)$ for m -vector models, and the permutation group S_p for p -state discrete spin models such as p -state Potts models. (This property will be used in Appendix A.) Note that this statement holds regardless of whether the Hamiltonian is invariant under the group, though it could be. Clearly this general setup can be extended to spins that take values in other spaces, S_0 , say, with a transitive action of a compact group as well. For discrete cases, the existence of a uniform measure that can be normalized as a probability distribution requires that S_0 be finite. For our purposes, the Ising ($m = 1$) and m -vector ($m > 1$) models provide sufficiently representative examples.

Next we will extend the class of disorder distributions beyond Gaussians (readers can continue to think solely of Gaussians if they prefer). We still require that the J_X [or $J_{(X,x)}$] are centered (except possibly for $p = 1$), have finite variance, are independent, and have a homogeneous joint distribution.

It will be convenient to assume also that the marginal distribution of $J_X/\sqrt{\text{Var } J_X}$ [$(J_X - \mathbf{E} J_X)/\sqrt{\text{Var } J_X}$ for $p = 1$] is the same for all X [or all (X, x)], though this will be rarely used, and can easily be relaxed. A class of distributions that we could use is that of subexponential random variables, for which the probability density (relative to Lebesgue measure on the line) decays like the exponential function in (a constant times) $|J_X|$ at large $|J_X|$, or faster (see, e.g., Ref. [18]). This condition ensures that the absolute moments of all orders are finite. We would then require that

$$|\mathbf{E} J_X^n| \leq n! k^n (\text{Var } J_X)^{n/2} \quad (2.4)$$

for all integers $n > 0$ and all X , where k is a constant, as in Refs. [15,19]. These conditions imply that $\mathbf{E} e^{t J_X}$ converges, and so is an analytic function of t , for sufficiently small t [see Ref. [19], Eq. (2.13)], which then implies that J_X is subexponential [18]. In practice, it will *not* be necessary to use the subexponential assumption, or conditions (2.4), in this paper; similarly to Refs. [13,20], the assumptions of independence, $\mathbf{E} J_X = 0$ for $p > 1$, $\text{Var } J_X < \infty$, homogeneity, and the same distribution of $J_X/\sqrt{\text{Var } J_X}$ for all X (modified for $p = 1$ as above), will be the only general conditions imposed on the distributions. (In Sec. III C 4, we will drop most of these conditions, except for independence and translation invariance, and impose a weaker condition in place of the others.)

Now we complete the definitions of the finite- and infinite-range SG models. (In fact, for vector spins it is convenient to impose two additional restrictions, for which see the following Sec. III A 1.) In both cases, we would like to have a nontrivial thermodynamic limit as $|\Lambda|$ (or N) $\rightarrow \infty$ through a suitable sequence of models, for both the usual thermodynamic functions, such as the free energy, and also for the state, or distribution functions for the spins (equivalently, for their correlation functions). (For results for thermodynamics, see especially Refs. [13,15,19–21], Ref. [13] for a review, and also Appendix C; for correlation functions and states in finite-range models at high temperature, see Refs. [22,23].) We will show that a sufficient condition (in addition to those of the preceding paragraph) for these to exist is

$$\lim_{|\Lambda| \rightarrow \infty} \sum_{p \geq 1} \sum_{X \subseteq \Lambda: i \in X, |X|=p} \text{Var } J_X < \infty \quad (2.5)$$

(which is independent of i), which we call the convergence condition. (Here for Ising spins the $p = 1$ term is irrelevant to the convergence of the sum, and can be dropped, as its variance is finite. Technical aspects of the definition and properties of such sequences of partial sums are given in Appendix A.) This condition is clearly equivalent to finiteness of the sums for each p , together with convergence of the sum over p . For the finite-range translation-invariant models, we prove in Appendix A that nontrivial limits for the states exist (in a sense to be explained) at all nonzero temperatures and are Gibbs states (see Sec. III A 1 below) when the convergence condition holds (a different approach that leads to a similar result but only at sufficiently high temperature can be found in Ref. [23]). We note that the sufficient condition for existence of a limit for thermodynamics involves convergence of a similar but (when more than one value of $p > 1$ occurs in the sum) different sum [13,20]; the above condition implies

that condition also (see Appendix C). Thus for the finite-range models, $\text{Var } J_X$ can be taken independent of $|\Lambda|$, but must fall off as the distances between the sites in X increase, such that the sum converges.

The models we call infinite range necessarily involve instead $\text{Var } J_X$ that depend explicitly on $N = |\Lambda|$ (that is, they tend to zero as $N \rightarrow \infty$), as already specified for SK-type p -spin models, such that they obey the same convergence condition (2.5), in which all terms contribute in the limit. Then again the thermodynamic limit of the free energy exists under a similar, though different, condition which is implied by this one, plus a convexity condition [13,21]. The $d = 1$, $p = 2$ power-law model defines a finite-range model for $\sigma > 1/2$, but for $0 \leq \sigma < 1/2$ $\text{Var } J_{ij}$ must scale as $|\Lambda|^{2\sigma-1}$ (or $1/\ln |\Lambda|$ for $\sigma = 1/2$) in order to produce correct behavior of the limit in this case; thus for those values of σ this model is infinite-range, but not strictly of SK type except when $\sigma = 0$. Finally, finite-range models, such as the EA model, in which for each i there are nonzero (and $|\Lambda|$ -independent) terms in H for only a finite number of X with $i \in X$ will be called strictly short range [some might use the term “finite” (i.e., bounded) range for those instead, but we will not]. In the following section we will also subdivide finite-range models into long range and short range, with strictly short range a subset of the short-range models.

III. COMPLEXITY BOUNDS FOR FINITE-RANGE SPIN GLASSES

Now we focus on finite-range SGs. In Sec. III A 1 we discuss Gibbs states and metastates in SGs from a somewhat informal point of view, including the cases of m -vector spins or long-range interactions. In Sec. III B 1 we explain some basic definitions and results from information theory that will be used in the paper. Finally, in Sec. III C 1 we combine the concepts to define complexity as mutual information, introduce basic results, and derive the upper bounds advertised in the Introduction. The basic methods used for the bounds in finite size are quite elementary, and are similar to methods for bounding surface free energies (see, e.g., Ref. [13]). The final (and essential) extension to results for infinite size uses some results from the Appendixes, especially in the long-range case. Section III C 1 ends with the extension to more general disorder distributions, and further physical discussion.

A. Gibbs states, pure states, and metastates

1. Gibbs states and short-range models

In a finite-size system, given a Hamiltonian $H = H(s)$, the Gibbs distribution on spin configurations $s = (s_i)_{i \in \Lambda}$ (Λ finite) at temperature T is given by the well-known formula

$$p_H(s) = e^{-H/T} / \sum_s e^{-H/T}. \quad (3.1)$$

(In general, our convention is to treat the spins as discrete variables, with evident generalizations to continuous variables such as m -vector spins.) This definition does not work when the set Λ becomes infinite, due to the infinite sum in H . Instead, the preliminary definition [1–4] of an infinite-size Gibbs state is as a probability distribution (also called a state)

$\Gamma(s)$ on spin configurations $s = (s_i)_{i \in \mathbf{Z}^d}$ such that for any finite subset Λ of \mathbf{Z}^d , the conditional probability distribution for spins $s|_{\Lambda}$ given $s|_{\Lambda^c}$ (for any Λ , $\Lambda^c = \mathbf{Z}^d - \Lambda$ is its complement) is

$$\Gamma(s|_{\Lambda} | s|_{\Lambda^c}) = p_{H'}(s|_{\Lambda}), \quad (3.2)$$

where

$$H' = - \sum_{X \in \mathcal{X}: \Lambda \cap X \neq \emptyset} J_X s_X, \quad (3.3)$$

so H' [or $H'_{\Lambda}(s)$] contains all interactions entirely within Λ , as well as all interaction terms involving at least one spin in Λ and at least one of the fixed spins in Λ^c . [Note that in our notation, $p_{H'}(s|_{\Lambda})$ depends implicitly on $s|_{\Lambda^c}$ through H' , while the sum in Eq. (3.1) is here over $s|_{\Lambda}$ only.] As \mathbf{Z}^d is infinite, the definition of conditional probability needs care; see Ref. [24,25]. In general, $\Gamma(s)$ may not be uniquely determined by this definition, and that may lead to interesting physics, namely, existence of many pure states. We note that we will write the (conditional) probabilities for a specific vector s value as we did here [24,25], but Γ is really a measure, and we will also write, for example, $\Gamma(A)$ for the thermal probability of a set A of values of s . We write \mathbf{E} or $\langle \dots \rangle$ for the operation of (possibly conditional) expectation over s in a finite- or infinite-size Gibbs state. A state Γ may be a Gibbs state for a specific J , but we will not show that explicitly in the notation.

The definition makes sense in the strictly short-range Ising case, because the sum over X is finite, and further the sum is finite for almost all J (i.e., with probability 1 with respect to ν). For finite-range models more generally, there is an obvious convergence issue regarding the sum over X in H' . In the remainder of this paragraph we restrict ourselves to Ising spins, with simple modifications for more general discrete spins with S_0 finite; we return to vector spins afterwards. The definition continues to make sense if the limits of partial sums (for fixed Λ),

$$\lim_{|\Lambda'| \rightarrow \infty} \sum_{X \subseteq \Lambda': \Lambda \cap X \neq \emptyset} J_X s_X, \quad (3.4)$$

exist and are finite for all s and all finite Λ (again, see Appendix A for the precise definition of the limit). For fixed Λ and J , convergence for all s is guaranteed if the sum converges absolutely for some s . In general, the corresponding partial sums $\sum_{X \subseteq \Lambda': \Lambda \cap X \neq \emptyset} |J_X|$ have a finite limit as $|\Lambda'| \rightarrow \infty$ for ν -almost all J if the sum of expectations

$$\sum_{X \subseteq \Lambda': \Lambda \cap X \neq \emptyset} \mathbf{E}|J_X| \quad (3.5)$$

converges. (This follows from the monotone convergence theorem, i.e., $\mathbf{E}f_n(J) \rightarrow \mathbf{E}f(J)$ for an increasing sequence of nonnegative functions $0 \leq f_n \nearrow f$ as $n \rightarrow \infty$ for almost all J [24,25], applied to the partial sums, together with the fact that, for $f \geq 0$, $\mathbf{E}f < \infty$ implies $f < \infty$ for almost all J .) Under our conditions on the distribution of J_X , that occurs if

$$\lim_{|\Lambda'| \rightarrow \infty} \sum_{X \subseteq \Lambda': \Lambda \cap X \neq \emptyset} \sqrt{\text{Var } J_X} \quad (3.6)$$

is finite. This condition holds for all finite Λ if and only if

$$\lim_{|\Lambda'| \rightarrow \infty} \sum_{p > 1} \sum_{X \subseteq \Lambda': i \in X, |X|=p} \sqrt{\text{Var } J_X} < \infty; \quad (3.7)$$

this is clearly equivalent to the sums for each $p > 1$ converging, together with convergence of the sum over p , and also implies that condition (2.5) holds. If a finite-range model satisfies this condition, then we call it short range; otherwise, it is long range [19]. (The interactions in short-range models are also called regular or absolutely summable [3]; those in the long-range models are also called square summable.) In the long-range case, the definition of a Gibbs state has to be modified somewhat, as given in Sec. 2 of Ref. [12], and requires further discussion; see Appendix A, and also Ref. [23] for a related treatment.

For m -vector spins, or other continuous spins, the sum of $|J_{(X,x)}|$ over (X, x) could diverge even when the system size or Λ' is finite. To avoid this possibility, we consider only finite-range models that obey, in addition to the previous general conditions and condition (2.5), both (1)

$$\sum_{p > 1} \sum_{(X,x): X \subseteq \Lambda', i \in X, |X|=p} \sqrt{\text{Var } J_{(X,x)}} < \infty \quad (3.8)$$

for all finite Λ' , and (2), for the $p = 1$ terms,

$$\sum_{(X,x): X = \{i\}} \mathbf{E}|J_X| \quad (3.9)$$

converges (for all i). Then as before, when the limit $\Lambda' \rightarrow \infty$ of the sum in (i) is finite, we define that to be the short-range case; if the $\Lambda' \rightarrow \infty$ limit is infinite, but the corresponding sum of variances converges, we define that to be the long-range case. In this way the names remain apt, the interactions among any given finite set Λ' of spins $s|_{\Lambda'}$ (and also the $p = 1$ single-site terms) are fairly well behaved, and only minimal changes are required for m -vector spins in the proofs of the results. We impose the same condition for infinite-range continuous spin models. The additional stipulations for continuous spins will usually be left implicit from here on. More general models than these are outside the scope of this paper.

It may be helpful here to consider the example of the 2-spin power-law models in dimension d . Then we have the well-known results that the models are finite range when $\sigma > 1/2$, with a well-defined limit for the thermodynamic properties, and short range when $\sigma > 1$. We emphasize that the distinction that we use in this paper between short-range and long-range models is important technically, but its physical significance for general d is less clear in general. A change in the behavior of our complexity bounds at a value of σ that is different from 1 for $d > 1$ will follow from the later results. For the $d = 1$ power-law 2-spin model [10], the critical value of σ at the short-long boundary is also the value above which the transition at $T > 0$ disappears; this is a consequence of being in one dimension. For analogous physical discussion of higher d , see Refs. [26,27].

2. Pure states

Once Gibbs states for a given Hamiltonian H (thus, for given J) are at hand, we notice that a convex combination (or mixture) of distinct Gibbs states for the same H is again a

Gibbs state; in general, a convex combination could involve an average taken using a probability measure on Gibbs states. A Gibbs state that cannot be expressed as such a combination of other Gibbs states is called an extremal or pure state. Any Gibbs state can be decomposed uniquely into a mixture of pure states [1–4], in the form

$$\Gamma(s) = \sum_{\alpha} w_{\alpha}(\Gamma) \Gamma_{\alpha}(s), \quad (3.10)$$

where Γ_{α} are pure states, and the weights $w_{\alpha} \geq 0$ sum to 1. (Again, in practice the decomposition might be continuous and require an integral over α using a measure $w(\alpha)$ in place of the discrete weights w_{α} , but we will usually not show this explicitly.) The complexity of this decomposition in a SG is one of the main topics investigated here.

3. Metastates and metastate-average state

In general, there may be many Gibbs states for a given H (especially at low temperature), and their relevance physically may not be obvious. To obtain states of physical relevance, we can use some sort of limit of finite-size systems; the metastate concept provides such a construction [7,8,11]. Here we will describe the operational content of the construction [28], leaving more rigorous discussion to Appendix A. The existence of a unique Gibbs state that is the limit of finite-size states is far from clear (again, particularly at low temperature), and the metastate provides instead a probability distribution on Gibbs states, for given J . To describe the pioneering Aizenman-Wehr (AW) construction [11], we first introduce hypercubes such as Λ_W of side W , which is centered at the origin in \mathbf{Z}^d (W is odd), and similar hypercubes Λ_R , Λ_L for odd R, L , where $W < R < L$. We start with a finite-size system in $\Lambda = \Lambda_L$ to produce a Gibbs state $\Gamma_L = p_H(s|_{\Lambda_L})$, where

$$H = - \sum_{X: X \subseteq \Lambda_L} J_X s_X \quad (3.11)$$

(one can use small variations on this choice of H at the boundary of Λ_L to obtain other boundary conditions, including the case of periodic boundary conditions). With this we can calculate thermal averages of functions of the spins $s|_{\Lambda_W}$ in Λ_W for the given J , or functions of such averages; these depend on J in general. An average over the finite-size analog of the metastate is now a disorder average (expectation) using $\nu(J|_{\mathcal{X}(\Lambda_L)-\mathcal{X}(\Lambda_R)})$, where the disorder is (at least partially) outside Λ_R (as the distribution is a product, it can be viewed as conditioned on $J|_{\mathcal{X}(\Lambda_R)}$). Finally, one can average (a function of) the metastate average using $\nu(J|_{\mathcal{X}(\Lambda_R)})$. We then take limits as $L \rightarrow \infty$, then $R \rightarrow \infty$, if these exist. In the limit, the final $\nu(J|_{\Lambda_R})$ is viewed (and written) as disorder average using $\nu(J)$, while the average over the original outer region is denoted as metastate average using the measure κ , which is a probability distribution $\kappa(\Gamma)$ on Gibbs states Γ for given J (i.e., those J in what was formerly the inner region Λ_R before $R \rightarrow \infty$). Likewise, the Gibbs states obtained, written $\Gamma(s)$, are distributions on spin configurations s , obtained as the limit as $W \rightarrow \infty$ subsequent to the prior limits. An AW metastate contains information about the extent to which the Gibbs state Γ (at given J) depends on the disorder asymptotically far away. The technical question of existence of the

limits can be handled by using subsequences, at the cost of possible nonuniqueness of the metastate, and the fact that the infinite-size states are Gibbs states has to be proved. κ or κ_J depends on the bonds J through the Gibbs states Γ , but we will leave that implicit. As in the case of the pure states in the pure state decomposition, we usually treat Γ as a discrete variable, though the distribution κ may in fact be continuous.

In the NS metastate construction [7,8], the use of the disorder average over $J|_{\mathcal{X}(\Lambda_L)-\mathcal{X}(\Lambda_R)}$ is replaced by an empirical average over a range of sizes, say, $R = L_0 < L_1 < \dots < L_K = L$. Again taking L, R , and K to infinity (possibly using a subsequence) produces the metastate average. With a suitable choice of such a subsequence of sizes, a NS metastate can be shown to be the same as a corresponding AW metastate. The NS metastate contains information about the extent to which, asymptotically, the Gibbs state Γ (at given J) varies with system size in the finite sizes from which it was constructed. It too will be denoted κ , as the following results apply equally to either metastate. In either construction, the question of uniqueness of the metastate in a SG remains open.

Finally, there is one additional construction that will be used. In either metastate construction, the Gibbs states can be averaged using the metastate, to produce the metastate-average state (MAS) or barycenter $\rho(s)$ [7,8,11], which is itself a Gibbs state (though a rather special one), and it still depends on the sample of J . That is, $\rho(s) = [\Gamma(s)]_{\kappa}$, where the square bracket $[\dots]_{\kappa}$ denotes metastate average. Then the MAS has pure state decomposition

$$\rho(s) = \sum_{\Gamma} \kappa(\Gamma) \sum_{\alpha} w_{\alpha}(\Gamma) \Gamma_{\alpha}(s) \quad (3.12)$$

$$= \sum_{\alpha} \mu_{\alpha} \Gamma_{\alpha}(s), \quad (3.13)$$

where the weights are $\mu_{\alpha} = \sum_{\Gamma} \kappa(\Gamma) w_{\alpha}(\Gamma)$. The MAS ρ can also be obtained more directly from finite size, by taking the average of the finite-size Gibbs state Γ_L over the disorder in the outer region using $\nu(J|_{\mathcal{X}(\Lambda_L)-\mathcal{X}(\Lambda_R)})$, and then taking limits along the same subsequences as before.

B. Information theory concepts

1. Mutual information and relative entropy

Next we define the concepts from information theory that we will use (see, e.g., Ref. [29]). First, we have the mutual information of two random variables, possibly conditioned on a third. Suppose that \mathcal{A} (\mathcal{B} , \mathcal{C}) is a random variable, which takes values that we generically call a (respectively, b , c). To lighten notation we will treat the values as discrete, and denote the joint probabilities $p(a, b, c)$; notation such as $p(a, b)$ will mean the corresponding marginal probabilities for a, b alone—that is, $p(a, b) = \sum_c p(a, b, c)$, and so on—and the conditional probabilities are given implicitly by, for example, $p(a, c) = p(a|c)p(c)$. Then the mutual information between \mathcal{A} and \mathcal{B} , given \mathcal{C} , is defined as

$$I(\mathcal{A}; \mathcal{B}|\mathcal{C}) = \sum_{a, b, c} p(a, b, c) \ln \frac{p(a, b|c)}{p(a|c)p(b|c)}. \quad (3.14)$$

(The convention $0 \ln 0 = 0$ is used.) The unconditional mutual information $I(\mathcal{A}; \mathcal{B})$ is given by the same formula with c and

conditioning on c deleted everywhere, or one may consider the case in which \mathcal{C} is a constant; the same applies to subsequent definitions also. Our random variables such as \mathcal{A} do not have to be real-valued; they could take other values, including vectors. Consequently, any of them could be interpreted as two or more random variables, leading to the definitions of $I(\mathcal{A}_1, \dots, \mathcal{A}_k; \mathcal{B}_1, \dots, \mathcal{B}_n | \mathcal{C}_1, \dots, \mathcal{C}_m)$ similarly. The (conditional) mutual information can be related to the (conditional) Shannon entropy,

$$S(\mathcal{A}|\mathcal{C}) = - \sum_{a,c} p(a,c) \ln p(a|c), \quad (3.15)$$

by

$$I(\mathcal{A}; \mathcal{B}|\mathcal{C}) = S(\mathcal{A}|\mathcal{C}) - S(\mathcal{A}|\mathcal{B}, \mathcal{C}). \quad (3.16)$$

We also define the relative entropy (or Kullback-Leibler divergence) of two (conditional) probability distributions, which will appear in calculations later. If $p(a, c)$, $q(a, c)$ are two distributions for the same random variables \mathcal{A}, \mathcal{C} , then the conditional relative entropy [i.e., of $p(a|c)$ relative to $q(a|c)$, and using $p(c)$ implicitly] is defined as [29]

$$D[p(a|c)||q(a|c)] = \sum_{a,c} p(a,c) \ln \frac{p(a|c)}{q(a|c)}. \quad (3.17)$$

(Here the left-hand side is a functional of the probability distributions, not a function of a, b, c .) $I(\mathcal{A}; \mathcal{B}|\mathcal{C})$ can be obtained from D by replacing \mathcal{A} by \mathcal{A}, \mathcal{B} , and taking $q(a, b|c) = p(a|c)p(b|c)$, the product distribution formed from the (conditional) marginals.

While we will continue to use the notation as for the discrete case, we note here that for continuous variables, the probabilities such as $p(a)$ must be replaced by the use of a probability measure P , say, where the probability that a is in a set A is $P(A)$ (here A is an element of a suitable σ algebra of measurable sets; in general, in the main text we will not need to concern ourselves with such issues), and expectation over a is expressed as the integral constructed from the measure P . In order to make use of definitions of I and D similar to the preceding ones for the discrete case, we can proceed in either of two equivalent ways. For the first, it may be that the probability $P(A)$ can, for all A , be expressed in terms of a density $p(a)$ (i.e., a measurable function of a) with respect to some reference measure P_0 , such that $P(A) = \int_A p(a)P_0[da]$, where the integral is taken over the subset A in the set of all a values. If this can also be done for another measure Q with density q relative to the same reference P_0 , for variables a, b , and c , then the definition of D takes the same form as above except that the sums are replaced with the corresponding integral with respect to P_0 ; a similar statement holds for I [29]. We note that I and D are invariant under a change of integration variable(s) of the form $a \rightarrow a'(a)$, $b \rightarrow b'(b)$, $c \rightarrow c'(c)$ or under a change of reference measure P_0 [both types of change involve changing p and q by a (common) Jacobian factor so that the measures $P(A)$, $Q(A)$ of a set A of a, b, c, \dots are unchanged]. The alternative, more general, approach uses a partition of the space of a, b, c, \dots , into a discrete set of disjoint subsets, for which the definitions above in the discrete case can be used; D or I can then be defined as the limit of such quantities as the partition becomes arbitrarily

fine [29,30]. Either approach gives well-defined results for those quantities.

As an aside, we note that for the Shannon entropy and conditional entropy in the case of a continuous distribution, the two corresponding approaches are not equivalent, and nor are they well defined. The (conditional) entropy of a partition diverges as the partition becomes arbitrarily fine, while the (conditional) differential entropy using the density relative to P_0 differs from the former by subtraction of a term that diverges in the limit [29], and also is not independent of P_0 . (These effects cancel in I and D .) This leads to additional difficulties for Palmer's definition of complexity as entropy [5], if the distribution w (or μ) of pure states is continuous.

An application of Jensen's inequality shows that I and D are both nonnegative, including in the continuous case; D is zero if and only if $p(a|c) = q(a|c)$ for all a (for all c that have nonzero probability), and hence I is zero if and only if \mathcal{A} and \mathcal{B} are (conditionally) independent, that is, $p(a, b|c) = p(a|c)p(b|c)$ (for all c that have nonzero probability) [29]. Thus D is a (nonsymmetric) nonnegative measure of the distance of the distribution $p(a|c)$ from $q(a|c)$, averaged over c , and I is a nonnegative measure of how far $p(a, b|c)$ is from the product of (conditional) distributions. Hence mutual information is a measure of how much information one gains about the value of \mathcal{B} from an observation of the value of \mathcal{A} (or *vice versa*), conditioned on the value of \mathcal{C} , with again the final c average. I is also a good way to measure the correlation between two subsystems, as it is independent of a choice of some representative random variable for each subsystem, as must be made in a conventional correlation function. There is a remaining issue of whether I and D are *finite*, i.e., nondiverging, to which we return later.

2. Chain rule and Markov property

An elementary calculation from the definitions leads to the chain rule formula [29],

$$I(\mathcal{A}; \mathcal{B}_1, \mathcal{B}_2 | \mathcal{C}) = I(\mathcal{A}; \mathcal{B}_1 | \mathcal{C}) + I(\mathcal{A}; \mathcal{B}_2 | \mathcal{B}_1, \mathcal{C}), \quad (3.18)$$

and also the same with $\mathcal{B}_1, \mathcal{B}_2$ interchanged. It can be extended to various multiterm formulas when, for example, \mathcal{B}_2 is replaced by $\mathcal{B}_2, \dots, \mathcal{B}_n$, by iterating the use of the above formula. [There is a similar version of the chain rule for (conditional) relative entropy.] Because the terms on the right-hand side are nonnegative, the chain rule for mutual information immediately implies various inequalities, such as

$$I(\mathcal{A}; \mathcal{B}_1, \mathcal{B}_2 | \mathcal{C}) \geq I(\mathcal{A}; \mathcal{B}_1 | \mathcal{C}), \quad (3.19)$$

$$I(\mathcal{A}; \mathcal{B}_1, \mathcal{B}_2 | \mathcal{C}) \geq I(\mathcal{A}; \mathcal{B}_1 | \mathcal{B}_2, \mathcal{C}). \quad (3.20)$$

An important application, which will be used repeatedly in the following, arises when there are three random variables, say, $\mathcal{A}, \mathcal{B}, \mathcal{C}$, that constitute a Markov chain, say, in the form $\mathcal{C} \rightarrow \mathcal{B} \rightarrow \mathcal{A}$ [29]. By definition, this means that \mathcal{A} and \mathcal{C} are independent when the value b of \mathcal{B} is given, in other words conditionally on \mathcal{B} : $p(a, c|b) = p(a|b)p(c|b)$. (As the definition is symmetric, the arrows could all be reversed, or replaced by double-headed arrows.) Equivalently, $p(a|b, c) = p(a|b)$ for all c , which is another standard definition of the Markov property. Then the chain rule can be applied in two ways: in

general,

$$I(\mathcal{A}; \mathcal{B}, \mathcal{C}) = I(\mathcal{A}; \mathcal{C}) + I(\mathcal{A}; \mathcal{B}|\mathcal{C}) \quad (3.21)$$

$$= I(\mathcal{A}; \mathcal{B}) + I(\mathcal{A}; \mathcal{C}|\mathcal{B}). \quad (3.22)$$

For the Markov chain $\mathcal{C} \rightarrow \mathcal{B} \rightarrow \mathcal{A}$, in the last line $I(\mathcal{A}; \mathcal{C}|\mathcal{B}) = 0$, so we have

$$I(\mathcal{A}; \mathcal{B}) = I(\mathcal{A}; \mathcal{C}) + I(\mathcal{A}; \mathcal{B}|\mathcal{C}). \quad (3.23)$$

As each term is nonnegative, this yields inequalities

$$I(\mathcal{A}; \mathcal{C}) \leq I(\mathcal{A}; \mathcal{B}), \quad (3.24)$$

$$I(\mathcal{A}; \mathcal{B}|\mathcal{C}) \leq I(\mathcal{A}; \mathcal{B}). \quad (3.25)$$

The first of these is known as the data-processing inequality [29] or as the pipeline inequality, the latter because the information that can be transmitted to \mathcal{A} from \mathcal{C} cannot be greater than that from \mathcal{B} , because it must “come through” \mathcal{B} . The second says that, for the Markov chain, conditioning on \mathcal{C} reduces (or cannot increase) the mutual information between \mathcal{A} and \mathcal{B} .

3. Infinite mutual information

Now we return to the question of whether the mutual information or the relative entropy can be infinite. In formal treatments, it is common to define the (conditional) relative entropy only when $p(a|c)$ is absolutely continuous with respect to $q(a|c)$. This condition requires that if $q(a|c)$ is zero for some a and c , then $p(a|c)$ is zero also. If this condition does not hold for some c with $p(c) > 0$, then one can see that D will be infinite, and that is how we will regard it, rather than imposing the condition of absolute continuity on p, q as part of the definition. (Here we have in mind the discrete case; we discuss the continuous case afterwards.) For the mutual information, $p(a|c) = \sum_b p(a, b|c)$, and so, when a and b are discrete, absolute continuity always holds relative to the product distribution $p(a|c)p(b|c)$. The relative entropy and mutual information could still diverge when the final sum is over an infinite set.

For the continuous case, absolute continuity means that if a set A of a values is given zero probability by Q , then A is also given zero probability by P . When P and Q are represented by functions (densities) relative to P_0 , this simply puts conditions on the behavior of p at the zeros of q . But the absolute continuity would certainly be violated if, for example, Q were represented by a genuine density q (an ordinary function of a) relative to P_0 , but the so-called density p for P contained a δ function of some coordinate, in which case of course p would not really be a function. Then D would again be regarded as infinite. For the mutual information I , the previous argument (with the sum replaced by an integral) for absolute continuity of p relative to the product distribution may now fail if $p(a, b|c)$ contains a δ function rather than being a function. The density for the product distribution, $p(a|c)p(b|c)$, may contain no δ function, even though $p(a, b|c)$ does, yielding infinite I , as we will show with a physical example in a moment. Before that, we state the general result, which is that, in all cases, if the relative entropy of P relative to Q is finite, then P is absolutely continuous with respect to Q [30]. The converse does not hold in general.

For an example, we consider two unit m -vector spins s_1, s_2 , with $m > 1$ and ferromagnetic Hamiltonian $H = -J s_1 \cdot s_2$ ($J > 0$), at temperature $T \geq 0$. At $T > 0$, the Gibbs distribution $p(s_1, s_2) \propto e^{-H(s_1, s_2)/T}$ is continuous, and so are the marginal distributions $p(s_i)$ ($i = 1, 2$), obtained by integrating out the other spin. At zero temperature, we have $s_1 = s_2$ with probability one, which means that the joint distribution contains a δ function in the relative angle coordinates between the two spins on the unit sphere. By symmetry, the marginal distributions $p(s_i)$ are uniform on the unit sphere for all $T \geq 0$. Hence at zero temperature the mutual information $I(\mathcal{S}_1; \mathcal{S}_2)$ of the two spins is infinite. This is quite physical; the mutual information increases steadily as $T \rightarrow 0$, due to the increasing degree of correlation between the two spins [it diverges as $I \sim (m-1) \ln(J/T)$ as $T \rightarrow 0$]. At $T = 0$, the infinite value of I reflects the fact that knowledge of one spin gives knowledge of the other with infinite precision. For the Ising case $m = 1$ with spin-flip symmetry, the mutual information instead tends to $\ln 2$ as $T \rightarrow 0$.

In SGs, when we examine the mutual information between two sets of spins (as discussed in Sec. III C 1 below), for vector spins in the absence of magnetic fields our models may have global $O(m)$ rotation symmetry (for $m = 1$, this reduces to \mathbf{Z}_2 , the spin-flip symmetry of some of our models). Then at $T = 0$ for $m > 1$ the joint distribution will include δ functions in the relative angle coordinates between one set and the other, though now these will be coordinates on $SO(m)$ rather than the unit sphere in \mathbf{R}^m . Hence a similar divergence in the mutual information will occur. This is clear for any Gibbs state in finite size, and again is replaced by $\ln 2$ for $m = 1$ in the presence of \mathbf{Z}_2 symmetry. For all $m \geq 1$, it should be possible to factor off this effect of symmetry to leave mutual information that tends to zero as $T \rightarrow 0$ in any given $O(m)$ -invariant Gibbs state drawn from the metastate. For the MAS, it may be that this is the only effect giving rise to such a divergence. That is, if the global rotation of the ground-state spin configuration were factored out, the remaining distributions might give finite mutual information. For SGs, it is a nontrivial question whether this will be true.

C. Complexity of Gibbs states and metastates

1. Complexity as mutual information

Next we apply mutual information to study the complexity of Gibbs states and metastates. In Ref. [9], Höller and the author introduced a definition of the complexity of a Gibbs state as the mutual information between the spin configuration and the pure states; it was motivated by the earlier idea of a difference of entropies [6]. A Gibbs state Γ drawn from the metastate has a decomposition as in Eq. (3.10), with weights $w_\alpha(\Gamma)$. As the pure states Γ_α are fixed for the given J , we can identify Γ with the set of $w_\alpha(\Gamma)$. The probability of Γ is $\kappa(\Gamma)$. Now let $\mathcal{S}, \mathcal{A}, \mathcal{G}$ be random variables with values s, α, Γ , respectively; then their joint distribution is

$$\kappa(\Gamma) w_\alpha(\Gamma) \Gamma_\alpha(s). \quad (3.26)$$

$\kappa(\Gamma)$ is the marginal distribution for \mathcal{G} , while the conditional distribution for $(\mathcal{A}, \mathcal{S})$ given $\mathcal{G} = \Gamma$ is $w_\alpha(\Gamma) \Gamma_\alpha(s)$.

Even without mentioning a metastate, for a given Γ we can form the mutual information $I(\mathcal{S}; \mathcal{A})_\Gamma$ [using the joint

distribution $w_\alpha(\Gamma)\Gamma_\alpha(s)$] which quantifies the amount of information obtained about which pure state \mathcal{A} the system is in from an observation of \mathcal{S} , the spin configuration of the whole system, for the given Γ . If $S(\mathcal{S})$ is the entropy of the distribution $\Gamma(s)$, and $S_\alpha(\mathcal{S})$ is that of $\Gamma_\alpha(s)$, and we ignore the fact that both are infinite in infinite size, then the mutual information is $I(\mathcal{S}; \mathcal{A})_\Gamma = S(\mathcal{S}) - \sum_\alpha w_\alpha(\Gamma)S_\alpha(\mathcal{S})$ [9]; cf. Eq. (3.20) in Ref. [6], where they use an additional assumption that any spin configuration s determines a unique α , and see HR for further discussion. If we average it over Γ as well, using the metastate, then we can identify the result as $I(\mathcal{S}; \mathcal{A}|\mathcal{G})$; we call this the (metastate average of the) complexity of a (typical) Gibbs state. As the MAS ρ is itself a Gibbs state, we can similarly define its complexity $I(\mathcal{S}; \mathcal{A})_\rho = I(\mathcal{S}; \mathcal{A})$, where the probability used is the marginal $\mu_\alpha\Gamma_\alpha(s)$ for $(\mathcal{S}, \mathcal{A})$ obtained by summing over Γ . As either of these mutual informations may diverge, a more conservative strategy is to replace \mathcal{S} in the definitions with \mathcal{S}_Λ with values $s|_\Lambda = (s_i)_{i \in \Lambda}$. We can then examine how this quantity grows with the size of Λ . This strategy can be viewed as calculating I using a partition, indexed by $s|_\Lambda$, of the space \mathcal{S} , as discussed in Sec. III B 1. This is what was actually used in HR; it can also be done in the following formulas of this subsection, and will be used below, but we leave it implicit for now.

In addition, we can also define the complexity of the metastate itself to be $I(\mathcal{S}; \mathcal{G})$, obtained by using the marginal $\kappa(\Gamma)\Gamma(s)$; this gives the amount of information about the Gibbs state \mathcal{G} obtained from an observation of the spins \mathcal{S} , ignoring the decomposition into pure states. It vanishes for a trivial metastate (i.e., again, one supported on a single Gibbs state for given J). For the AW metastate construction, it is an infinite-size limit of the mutual information between the spins \mathcal{S} (alternatively, the spins in a finite window, \mathcal{S}_Λ) and the disorder in the outer region.

Summarizing these definitions, we introduce symbols for the complexity K_Γ of a (typical) Gibbs state, K_κ of a metastate, and K_ρ of a MAS:

$$K_\Gamma = I(\mathcal{S}; \mathcal{A}|\mathcal{G}), \quad (3.27)$$

$$K_\kappa = I(\mathcal{S}; \mathcal{G}), \quad (3.28)$$

$$K_\rho = I(\mathcal{S}; \mathcal{A}). \quad (3.29)$$

All of these depend implicitly on J ; the first two depend on the metastate κ from which a Γ is drawn, while the last depends only on the MAS ρ . For the versions with a finite window Λ for the spins \mathcal{S}_Λ , we write $K_\Gamma(\Lambda)$ and so on. In the case of discrete spins and strictly short-range interactions, and for $T \geq 0$, a bound proportional to the surface area of Λ , similar to that in HR [9], applies to all three of these, including the complexity of the metastate.

An important observation is that in general the three random variables form a Markov chain,

$$\mathcal{G} \rightarrow \mathcal{A} \rightarrow \mathcal{S}. \quad (3.30)$$

This is because the conditional distribution for \mathcal{G} and \mathcal{S} , given $\mathcal{A} = \alpha$, is

$$\frac{\kappa(\Gamma)w_\alpha(\Gamma)}{\sum_\Gamma \kappa(\Gamma)w_\alpha(\Gamma)}\Gamma_\alpha(s) \quad (3.31)$$

[using $\sum_s \Gamma_\alpha(s) = 1$], which is a product of the marginals conditioned on $\mathcal{A} = \alpha$. [Perhaps more simply, the conditional distribution of \mathcal{S} given $\mathcal{A} = \alpha$ and $\mathcal{G} = \Gamma$ is $\Gamma_\alpha(s)$, which depends on α but is independent of Γ .] The general result for a Markov chain then gives for the present case

$$I(\mathcal{S}; \mathcal{A}) = I(\mathcal{S}; \mathcal{A}|\mathcal{G}) + I(\mathcal{S}; \mathcal{G}). \quad (3.32)$$

This equation is among the main results of this paper. It says that the complexity of the MAS results from (the metastate average of) that of a Gibbs state, plus that of the metastate itself: $K_\rho = K_\Gamma + K_\kappa$, an intuitively appealing statement (we emphasize that a κ expectation is included in the definition of each one). It should not be confused with the simple chain rule; the Markov property was crucial to reduce $I(\mathcal{S}; \mathcal{A}, \mathcal{G})$ to $I(\mathcal{S}; \mathcal{A})$. It implies inequalities as before; of the three complexities, the largest is that of the MAS, K_ρ . Either or both of the terms on the right could be zero, and both will be at high temperature.

2. Monotonicity and bound by mutual information with nearby spins

Now and for the remainder of the paper, we usually consider the mutual information for \mathcal{S}_Λ rather than for \mathcal{S} , and sometimes take the $\Lambda \rightarrow \infty$ limit; Λ could be an arbitrary finite set, but one may think of the example of a hypercube Λ_W of side W , centered at the origin (we take W odd), with sides parallel to the coordinate axes of \mathbb{Z}^d .

For this case, we may first notice a general result in the setting of Gibbs states and metastates. Suppose that $\Lambda_1 \subseteq \Lambda_2$. Then we can identify \mathcal{S}_{Λ_2} with the pair $(\mathcal{S}_{\Lambda_1}, \mathcal{S}_{\Lambda_2 - \Lambda_1})$, and the general inequalities that result from the chain rule imply, for example, that

$$I(\mathcal{S}_{\Lambda_1}; \mathcal{A}|\mathcal{G}) \leq I(\mathcal{S}_{\Lambda_2}; \mathcal{A}|\mathcal{G}) \quad (3.33)$$

and similarly for the other complexities relativized to the windows $\Lambda_1 \subseteq \Lambda_2$. Hence the mutual information is a monotonically increasing function of the window Λ .

We further notice that the random variables \mathcal{A} and \mathcal{G} reflect dependence on spins or bonds far away, effectively at infinity. We expect that the mutual information of \mathcal{S}_Λ with either of them must be mediated by the spins in between, due to the finite range of the interactions. That is, we should have a refined Markov chain, in the form

$$\mathcal{G} \rightarrow \mathcal{A} \rightarrow \mathcal{S}_{\Lambda^c} \rightarrow \mathcal{S}_\Lambda, \quad (3.34)$$

or we could omit \mathcal{A} or \mathcal{G} . (The displayed form means that any three successive terms form a Markov chain as above.) Indeed, the Markov property does hold at \mathcal{S}_{Λ^c} , because for any Gibbs state (for the given J) the conditional distribution of \mathcal{S}_Λ given \mathcal{S}_{Λ^c} is given by the fixed formula (3.2) independent of which Gibbs (or pure) state is considered. [Equation (3.2) is discussed further for the long-range case in Appendix A.]

Consequently, for each complexity that we consider, we have an inequality, for example,

$$K_\rho(\Lambda) \leq I(\mathcal{S}_\Lambda; \mathcal{S}_{\Lambda^c}), \quad (3.35)$$

that results directly from the pipeline inequality above where, to be completely explicit, the right-hand side is the mutual information in the MAS, $I(\mathcal{S}_\Lambda; \mathcal{S}_{\Lambda^c}) = I(\mathcal{S}_\Lambda; \mathcal{S}_{\Lambda^c})_\rho$. The same

upper bound holds for the other complexities $K_\Gamma(\Lambda)$, $K_\kappa(\Lambda)$ also. Hence it will be sufficient to upper bound the right-hand side of this expression, and we will do so for its disorder average over J .

In a strictly short-range model, we can replace Λ^c in these formulas by $\Lambda_2 - \Lambda$ for a sufficiently large finite set Λ_2 , with $\Lambda \subseteq \Lambda_2$, such that the Markov property still holds. For Ising or other discrete spins, this gives one way to recover the bound on any of the complexities by a (model-dependent) constant times the surface area of Λ , thus generalizing the result of HR to the complexity of the metastate in addition to the other two.

3. Bounds on expected mutual information of a partition

In light of the inequality (3.35), we will now consider the mutual information between the spins in two parts of our system. That is, we partition the sites into two or more parts, and we will find upper bounds on the expectation over J of the mutual information between the spins in the two chosen parts. This can be done straightforwardly for a finite system, and then we can consider infinite size. We first consider three useful finite-size examples.

We consider finite systems with sites in a finite set, say, Λ'' . For an arbitrary Hamiltonian $H'' = H''(s|_{\Lambda''})$, say, we will define the free energy for a sum over a subset $\Lambda'^c \subseteq \Lambda''$ of the spin variables on which H'' depends ($\Lambda' = \Lambda'' - \Lambda'^c$),

$$e^{-F_{H''}(s|_{\Lambda'})/T} = \sum_{s|_{\Lambda'^c}} e^{-H''(s|_{\Lambda''})/T}, \quad (3.36)$$

so the free energy $F_{H''}(s|_{\Lambda'})$ will usually depend on $s|_{\Lambda'}$. If $\Lambda' = \emptyset$, we write $F_{H''}$ for $F_{H''}(\emptyset)$; often $\Lambda'' = \Lambda$, the full system, and then we write F for F_H , where H is the full Hamiltonian on Λ .

First, consider the simplest case of a system with two parts, in which the two parts are Λ_1 , Λ_2 , with $\Lambda_1 \cup \Lambda_2 = \Lambda$, $\Lambda_1 \cap \Lambda_2 = \emptyset$. We will write $s = s|_{\Lambda}$ as before, but write $s|_{\Lambda_1}$, $s|_{\Lambda_2}$ as s_1 , s_2 for brevity (no confusion with the strict notation that s_i is the spin at a single site should arise). We simplify the notation for finite-size Gibbs distributions, as follows. For the two-part system, we have

$$p(s_1, s_2) = e^{-H/T+F/T}, \quad (3.37)$$

where $H = H(s)$ is the Hamiltonian of the whole system, and the marginal distribution $p(s_1) = \sum_{s_2} p(s_1, s_2)$ is

$$p(s_1) = e^{-H_1/T-F_{H-H_1}(s_1)/T+F/T}, \quad (3.38)$$

and similarly for $p(s_2)$ by replacing 1 with 2. Here $H_1 = H_1(s_1)$ is the Hamiltonian H with all terms involving spins in Λ_2 omitted, or alternatively it can be thought of as H projected into Λ_1 by summing H over s_2 , as that annihilates the unwanted terms as in the Fourier-Walsh expansion (recall that the constant term in the Hamiltonian is always set to zero). Then the mutual information of the two parts is

$$I(1;2) = \sum_{s_1, s_2} p(s_1, s_2) \ln \frac{p(s_1, s_2)}{p(s_1)p(s_2)}. \quad (3.39)$$

Now consider the same system with all terms that involve spins in both Λ_1 and Λ_2 set to zero, so the parts are decoupled. Denote the corresponding Gibbs distributions by $p^{(0)}(s_1, s_2)$,

$p^{(0)}(s_1)$, and $p^{(0)}(s_2)$, so $p^{(0)}(s_1, s_2) = p^{(0)}(s_1)p^{(0)}(s_2)$. Multiplying and dividing by the latter function inside the logarithm gives

$$I(1;2) = D[p(s_1, s_2)||p^{(0)}(s_1, s_2)] - D[p(s_1)||p^{(0)}(s_1)] \quad (3.40)$$

$$\leq D[p(s_1, s_2)||p^{(0)}(s_1, s_2)] \quad (3.41)$$

by nonnegativity of relative entropy D . The right-hand side is

$$\begin{aligned} D[p(s_1, s_2)||p^{(0)}(s_1, s_2)] \\ = -\frac{1}{T} \sum_{s_1, s_2} p(s_1, s_2)(H - H_1 - H_2) \\ + \frac{F - F_{H_1} - F_{H_2}}{T}. \end{aligned} \quad (3.42)$$

Here F_{H_1} is defined as above, and is independent of s_2 , and similarly for F_{H_2} ; this allowed the sum over s_1, s_2 in the second term to be carried out already.

Next we take the expectation \mathbf{E} of the preceding D over J using the distribution $\nu(J)$ which was specified in Sec. II. In fact, we only need to take the expectation \mathbf{E}' with respect to the terms that involve spins in both parts 1 and 2, in other words X such that $X \cap \Lambda_1 \neq \emptyset$, $X \cap \Lambda_2 \neq \emptyset$, as that will produce the desired bounds which do not depend on the other bonds. The first term in D contains $H - H_1 - H_2$, which consists precisely of all the terms that involve both parts. Its expectation is

$$\frac{1}{T} \mathbf{E}' \sum_{X \subseteq \Lambda: X \cap \Lambda_1 \neq \emptyset, X \cap \Lambda_2 \neq \emptyset} J_X \langle s_X \rangle \quad (3.43)$$

[the expectation is in the Gibbs state $p(s_1, s_2)$].

Here it will be useful to introduce a version of a technique (see, e.g., Ref. [13]) that will be used repeatedly. For given X , consider the modified Hamiltonian with $-J_X s_X$ replaced by $-\lambda J_X s_X$ to form the interpolating Hamiltonian $H - (\lambda - 1)J_X s_X$, and examine what happens as λ changes from 0 to 1. We will write $\langle \dots \rangle_\lambda$ for the thermal average in the interpolating finite-size Gibbs distribution ($\langle \dots \rangle = \langle \dots \rangle_{\lambda=1}$ will continue to mean the usual full thermal average with H). Then we have

$$\mathbf{E}' J_X \langle s_X \rangle = \mathbf{E}' J_X \int_0^1 d\lambda \frac{d}{d\lambda} \langle s_X \rangle_\lambda, \quad (3.44)$$

where the term from the lower limit $\lambda = 0$ is zero because the J_X s are independent, which implies that the thermal average $\langle s_X \rangle_{\lambda=0}$ is independent of J_X , and centered. Then for Ising spins

$$\mathbf{E}' J_X \langle s_X \rangle = \frac{1}{T} \mathbf{E}' J_X^2 \int_0^1 d\lambda (1 - \langle s_X \rangle_\lambda^2) \quad (3.45)$$

$$\leq \frac{1}{T} \mathbf{E}' J_X^2 \quad (3.46)$$

$$= \frac{1}{T} \text{Var} J_X. \quad (3.47)$$

[For Gaussian bonds, this can also be obtained by integration by parts, without introducing λ at this stage. The same bound for m -vector spins with (X, x) in place of X is obtained similarly. An alternative to this upper bound is to use $\mathbf{E}' J_X \langle s_X \rangle \leq$

$\mathbf{E}'|J_X|$, which still leads to useful results in the short-range case; this variation can also be made in the following calculations.]

Hence the first term in $\mathbf{E}D$ is bounded above by

$$\frac{1}{T^2} \sum_{X \subseteq \Lambda: X \cap \Lambda_1 \neq \emptyset, X \cap \Lambda_2 \neq \emptyset} \text{Var } J_X. \quad (3.48)$$

For m -vector spins we obtain similarly an upper bound on the first term in D , which is the same except that the sum over X becomes a sum over (X, x) .

Under our current condition that all J_X are centered, the second term in $\mathbf{E}D$ is in fact nonpositive and can be dropped to obtain the upper bound on $\mathbf{E}D$ and on $\mathbf{E}I(1; 2)$, as follows. It is the difference of the free energies F , with bonds between 1 and 2 included, and $F_{H_1} + F_{H_2}$ with those bonds omitted (set to zero). For the interpolating Hamiltonian $H + (\lambda - 1)(H - H_1 - H_2)$, the corresponding free energy $F(\lambda) = F_{H + (\lambda - 1)(H - H_1 - H_2)}$ obeys $F(1) = F$, $F(0) = F_{H_1} + F_{H_2}$, and is a concave function of λ (its second derivative with respect to λ is ≤ 0). Note that $F(1) - F(0) = \int_0^1 d\lambda dF(\lambda)/d\lambda$. The first derivative of $F(\lambda)$, taken at $\lambda = 0$, is the thermal average of $H - H_1 - H_2$ in the Gibbs distribution $p^{(0)}(s_1, s_2)$ in which the couplings between 1 and 2 are set to zero, so again the \mathbf{E}' expectation of it is zero, as the J_X have mean zero and are independent; then by concavity of $F(\lambda)$ and hence of $\mathbf{E}F(\lambda)$, the first derivative of $\mathbf{E}F(\lambda)$ is nonpositive for all $\lambda \geq 0$. It follows that $\mathbf{E}F(\lambda) \leq \mathbf{E}F(0)$ for all $\lambda \geq 0$, so $\mathbf{E}(F - F_{H_1} - F_{H_2}) \leq 0$, which is what we set out to prove. [If the J_X s are not centered, then the preceding upper bound ≤ 0 can be replaced by $\leq |\mathbf{E}J_X| \leq \mathbf{E}|J_X|$, summed over the X that involve spins in both 1 and 2.]

Thus we have proved that

$$\mathbf{E}I(1; 2) \leq \frac{1}{T^2} \sum_{X \subseteq \Lambda: X \cap \Lambda_1 \neq \emptyset, X \cap \Lambda_2 \neq \emptyset} \text{Var } J_X, \quad (3.49)$$

for the Ising case, with the usual change $X \rightarrow (X, x)$ for more general models. This is another of the main results of this paper, and similar formulas will be found more generally, not just for the case considered here. Note that, from the derivation, the upper bound was a bound on (minus) the expectation of the “surface energy” (the thermal average of $H - H_1 - H_2$), divided by T ; compare Ref. [31]. We can already take $\Lambda_1 = \Lambda_W$ for W fixed, and let Λ increase to infinity, and then we expect the corresponding limit of the upper bound on $\mathbf{E}I(\mathcal{S}_{\Lambda_W}; \mathcal{S}_{\Lambda_W^c})$ to hold for the limit; we will return to this later.

Next we consider the more general case of a partition into three disjoint regions $\Lambda_1, \Lambda_2, \Lambda_3$, and bound the mutual information between two of them, say, regions 1 and 2. We choose to compare with the Gibbs distributions in which all interactions involving spins in both Λ_1 and $\Lambda_2 \cup \Lambda_3$ are set to zero; then we have again

$$I(1; 2) \leq D[p(s_1, s_2) || p^{(0)}(s_1, s_2)], \quad (3.50)$$

where now, with $p^{(0)}(s_1, s_2) = p^{(0)}(s_1)p^{(0)}(s_2)$,

$$p(s_1, s_2) = e^{-H_{12}/T - F_{H - H_{12}}(s_1, s_2)/T + F/T}, \quad (3.51)$$

$$p^{(0)}(s_1) = e^{-H_1/T + F_{H_1}/T}, \quad (3.52)$$

$$p^{(0)}(s_2) = e^{-H_2/T - F_{H_{23} - H_2}(s_2)/T + F_{H_{23}}/T}. \quad (3.53)$$

Here H_{ij} , for $i, j \in \{1, 2, 3\}$, are H projected so that it depends only on s_i, s_j , H_i are H projected so it depends on s_i only, and the F s depend only on the spin variables displayed. This gives

$$\begin{aligned} I(1; 2) &\leq \frac{1}{T} \sum_{s_1, s_2} p(s_1, s_2) \{ -(H_{12} - H_1 - H_2) \\ &\quad - [F_{H - H_{12}}(s_1, s_2) - F_{H_{23} - H_2}(s_2)] \} \\ &\quad + \frac{F - F_{H_1} - F_{H_{23}}}{T}. \end{aligned} \quad (3.54)$$

Taking the expectation value, the first term is bounded above by almost the same sum as when Λ_3 was empty, except that now $X \subseteq \Lambda$ is replaced by $X \subseteq \Lambda_1 \cup \Lambda_2$, and the third term is again nonpositive and can be dropped.

The second term contains

$$F_{H - H_{12}}(s_1, s_2) - F_{H_{23} - H_2}(s_2), \quad (3.55)$$

which again is a free energy, this time with s_1, s_2 fixed, minus the same quantity with the bonds that connect 1 to 3 removed, so 3, where the summation over s_3 occurs, is completely decoupled from 1. As this combination occurs with a negative overall sign, we cannot use the same argument to drop its expectation, and must find a bound on its expectation from the other side instead. We can write this free energy difference as the integral of the derivative of $F(\lambda) = F_{H - H_{12} + (\lambda - 1)(H - H_{12} - H_{23} + H_2)}$, where dependence on s_1, s_2 is implicit. The λ -dependent term of the interpolating Hamiltonian $H - H_{12} + (\lambda - 1)(H - H_{12} - H_{23} + H_2)$ contains only the interactions that involve spins in both Λ_1 and Λ_3 , and possibly also some in Λ_2 . In writing those terms, we will use notation $s_X = s_{X^{(1)}}s_{X^{(2)}}s_{X^{(3)}}$ (for Ising spins) which shows how s_X , and X itself, is split between the three parts ($X^{(i)} = X \cap \Lambda_i \subseteq \Lambda_i$); $X^{(2)}$ could be empty, and $s_\emptyset = 1$, of course.

The minor subtlety with finding the upper bound is that the outer thermal average depends on the bonds whose strength we vary in $F(\lambda)$, though they do not vary in the outer average. We have

$$-\frac{1}{T} \frac{dF}{d\lambda} = \frac{1}{T} \sum_{X \subseteq \Lambda: X^{(1)} \neq \emptyset, X^{(3)} \neq \emptyset} J_X s_{X^{(1)}} s_{X^{(2)}} \langle s_{X^{(3)}} \rangle_{s_1, s_2, \lambda}, \quad (3.56)$$

where the average $\langle \dots \rangle_{s_1, s_2, \lambda}$ is taken using the conditional Gibbs distribution for the interpolating Hamiltonian, with parameter λ , and s_1 and s_2 fixed, not summed over, as indicated. Then in order to bound the \mathbf{E}' expectation, in the X th term in the X sum we must introduce a second parameter λ' into $-J_X s_X$ only, with an integral from 0 to 1 of the λ' derivative, in both the outer thermal average and the inner one; in the latter, this means replacing λ with $\lambda\lambda'$ in the X th term of the

interpolating Hamiltonian. This leads to

$$\begin{aligned} & -\frac{1}{T} \mathbf{E} \left\langle \frac{dF}{d\lambda} \right\rangle \\ &= \frac{1}{T^2} \mathbf{E} \sum'_X J_X^2 \int_0^1 d\lambda' [\lambda (1 - \langle \langle s_{X^{(3)}} \rangle_{s_1, s_2, \lambda, \lambda'}^2 \rangle_{\lambda'}) \\ & \quad + \langle s_{X^{(3)}} \rangle_{s_1, s_2, \lambda, \lambda'} \rangle_{\lambda'} \\ & \quad - \langle s_X \rangle_{\lambda'} \langle s_{X^{(1)}} s_{X^{(2)}} \rangle_{s_1, s_2, \lambda, \lambda'}] \quad (3.57) \end{aligned}$$

$$\leq \frac{1}{T^2} \sum_{X \subseteq \Lambda: X^{(1)} \neq \emptyset, X^{(3)} \neq \emptyset} \text{Var } J_X [\lambda + 2], \quad (3.58)$$

where in the first line the sum \sum'_X is over the same X as in the last line. Then integrating this bound from zero to one, we obtain

$$-\frac{1}{T} \mathbf{E} \langle [F(1) - F(0)] \rangle \leq \frac{5}{2T^2} \sum_{X \subseteq \Lambda: X^{(1)} \neq \emptyset, X^{(3)} \neq \emptyset} \text{Var } J_X, \quad (3.59)$$

so finally we have the desired bound

$$\begin{aligned} \mathbf{E} I(1; 2) & \leq \frac{1}{T^2} \sum_{X \subseteq \Lambda_1 \cup \Lambda_2: X \cap \Lambda_1 \neq \emptyset, X \cap \Lambda_2 \neq \emptyset} \text{Var } J_X \\ & \quad + \frac{5}{2T^2} \sum_{X \subseteq \Lambda: X \cap \Lambda_1 \neq \emptyset, X \cap \Lambda_3 \neq \emptyset} \text{Var } J_X. \quad (3.60) \end{aligned}$$

Again, the same bound (with changes as before) is obtained for m -vector spins. Essentially, the first term is a sum of interactions between 1 and 2 only, while the second is a sum of interactions that involve 1 and 3 and possibly also 2 (the latter only for $p \geq 3$ interactions).

The preceding bound can be applied in more than one way. If 1 and 2 are adjacent regions, while 3 is distant from 1, then as the distance from 1 to 3 goes to infinity, those interactions should go to zero to satisfy the convergence condition, and it reduces to the same bound as when there were only two parts. If instead 3 is close to 1, and 2 goes off to infinity, then the present bound is weaker, so the other case gives a better bound.

As a third example of a bound on mutual information, we consider the conditional mutual information between two parts 1 and 2, conditioned on 3, $I(1; 2|3)$. The application we have in mind is to regions 1 and 2 far apart, and separated by part 3, and we expect that 1 and 2 become conditionally independent in the limit. Thus we consider

$$I(1; 2|3) = \sum_{s_1, s_2, s_3} p(s_1, s_2, s_3) \ln \frac{p(s_1, s_2|s_3)}{p(s_1|s_3)p(s_2|s_3)}. \quad (3.61)$$

In this case we can compare with the Gibbs distribution with the interaction terms that involve at least one spin in each of 1 and 2 set to zero, because the corresponding conditional probability factorizes, $p^{(0)}(s_1, s_2|s_3) = p^{(0)}(s_1|s_3)p^{(0)}(s_2|s_3)$. As before, we then have

$$I(1; 2|3) \leq D[p(s_1, s_2|s_3) || p^{(0)}(s_1, s_2|s_3)], \quad (3.62)$$

where now

$$p(s_1, s_2|s_3) = e^{-(H - H_3)/T + F_{H - H_3}(s_3)/T}, \quad (3.63)$$

$$p^{(0)}(s_1|s_3) = e^{-(H_{13} - H_3)/T + F_{H_{13} - H_3}(s_3)/T}, \quad (3.64)$$

$$p^{(0)}(s_2|s_3) = e^{-(H_{23} - H_3)/T + F_{H_{23} - H_3}(s_3)/T}. \quad (3.65)$$

The calculation of the upper bound on the expectation of this D is similar to the previous example, and the result is

$$\mathbf{E} I(1; 2|3) \leq \frac{3}{T^2} \sum_{X \subseteq \Lambda: X \cap \Lambda_1 \neq \emptyset, X \cap \Lambda_2 \neq \emptyset} \text{Var } J_X. \quad (3.66)$$

(In this case, a term that resembles the one that produced the $1/2$ before now has the opposite sign, so gives zero in the upper bound instead; the other terms combine to give the factor 3.) Thus if the direct interactions involving both 1 and 2 are weak, the conditional mutual information is small, because for fixed spins in 3, the mutual information or correlation arises only from the direct interactions.

So far in this subsection, we have produced upper bounds on mutual information in finite-size systems. We still have to obtain bounds for infinite size, to make contact with the previous discussion of complexity. If we were concerned only with short-range systems, there would not be much of a problem. In that case, if a Gibbs state (say, the MAS ρ) that depends on the bonds J is given in infinite size, then the preceding bounds continue to hold, because the short-range condition ensures absolute convergence of the relevant sums in infinite size, and in particular allows the term-by-term use of bound (3.47). (We discuss this further, though with different motivation, and showing how to avoid reference to the free energy of an infinite portion of the system during the derivation, in Appendix B. Thus here we are *not* saying “take the thermodynamic limit,” as the existence of a limit of the finite-size Gibbs states is not clear even in the strictly short-range case, and that is the reason for introducing metastates.) Then the same upper bound (3.49) with $\Lambda_1 = \Lambda$, $\Lambda_2 = \Lambda^c$ [or the limit $\Lambda_2 \rightarrow \Lambda^c$ of bound (3.60)] applies to the expected mutual information of \mathcal{S}_Λ and \mathcal{S}_{Λ^c} in ρ , that is, to $\mathbf{E} I(\mathcal{S}_\Lambda; \mathcal{S}_{\Lambda^c})_\rho = \mathbf{E} I(\mathcal{S}_\Lambda; \mathcal{S}_{\Lambda^c})$. By the Markov chain arguments, the same upper bound also applies to the expected complexity of the MAS, giving

$$\mathbf{E} K_\rho(\Lambda) \leq \frac{1}{T^2} \sum_{X \in \mathcal{X}: X \cap \Lambda \neq \emptyset, X \cap \Lambda^c \neq \emptyset} \text{Var } J_X. \quad (3.67)$$

As the complexity of the MAS is at least as large as either of the other two complexities $K_\Gamma(\Lambda)$, $K_\kappa(\Lambda)$ by relation (3.32), the same upper bound as in (3.67) applies to the ν -expectation of those also.

But in fact we wish to include the long-range cases, and then we need to take some limits carefully (some readers may prefer to skip the following more technical discussion). Again, it is useful to consider the expected complexity of the MAS. Even in finite size, if we average the state over disorder J_X such that X has zero intersection with $\Lambda_1 \cup \Lambda_2$ [in the notation of the bound (3.60) above], additional effective interactions with the spins in region 1 are generated (the effective Hamiltonian can be obtained by applying Fourier-Walsh expansion to the logarithm of the average state), so a bound may not take the same form as before. It would be better to remove the interactions of region 1 with the outer region before proceeding

with the average and the infinite-size limit. The most efficient way to do so is to use the truncated-interaction metastate κ_{ext} introduced in the course of the proof of Proposition 3 in Appendix A. First, in finite size, we consider Gibbs states $\Gamma_{(\Lambda, \Lambda')}^{n'}$, where n' refers to a finite system on $\Lambda'_{n'}$ ($= \Lambda_L$, say, for now) and Λ, Λ' are smaller finite sets of spins, with $\Lambda \subseteq \Lambda' \subseteq \Lambda'_{n'}$ (thus we will now revert to the notation of the preceding subsections). This is the Gibbs state for the usual Hamiltonian on $\Lambda'_{n'}$ but with all terms that involve spins in both Λ and Λ'^c dropped. For Λ_R such that $\Lambda' \subseteq \Lambda_R \subseteq \Lambda'_{n'}$, we now take the average of $\Gamma_{(\Lambda, \Lambda')}^{n'}$ over disorder (partially) outside Λ_R using $\nu(J|_{\mathcal{X}(\Lambda'_{n'})} - \mathcal{X}(\Lambda_R))$ to obtain $\rho_{(\Lambda, \Lambda')R}^{n'}$. The average state $\rho_{(\Lambda, \Lambda')R}^{n'}$ can be expressed as a Gibbs state for an effective Hamiltonian, and the latter can be obtained by using the Fourier-Walsh expansion of $\ln \rho_{(\Lambda, \Lambda')R}^{n'}$. Due to the truncated interaction, the terms in the effective Hamiltonian that involve the spins in Λ are unaffected by the averaging; they are the same as in H . Then we can bound the expected mutual information between Λ and $\Lambda' - \Lambda$ in the state $\rho_{(\Lambda, \Lambda')R}^{n'}$ exactly as in the bound (3.60) above, where Λ and $\Lambda' - \Lambda$ replace Λ_1 and Λ_2 , and the term involving interactions with region 3 (i.e., $\Lambda_3 = \Lambda'_{n'} - \Lambda'$) drops out as those interactions are zero.

Next, as $\Lambda'_{n'} \rightarrow \infty$, that is, as $n' \rightarrow \infty$, and then $R \rightarrow \infty$ (along the subsequences that produce the extended or truncated-interaction metastate), $\rho_{(\Lambda, \Lambda')R}^{n'}$ tends to the MAS $\rho_{(\Lambda, \Lambda')}$ [7,8,11] in distribution [24,25]; $\rho_{(\Lambda, \Lambda')}$ is the average under the extended metastate of an infinite-size state $\Gamma_{(\Lambda, \Lambda')}$ with truncated interactions, drawn from the extended metastate. The upper bound is still finite when $n' \rightarrow \infty$, so as in the proof of Proposition 2 in Appendix A we have an upper bound on the expected mutual information $EI(\mathcal{S}_\Lambda; \mathcal{S}_{\Lambda'^c - \Lambda})_{\rho_{(\Lambda, \Lambda')}}$ in infinite size [where in this paragraph E means expectation under $\nu(J)$]. In the proof of Proposition 3, it was proved that in any given sample of the indexed set of $(\Gamma, (\Gamma_{(\Lambda, \Lambda')})_{(\Lambda, \Lambda')})$ [that is, a state Γ without truncation, and the collection of states with interactions truncated for all pairs (Λ, Λ')] drawn (simultaneously) from the extended metastate, $\Gamma_{(\Lambda, \Lambda')} \rightarrow \Gamma$ in the strongest possible sense as $\Lambda' \rightarrow \infty$ (along a certain sequence Λ_n) for any Λ , $\nu\kappa_{\text{ext}}$ -almost surely. It then follows that also $\rho_{(\Lambda, \Lambda')} \rightarrow \rho$. Then again the same upper bound applies to the expected mutual information of \mathcal{S}_Λ and $\mathcal{S}_{\Lambda'^c}$ in ρ , and hence the same upper bound as in (3.67) applies to the three complexities $EK_\Gamma(\Lambda)$, $EK_\kappa(\Lambda)$, $EK_\rho(\Lambda)$ in all the finite-range models. These are the final forms of the three inequalities, and are among the main results of this paper.

4. Other distributions for disorder

So far in this paper, we have considered only random variables J_X that are independent with mean zero (except that we can allow nonzero mean when $|X| = 1$) and which all have the same distribution when rescaled by $\sqrt{\text{Var } J_X}$; the discussion also applies to $J_{(X, x)}$ for more general models with the usual modifications. Here we will briefly mention the extension to some other distributions. (Elsewhere in the paper, we continue to use the models of Sec. II.) An easy extension of the results is to the case of J_X with a distribution that is a weighted mixture (i.e., a convex combination) of, say, Gaussian distributions, independently for each X . For these, it can easily be seen that the complexity upper bounds expressed

in terms of $\text{Var } J_X$ are unchanged, because each such variance is just the weighted sum of variances of the Gaussians in the mixture.

In particular, the use of diluted bonds is popular in simulations [32]. In these models, each J_X is either a Gaussian random variable of mean zero and variance 1, with weight (probability) p_X , or zero with probability $1 - p_X$, and the J_X s for each X are independent. A δ -function distribution at zero can be viewed as a Gaussian with zero variance. Then for the $d = 1$ $p = 2$ power-law model, one can take $p_{ij} = |i - j|^{-2\sigma}$, so that $\text{Var } J_{ij}$ takes its usual form $\sim |i - j|^{-2\sigma}$. These models are expected to be in the same universality class for the behavior at $T = T_c$ as the previous power-law model, and are expected to exhibit similar behavior also more generally, for the range $1/2 < \sigma < 1$ (though not when $T = 0$ and $\sigma > 1$ [34]). Then as we said, the upper bounds on the expected complexities take the same form as in the usual model also.

We note, however, that in these diluted-bond models the sums like $\sum_X |J_X s_X| = \sum_X |J_X|$ converge if $\sum_X p_X$ does, and the latter is the same as the corresponding sum of variances. Hence the sums of interest (for, e.g., the existence of Gibbs states) converge *absolutely* whenever the convergence condition (2.5) holds, and not only when the more restrictive condition (3.7) holds. Consequently, for the purposes of this paper the diluted models can be handled with the easier methods that apply to the short-range models, even in cases like the $d = 1$ $p = 2$ models with $1/2 < \sigma < 1$ that for the models of Sec. II we earlier classed as long-range. This illustrates again that, while the condition (2.5) is the important one in general, the technical distinction between short- and long-range cases, which was based on absolute convergence properties of sums, may not be so important physically.

More generally, if the scaling assumption on the distributions is dropped, bounds based on sums of $E|J_X|/T$ in place of $\text{Var } J_X/T^2$ can be used, as mentioned in Sec. III C 3 (for the diluted models, this gives the same result just mentioned), and may be more effective, for example for very broad distributions where $\text{Var } J_X$ may be infinite or its sum over X may be poorly convergent. This leads to the following general form of disorder for which our results on complexity and existence of Gibbs states hold: the bond J_X for each X can be a sum $J_X = J_X^{(1)} + J_X^{(2)}$, where $(J_X^{(1)})_X, (J_X^{(2)})_X$ are all independent, the joint distribution is homogeneous (i.e., translation invariant), the $(J_X^{(1)})_X$ are not necessarily centered but obey the condition for absolute convergence,

$$\lim_{|\Lambda'| \rightarrow \infty} \sum_{p \geq 1} \sum_{X \subseteq \Lambda': i \in X, |X|=p} E|J_X^{(1)}| < \infty, \quad (3.68)$$

while $(J_X^{(2)})_X$ are centered and obey the convergence condition

$$\lim_{|\Lambda'| \rightarrow \infty} \sum_{p \geq 1} \sum_{X \subseteq \Lambda': i \in X, |X|=p} \text{Var } J_X^{(2)} < \infty. \quad (3.69)$$

(These forms of disorder resemble, but the conditions are somewhat more restrictive than, those used in Ref. [20] for thermodynamics, and are mentioned there; see also Appendix C in this paper.) Then in the upper bounds in the proofs of results, the corresponding form of bound is applied to terms containing $J_X^{(1)}$ or $J_X^{(2)}$, respectively, which results in simple

modifications to the bounds on complexity; they become

$$\begin{aligned} \mathbf{E} K_\rho(\Lambda) &\leq \frac{2}{T} \sum_{X \in \mathcal{X}: X \cap \Lambda \neq \emptyset, X \cap \Lambda^c \neq \emptyset} \mathbf{E} |J_X^{(1)}| \\ &\quad + \frac{1}{T^2} \sum_{X \in \mathcal{X}: X \cap \Lambda \neq \emptyset, X \cap \Lambda^c \neq \emptyset} \text{Var} J_X^{(2)}, \end{aligned} \quad (3.70)$$

and the same bound for the other complexities. We note that this form applies even when bonds with $p > 1$ have nonzero mean, so more generally than other results of this paper, and includes models without disorder as special cases. For m -vector models, we can also require that the Hamiltonian restricted to any finite set Λ of spins be almost-surely finite for all $s|_\Lambda$, as before.

There are also versions of Theorem 2 (see Appendix B) for one-dimensional models with this more general form of disorder. In place of a short-range model, we take bonds with $J_X^{(1)} = 0$ for all X . Then we write $J_X^{(1)} = \mathbf{E} J_X^{(1)} + (J_X^{(1)} - \mathbf{E} J_X^{(1)})$, and now a similar argument as in the proof of Theorem 2 shows that a sufficient condition for the Gibbs state at $T > 0$ to be unique is that both

$$\sum_{X \in \mathcal{X}: X \cap \mathbb{Z}_- \neq \emptyset, X \cap \mathbb{Z}_+ \neq \emptyset} |\mathbf{E} J_X^{(1)}| < \infty \quad (3.71)$$

and

$$\sum_{X \in \mathcal{X}: X \cap \mathbb{Z}_- \neq \emptyset, X \cap \mathbb{Z}_+ \neq \emptyset} \text{Var} J_X^{(1)} < \infty \quad (3.72)$$

hold. The second of these is the same as the condition in the original version of Theorem 2, while the first has the same form as the condition in a nonrandom model [33], which is thus a special case of this result. An alternative sufficient condition is to replace that above by the single condition that the same sum, but now of $|\mathbf{E} J_X^{(1)}|$, be finite; this may be weaker if the distribution of each $J_X^{(1)}$ is very broad. There are also other valid conditions that combine both forms.

5. Discussion

As basic examples, if we use the upper bound (3.67) for $\Lambda = \Lambda_W$ a hypercube of side W , for short-range models the sum gives a result proportional to the surface area $\propto W^{d-1}/T^2$ (for fixed $T > 0$ and for a strictly short-range model, this is similar to Ref. [9]). For long-range models, it can be viewed as a definition of surface area for these cases, and generally grows faster than W^{d-1}/T^2 as W increases. For the example of the one-dimensional 2-spin power-law model, which is long range when $1/2 < \sigma \leq 1$, the sum behaves as $W^{2-2\sigma}/T^2$ (for $\sigma < 1$), as stated in the introduction. (The same sum playing the role of the surface area of the window also arose in another distinct bound in the Appendix of Ref. [34].) In these models, for $\sigma > 1$, the bound on any of the complexities is order one as $W \rightarrow \infty$, indicative of low complexity, and consistent with the absence of a spin glass phase at $T > 0$ in these models. We consider these cases further in Appendix B.

We also comment that the third bound (3.66), where we take Λ_3 separating Λ_1 from Λ_2 as that separation becomes large, tells us that distant sets of spins are conditionally independent, if we condition on the spins in between, when

the model is in the finite-range class. To some extent, this justifies the term finite-range, as it means that direct interactions between distant spins really are negligible in their effect; indeed, the relative entropy that we bounded above for this case represents the distance from the state with those interactions dropped.

In a general Gibbs state, the same would not be true if we did not condition on the spins in the region separating Λ_1 from Λ_2 , however, in a *pure* state it would hold without the conditioning, because of the correlation decay property that characterizes pure states (see, e.g., Georgii [1], Ch. 7). Hence in a Gibbs state Γ , the mutual information $I(\mathcal{S}_{\Lambda_1}; \mathcal{S}_{\Lambda_2} | \mathcal{A})_\Gamma \rightarrow 0$ in the limit, and then from a Markov chain argument like those above, applied to $\mathcal{A} \rightarrow \mathcal{S}_{\Lambda^c} \rightarrow \mathcal{S}_{\Lambda_1}$ for $\Lambda_1 \subseteq \Lambda'$ both finite, and Λ^c far distant from Λ_1 , we find that

$$I(\mathcal{S}_{\Lambda_1}; \mathcal{S}_{\Lambda_2})_\Gamma \rightarrow I(\mathcal{S}_{\Lambda_1}; \mathcal{A})_\Gamma \quad (3.73)$$

in the limit $\Lambda_2 \rightarrow \Lambda^c$ and $\Lambda' \rightarrow \infty$. That is, the mutual information in Γ between \mathcal{S}_{Λ_1} and the spins \mathcal{S}_{Λ_2} in a very large region Λ_2 very distant from Λ_1 can serve as a proxy for the complexity of Γ , and this extends immediately to K_Γ and K_ρ .

There are long-standing controversies surrounding the nature of the SG phase in classical SGs. It is our general goal to shed light on these matters. To relate these to complexity, it will be useful first to review briefly some main results of NS [7,8]. First, a metastate, which is a probability distribution on Gibbs states that contains information about behavior in finite size, can be either trivial, that is, it consists of a δ function on a single Gibbs state, or nontrivial, meaning it is spread (or “dispersed”) over more than one Gibbs state. Next, we again discuss the form of the pure state decomposition of a Gibbs state drawn from a given metastate. If the Hamiltonian has a global symmetry, say, under $O(m)$, as it can in the m -vector models (including the Ising $m = 1$ case), for example, then first we note that the metastate constructions preserve the symmetry, and a Gibbs state drawn from the metastate will possess the full symmetry of the Hamiltonian.

If spontaneous symmetry breaking (SSB) occurs in a low-temperature phase, then the Gibbs state will have a decomposition into pure states, at least some of which will not be invariant under the symmetry, and those will map to other pure states under the action of the symmetry, resulting in sets of pure states that each form an orbit under the symmetry action. (An orbit is defined by the property that the group acts transitively on it, that is, any point can be mapped to any other point by a symmetry. If there is no global symmetry, a symmetry orbit is of course a single point.) The invariance of the Gibbs state is then preserved because its decomposition is uniform on each of its symmetry orbits of pure states. Note that for the Ising case, a nontrivial orbit has exactly two pure states in it, while for $SO(m)$ symmetry, $m > 1$, a nontrivial orbit must be a continuum [and in a SG will be a copy of $SO(m)$]. Thus the Gibbs state must have a decomposition into more than one pure state if SSB occurs. But this consequence of SSB is not so interesting for our purposes. Hence we will consider a Gibbs state to be trivial if it consists of a single orbit under the symmetry action, that is, it decomposes into either a single invariant pure state, or into a single orbit consisting of more than one pure state; otherwise it is nontrivial.

Then as both the metastate, and a typical Gibbs state drawn from it, can be either trivial or nontrivial, there are in principle four possible combinations of cases (in this discussion we assume that all the Gibbs states have the same character in this sense, as seems plausible). A further distinction that arises in NS's work is that the cardinality of the set of Gibbs states that may be obtained as a sample drawn from the metastate, and of the set of symmetry orbits of pure states that occur in the pure state decomposition of a Gibbs state, not only could be either one or greater than one, but also could be either countable (i.e., either finite or countably infinite), or uncountable. (Again, we assume that the answer to the last question for the pure state decomposition does not depend on the Gibbs state.) NS [35] proved a theorem that states that, when the Gibbs states are nontrivial and their decomposition is finite or countable, the metastate must be supported on an uncountable number of Gibbs states, and in particular must be nontrivial. NS later showed [36] that in fact, almost surely, a Gibbs state drawn from the metastate is either trivial or consists of infinitely many symmetry orbits, eliminating the finite case just mentioned (see also Ref. [37]). We should note that the results mentioned here involve the use of translation invariance, obtained by using periodic boundary conditions in finite size.

The complexities we have defined are useful as a way to further quantify the degree of dispersal of the metastate and the number of pure states in the decomposition of a Gibbs state. The ν -expectation of the average mutual information between the spins in a window and the pure state in a Gibbs state, $EI(\mathcal{S}_{\Lambda_W}; \mathcal{A}|\mathcal{G})$ (where “average” refers to the average over the Gibbs state), is the (expected) complexity $EK_{\Gamma}(\Lambda_W)$ of a typical Gibbs state Γ , relativized to the window Λ_W ; we will denote this more simply as $EK_{\Gamma}(W)$, and similarly for the others. Likewise, the ν -expectation of the mutual information between the spins in a window and the Gibbs state, $EI(\mathcal{S}_{\Lambda_W}; \mathcal{G})$, is the (expected) complexity $EK_{\kappa}(W)$ of the metastate κ , relativized to the window Λ_W . [We recall that the complexity of the MAS ρ is the sum of these two complexities, $EK_{\rho}(W) = EK_{\Gamma}(W) + EK_{\kappa}(W)$. We will leave the ν -expectation implicit in the remainder of this discussion.] We have seen in Sec. III C 2 that each of the complexities increases monotonically with W , so we can consider how fast they grow as $W \rightarrow \infty$. Typical forms that we may expect are growth as a power law in W , as a logarithm (or more generally perhaps a power of a logarithm), or bounded and tending to a nonnegative constant (possibly zero). As we have seen, the upper bounds generally take a power-law form, so they place upper bounds on the exponent of the power (in one dimension, the bound may instead be constant, ruling out any form of unbounded growth). A complexity that grows as a power of W should correspond to uncountable cardinality, while logarithmic or similar growth may correspond to a countable infinity. Of course, finite cardinality implies bounded complexity, tending to a constant, but not conversely. In applying these remarks to the questions of pure-state decomposition, it would be necessary to remove the contribution to the complexity from SSB, which we have seen implies that the cardinality is that of the continuum in the case of breaking a continuous symmetry. We have not considered how to do that. [Such an infinity was discussed at $T = 0$ in Sec. III B 3.

At nonzero temperature, SSB would imply uncountable cardinality, and infinite complexity as $W \rightarrow \infty$, but for finite W thermal fluctuations would render the complexity finite.]

In many systems, even some with disorder, one does not expect to find infinite complexity or an uncountable number of pure or Gibbs states in a physical (e.g., a metastate) construction, other than when there is breaking of continuous symmetry. Hence the main applications of these ideas and results may be in SG theory. We can illustrate these applications by using various scenarios or models of SGs with Ising spins. First, in the scaling-droplet (SD) theory [27,38,39], the metastate and the Gibbs state were assumed (implicitly or explicitly) to be trivial, and so the complexities are zero (after subtracting the $\ln 2$ due to SSB).

Replica symmetry breaking (RSB), in its (now) standard interpretation and its presumed application to finite-range Gibbs state [40–42], involves nontrivial Gibbs states, and also a nontrivial metastate [7,28,35]. Although in RSB the number of pure states in a given Gibbs state is countably infinite, the complexity of the Gibbs states is finite (as shown explicitly in Ref. [43]), and so this is a case in which, as $W \rightarrow \infty$, $EK_{\Gamma}(W)$ tends to a positive finite limit, which is the expected entropy of $\{w_{\alpha}\}$ [9]. The complexity of the metastate $EK_{\kappa}(W)$ should then grow as a power law. We note that the exponent for the growth of complexity of the MAS is the maximum of the exponents for the other two. In Ref. [9], the complexity of the MAS was used to define an exponent ζ' by

$$EK_{\rho}(W) \sim cW^{d-\zeta'} \quad (3.74)$$

as $W \rightarrow \infty$, for some constant c . (The notation ζ' references another exponent ζ that was defined [28] using a correlation function in the MAS, and argued there to equal ζ' . In RSB, $\zeta = 4$ at zero magnetic field for $d > 6$ [28].) We have the bound $\zeta' \geq 1$ in short-range models, and $\zeta' \geq 2\sigma - 1$ for the one-dimensional $p = 2$ power-law model when $\sigma \leq 1$, and these are valid for $T > 0$, and for m -vector as well as Ising models. (This bound was mentioned recently in Ref. [44].) Hence in this case $d - \zeta'$ describes the growth of complexity of the metastate.

Another scenario, introduced by NS [7,8] for spin-flip invariant Hamiltonians, is called chaotic pairs. In this case the Gibbs states are trivial (a flip-related pair of pure states), while the metastate is nontrivial. The analogous case without spin-flip symmetry has been called chaotic singles [28]. In any of these cases the metastate might not contain an uncountable number of Gibbs states, or might even be finite. The random-field Ising ferromagnet [11] exhibits chaotic singles behavior, apparently with just two pure states (“up” and “down”). A number of short-range SG models with spin-flip symmetry, which differ somewhat from those considered in this paper, were constructed by White and Fisher [45], and appear to possess chaotic-pairs SG phases, with power-law growth of $EK_{\kappa}(W)$.

The remaining member of the set of four classes of combinations of trivial and nontrivial would be a phase in which the metastate is trivial, and the (single) Gibbs state is nontrivial (with an uncountable pure state decomposition). A similar phase arose in infinite-range SG models and was associated with a transition into a dynamically frozen phase (breaking

ergodicity) [46]. The complexity $EK_\Gamma(W)$ of the Gibbs state appeared to be extensive, and the metastate appears to be trivial. It was then argued [47,48] that in a short-range analog, those states are instead “metastable” states, not pure, and the actual state is a “mosaic” of regions of those states, resulting in a single pure Gibbs state in place of the dynamically frozen phase (the state is thus not distinct from the high temperature phase). As stated previously [9], it is not clear to us why such effects must lead to a single pure state, rather than to a distinct phase in which there is a single Gibbs state with subextensive $EK_\Gamma(W)$ with growth exponent $d - \zeta'$ and $1 \leq \zeta' < d$; the latter phase would then be an example of the final class.

Clearly the use of quantitative complexities and their rates of growth sharpens the discussion of the universal properties of the phases beyond simply the trivial-nontrivial and countable-uncountable distinctions. At present, we know of no further rigorous results that would eliminate any of the possible behaviors discussed above. Future work might produce stronger upper bounds on complexity than those found here, which could eliminate some possibilities. Of course, lower bounds would be of great interest also.

IV. CONCLUSION

Let us begin by summarizing some highlights of this paper. For finite-range models (see Sec. II), we defined three complexities (3.27), (3.28), and (3.29), which are respectively the complexity of a (typical) Gibbs state, of a metastate, and of the MAS. The last of these is the sum of the first two by Eq. (3.32). Each of the complexities can be relativized to a finite window Λ_W of size W , giving finite quantities $K_\Gamma(W)$, $K_\kappa(W)$, and $K_\rho(W)$ that increase monotonically with W ; these obey the same relation $K_\Gamma(W) + K_\kappa(W) = K_\rho(W)$. A Markov chain argument showed that each of the relativized complexities is bounded by the mutual information between the spins in Λ_W and those in its complement Λ_W^c , inequality (3.35). From this point of view, what is important is that each complexity involves mutual information with random variables effectively at spatial infinity, and the locality of the finite-range models [expressed as the DLR definition of a Gibbs state; see Eq. (3.2)] implies that the information must be transmitted through the spins in Λ_W^c ; this gives a bound that does not depend on the nature of the variables at infinity. With this, for $T > 0$, we obtained bounds (3.67) or more generally (3.70) on the expectation value of each complexity, which have been discussed further in Secs. I and III C 5. In order to do this rigorously for the general finite-range models it was necessary to provide proofs of the existence and nature of Gibbs states for the long-range cases, which was done in Appendix A. A crucial (and nonelementary) step there was the use of the joint lower semicontinuity of the relative entropy, (A17), to transport bounds from finite size to the infinite-size limit.

A number of issues remain open for further study. The bounds on the expected complexity give us no sense of how large the statistical fluctuations of each complexity may be; as the complexities are not proportional to the volume of the window, the size of fluctuations is not obvious (though square root of the surface area is a possible bound in the short-range cases). Finally, the growth exponent for each complexity is presumably a universal property, independent of most details

of the model, and independent of T at least for $T > 0$ and within a given phase. There are very few circumstances in which we can calculate explicitly the expected complexity, or its growth exponent, in a SG model; the exceptions are scenarios and certain special models [see Sec. III C 5]. This question, which is a more quantitative form of the basic question of whether or not there are many pure states or Gibbs states, gets to the heart of the SG problem in finite-range models.

ACKNOWLEDGMENTS

I am grateful to J. Höller, S. Jensen, and A.P. Young for stimulating collaborations, to C. M. Newman, D. L. Stein, and V. Vu for discussions, to A. C. D. van Enter for discussions and references, to A. Barron for information on the lower semicontinuity of relative entropy under weak* convergence, and to J. Fröhlich for a remark during a talk by the author. The author is grateful for the support of NSF Grant No. DMR-1724923.

APPENDIX A: GIBBS STATES, METASTATES, AND PURE STATES FOR LONG-RANGE MODELS

Here we give, for the finite-range models, a somewhat detailed account of Gibbs [i.e., Dobrushin-Lanford-Ruelle (DLR)] states, which will be defined more stringently here, and discuss metastates also. Some of the discussion will be more technical than that in the main text. We discuss an extension of the DLR definition of a Gibbs state for long-range mixed p -spin models, and show that, within a metastate construction, the states produced satisfy that definition. We focus on Ising spins, but for m -vector spins we make some parenthetical side remarks when the necessary changes are more than merely notation. Some of the constructions (in particular, part of the proof of Proposition 3) will be utilized in the main text also (without circularity).

1. Partial sums and infinite-size limits

The definition of a Gibbs state was begun in Sec. III A 1. The discussion frequently involves partial sums like

$$\lim_{|\Lambda| \rightarrow \infty} \sum_{X \subseteq \Lambda: \dots} \dots, \quad (A1)$$

where “ \dots ” is a function of X and there may be additional stipulations \dots on the X included in the sum (sometimes it appears as $X \subseteq \Lambda'$ instead of Λ). These will be appear throughout the discussion in this Appendix, and are defined as follows: we evaluate the sum for $\Lambda = \Lambda_n$ (or $\Lambda' = \Lambda_n$), where $(\Lambda_n)_n$ is a cofinal sequence of finite subsets of \mathbb{Z}^d , that is, for any X , there is an n such that $X \subseteq \Lambda_n$, and $\Lambda_n \subset \Lambda_{n+1}$ (strictly increasing); then the limit is just $n \rightarrow \infty$. (It may be unconventional to include strictness of increasingness in the definition, but it will be convenient.) These conditions ensure that eventually all X (or all X satisfying additional conditions) are included in the sum. In the short-range case of the sum (3.4), convergence is absolute, and so does not depend on the choice of the cofinal sequence $(\Lambda_n)_n$ for Λ' , which appears there. Indeed, the set \mathcal{X} of finite subsets of \mathbb{Z}^d is countable, so here we could use any enumeration of \mathcal{X} to define the sum,

and by absolute convergence it always gives the same sum. In the sums in the following, in principle we must specify the cofinal sequence $(\Lambda_n)_n$ in the notation whenever some terms in the sum could be negative; if not, or when we have reason to believe that the choice of cofinal sequence makes no difference, we write simply Λ or $\Lambda' \rightarrow \infty$ to show the use of an unspecified cofinal sequence. [For m -vector or other models, there may be sums over interaction types x as well as over X . When the sums over x are infinite for some X , we would need a cofinal sequence of finite sets of (X, x) . Under the additional condition imposed in Sec. III A 1, this will not be necessary in practice.] When discussing the existence of the thermodynamic limit for thermodynamics, it is necessary to specify more precisely how a cofinal sequence $\Lambda_n \rightarrow \mathbf{Z}^d$; for those we will stipulate (without further comment) convergence “in the sense of Fisher” [13,19], which holds when $(\Lambda_n)_n$ is a cofinal sequence of hypercubes, for example.

2. Gibbs states

For the finite-range models in general, the limit of a sum of the form (3.4), which we reproduce here and which is involved in the definition of a Gibbs state as in Eq. (3.2),

$$\lim_{n \rightarrow \infty} \sum_{X \subseteq \Lambda_n: X \cap \Lambda \neq \emptyset} J_X s_X, \quad (\text{A2})$$

may not exist at all, and if it does, it could depend on the cofinal sequence Λ_n that is used. Before addressing that issue, we make further progress by noticing an equivalent form of the definition of a Gibbs state [1–3], in which the normalizing factor in $p_{H'}$ in Eq. (3.2) is removed. This is done simply by taking ratios for distinct $s|_\Lambda$, for given $s|_{\Lambda^c}$:

$$p_{H'}(s'|_\Lambda) / p_{H'}(s|_\Lambda) = \exp[-(H'(s') - H'(s))/T]; \quad (\text{A3})$$

where $s'|_{\Lambda^c} = s|_{\Lambda^c}$ and $H' = H'_\Lambda(s)$ was defined in Eq. (3.3). (This applies for $T \geq 0$; for $T = 0$, it can be viewed as defined by the limit $T \rightarrow 0^+$, and the exponential may have to be interpreted as 0 or ∞ , even in a finite-size system.) Here one sees that this ratio is independent of Λ , provided Λ is finite and contains $\{i : s'_i \neq s_i\}$ as a subset. Hence it is sufficient to consider Λ as small as possible, that is, the case $\Lambda = \{i : s'_i \neq s_i\}$, in which s_i is changed, $s_i \rightarrow -s_i = s'_i$, for all i in Λ , and $s'_i = s_i$ for $i \in \Lambda^c$; we then have (formally, because we have not yet discussed convergence of the infinite sum)

$$H'(s') - H'(s) = H(s') - H(s) \quad (\text{A4})$$

$$= \sum_{X: X \cap \Lambda \neq \emptyset} (-J_X s'_X + J_X s_X) \quad (\text{A5})$$

$$\equiv 2h_\Lambda(s) \quad (\text{A6})$$

(of course, $s'_X = \prod_{i \in X} s'_i$), which defines the formal sum $h_\Lambda(s)$ up to the convergence issue. [For m -vector spins, we must consider s'_i differing from s_i by an arbitrary rotation for each $i \in \Lambda$, not simply reversal, so $h_\Lambda(s)$ is also a function of those rotations; for later use, the rotations should be constants, independent of $s|_{\Lambda^c}$.] In this form, the DLR definition of a Gibbs state says that

$$\frac{\Gamma(s'_\Lambda | s|_{\Lambda^c})}{\Gamma(s_\Lambda | s|_{\Lambda^c})} = \exp[-2h_\Lambda(s)/T] \quad (\text{A7})$$

for Γ -almost every s and for all Λ . Given the ratios for all $\Lambda \subseteq \Lambda'$ (Λ' finite), the general conditional probabilities $\Gamma(s_{\Lambda'} | s|_{\Lambda^c})$ can be recovered (which involves finding the normalizing factor). This then gives the full alternate version of the preliminary definition of Gibbs states, if the values of $h_\Lambda(s)$ are known.

In this form, for $T > 0$ it might be that, for some s , $h_\Lambda(s)$ should be viewed as taking one of the values $\pm\infty$, which would mean that one of the two conditional probabilities is zero for this s . (In fact, we will soon see that $-\infty$ cannot occur.) This means that, on a set of s of nonzero Γ probability, s' cannot occur: the spins in Λ cannot all be reversed, and at least one is “locked,” or not reversible, when starting from s . In the short-range case, $h_\Lambda(s)$ converges absolutely for all s (and so is finite, $|h_\Lambda(s)| < \infty$), with ν probability 1; hence in this Appendix we mostly focus on the long-range case (most statements also apply to short-range ones, some with trivial proofs). Thus we arrive at the question for long-range models: do locked spins occur in the Gibbs states when $T > 0$? Gandolfi, Newman, and Stein (GNS) [12] showed that in long-range models, there exist distributions for the spins (i.e., states) in which some spins are locked, however those constructions are, by their own account, unphysical as they involve J -dependent boundary conditions. They gave a modified definition of Gibbs states in which locked spins (in the above sense) do not occur. (A related construction of Gibbs states in Ref. [23] is more closely related to the specification and DLR equations in Sec. A 5 below; the results show that, when T is sufficiently high, the thermodynamic limit from finite size exists, and locked spins do not occur.)

The idea contained in the definition of a Gibbs state that appears in the work of GNS is that spin configurations s that produce difficulties with the sums $h_\Lambda(s)$ occur with Γ probability zero (we will extend their definitions from $p = 2$ to general mixed p -spin models). We will need a definition for the meaning of the infinite sum. GNS took it to be defined using partial sums along a cofinal sequence $(\Lambda_n)_n$ (fixed independently of s , Λ , and J) as follows: retaining the same notation as above, we define, for Λ' another finite set,

$$2h_\Lambda(s|_{\Lambda'}) = \sum_{X \subseteq \Lambda': X \cap \Lambda \neq \emptyset} J_X (-s'_X + s_X), \quad (\text{A8})$$

and then define h_Λ by

$$h_\Lambda(s) = \lim_{n \rightarrow \infty} h_\Lambda(s|_{\Lambda_n}) \quad (\text{A9})$$

for finite sets Λ and configurations s if this limit exists as an extended real number, that is, either real or $\pm\infty$. (It is not obvious that this is the most appropriate procedure.) Thus, at this stage we could allow the possibility $h_\Lambda = \pm\infty$.

We can also define a version of $h_\Lambda(s)$ when $T > 0$ directly from Eq. (A7):

$$2\tilde{h}_\Lambda(s) \equiv -T \ln \frac{\Gamma(s'_\Lambda | s|_{\Lambda^c})}{\Gamma(s_\Lambda | s|_{\Lambda^c})} \quad (\text{A10})$$

provided the right-hand side exists. (Alternatively, we can rewrite this as a formula for $\tilde{h}_\Lambda(s)/T$, and then it is valid at $T = 0$ also). In these terms, we would want to show (and eventually, we will show) both that $h_\Lambda(s)$ exists and that it equals $\tilde{h}_\Lambda(s)$ for Γ -almost all s . Here, however, as the idea

is to consider $\tilde{h}_\Lambda(s)$ as a random variable that depends on s , and s is drawn from some distribution $\Gamma(s)$, $\tilde{h}_\Lambda(s) = -\infty$ would imply that, conditionally on $s|_{\Lambda^c}$, the spins are locked in configuration s' , and s is forbidden. But, just as for elementary probability theory of discrete variables, $p(a|c)$ does not vanish unless also $p(a, c) = 0$, so in general the Γ probability of the set of s such that $\Gamma(s|_\Lambda \mid s|_{\Lambda^c}) = 0$ must be zero. Thus with Γ probability 1, $\tilde{h}_\Lambda(s)$ cannot be $-\infty$, when it exists.

$h_\Lambda(s)$ [and $h(s)$] is a generalization of a local effective magnetic field $h_i(s)$: by taking $\Lambda = \{i\}$, we define $h_i(s)$ by $h_{\{i\}}(s) = h_i(s)s_i$. For $p = 2$ p -spin interaction models, h_i was used by GNS, and in that case, for general Λ , $h_\Lambda(s) = \sum_{i \in \Lambda} h_i(s)s_i - 2 \sum_{i, j \in \Lambda} J_{ij}s_i s_j$, where the last sum is finite. In this case, for fixed s , existence of $h_\Lambda(s)$ for all finite Λ is equivalent to existence of $h_i(s)$ for all i . [Here we assumed that, when $h_\Lambda(s)$ exists, it cannot be $-\infty$.]

Now, using the GNS definition of $h_\Lambda(s)$ as the limit of partial sums, following GNS we define the *allowed configurations* as those s for which $h_\Lambda(s)$ exists, and is finite, for all finite Λ (with the given J). (Configurations that are not allowed will be called nonallowed. Also, for m -vector models, we must say “for all $s'|_\Lambda$ ” as well as “for all finite Λ .”) Finally, again following GNS, we make

Definition 1: A Gibbs state at $T \geq 0$ is a probability distribution Γ on S , for given J , with the following two properties: (1) the distribution is supported on allowed configurations only, that is, $\Gamma(\{\text{allowed } s\}) = 1$; (2) for all finite subsets Λ , the conditional probabilities are given by the Gibbs form

$$\Gamma(s|_\Lambda \mid s|_{\Lambda^c}) = p_{H'}(s|_\Lambda) \quad (\text{A11})$$

[or alternatively by the formula for the ratios, so that $h_\Lambda(s) = h_\Lambda(s')$, for Γ -almost every s as before.]

We comment that (a) the evaluation of the right-hand side in (2) is well defined as a consequence of property 1); (b) this characterization of a given Γ as a Gibbs state obviously depends on the definition of h_Λ , and might depend on the choice of the sequence $(\Lambda_n)_n$ when partial sums are used; (c) for $T > 0$, vanishing conditional probabilities [i.e., $h_\Lambda(s) = +\infty$] for s' given $s|_{\Lambda^c}$ do not occur, so there are no locked spins. While this certainly holds in the short-range models of this paper and in other short-range models, the motivation in general for this part of the definition may not be obvious at this stage. We will see that it is satisfied in the finite-range models considered here; and (d) the definition makes sense at $T = 0$ as well as $T > 0$, by taking $T \rightarrow 0^+$ as above. In this case, (2) implies that conditionally on s , for any finite set of sites, reversing those spins does not decrease the energy: $h_\Lambda(s) \geq 0$ for all Λ , and (1) implies that this energy change is well defined and finite. We call an allowed configuration s a *ground state* if $h_\Lambda(s) \geq 0$ for all Λ ; then the definition implies that, for $T = 0$, $\Gamma(\{\text{ground states}\}) = 1$ (i.e., all configurations that occur are ground states). This is clear physically.

We note that at $T = 0$, for nondegenerate J a Gibbs state might be supported on a unique ground state (which would be a pure state), or if spin-flip symmetry is present, on the flip-invariant mixture of a pair of flip-related ground states (each of which is pure). A Gibbs state could also be a more general mixture of such states, implying nonuniqueness of ground states (or of ground state pairs). For degenerate J , as

is well known (for example, for $p = 2$ interactions only, with $J_{ij} = \pm 1$ for all $\{i, j\} \in \mathcal{E}$, the set of nearest neighbor pairs), there could be extensive entropy at $T = 0$, and individual spin configurations drawn from such a Gibbs state would still be ground states, but pure states would not be supported on a single allowed configuration, because a single configuration would not be a Gibbs state. Thus it would be a mistake to think that the general such Gibbs state necessarily possesses an extensive complexity of pure states (cf. Ref. [5]).

If $h_\Lambda(s)$ is defined in the manner of GNS as the limit of partial sums, then the allowed configurations are defined without reference to temperature. We note that whether or not a series (with given coefficients J_X) converges is a *tail event* in s , that is, the convergence (or not) of $h_\Lambda(s)$ is unaffected if s_i is changed on a finite set of i . (It is also a tail event in J , similarly.) For similar reasons, the allowed configurations form a dense set in S in the product topology, and so do the nonallowed (in the long-range case), because a base of open sets for S is given by the collection of sets that are a product of open sets for a finite collection of i , times a copy of S_0 for all other sites [24,25,49,52], and so any nonempty open set contains both allowed and nonallowed configurations (in the long-range case). Turning to measure-theoretic properties, a $T = \infty$ Gibbs state would be the uniform distribution on s . That suggests that under that distribution, almost every configuration is allowed: for a configuration drawn from the uniform distribution, the sum $h_\Lambda(s)$ would converge (in the sense of partial sums along a cofinal sequence) with probability one. This is reminiscent of the textbook “random signs” problem of a series with independent uniformly random ± 1 factors in each term [24,25], and indeed, for $p = 2$, $h_i(s)$ reduces to exactly that form. In that case the series converges almost surely if and only $\sum_j J_{ij}^2$ converges, which holds ν -almost surely when the convergence condition holds. We may expect similar results for the more general mixed p -spin case for the uniform distribution on s . For general $T < \infty$, a state will not be absolutely continuous with respect to the uniform $T = \infty$ state, and so it is not ruled out that it is not a Gibbs state under the above definition. Nonetheless, one would expect that cancellations between terms, due to spins far apart being almost independent, could make the series converge for typical configurations drawn from the state, and that is the type of result at which we are aiming. (See also Refs. [22,23] for results at high T .)

Of course, so far we have only made a definition; the question now is whether the physically relevant states of finite-range SGs are Gibbs states in the sense of Definition 1. To address this, we will examine the states produced by the metastate construction.

3. Metastates

For the metastate construction, we follow AW [11] and NS, Ref. [50] or [51] (see also Ref. [3]), and sketch the main steps. First, we consider finite-size models on finite $\Lambda'_{n'}$ for a cofinal sequence $\Lambda'_{n'} \subset \mathbb{Z}^d$ indexed by n' [not necessarily the sequence $(\Lambda_n)_n$ that might be used in defining allowed configurations]. Given J , we construct a finite-size Gibbs state $\Gamma^{n'} = p_{H_{\Lambda'_{n'}}}(s|_{\Lambda'_{n'}})$, where $H_\Lambda = -\sum_{X \subseteq \Lambda} J_X s_X$ for any $\Lambda \in$

\mathcal{X} , and $T \geq 0$. $\Gamma^{n'}$ is a probability distribution on $s|_{\Lambda'_{n'}}$, and depends only on J in $J|_{\mathcal{X}(\Lambda'_{n'})}$; it can be viewed as a consistent system of marginal distributions $\Gamma^{n'} = (\Gamma^{n'}(s|_{\Lambda}))_{(\Lambda, s|_{\Lambda})}$ for $s|_{\Lambda}$, for all finite subsets $\Lambda \subseteq \Lambda'_{n'}$. Then for the joint distribution of the bonds and the state [that is, for the distribution of the pair (J, Γ)], by tightness and sequential compactness arguments [24,25,49], there is a subsequence n'_k of n' such that all the final-dimensional marginal distributions for $(J|_{\mathcal{X}(\Lambda'_{n'_k})}, \Gamma^{n'_k})$ have $k \rightarrow \infty$ limits that form a consistent family of marginals; these define a joint probability distribution κ^{\dagger} on (J, Γ) , where $\Gamma(s)$ (or $\Gamma_J(s)$ to show its dependence on J) is a probability distribution for s in infinite size [49]. [That is, $(J|_{\mathcal{X}(\Lambda'_{n'_k})}, \Gamma^{n'_k}) \rightarrow (J, \Gamma)$ as $k \rightarrow \infty$ in the sense of convergence in distribution, also known as convergence in law, or as vague or (in functional analysis) weak* convergence of the joint probability distributions for these random variables. The definition of weak* convergence for probability distributions $p_n(y)$, $p(y)$ on a space Y is that $p_n \rightarrow p$ in the weak* sense as $n \rightarrow \infty$ if for all bounded continuous functions $f(y)$, $\int_Y f p_n \rightarrow \int_Y f p$. In the present case there is no need for a proof of tightness, because the space of probability distributions $\Gamma(s)$ on S is compact, while the marginal $\nu(J)$ for the bonds is certainly tight.] The marginal distribution for J is the original $\nu(J)$, and the conditional distribution on Γ given J is the AW metastate κ (or κ_J to show its dependence on J), or $\kappa(\Gamma)$ if we (probably inaccurately) imagine that it can be described by a probability density on states Γ . In other words, we can draw from the metastate κ a (random) distribution or state $\Gamma(s)$ for the spins s .

In this step, for the case of the finite-range mixed p -spin models there is almost no change in the argument, compared with the references; we note that the set \mathcal{X} of all finite subsets of \mathbb{Z}^d , which indexes J , is countable, so that the spaces involved remain separable as $n' \rightarrow \infty$, allowing the argument to go through. As mentioned already, κ^{\dagger} and κ may not be unique: in addition to the choice of the sequence $(\Lambda'_{n'})_{n'}$ of finite sizes, a choice of a subsequence may have been required in order to obtain a limit. For notation, we will write formally $\kappa^{\dagger} = \nu\kappa$, which when the distributions can be represented by densities can be interpreted literally as $\kappa^{\dagger}(J, \Gamma) = \nu(J)\kappa_J(\Gamma_J)$, and similarly when we also use Γ to take expectations over s we write $\kappa^{\dagger}\Gamma$ which (for densities) means $\kappa^{\dagger}(J, \Gamma)\Gamma(s) = \nu(J)\kappa_J(\Gamma_J)\Gamma_J(s)$. Conditioning that on J gives $\kappa_J(\Gamma_J)\Gamma_J(s)$, which summed over Γ gives by definition $\rho_J(s)$, the metastate-average state (MAS). This completes the construction of a metastate.

4. Proof that the states are Gibbs states

Next, we want to argue that the states Γ (and not only the finite-size $\Gamma^{n'}$) drawn from a metastate are Gibbs states in the sense of Definition 1; this is where, in comparison with Refs. [11,50,51], additional work is required. We recall that a cofinal sequence $(\Lambda_n)_n$ will be used; its form, which will depend on the distribution $\nu(J)$ (and is arbitrary for the short-range case) but will be independent of J , Γ , and s , and of the construction of κ^{\dagger} [i.e., it is not necessarily the same as $(\Lambda'_{n'})_{n'}$ used in the construction of κ^{\dagger}], will be determined in the course of the proof.

Theorem 1: In the finite-range mixed p -spin models with a translation-invariant distribution (as defined in Secs. II and III A 1), for $T > 0$ a state Γ drawn from a metastate κ is a Gibbs state (as in Definition 1), κ^{\dagger} -almost surely.

Proof: The proof of Theorem 1 will take the remainder of this subsection, and follows immediately from the forthcoming Propositions 1, 2, and 3; in those propositions, the hypotheses of Theorem 1 remain in force. QED.

Proposition 1: For each finite Λ , the limit $\lim_{\Lambda' \rightarrow \infty} \Gamma(s|_{\Lambda} | s|_{\Lambda'}) = \Gamma(s|_{\Lambda} | s|_{\Lambda^c})$ exists $\kappa^{\dagger}\Gamma$ -almost surely.

Proof: (The proof is essentially standard [11,50,51], but included for completeness.) First, for $T \geq 0$, we consider the conditional probabilities $\Gamma(s|_{\Lambda} | s|_{\Lambda' - \Lambda})$ where Λ, Λ' are both finite, and $\Lambda \subseteq \Lambda'$. [Strictly, one should say the conditional probabilities $\Gamma(s|_{\Lambda} | \mathcal{F}_{\Lambda' - \Lambda})$, where $\mathcal{F}_{\Lambda' - \Lambda}$ is the σ algebra generated by $s|_{\Lambda' - \Lambda}$.] We note that, like any conditional probabilities, these are viewed as random variables, due to their dependence on $s|_{\Lambda^c}$, with distribution induced from $\Gamma(s)$. We fix Λ and let Λ' increase along any cofinal sequence $(\Lambda_n)_n$; then define $P_n = (\Gamma(s|_{\Lambda} | s|_{\Lambda_n - \Lambda}))_{s|_{\Lambda}}$, which is a sequence of random vectors, where each vector has components indexed by $s|_{\Lambda}$. (For m -vector models, this is better viewed as a regular conditional probability distribution on $s|_{\Lambda}$ [25].) Then the standard properties of conditional expectation [24,25] imply that we have

$$E(P_{n+1} | s|_{\Lambda_n - \Lambda}) = P_n; \quad (A12)$$

in other words, taking the Γ average over the spins in $\Lambda_{n+1} - \Lambda_n$ just produces the conditional probability conditioned on $s|_{\Lambda_n - \Lambda}$. This means the conditional probabilities P_n form a martingale. (More formally, letting $\mathcal{F}_n = \mathcal{F}_{\Lambda_n - \Lambda}$, the σ algebras form a filtration $\mathcal{F}_n \subseteq \mathcal{F}_{n+1}$, used to define the martingale; as $n \rightarrow \infty$, $\mathcal{F}_n \rightarrow \mathcal{F}_{\infty}$, the σ algebra generated by $s|_{\Lambda^c}$ [24,25].) The martingale convergence theorem [24,25] then tells us that, as $n \rightarrow \infty$, the conditional probabilities P_n tend Γ -almost surely to a limit $P = (\Gamma(s|_{\Lambda} | s|_{\Lambda^c}))_{s|_{\Lambda}}$. As the notation suggests, the limit is independent of the sequence $(\Lambda_n)_n$ (or filtration \mathcal{F}_n) that was used in this part of the construction. This concludes the proof of Proposition 1. QED.

By Proposition 1, the conditional probabilities have limits as $\Lambda' \rightarrow \infty$, and hence $\tilde{h}_{\Lambda}(s)/T$ is well defined as an extended real number, and in fact for $T > 0$ $\tilde{h}_{\Lambda}(s) > -\infty$ $\kappa^{\dagger}\Gamma$ -almost surely, as explained already.

Next we will prove the following.

Proposition 2: For $T > 0$ and with $\kappa^{\dagger}\Gamma$ probability 1, $\tilde{h}_{\Lambda}(s) < \infty$.

Proof: First take the average

$$E_{\kappa^{\dagger}} E \ln \frac{\Gamma(s)}{\Gamma(s')} = E_{\kappa^{\dagger}} D[\Gamma(s) || \Gamma_{\theta_{\Lambda}}(s)] \quad (A13)$$

of the logarithm of the ratio of probability distributions, or of corresponding conditional probabilities, which give the same result when conditioned on $s|_{\Lambda^c}$. (Here and in similar expressions below we abuse notation slightly on the left-hand side; it is an average over, not a function of, s, s' , just as the right-hand side is.) We have identified this with a relative entropy, as follows. If we consider the transformation $s \rightarrow s'$, where s_i and s'_i differ only in the finite set Λ , and express it as $s' = \theta_{\Lambda}s$ where θ_{Λ} stands for the operation of reversing all spins in Λ (or more general rotations for m -vector models and so on), we

can instead consider this as a change in the state from $\Gamma(s)$ to $\Gamma_{\theta_\Lambda}(s) = \Gamma(s')$. (If Γ is a Gibbs state, then Γ_{θ_Λ} is a Gibbs state in which H has been transformed to $H_{\theta_\Lambda}(s) = H(\theta_{\Lambda^c})$. This uses the symmetry property of the uniform reference measure on S_0 at each site that was mentioned in Sec. II; no symmetry of H is assumed. The symmetry also implies that the following marginals are equal, namely $\Gamma_{\theta_\Lambda}(s|_{\Lambda^c}) = \Gamma(s|_{\Lambda^c})$, and hence it makes no difference whether we use the relative entropy of the distributions or of their conditionals on $s|_{\Lambda^c}$ or a subset thereof.) Then the above E expectation is indeed, as displayed, the relative entropy of Γ relative to Γ_{θ_Λ} , with expectation using κ^\dagger over the states Γ and bonds J . This shows at once that $E\tilde{h}(s)/T \geq 0$, including for the case $T = 0$. Further, it can be shown [30] that, in the thermal average $E \ln[\Gamma(s)/\Gamma_{\theta_\Lambda}(s)]$, the contribution of the negative part of $\ln[\Gamma(s)/\Gamma_{\theta_\Lambda}(s)]$ is $> -\infty$, which in particular again implies that $h_\Lambda(s)/T > -\infty \kappa^\dagger \Gamma$ -almost surely.

To arrive at an upper bound, we first recall the results for relative entropy of the Soviet authors, who included Gelfand, Kolmogorov, Yaglom, Perez, Dobrushin, and Pinsker, as described in Pinsker's book [30]. The relative entropy can be defined in general by partitioning configuration space S into regions, finding the probability assigned to each region by each of the two probability distributions, and calculating the relative entropy for a partition using the formulas for discrete probabilities; then the relative entropy is the supremum, over partitions, of the relative entropy for a partition. In the present case we can define partitions by assigning each s to the corresponding $s|_{\Lambda'}$, which indexes the regions, giving the relative entropies $E \ln[\Gamma(s|_{\Lambda'})/\Gamma_{\theta_\Lambda}(s|_{\Lambda'})]$. (For continuous spins, the latter relative entropies must themselves be defined as the supremum over partitions of $S_{\Lambda'}$ into a discrete set.) As $\Lambda' \rightarrow \infty$, the relative entropy associated to such a partition tends to the one we need. In fact, for all Λ' , $E \ln[\Gamma(s|_{\Lambda'})/\Gamma_{\theta_\Lambda}(s|_{\Lambda'})]$ is nonnegative and increases monotonically to the limit [30]. (The monotonicity for relative entropy follows from the chain rule, similarly to the argument in Sec. III C 2.) These facts allow the use of the monotone convergence theorem [52], so the limit can instead be taken after the κ^\dagger expectation. Thus we have

$$E_{\kappa^\dagger} E \ln \frac{\Gamma(s)}{\Gamma_{\theta_\Lambda}(s)} = E_{\kappa^\dagger} \lim_{\Lambda' \rightarrow \infty} E \ln \frac{\Gamma(s|_{\Lambda'})}{\Gamma_{\theta_\Lambda}(s|_{\Lambda'})} \quad (A14)$$

$$= \lim_{\Lambda' \rightarrow \infty} E_{\kappa^\dagger} E \ln \frac{\Gamma(s|_{\Lambda'})}{\Gamma_{\theta_\Lambda}(s|_{\Lambda'})}. \quad (A15)$$

Here the first equality is from Ref. [30], and the second is from the monotone convergence theorem.

Now, using the second line, the expectation of the relative entropy for finite Λ' can be upper bounded, using finite-size systems. For fixed Λ' , the thermal E average is a continuous function of Γ , and the κ^\dagger expectation of that function calculated in finite size (i.e., its disorder average) tends to the infinite size version (κ^\dagger average) by definition of convergence in distribution. In a finite system on $\Lambda'_{n'}$, where $\Lambda \subseteq \Lambda' \subseteq \Lambda'_{n'}$, we have, using notation from Sec. III C 3 and

this Appendix (here $E^{n'}$ is thermal expectation in $\Gamma^{n'}$),

$$\begin{aligned} E E^{n'} \ln \frac{\Gamma^{n'}(s|_{\Lambda'})}{\Gamma^{n'}(s'|_{\Lambda'})} \\ = \frac{1}{T} E E^{n'} \{ -[F_{H-H_{\Lambda'}}(s|_{\Lambda'}) - F_{H-H_{\Lambda'}}(s'|_{\Lambda'})] \\ + 2h_\Lambda(s|_{\Lambda'}) \}. \end{aligned} \quad (A16)$$

The free energy difference would be zero if all interactions that involve spins in both Λ and Λ'^c were set to zero. Then by methods similar to those for finite-size bounds in Sec. III C 3, in particular the bound (3.60), we can obtain an upper bound at given n' , which under the convergence condition (2.5) is finite, including for $n' \rightarrow \infty$.

Next we wish to pass to the $n' \rightarrow \infty$ limit, in which the marginal states $\Gamma^{n'}(s|_{\Lambda'})$, $\Gamma^{n'}(s'|_{\Lambda'})$ converge in distribution (weak* convergence of their distribution) to limits $\Gamma(s|_{\Lambda'})$, $\Gamma(s'|_{\Lambda'})$, and we would like the $n' \rightarrow \infty$ limit of the bound to still hold for the expected relative entropy of $\Gamma(s|_{\Lambda'})$, $\Gamma(s'|_{\Lambda'})$. Such a limit of relative entropy under weak* convergence may not even exist, even when the probabilities are discrete as for the case of Ising spins. However, we can make progress by first using the following general result [53]: on a fixed complete separable metric space, the relative entropy of probability distributions is “jointly lower semicontinuous” under weak* convergence, that is, if as $n' \rightarrow \infty$, $P_{n'} \rightarrow P$, $Q_{n'} \rightarrow Q$ in the sense of weak* convergence, then the relative entropy obeys

$$D[P||Q] \leq \liminf_{n' \rightarrow \infty} D[P_{n'}||Q_{n'}] \quad (A17)$$

(see Ref. [52] for \liminf and for semicontinuity). Thus the relative entropy of a pair of states, such as $\Gamma^{n'}(s|_{\Lambda'})$, $\Gamma^{n'}(s'|_{\Lambda'})$, is a lower semicontinuous function on pairs of states in the appropriate topology of the space of such pairs. Next, any lower semicontinuous function f on a compact metric space can be obtained as the pointwise limit of a pointwise-increasing sequence of bounded continuous functions f_n , that is $f_{n+1} \geq f_n$ for all n . Suppose then that $(f_n)_n$ is such an increasing sequence of bounded continuous functions of a pair of marginal probability distributions on $S_{\Lambda'}$ (f_n independent of n' for all n) that tends to the relative entropy; as relative entropy is non-negative, we can also assume that $f_n \geq 0$ for all n . First applying f_n to $\Gamma^{n'}(s|_{\Lambda'})$, $\Gamma^{n'}(s'|_{\Lambda'})$, we take $\lim_{n' \rightarrow \infty} E f_n = E_{\kappa^\dagger} f_n$ by weak* convergence (along the subsequence n'_κ as always), which obeys the same $n' \rightarrow \infty$ limit of the upper bound as the right-hand side of Eq. (A16), and then the $n \rightarrow \infty$ limit exists by the monotone convergence theorem, and gives the desired upper bound on $E_{\kappa^\dagger} D[\Gamma(s|_{\Lambda'})||\Gamma_{\theta_\Lambda}(s|_{\Lambda'})]$. Then taking the $\Lambda' \rightarrow \infty$ limit, the expression is bounded above by $C_2 T^{-2} \sum_{X: X \cap \Lambda \neq \emptyset} \text{Var } J_X$ for a constant $C_2 > 0$ (the terms in the bound involving sites in both Λ , Λ'^c go to zero in this limit), which for $T > 0$ is finite. We conclude that $E_{\kappa^\dagger} E \tilde{h}_\Lambda(s) < \infty$.

As $E h(s) \geq 0$, the fact that $E_{\kappa^\dagger} E \tilde{h}_\Lambda(s)$ is finite implies that $E \tilde{h}(s) < \infty \kappa^\dagger$ -almost surely. (This result can also be viewed as saying that Γ is absolutely continuous with respect to Γ_{θ_Λ} , κ^\dagger -almost surely, and also holds *vice versa* similarly.) Then as the contribution of the negative part of $\ln[\Gamma(s)/\Gamma_{\theta_\Lambda}(s)]$ is $> -\infty$, the contribution of the positive part must be $< +\infty$, and

so, for $T > 0$, $\tilde{h}_\Lambda(s) < \infty$, $\kappa^\dagger\Gamma$ -almost surely. This concludes the proof of Proposition 2. QED.

In the proof of the following Proposition 3, we will use the notion of an extended metastate. The meaning of extending a metastate is as follows. It is sometimes desirable to carry additional information, other than the J and the state $\Gamma^{n'}$, from finite size systems through to the limit of infinite size. From the probability distributions in finite size, this is carried out as convergence in distribution starting from finite-dimensional marginal distributions, exactly as for the basic construction of κ^\dagger (in either AW or NS versions). If some κ^\dagger is already given, then a (sub-) sequence $(n'_k)_k$ was used, and (assuming tightness holds for the distributions of the additional data) one can find a subsequence of $(n'_k)_k$ along which the distribution on the higher-dimensional data, as well as on J , Γ , converges, to obtain an extended distribution $\kappa_{\text{ext}}^\dagger$, which conditioned on J gives an extended metastate κ_{ext} . As the sequence of finite sizes used is a subsequence of that used to obtain κ , the marginal distribution of the extended metastate on only the states Γ is the original metastate κ . It is in this sense that κ_{ext} is an extension of κ (or similarly for $\kappa_{\text{ext}}^\dagger$ and κ^\dagger). Our use of extended metastates is inspired by the excitation metastate used by NS, and also by the “natural” metastate used in Ref. [34] to handle problems similar to the present ones. The form of the extended metastate that we use here will be explained during the proof, but we emphasize that the statement of the proposition does not involve the extension.

Proposition 3: There is a cofinal sequence $(\Lambda_n)_n$ such that, for $T > 0$, for all finite Λ , and for $\kappa^\dagger\Gamma$ -almost every s ,

$$\tilde{h}_\Lambda(s) = \lim_{n \rightarrow \infty} h_\Lambda(s|_{\Lambda_n}). \quad (\text{A18})$$

Proof: The basic idea behind the proof is to compare $\ln \Gamma(s|_\Lambda | s|_{\Lambda^c})$ with the same quantity evaluated for a corresponding state $\Gamma_{(\Lambda, \Lambda_n)}$ in a system in which the interactions that connect spins in Λ with those in Λ_n^c are set to zero. In the long-range cases, that is, an infinite set of bonds, and in trying to bound the effect of that change by the methods of this paper, which involve expectation over those bonds, sooner or later one reaches an expectation of an infinite sum of terms corresponding to those bonds, and the use of bound (3.47) would be the next step. However, in general it is not clear if the expectation can be taken term by term. In order to render the series a finite sum, we can consider a finite size system (as was done in the proof of Proposition 2). Then we can use the metastate construction to take the limit, so that $\Gamma(s)$ is drawn from the metastate. However, it is not clear if the basic metastate construction implies that the state $\Gamma_{(\Lambda, \Lambda_n)}$ with truncated interactions also exists in the limit. To deal with that, we introduce the truncated-interaction metastate, an extension of the metastate κ , as follows.

In finite size Λ'_n , we construct for given J the state $\Gamma^{n'}$, and also the states $\Gamma_{(\Lambda, \Lambda')}^{n'}$ for pairs of finite sets Λ and Λ' , $\Lambda \subseteq \Lambda'$, in which all terms in the Hamiltonian that connect Λ and Λ'^c are dropped. Clearly, it is sufficient to do this for Λ and $\Lambda' \subseteq \Lambda'_n$. It would be sufficient to use for Λ' only the members of the eventual cofinal sequence Λ_n (which will be characterized later), but use of all Λ' (eventually without restriction, except $\Lambda \subseteq \Lambda'$) would work just as well. We then take the limit in distribution as $n' \rightarrow \infty$ of the joint distribution of $J|_{\mathcal{X}(\Lambda'_n)}$,

$\Gamma^{n'}$, and the array $(\Gamma_{(\Lambda, \Lambda')}^{n'})_{(\Lambda, \Lambda')}$ in the same way as was done to obtain κ^\dagger , by choosing a subsequence of the sequence $(n'_k)_k$ used there, and call the resulting distribution $\kappa_{\text{ext}}^\dagger$. As the finite subsets of \mathbf{Z}^d are countable, the space of all these variables is again separable, allowing the procedure to work. We denote the resulting truncated-interaction metastate by κ_{ext} .

Then in finite size, we define

$$\Delta = E \ln \frac{\Gamma^{n'}(s)}{\Gamma_{(\Lambda, \Lambda')}^{n'}(s)} \quad (\text{A19})$$

$(s = s|_{\Lambda'_n})$ which, being a relative entropy, is nonnegative. In the same way as before, we can derive bounds

$$E\Delta \leq \frac{C_3}{T^2} \sum_{X \in \mathcal{X}: X \cap \Lambda \neq \emptyset, X \cap \Lambda'^c \neq \emptyset} \text{Var } J_X, \quad (\text{A20})$$

in which we took $n' \rightarrow \infty$ on the right-hand side, and $C_3 > 0$ is a constant. As these bounds are finite under the convergence condition (2.5), and independent of n' , they also apply to the limit, and so hold for the $\kappa_{\text{ext}}^\dagger$ expectation of the relative entropy of the states in infinite size, which are denoted Γ , $\Gamma_{(\Lambda, \Lambda')}$. To be more careful here, we can instead begin with the relative entropy of the two marginal distributions for $s|_{\Lambda''}$ for $\Lambda'' \subseteq \Lambda'_n$ in finite size, which is less than that above as discussed in the proof of Proposition 2, so obeys the same bound. Let $n' \rightarrow \infty$ with Λ'' fixed (the lower semicontinuity property must be used again here; see the proof of Proposition 2), and then $\Lambda'' \rightarrow \infty$ which, again as discussed in the proof of Proposition 2, gives the ($\kappa_{\text{ext}}^\dagger$ expectation of the) desired relative entropy of Γ , $\Gamma_{(\Lambda, \Lambda')}$, and the upper bound again still holds. A similar bound applies when the thermal average is taken in $\Gamma_{(\Lambda, \Lambda')}$, and also for the relative entropy of Γ_{θ_Λ} , $\Gamma_{(\Lambda, \Lambda')\theta_\Lambda}$. The finite upper bound implies that, $\kappa_{\text{ext}}^\dagger$ -almost surely, the two states are mutually absolutely continuous for any pair $\Lambda \subseteq \Lambda'$. (We can do the same for the relative entropy of $\Gamma_{(\Lambda, \Lambda')}$ relative to $\Gamma_{(\Lambda, \Lambda'')}$, so any such pair are also mutually absolutely continuous.)

Call the right-hand side of the upper bound V' . We notice that

$$V' \leq \sum_{i \in \Lambda} V'_i, \quad (\text{A21})$$

where V'_i is V' for $\Lambda = \{i\}$. Hence $V' \rightarrow 0$ as $\Lambda' \rightarrow \infty$ (through some cofinal sequence) for all Λ if and only if the convergence condition (2.5) holds. As relative entropy is nonnegative, this implies that the $L^1(\kappa_{\text{ext}}^\dagger)$ -norm of the relative entropy tends to zero, and hence the latter also tends to zero in probability.

Next we will show that there is a sequence of sets Λ' along which the relative entropy tends to zero $\kappa_{\text{ext}}^\dagger$ -almost surely. This subsequence is what we will take for Λ_n . We will describe such a sequence in detail to show that its properties are independent of Λ . We use hypercubes Λ_L of side L , L an odd integer, centered at the same point in \mathbf{Z}^d for all L ; this sequence is clearly cofinal. Here we will take the center of all the cubes to be $i = 0$, the origin ($\mathbf{x}_0 = \mathbf{0}$). Consider the sequences of $V'_i = V'_i(L)$ for each i , where $\Lambda' = \Lambda_L$, indexed by L increasing through the odd integers; each sequence $(V'_i(L))_L$ decreases to zero. We intend to use V'_0 to define the sequence

$(\Lambda_n)_n$. We need to show that $[V'_i(L) - V'_0(L)]/V'_0 \rightarrow 0$ as $L \rightarrow \infty$. This is straightforward, using the invariance of $\text{Var } J_X$ under translations of X . We can show that $|V'_i(L) - V'_0(L)| \leq V'_0(L - |\mathbf{x}_i|) - V'_0(L + |\mathbf{x}_i|)$, and as Λ is fixed, $|\mathbf{x}_i|$ is fixed and bounded for all $i \in \Lambda$. From basic calculus the last difference is eventually much smaller than $V'_0(L)$. Then V' is less than of order $|\Lambda|V'_0$, or we can say that for any $\varepsilon > 0$, $V'(L) \leq (1 + \varepsilon)|\Lambda|V'_0(L)$ for sufficiently large L (which depends on Λ and ε).

Now we use (the simplest form of) Chebyshev's inequality, that is, if X is a nonnegative random variable and $t > 0$, then $\mathbf{P}[X \geq t] \leq \mathbf{E}[X]/t$; we apply this to the probability $\mathbf{P}_{\kappa_{\text{ext}}^{\dagger}}$ in the truncated-interaction metastate. From this we obtain that, for any $\varepsilon \geq 0$, for sufficiently large L

$$\mathbf{P}_{\kappa_{\text{ext}}^{\dagger}}[\Delta(L) \geq t] \leq \frac{(1 + \varepsilon)|\Lambda|V'_0}{t}. \quad (\text{A22})$$

Now, for each n , define L_n to be the smallest L such that $V'_0(L) \leq 1/n^2$ for all $L \geq L_n$, and $\Lambda_n = \Lambda_{L_n}$. Then

$$\sum_n \mathbf{P}_{\kappa_{\text{ext}}^{\dagger}}[\Delta(L_n) \geq t] \leq \frac{(1 + \varepsilon)|\Lambda| \sum_n \frac{1}{n^2}}{t} \quad (\text{A23})$$

(at least for the tail of the sum at large n on both sides), and the sum converges, so

$$\sum_n \mathbf{P}_{\kappa_{\text{ext}}^{\dagger}}[\Delta(L_n) \geq t] < \infty. \quad (\text{A24})$$

Hence for any Λ , by the Borel-Cantelli lemma (see Ref. [24], p. 77), the $\kappa_{\text{ext}}^{\dagger}$ probability that $\Delta(L_n) \geq t$ for infinitely many $n > 0$ is zero, and as $t > 0$ is arbitrary it follows that $\Delta(L_n) \rightarrow 0$ as $n \rightarrow \infty$, $\kappa_{\text{ext}}^{\dagger}$ -almost surely.

Relative entropy that tends to zero implies that $\Gamma_{(\Lambda, \Lambda_n)} \rightarrow \Gamma$ as $n \rightarrow \infty$ in a sense we now describe. Relative entropy obeys the Csiszár-Kullback-Kemperman inequality for two probability distributions P_1, P_2 (defined on the same σ algebra on the same space S),

$$d_{\text{TV}}(P_1, P_2)^2 \leq 2D[P_1 || P_2], \quad (\text{A25})$$

where the total variation (or L^1 -) distance (i.e., metric) between any two signed measures P_1, P_2 is defined as

$$d_{\text{TV}}(P_1, P_2) = \int_S |dP_1 - dP_2| \quad (\text{A26})$$

(see, e.g., Ref. [29] for a simple proof; some authors define the total variation distance to be half of ours). The total variation $d_{\text{TV}}(P, 0)$ is the norm on the space of finite signed measures P that arises when that space is viewed as the dual of the space $C(S)$ of bounded continuous functions on the compact configuration space S . It defines a topology on the space of probability measures on S that is (much) stronger than the weak* topology. Convergence in d_{TV} implies that the probability distributions become equal in the limit. Thus we have $\Gamma_{(\Lambda, \Lambda_n)} \rightarrow \Gamma$ in total variation distance (and hence also in the weak* topology, so that correlation functions also converge), $\kappa_{\text{ext}}^{\dagger}$ -almost surely.

Then, from Proposition 2 and its analog for $\Gamma_{(\Lambda, \Lambda')}$, we have mutual absolute continuity of all the probability distributions $\Gamma, \Gamma_{(\Lambda, \Lambda_n)}, \Gamma_{\theta_\Lambda}, \Gamma_{(\Lambda, \Lambda_n)\theta_\Lambda}$ for all $n, \kappa_{\text{ext}}^{\dagger}$ -almost surely. We note that, while absolute continuity is not a symmetric

relation, mutual absolute continuity is an equivalence relation, and distributions that are mutually absolutely continuous have exactly the same null sets; therefore, we can here say simply $\kappa_{\text{ext}}^{\dagger}\Gamma$ -almost surely, regardless of which of these distributions we are using. Then we can take the Radon-Nikodym (RN) derivative [52], for example $d\Gamma/d\Gamma_{\theta_\Lambda}$ of Γ with respect to Γ_{θ_Λ} , which is a function of s that, because of absolute continuity, exists $\kappa_{\text{ext}}^{\dagger}\Gamma$ -almost surely; this is the true meaning of the ratio $\Gamma(s)/\Gamma_{\theta_\Lambda}(s)$ in the definition of the relative entropy, when the two distributions are mutually absolutely continuous [30]. We do the same with $\Gamma_{(\Lambda, \Lambda_n)}$ with respect to $\Gamma_{(\Lambda, \Lambda_n)\theta_\Lambda}$. For a collection of mutually absolutely continuous measures, the RN derivatives of one with respect to another enjoy properties similar to those of derivatives of a collection of functions with respect to one another, in particular the chain rule [52]. As $\Gamma_{(\Lambda, \Lambda_n)} \rightarrow \Gamma$ and $\Gamma_{(\Lambda, \Lambda_n)\theta_\Lambda} \rightarrow \Gamma_{\theta_\Lambda}$ as probability measures, the RN derivatives tend to the $n \rightarrow \infty$ limit $\kappa_{\text{ext}}^{\dagger}\Gamma$ -almost surely, for example, $d\Gamma_{(\Lambda, \Lambda_n)}/d\Gamma \rightarrow 1$. Taking logarithms, we have

$$\ln \frac{d\Gamma_{(\Lambda, \Lambda_n)}}{d\Gamma_{(\Lambda, \Lambda_n)\theta_\Lambda}} = 2h_\Lambda(s|_{\Lambda_n})/T, \quad (\text{A27})$$

and from the definition above,

$$\ln \frac{d\Gamma}{d\Gamma_{\theta_\Lambda}} = 2\tilde{h}_\Lambda(s)/T. \quad (\text{A28})$$

It follows that $h_\Lambda(s|_{\Lambda_n}) \rightarrow \tilde{h}_\Lambda(s)$ as $n \rightarrow \infty$, $\kappa_{\text{ext}}^{\dagger}\Gamma$ -almost surely and for all Λ . As the final statement does not refer to $\Gamma_{(\Lambda, \Lambda_n)}$, it in fact holds $\kappa^{\dagger}\Gamma$ -almost surely; the proof is complete. QED.

We note that, once a metastate κ has been obtained, a $\Gamma(s)$ drawn from it is well defined as a (random) state (a distribution on S). Only the characterization of the allowed configurations, and of Γ as a *Gibbs* state, involved the sequence $(\Lambda_n)_n$. The proof of Proposition 3 shows that there is considerable leeway in the choice of the sequence, for example the choice of the common center for the hypercubes is arbitrary, or the hypercubes could be replaced by other compact shapes, and so on. A configuration s that is allowed under the definition using one sequence will, with Γ probability one, also be allowed under the definition using *any* other sequence such that the steps in the proof can be carried through.

5. Reformulation and pure state decomposition

The GNS definition of a Gibbs state seems acceptable, but if we also wish to use some notion of pure Gibbs states, which preferably should possess the same properties as in the short-range cases, then some reformulation is required. The usual description of a Gibbs state [1–3] begins from the notion of a family of so-called specifications, which normally correspond to the same family of conditional probabilities $\Gamma(s|_{\Lambda} | s|_{\Lambda^c})$ (for Λ finite) with which we began here. The specifications $\gamma = (\gamma_\Lambda(A | s|_{\Lambda^c}))_\Lambda$ should for each Λ be defined for all values of $s|_{\Lambda^c}$; here in most general form, these are probability kernels, so they are both a probability distribution on sets A of s , and a measurable function of $s|_{\Lambda^c}$, and are assumed to be proper (see Ref. [1], Chapter 1). They are defined *a priori*, meaning without specifying an unconditional Γ first; they are independent of such Γ . They can be defined abstractly as

possessing the properties of conditional probabilities, and to be Gibbsian they should be related to the Hamiltonian $H'_\Lambda(s)$. (Somewhat similar probability kernels, but conditioned on s , also arise as the transition probabilities for the dynamics of our systems [4].) Then the conventional definition is

Definition 2: A Gibbs state admitted by a specification γ is a probability distribution on S whose conditional probabilities for each finite Λ are

$$\Gamma(A \mid s|_{\Lambda^c}) = \gamma_\Lambda(A \mid s|_{\Lambda^c}) \quad (\text{A29})$$

Γ -almost surely.

By properties of conditional probabilities, this also implies that a Gibbs state is mapped to itself by each of a family (indexed by Λ) of linear maps of measures defined by the specifications (the DLR equations) [1–3]:

$$\Gamma(A) = \int \gamma_\Lambda(A \mid s|_{\Lambda^c}) \Gamma(s) ds \quad (\text{A30})$$

for all A and Λ . (In dynamics, corresponding conditions involving the transition probabilities say that Γ is a stationary state [4].) Now to handle the long-range case, we only have to define the specifications to be used. Our proposal is to use the (generalized) specification, which in spirit is Gibbsian, defined for all finite Λ and all s by

$$\gamma_\Lambda(s|_\Lambda \mid s|_{\Lambda^c}) = \begin{cases} p_{H'}(s|_\Lambda) & \text{if } s \text{ is allowed,} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A31})$$

We use the term “generalized” here (but will drop it hereafter) because if s is not allowed, γ_Λ is not in fact a probability; it gives zero for any set of $s|_\Lambda$! However, the set on which it is zero is a tail event. Thus, rather than modifying the definition of a Gibbs state, we have extended the definition of what can serve as a specification. We remark that (a) once again, the definition makes sense because of the definition of allowed s ; (b) this definition is consistent with the conditional probabilities of the Gibbs states under part 2 of Definition 1, because for any set of configurations that has zero probability, the conditional can be defined arbitrarily; (c) conversely, by Eqs. (A30), Gibbs states under the current definition assign probability 1 to the set of allowed s , as required by part 1 of Definition 1. Hence, for these specifications, Definition 2 is equivalent to Definition 1.

A conventional route to proving the existence of a decomposition of any Gibbs state into a mixture of pure (i.e., extremal) states involves showing that the Gibbs states admitted by the specification γ form a set $\mathcal{G}(\gamma)$, which as a subset of the space of probability measures on S (with, say, the weak* topology) is clearly convex, and also closed. Closure is guaranteed if the maps of measures defined by the specification are continuous. That then gives a compact convex set of probability measures, and the Choquet theory of such sets [54] leads to the desired results. In the present case, for the long-range models, continuity of the above γ is not obvious. [Nor is it “quasilocal”; see remark (2.22) in Ref. [1].] However, Georgii’s book [1], in particular Secs. 7.1 and 7.3, leading up to Theorem 7.26, characterizes the pure states and gives results of Dynkin and Föllmer that establish that any Gibbs state has a unique decomposition into pure states, without the use of any such topological properties of γ . That decomposition is then the starting point for the analysis

of the complexity of Gibbs states discussed in this paper. In the present case, there are some s for which γ fails to be a probability distribution, but one can verify that the proofs of the main results from Chapter 7 of Georgii [1] still go through in this case with only minor modifications (in particular, formulas involving γ that hold Γ -almost surely for Gibbs states Γ are unaffected).

Finally, we should note that Ref. [23] uses a specification defined on the set of the allowed configurations s only, and shows for any $T > 0$ that any state obtained as a limit from finite size (with boundary conditions) must satisfy the DLR equations (A30) (however, a proof there that the state puts measure 1 on the set of allowed s holds only at sufficiently large T). This does not seem to be sufficient for our purposes.

APPENDIX B: UNIQUENESS OF GIBBS STATES AT $T > 0$ IN SHORT-RANGE CASE IN ONE DIMENSION

In this Appendix we provide a short and fairly simple proof that at $T > 0$ there is a unique Gibbs state (a pure state) in any short-range mixed p -spin SG model in dimension $d = 1$ that satisfies one simple condition. (Under the same conditions, this rules out a nontrivial metastate, and also rules out any phase transition that would imply a change in the number of pure states.) For the $p = 2$ power-law model of Ising spins, arguments for similar results were given in earlier work [10,31,55,56] (see also Ref. [57] for a similar result in the case of short-range m -vector models with $O(m)$ symmetry and $d \leq 2$). Reference [55] gives a complete proof for $\sigma > 3/2$. Following the proposal of Kotliar *et al.* that there would be no transition at $T > 0$ for $\sigma > 1$ [10], van Enter and van Hemmen [31] employed a simple approach based on relative entropy to show the absence of spontaneous breaking of spin-flip symmetry in that region, but their paper and Ref. [57] were criticized for some technical issues in Ref. [56]. The latter [56] employs a very different approach and arrives at a full proof for $\sigma > 1$ in a setup using fixed-spin boundary conditions, but the proof is rather long and some may find it difficult. We note that similar results for short-range nondisordered spin systems, such as Ising ferromagnets, are well-known folklore (and for strictly short-range cases can be proved easily using a transfer matrix) and were proved rigorously in Ref. [33]; see also Ref. [2] (p. 303) and references therein for a simpler proof. The strategy of our proof is to show that the relative entropy of two distinct pure states is bounded, which gives a contradiction; this is similar to that of Ref. [31], but we implement it in a form that avoids some technical questions, using an upper bound exactly like those elsewhere in this paper. The statement in the theorem is more general than in Refs. [31,55,56]; in particular, other than existence of the first two absolute moments, we do not use a condition on the tail of the distribution of J_X .

The statement is the following.

Theorem 2: Consider a short-range SG model (as defined in Secs. II and III A 1) in $d = 1$, with sites $\mathbf{x}_i = i \in \mathbf{Z}$, and $T > 0$. If the bonds satisfy

$$\sum_{X \in \mathcal{X}: X \cap \mathbf{Z}_- \neq \emptyset, X \cap \mathbf{Z}_+ \neq \emptyset} \text{Var } J_X < \infty \quad (\text{B1})$$

(where $\mathbf{Z}_- = \{i \leq 0\}$ and $\mathbf{Z}_+ = \{i > 0\}$), then there is at most one pure Gibbs state, ν -almost surely. (For m -vector models with $m > 1$, or other non-Ising cases, the sum ranges over x as well as X , with no additional conditions on x .)

Before starting the proof, we discuss some general points. First, as we consider only short-range cases, most of the technicalities of Appendix A will not be needed here; in particular, as we begin simply from pure states, the metastate construction is not required, though it is useful in that it establishes that some Gibbs states actually exist.

Second, the fact that the Gibbs states considered in the proof are pure is not used until the end of the proof. We will need the fact that distinct pure states are supported on disjoint sets of spin configurations [1], and so are mutually singular. (This may be physically obvious if the supports of the distinct pure states are viewed as sets of configurations that are mutually inaccessible in dynamics, i.e., as ergodic components.) This behavior is the extreme opposite of absolute continuity, and implies that the relative entropy of either with respect to the other must be infinite. If we examine the relative entropy of one with respect to the other for the marginal distributions in a finite region (window), then (for Gibbs states) the result will be finite, but it will increase monotonically (see proof of Proposition 2 in Appendix A) to infinity as the size of the region expands until it eventually includes all the sites (i.e., along a cofinal sequence). That is, for any choice of a bound M , the relative entropy must eventually be greater than M for all sufficiently large window sizes. Of course, in our case, the relative entropy will be a random variable, and the statements must be made probabilistically. We will also use in the next paragraph the fact that, in a pure state, connected thermal correlations of functions of spins in well-separated regions tend to zero as the separation goes to infinity (see Refs. [1–3] for the precise statement).

Third (some readers may prefer to skip this on a first reading), we will be interested in an expectation of a relative entropy, which is a thermal average, here in a pure state. The expectation involves two pure states as (some, at least, of) the J are varied, and the question may arise whether we can do this here: can we be sure that we obtain the “same” two pure states as some of the J_X are changed? (This is related to the concerns in Ref. [56], that construction of pure states might require J -dependent boundary conditions, preventing naive manipulations of an expectation over all J .) In the present case, we will be interested in the expectation \mathbf{E}' over only the bonds J_X such that X has nonempty intersection with a finite interval, say,

$$\Lambda_W = \{-(W-1)/2, -(W-1)/2 + 1, \dots, (W-1)/2\} \quad (\text{B2})$$

($W > 0$ odd); call that expectation \mathbf{E}'_W . Here we discuss this question in general, for any dimension d . If we begin with a given Gibbs state, say, Γ , for some given J , we can actually construct a corresponding Gibbs state Γ' with other values of some of the J_X , as follows. For a change in Hamiltonian ΔH , the Gibbsian formulation suggests a definition of a perturbed state, such that expectation of any function of s in Γ' is related to that in Γ by

$$\langle \dots \rangle_{\Gamma'} = \frac{\langle \dots e^{-\Delta H/T} \rangle_{\Gamma}}{\langle e^{-\Delta H/T} \rangle_{\Gamma}}, \quad (\text{B3})$$

where $\langle \dots \rangle_{\Gamma'}(\langle \dots \rangle_{\Gamma})$ denotes expectation in Γ' (respectively, Γ). If $\Delta H = -\sum_X \Delta J_X s_X$ has the form of a general mixed p -spin Hamiltonian, but includes nonzero terms for only a *finite* set of X , then the perturbed expectation can be expanded out in terms of averages in Γ (with coefficients based on those in ΔH) and the change in its value is finite. This fact extends to the case of ΔH containing an infinite number of terms, if ΔH also satisfies the absolute convergence condition for a short-range Hamiltonian like that which precedes Eq. (3.7). Namely, if we define

$$\|J\|_1 \equiv \sum_X |J_X| \quad (\text{B4})$$

[the l^1 norm on J ; in models other than Ising, the sum must range over the pairs (X, x) as before], then using $|\Delta H| \leq \|\Delta J\|_1$ a sufficient condition is that $\|\Delta J\|_1$ be finite. In that case, the infinite sum converges because the terms in its tail decrease sufficiently rapidly, and their effect on thermal averages converges also, by use of an easily proved inequality such as

$$\left| \frac{\langle f(s) e^{-\Delta H/T} \rangle_{\Gamma}}{\langle e^{-\Delta H/T} \rangle_{\Gamma}} - \langle f(s) \rangle_{\Gamma} \right| \leq (e^{2\|\Delta J\|_1/T} - 1) \sup_s |f(s)|, \quad (\text{B5})$$

where f is any function of s [cf. Eq. (A.1.6) of AW [11]; this topic forms part of their discussion of properties of metastates]. This means that (in models of the general form defined in Sec. II), Gibbs states (in the weak* topology) are continuous functions of J (in the topology determined by the norm $\|\dots\|_1$) [11]. Further, if Γ is a pure state, the asymptotic behavior of thermal averages is unaffected by such a change in the Hamiltonian, so Γ' is also pure. This is not difficult to show if ΔH contains only a finite number of terms, by using the asymptotic decay of connected correlations in Γ [1], and it extends to all cases in which $\|\Delta J\|_1 < \infty$ by a simple approximation argument. In effect, the change in the Hamiltonian is only a local one, without detrimental long-range effects. In the case of interest, $\Delta H \propto H'_{\Lambda_W}(s) - H_{\Lambda_W}(s)$, and $\|\Delta J\|_1$ is finite ν -almost surely for a short-range model. We will now adopt the corresponding perturbations of the original two pure states with which we began, and so view them as functions of the bonds in question.

Proof of Theorem 2: Suppose that there are two distinct pure Gibbs states Γ_{α} and Γ_{β} . We will bound their expected relative entropy. Consider their marginal distributions on, without loss of generality, Λ_W . Form the relative entropy

$$D_{\alpha\beta,W} = \mathbf{E}_{\alpha} \ln \frac{\Gamma_{\alpha}(s|_{\Lambda_W})}{\Gamma_{\beta}(s|_{\Lambda_W})} \quad (\text{B6})$$

of $\Gamma_{\alpha}(s|_{\Lambda_W})$ with respect to $\Gamma_{\beta}(s|_{\Lambda_W})$, and take the \mathbf{E}'_W expectation (for which, see the discussion before this proof). It can be bounded by the method that by now should be familiar, using the (formal) interpolating Hamiltonian $H + (\lambda - 1)(H'_{\Lambda_W} - H_{\Lambda_W})$ where λ runs from zero to one, and the perturbation of the states proportional to $\lambda - 1$ can be handled as explained before this proof; we call the resulting pure states $\Gamma_{\alpha}^{(\lambda)}$, $\Gamma_{\beta}^{(\lambda)}$. We note that, at $\lambda = 0$, $\ln \Gamma_{\alpha}^{(0)}(s|_{\Lambda_W})/\Gamma_{\beta}^{(0)}(s|_{\Lambda_W}) = 0$, as Λ_W is decoupled from the rest of the system. Hence

$\ln \Gamma_\alpha(s|_{\Lambda_W})/\Gamma_\beta(s|_{\Lambda_W})$ is equal to the integral $\int_0^1 d\lambda$ of

$$\frac{d}{d\lambda} \ln \Gamma_\alpha^{(\lambda)}(s|_{\Lambda_W}) = -\frac{1}{T} (\langle H'_{\Lambda_W} - H_{\Lambda_W} \rangle_{\alpha,s|_{\Lambda_W},\lambda} - \langle H'_{\Lambda_W} - H_{\Lambda_W} \rangle_{\alpha,\lambda}), \quad (B7)$$

minus the similar expression with β in place of α . [The notation here is similar to that in Sec. III C 3, though the (conditional) thermal expectations are taken using $\Gamma_\alpha^{(\lambda)}$, as indicated.] We have to be careful about taking the expectation (doing integrals) term by term on the sum, because the method of returning to a finite-size system is not available here. There are in fact up to three integrals or expectations (namely, those implied by \mathbf{E}'_W , \mathbf{E}_α , and $\int_0^1 d\lambda$), and also the sum. If we replace each term $J_X s_X$ (or in some places a conditional \mathbf{E}_α or \mathbf{E}_β average of such a term) by its absolute value, then the short-range condition (3.7) implies that the sum converges, and so by part of the Fubini-Tonelli Theorem [52] it is legitimate to carry out the integrals and sum in any order, and the result is finite. That proves that the integrand summand is integrable, and so the other part of Fubini-Tonelli tells us that the integrals and sum of the original series can be carried out in any order also. Then we obtain

$$\mathbf{E}'_W D_{\alpha\beta,W} \leq C \frac{1}{T^2} \sum_{X: X \cap \Lambda_W \neq \emptyset, X \cap \Lambda_W^c \neq \emptyset} \text{Var } J_X, \quad (B8)$$

where $C > 0$ is a constant, and of course also the same with α and β interchanged. (The m -vector models involve summation over x as well as X .) The upper bound increases with W . As $W \rightarrow \infty$, it is finite if and only if the hypothesis (B1), which arises from the contribution of each of the two ends of the interval, holds.

Going back to W finite, the same bound applies for each W if we take the conditional expectation over a larger set of bonds than those involved in \mathbf{E}'_W , say, those for X intersecting some finite Λ_n , $\Lambda_W \subseteq \Lambda_n$. As we take such a set larger and larger (i.e., $\Lambda_n \rightarrow \infty$ along a cofinal sequence), effectively removing the conditioning on more and more bonds, because of the upper bound (and positivity of the relative entropy) the backward martingale convergence theorem [24,25] tells us that, for each W , the limit of the conditional expectation exists as a random variable that is measurable with respect to the tail σ algebra of J , and that it obeys the same bound above. We now work in this limit, where expectations are the full \mathbf{E} , and similarly for the distribution ν (strictly speaking, they are still conditioned on the tail σ algebra of J).

If we consider the infinite sequence of relative entropies $D_{\alpha\beta,W}$ for all finite odd $W > 0$, and if the hypothesis (B1) holds, then using Chebyshev's inequality the family, indexed by W , of distributions (induced from ν) of $D_{\alpha\beta,W}$ for finite W is tight: no weight goes off to infinity as $W \rightarrow \infty$. Because $D_{\alpha\beta,W}$ is an increasing function of W , it tends either to a finite limit or to infinity as $W \rightarrow \infty$, so the preceding result shows that, ν -almost surely, it tends to a finite limit: the probability that the two states are mutually singular is zero. As the Gibbs states were assumed to be pure and distinct, this contradiction shows that any two pure states are in fact identical, so there is a unique pure state, and a unique Gibbs state. That concludes the proof. QED.

We comment that in the $p = 2$ $d = 1$ power-law models both the short-range condition and condition (B1) imply that $\sigma > 1$. In general, a comparison of the two conditions involves the dependence of $\text{Var } J_X$ on both $|X| = p$ and $\text{diam } X > 0$ (in one dimension, we can define the diameter of X as $\text{diam } X = \max\{i : i \in X\} - \min\{i : i \in X\}$).

An information-theoretic interpretation of the proof of Theorem 2 is that the short-range interactions in $d = 1$ are not strong enough to convey to a finite window Λ_W sufficient information concerning in which pure state the system is supposed to be; for pure states, that information (the relative entropy) would have to tend to infinity with W .

An alternative, slightly weaker, statement of the result is that, for any two pure states in the decompositions of two respective MASSs, they must be equal, implying the triviality and uniqueness of the metastate and of the Gibbs states drawn from it. In this case the proof can use the expectation $\mathbf{E}_{\nu(\kappa_1 w_1 \times \kappa_2 w_2)}$ over pairs of pure states in the decompositions of respective Gibbs states drawn from respective metastates κ_1 and κ_2 (where for $\tau = 1, 2$, $w_\tau = w_\alpha(\Gamma_\tau)$ denotes the decomposition of a Gibbs state Γ_τ , say, drawn from κ_τ , into pure states Γ_α for the given J), and over bonds.

We emphasize that the proof holds for pure states in any m -vector model, not only those produced by applying $O(m)$ symmetry with an $O(m)$ -invariant Hamiltonian, but also those produced by non- $O(m)$ -invariant Hamiltonians. If we wish only to show the nonexistence of $T > 0$ pure states that spontaneously break the global spin-flip symmetry [$O(1) = \mathbf{Z}_2$] in the Ising case, for a flip-invariant Hamiltonian that contains terms with even p only, then we can consider a pure state Γ_α and the state obtained by flipping all spins in Λ_W , as in Proposition 2 above, and closer to Ref. [31]. The argument is similar to before, and involves marginalizing to a region Λ' containing Λ_W , and taking $\Lambda' \rightarrow \infty$, so essentially we bound $\mathbf{E}\mathbf{E}_\alpha h_{\Lambda_W}(s)$ (see Appendix A), then finally take $W \rightarrow \infty$. In the resulting bound, the domain wall sum in condition (B1) is modified to include only terms indexed by X such that $|X \cap \mathbf{Z}_-|$ and $|X \cap \mathbf{Z}_+|$ are both odd, and there is no symmetry breaking if that sum is finite (thus when terms with $p > 2$ are present, this condition is slightly less restrictive than the direct application of that found above to the present models). The same sum is also found if we extend to mixed even- p -spin models the proof for $d = 1$ of an upper bound [34] that leads to a bound on the exponent θ in the scaling-droplet theory [27]; when the sum is finite, the scaling-droplet arguments predict that there is no transition that spontaneously breaks spin-flip symmetry at $T > 0$.

Very similar proofs as for Theorem 2 and for the absence of symmetry breaking at $T > 0$ hold for the case of nonrandom p -spin interactions [33] (without the need for any disorder average, of course), and reproduce the well-known classic result for only $p = 2$ interactions, which has $|J_X|$ in place of $\text{Var } J_X$ in the domain-wall sum, and which is then indeed a bound on the energy of a domain wall in the ground state. For the case of proving the absence of spontaneous breaking of continuous symmetry in short-range SGs in $d \leq 2$, an upper bound on the relative entropy stronger than one like that used above is required; see Ref. [57]. For that case, there are also other arguments [58] in the style of the Bogoliubov inequality approach, and we will not pursue it here.

APPENDIX C: THERMODYNAMIC LIMIT, CONVERGENCE CONDITIONS, AND INFINITE-RANGE MODELS

In this Appendix we briefly discuss some questions of the existence of the thermodynamic limit for thermodynamic properties, such as the free (and ground state) energy density, and some other convergence conditions related to these; we also add some comments about infinite-range models. Here we do not assume all of the conditions of Secs. II and III A 1 on the distributions until later; instead, we generally assume only that the bonds J_X are independent and (for simplicity) centered for all X .

The proof of the existence for ν -almost every J of the thermodynamic limit of the free energy per spin at $T > 0$ in a finite-range model is in principle analogous to that of the strong law of large numbers [24,25], if the free energy is roughly the sum of almost independent free energies of distinct regions of the system. A proof, following Ref. [13,19], is based on two ingredients, from which the result follows by a subadditive ergodic theorem. The ingredients are the subadditivity of the free energy when two or more disjoint finite parts are coupled together, and a lower bound on the expected free energy per spin that is uniform in the system size. For the case of independent, centered J_X s, subadditivity is straightforward to show for the *expectation* of the free energy [13], while the proof of the lower bound [19] has been extended to this case in Ref. [13] (and references therein). We will show briefly how the latter proof can be improved and simplified using methods from the body of this paper. We note that Ref. [20] directly proves almost-sure existence of the limit under the same conditions that we discuss below.

First, consider a finite system of sites in a set Λ . Using the Gibbs distribution, we consider $\ln \sum_s e^{-H(s)/T} = -F/T$, introduce a factor λ by replacing $H = H_\Lambda$ by λH , and then consider the integral from 0 to 1 of the derivative with respect to λ . This gives (for Ising spins) the identity

$$-\frac{1}{T}F = |\Lambda| \ln 2 + \frac{1}{T} \int_0^1 d\lambda \sum_{X:X \subseteq \Lambda} J_X \langle s_X \rangle_\lambda, \quad (\text{C1})$$

and then clearly this is bounded above by

$$\leq |\Lambda| \ln 2 + \frac{1}{T} \sum_{X:X \subseteq \Lambda} |J_X| \quad (\text{C2})$$

for any $J = (J_X)_X$. This upper bound on (minus) the free energy F (divided by T) can be obtained directly: the first term is the maximum entropy, and the second is (minus) a lower bound on the internal energy $-\sum_X J_X \langle s_X \rangle$ (divided by T), which would be attained if every bond were satisfied (usually that is not possible). Then for independent, centered bonds, a sufficient condition for the existence of the thermodynamic limit for the expected free energy per spin for $T > 0$ is

$$\lim_{\Lambda \rightarrow \mathbb{Z}^d} \frac{1}{|\Lambda|} \sum_{X:X \subseteq \Lambda} \mathbf{E}|J_X| < \infty. \quad (\text{C3})$$

For homogeneous distributions, we can also express this as

$$\lim_{\Lambda \rightarrow \mathbb{Z}^d} \sum_{p \geq 1} \frac{1}{p} \sum_{X \subseteq \Lambda: |X|=p} \mathbf{E}|J_X| < \infty \quad (\text{C4})$$

for any i . For the models we defined in Sec. II, if the set of p that contribute to the sum is finite, this is equivalent to the condition (3.7) for the model to be short range, but not when the set of p is infinite; in that case the present condition is weaker.

If we return to the identity (C1) and take its \mathbf{E} expectation, then for independent, centered J_X , we can use inequality (3.47), applied to each $\mathbf{E} J_X \langle s_X \rangle_\lambda$, to obtain

$$-\frac{1}{T} \mathbf{E} F \leq |\Lambda| \ln 2 + \frac{1}{2T^2} \sum_{X:X \subseteq \Lambda} \text{Var } J_X. \quad (\text{C5})$$

[As in Sec. III C 3, for Gaussian bonds the same bound can be obtained by integration by parts, or otherwise [13]; for distributions that are not necessarily Gaussian, these bounds (C2) and (C5) are stronger than the corresponding ones in Sec. 3.4 of Ref. [13]. We also note that the variance of the Hamiltonian is given by the same sum,

$$\mathbf{E} H_\Lambda^2 = \sum_{X \subseteq \Lambda} \text{Var } J_X. \quad (\text{C6})$$

This is for the Ising case; for m -vector models, the equality should be replaced by \leq , and as usual the sum should range over x as well as X .]

If we now divide the bound (C5) by $|\Lambda|$, then the condition that the right-hand side be finite as $|\Lambda| \rightarrow \infty$ is another sufficient condition for the existence of the thermodynamic limit for the expected free energy density when $T > 0$ [13,20]. For homogeneous distributions, it reduces to

$$\lim_{\Lambda \rightarrow \mathbb{Z}^d} \sum_{p \geq 1} \frac{1}{p} \sum_{X \subseteq \Lambda: i \in X, |X|=p} \text{Var } J_X < \infty. \quad (\text{C7})$$

For the models we defined in Sec. II, if the set of p that contribute to the sum is finite, this is equivalent to the convergence condition (2.5) for the model. If the set of p is infinite, then the convergence condition implies this one, but not conversely. Hence we see that one can construct models on \mathbb{Z}^d in which the free energy density possesses a thermodynamic limit, but in which Gibbs states presumably do not exist, and there may be locked spins (cf. Appendix A). These models are not finite range, nor are they infinite range in the sense defined in Sec. II. They exist even when the number of terms (indexed by X) for which $i \in X$ and $|X|=p$ is *finite* for all p , simply because of a divergence of the sum over p in (2.5).

The case of the infinite-range models runs parallel to the long-range models. Although they never have Gibbs states in a strict (DLR) sense, the same condition (2.5), where now $\text{Var } J_X$ for $p > 1$ depends on $|\Lambda|$, implies that the molecular field on a given spin is finite including when $T \rightarrow 0$, as one can see heuristically, for example by using a replica symmetric ansatz [16,43]; a related function involving the overlaps arises in the Parisi formula for SK-type models (see, e.g., Ref. [59], where again it is assumed that the sum in the convergence condition converges sufficiently rapidly). Hence there are also SK-type models that have a limit for the free energy density [13,21], but not for the molecular field, and which are thus more singular than what we called infinite-range models. An example of this phenomenon is Derrida's random-energy model, when it is viewed as the $p \rightarrow \infty$ limit of the pure p -spin models [17]. The scaling of $\text{Var } J_X$ is such that the thermodynamic limit of the free energy per spin exists, which

implies that the (p times larger) convergence condition sum tends to infinity. Hence it is not surprising that in the thermodynamic limit of the random-energy or $p \rightarrow \infty$ p -spin model there is a low-temperature region in which the entropy per spin is zero [17], implying that the spins are (in effect) locked into a small number of configurations. We expect similar locking phenomena in other models that satisfy the weaker condition above but not the stronger condition (2.5), including in the models on \mathbb{Z}^d that we mentioned in the preceding paragraph.

Except in the case of short-range models, the bounds so far on the free energy per site are not effective at $T = 0$. However, the variance of the Hamiltonian per site enters a general lower bound on the expectation of the ground state energy $E_0(\Lambda) = \min_s H_\Lambda(s)$ in a finite system Λ (and hence also of the internal energy at $T \geq 0$) for any SG of Ising spins in which the variance of H_Λ is independent of $s = s|_\Lambda$, such as the mixed p -spin models in this paper; the bound does not seem well known in the physics literature. Here we assume the J_X are independent, centered, and Gaussian. The bound is (see, e.g., Ref. [60])

$$\mathbf{E}E_0(\Lambda) \geq -\sqrt{\mathbf{E}H_\Lambda^2} \sqrt{2|\Lambda| \ln 2}, \quad (\text{C8})$$

so the expected ground state energy per spin has a finite limit if the variance per spin does. [If $\mathbf{E}H_\Lambda^2$ is not independent of s , it can be replaced by its maximum to obtain the bound. Of course a similar bound applies to $\mathbf{E} \max_s H_\Lambda(s)$.] This bound can be used to simplify an argument in Ref. [34] (see Proposition 6 there) that locked spins do not occur in the $p = 2$ model at $T = 0$; that argument can be generalized to give the same statement for any finite-range *pure* p -spin model at $T = 0$.

The methods here and in Ref. [13] suffice to prove the existence of a limit for the expected free (or ground state) energy per spin under the conditions stated. A proof of almost sure convergence of the free (or ground state) energy per spin as in Refs. [13,19] can be obtained if one can prove either the subadditivity of the free energy without taking the expectation [19], or else that the free energy per spin concentrates at its expectation as the limit is taken [13]. For either of these, some additional conditions may be necessary, but that lies outside the scope of this paper. Reference [20] proves almost-sure existence of the limit for $T > 0$ without such additional conditions.

[1] H.-O. Georgii, *Gibbs Measures and Phase Transitions*, de Gruyter Studies in Mathematics Vol. 9 (Walter de Gruyter, Berlin, 1988).

[2] B. Simon, *The Statistical Mechanics of Lattice Gases* (Princeton University Press, Princeton, NJ, 1993).

[3] A. Bovier, *Statistical Mechanics of Disordered Systems: A Mathematical Perspective* (Cambridge University Press, Cambridge, 2006).

[4] T. M. Liggett, *Interacting Particle Systems* (Springer, New York, 1985).

[5] R. G. Palmer, *Adv. Phys.* **31**, 669 (1982).

[6] A. C. D. van Enter and J. L. van Hemmen, *Phys. Rev. A* **29**, 355 (1984).

[7] C. M. Newman and D. L. Stein, *Phys. Rev. Lett.* **76**, 4821 (1996); *Phys. Rev. E* **55**, 5194 (1997).

[8] C. M. Newman and D. L. Stein, *J. Phys.: Condens. Matter* **15**, R1319 (2003).

[9] J. Höller and N. Read, *Phys. Rev. E* **101**, 042114 (2020).

[10] G. Kotliar, P. W. Anderson, and D. L. Stein, *Phys. Rev. B* **27**, 602 (1983).

[11] M. Aizenman and J. Wehr, *Commun. Math. Phys.* **130**, 489 (1990).

[12] A. Gandolfi, C. M. Newman, and D. L. Stein, *Commun. Math. Phys.* **157**, 371 (1993).

[13] P. Contucci and C. Giardina, *Perspectives on Spin Glasses* (Cambridge University Press, Cambridge, 2013).

[14] S. F. Edwards and P. W. Anderson, *J. Phys. F* **5**, 965 (1975).

[15] K. M. Khanin and Ya. G. Sinai, *J. Stat. Phys.* **20**, 573 (1979).

[16] D. Sherrington and S. Kirkpatrick, *Phys. Rev. Lett.* **35**, 1792 (1975).

[17] B. Derrida, *Phys. Rev. B* **24**, 2613 (1981).

[18] R. Vershynin, *High Dimensional Probability: An Introduction with Applications in Data Science* (Cambridge University Press, Cambridge, 2018), Sec. 2.7.

[19] A. C. D. van Enter and J. L. van Hemmen, *J. Stat. Phys.* **32**, 141 (1983).

[20] B. Zegarlinski, *Commun. Math. Phys.* **139**, 305 (1991).

[21] F. Guerra and F. L. Toninelli, *Commun. Math. Phys.* **230**, 71 (2002).

[22] J. Fröhlich and B. Zegarlinski, *Commun. Math. Phys.* **110**, 121 (1987).

[23] B. Zegarliński, *J. Stat. Phys.* **47**, 911 (1987).

[24] K. L. Chung, *A First Course in Probability Theory*, 3rd ed. (Academic, San Diego, 2001).

[25] L. Breiman, *Probability* (Society for Industrial and Applied Mathematics, Philadelphia, 1992).

[26] A. J. Bray, M. A. Moore, and A. P. Young, *Phys. Rev. Lett.* **56**, 2641 (1986).

[27] D. S. Fisher and D. A. Huse, *Phys. Rev. Lett.* **56**, 1601 (1986); *Phys. Rev. B* **38**, 386 (1988).

[28] N. Read, *Phys. Rev. E* **90**, 032142 (2014).

[29] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (John Wiley & Sons, Hoboken, NJ, 2006), Chs. 2 and 8.

[30] M. S. Pinsker, *Information and Information Stability of Random Variables and Processes*, translated by A. Feinstein (Holden-Day, San Francisco, 1964), Chaps. 2 and 3.

[31] A. C. D. van Enter and J. L. van Hemmen, *J. Stat. Phys.* **39**, 1 (1985).

[32] L. Leuzzi, G. Parisi, F. Ricci-Tersenghi, and J. J. Ruiz-Lorenzo, *Phys. Rev. Lett.* **101**, 107203 (2008).

[33] D. Ruelle, *Commun. Math. Phys.* **9**, 267 (1968).

[34] N. Read, *Phys. Rev. E* **97**, 012134 (2018).

[35] C. M. Newman and D. L. Stein, in *Spin Glasses*, Lecture Notes in Mathematics Vol. 1900 edited by E. Bolthausen and A. Bovier (Springer, New York, 2007), pp. 159–175.

[36] C. M. Newman and D. L. Stein, in *New Trends in Mathematical Physics: Proceedings of the 2006 International Congress of*

Mathematical Physics, edited by V. Sidoravicius (Springer, New York, 2009), pp. 643–652.

- [37] L.-P. Arguin and M. Damron, *J. Stat. Phys.* **143**, 226 (2011).
- [38] A. J. Bray and M. A. Moore, *J. Phys. C: Solid State Phys.* **17**, L463 (1984); *Phys. Rev. B* **31**, 631 (1985); *Heidelberg Colloquium on Glassy Dynamics and Optimization*, Lecture Notes in Physics Vol. 275, edited by J. L. van Hemmen and I. Morgenstern (Springer-Verlag, Berlin, 1986), p. 121.
- [39] W. L. McMillan, *J. Phys. C: Solid State Phys.* **17**, 3179 (1984); *Phys. Rev. B* **30**, 476 (1984).
- [40] G. Parisi, *Phys. Rev. Lett.* **43**, 1754 (1979); *J. Phys. A: Math. Gen.* **13**, L115 (1980); **13**, 1101 (1980); **13**, 1887 (1980).
- [41] G. Parisi, *Phys. Rev. Lett.* **50**, 1946 (1983).
- [42] M. Mezard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond*, Lecture Notes in Physics Vol. 9 (World Scientific, Singapore, 1987).
- [43] D. J. Gross and M. Mezard, *Nucl. Phys. B* **240**, 431 (1984).
- [44] S. Jensen, N. Read, and A. P. Young, *Phys. Rev. E* **104**, 034105 (2021).
- [45] O. L. White and D. S. Fisher, *Phys. Rev. Lett.* **96**, 137204 (2006).
- [46] T. R. Kirkpatrick and D. Thirumalai, *Phys. Rev. Lett.* **58**, 2091 (1987); *Phys. Rev. B* **36**, 5388 (1987); **37**, 5342 (1988); **38**, 4881 (1988).
- [47] T. R. Kirkpatrick, D. Thirumalai, and P. G. Wolynes, *Phys. Rev. A* **40**, 1045 (1989).
- [48] J.-P. Bouchaud and G. Biroli, *J. Chem. Phys.* **121**, 7347 (2004).
- [49] P. Billingsley, *Convergence of Probability Measures*, 2nd ed. (John Wiley, New York, 1999).
- [50] C. M. Newman and D. L. Stein, in *Mathematics of Spin Glasses and Neural Networks*, edited by A. Bovier and P. Picco (Birkhäuser, Boston, 1997), Appendix A.
- [51] C. M. Newman, *Topics in Disordered Systems* (Birkhäuser, Basel, 1997), Appendix B.
- [52] H. L. Royden, *Real Analysis*, 3rd ed. (Prentice Hall, Upper Saddle River, NJ, 1988).
- [53] E. C. Posner, *IEEE Trans. Inform. Theory*, **21**, 388 (1975).
- [54] R. R. Phelps, *Lectures on Choquet's Theorem*, 2nd ed. (Springer, New York, 2001).
- [55] K. M. Khanin, *Theor. Math. Phys.* **43**, 445 (1980).
- [56] M. Campanino, E. Olivieri, and A. C. D. van Enter, *Commun. Math. Phys.* **108**, 241 (1987).
- [57] A. C. D. van Enter and J. Fröhlich, *Commun. Math. Phys.* **98**, 425 (1985).
- [58] M. Schwartz and A. P. Young, *Europhys. Lett.* **15**, 209 (1991).
- [59] D. Panchenko, *The Sherrington-Kirkpatrick Model* (Springer, New York, 2013).
- [60] S. Chatterjee, *Superconcentration and Related Topics* (Springer, New York, 2014), Appendix A.