Metastate approach to thermodynamic chaos

C. M. Newman
Courant Institute of Mathematical Sciences, New York University, New York, New York 10012

D. L. Stein
Department of Physics, University of Arizona, Tucson, Arizona 85721

(Received 10 December 1996)

In realistic disordered systems, such as the Edwards-Anderson (EA) spin glass, no order parameter, such as the Parisi overlap distribution, can be both translation-invariant and non-self-averaging. The standard mean-field picture of the EA spin glass phase can therefore not be valid in any dimension and at any temperature. Further analysis shows that, in general, when systems have many competing (pure) thermodynamic states, a single state which is a mixture of many of them (as in the standard mean-field picture) contains insufficient information to reveal the full thermodynamic structure. We propose a different approach, in which an appropriate thermodynamic description of such a system is instead based on a *metastate*, which is an *ensemble* of (possibly mixed) thermodynamic states. This approach, modeled on chaotic dynamical systems, is needed when chaotic size dependence (of finite volume correlations) is present. Here replicas arise in a natural way, when a metastate is specified by its (meta)correlations. The metastate approach explains, connects, and unifies such concepts as replica symmetry breaking, chaotic size dependence and replica nonindependence. Furthermore, it replaces the older idea of non-self-averaging as dependence on the bulk couplings with the concept of dependence on the state within the metastate at *fixed* coupling realization. We use these ideas to classify possible metastates for the EA model, and discuss two scenarios introduced by us earlier—a nonstandard mean-field picture and a picture intermediate between that and the usual scaling-droplet picture. [S1063-651X(97)06905-5]

PACS number(s): 05.50.+q, 75.10.Nr, 75.50.Lk, 05.45.+b

I. INTRODUCTION

The nature of the spin glass phase remains a fundamental and unsolved problem in both condensed matter physics and statistical mechanics despite over 20 years of intensive research. At a more basic level, the proper theoretical treatment of systems with quenched disorder and frustration remains open. More recent experiments exhibiting intriguing properties such as aging have not helped to resolve matters, but have instead intensified the ongoing debate [1].

Spin glasses can be metallic or insulating, uniaxial or isotropic, mostly crystalline, or completely amorphous; in general they are not confined to a single set of materials. The microscopic interactions which give rise to spin glass behavior may differ considerably from one material to another. (For a more extensive discussion, see the review article by Binder and Young [2].) Nevertheless, in 1975 Edwards and Anderson (EA) [3] proposed a simple (and unifying) Hamiltonian to describe the thermodynamic, magnetic, and dynamical properties of realistic spin glasses. Their basic assumption was that the essence of spin glass behavior arose from a competition between quenched ferromagnetic and antiferromagnetic interactions, randomly distributed throughout the system.

While the EA model and its mean-field version, the Sherrington-Kirkpatrick (SK) model [4], remain the primary focus of theoretical treatments of spin glasses, a number of other models have also been proposed [2]. It is not our aim in this paper to compare the suitability of these models for the description of all, or some subset of, laboratory spin glasses. Here we are concerned instead with presenting the correct thermodynamic approach to understanding macroscopic properties of not only spin glasses, but, more generally, systems which may have many competing pure states. Throughout this paper we will often apply our ideas and methods to the EA spin glass model—in particular, in its Ising form—but our scope is more general and is not confined to a particular model or a single condensed matter system. We will begin, however, by considering some of the very basic open questions which arise in connection with the EA Ising spin glass and related models.

These open problems cover both thermodynamic and dynamical questions. It is somewhat discouraging that they persist at such a basic level. Very slow equilibration times make the analysis of both laboratory experiments and numerical simulations difficult; and techniques for the theoretical analysis of systems with quenched disorder and frustration remain primitive. So, for example, even though there has been a steady accumulation of evidence that there exists a true thermodynamic phase transition in the EA model (and in real spin glasses), it is fair to say that the issue is not yet closed (and from the standpoint of a mathematical proof, or even a convincing heuristic argument, remains wide open). If an equilibrium phase transition does exist, the lower critical dimension—and in particular, whether it is above or below three—is similarly unknown (see, for example, Refs. [5–7]).

Assuming the existence of a phase transition in some dimension, other open thermodynamic issues include the effect of a magnetic field on the transition; the number of pure states below the transition; the correct description of broken symmetry and the nature of the order parameter; critical properties at the transition; the role of quenched disorder and/or frustration, separately and together, in determining ground state structure and multiplicity [8,9]; and the relationship between the properties of mean-field models and realistic spin glasses. These are only a subset of such very basic questions, which remain the subject of intense debate.

Dynamical properties are central to spin glass physics—here, open problems include the origin of long relaxation times, the understanding of frequency-dependent susceptibility experiments, the origin and interpretation of aging, and the nature of metastable states. As before, this is only a small sample of outstanding questions. Tying together both the thermodynamic and dynamical problems is the general issue of the nature of broken ergodicity (BE) in spin glasses [10–12]. BE may also serve as a bridge toward an investigation of the relationships, if any, between spin glasses and other disordered systems—structural glasses, electric dipole glasses, quadrupolar glasses, and so on [2].

Some thermodynamic questions are of direct experimental relevance: low-temperature properties cannot be understood without knowing the nature of low-energy excitations above the ground state(s); a knowledge of the critical behavior is required before properties near T_c (if it exists) can be explained. It should also be emphasized that many (though not all) important dynamical questions cannot be properly understood, or in some cases even posed, without a correct thermodynamic theory of spin glasses. For example, what is the relationship, if any, between the metastable states of a spin glass and the pure thermodynamic states [13,14]? Moreover, many experiments, e.g., aging, have been explained using conflicting theoretical pictures [15–24]. In the absence of conclusive experimental (or even numerical) data deciding the matter, how does one decide among different theories of the spin glass state, much less explain the experimental observations?

Our concern in this paper is therefore with the thermodynamic nature of spin glasses. In two recent papers [25,26], we introduced several concepts that we believe are crucial for providing a correct and complete description of the equilibrium statistical mechanics of spin glasses and other disordered and/or inhomogeneous systems. Our approach, modeled on chaotic dynamical systems, is necessary in particular for understanding systems with competing thermodynamic states. The unifying idea is that of the metastate [26,27], which enables us to explain and relate chaotic size dependence [28], replica symmetry breaking [29], replica nonindependence, and overlap (non-)self-averaging.

Using the notion of the metastate, we have classified allowable thermodynamic ''solutions'' of the spin glass phase and ruled out others, including one which has long dominated the theoretical literature. In this paper, we will expand and clarify the ideas presented in Refs. [25] and [26], and use them to present a coherent approach to the thermodynamics of spin glasses and, more generally, to disordered and other systems with many competing states. We will begin by discussing a long-standing controversy over the thermodynamic nature of the spin glass phase.

This controversy focuses on the multiplicity and ordering of pure states in realistic spin glasses in finite dimensions, at temperatures below some $T_c>0$ (which, supported by various arguments, is assumed to exist). One approach, which has dominated the spin glass literature for over a decade, assumes that the main features of Parisi and co-workers solution [30–33] of the infinite-ranged SK model—an infinite number of pure states, organized by an ultrametric overlap structure [33], and whose pairwise overlaps are non-self-averaging even in the thermodynamic limit [33]—apply also to realistic spin glasses. In this scenario, the number of order parameters is infinite—i.e., the order parameter is a distribution that is a function of a continuous variable, and this distribution has a characteristic structure, both for a single realization of the couplings, and for the average over all such realizations. The nature of the symmetry breaking here differs from more conventional kinds, familiar from studies of various nondisordered systems: in Parisi's solution, the spin glass phase(s) exhibit *spontaneously broken replica symmetry* of a nontrivial kind.

An alternative point of view arises from a scaling ansatz due to MacMillan [34], Bray and Moore [35], and Fisher and Huse [36–39]. This gave rise to a thermodynamic picture very different from that implied by the Parisi solution (although some features, such as chaotic dependence of correlation functions on temperature, are similar in the two pictures). In particular, the droplet analysis of Fisher and Huse [38,39] led to the conclusion that there exists, at any temperature and in any finite dimension, at most a pair of pure states (see, however, [40] for a critique of this prediction). Here the order parameter and the nature of symmetry breaking is markedly different from that of the Parisi picture.

These two pictures reach opposite conclusions on a number of other thermodynamic issues; for example, any external magnetic field destroys the phase transition in the droplet picture, while that based on the Parisi solution displays an Almeida-Thouless line [41]. (For discussions on whether numerical evidence supports such a transition, see Refs. [42,43].)

Both pictures also imply certain dynamical behavior for spin glasses. However, although the physical origins behind various dynamical mechanisms differ markedly in the two pictures, their observable consequences are often similar (see, for example, the experimental and theoretical discussions on aging in Refs. [15–24]), and most experiments have so far been unable to distinguish between the two pictures. (One possible exception, however, is the set of experiments on noise in mesoscopic spin glass samples by Weissman [44].)

In addition to these pictures, there also exist scaling approaches which predict many pure state pairs at low temperature above three dimensions [45]. A number of other speculative pictures of the spin glass state have also appeared (see [2] for a more thorough presentation), but it is probably fair to say that the scaling-droplet and Parisi pictures presented above have dominated the discussion of the nature of the spin glass phase(s).

The droplet picture of Fisher and Huse makes a number of clear predictions, and is relatively easy to interpret for realistic spin glasses. This has not generally been the case for the Parisi ansatz, and indeed an important issue—although not always recognized as such—is to interpret the implications of the Parisi solution, both thermodynamically and dynamically, for the spin glass phase in finite dimensions. A large literature (see below) exists on the subject, and as a result a reasonably clear consensus has emerged on the thermodynamic structure of short-ranged spin glasses given that the Parisi ansatz holds for them also. We have called this the ''standard SK picture'' in [25] and [26], and will use that terminology also throughout this paper.

The main result of [25] was to prove, however, that the standard SK picture cannot apply to realistic spin glasses in any dimension and at any temperature. This result then led in [26] to an observation which is central to understanding any system with many competing thermodynamic states: one should not focus on any particular (mixed) thermodynamic state, which cannot provide sufficient information to describe the thermodynamic structure; instead, one must consider the *metastate,* which is essentially a probability distribution over the thermodynamic states. One important consequence of [25] and [26] is to redefine the meaning of non-self-averaging, and to show that most quantities of interest can be defined for a *single* realization of the disorder (including those which had been thought to be non-self-averaging in the thermodynamic limit). One can then focus on, and make meaningful statements about, a particular sample rather than an ensemble of samples. This feature should hold also for nondisordered (e.g., inhomogeneous) systems in general.

Using the metastate approach, we were able to narrow down the possible thermodynamic structures for realistic spin glasses. One of these is the scaling-droplet picture; some are new, to our knowledge. Finally, we proposed a possible picture which incorporates some of the features of the Parisi solution for the SK model. In fact, this is the ''maximal'' mean-field picture allowable for realistic spin glasses, but it differs considerably from the familiar standard SK picture presented in the literature. We call this scenario the ''nonstandard SK picture,'' and will discuss it in Sec. VII. One important lesson from our analysis is that, for disordered systems, the features which characterize the system in very large but finite volumes may lead to a misleading thermodynamic picture if straightforwardly extrapolated to infinite volumes. This is of potential importance, for example, in interpreting numerical results. There are previously unsuspected intricacies involved in taking the thermodynamic limit for certain disordered systems.

The plan of the paper is as follows. In Sec. II we review some basic features of the EA model and discuss its finite- and infinite-volume Gibbs states. We discuss the problem of whether many pure states may exist at some dimension and temperature, and show that the answer is independent of coupling realization. In Sec. III we introduce the SK model and the Parisi ansatz for its thermodynamic structure. In Sec. IV we discuss the standard mean-field picture for realistic spin glass models in finite dimension. We then show that this picture cannot hold in any dimension and at any temperature. We also provide an explicit construction of a non-self-averaged thermodynamic state whose overlap distribution function must be self-averaged. In Sec. V we describe an approach to the thermodynamics of systems with many competing states, based on the idea of the *metastate,* an *ensemble* of thermodynamic states. We also present some of the possible scenarios for the metastate of the EA model, including one that is intermediate between the scaling-droplet and mean-field pictures. In Sec. VI we show how replicas arise naturally within this approach, and how the older idea of replica symmetry breaking is understood and unified with newer ideas of dispersal of the metastate and replica non-independence. A replacement for the usual definition of non-self-averaging is also presented. In Sec. VII we introduce the maximal mean-field picture allowable for realistic spin glasses, and show that its thermodynamic features are considerably different from those of the more familiar picture (which cannot hold). In Sec. VIII, we summarize our main results, and discuss their implications for the study of spin glasses and, in a larger framework, systems with many competing states in general. Finally, in the Appendix, we discuss the importance of distinguishing among differing constructions of overlap distributions in the literature.

II. EDWARDS-ANDERSON MODEL

The EA model [3] on a cubic lattice in d dimensions is described by the Hamiltonian

$$\mathcal{H}_{\mathcal{J}}(\sigma) = -\sum_{\langle x,y\rangle} J_{xy}\sigma_x\sigma_y, \tag{1}$$

where $\mathcal{J}$ denotes the set of couplings J_{xy}, and where the brackets indicate that the sum is over nearest-neighbor pairs only, with the sites $x,y \in Z^d$. We will take the spins σ_x to be Ising, i.e., $\sigma_x = \pm 1$; although this will affect the details of our discussion, it is unimportant for our main conclusions. The couplings J_{xy} are quenched, independent, identically distributed random variables; throughout the paper we will assume their common distribution to be symmetric about zero (and usually with the variance fixed to be 1). The most common examples are the Gaussian and $\pm J$ distributions.

Equation (1) is the EA Ising Hamiltonian for an infinite-volume spin glass on Z^d; it is important also to define the EA model on a finite volume, given specified boundary conditions. Let Λ_L be a cube of side $2L+1$ centered at the origin; i.e., $\Lambda_L = \{-L, -L+1, \ldots, L\}^d$. The finite-volume EA Hamiltonian is then just that of Eq. (1) confined to the volume Λ_L, with the spins on the boundary $\partial\Lambda_L$ of the cube obeying the specified boundary condition. (The boundary $\partial\Lambda_L$ of the volume Λ_L consists of all sites not in Λ_L with one nearest neighbor belonging to Λ_L.) For example, the Hamiltonian with *free* boundary conditions is simply

$$\mathcal{H}^f_{\mathcal{J},L}(\sigma) = -\sum_{\langle x,y\rangle \in \Lambda_L} J_{xy}\sigma_x\sigma_y. \tag{2}$$

Another important boundary condition, called a fixed b.c., is where the value of each spin on the boundary is specified. If we denote by $\overline{\sigma}$ the specified boundary spins, then the Hamiltonian becomes

$$\mathcal{H}^{\overline{\sigma}}_{\mathcal{J},L}(\sigma) = \mathcal{H}^f_{\mathcal{J},L}(\sigma) - \sum_{\substack{\langle x,y\rangle \\ x\in\Lambda_L,\, y\in\partial\Lambda_L}} J_{xy}\sigma_x\overline{\sigma_y}. \tag{3}$$

We will frequently employ a familiar and commonly used boundary condition, namely, periodic boundary conditions, where each face of the cube Λ_L is identified with its opposite face. These are generally thought of as minimizing the effects of the boundary (but see van Enter [40]), and allow us to construct manifestly translation-covariant states.

Given the EA Hamiltonian $\mathcal{H}_{\mathcal{J},L}$ on a finite volume Λ_L with a specified boundary condition (e.g., free or fixed or periodic, but without the boundary condition superscript here), we can now define the finite-volume Gibbs distribution $\rho_{\mathcal{J},\beta}^{(L)}$ on Λ_L at inverse temperature β:

$$\rho_{\mathcal{J},\beta}^{(L)}(\sigma)=Z_{L,\beta}^{-1}\exp\{-\beta\mathcal{H}_{\mathcal{J},L}(\sigma)\},\tag{4}$$

where the partition function $Z_{L,\beta}$ is such that the sum of $\rho_{\mathcal{J},\beta}^{(L)}$ over all spin configurations in Λ_L yields 1. In addition to the boundary conditions mentioned so far, one also considers so-called mixed boundary conditions where the Gibbs distribution is a convex combination of the fixed boundary condition distributions for a given L with the weights for the different $\bar{\sigma}$'s adding up to 1.

$\rho_{\mathcal{J},\beta}^{(L)}(\sigma)$ is a finite-volume probability measure, describing at fixed β the likelihood of appearance, within the volume Λ_L, of a given spin configuration obeying the specified boundary condition. Equivalently, the measure is specified by the set of all correlation functions within Λ_L, i.e., by the set of all $\langle\sigma_{x_1}\ldots\sigma_{x_m}\rangle$ for arbitrary m and arbitrary $x_1,\ldots,x_m\in\Lambda_L$.

Thermodynamic states are described by *infinite*-volume Gibbs measures, and therefore can be thought of (and indeed, constructed) as a limiting measure of a sequence, as $L\to\infty$, of such finite-volume measures (each with a specified boundary condition, which may remain the same or may change with L) [46]. The idea of a limiting measure can be made precise by requiring that every m-spin correlation function, for $m=1,2,\ldots$, possesses a limit as $L\to\infty$.

It is clear that, if there is more than one thermodynamic state (at a given temperature) and if arbitrary boundary conditions are allowed for each L, different sequences (of volumes and/or boundary conditions) can have different limiting measures. What is less obvious, but will have important consequences for spin glasses, is that if many thermodynamic states exist, a sequence of measures each having the *same* (e.g., periodic or free) boundary condition may not even *have* a limit [28]. This phenomenon, which we call *chaotic size dependence*, will be more fully described in Sec. V. Because of compactness (i.e., because each of the correlations determining the measure takes values in $[-1,1]$, a bounded closed interval), it follows, however, that every such infinite sequence will have some subsequence(s) with a single limit, so that we are guaranteed the existence of at least one thermodynamic state (i.e., infinite-volume Gibbs distribution). At sufficiently high temperatures it is rigorously known (see below) that there exists only one such state (limiting Gibbs measure), which of course is the paramagnetic state. If the spin-flip symmetry present in the EA Hamiltonian Eq. (1) is spontaneously broken above some dimension d_0 and below some temperature $T_c(d)$, there will be at least a *pair* of limiting measures, such that their even-spin correlation functions will be identical, and their odd-spin correlation functions will have the opposite sign. Assuming that such broken spin-flip symmetry indeed exists for $d>d_0$ and $T<T_c(d)$, the question of whether there exists *more than one* such limiting pair (of spin-flip-related infinite-volume Gibbs distributions) is a central unresolved issue for the EA and related models.

Thermodynamic states may or may not be mixtures of other states. If a Gibbs state $\rho_{\mathcal{J},\beta}$ can be decomposed according to

$$\rho_{\mathcal{J},\beta}=\lambda\rho_{\mathcal{J},\beta}^1+(1-\lambda)\rho_{\mathcal{J},\beta}^2,\tag{5}$$

where $0<\lambda<1$ and ρ^1 and ρ^2 are also infinite-volume Gibbs states (distinct from ρ), then we say that $\rho_{\mathcal{J},\beta}$ is a *mixed* thermodynamic state or simply, mixed state. (A mixed state may, of course, have many, perhaps infinitely many, states in its decomposition.) The meaning of Eq. (5) can be understood as follows: Any correlation function computed using the Gibbs distribution $\rho_{\mathcal{J},\beta}$ can be decomposed as

$$\langle\sigma_{x_1}\cdots\sigma_{x_m}\rangle_{\rho_{\mathcal{J},\beta}}=\lambda\langle\sigma_{x_1}\cdots\sigma_{x_m}\rangle_{\rho_{\mathcal{J},\beta}^1}+(1-\lambda)\langle\sigma_{x_1}\cdots\sigma_{x_m}\rangle_{\rho_{\mathcal{J},\beta}^2}.\tag{6}$$

If a state cannot be written as a convex combination of any other infinite-volume Gibbs states, it is called a *pure* (or extremal) state. As an illustration, the paramagnetic state is a pure state, as are each of the positive and negative magnetization states in the Ising ferromagnet. In that same system, the Gibbs state produced by a sequence of increasing volumes, at $T<T_c$, using only periodic or free boundary conditions is a mixed state, decomposable into the positive and negative magnetization states, with $\lambda=\frac{1}{2}$. A pure state ρ_P can be intrinsically characterized by a *clustering property* (see, e.g., [47,48]), which implies that, for any fixed x,

$$\langle\sigma_x\sigma_y\rangle_{\rho_P}-\langle\sigma_x\rangle_{\rho_P}\langle\sigma_y\rangle_{\rho_P}\to0,\quad |y|\to\infty,\tag{7}$$

and similar clustering for higher-order correlations.

Let $\eta(\mathcal{J},d,\beta)$ now denote the number of pure states in the EA model for a specific coupling realization $\mathcal{J}$. For any d and $\mathcal{J}$ this equals one at sufficiently low β (except for a set of $\mathcal{J}$'s with zero probability according to the underlying disorder distribution—see, e.g., Chapter 3 of [50] and the references cited there). Recall that the droplet picture predicts that $\eta(\mathcal{J},d,\beta)\leqslant2$ for all d and β, while the SK picture assumes that $\eta(\mathcal{J},d,\beta)=\infty$ for d above some (unknown) d_0 and $\beta>\beta_c(d)$.

A reasonable question might then be, could the answer (at fixed β and d) depend on $\mathcal{J}$? What if $\eta=2$ for half the coupling realizations (i.e., for a set of coupling realizations with probability $\frac{1}{2}$), and infinity for the other half? As it turns out, this cannot happen: for a fixed coupling distribution, η at some (d,β) must have the same value for *all* instances $\mathcal{J}$ chosen from the disorder distribution (or more precisely, for *almost every* $\mathcal{J}$—i.e., except for a set with probability zero). In other words, $\eta(\mathcal{J},d,\beta)$ is *self-averaged*; for fixed d and β, it is a constant almost surely as a function of $\mathcal{J}$.

The above statement is mathematically rigorous, but since its proof, and that of all other theorems which appear in this paper, have appeared elsewhere (see, e.g., [49,50]), we here recount only the central arguments. (These arguments will be useful later when we discuss possible scenarios for the thermodynamic structure of the spin glass phase in the EA model.) We first note that $\eta(\mathcal{J},d,\beta)$ is clearly translation-invariant; i.e., if all couplings are translated by any lattice vector a, so that each $J\to J^a$ (i.e., $J_{xy}\to J_{x+a,y+a}$), the function is unchanged: $\eta(\mathcal{J},d,\beta)=\eta(\mathcal{J}^a,d,\beta)$. We next note that the disorder distribution $\nu(\mathcal{J})$ (e.g., an independent Gaussian distribution of mean zero and variance one at each bond on the lattice) is both translation-invariant (trivially) and *translation-ergodic*. Translation ergodicity means that for ν-almost every $\mathcal{J}$, the (spatial) average of translates $\hat{f}(\mathcal{J}^a)$ of any (measurable) function $\hat{f}$ on $\mathcal{J}$ equals the ν average of $\hat{f}$. [As a trivial example, consider a $1d$ problem where the function $\hat{f}(\mathcal{J})$ is just the coupling value J_{01} at a given location on the line. The average of $\hat{f}(\mathcal{J}^a)$ along the line is clearly 0; so is the distribution average at any site. Similarly, for the function $\hat{f}(\mathcal{J})=(J_{01})^k$, the spatial average along the line equals the distribution average at a site, i.e., the kth moment of the random variable J_{01}.] That translation ergodicity in several dimensions holds analogously to the more familiar one-dimensional case seems to have first been shown by Wiener [51].

The assumption that $\hat{f}$ be measurable is a necessary, but somewhat technical requirement. A proof that $\eta(\mathcal{J},d,\beta)$ satisfies the necessary measurability properties appears in [49,50], and will not be discussed further here, except to note that measurability of a function $\hat{f}$ is the minimal requirement for having a well-defined meaning for the average of $\hat{f}$ over the disorder distribution ν.

Returning to the main argument, we note that because $\eta(\mathcal{J},d,\beta)$ is a translation-invariant function of random variables $\mathcal{J}$ whose distribution is translation-ergodic, by averaging $\eta(\mathcal{J}^a,d,\beta)=\eta(\mathcal{J},d,\beta)$ over translates, it follows that $\eta(\mathcal{J},d,\beta)$ equals a constant $\eta(d,\beta)$ almost surely (i.e., for almost every $\mathcal{J}$). $\eta(d,\beta)$ is the distribution average of $\eta(\mathcal{J},d,\beta)$ and it could of course depend on the *distribution* from which the couplings are chosen, but not on any specific realization chosen from a *fixed* distribution.

The same line of reasoning used here to show that the number of pure states at fixed d and β is the same for almost every realization $\mathcal{J}$ was used in [25] to rule out the standard SK picture. This will be discussed later in Sec. IV, but first we present a discussion of the infinite-ranged SK model and the Parisi solution.

III. MEAN-FIELD THEORY AND THE PARISI SOLUTION

The SK model has played an important role in spin glass physics for several reasons. First among these is that it is one of the few (nontrivial) spin glass models which is (generally) believed to have been solved. Moreover, the proposed solution, due to Parisi and co-workers [30,31,33] admits a striking type of symmetry breaking, called *replica symmetry breaking* (RSB), of a form previously unseen in other (nondisordered) systems. The possibility that RSB plays an important role in the physics of *realistic* (i.e., finite dimensional) spin glasses and, possibly, other complex systems has generated a substantial literature (see, for example, Refs. [6,13,14,33,52–61]), and remains controversial.

The SK model is simply the infinite-ranged version of the EA model and thus has no spatial structure. The volume L^d is replaced by N, the number of spins, and the Hamiltonian is

$$\mathcal{H}_{\mathcal{J},N}(\sigma)=-\frac{1}{\sqrt{N}}\sum_{i>j=1}^{N} J_{ij}\sigma_i\sigma_j, \tag{8}$$

where the factor $1/\sqrt{N}$ ensures that the energy per spin remains finite (and nonzero) as $N\to\infty$ (given that, as before, the distribution of each J_{ij} is symmetric about zero and has variance 1). Because there is no natural sense of a boundary in this model, one usually considers simply a sequence of Hamiltonians of the form (8) with increasing N. The probability measure on spin configurations in this model is given by

$$\rho_{\mathcal{J},\beta}^{(N)}(\sigma)=Z_{N,\beta}^{-1}\exp\{-\beta\mathcal{H}_{\mathcal{J},N}(\sigma)\}. \tag{9}$$

It has been known for many years [62] that a correct treatment of quenched disorder involves an averaging (over the couplings) of the free energy and other extensive variables rather than of the partition function. The replica trick [3,63,64] was introduced as a tool for doing such an average; because of the lack of spatial structure in the SK model, it is especially well suited for this approach. Using the replica trick, SK demonstrated the existence of a phase transition, but found that the resulting low-temperature phase was unphysical [4]. It is currently believed that their solution was unstable because it was replica symmetric. Several attempts were made to introduce solutions which broke the replica symmetry [65], but it is now thought that the correct procedure to break replica symmetry in the low-temperature phase of the SK model was the one introduced by Parisi [30].

The Parisi solution to the SK model is both stable and agrees well with numerical results [29]; moreover, some of its essential features can be rederived without the use of replicas, primarily through a cavity method [66,67]. Parisi's approach suggests that there are many pure states of the infinite-ranged model, organized in a highly specific manner which characterizes the SK spin glass phase and its mode of symmetry breaking. Although Parisi's solution predicts many other important features of the spin glass phase [2,29], we will focus here only on its aspects regarding symmetry breaking and order parameters.

We first need to comment on what is meant by ''pure state'' in the SK model, since a precise definition is not available and its meaning remains quite unclear. Other approaches to spin glass mean-field theory (e.g., the Thouless-Anderson-Palme equations [68]) had already suggested the existence of many states at low temperature, in the sense that many solutions could be found which were extrema of the free energy, some subset of which were believed to be minima [69]. It had further been argued that they were separated by barriers which diverged in the thermodynamic limit [70,71]. These are what have come to be called [29] the ''pure states'' of the SK model. The clustering property described by Eq. (7) cannot be used in an infinite-ranged model, which has no spatial structure, but it has been suggested [2,29] that it can be replaced by

$$\lim_{N\to\infty}\langle\sigma_i\sigma_j\rangle_{\beta,N}-\langle\sigma_i\rangle_{\beta,N}\langle\sigma_j\rangle_{\beta,N}=0 \quad (i\neq j), \tag{10}$$

where averages are taken using the distribution corresponding to one of these pure states. The meaning of the averaging in Eq. (10) is poorly defined, however. Because the strength of the random couplings scales to zero as $N \to \infty$, it is unclear what meaning, if any, can be ascribed to the notion of nontrivial thermodynamic states, pure or mixed. In the EA model, on the other hand, methods do exist, as will be briefly discussed in Sec. IV, to construct just such states, for almost every $\mathcal{J}$. This contrast between the SK and EA models will be seen to be significant.

However, in accordance with the usual practice, we will ignore these complications in what follows, though keeping in mind that the meaning of pure state in the SK context remains vague. Using a replica analysis, Parisi found that the SK spin glass state could be described properly only with an *infinite* number of order parameters, describing the relations among the many pure states. This requires the introduction of a new random variable which describes the *replica overlap*,

$$Q_N = \frac{1}{N} \sum_{i=1}^{N} \sigma_i \sigma_i', \tag{11}$$

where the spin configurations σ and σ' are chosen independently from the distribution $\rho_{\mathcal{J},\beta}^{(N)}(\sigma)$ given by Eq. (9). [Technically, σ and σ' are said to be chosen from the *product distribution* $\rho_{\mathcal{J},\beta}^{(N)}(\sigma)\rho_{\mathcal{J},\beta}^{(N)}(\sigma')$.] It is clear that $-1 < Q_N < 1$ for any N. (The subscript β will be suppressed in expressions related to replica overlaps and their distributions. It is understood that all calculations take place at fixed β, and the results depend on β.)

The role of order parameter in the Parisi theory is played not by a single variable, but rather by the probability density $P_{\mathcal{J},N}(q)$ of the random variable Q_N (or functions closely related to it); i.e., $P_{\mathcal{J},N}(q)dq$ is the probability that the random variable Q_N takes on a value between q and $q+dq$ (for fixed $\mathcal{J}$, N and β). Above the critical temperature (i.e., in the paramagnetic state), the distribution of Q_N converges to a δ function at zero as $N \to \infty$. Below this temperature, however, the presumed existence of many states gives rise to a rich and nontrivial behavior of $P_{\mathcal{J},N}(q)$. In particular, Parisi found that in the spin glass phase, $P_{\mathcal{J},N}(q)$ approximates a sum of many δ functions, with weights and locations depending on $\mathcal{J}$ even in the limit $N \to \infty$. This is the first appearance of non-self-averaging (NSA), which plays a central role in the Parisi theory of the spin glass phase.

The usual explanation given for this behavior is that for large N the Gibbs measure $\rho_{\mathcal{J},\beta}^{(N)}$ given by Eq. (9) (for $\beta > \beta_c$) has a decomposition into many pure states $\rho_{\mathcal{J}}^{\alpha}$, where α indexes the pure states

$$\rho_{\mathcal{J}}^{(N)}(\sigma) \approx \sum_{\alpha} W_{\mathcal{J}}^{\alpha} \rho_{\mathcal{J}}^{\alpha}(\sigma), \tag{12}$$

where $W_{\mathcal{J}}^{\alpha}$ denotes the weight of pure state α and the dependence on the inverse temperature β has been suppressed. (The use of the approximation sign is necessary because of the haziness of the meaning of pure state, as discussed above.) Granted the existence of these pure states, one can then consider the case where σ is drawn from the distribution $\rho_{\mathcal{J}}^{\alpha}$ and σ' independently from $\rho_{\mathcal{J}}^{\gamma}$; then the expression in Eq. (11) equals its thermal mean

$$q_{\mathcal{J}}^{\alpha\gamma} \approx \frac{1}{N} \sum_{i=1}^{N} \langle \sigma_i \rangle^{\alpha} \langle \sigma_j \rangle^{\gamma}. \tag{13}$$

Finally, the density $P_{\mathcal{J},N}(q)$ is then given by

$$P_{\mathcal{J},N}(q) \approx \sum_{\alpha,\gamma} W_{\mathcal{J}}^{\alpha} W_{\mathcal{J}}^{\gamma} \delta(q - q_{\mathcal{J}}^{\alpha\gamma}). \tag{14}$$

These expressions can be made precise for the EA model, as will be seen in Sec. IV.

The $q_{\mathcal{J}}^{\alpha\gamma}$'s also exhibit NSA for arbitrarily large N, except for the trivial cases $\alpha = \gamma$ and $\alpha = -\gamma$ (where the minus sign denotes a global spin flip), which correspond respectively to the self-overlaps q_{EA} and $-q_{\mathrm{EA}}$ (with no dependence on $\mathcal{J}$ or α). Why do we not then simply examine the $N \to \infty$ limit of the $q_{\mathcal{J}}^{\alpha\gamma}$'s and their distribution? *A priori* it might seem that, even though the states themselves are not well defined for infinite N, their *overlaps* might still have a well-defined limit. It can be proved, however, that *the existence of the $N \to \infty$ limit (where the limit is taken in a $\mathcal{J}$-independent manner) of the distribution of the $q_{\mathcal{J}}^{\alpha\gamma}$'s is inconsistent with NSA* [28]. This is the first appearance of chaotic size dependence, which will be seen later to play a central role in the analysis of systems with many competing states.

Even though the decomposition of Eq. (12) is presumed to consist of infinitely many states as $N \to \infty$, it is believed [29] that relatively few of them have non-negligible weight (and are therefore thermodynamically significant). These lowest-lying states are believed to have free-energy differences of order 1 (for arbitrarily large N), and their density rises exponentially [66,67] at the lowest energies.

So far we have only discussed the overlaps among pairs of pure states. The relationships among triples of pure states were also investigated [33], and were found to have an ultrametric structure. That is, the Hamming distances (determined by the overlaps) among any three pure states are such that the largest two are always equal (with the third smaller than or equal to the other two).

The main features of the Parisi analysis of the SK model, relevant to the ordering of the spin glass phase, are then the following: (1) The spin glass phase consists of a mixture of infinitely many pure states. Two replicas have non-negligible probability of appearing in different pure states (not related by a trivial global spin flip). This is one interpretation of spontaneous replica symmetry breaking (RSB). (2) For fixed $\mathcal{J}$, the probability density $P_{\mathcal{J}}(q)$ consists of a sum of (approximate) δ functions at discrete locations q such that $-q_{\mathrm{EA}} < q < q_{\mathrm{EA}}$, and a pair always at $\pm q_{\mathrm{EA}}$. The weights and locations of the δ functions, excluding the pair at $\pm q_{\mathrm{EA}}$, depend on $\mathcal{J}$, even as $N \to \infty$ (NSA). (3) Because of this variation with $\mathcal{J}$ of the "interior" δ functions, the average $P(q) = \overline{P_{\mathcal{J}}(q)}$ over all (uncountably many) coupling realizations has a continuous (and nonzero) component for q between the δ functions at $\pm q_{\mathrm{EA}}$, for $0 < T < T_c$. (4) The locations of the δ functions in (2) have an ultrametric structure. In Sec. IV, we examine the meaning of the Parisi picture applied to the EA model.

IV. STANDARD SK PICTURE

If the Parisi solution of the SK model describes the nature of the spin glass phase in realistic spin glasses, as frequently supposed [6,13,14,33,52–55,57–61], what should be its thermodynamic properties? A description along these lines has emerged in the literature (see, for example, Refs. [29,55,59,61,72]) over the past decade and a half. This scenario, which we have called the ''standard SK picture'' [26], is the most straightforward extrapolation of the main features of the Parisi solution to infinite volume spin glasses in finite dimension, and is presented in this section as a precise description of the usual presentations in the literature (see, for example, Refs. [2,29,42,72]).

As discussed in Sec. II, the meaning of pure states in the EA model (and other realistic models) is clear—they are extremal infinite volume Gibbs states [i.e., thermodynamic states which cannot be decomposed as in Eq. (5); equivalently, they satisfy the clustering property as in Eq. (7)]. It is natural then to replace the approximate relation Eq. (12) with an equality

$$\rho_{\mathcal{J}}(\sigma)=\sum_{\alpha} W_{\mathcal{J}}^{\alpha}\,\rho_{\mathcal{J}}^{\alpha}(\sigma), \tag{15}$$

where $\rho_{\mathcal{J}}(\sigma)$ is an infinite volume mixed Gibbs state (at fixed temperature) for a particular coupling realization $\mathcal{J}$, and the $\rho_{\mathcal{J}}^{\alpha}$ are pure states for that $\mathcal{J}$. There may be many such mixed states, so we specify the one above as that produced in some natural way (to be specified later) by a sequence as $L\to\infty$ of finite volume Gibbs distributions on cubes Λ_L with boundary conditions, such as periodic, not depending on the coupling realization. Periodic boundary conditions minimize, in some sense, the effect of a boundary, and are thus a natural analog to the lack of boundary conditions in the SK model. It should be noted, however, that there is some possibility of different behavior for periodic as opposed to, say, free boundary conditions [40].

We digress momentarily to discuss briefly two important points. The first is that, while the notion of an infinite volume (pure or mixed) state is well-defined for nearest-neighbor models, it is less so for systems with very long-ranged interactions, such as Ruderman-Kittel-Kasiya-Yosida. Our arguments that are based on the homogeneity of the disorder, presented below, will still apply to these systems, but this point should be kept in mind.

The second point is that it is necessary that $\rho_{\mathcal{J}}$, obtained from the natural limit procedure discussed above, be defined for almost every $\mathcal{J}$ chosen from the disorder distribution ν (and be measurable in its dependence on $\mathcal{J}$—see Sec. II). While this may seem like a technical point of little physical consequence, it actually plays a crucial role in any thermodynamic treatment of systems with many competing states [25,26,28]. We will come back to this point in a little while.

Returning to the standard SK picture, we note that the other equations in Sec. III are similarly replaced with their exact EA counterparts. The overlap random variable becomes

$$Q=\lim_{L\to\infty}|\Lambda_L|^{-1}\sum_{x\in\Lambda_L}\sigma_x\sigma_x', \tag{16}$$

where σ and σ' are chosen, similarly to before, from the product distribution $\rho_{\mathcal{J}}(\sigma)\rho_{\mathcal{J}}(\sigma')$. If σ is drawn from $\rho_{\mathcal{J}}^{\alpha}$ and σ' from $\rho_{\mathcal{J}}^{\gamma}$, then it follows that the overlap is the constant

$$q_{\mathcal{J}}^{\alpha\gamma}=\lim_{L\to\infty}|\Lambda_L|^{-1}\sum_{x\in\Lambda_L}\langle\sigma_x\rangle^{\alpha}\langle\sigma_x\rangle^{\gamma}. \tag{17}$$

The probability distribution $P_{\mathcal{J}}(q)$ of Q is therefore given by

$$P_{\mathcal{J}}(q)=\sum_{\alpha,\gamma} W_{\mathcal{J}}^{\alpha}W_{\mathcal{J}}^{\gamma}\delta(q-q_{\mathcal{J}}^{\alpha\gamma}). \tag{18}$$

So this SK picture for the EA model (or realistic spin glasses in general) includes the same four features presented at the end of Sec. III (except the word ''approximate'' in the second of these should be deleted); their meanings are now precise. There are other elements of the standard SK picture—e.g., energy gaps of order one separating the lowest-lying states in any large volume and an exponentially increasing density at the lowest energies—but these will not play a central role in what follows.

We turn now to the question of whether the standard SK picture can be valid in any dimension or at any temperature. This question has two parts.

First, does there exist some natural construction which begins with the finite-volume Gibbs states $\rho_{\mathcal{J},\beta}^{(L)}(\sigma)$ of Eq. (4), takes $L\to\infty$, and ends with a (non-self-averaged) infinite-volume state, $\rho_{\mathcal{J},\beta}$ (possibly mixed), and its accompanying overlap distribution $P_{\mathcal{J}}(q)$? By ''natural'' we mean not only the usual sense of the term but also that the construction result in a thermodynamic state $\rho_{\mathcal{J}}$ for almost every $\mathcal{J}$. In particular, we want the limit procedure (e.g., choice of boundary conditions or sequence of cube sizes) to be chosen independently of any specific $\mathcal{J}$. This will help guarantee that the $\mathcal{J}$ dependence generated by this procedure is measurable, and therefore that averages (e.g., of the moments of Q) can be taken with respect to the disorder distribution. We also emphasize that we are interested only in procedures which result in non-self-averaged infinite-volume states (i.e., at least some correlation functions computed within such a state depend on $\mathcal{J}$). Recall that for the SK model, the very notion of such a $\mathcal{J}$-dependent infinite-volume state is unclear. Second, can such a $P_{\mathcal{J}}(q)$ exhibit all the essential features of the SK picture, including those described by the four features above?

The answers to these two parts, given in [25] are, respectively, yes and no. We will explain our construction of $\rho_{\mathcal{J}}$ (which is somewhat technical) and thus the ''yes'' answer to part one later in this section. Meanwhile, we mention one crucial feature of the resulting $\rho_{\mathcal{J}}$, which plays a key role in the ''no'' to part two. That feature is translation covariance; i.e., under the translation of $\mathcal{J}$ to $\mathcal{J}^a$, where $J_{xy}^a=J_{x+a,y+a}$ for each J_{xy}, $\rho_{\mathcal{J}}$ transforms so that

$$\begin{aligned}\rho_{\mathcal{J}^a}(\sigma_{x_1}=\sigma_1,\ldots,\sigma_{x_m}=\sigma_m)\\ =\rho_{\mathcal{J}}(\sigma_{x_1-a}=\sigma_1,\ldots,\sigma_{x_m-a}=\sigma_m).\end{aligned} \tag{19}$$

The conceptual significance of translation covariance is that the mapping from $\mathcal{J}$ to $\rho_{\mathcal{J}}$, being a natural one, should not (and in our construction does not) depend on the choice of an origin. It also implies the technical conclusion [49,50] that this procedure leads to a limiting infinite-volume overlap distribution function $P_{\mathcal{J}}(q)$, which exists for almost every $\mathcal{J}$ and depends measurably on $\mathcal{J}$ (guaranteeing that averages of q-dependent functions can be taken over the couplings). To begin our answer to the second part of the question, we see what the translation covariance of $\rho_{\mathcal{J}}$ implies about $P_{\mathcal{J}}$.

By translation covariance of $\rho_{\mathcal{J}}$, the overlap random variable $Q_{\mathcal{J}^a}$ has the same distribution as the random variable $Q_{\mathcal{J}}^{-a}$, where

$$Q_{\mathcal{J}}^{-a} \equiv \lim_{L\to\infty} |\Lambda_L|^{-1} \sum_{x\in\Lambda_L} \sigma_{x-a}\sigma'_{x-a} = Q_{\mathcal{J}}. \tag{20}$$

Thus the overlap distribution function $P_{\mathcal{J}^a} = P_{\mathcal{J}}$ for almost every $\mathcal{J}$ and all $a \in Z^d$; i.e., *$P_{\mathcal{J}}$ is translation-invariant*.

As in the case of the translation covariance of $\rho_{\mathcal{J}}$, the translation invariance of $P_{\mathcal{J}}$ has the conceptual significance that a natural object like the Parisi order parameter distribution should not (and in our construction does not) depend on the choice of an origin. But it also has the very important technical (and physical) significance that *$P_{\mathcal{J}}$ must be self-averaged* because, as already noted in Sec. II (in the discussion on the number of pure states), a (measurable) translation-invariant function of random variables whose distribution is translation-ergodic is a constant almost surely, by the ergodic theorem. (We remark that the fact that we are dealing with a function of $\mathcal{J}$ whose value for each $\mathcal{J}$ is an entire distribution is not a problem, since any particular moment of that distribution is a real-valued function of $\mathcal{J}$.)

This answers the second part of our question: the overlap distribution function $P_{\mathcal{J}}(q) = P(q)$ is independent of $\mathcal{J}$. It therefore does not exhibit non-self-averaging [property (1)], and can exhibit *at most* one of the two properties (2) and (3) discussed at the end of Sec. III. While property (2) [discreteness of the locations of the δ functions which appear in $P(q)$] is not rigorously ruled out, it now seems like a highly implausible possibility, since it would imply that the locations (and weights) of the δ functions (and consequently the values of q which correspond to *no* overlap value) are all independent of $\mathcal{J}$. If property (2) is then eliminated as a realistic possibility, then one can also rule out property (4) of the SK picture, i.e., ultrametricity of all of the pure state overlaps for fixed $\mathcal{J}$ [25].

Consequently, we have proved that *the standard SK picture cannot be valid in any dimension and at any temperature*. This result goes beyond our specific construction of the Gibbs state $\rho_{\mathcal{J}}$ and overlap distribution $P_{\mathcal{J}}$, since *any* infinite-volume translation-invariant overlap distribution function would be self-averaging. It would be quite peculiar if the overlap distribution depended on the choice of origin of the coordinate system, and we therefore regard the property of translation-invariance for $P_{\mathcal{J}}$ (or translation covariance for $\rho_{\mathcal{J}}$) as not specific to our particular construction.

Our conclusion is that *nearest-neighbor (and in general realistic) spin glasses exhibit non-mean-field behavior, because for those systems one can construct a non-self-averaged Gibbs state $\rho_{\mathcal{J}}$ whose overlap distribution $P_{\mathcal{J}}$ is self-averaged.* The standard SK picture therefore cannot describe realistic spin glasses at *any* dimension or temperature. It is important to note that these conclusions apply to the *thermodynamics* of spin glasses. What might or might not occur in *finite* volumes involves several subtle issues and will be discussed in Sec. VII.

While the demise of the standard SK picture is interesting in itself, and has important consequences for our understanding of spin glasses, the methods used above and in our construction of $\rho_{\mathcal{J}}$ lead to perhaps more significant conclusions that might affect our thinking on not only spin glasses, but disordered—and more widely, inhomogeneous—systems at a deeper level. Indeed, these methods indicate a path to a general approach for studying the thermodynamics of systems with many pure states. One consequence will be the emergence of a replacement for the standard SK picture, a picture which retains some mean-field flavor. The general formulation introduces several interesting concepts, among them replica nonindependence and a definition of non-self-averaging, and relates them to replica symmetry breaking, overlaps, and chaotic size dependence. The unifying theme is the concept of the metastate, which is introduced in Sec. V. Before that, we complete this section with a discussion of how we construct our thermodynamic state $\rho_{\mathcal{J}}$.

We begin by noting that we cannot simply fix $\mathcal{J}$ and take an ordinary limit (i.e., through a sequence of cube sizes L chosen independently of $\mathcal{J}$) of the finite cube, periodic b.c. state $\rho_{\mathcal{J},\beta}^{(L)}$, as $L\to\infty$. Unlike, say, the $d=2$ homogeneous Ising ferromagnet, where such a limit exists (and equals $\frac{1}{2}\rho^{+} + \frac{1}{2}\rho^{-}$) by spin flip symmetry considerations (and the fact that ρ^{+} and ρ^{-} are the only pure states [73,74]), there is no guarantee for a spin glass that there is a well-defined limit. (In fact, if such a limit *does* exist for the spin glass, one can then prove [28] that the *same* limiting state is obtained through the use of *antiperiodic* boundary conditions—a feature that already seems incompatible with a SK picture.)

It is true that one can easily prove, using compactness arguments, convergence along *subsequences* of L's for each $\mathcal{J}$. But these subsequences should (in a SK picture) be $\mathcal{J}$ dependent. The inconsistency between the existence of many pure states and the existence of a thermodynamic limit for a sequence of finite-volume Gibbs states using coupling-independent boundary conditions (such as periodic) and cube sizes has serious consequences not only for spin glasses but also for systems in general with many competing states. It suggests in the present case that, if many pure states exist, such a sequence of finite-volume Gibbs state exhibits *chaotic size dependence* (CSD) and does not converge to a limit. The convergent ($\mathcal{J}$-dependent) subsequences would then give rise to different (non-self-averaged) pure states or mixed states with no way to make a (measurable) choice of a limit state for each $\mathcal{J}$.

There is, however, a natural limit procedure which does give rise to an infinite-volume Gibbs state $\rho_{\mathcal{J}}$, while avoiding such difficulties. Here, one considers the *joint* distribution on the spins *and* the couplings; i.e., one considers the distribution $\nu(\mathcal{J}) \times \rho_{\mathcal{J},\beta}^{(L)}$ on the periodic cube Λ_L [25]. Then (again using compactness arguments) some subsequence of L's converges to a limiting infinite-volume joint distribution $\mu(\mathcal{J},\sigma)$. From this joint distribution, $\rho_{\mathcal{J}}$ results when the spin configurations are chosen conditioned on $\mathcal{J}$, which is chosen from $\nu(\mathcal{J})$ in the usual way; i.e., $\rho_{\mathcal{J}}$ is determined by the identity $\mu(\mathcal{J},\sigma)=\nu(\mathcal{J})\times\rho_{\mathcal{J}}(\sigma)$. The important difference with the earlier limit procedure is that this one is valid for almost every $\mathcal{J}$, i.e., the subsequence of L's is $\mathcal{J}$ *independent*. (The discussion so far has been based on mathematically rigorous arguments. At this point however, we would suggest—but cannot rigorously prove—that it is probably the case that using a subsequence of L's is not needed for convergence, because the use of a joint distribution for $\mathcal{J}$ and σ should avoid CSD.) A proof that the resulting limiting distribution is indeed a Gibbs state may be found in [27,49,50]. We note that such joint distribution limits were considered, implicitly or explicitly, in Refs. [75–78].

To obtain a clearer idea of this construction, consider the following procedure. Start with three cubes (labeled a, b, and c), all centered at the origin, with volumes L_a^d, L_b^d, and L_c^d, with $1\ll L_a\ll L_b\ll L_c$. On the outermost box we impose periodic boundary conditions. The couplings are fixed inside the intermediate box (and averaged over between the intermediate and largest box); and in the innermost box the overlap computation is done. The average over couplings between the intermediate and large boxes is equivalent to an average over many boundary conditions (consistent with the outer periodic b.c.) on the boundary of the intermediate box.

Now let $L_c\to\infty$ while keeping L_a and L_b fixed; then let $L_b\to\infty$, while still keeping L_a fixed. This sends our ''average over boundary conditions'' off to infinity, and results in an infinite-volume $\rho_{\mathcal{J}}$ which is conditioned on *all* of the couplings and is therefore non-self-averaged. Finally, let $L_a\to\infty$; this gives finally the overlap distribution $P_{\mathcal{J}}(q)$ between infinite-volume pure states appearing in the $\rho_{\mathcal{J}}$.

It is important to note that any analog of this procedure for the SK model will result in an infinite-spin Gibbs state, but a trivial one; i.e., it will already be self-averaged and therefore uninteresting. This is because fixing only finitely many couplings in the SK model and averaging over the remainder is equivalent to averaging over *all* of the couplings when $N\to\infty$. This difference between finite-dimensional and mean-field models is crucial.

We conclude by pointing out why our construction (for the EA model) of the limiting joint distribution $\mu(\mathcal{J},\sigma)$ yields translation covariance for $\rho_{\mathcal{J}}$. This is so because taking periodic b.c.'s on the cube Λ_L really means that the couplings and spins are defined on a (discrete) torus of size L, with a finite-volume joint distribution invariant under torus translations. This implies that any (subsequence) limit joint distribution is invariant under translations of Z^d, which in turn implies that the infinite-volume Gibbs state $\rho_{\mathcal{J}}$ is translation covariant. In Sec. V, we go beyond the construction of a single limiting thermodynamic state by introducing the notion of metastates.

V. CHAOTIC SIZE DEPENDENCE AND METASTATES

In this section we will describe an approach, introduced in Ref. [26], to studying inhomogeneous and other systems with many competing pure states. This approach is based on an analogy to chaotic dynamical systems, and involves the replacement of the study of a single thermodynamic state with an *ensemble* of (pure or mixed) thermodynamic states.

In Sec. IV we were forced (by chaotic size dependence) to replace a simple sequence of states on cubes with periodic boundary conditions with a more complicated sequence which involved an averaging over boundary conditions, followed by sending this average off to infinity. This avoids chaotic size dependence (at least for a $\mathcal{J}$-independent subsequence of volumes, but probably altogether). In this section, we will pursue the opposite strategy—we will forego the end run around CSD, and instead use it to gather maximal information about the disordered system. The price will be to abandon the usual procedure of constructing and studying a single infinite volume Gibbs state $\rho_{\mathcal{J}}$.

The central observation behind this is that, at any (large) fixed L (and with periodic boundary conditions), the existence of multiple pure states should generally require an approximate decomposition as in Eq. (12) [see also Eq. (15)]:

$$\rho_{\mathcal{J}}^{(L)}(\sigma)\approx\sum_{\alpha} W_{\mathcal{J},L}^{\alpha}\rho_{\mathcal{J}}^{\alpha}(\sigma). \tag{21}$$

For each L, the pure states appearing with the largest weights in the sum will be those whose configurations within the volume of size L are best adapted to the boundary condition. Chaotic size dependence requires that the pure states and weights appearing within this decomposition vary persistently as L is increased (though it says nothing about the rate at which this variation occurs).

The analogy with the chaotic orbit of a dynamical system follows from the identification of cube size L with time t along such an orbit. A (chaotic) dynamical system's trajectory will appear random if one considers the sequence of points along its orbit, but one can describe its long-time behavior by studying the appropriate probability measure κ_{dyn} on state space. That is, one can construct a histogram at times $t_1,t_2,\ldots,t_N$, with N increasing to infinity, and study the fraction of times spent by the orbit in different parts of state space (in a continuous space this would require breaking the space up into bins). The $N\to\infty$ limit of this process yields a well-defined κ_{dyn}.

Similarly, we consider at fixed $\mathcal{J}$ a histogram of finite volume Gibbs states $\rho_{\mathcal{J}}^{(L_1)},\rho_{\mathcal{J}}^{(L_2)},\ldots,\rho_{\mathcal{J}}^{(L_N)}\to\kappa_{\mathcal{J}}$ as $N\to\infty$. The information contained in $\kappa_{\mathcal{J}}$ provides the fraction of cube sizes which the system spends in different (possibly mixed) thermodynamic states Γ. We refer to $\kappa_{\mathcal{J}}$, which is a probability measure on thermodynamic states Γ at fixed $\mathcal{J}$, as the *metastate*.

To simplify notation, it will be assumed in the ensuing discussion that convergence to the metastate is valid without need for a subsequence of N's or a subsequence $L_1,L_2,\ldots$ of the cube sizes. We point out however that Külske has studied some models, e.g. the Curie-Weiss random-field model, in which choosing a sparse subsequence of sizes is necessary for the empirical distribution (i.e., the histogram) to converge (for almost every disorder configuration) to the metastate. We will not discuss these issues here, but refer the reader to Ref. [79] for details.

Our empirical distribution approach to construction of a metastate, based on CSD for fixed $\mathcal{J}$, constructs $\kappa_{\mathcal{J}}$ as the limit of $\kappa_{\mathcal{J}}^{(L)}$, a type of microcanonical ensemble in which each of the finite-volume states $\rho_{\mathcal{J}}^{(1)},\rho_{\mathcal{J}}^{(2)},\ldots,\rho_{\mathcal{J}}^{(L)}$ has weight L^{-1}. This limit can be understood in the following way: consider a (nice) function on the states, such as the correlation $\langle \sigma_x \sigma_y \rangle(\cdot)$; i.e., $\langle \sigma_x \sigma_y \rangle(\Gamma)$ is this correlation computed using a particular infinite-volume Gibbs distribution Γ, and $\langle \sigma_x \sigma_y \rangle(\rho^{(L)})$ is the same correlation computed using the finite-volume Gibbs distribution $\rho^{(L)}$ (we suppress the $\mathcal{J}$ index here, which is understood). If $[\cdot]_\kappa$ denotes an average of a state-dependent function over the *metastate* (i.e., the function of each state is weighted using the weight of the corresponding state within the metastate), then

$$\lim_{L\to\infty} (1/L) \sum_{l=1}^{L} \langle \sigma_x \sigma_y \rangle(\rho^{(l)}) = [\langle \sigma_x \sigma_y \rangle(\Gamma)]_\kappa . \tag{22}$$

Furthermore, such an equation similarly holds for any other (nice) function of finitely many correlations (regarded as a function on states).

There is another approach to constructing the metastate, due to Aizenman and Wehr [27], which uses the randomness of the couplings directly, in a manner similar to that of the construction of $\rho_{\mathcal{J}}$ in Sec. IV. There we studied the limiting joint distribution $\mu(\mathcal{J},\sigma)$ of the random pairs $(\mathcal{J},\sigma^{(L)})$ distributed for finite L by $\nu(\mathcal{J}) \times \rho_{\mathcal{J}}^{(L)}$. Here one considers instead the random pair $(\mathcal{J},\rho_{\mathcal{J}}^{(L)})$ at finite L. We will not discuss various technicalities associated with this method; details can be found in Refs. [27,49,50]. We will simply note here that the two approaches (at the very least along common $\mathcal{J}$-independent subsequences) yield the same limiting metastate.

The metastate $\kappa_{\mathcal{J}}$ contains all of the thermodynamic information about a system, in this case the EA spin glass with coupling realization $\mathcal{J}$. As such, it contains far more information than the single thermodynamic state $\rho_{\mathcal{J}}$ generated by the construction of Sec. IV (or any other single state). In fact, it can be seen [27,49,50] that the $\rho_{\mathcal{J}}$ of Sec. IV is the *average* thermodynamic state of the ensemble of states within the metastate $\kappa_{\mathcal{J}}$, in the following sense: consider any spin correlation in the state $\rho_{\mathcal{J}}$, e.g., $\langle \sigma_{x_1}, \ldots, \sigma_{x_n} \rangle_{\rho_{\mathcal{J}}}$. This equals the average (over the metastate) of the correlation function of the same set of spins over *all* thermodynamic states Γ of the ensemble. So if $\kappa_{\mathcal{J}}(\Gamma) d\Gamma$ denotes (formally) the probability of appearance of the states within a region of state space centered on Γ of state-space volume $d\Gamma$, then

$$\begin{aligned}\langle \sigma_{x_1} \cdots \sigma_{x_n} \rangle_{\rho_{\mathcal{J}}} &= [\langle \sigma_{x_1} \cdots \sigma_{x_n} \rangle_\Gamma]_{\kappa_{\mathcal{J}}} \\ &= \int \langle \sigma_{x_1} \cdots \sigma_{x_n} \rangle_\Gamma \kappa_{\mathcal{J}}(\Gamma) d\Gamma, \end{aligned}\tag{23}$$

and similarly for all other correlation functions.

We see that one problem with the standard SK picture (and with other standard thermodynamic treatments of systems with many competing states) is that the state $\rho_{\mathcal{J}}$ (or any other single state, pure or mixed) is simply not a rich enough description of the $L \to \infty$ behavior of a thermodynamic system where CSD occurs. In these approaches, one is in effect replacing with a single average all of the information contained in an entire distribution.

To illustrate the nature of the metastate, we now present some simple examples. The first is the trivial case of a single pure phase, e.g., the paramagnetic state. Then $\lim_{L\to\infty} \rho_{\mathcal{J}}^{(L)} = \rho_{\mathcal{J}}$ is a single pure state, there is no CSD, and $\kappa_{\mathcal{J}}(\Gamma) = \delta(\Gamma - \rho_{\mathcal{J}})$.

In the second example we suppose that the scaling-droplet picture is correct, so that only two pure states $\rho'_{\mathcal{J}}$ and $\rho''_{\mathcal{J}}$ exist, related by a global spin flip. Then (as in the $d=2$ homogeneous Ising ferromagnet with periodic or free b.c.'s)

$$\lim_{L\to\infty} \rho_{\mathcal{J}}^{(L)} = \tfrac{1}{2} \rho'_{\mathcal{J}} + \tfrac{1}{2} \rho''_{\mathcal{J}}, \tag{24}$$

and again there would be no CSD. Indeed, the analogy here is to a dynamical system with a simple fixed point. The metastate is simply

$$\kappa_{\mathcal{J}}(\Gamma) = \delta(\Gamma - [\tfrac{1}{2} \rho'_{\mathcal{J}} + \tfrac{1}{2} \rho''_{\mathcal{J}}]). \tag{25}$$

However, we can introduce a slight variation of this procedure to illustrate the potential sensitivity of the metastate to the boundary conditions used in the limiting procedure. Suppose that, instead of periodic boundary conditions for each L, we use *fixed* boundary conditions, e.g., all spins are $+1$ on the boundary of each L^d cube appearing in the sequence. This of course breaks the spin flip symmetry, and for some L's the state $\rho'_{\mathcal{J}}$ will be preferred, while others will prefer $\rho''_{\mathcal{J}}$, depending in each case on whether the sum along the boundary of $\langle \sigma_x \rangle_{\rho'_{\mathcal{J}}}$ is (substantially) positive or negative. (There may be occasional L's where the preference for each state is roughly equal, but this should be a negligible fraction of L's and so would not show up in the limiting histogram which yields the metastate.)

So in this case we obtain chaotic size dependence, albeit of a trivial sort: $\rho_{\mathcal{J}}^{(L)} \approx \rho'_{\mathcal{J}}$ for roughly half of the L's, and $\rho_{\mathcal{J}}^{(L)} \approx \rho''_{\mathcal{J}}$ for the remainder. The metastate is now

$$\kappa_{\mathcal{J}}(\Gamma) = \tfrac{1}{2} \delta(\Gamma - \rho'_{\mathcal{J}}) + \tfrac{1}{2} \delta(\Gamma - \rho''_{\mathcal{J}}). \tag{26}$$

Here the metastate is a rough analog to the $\kappa_{\rm dyn}$ obtained from a discrete time dynamical system with an attractive orbit of period 2.

This is our first example in which the metastate is not simply a δ function on thermodynamic states. We call this behavior *dispersal of the metastate*, and it is intimately connected with CSD. From this and the previous example, it should be clear that dispersal of the metastate is quite different from the mere existence of multiple states; while the existence of more than one state is necessary for dispersal to occur, it by no means guarantees it.

The above discussion leads naturally to the following possibility, first proposed in [26] as a possible candidate for the EA metastate, based in turn upon earlier work in [8]. Suppose that the EA spin glass has many pure states in some d and at some β, but unlike in the mean-field picture each volume ''sees'' essentially only one pair at a time. In other words, for every L (and once again using periodic boundary conditions), one finds that

$$\rho_{\mathcal{J}}^{(L)} \approx \tfrac{1}{2} \rho_{\mathcal{J}}^{\alpha_L} + \tfrac{1}{2} \rho_{\mathcal{J}}^{-\alpha_L} \tag{27}$$

where $-\alpha$ refers to the global spin flip of pure state α. In any volume, this looks very much like the droplet-scaling picture, but its thermodynamic behavior is considerably different: there are infinitely many pure states and which pair appears in any finite volume depends chaotically on L. Unlike the droplet-scaling picture, this possibility exhibits CSD and dispersal of the metastate. In this ''chaotic pairs'' picture the (periodic b.c.) metastate is dispersed over (infinitely) many Γ's, of the form $\Gamma=\Gamma^{\alpha}=\frac{1}{2}\rho_{\mathcal{J}}^{\alpha}+\frac{1}{2}\rho_{\mathcal{J}}^{-\alpha}$.

It is interesting to note that just this type of behavior is observed for the many ground states of a simple highly disordered spin glass model in high dimension ([8]; see also [80]). (Indeed, for the EA model itself, we would expect this type of behavior to occur at $T=0$ if infinitely many ground states exist.)

One can study metastates in models more complicated than those above, but still simpler than the EA spin glass, and a discussion of some of these (e.g., random-field Ising models, the highly disordered spin glass of Ref. [8], and the homogeneous XY model with random b.c.'s) appears in [49], to which the interested reader is referred for details. At this point, however, we will proceed and use the ideas introduced in this section to revisit the concepts of replica symmetry breaking and non-self-averaging, and will introduce some concepts such as replica nonindependence. The idea of the metastate will enable us to relate, explain, and unify these concepts. We will then return to the EA model and discuss the remaining possibilities for its metastate, and therefore its low-temperature thermodynamic structure.

VI. REPLICA SYMMETRY BREAKING, REPLICA NONINDEPENDENCE, AND OVERLAP DISTRIBUTIONS

In Sec. IV we discussed the Parisi order parameter distribution $P_{\mathcal{J}}(q)$ in the standard SK picture, whose counterpart $P_{\mathcal{J},N}(q)$ for the SK model was successful in describing mean-field spin glass ordering. In the standard SK model $P_{\mathcal{J}}(q)$ is constructed as the distribution of the overlap random variable Q, which in turn is constructed according to Eq. (16), where the spin configurations σ and σ' are chosen from the product distribution $\rho_{\mathcal{J}}(\sigma)\rho_{\mathcal{J}}(\sigma')$; i.e., each is chosen independently from the same (thermodynamic) state $\rho_{\mathcal{J}}$.

But now, given the metastate point of view discussed in Sec. V, we know that the state $\rho_{\mathcal{J}}(\sigma)$ is really the average over the metastate, in the sense described by Eq. (23). Equivalently,

$$\rho_{\mathcal{J}}(\sigma)=\int \Gamma(\sigma)\kappa_{\mathcal{J}}(\Gamma)d\Gamma. \tag{28}$$

So each time a pair of spin configurations, say (σ^1,σ^2), is chosen from $\rho_{\mathcal{J}}(\sigma^1)\rho_{\mathcal{J}}(\sigma^2)$, an independent Γ is used for each configuration. That is, σ^1 is chosen from Γ^1, and σ^2 from Γ^2 with (Γ^1,Γ^2) chosen from $\kappa_{\mathcal{J}}(\Gamma^1)\kappa_{\mathcal{J}}(\Gamma^2)$; Γ^1 and Γ^2 will in general be distinct if the metastate is dispersed. This in turn means in essence [see Eq. (22)] that the spin configuration σ^1 is chosen using the distribution $\rho_{\mathcal{J}}^{(L_1)}$ and σ^2 from $\rho_{\mathcal{J}}^{(L_2)}$, with $L_1\neq L_2$. It seems more natural instead to take the two replicas from the same distribution, i.e., for a *single* L, and therefore from the *same* Γ. As Guerra has pointed out [81], this order of operations (in which replicas are taken *before* $L\to\infty$) could yield a different result than that obtained by first letting $L\to\infty$ to obtain an infinite volume state $\rho_{\mathcal{J}}$ and then taking replicas. The noncommutativity of these operations will be shown to follow from a phenomenon we call *replica nonindependence*, which is not the same as replica symmetry breaking, as will be seen below. But first we will explore the meaning of this way of taking replicas.

Taking replicas first (i.e., from the same L) really means, in terms of the metastate, that they are being taken from the same Γ, i.e., from $\Gamma(\sigma^1)\Gamma(\sigma^2)$ for some Γ chosen from $\kappa_{\mathcal{J}}(\Gamma)$. (Without metastates, it would be difficult to assign a clear meaning to this statement.) For n replicas (where n can be any positive integer, or infinity) we take n uncoupled (but identical) Hamiltonians (and boundary conditions) in the cube L^d. We then use for finite L the product measure $\rho_{\mathcal{J}}^{n(L)}=\rho_{\mathcal{J}}^{(L)}(\sigma^{1(L)})\rho_{\mathcal{J}}^{(L)}(\sigma^{2(L)})\cdots\rho_{\mathcal{J}}^{(L)}(\sigma^{n(L)})$. The limiting joint distribution for $(\mathcal{J},\sigma^{1(L)},\sigma^{2(L)},\ldots,\sigma^{n(L)})$ is then of the form $\nu(\mathcal{J})\rho_{\mathcal{J}}^{n}(\sigma^1,\sigma^2,\ldots,\sigma^n)$ for some $\rho_{\mathcal{J}}^{n}$ that we call the infinite-volume n-replica measure. (The mathematical analysis of this limit procedure is essentially the same as was discussed for $n=1$ in Sec. IV above and, with more detail, in [49,50].) In this approach, replicas in the infinite volume limit are described by $\Gamma(\sigma^1)\Gamma(\sigma^2)\cdots$ with Γ distributed by $\kappa_{\mathcal{J}}$; replicas in finite volume are taken from the *same* L, and $\kappa_{\mathcal{J}}$ describes the sampling of states as L varies.

A crucial point, as emphasized by Guerra [81], is that *in the infinite-volume replica measure* $\rho_{\mathcal{J}}^{n}$, *the replicas need not be independent*, although they of course *are* independent in the finite-volume measure $\rho_{\mathcal{J}}^{n(L)}$. The replicas in infinite volume can be thought of as coupled through ''boundary conditions at infinity.''

If this occurs, we say [26] that the system displays replica nonindependence (RNI). The presence of RNI means that $\rho_{\mathcal{J}}^{n}$, which is a thermodynamic state for the uncoupled replica Hamiltonians, is *not* simply equal to the product of the individual Gibbs states $\rho_{\mathcal{J}}(\sigma^i)$. In general, we have from the above description that

$$\rho_{\mathcal{J}}^{n}(\sigma^1,\sigma^2,\ldots,\sigma^n)=\int [\Gamma(\sigma^1)\Gamma(\sigma^2)\cdots\Gamma(\sigma^n)]\kappa_{\mathcal{J}}(\Gamma)d\Gamma. \tag{29}$$

This makes it clear that RNI is equivalent to dispersal of the metastate. If the metastate is nondispersed, its weight is concentrated entirely on a single thermodynamic state, so $\kappa_{\mathcal{J}}$ is a delta function, and the RHS of Eq. (29) reduces to a simple product of Gibbs states (each of which is the single state on which the metastate is concentrated). Otherwise, the product decomposition of $\rho_{\mathcal{J}}^{n}$ is as a mixture over $\kappa_{\mathcal{J}}$. This also shows that RNI is equivalent to the noninterchangeability of the operations of taking replicas and the thermodynamic limit.

In Ref. [26], these points were explained using the idea of ''metacorrelations.'' Just as the usual correlations $\langle\sigma_{i_1}\cdots\sigma_{i_m}\rangle_\Gamma$ are moments (in this case, of order m) characterizing the thermodynamic state Γ, metacorrelations are moments that characterize the metastate κ. I.e., they are the averages (over the metastate) of functions $g(\Gamma)$ on the states that are monomials (of order m) in the correlations (of varying orders):

$$[g(\Gamma)]_\kappa=[\langle\sigma_{A_1}\rangle_\Gamma\cdots\langle\sigma_{A_m}\rangle_\Gamma]_\kappa, \tag{30}$$

where σ_A denotes $\sigma_{i_1}\cdots\sigma_{i_k}$ for the set $A=\{i_1,\ldots,i_k\}$. As noted in Ref. [26], restriction to metacorrelations of order $m=1$ yields $\rho_{\mathcal{J}}$, to $m\leqslant 2$ yields the two-replica measure $\rho^2_{\mathcal{J}}(\sigma^1,\sigma^2)$, which corresponds to ''integrating out'' all the other replicas in $\rho^\infty_{\mathcal{J}}$, and so on. The measure $\rho^\infty_{\mathcal{J}}$ therefore not only contains information about arbitrarily many replicas, but since it determines all the metacorrelations, it also contains all information about the metastate $\kappa_{\mathcal{J}}$.

Replica symmetry breaking (RSB) occurs when individual thermodynamic states Γ (chosen from $\kappa_{\mathcal{J}}$) are mixtures of multiple pure states, so that even, when restricted to a single Γ, replicas can come from different pure states, in the spirit of the Parisi ansatz. This definition allows for what we call trivial RSB (e.g., in a two-state picture), but corresponds to the more familiar meaning when many pure states are present in Γ. The presence of RSB means that when one decomposes each mixed Γ in Eq. (29) into pure states, then the permutation symmetry between different replicas is lost in each of the products where a pure state is chosen for each replica. It follows that RSB and RNI are distinct phenomena, and either can occur without the other.

Although we have a new way of constructing a replica measure, we can still take overlaps in the usual way, i.e., according to Eq. (16). The distribution of an overlap Q, though, depends on how σ and σ' are chosen. Because of the possibility of RNI, we no longer take overlaps (between one or more pairs of replicas) from the product measure $\rho_{\mathcal{J}}(\sigma^1)\rho_{\mathcal{J}}(\sigma^2)\cdots$, but instead from the more suitable replica measure $\rho^\infty_{\mathcal{J}}$. Because of this, the nature of the overlaps changes. For example, the distribution of a single overlap Q is no longer the $P_{\mathcal{J}}(q)$ obtained from $\rho_{\mathcal{J}}(\sigma)\rho_{\mathcal{J}}(\sigma')$, but rather is $\int P_\Gamma(q)\kappa_{\mathcal{J}}(\Gamma)d\Gamma$, where $P_\Gamma(q)$ denotes the overlap distribution obtained from $\Gamma(\sigma)\Gamma(\sigma')$. When $\kappa_{\mathcal{J}}$ is dispersed, P_Γ may or may not depend on the Γ chosen from $\kappa_{\mathcal{J}}$. (It does not in the chaotic pairs picture but does in the nonstandard SK picture discussed below.) Information on this dependence is contained in the overall ''overlap structure,'' by which we mean the joint distribution of all overlaps Q^{ij} between all pairs of replicas (σ^i,σ^j) from $\rho^\infty_{\mathcal{J}}$. This (possible) dependence on Γ is significant because, as in our analysis above of the standard SK picture, *the overlap structure still does not depend on* $\mathcal{J}$, by essentially the same arguments using translation invariance of the overlaps and translation ergodicity of the coupling distribution $\nu(\mathcal{J})$. More specifically, regarding P_Γ as random because of its dependence on Γ for fixed $\mathcal{J}$, the probability of appearance of a particular set of weights *and* corresponding locations of the overlap δ functions will not depend on $\mathcal{J}$. (Here, we are describing the situation, discussed at length in Sec. VII of the paper, in which P_Γ, for each Γ, has an SK-type form.)

In realistic systems, thermodynamic state observables can depend on the bulk couplings and/or on the couplings at infinity. Thus we observe that *there are two distinct types of dependence: (i) on* $\mathcal{J}$, *and (ii) on the state* Γ *within the metastate* κ *for fixed* $\mathcal{J}$. We have seen that replica overlaps cannot have the first type of dependence, but can in principle have the second kind. In that case, if one examines the same (finite) volume for two different coupling realizations, it could happen that two different sets of weights and overlap locations are seen (in the approximate sense corresponding to the fact that we're restricted here to finite volumes, so that, e.g., the finite-volume overlap distribution $P^{(L)}_{\mathcal{J}}$ is not a sum of exact δ functions). It is logically possible that in such a case, fluctuations in $P^{(L)}_{\mathcal{J}}$ persist for arbitrarily large L. From our previous discussions, however, it would be *incorrect* to conclude that there is an infinite-volume overlap distribution that is non-self-averaging (i.e., that depends on $\mathcal{J}$). Rather, it would imply that the limit $\lim_{L\to\infty}P^{(L)}_{\mathcal{J}}$ does not exist; i.e., that $P^{(L)}_{\mathcal{J}}$ exhibits chaotic size dependence. (See the Appendix for a discussion of the distinctions among differing constructions of overlap distributions.)

Our conclusion is that if overlap fluctuations (due to coupling dependence) do not vanish as $L\to\infty$ [82], this does *not* mean that the standard SK picture of overlap non-self-averaging holds; rather, it is a signal that the *second* kind of dependence holds for infinite volume.

With the approach outlined above, a replacement for the standard SK picture suggests itself. This replacement at first may seem very unusual and different from previous understandings of thermodynamic spin glass structure, but it falls out naturally from the ideas presented in this and in Sec. V. In Sec. VII we ask how can at least some of the familiar characteristics of the Parisi version of spin glass ordering be retained in realistic spin glasses? We will see that the ''maximal'' mean-field picture allowed, given our understanding of metastates and their consequences, has an intricate and novel thermodynamic structure.

VII. NONSTANDARD SK PICTURE

We saw in Sec. IV that the familiar thermodynamic picture usually associated with the Parisi ansatz applied to the EA model, which we called the standard SK picture, could not be valid in any dimension and at any temperature. Any thermodynamic theory of realistic spin glasses will differ considerably from this picture. The question then is whether and how any aspects of mean-field behavior can survive in such a theory. We now address this question.

We begin by asking what a maximal mean-field picture would look like in *finite* volume. There have been a number of numerical simulations (e.g., [42,61,72]) which appear to see a Parisi-like structure of finite-volume states, i.e., the appearance of several states with nonnegligible weight, several (approximate) δ function overlaps whose positions depend on coupling realization, and a Parisi-like $P(q)$ (i.e., δ functions at $\pm q_{\mathrm{EA}}$ connected by a continuous part) after averaging over the couplings. (See, however, [43] for a criticism of [42].) We will not attempt to resolve controversies associated with these or other simulations, nor will we speculate whether, if correct, these results persist for larger volumes. Rather, we ask *if* such results should hold for arbitrarily large volumes, what does that imply about the thermodynamics of spin glasses, given that the usual thermodynamic extrapolation of these finite-volume results (i.e., the standard SK picture) is incorrect?

We will see that the metastate approach allows us to construct such a thermodynamic scenario, which we call the *nonstandard SK picture*. This picture, or one closely related to it, must describe the thermodynamics of realistic spin glasses if the above *finite*-volume picture is correct. That is, the nonstandard SK picture allows for properties (1)–(4) appearing at the end of Sec. III (or more precisely, finite-volume versions of these properties) to hold in any fixed finite volume (with, e.g., periodic boundary conditions). It is therefore a maximal mean-field picture, as promised at the end of Sec. VI. However, the thermodynamics to which it gives rise is unconventional. It displays RSB and RNI (equivalently, CSD) and a type of non-self-averaging, suitably redefined (as described in Sec. VI). It does not have the features commonly thought of as associated with the Parisi ansatz, e.g., ultrametricity of all of the pure states [56,57], but displays some of its properties in a more limited fashion.

As a starting point, then, we require that in any (large) finite volume, the Gibbs state is an approximate decomposition over many pure states,

$$\rho_{\mathcal{J}}^{(L)}(\sigma) \approx \sum_{\alpha} W_{\mathcal{J}}^{\alpha,L} \rho_{\mathcal{J}}^{\alpha}(\sigma), \tag{31}$$

where a few states dominate the sum. From the metastate point of view, this implies that each Γ (chosen from $\kappa_{\mathcal{J}}$) is a mixed state with a nontrivial decomposition into pure states, namely,

$$\Gamma = \sum_{\alpha} W_{\Gamma}^{\alpha} \rho_{\mathcal{J}}^{\alpha}(\sigma), \tag{32}$$

and this decomposition is discrete but with many nonzero weights W_{Γ}^{α} [83].

In order that this scenario correspond to the usual expectations of the Parisi-SK picture in *finite* volumes (and at fixed $\mathcal{J}$), we require that the fixed-Γ overlap distribution

$$P_{\Gamma}(q) = \sum_{\alpha,\gamma} W_{\Gamma}^{\alpha} W_{\Gamma}^{\gamma} \delta(q - q^{\alpha\gamma}) \tag{33}$$

display the form consistent with property (2) listed at the end of Sec. III; that is, a countable sum of many δ functions. [Note that the occurrence of many (distinct) $q_{\alpha\gamma}$'s is an additional requirement, and does not follow automatically from Eq. (32).] The metastate must be an ensemble of many such Γ's (in fact, a continuum of them, as we explain below), each of which yields a pair of δ functions at $\pm q_{\mathrm{EA}}$, but with the locations of the remaining δ functions being Γ dependent. We further require that the locations of the δ functions within a specific $P_{\Gamma}(q)$ be ultrametric.

The above requirements are consistent with properties (1)–(4) of Sec. III holding for any finite volume, including (conventional) non-self-averaging for arbitrarily large L. However, instead of the straightforward extrapolation to infinite volumes characteristic of the standard SK picture, the thermodynamic properties of this nonstandard SK picture are considerably different. We now discuss what these properties are.

The crucial conceptual point is that the translation covariance of the metastate $\kappa_{\mathcal{J}}$ still requires that the resulting ensemble of overlap distributions is independent of $\mathcal{J}$. The metastate in this picture must be an ensemble of many Γ's, with a single Γ appearing in any fixed cube L^d (with, e.g., periodic boundary conditions). The dependence on Γ (as Γ varies within the metastate ensemble) is the new sort of non-self-averaging discussed at the end of Sec. VI. It is clear then that this picture must have both nontrivial RSB (because each Γ is a sum over many pure states), and CSD (and RNI) since the metastate is dispersed.

Finally, we require that the (averaged) Parisi order parameter $P(q)$ have the usual form, that is, two δ functions at $\pm q_{\mathrm{EA}}$, connected by a continuous component with nonzero weight everywhere in between; however, the averaging *must now be done over the states* Γ *within the metastate* $\kappa_{\mathcal{J}}$, all at fixed $\mathcal{J}$, rather than over $\mathcal{J}$ itself:

$$P(q) = [P_{\Gamma}(q)]_{\kappa_{\mathcal{J}}} = \int P_{\Gamma}(q) \kappa_{\mathcal{J}}(\Gamma) d\Gamma. \tag{34}$$

In order for this requirement to be valid along with discreteness of the individual P_{Γ}'s, it must be that there is a *continuum* of Γ's in the metastate ensemble. So we have replaced dependence on coupling realization $\mathcal{J}$ with dependence on the state Γ within the metastate for *fixed* $\mathcal{J}$.

We see that the nonstandard SK picture differs from the usual mean-field picture in several important respects. One is the lack of dependence of overlap distributions on $\mathcal{J}$, and the replacement of the usual sort of non-self-averaging with the concept of dependence on states within the metastate. Another important difference is that, in the nonstandard SK picture, a continuum of pure states *and* their overlaps must be present; therefore, *ultrametricity would not hold in general among any three pure states chosen at fixed* $\mathcal{J}$, unlike in the standard SK picture (see, for example, [56,57]). (The argument supporting this conclusion is presented in Ref. [25].) Rather, the pure states at fixed $\mathcal{J}$ are split up into (a continuum of) families, where each family consists of those pure states occuring in the decomposition of a particular Γ, and only within each such family would ultrametricity hold.

We have presented the nonstandard SK picture as a replacement for the more standard mean-field picture; if realistic spin glasses display any mean-field features, something like it must occur. However, this leaves open the question of what actually happens in realistic spin glasses. In particular, does the nonstandard SK picture actually occur? It turns out to have an important covariance property which may provide a clue.

For specificity, consider the EA model with a (mean zero, variance 1) Gaussian coupling distribution. Suppose that we change a *finite* number of couplings. The metastate $\kappa_{\mathcal{J}}(\Gamma)$, in addition to translation covariance, is also covariant with respect to this change [27]; that is, the ensemble transforms (as would any probability measure) under a change of variables $\Gamma \rightarrow \Gamma'$. Here, Γ' is the thermodynamic state with correlations $\langle \sigma_A \rangle_{\Gamma'} = \langle \sigma_A e^{-\beta \Delta H} \rangle_{\Gamma} / \langle e^{-\beta \Delta H} \rangle_{\Gamma}$, where ΔH is the change in the Hamiltonian. Under this change of variables, pure states remain pure and their overlaps do not change. However, the weights which appear in Eq. (33) *will* in general change. Nevertheless, the overall overlap structure (i.e., the probability of appearance of a given set of weights and overlap locations) must remain invariant.

We propose [49] this covariance property under coupling changes as an appropriate analog to that of of dynamical systems having a probability measure invariant under the dynamics. Our reasoning is as follows. Suppose we consider free b.c.'s. Changing from a cube of size L to one of size $L+1$ corresponds to taking a certain layer of couplings and changing them from zero to nonzero values. Having already made an analogy between L and the time t for the dynamical system, it seems appropriate to extend it to one between dynamical invariance ($t \rightarrow t+1$) and coupling covariance ($\mathcal{J} \rightarrow \mathcal{J}+\Delta\mathcal{J}$). The analogy is even clearer if we consider increasing volumes not by a whole layer at a time but by a single site at a time. Exploitation of this covariance property could result in a type of cavity method [29,66,67] for studying the properties of realistic spin glass models.

In the nonstandard SK picture, there seems every reason to expect nontrivial dependence of, e.g., $\langle e^{-\beta\Delta H}\rangle_{\mathcal{J}}^{\alpha}$ on the many α's appearing for each Γ. Thus, under changes of finitely many couplings, each P_Γ would be changed to a $P_{\Gamma'}$ with the same $q_{\alpha\gamma}$'s but with *different* weights. Nevertheless, by the translation-invariance–ergodicity argument mentioned earlier in this section, the *distribution* of the P_Γ's (as Γ varies over the metastate) in fact does not depend on $\mathcal{J}$ and hence is unchanged by $\mathcal{J} \rightarrow \mathcal{J}+\Delta\mathcal{J}$.

Thus the above covariance property under changes of couplings places a large number of constraints on the distribution of the P_Γ's that can arise in the nonstandard SK picture. We wonder whether all these constraints (which do *not* arise either in the droplet-scaling or in the chaotic pairs pictures) can actually be satisfied. Clearly, more study of this issue is needed.

VIII. CONCLUSIONS

We have shown that the traditional picture of spin glass thermodynamics, based on the Parisi ansatz as applied to finite-dimensional models, cannot hold for realistic spin glasses in any dimension and at any temperature. This standard SK picture is a natural and straightforward extrapolation to infinite volumes of the main features of spin glass ordering uncovered by Parisi and others for the SK model. It assumes a single infinite-volume overlap distribution function $P_{\mathcal{J}}(q)$ which is non-self-averaging, i.e., dependent on $\mathcal{J}$. This picture proposes that the pure states are chosen independently from some mixed (and, of course, non-self-averaged) thermodynamic state $\rho_{\mathcal{J}}$ with a decomposition of the form of Eq. (15) and that the resulting $P_{\mathcal{J}}(q)$ will consist of (many) discrete δ functions lying between a pair at $\pm q_{\mathrm{EA}}$. The locations of the δ functions (except for the pair at $\pm q_{\mathrm{EA}}$) and their weights depend on the coupling realization $\mathcal{J}$, but for any fixed $\mathcal{J}$ their locations are ultrametric. When averaged over the (uncountably many) coupling realizations chosen from the coupling distribution, the order parameter distribution $P(q)=\underline{P_{\mathcal{J}}(q)}$ shows the characteristic form of a continuous component connecting the δ functions at $\pm q_{\mathrm{EA}}$, and nonzero everywhere in between (at least for $0<T<T_c$).

We have shown, however, that this picture can never hold: *any $P_{\mathcal{J}}(q)$ with the weak (and physically reasonable) property of translation invariance must be self-averaging*, due to the underlying translation ergodicity of the coupling distribution. In Sec. IV we presented an explicit construction of such a non-self-averaged thermodynamic state $\rho_{\mathcal{J}}$, which obeyed the physically important requirement of translation covariance, and whose overlap distribution function was thus translation-invariant. We know of no other (natural) construction of a thermodynamic state for the EA model (which is measurable with respect to the disorder configuration), in the event that the spin glass does indeed possess many states (in which case chaotic size dependence must be taken into account) at some dimension and temperature.

Although we presented these results (and much of our other discussion on spin glasses) in the context of the EA model, we stress that they apply quite generally to most realistic spin glass models, because they depend only on general properties such as translation-invariance of the overlap function and translation-ergodicity of the underlying disorder distribution. These results lead, however, to an interesting approach to the thermodynamics of systems with many competing states that is far more general than considerations of spin glasses alone might indicate. The failure of the standard SK picture arises from the fact that if many pure states do exist for a particular system (at some dimension and temperature), then chaotic size dependence generally follows, and it becomes unreasonable to describe the thermodynamics through a single state—even though this state may be a mixture of infinitely many pure states—as in the standard approach. As an alternative, and based on the example of a chaotic dynamical system, we describe the thermodynamics through an *ensemble of states* (which may themselves be mixtures of pure states) that we call the metastate. Within that approach, the idea of replicas (whose correlations determine the metacorrelations of the metastate) becomes natural and formerly mysterious concepts—such as replica symmetry breaking—become clear. Further, the connections between these and more recent concepts such as replica nonindependence and dispersal of the metastate can be easily uncovered.

A crucial issue is the replacement of the old concept of non-self-averaging (as dependence on the bulk coupling realization) with a version of dependence on boundary conditions at infinity. This allows for the possibility that moments of Q, for example, as computed through the distribution $P_{\mathcal{J}}^{(L)}(q)$ in any finite volume can depend on $\mathcal{J}$ for arbitrarily large L—even though $P_{\mathcal{J}}(q)$ itself is independent of $\mathcal{J}$. Within the context of the nonstandard SK model, we replace the idea of dependence on $\mathcal{J}$ with dependence on the state Γ within the metastate for *fixed* $\mathcal{J}$. This notion corresponds, roughly speaking, to dependence on couplings at infinity (which yield a kind of annealed boundary condition at infinity) or to dependence on L, all for fixed $\mathcal{J}$.

Applying these results to the EA model, we find that several scenarios for the metastate remain as logical possibilities in various dimensions and temperatures. One of course is the trivial paramagnetic phase. Another is the scaling-droplet model. Two other possibilities, mentioned in Ref. [26], involve states Γ consisting of a continuum of pure states; in one of these scenarios the metastate is dispersed and in the other it is not, although both exhibit replica symmetry breaking. However, we see no evidence that either of these apply to realistic spin glasses, and so do not discuss them further here.

An intriguing new possibility, also discussed in Ref. [26], is the chaotic pairs picture, which is different from both droplet-scaling and mean-field pictures. This picture follows naturally from our earlier discussion on the metastate; it has infinitely many pure states, but with weights so mismatched in any finite volume (with, say, periodic boundary conditions) that only a pair of pure states (related by a global spin flip) appear. So in finite volumes this picture resembles droplet or scaling, but it has a very different thermodynamics; in particular, there are infinitely many pure state pairs and which of these appears in a given volume depends chaotically on L. It is interesting to note that this scenario actually arises in high dimensions in a highly disordered ground-state model of spin glasses [8].

Finally, we discussed a maximal mean-field picture called the nonstandard SK picture. This picture has features which resemble some of the familiar aspects of Parisi-type spin glass ordering in finite volume—and is consistent with various numerical simulations which claim to see this type of ordering—but has an unfamiliar thermodynamic structure and does not correspond to the usual picture presented in the literature. In particular, it does not possess nontrivial ultrametricity of all of the pure states corresponding to a fixed coupling realization $\mathcal{J}$; indeed, one of our results is that such ultrametricity cannot occur in any reasonable spin glass picture. It also does not possess non-self-averaging (in the sense of $\mathcal{J}$ dependence) of thermodynamic quantities related to the order parameter, such as the overlap distribution function.

Nevertheless, the features of non-self-averaging and ultrametricity could appear in any *finite* volume if this picture should hold. This leads to a further conclusion, namely, that for systems with quenched disorder (and for inhomogeneous systems in general) with many competing thermodynamic states, *properties which persist in large finite volumes cannot be straightforwardly extrapolated to a description of the thermodynamics*. In these cases, the metastate approach is indispensible for sorting out the thermodynamics.

In any case, we have serious reservations about the viability of the nonstandard SK picture. Although we cannot at this point rule it out on purely logical grounds, it requires an enormous number of constraints to be simultaneously satisfied. Further arguments suggesting that the nonstandard SK picture of Sec. VII is *not* valid (in any dimension) will appear in a future publication.

Further work is needed to determine which of these remaining pictures does hold for real spin glasses. Work is also needed to study the connections between the approach presented in this paper to systems in equilibrium and the dynamical behavior of systems out of equilibrium. Such investigations are currently in progress. We conclude by again pointing out that although we have concentrated in this paper on spin glasses, the phenomenon of thermodynamic chaos and the metastate approach to its analysis are potentially applicable to any thermodynamic system (disordered or not, inhomogeneous or not) in which there are many competing pure states and the finite-volume boundary conditions (or fields) are not (or cannot be) carefully chosen to favor one (or a very few) of them.

ACKNOWLEDGMENTS

The authors thank Aernout van Enter for many valuable comments on this work, and in particular for his illuminating example discussed in Ref. [88]. This research was partially supported by NSF Grant No. DMS-95-00868 (C.M.N.) and by DOE Grant No. DE-FG03-93ER25155 (D.L.S.).

APPENDIX: CONSTRUCTIONS OF OVERLAP DISTRIBUTIONS

In this appendix we briefly discuss different methods for constructing overlap distributions. Several approaches presently exist, and it may be the case that the use of different boundary conditions and/or limit procedures can lead to different $P_{\mathcal{J}}(q)$'s. This has led to some confusion in the literature (see, for example, the discussion in Refs. [85,86]) over the ''correct'' way to compute $P_{\mathcal{J}}(q)$. We emphasize at the outset that, although the actual form of the overlap distribution in various constructions can differ, our self-averaging arguments of Sec. IV apply to *all* of them. That is, if a particular construction of an overlap distribution has a well-defined thermodynamic limit at all (that does not depend on the choice of an origin), then it is necessarily self-averaging.

The fact that $P(q)$ for various models, including short-ranged spin glasses, can depend sensitively on the choice of boundary conditions was pointed out by Huse and Fisher in Ref. [38]. They further argued that $P(q)$ was too global a measurement to give reliably accurate information about numbers of pure states in many models (including relatively ''simple'' systems like conventional Ising ferromagnets). They provided examples where $P(q)$ could have a trivial structure in spite of the presence of many states, and other examples where it had a nontrivial structure in spite of the *absence* of many states. Notwithstanding these examples (which we believe to be correct), much of the literature on the subject persists in using $P(q)$ as an order parameter for spin glasses, and so it is necessary to sort out various subtleties which may arise in its use. [Arguments supporting the use of $P(q)$ to gain interesting thermodynamic information on spin glasses are presented in Ref. [72].]

In Ref. [25] we presented two different constructions (each of which yields a well-defined limit for the overlap distribution as $L\to\infty$). In these constructions, there are boxes, of sizes L_a, L_b, and L_c, with periodic b.c.'s imposed on the L_c box, fixed couplings in the L_b box, and overlaps computed in the L_a box. The first construction takes $1\ll L_b\ll L_a=L_c$, while the second construction, which is the one described in Sec. IV of this paper, takes $1\ll L_a\ll L_b\ll L_c$. The averaging (over the couplings in the region between L_b and L_c) is necessary to obtain a thermodynamic limit only when many pairs of states are present.

The first of these constructions is related to numerical computations which appear in the literature; it computes $P_{\mathcal{J}}(q)$ directly without prior construction of thermodynamic states. The second of these is of greater theoretical importance in the sense that it first computes a thermodynamic state (the average of the periodic b.c. metastate, as discussed in Sec. VI) which, if many pure states are present, will be a mixture of them. This second construction first takes the thermodynamic limit, and then takes replicas and overlaps; thus it averages over the couplings between L_b and L_c separately for each of the replicas. The first construction, as well as the metastate procedure of Sec. VI for constructing overlaps (which takes L_a, L_b, and L_c as in the second construction), uses the opposite order, taking replicas first and then the thermodynamic limit; thus replicas are taken with the same couplings in the entire L_c box. If RNI is present, this opposite order of operations could lead to a different $P_{\mathcal{J}}(q)$, as discussed in Sec. VI and Ref. [81]. The form of the periodic b.c. overlap distribution (using the construction of Sec. VI) for the nonstandard SK picture was discussed in Sec. VII; here we will mostly focus on the two constructions discussed above.

In agreement with the arguments of Huse and Fisher [38], we noted in Ref. [25] that these two constructions probably give different limit overlap distributions for the random-field Ising model [87] and for the highly disordered spin glass model of Ref. [8] in high dimensions. Furthermore, they possibly would also yield different distributions for the usual finite-dimensional spin glasses if something like the nonstandard SK picture, described in Sec. VII, holds.

In the case of the random-field Ising model, the first construction should yield a single δ-function spike even for $T<T_c$ because, for each volume, only one of the two pure states will appear (although which one will depend on the volume). This situation is illustrated in Fig. 1(b) of Ref. [38]. The second construction, however, should yield the expected pair of δ functions symmetrically placed with respect to the origin, because here the replicas are drawn from the thermodynamic mixed state which is an equal mixture of the magnetization up and down pure states.

Our second example uses the highly disordered ground-state model of Ref. [8] in dimensions high enough so that there exists an uncountable number of ground states. The situation here is essentially the ''chaotic pairs'' picture of Sec. V. In any specified volume (as usual, with periodic boundary conditions), only a single pair of ground states, related by a global spin flip, appears, but which pair is present changes chaotically with volume. Therefore, we expect the first construction to yield only a single pair of δ functions at $\pm q_{\mathrm{EA}}$, as in the Fisher-Huse droplet picture. (This would also be the case for a construction like that of Sec. VI.) However, we believe that the second construction will yield a *single* δ-function spike at zero. This is because each ground state can be thought of as consisting of an infinite set of invasion ''trees'' (see Refs. [8] and [9] for details) of rigidly coupled spins. Each tree can exist in one of two configurations which are global flips of each other. A ground state is therefore an assignment of a particular choice to each tree; the set of ground states is the collection of all possible assignments. The average $\rho_{\mathcal{J}}$ of the metastate should be a uniform distribution over the uncountably many ground states (in the sense of independent tosses of a fair coin for each tree).

The overlap of any two ground states chosen at random (i.e., according to this uniform distribution and independently of each other) can therefore be related to the overlap of two Bernoulli coin-tossing processes. Since half of all flips will agree and half will disagree for almost every pair of process realizations, the overlap should be zero and therefore we expect that $P_{\mathcal{J}}(q)=\delta(q)$ for almost every $\mathcal{J}$. (We emphasize that this argument is only heuristic, and in particular assumes that no trees contribute too much to the overlap, which is reasonable but has not been proven at this time.)

This last example provides a nice illustration of the contention of Huse and Fisher that a trivial overlap function (in this example, a single δ function) can mask the existence of more than one pure state (in this example, the extreme case of uncountably many). It also illustrates the contention of the authors in Refs. [25] and [86] that $P(q)$ may be a poor choice of order parameter in finite dimensional models.

The nonstandard SK picture presented in Sec. VII provides a more interesting possibility. Here there may also be two forms of $P_{\mathcal{J}}(q)$, depending on which construction is used (but both forms are still self-averaging). The picture was created so that the overlap distribution constructed according to Sec. VI consists of a collection of δ functions of varying weights lying between a pair at $\pm q_{\mathrm{EA}}$. While it is less clear what will occur in this case in each of the two constructions under discussion, a likely outcome (if this picture were to hold) is that the first construction yields the familiar Parisi $P(q)$—i.e., a continuous component connecting the pair of δ functions at $\pm q_{\mathrm{EA}}$. However, the second construction should then result either in $\delta(q)$ or else in a continuous distribution with nonzero density between $\pm q_{\mathrm{EA}}$, but no δ functions at $\pm q_{\mathrm{EA}}$. The reason for no δ functions at $\pm q_{\mathrm{EA}}$ is that here states are chosen from a distribution (i.e., from a metastate) containing uncountably many such states, so that the probability of two states being the same or related by a global spin flip is zero.

So far we have discussed the consequences of the two constructions proposed in Ref. [25], mostly using periodic boundary conditions. (Different boundary conditions are discussed in Secs. V and VI.) In a comment on Ref. [25], Parisi [85] proposed two constructions, which are further discussed in a reply from the authors [86]. Parisi's constructions (denoted as $P_I^{(1)}$ and $P_I^{(2)}$ in his comment) are modifications of the two constructions of Ref [25]. $P_I^{(1)}$ is similar to the first construction (with replicas having the same couplings in the entire L_c box), but takes $L_a=L_b \ll L_c$. $P_I^{(2)}$ is likewise similar to the second construction (with separate averaging of couplings between L_b and L_c for each of the replicas) but again takes $L_a=L_b \ll L_c$. Once more, finite-size effects might lead, in each case, to overlap distributions different from the corresponding constructions in Ref. [25]. Parisi asserts that $P_I^{(2)}=\delta(q)$ for realistic spin glasses, which may be the case for some of the pictures presented in this paper, although it is not clear whether it holds for all (particularly the nonstandard SK picture). For a more detailed discussion, see Refs. [85] and [86].

In our work, we have mostly chosen constructions where overlaps are computed in boxes which are small compared to the box on which boundary conditions are placed. We believe this to be essential if one is trying to understand the microscopic structure of the pure states, in particular, what the spin configurations and their correlations and overlaps look like in the neighborhood of the origin [88]. In this appendix, however, we have emphasized that different procedures may well lead to different overlap distributions (and in models better understood than the EA and related models, this appears often to be the case.) It remains an important issue, to be eventually resolved, whether in the EA model the different overlap procedures discussed do in fact lead to different overlap distributions—indeed, this may be a better signature of many states than whether a *single* procedure gives rise to a complicated or trivial $P(q)$. In all cases, however, a given procedure leading to a thermodynamic overlap distribution function which has the weak property of translation invariance will always be self-averaging.

[1] For a recent review of a variety of experimental tests, see *Recent Progress in Random Magnets*, edited by D. H. Ryan (World Scientific, Singapore, 1992).
[2] K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).
[3] S. Edwards and P. W. Anderson, J. Phys. F **5**, 965 (1975).
[4] D. Sherrington and S. Kirkpatrick, Phys. Rev. Lett. **35**, 1972 (1975).
[5] M. E. Fisher and R. R. P. Singh, in *Disorder in Physical Systems*, edited by G. Grimmett and D. J. A. Welsh (Clarendon, Oxford, 1990), p. 87.
[6] E. Marinari, G. Parisi, and F. Ritort, J. Phys. A **27**, 2687 (1994).
[7] M. J. Thill and H. J. Hilhorst, J. Phys. (France) I **6**, 67 (1996).
[8] C. M. Newman and D. L. Stein, Phys. Rev. Lett. **72**, 2286 (1994).
[9] C. M. Newman and D. L. Stein, J. Stat. Phys. **82**, 1113 (1996).
[10] R. G. Palmer, Adv. Phys. **31**, 669 (1982).
[11] P. Sibani and J.-O. Andersson, Physica A **206**, 1 (1994).
[12] D. L. Stein and C. M. Newman, Phys. Rev. E **51**, 5228 (1995).
[13] M. Lederman, R. Orbach, J. M. Hamann, M. Ocio, and E. Vincent, Phys. Rev. B **44**, 7403 (1991).
[14] Y. G. Joh, R. Orbach, and J. M. Hamann (unpublished).
[15] P. Refrigier, E. Vincent, J. Hamman, and M. Ocio, J. Phys. (Paris) **48**, 1533 (1987).
[16] G. J. M. Koper and H. J. Hilhorst, J. Phys. (Paris) **49**, 249 (1988).
[17] P. Sibani and K.-H. Hoffmann, Phys. Rev. Lett. **63**, 2853 (1989).
[18] K.-H. Hoffmann and P. Sibani, Z. Phys. B **80**, 429 (1990).
[19] P. Svedlinh, K. Gunnarson, J.-O. Andersson, H. A. Katori, and A. Ito, Phys. Rev. B **46**, 13 687 (1992).
[20] J. P. Bouchaud, J. Phys. (France) I **2**, 1705 (1992).
[21] F. Lefloch, J. Hamann, M. Ocio, and E. Vincent, Europhys. Lett. **18**, 647 (1992).
[22] H. Rieger, J. Phys. A **26**, L615 (1993).
[23] S. Franz and M. Mézard, Physica A **210**, 48 (1994).
[24] E. Vincent, J. P. Bouchaud, D. S. Dean, and J. Hamann, Phys. Rev. B **52**, 1050 (1995).
[25] C. M. Newman and D. L. Stein, Phys. Rev. Lett. **76**, 515 (1996).
[26] C. M. Newman and D. L. Stein, Phys. Rev. Lett. **76**, 4821 (1996).
[27] M. Aizenman and J. Wehr, Commun. Math. Phys. **130**, 489 (1990).
[28] C. M. Newman and D. L. Stein, Phys. Rev. B **46**, 973 (1992).
[29] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
[30] G. Parisi, Phys. Rev. Lett. **43**, 1754 (1979).
[31] G. Parisi, Phys. Rev. Lett. **50**, 1946 (1983).
[32] A. Houghton, S. Jain, and A. P. Young, J. Phys. C **16**, L375 (1983).
[33] M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, Phys. Rev. Lett. **52**, 1156 (1984).
[34] W. L. McMillan, J. Phys. C **17**, 3179 (1984).
[35] A. J. Bray and M. A. Moore, Phys. Rev. Lett. **58**, 57 (1987).
[36] D. S. Fisher and D. A. Huse, Phys. Rev. Lett. **56**, 1601 (1986).
[37] D. S. Fisher and D. A. Huse, Phys. Rev. B **38**, 386 (1988).
[38] D. A. Huse and D. S. Fisher, J. Phys. A **20**, L997 (1987).
[39] D. S. Fisher and D. A. Huse, J. Phys. A **20**, L1005 (1987).
[40] A. van Enter, J. Stat. Phys. **60**, 275 (1990).
[41] J. R. L. de Almeida and D. J. Thouless, J. Phys. A **11**, 983 (1978).
[42] S. Caracciolo, G. Parisi, S. Patarnello, and N. Sourlas, J. Phys. (Paris) **51**, 1877 (1990).
[43] D. A. Huse and D. S. Fisher, J. Phys. (France) I **1**, 621 (1991).
[44] M. Weissman, Rev. Mod. Phys. **65**, 829 (1993).
[45] A. Bovier and J. Fröhlich, J. Stat. Phys. **44**, 347 (1986).
[46] Infinite-volume Gibbs measures $\rho_{\mathcal{J},\beta}$ can also be characterized, without such a limiting process, as probability measures (on infinite-volume spin configurations) which satisfy the Dobrushin-Lanford-Ruelle (DLR) equations. For a mathematically detailed presentation, see [47].
[47] H. O. Georgii, *Gibbs Measures and Phase Transitions* (de Gruyter Studies in Mathematics, Berlin, 1988).
[48] A. C. D. van Enter and J. L. van Hemmen, Phys. Rev. A **29**, 355 (1984).
[49] C. M. Newman and D. L. Stein, in *Mathematics of Spin Glasses and Neural Networks*, edited by A. Bovier and P. Picco (Birkhäuser, Boston, in press).
[50] C. M. Newman, *Topics in Disordered Systems* (Birkhäuser, Basel, in press).
[51] N. Wiener, Duke Math. J. **5**, 1 (1939).
[52] M. Mézard and G. Parisi, J. Phys. (France) I **1**, 809 (1991).
[53] J.-P. Bouchaud, M. Mezard, and J. S. Yedidia, Phys. Rev. Lett. **67**, 3840 (1991).
[54] J. S. Yedidia, in *1992 Lectures in Complex Systems*, edited by D. L. Stein (Addison-Wesley, Reading, MA, 1993), p. 299.
[55] G. Parisi, Physica A **194**, 28 (1993).
[56] E. Vincent, J. Hammann, and M. Ocio, in *Recent Progress in Random Magnets* Ref. [1], p. 207.
[57] D. Badoni, J. C. Ciria, G. Parisi, F. Ritort, J. Pech, and J. J. Ruiz-Lorenzo, Europhys. Lett. **21**, 495 (1993).
[58] S. Franz, G. Parisi, and M. A. Virasoro, J. Phys. (France) I **4**, 1657 (1994).
[59] F. Ritort, Phys. Rev. B **50**, 6844 (1994).
[60] P. Le Doussal and T. Giamarchi, Phys. Rev. Lett. **74**, 606 (1995).
[61] E. Marinari, G. Parisi, J. J. Ruiz-Lorenzo, and F. Ritort, Phys. Rev. Lett. **76**, 843 (1996).
[62] R. Brout, Phys. Rev. **115**, 824 (1959).
[63] M. Kac, NORDITA, Report No. 286, 1968 (unpublished).
[64] S. F. Edwards, in *Proceedings of the Third International Conference on Amorphous Materials*, edited by R. W. Douglas and B. Ellis (Wiley, New York, 1970), p. 279; also in *Polymer Networks*, edited by A. J. Chompff and S. Newman (Plenum, New York, 1971), p. 83.
[65] A. Blandin, J. Phys. (Paris) Colloq. **39**, C6-1499 (1978).
[66] B. Derrida and G. Toulouse, J. Phys. (Paris) Lett. **46**, L223 (1985).
[67] M. Mézard, G. Parisi, and M. A. Virasoro, Europhys. Lett. **1**, 77 (1986).
[68] D. J. Thouless, P. W. Anderson, and R. G. Palmer, Philos. Mag. **35**, 593 (1977).
[69] A. J. Bray and M. A. Moore, J. Phys. C **13**, L469 (1980).
[70] A. J. Bray and M. A. Moore, J. Phys. A **14**, L377 (1981).
[71] N. D. Mackenzie and A. P. Young, Phys. Rev. Lett. **49**, 301 (1982).
[72] J. D. Reger, R. N. Bhatt, and A. P. Young, Phys. Rev. Lett. **64**, 1859 (1990).

[73] M. Aizenman, Commun. Math. Phys. **73**, 83 (1980).

[74] Y. Higuchi, in *Random Fields*, edited by J. Fritz, J. L. Lebowitz, and D. Szász (North-Holland, Amsterdam, 1979), Vol. I, p. 517.

[75] A. Gandolfi, M. Keane, and C. M. Newman, Prob. Theory Relat. Fields **92**, 511 (1992).

[76] F. Ledrappier, Commun. Math. Phys. **56**, 297 (1977).

[77] F. Comets, Prob. Theory Relat. Fields **80**, 407 (1989).

[78] T. Seppäläinen, Commun. Math. Phys. **171**, 233 (1995).

[79] C. Külske, in *Mathematics of Spin Glasses and Neural Networks*, edited by A. Bovier and P. Picco (Birkhäuser, Boston, in press).

[80] M. Cieplak, A. Maritan, and J. R. Banavar, Phys. Rev. Lett. **72**, 2320 (1994).

[81] F. Guerra (private communication).

[82] This was proved for the SK model by L. Pastur and M. Shcherbina, J. Stat. Phys. **62**, 1 (1991). See F. Guerra, Int. J. Mod. Phys. B **10**, 1675 (1996) for further rigorous results on the SK model overlap distribution.

[83] We note, though, that what is actually required is discreteness of the *overlap distribution*, and it has been pointed out [84] that this could be the case without discreteness of the pure state decomposition Eq. (32). However, there is another feature of the standard SK picture which seems to require discreteness of at least the low-lying part of the energy (or free energy) spectrum of pure states. This is the occurrence of energy (or free energy) gaps of order unity between the low-lying states in any (large) volume, accompanied by an exponentially increasing density of states near the bottom of the spectrum [66,67]. For the remainder of this paper, we will assume a countable pure state decomposition.

[84] A. C. D. van Enter, A. Hof, and J. Miękisz, J. Phys. A **25**, L1133 (1992).

[85] G. Parisi (unpublished).

[86] C. M. Newman and D. L. Stein (unpublished).

[87] A. C. D. van Enter (private communication).

[88] An interesting and simple illustration of this has been suggested by A. van Enter (private communication) that extends an earlier example of Huse and Fisher [38]. He considers the overlap distribution of an Ising antiferromagnet in two dimensions with periodic boundary conditions. For odd-sized squares the overlap is equivalent to that of the ferromagnet with periodic boundary conditions, and for even-sized squares it is equivalent to that of the ferromagnet with antiperiodic boundary conditions. If the overlap distribution were computed in the full volume, it would therefore oscillate between two different answers, an example of CSD for *overlap* distributions. On the other hand, our choice of computing overlaps in boxes which are much smaller than the system size would give rise in this example to a well-defined answer—i.e., the two-δ-function overlap distribution of the periodic ferromagnet.