Simplicity of state and overlap structure in finite-volume realistic spin glasses

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

D. L. Stein
Departments of Physics and Mathematics, University of Arizona, Tucson, Arizona 85721

(Received 25 September 1997)

We present a combination of heuristic and rigorous arguments indicating that both the pure state structure and the overlap structure of realistic spin glasses should be relatively simple: in a large finite volume with coupling-independent boundary conditions, such as periodic, at most a pair of flip-related (or the appropriate number of symmetry-related in the non-Ising case) states appear, and the Parisi overlap distribution correspondingly exhibits at most a pair of δ functions at $\pm q_{\mathrm{EA}}$. This rules out the nonstandard mean-field picture introduced by us earlier, and when combined with our previous elimination of more standard versions of the mean-field picture, argues against the possibility of even limited versions of mean-field ordering in realistic spin glasses. If broken spin-flip symmetry should occur, this leaves open two main possibilities for ordering in the spin glass phase: the droplet-scaling two-state picture, and the chaotic pairs many-state picture introduced by us earlier. We present scaling arguments which provide a possible physical basis for the latter picture, and discuss possible reasons behind numerical observations of more complicated overlap structures in finite volumes. [S1063-651X(98)07202-X]

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

I. INTRODUCTION

Prevalent scenarios [1,2] concerning realistic spin glasses require that the nature of the spin glass order parameter (i.e., the Parisi overlap distribution) and the structure of the thermodynamic states from which it is obtained be highly complex; see, for example, Refs. [3–19]. This complexity is asserted to be a consequence of the existence of many competing pure states. In previous papers [20–25] we showed that the standard picture of this complex structure (including non-self-averaging of the thermodynamic overlap distribution function, ultrametricity of distances among all pure states for fixed coupling realization, etc.) cannot hold in any finite dimension. However, at the same time we presented (as a logical possibility) a *nonstandard* mean-field picture in which some of these features appear in finite-dimensional spin glasses but in a more limited sense — i.e., in large finite volumes with coupling-independent boundary conditions such as periodic. In this picture, only a *subset* of all the pure states appears in each finite-volume mixed state (which varies with volume); those pure states along with their weights and overlaps retain some mean-field structure.

In this paper, however, we provide both heuristic and rigorous arguments that indicate the state and overlap structure in finite volumes must in fact be relatively simple. This is so even if there are many pure states overall. These arguments preclude the possibility of any type of mean-field structure — even the nonstandard, limited type — for the spin glass phase in finite dimensions.

Although the arguments and conclusions of this paper are applicable to fairly general examples of disordered systems, we will focus on the Edwards-Anderson (EA) Ising spin glass [26]. When there are many pure (infinite volume) states ρ^{α}, it has been generally believed [1] that the finite-volume Gibbs state $\rho_{\mathcal{J}}^{L}$ (for a coupling configuration $\mathcal{J}$ in the cube Λ_L of side L centered at the origin with, say, periodic boundary conditions) is (approximately) a mixture of many pure states:

$$\rho_{\mathcal{J}}^{(L)} \approx \sum_{\alpha} W_{\mathcal{J},L}^{\alpha} \rho_{\mathcal{J}}^{\alpha} \tag{1}$$

and the finite-volume overlap distribution $P_{\mathcal{J}}^{L}(q)$ is (approximately) the corresponding mixture of many δ functions:

$$P_{\mathcal{J}}^{L}(q) \approx \sum_{\alpha,\gamma} W_{\mathcal{J},L}^{\alpha} W_{\mathcal{J},L}^{\gamma} \delta(q - q_{\mathcal{J}}^{\alpha\gamma}), \tag{2}$$

where $q_{\mathcal{J}}^{\alpha\gamma}$ is the overlap between the pure states α and γ:

$$q_{\mathcal{J}}^{\alpha\gamma} = \lim_{L' \to \infty} |\Lambda_{L'}|^{-1} \sum_{x \in \Lambda_{L'}} \langle \sigma_x \rangle^{\alpha} \langle \sigma_x \rangle^{\gamma}. \tag{3}$$

Of course, if there is only a single pair of pure states (related by a global spin flip) as in the droplet-scaling picture of Refs. [27–30] (see also [31,32]), then for each L, $P_{\mathcal{J}}^{L}(q)$ will simply approximate a sum of two δ functions at $\pm q_{\mathrm{EA}}$. We will argue here that *the same conclusion is true for the finite-volume overlap distributions even when there are many pure states*. This is because $\rho_{\mathcal{J}}^{L}$ will still be approximately a mixture of a *single* pair of pure states, although now the choice of the pair will depend upon L. This scenario was previously proposed in Refs. [21–23] as a logical possibility that followed from the metastate approach introduced in those papers. Here we will argue that it is the only *reasonable* possibility consistent with many pure states, and we will also present scaling arguments that provide a physical basis for it and at the same time explain its relation to the earlier and simpler two-state droplet-scaling picture.

It is important to note that in computing overlap distributions as in Eq. (2), the region in which the computation is done should be small compared to the overall size of the system — i.e., the system boundaries should be far from the region of interest. The reasons for this were discussed at some length in the Appendix to Ref. [22], and will be returned to in Sec. VI. This guarantees that one is focusing on the thermodynamic states of the system [22,28] and avoiding extraneous finite size and boundary effects.

With this understanding, our arguments indicate that the nonstandard Sherrington-Kirkpatrick (SK) picture, introduced by us previously as the only remaining viable mean-field-like picture, is not valid in any dimension. The reader may wish to glance ahead at Sec. IV in which this conclusion, one of the main results of the paper, is presented.

The plan of the paper is as follows. In Sec. II we review the concept of metastates. In Sec. III we discuss previously proposed scenarios for the spin glass phase, including the newer chaotic pairs and nonstandard SK pictures. In Sec. IV we present the first of our main results, a theorem on the invariance of the metastate with respect to flip-related boundary conditions, and then discuss the consequences of the theorem. We will see why this result should be incompatible with any but the simplest spin glass ordering, and in particular how that argues against the nonstandard SK picture. In Sec. V we will provide a scaling basis for the chaotic pairs picture, and present one possible physical scenario under which it would occur. In Sec. VI we discuss, in light of our results, the question of why some numerical experiments appear to see a complicated overlap structure. We further discuss appropriate procedures for computing overlap structures in finite volumes as a means of extracting at least partial information on ordering in the low-temperature phase. Finally, in Sec. VII we present our conclusions.

II. METASTATES

For specificity, we will focus on the Edwards-Anderson model [26] which, on a cubic lattice in d dimensions, is described by the Hamiltonian

$$\mathcal{H}_{\mathcal{J}}(\sigma) = -\sum_{\langle x,y\rangle} J_{xy}\sigma_x\sigma_y\,, \tag{4}$$

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 one). The most common examples are the Gaussian and $\pm J$ distributions. The infinite-ranged version of the EA model was introduced by Sherrington and Kirkpatrick [33] and is commonly referred to as the SK model.

Numerical studies of spin glass overlap structure in the EA model study finite-volume cubes with (usually) periodic boundary conditions [19,34,35]. A crucial property of disordered systems with many competing states is that, although particular pure states may be picked out by a special choice of boundary conditions depending on the disorder realization, such boundary conditions are not relevant for comparison either to experiments on physical spin glasses or to numerical simulations. In all of these cases boundary conditions are chosen *independently* of the coupling realization.

In this paper we will therefore focus on either fixed or periodic boundary conditions (BC's) (and their flip-related BC's; see Sec. IV) chosen independently of the couplings. From a theoretical point of view, observable properties in this situation are amenable to analysis by means of the metastate approach [21–25].

Metastates enable us to relate the observed behavior of a system in large but finite volumes with its thermodynamic properties. This relation is relatively straightforward for systems with few pure states or for those whose states are related by well-understood symmetry transformations; but in the presence of many pure states not related by any obvious transformations, this relation may be subtle and complex. In these cases the metastate approach may be highly useful.

One reason for this is that, in the presence of many competing pure states, a sequence as $L\to\infty$ of finite-volume Gibbs measures on cubes Λ_L with coupling-independent BC's will generally *not* converge to a single limiting thermodynamic state [36]. We call this phenomenon *chaotic size dependence* (CSD). In the metastate approach, we exploit the presence of CSD by replacing the study of single thermodynamic states (as is conventionally done) with an *ensemble* of (pure or mixed) thermodynamic states. This approach, based on an analogy to chaotic dynamical systems, enables us to construct a limiting measure, but it is a measure on the thermodynamic states themselves.

This (infinite-volume) measure contains far more information than any single thermodynamic state. It has a particular usefulness in the context of the study of finite volumes because it carries — among other things — the following information. Suppose that there exist many thermodynamic states in some (fixed) dimension and at some (fixed) temperature. Then (for example) the periodic BC metastate (constructed from an infinite sequence of finite-volume Gibbs measures on cubes with periodic boundary conditions) tells us the likelihood of appearance of any specified thermodynamic state, pure or mixed, in a typical large volume. More precisely, it provides a probability measure for all possible $1,2,\ldots,n$-point correlation functions contained in a box centered at the origin whose sides are sufficiently far from any of the boundaries so that finite size effects do not appreciably affect the result.

Details on the construction and properties of the metastate were given in previous papers [21–23]. Here we simply recount some central features. The histogram, or empirical distribution approach, is a type of microcanonical ensemble which considers at fixed $\mathcal{J}$ a sequence of volumes with specified BC, such as periodic. The resulting sequence of finite-volume Gibbs states $\rho_{\mathcal{J}}^{(L_1)},\rho_{\mathcal{J}}^{(L_2)},\ldots,\rho_{\mathcal{J}}^{(L_N)}$ each is given weight N^{-1}. This ''histogram'' of finite-volume Gibbs states converges to some $\kappa_{\mathcal{J}}$ as $N\to\infty$. The (periodic BC, in this example) metastate $\kappa_{\mathcal{J}}$ is a probability measure on thermodynamic states Γ at fixed $\mathcal{J}$, and specifies the fraction of cube sizes that the system spends in each different (possibly mixed) thermodynamic state Γ [37].

An alternative (and earlier) construction of the metastate, in which the randomness of the couplings is used directly to generate an ensemble of states, was provided by Aizenman and Wehr [38]. In this approach one considers the limiting joint distribution $\mu(\mathcal{J},\rho_{\mathcal{J}}^{(L)})$ as $L\to\infty$. Technical details can be found in [22–24,38].

It can be proved that there exists at least a $\mathcal{J}$-independent subsequence of volumes along which the two approaches (empirical distribution and Aizenman-Wehr) yield the *same* limiting metastate [22–24]. This will be important in what follows [39].

Occasionally a distinction is drawn between finite- and infinite-volume states (see, for example, [19]), where it is argued that the first is more physical and the second merely mathematical in nature. While we have shown [22] that the relation between the two may be more subtle than previously realized — at least in the case where many competing states are present — we also argue that the distinction drawn above is misleading. Indeed, it should be clear from the discussion above that the metastate approach is specifically constructed to consider both finite and infinite volumes together and to unify the two cases. In the next section, guided by this approach, we review various allowable scenarios for the EA spin glass phase.

III. THE FINITE-DIMENSIONAL SPIN GLASS PHASE

Of the possible scenarios for spin glasses at low temperature, the simplest is that spin-flip symmetry is not broken at positive temperatures in any dimension. This would be the case if there were no phase transition at all and the paramagnetic state persisted to arbitrarily low temperatures. It would also be the case if there *were* a phase transition but the EA order parameter $q_{\rm EA}$ [corresponding to the self-overlap of a pure state, i.e., $q_{\mathcal{J}}^{\alpha\alpha}$ in Eq. (3)] remained zero. Such a phase might have, e.g., single-site magnetizations equalling zero at low temperatures but two-spin correlation functions decaying as a power law at large distances.

More likely, however, is that spin-flip symmetry *is* broken for $d>d_0$ and $T<T_c(d)$ [2]. In that case the simplest scenario for the low-temperature spin glass phase is the Fisher-Huse scaling-droplet picture [27–30] (see also [31,32]), in which a single pair of pure states is present. In that case, with periodic BC's, CSD is absent, and the metastate is concentrated on a single mixed thermodynamic state, with each of the two pure states having weight 1/2. This picture seems internally consistent.

We now consider possible many-state pictures. In the standard SK picture, there is an overlap distribution $P_{\mathcal{J}}(q)$ that exhibits non-self-averaging (NSA) even after the thermodynamic limit has been taken [6–8]; that is, it fluctuates with $\mathcal{J}$ even though it is a thermodynamic quantity. Other features of this picture include ultrametricity among *all* pure state overlaps and a continuous part of $P(q)$ [the average of $P_{\mathcal{J}}(q)$ over all $\mathcal{J}$] between $\pm q_{\rm EA}$. For details, see [1].

However, this standard SK picture *cannot* hold (in any dimension and at any temperature) [20] because the translation invariance of $P_{\mathcal{J}}(q)$ combined with the translation ergodicity of the underlying distribution of couplings implies that $P_{\mathcal{J}}(q)$ must be self-averaged [40].

This problem with the standard SK picture might sound like a mere mathematical technicality — for which one might hope to find a technical solution. But in fact this picture has an inherent conceptual flaw — namely, the basic problem that a single-state $\rho_{\mathcal{J}}$ is simply not a rich enough description of the $L\to\infty$ behavior of a thermodynamic system where CSD occurs. In such a picture, one is in effect replacing with a single state all of the information contained in an entire distribution of states, i.e., the metastate. We now consider two nonstandard pictures, each of which arises naturally in the context of the metastate approach and the possible presence of CSD.

The first of these resembles the Fisher-Huse picture in finite volumes, but has a very different thermodynamic structure. It is a many-state picture, but unlike in the mean-field picture each large volume (with periodic boundary conditions) ''sees'' essentially only one pair of states at a time (in Sec. VI we discuss what it means for a finite volume to ''see'' a thermodynamic state, pure or mixed). In other words, for large L, one finds that

$$\rho_{\mathcal{J}}^{(L)}\approx\frac{1}{2}\rho_{\mathcal{J}}^{\alpha_L}+\frac{1}{2}\rho_{\mathcal{J}}^{-\alpha_L}, \tag{5}$$

where $-\alpha$ refers to the global spin flip of pure state α. Here, the pure state pair (of the infinitely many present) appearing in finite volume depends chaotically on L. Unlike the droplet-scaling picture, this new possibility exhibits CSD with periodic BC's. In this ''chaotic pairs'' picture the (periodic BC) metastate is dispersed over (infinitely) many Γ's, of the form $\Gamma=\Gamma^{\alpha}=\frac{1}{2}\rho_{\mathcal{J}}^{\alpha}+\frac{1}{2}\rho_{\mathcal{J}}^{-\alpha}$. The overlap distribution for each Γ is the same: $P_{\Gamma}=\frac{1}{2}\delta(q-q_{\rm EA})+\frac{1}{2}\delta(q+q_{\rm EA})$. Like the Fisher-Huse picture, this scenario also seems internally consistent. It is interesting to note that a highly disordered spin glass model [41,42] (see also [43]) appears to display just this behavior in its ground state structure in sufficiently high dimension.

The last picture we discuss is a nonstandard SK-like picture that resembles the standard SK picture in finite volumes, but has an altogether different thermodynamic structure. This picture, which also assumes infinitely many pure states, organizes them such that each $\Gamma=\Sigma_{\alpha}W_{\mathcal{J}}^{\alpha}\rho_{\mathcal{J}}^{\alpha}$. The metastate $\kappa_{\mathcal{J}}$ is dispersed over many such Γ's, so that different Γ's again appear in different volumes, leading to CSD. Unlike the chaotic pairs picture, each P_{Γ} depends on Γ (because each Γ is now itself a nontrivial mixture of infinitely many pure states). However, the ensemble of P_{Γ}'s (like the single $P_{\mathcal{J}}$ of the standard SK picture) does *not* depend on $\mathcal{J}$ (again because of translation invariance and ergodicity). So the conventional meaning of NSA — thermodynamic quantities such as the overlap distribution depending on $\mathcal{J}$ — is replaced by a new notion: not dependence on $\mathcal{J}$ but rather dependence on the state Γ within the metastate for fixed $\mathcal{J}$. Moreover, ultrametricity of overlaps among pure states may be present within individual Γ's, but not for all of the pure states taken together. A more detailed description of this nonstandard SK picture is given in Refs. [21–23].

Given the results of [20], the nonstandard SK picture is the only remaining viable mean-field-like picture. We have presented preliminary arguments (based on the invariance of the ensemble of P_Γ's with respect to $\mathcal{J}$; we refer the reader to Ref. [22] for details) that already cast some doubt on its validity, by demonstrating that the nonstandard SK picture requires an enormous number of constraints to be simultaneously satisfied. In the next section we present further arguments that more strongly rule it out as a viable possibility.

IV. INVARIANCE OF THE METASTATE

The main result of this section is a theorem on the invariance of the metastate $\kappa_{\mathcal{J}}$ with respect to boundary conditions that are flip related. Two (sequences of) BC's are flip related if, for each finite L, there is some subset of the boundary $\partial\Lambda_L$ whose flip transforms one BC for that L into the other. An obvious example of flip-related boundary conditions are periodic and antiperiodic; a second example is any two fixed boundary conditions, i.e., where each spin on the boundary is specified. On the other hand, periodic and fixed BC's are not flip related.

In the following theorem we continue to assume that the common distribution of the couplings J_{xy} is symmetric about zero, i.e., that J_{xy} has the same distribution as $-J_{xy}$, and that the external field is zero.

Theorem. Consider two metastates constructed (at fixed, arbitrary dimension and temperature, and using either the histogram method or the Aizenman-Wehr method) using two different boundary conditions, with neither depending on $\mathcal{J}$, on an infinite ($L_N \to \infty$) sequence of cubes Λ_{L_N}. If the two different sequences of boundary conditions are flip related, then the two metastates are the same (with probability one — i.e., for almost every $\mathcal{J}$).

Proof. We use the fact, discussed above, that along some $\mathcal{J}$-independent subsequence of volumes both the histogram construction of metastates and the Aizenman-Wehr construction have a limit, and that limit is the same. Because the Aizenman-Wehr construction averages over couplings ''at infinity'' (for details, see Refs. [21–24,38]), it rigorously follows (using gauge transformation arguments like those used in the proof of Theorem 3 in Ref. [36]) that the two metastates must be the same.

This is a striking result (despite the brevity of the proof), with important physical consequences. It says, for example, that the periodic BC metastate $\kappa_{\mathcal{J}}$ must be the same as the *antiperiodic* BC metastate. In fact, if one were to choose (independently of $\mathcal{J}$) two *arbitrary* sequences of periodic and antiperiodic BC's, the metastates (with probability one) would *still* be identical. In other words, the metastate (and corresponding overlap distributions constructed from it) at fixed temperature and dimension is highly *insensitive* to boundary conditions.

To appreciate the implications of this, consider the histogram construction of the metastate. The invariance of the metastate with respect to different sequences of periodic and antiperiodic BC's means that the frequency of appearance (in finite volumes) of various thermodynamic states is (with probability one) *independent* of the choice of boundary conditions. Moreover, this same invariance property holds (with probability one) among any two sequences of *fixed* boundary conditions (and the fixed boundary condition of choice may even be allowed to vary arbitrarily along any single sequence of volumes). It follows that, with respect to changes of boundary conditions, the metastate is highly robust.

Of course, the insensitivity of the metastate with respect to changes of boundary conditions would be unsurprising if there were only a single thermodynamic state (e.g., paramagnetic) or a single pair of flip-related states as in the droplet picture. But it is difficult to see how our result can be reconciled with the presence of *many* thermodynamic states; indeed, at first glance it would appear to rule them out.

Nevertheless, we argue below that our theorem does *not* rule out the existence of many states, but clearly puts severe constraints on the form of the metastate (and overlap distribution function, which also possesses this invariance property). Our heuristic conclusion is that, in light of this strong invariance property, any metastate constructed via coupling-independent BC's can support only a very simple structure. As a consequence, we will argue that this theorem effectively rules out the nonstandard SK picture.

To see that an *uncountable* set of pure states is not ruled out (we will discuss countably infinite sets below), consider the highly disordered ground state model [41] in high dimensions, which is believed to exhibit a version of the chaotic pairs picture with uncountably many states. Our invariance theorem applies to this model also, and so (e.g.) the periodic and antiperiodic metastates must be the same, even though we might *a priori* expect them to be different. By what mechanism could this happen? The most natural possibility is that both the periodic and antiperiodic BC metastates are the same as the free BC metastate [44] in which all relative signs between the different trees in the invasion forest (see Refs. [41,42] for details) are equally likely. That is, each of these metastates consists of a *uniform distribution* on the ground state pairs. Given that, it does not seem unreasonable that all sorts of different BC's should give rise to a similar uniform distribution. Indeed, any fixed BC *does* give a uniform distribution on all *single* ground states [41,42].

But this line of reasoning does appear to rule out the chaotic pairs picture with a *countable* infinity of states. In that case, of course, one cannot have a uniform distribution (i.e., all equal, positive weights within the metastate). So now suppose that for some $\mathcal{J}$ the periodic BC metastate assigns, for example, probability 0.39 to one pair of pure states, 0.28 to another, and so on. In other words, with periodic BC's 39% of the finite cubes prefer pair number 1, 28% prefer pair number 2, etc. So pair number 1 is the overall ''winner'' (among different finite volumes) in the periodic BC popularity vote.

It now seems clear heuristically, though, that the popularity vote by *antiperiodic BC's* should come out differently; it is unreasonable to suppose that pair number 1 be preferred by 39% of the periodic BC cubes and at the same time by 39% of the antiperiodic BC cubes. The uniform distribution conclusion seems even more inevitable when one considers that analogous arguments also apply to pairs of arbitrarily chosen sequences of *fixed* boundary conditions.

We conclude that consistency between our invariance theorem and the existence of (uncountably) many states requires, in some sense, an equal likelihood of the appearance (in the metastate) of all states, i.e., some sort of uniform distribution on them. Let us examine this further. We have already noted that different sequences of volumes with fixed BC's — i.e., all volumes having plus boundary conditions, all volumes having plus on some boundary faces and minus on others, all volumes with each boundary spin chosen by the flip of a fair coin, and so on — result in the same metastate. We note for future reference that the term ''chaotic pairs,'' which was chosen in reference to spin-symmetric BC's (such as periodic) should be replaced here by ''chaotic pure states''; i.e., in this picture, the Gibbs state in a typical large volume Λ_L with fixed BC's will be (approximately) a single pure state that varies chaotically with L. But we expect that the mixed state $\rho_{\mathcal{J}}$, which is the *average* over the metastate [20–23],

$$\rho_{\mathcal{J}}(\sigma)=\int \Gamma(\sigma)\kappa_{\mathcal{J}}(\Gamma)d\Gamma, \tag{6}$$

would be the same for periodic and fixed BC's. One can also think of this $\rho_{\mathcal{J}}$ as the average thermodynamic state, $N^{-1}(\rho_{\mathcal{J}}^{(L_1)}+\rho_{\mathcal{J}}^{(L_2)}+\cdots,\rho_{\mathcal{J}}^{(L_N)})$, in the limit $N\to\infty$.

Now consider the mixed boundary condition in which *every* fixed BC on the boundary of each Λ_L is given equal weight. If there are (uncountably) many pure states present, then in a typical large volume one would expect to see a Gibbs state which approximates a continuous mixture over the pure states (cf. Possibility 3 or 4 discussed in Ref. [21]). But we still expect that the average over the mixed BC metastate would be the same $\rho_{\mathcal{J}}$ as for the fixed BC metastate, the periodic BC metastate, and so on. That is, the average over the metastate should be even more robust than the metastate itself, i.e., it should be the same for metastates constructed through *any* two sequences of (coupling-independent) BC's, not just flip-related ones.

Although logically possible, it seems unreasonable that this last (mixed BC with all fixed BC's given equal weight) metastate, chosen from a maximally uniform mixture of boundary conditions, can have anything other than a uniform distribution over the pure states. But, as just pointed out, this distribution should be the same for this as for all the other metastates under discussion. (We caution the reader that, unlike the case of the strongly disordered model [45], we do not have a precise sense in which this distribution can be defined to be uniform. For that reason, this part of the argument must be regarded as heuristic.)

With these points in mind, we now turn to a discussion of the nonstandard SK picture, and other possible mixed state scenarios.

The nonstandard SK picture requires [cf. Eq. (1)] that the Γ's appearing in the metastate be of the form $\Sigma_\alpha W_{\mathcal{J}}^\alpha \rho_{\mathcal{J}}^\alpha$, with at least some subset of the weights $W_{\mathcal{J}}^\alpha$ in each Γ nonzero and unequal. We would then have a situation like the following. With periodic BC's, say, the fraction of L_j's for which the finite-volume Gibbs state in Λ_{L_j} puts (e.g.) at least 84% of its weight in one pair of pure states (but with that pair not specified) is 0.39. But then it must also be the case that with antiperiodic BC's the fraction of volumes for which the finite-volume Gibbs state puts at least 84% of its weight in some unspecified pair is still exactly 0.39. Moreover, the same argument must apply to any ''cut'' one might care to make; i.e., one constructs the periodic BC metastate and finds that x% of all finite volumes have put y% of their weight in z states, with z depending on the (arbitrary) choice of x and y. Then this must be true also for all volumes with antiperiodic BC's; and similarly (but possibly separately) among all pairs of fixed BC states.

Once again, the only sensible way in which this could happen would be for the selection of states to be relatively insensitive (in some global sense) to the choice of boundary conditions, i.e., for the BC's to choose the states in some ''democratic'' fashion without favoritism so that $\rho_{\mathcal{J}}$, the average over the metastate, should be some sort of uniform mixture of the pure states, as before. However, unlike in the chaotic pairs picture discussed earlier, we claim that this *cannot* happen when the Γ's are (nontrivial) mixed states.

The reason for this is that the metastate has a strong covariance property [38] (see also [22]) in which the Γ's must transform in a specified way under an arbitrary finite change in the coupling realization. Under this finite change, the ensemble $\kappa_{\mathcal{J}}(\Gamma)$ transforms (as would any probability measure) according to the change of variables $\Gamma\to\Gamma'$. 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 each Γ *will* in general change, as one would expect. To see this, consider a particular Γ having a discrete pure state decomposition

$$\Gamma=\sum_\alpha W_\Gamma^\alpha \rho_{\mathcal{J}}^\alpha(\sigma), \tag{7}$$

with many nonzero weights W_Γ^α. Suppose that one chooses a particular coupling J_{xy} and imposes the transformation $J_{xy}\to J'_{xy}=J_{xy}+\Delta J$. Then the weight W^α (within Γ) of the pure state α will transform for each α as

$$W^\alpha\to W'^\alpha=r_\alpha W^\alpha \Big/ \sum_\gamma r_\gamma W^\gamma, \tag{8}$$

where

$$r_\alpha=\langle\exp(\beta\Delta J\sigma_x\sigma_y)\rangle_\alpha=\cosh(\beta\Delta J)+\langle\sigma_x\sigma_y\rangle_\alpha\sinh(\beta\Delta J). \tag{9}$$

In either the droplet-scaling or the chaotic pairs picture, there are in each Γ only two pure states (depending on Γ in chaotic pairs), each with weight 1/2. Because all even correlations are the same in each pair of (flip-related) pure states, the transformation of Eq. (8) leaves the weights unchanged.

However, in nonstandard SK there exist pure states within each (mixed) Γ with relative domain walls, so that they differ in at least some even correlation functions. But this then rules out that $\rho_{\mathcal{J}}$ must always be a uniform mixture of the pure states, because a suitable change of couplings will shift the weights for each Γ in such a way that the distribution over pure states of $\rho_{\mathcal{J}}$ also shifts. (This reasoning can be made rigorous, but because other parts of the argument are heuristic, we omit a proof.)

In other words, we argued above that the invariance of the metastate with respect to boundary conditions left open, as the only reasonable possibility for the presence of many pure states, that $\rho_{\mathcal{J}}$, the average over the metastate, be some sort of uniform mixture over the pure states. This must be true for any $\mathcal{J}$ (with probability one), so the weight distribution over all pure states must also be invariant with respect to changes in $\mathcal{J}$. But this invariance is inconsistent with the transformation properties of the Γ's with respect to finite changes in $\mathcal{J}$: if there are multiple pure states in the Γ's, with the pure states in each Γ not having the same even correlations (i.e., they have relative domain walls), then their relative weights must vary (as expected) with changes in the coupling realization. This leads to a contradiction, and therefore rules out not only nonstandard SK but any picture in which the Γ's are a nontrivial mixture of pure states.

Our conclusion, based on the above combination of both rigorous results and heuristic arguments, is that the nonstandard SK picture cannot be valid in any dimension and at any temperature. More generally, the many invariances of the spin glass metastate cannot support *any* picture in which thermodynamic mixed states (other than a single flip-related pair) are seen in finite volumes.

Given that the only reasonable possibilities remaining (that display broken spin-flip symmetry) are the droplet-scaling picture and the chaotic pairs picture, we conclude that the overlap distribution function P_Γ,

$$P_\Gamma(q)=\sum_{\alpha,\gamma} W_\Gamma^\alpha W_\Gamma^\gamma \delta(q-q^{\alpha\gamma}) \tag{10}$$

can at most be a pair of δ functions at $\pm q_{\rm EA}$ for each Γ; i.e., in each finite volume the overlap between pure states that appear in that volume is just that pair of δ functions. This will be the case regardless of whether there is only a single pair or uncountably many pairs of pure states. We will discuss this further in Sec. VI, but first we turn to another topic.

In the next section we present a simple scaling approach that provides both a plausibility argument and also a physical starting point for understanding the ''chaotic pairs'' many-state picture introduced in Refs. [21–23]. It is important to note that this scaling picture is consistent with the Fisher-Huse droplet picture [27,30] for appropriate values of the new scaling exponents, but for other values can give rise to a different thermodynamic picture.

V. A SCALING APPROACH TO THE CHAOTIC PAIRS PICTURE

We have argued above that with periodic boundary conditions, one should see at most a single pair of flip-related pure states in a large volume. As already discussed, this leaves open the possibilities of either a single pure state (e.g., but not necessarily, a paramagnet), a single pair of pure states (as in the droplet picture), or the chaotic pairs many-state picture discussed above. We now present a simple extension of earlier scaling-droplet arguments [27,30] which is consistent with this last possibility, and also provides a possible scenario for the spatial structure of domain wall configurations among the ground states.

The object here is to obtain estimates on the difference in energy or free energy between the lowest-lying state in a fixed volume and the next higher one. The appearance at nonzero temperature of multiple (non-spin-flip related) states in a single (large) volume requires that the energies of the lowest-lying states differ by order one. If, on the other hand, the ''minimal'' energy difference scales as some positive power of the system size, then one will see at most a single pair of states in any given box (with spin-symmetric boundary conditions, such as periodic).

To analyze the appearance in finite volumes, and at very low temperature, of infinite-volume pure states, as in Eq. (1), we will consider infinite-volume *ground* states restricted to the cube of size L, with a fixed boundary condition $\hat{\sigma}$ chosen independently of the couplings. In our analysis below we will treat the boundary spins as chosen randomly and independently of the couplings — but for a nonrandom fixed BC such as plus, the same arguments go through with minor modifications.

Although there may *a priori* be infinitely many infinite-volume ground states, the number of distinct restrictions to the cube is finite and its logarithm should be of order $L^{d-1-\phi}$ for some ϕ. The scaling exponent ϕ (with $0\leq\phi\leq d-1$) may be understood in another way: the minimum number of spins on the surface of the cube that differ between two infinite-volume ground states, whose spins disagree at (or near) the origin [46], should scale as L^ϕ. These two states should correspondingly differ in the bulk by a number of spins of (at least) order $L^{\phi+1}$.

If there exists only a single pair of flip-related ground states (as argued in Refs. [28,29]), then $\phi=d-1$. In the highly disordered spin glass model of Refs. [41,42] (see also Ref. [43]), it appears that $\phi=d-1$ below eight dimensions while $\phi=3$ above eight dimensions.

Let us examine the exponent ϕ more closely. Although *a priori* there seems to be no reason to exclude the possibility that $\phi=0$, there are several arguments indicating otherwise. (Note also that $\phi=0$ would saturate the possible growth rate of the number of distinguishable ground states in any finite volume since the logarithm of this number cannot exceed order L^{d-1}.) If $\phi=0$, then spins in regions between domain walls would exist in one-dimensional tubelike objects. It seems very unlikely that such tubes could be stable; i.e., eventually such a tube should encounter a fluctuation which destroys its structure. A second and somewhat different argument uses the fact that ϕ should be bounded from below by the exponent θ introduced by Fisher and Huse [27,30], which governs the minimal interface free energy between different pure states on a length scale L; i.e., this minimal free energy is presumed to grow as L^θ. It is not difficult to see, then, that $\phi\geq\theta$. However, it was argued in Refs. [27,30] that the inequality $\theta>0$ is necessary in order for spin-flip symmetry to be broken at positive temperature. In what follows we therefore always assume that $\phi>0$.

Before considering the EA model itself, we first treat the much simpler case of a homogeneous Ising ferromagnet with fixed BC's chosen at random. First we consider the energy difference between the plus and minus ground states (with interface ground states temporarily not considered). Here there is no bulk energy difference, and $\phi=d-1$. Because of the randomness of the BC, the boundary energy difference is of order $L^{\phi/2}$. The conclusion in this case [47] (see also [23]) is that the total energy difference is also $L^{\phi/2}$ and thus with random BC's one does not see a mixture of the plus and minus states but only one of them (chosen by the sign of the boundary energy) chaotically changing with L.

What about seeing interface states? Here, the appropriate bulk energy difference between the constant ground states and the interface states scales as L^{d-1} (with ϕ the same as before) and so the bulk energy difference dominates the boundary energy difference. In this case the total energy difference between the homogeneous state and the lowest-lying interface state is of order L^{d-1}. As a result, all interface states are ''invisible'' in the random BC finite-volume ferromagnet [23,47].

We now consider the EA Ising spin glass from this point of view. That is, we consider the energies of the restrictions of all infinite-volume ground states to the L^d cube centered at the origin. As before, we divide the energy into a bulk and a boundary part, and ask how the energy difference between the lowest-energy and next-lowest-energy state scales with L. Consider the state ρ^* with minimum total energy (subject to the fixed boundary condition) and the state of next lowest energy that differs from ρ^* near the origin. By the definition of ϕ, the two states differ by at least $L^{\phi+1}$ spins in the bulk and by L^{ϕ} spins on the boundary.

To estimate the energy differences between low-lying states, we will separately consider the boundary energy coming from the couplings between $\hat{\sigma}$ and the adjacent spins in the cube, and the bulk energy difference (from the remainder of the finite-volume Hamiltonian). If there were no bulk energies to consider, then one might expect that two states which differ by L^{ϕ} spins on the boundary would typically differ by an overall energy of order $L^{\phi/2}$. If this were indeed the case for the two lowest-lying states in almost any volume, then one would see only one state per volume (for fixed boundary conditions). However, since one is doing a minimization problem which includes bulk energies as well, it is not at all obvious *a priori* that this will happen. In particular, there might be some delicate cancellation between bulk and boundary energies.

We will now, however, present a specific scenario in which an explicit calculation shows that the lowest-lying states, in a volume with fixed boundary conditions chosen independently of the couplings, do indeed have an energy difference of order $L^{\phi/2}$. This example is presented as a plausibility argument and demonstrates one way in which this can occur, but is not meant to imply that it can occur in *only* this way.

Consider then a scenario in which the spin at the origin belongs to a cluster, not intersected by *any* domain walls, whose intersection with the boundary as before is of size L^{ϕ}. We denote that cluster $\mathcal{S}_0$. Suppose further that ρ^{α} is a general infinite-volume ground state, and that $E_L(\alpha)$ is the energy — including both the boundary and bulk components — of ρ_{α} restricted to Λ_L, the L^d cube centered at the origin.

The energy $E_L(\alpha)$ can therefore be written

$$E_L(\alpha) = -\sum_{x\in\mathcal{S}_0\cap\partial\Lambda_L}\sigma_x^{\alpha}\overline{\sigma_x} - \sum_{x\in\partial\Lambda_L\setminus\mathcal{S}_0}\sigma_x^{\alpha}\overline{\sigma_x} + E_L^b(\alpha), \tag{11}$$

where the first term is the contribution from the spins in the cluster $\mathcal{S}_0$ on the boundary $\partial\Lambda_L$, the second term is the surface energy contribution from all other boundary spins, and the final term is the energy contribution of the bulk spins. More precisely, $\partial\Lambda_L$ is the set of sites x inside Λ_L with a nearest neighbor y outside $\partial\Lambda_L$ and $\overline{\sigma_x}$ is the boundary spin $\hat{\sigma}_x$ times J_{xy}. Equation (11) can be rewritten as

$$E_L(\alpha) = -\eta(\alpha) Z_L \sqrt{|\mathcal{S}_0\cap\partial\Lambda_L|} + Y_L(\alpha), \tag{12}$$

where three new variables have been introduced: $\eta(\alpha) = \pm 1$ represents the sign of the spin at the origin in ground state α, Z_L is (approximately) a Gaussian random variable with zero mean and unit variance, and $Y_L(\alpha)$ depends both on the bulk energy of α and on the rest of the boundary spins (i.e., those not included in the first term).

In going from Eq. (11) to Eq. (12) we used the fact that the boundary condition consists of fixed random spins, chosen independently of α. The crucial observation is that the random variables Z_L, which arise from the random boundary conditions, are independent of the spectrum of the (mostly) bulk energies $Y_L(\alpha)$. We now show that, regardless of the number and distribution of the $Y_L(\alpha)$'s as α varies, there will be no strong cancellations between the two terms (with probability close to one).

Consider the ground state whose energy in Eq. (11) is the minimum, and also the ground state which has the next higher energy, and is *required* to have a relative spin flip with respect to the lowest-energy state at the origin. We then have

$$\left|\min_{\gamma:\sigma_0^{\gamma}=-1} E_L(\gamma) - \min_{\alpha:\sigma_0^{\alpha}=+1} E_L(\alpha)\right| = \left|2Z_L\sqrt{|\mathcal{S}_0\cap\partial\Lambda_L|} + Y_L^- - Y_L^+\right|, \tag{13}$$

where Y_L^- and Y_L^+ are the bulk plus remainder boundary energies of the two lowest-lying states with a relative spin flip at the origin.

Since Z_L and $Y_L^- - Y_L^+$ are functions of *disjoint* sets of the random boundary spins, they are independent random variables. Hence, variances add and the effect of $Y_L^- - Y_L^+$ on the random variable $2Z_L\sqrt{|\mathcal{S}_0\cap\partial\Lambda_L|}$ can only be to *increase* the spread of its distribution. This allows us to conclude that with probability close to one (i.e., for most choices of the boundary spins) the expression on the right-hand side of Eq. (13) is of order (at least) $\sqrt{|\mathcal{S}_0\cap\partial\Lambda_L|}$, i.e., of order $L^{\phi/2}$. As long as $\phi>0$, which is part of our scenario, this growth with L in the spacing of the low-lying spectrum of ground states argues for the appearance at small positive temperature of only a single pure state in large finite-volume Gibbs states with fixed BC's (that are independent of the couplings).

The above argument is instructive in several respects. It demonstrates that, given the condition that no domain wall separates the origin from the boundary of the box, there can be no miraculous ''conspiracy'' under which bulk and boundary energies cancel out to order one. It does require a strong condition, namely, that all domain walls, in the union of all symmetric differences over all ground states, do not form any closed and bounded regions. As stated above, this is a *sufficient* condition for the scaling argument given above to work, but we see no reason at this point why it should be a *necessary* condition in order for the conclusions to be valid.

Nevertheless, it provides one interesting scenario for the spatial structure of ground states and domain walls if many states should exist. Interestingly, in the only example of which we are aware in which a finite-dimensional spin glass apparently *does* possess many states in high dimensions — the highly disordered ground state model of Refs. [41,42] — exactly this structure occurs. These considerations provide a possibly fruitful avenue for future investigations.

VI. PURE STATES IN FINITE VOLUMES: WHAT IS GOING ON HERE?

In this section we address what it actually means, in an operational sense, to ''see'' a pure state — which formally is an infinite-volume object — inside a finite volume. We then use that analysis to answer a glaring question: if states and overlaps in finite volumes are restricted to, at most, a single pair of flip-related pure states and a pair of δ functions at $\pm q_{\mathrm{EA}}$, respectively, then what are the many numerical simulations (e.g., [12,16,18,19,34,35]) and experiments [48,49] that appear to see a more complicated state and overlap structure actually seeing?

Our main point will be that pure state structure can and does manifest itself in finite volumes, and governs the physics at finite length scales. Conversely, observations made in large, finite volumes must in turn reveal the thermodynamic structure and the nature of ordering of the system — if sufficient care is given to the analysis of those observations. Indeed, were both the above statements not true, it would be difficult to see why the study of thermodynamics would be of any interest to physics.

While the above assertions have long been noncontroversial for most statistical mechanical systems and models, there remains considerable confusion in the case of spin glasses [50]. At least part of the problem is that reliance on the overlap structure alone can at best give only partial — and sometimes misleading — information on the thermodynamics of realistic spin glass models [21,22,28,29]. A second problem is that, as we have emphasized in previous papers [21,22], the connection between finite- and infinite-volume behavior may be more complex and subtle in spin glasses than in simpler systems. An analysis of this connection thus deserves more thought than a simple attempt to sever the link altogether between the two behaviors (as in Appendix I of [19]). So in this section we will expand on previous discussions [22] to further clarify these issues.

A thermal state, whether pure or mixed, is completely specified by the set of all of its (one-point, two-point, three-point, . . .) correlation functions. In a finite volume, a state will manifest itself through the appearance of a particular set of such correlations. Because boundary effects will invariably alter or distort (compared to an infinite-volume state) these correlations in some region (whose size will depend on the specifics of the Hamiltonian, temperature, dimension, etc.), one must always be careful to examine them in a volume small enough so that these ''distortion'' effects are negligible. In other words, the boundary should be sufficiently far from the region under examination so that an accurate picture of the thermodynamics can be obtained [51].

So, for example, even in the paramagnetic state, one would measure nonzero magnetizations at interior sites in the vicinity of a boundary on which all spins are fixed (e.g., to be $+1$). As the boundary moves farther out, subsequent measurements at those same sites would find their magnetization tending to zero.

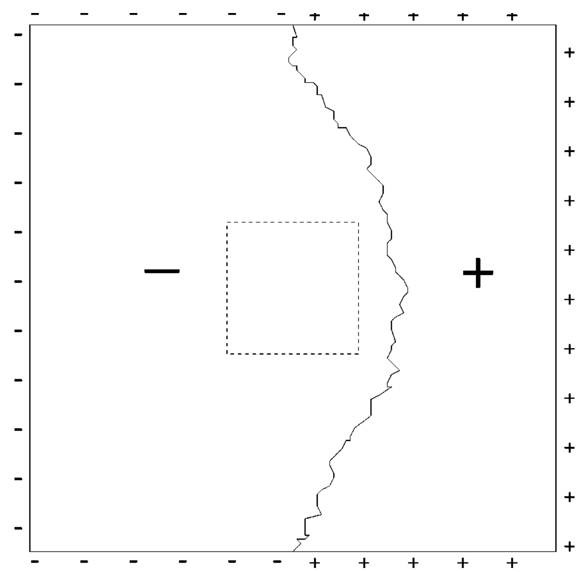

FIG. 1. A typical spin configuration in a two-dimensional Ising ferromagnet at positive temperature below T_c, with fixed spin boundary conditions that are $+1$ on the right half of the boundary and -1 on the left half. The maximum (and typical) deviation of the induced domain wall from the vertical line through the origin is $O(L^{1/2})$. This domain wall persists on all length scales but is unrelated to the low-temperature ordering. It will miss a sufficiently small $[o(L^{1/2})]$ window about the origin; examination of the order parameter inside only this window will correctly capture the thermodynamics. (In particular, one can examine any fixed finite region as the boundaries move far away.) This sketch depicts a relatively small square; for large L, the domain wall would be virtually indistinguishable from a straight line through the origin (on the scale L of the entire square), and the window would be extremely small (on that scale).

It is not unusual, even for comparatively simple systems, for boundary effects to penetrate more deeply into the interior than a shallow ''boundary layer.'' Consider the example of the two-dimensional uniform Ising ferromagnet. It is known [52,53] that this system has only two pure states — the translationally invariant positive and negative magnetization states — for all $0<T<T_c$. Suppose now that on a square of side L one were to impose fixed boundary conditions such that all spins on the right half of the boundary are $+1$ and all spins on the left are -1. This will impose a domain wall on the system, whose maximum (and typical) deviation (from the vertical line passing through the origin) will scale as $L^{1/2}$ (see Fig. 1). So for all large L the system gives the *appearance* of having a pure state with a domain wall [54]; indeed, the domain wall always stays quite far from the (vertical) boundaries. However, if one were to look at *any* fixed, finite region, then as the size L of the square grows, the domain wall eventually moves outside the fixed region, and one would see only a mixture of the positive and negative translationally invariant states. The (equal, in this case, as $L\rightarrow\infty$) weights in the mixture correspond to the probabilities of the domain wall thermally fluctuating to the left or to the right of the fixed region.

So in this example the domain wall is an artifact of the imposed boundary condition, and has nothing to do with any thermodynamic structure or low-temperature ordering properties of the system. Moreover, consideration of the spin configurations over the entire square would lead to incorrect conclusions about the pure state structure. This illustrates our contention that *in order to arrive at an accurate picture of the thermodynamic structure and the nature of ordering of a system, one must focus attention on a fixed ''window'' near the origin (which may be arbitrarily large, but is small compared to the entire volume under consideration).*

This conclusion is especially important when evaluating, and drawing inferences from, overlap functions. A more detailed discussion is given in the Appendix of Ref. [22], to which we refer the reader; here we will only reiterate an illuminating example due to van Enter [55], which in turn extends an earlier example due to Huse and Fisher [28]. Consider the overlap distribution of an Ising antiferromagnet in two dimensions with periodic boundary conditions. For odd-sized squares the overlap is equivalent (by the obvious gauge transformation) 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* square, it would therefore oscillate between two different answers [one a sum of two δ functions at plus or minus the square of M^*, the spontaneous magnetization, and the other a continuous distribution between $\pm(M^*)^2$]. On the other hand, 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 — which provides a more accurate picture of the nature of ordering in this system.

With these remarks in mind, we now turn to the finite-dimensional Ising EA spin glass. Essentially all the simulations of which we are aware compute the overlap distribution in the *entire* box. Boundary conditions are chosen independently of the couplings, and are usually periodic. Given our conclusion that, under these circumstances, at most a pair of flip-related pure states will appear in almost any finite volume, we suspect that the overlaps computed over the entire box are observing domain wall effects arising solely from the imposed boundary conditions, rather than revealing the actual spin glass ordering. (This is the reason why in Sec. V we looked only at states with relative domain walls in the vicinity of the origin.)

In other words, if overlap computations were measured in ''small'' windows far from any boundary, one should find only a pair of δ functions. One way to test this would be to fix a region at the origin, and do successive overlap computations in that fixed region for increasingly larger boxes with imposed periodic boundary conditions; as the boundaries move farther away, the overlap distribution within the fixed region should tend toward a pair of δ functions [56].

It is important to clear up one other misconception. It was asserted at the end of Sec. 2 in Ref. [19] that ''after Ref. [34] one has to argue that the physics must change after some very large length scale ... in order to claim that the mean field limit is not a good starting point to study the realistic case of finite dimensional models... .'' Although, of course, this changeover may well occur, it is at least as likely that it does not [57], and that nontrivial overlaps will be seen for all large L (as the uniform ferromagnet domain wall example illustrates). The real problem is in some sense the opposite: namely, that overlap computations are not being done in *small* enough regions to provide an accurate picture of spin glass ordering.

VII. CONCLUSIONS

In our previous papers [20–22], we showed that spin glasses may be more complex — in the relation between their behavior in finite and infinite volumes — than had previously been noted in the literature. In the present paper, we have presented arguments indicating that, in a different sense, spin glasses are more *simple* than had previously been claimed in much of the literature.

Our main conclusion is that, for realistic spin glass models such as Ising Edwards-Anderson, any large finite volume (with, say, spin-symmetric BC's, such as periodic, chosen independently of the couplings) will display at most a single pair of flip-related pure states. This may correspond to either a single pair of pure states in total, as in the droplet-scaling picture [27,29,30], or to the ''chaotic pairs'' picture introduced in Ref. [21] and elaborated upon in Refs. [22,23].

This rules out the nonstandard SK picture also introduced in Ref. [21] and elaborated upon in Refs. [22,23]. Combined with our earlier result [20] ruling out the standard SK picture, we conclude that *the thermodynamic structure and the nature of spin glass ordering, whether in finite or infinite volumes, cannot be mean-field-like in any dimension and at any temperature*.

The argument leading to this conclusion followed a theorem, presented in Sec. IV, that the metastate for fixed $\mathcal{J}$ is invariant with respect to arbitrary choices of flip-related boundary conditions (such as periodic and antiperiodic). It was then argued that only the simplest pure state (and corresponding overlap) structures could be so robust [58]. The only reasonable scenario under which (uncountably) many states could then appear is that, statistically, the states are insensitive to the boundary conditions. That is, the metastates would be generated (as in the highly disordered ground state model) through some kind of random fair-coin-tossing process.

We argued in Sec. VI that overlap computations should be done in small interior boxes (surrounded by much larger boxes where the boundary conditions are actually imposed) in order to remove boundary effects and get a picture of spin glass ordering that is not misleading. We expect that (with periodic BC's) for those dimensions and temperatures where $q_{\rm EA}\neq 0$, this procedure would result in a single pair of δ functions at $\pm q_{\rm EA}$ [59].

We also presented in Sec. V a scaling argument that shows how a ''chaotic pairs'' (or chaotic pure states, under fixed BC's) picture can arise. We provided an explicit calculation that supported this picture under the sufficient (but not necessary) condition that the union of domain walls between *all* pairs of pure states form no closed and bounded regions. Interestingly, exactly such a structure is present in the only example of a nontrivial short-ranged spin glass model known to have many ground states — i.e., the highly disordered spin glass model of Refs. [41,42] (see also [43]).

Given that an overlap structure computed in an entire finite volume (as opposed to that computed within a smaller window) might be nontrivial due only to boundary effects, it cannot yield definitive information on the ordering of the spin glass phase. Furthermore, there is no *a priori* reason to expect that it would display any exotic or intricate properties such as ultrametricity, or in general bear any particular resemblance to the mean-field structure observed for the SK model. However, the domain walls responsible for this overlap structure (if present) could have an observable, although perhaps nonuniversal, effect on dynamics. We plan to explore this issue in the future.

ACKNOWLEDGMENTS

This research was partially supported by N.S.F. Grant No. DMS-95-00868 (C.M.N.) and by U.S.D.O.E. Grant No. DE-FG03-93ER25155 (D.L.S.).

[1] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).

[2] K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).

[3] G. Parisi, Phys. Rev. Lett. **43**, 1754 (1979).

[4] G. Parisi, Phys. Rev. Lett. **50**, 1946 (1983).

[5] A. Houghton, S. Jain, and A. P. Young, J. Phys. C **16**, L375 (1983).

[6] M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, Phys. Rev. Lett. **52**, 1156 (1984).

[7] M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, J. Phys. (France) **45**, 843 (1984).

[8] R. Rammal, G. Toulouse, and M. A. Virasoro, Rev. Mod. Phys. **58**, 765 (1986).

[9] M. Mézard and G. Parisi, J. Phys. I **1**, 809 (1991).

[10] J.-P. Bouchaud, M. Mézard, and J. S. Yedidia, Phys. Rev. Lett. **67**, 3840 (1991).

[11] J. S. Yedidia, in *1992 Lectures in Complex Systems*, edited by D. L. Stein (Addison-Wesley, Reading, MA, 1993), p. 299.

[12] G. Parisi, Physica A **194**, 28 (1993).

[13] D. Badoni, J. C. Ciria, G. Parisi, F. Ritort, J. Pech, and J. J. Ruiz-Lorenzo, Europhys. Lett. **21**, 495 (1993).

[14] S. Franz, G. Parisi, and M. A. Virasoro, J. Phys. I **4**, 1657 (1994).

[15] F. Ritort, Phys. Rev. B **50**, 6844 (1994).

[16] E. Marinari, G. Parisi, and F. Ritort, J. Phys. A **27**, 2687 (1994).

[17] P. Le Doussal and T. Giamarchi, Phys. Rev. Lett. **74**, 606 (1995).

[18] E. Marinari, G. Parisi, J. J. Ruiz-Lorenzo, and F. Ritort, Phys. Rev. Lett. **76**, 843 (1996).

[19] E. Marinari, G. Parisi, and J. J. Ruiz-Lorenzo, in *Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, Singapore, in press).

[20] C. M. Newman and D. L. Stein, Phys. Rev. Lett. **76**, 515 (1996).

[21] C. M. Newman and D. L. Stein, Phys. Rev. Lett. **76**, 4821 (1996).

[22] C. M. Newman and D. L. Stein, Phys. Rev. E **55**, 5194 (1997).

[23] 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).

[24] C. M. Newman, *Topics in Disordered Systems* (Birkhäuser, Basel, 1997).

[25] C. M. Newman and D. L. Stein, in *Proceedings of the 1997 International Congress of Mathematical Physics*, edited by D. De Wit, M. D. Gould, P. A. Pearce, and A. J. Bracken (World Scientific, Singapore, in press).

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

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

[28] D. A. Huse and D. S. Fisher, J. Phys. A **20**, L997 (1987).

[29] D. S. Fisher and D. A. Huse, J. Phys. A **20**, L1005 (1987).

[30] D. S. Fisher and D. A. Huse, Phys. Rev. B **38**, 386 (1988).

[31] W. L. McMillan, J. Phys. C **17**, 3179 (1984).

[32] A. J. Bray and M. A. Moore, Phys. Rev. Lett. **58**, 57 (1987).

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

[34] J. D. Reger, R. N. Bhatt, and A. P. Young, Phys. Rev. Lett. **64**, 1859 (1990).

[35] S. Caracciolo, G. Parisi, S. Patarnello, and N. Sourlas, J. Phys. (France) **51**, 1877 (1990).

[36] C. M. Newman and D. L. Stein, Phys. Rev. B **46**, 973 (1992).

[37] We reemphasize a point made earlier in Sec. II, which in the current context implies that only boundary conditions chosen independently of the couplings should be used to construct the metastate. We note that there might exist pure states for a given $\mathcal{J}$ that would not appear in a metastate so constructed; i.e., they would not appear in the pure state decomposition of any of the Γ's appearing in the metastate. A situation where such "invisible" pure states occur in the context of Ising ferromagnets is discussed in Sec. V. In the spin glass, such states would require special coupling-dependent boundary conditions in order to appear in the metastate. As indicated in our remarks at the beginning of Sec. II, these states, should they exist, could be of mathematical interest but would almost certainly play no physical role. [Arguments along these lines can also be found in Sec. 3 of A. C. D. van Enter and J. Frölich, Commun. Math. Phys. **98**, 425 (1985)]. It should be understood throughout that when "pure states" are referred to (e.g., in Sec. IV), we generally mean only those that appear in a metastate constructed using coupling-independent boundary conditions.

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

[39] It should be noted that currently existing proofs require not only a subsequence $L_1, L_2, \ldots, L_N, \ldots$ of cube sizes, but possibly also a subsequence of N's when taking the histogram limit. However, the crucial point is that this subsequence of cube sizes, even if necessary in some instances, remains independent of $\mathcal{J}$.

[40] As shown in Ref. [20], the rigorous exclusion of the non-self-averaging property for $P_{\mathcal{J}}(q)$ also implies a lack of ultrametricity of distances among all of the pure states. That is, although not also rigorously excluded, the ultrametricity property was shown to be highly implausible. It is worth noting that an ultrametric structure in state space can appear in the ground

state structure of models with *deterministic* (and hence trivially self-averaged) interactions, although in order to obtain this structure one has to make a very artificial choice of interaction. For details, see A. C. D. van Enter, A. Hof, and J. Miekisz, J. Phys. A **25**, L1133 (1992).

[41] C. M. Newman and D. L. Stein, Phys. Rev. Lett. **72**, 2286 (1994).

[42] C. M. Newman and D. L. Stein, J. Stat. Phys. **82**, 1113 (1996).

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

[44] Note that free boundary conditions are *not* flip related to periodic and antiperiodic boundary conditions.

[45] The uniform distribution in the case of the strongly disordered model corresponds to the sign of each tree in the invasion forest being chosen by the flip of a fair coin.

[46] This is clearly a subset of all pairs of the ground states that are distinct within the cube Λ_L; the reason for this restriction will be discussed in Sec. VI.

[47] A. C. D. van Enter, J. Stat. Phys. **60**, 275 (1990).

[48] M. Lederman, R. Orbach, J. M. Hamann, M. Ocio, and E. Vincent, Phys. Rev. B **44**, 7403 (1991).

[49] 0Y. G. Joh, R. Orbach, and J. M. Hamann, Phys. Rev. Lett. **77**, 4648 (1996).

[50] In Appendix 1 of Ref. [19], for example, it is argued that the pure state, or thermodynamic, structure is merely a mathematical infinite-volume construct that has little or no physical relevance to real (finite-volume) systems such as spin glasses. We believe those arguments to be misleading, and indeed, misdirected in that the metastate approach precisely does connect the behavior of observable quantities in finite volumes with the thermodynamic structure of the system. (Moreover, the suggestion in that same reference that the Boltzmann-Gibbs probability distribution does not even exist in the infinite-volume limit for many disordered systems is simply incorrect.) It is, for example, a misconception that the behavior of correlation functions is more physical or less ''metaphorical'' (cf. Appendix 1 of Ref. [19]) than thermodynamic states. Indeed, the two are simply different labels for the same object, in the same way that one can talk either of the probability distribution of a random variable or the set of its moments.

[51] This should not be confused with the fact that, if many pure states are present, then changes in boundary conditions can change the state everywhere in the volume, including the region about the origin. In this situation, boundary conditions can select the thermodynamic state in the interior; but in order to see which state has been selected, one must still measure correlations in a region about the origin sufficiently far from the boundaries.

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

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

[54] Such a non-translation-invariant pure state *will* occur in higher dimensions than two, below the roughening temperature.

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

[56] Although all direct numerical computations of $P_{\mathcal{J}}(q)$ [and $P(q)$] of which we are aware compute overlaps in the full volume, at least one computation has been reported [18,19] that does examine a type of overlap measure, called the Binder cumulant, constructed on restricted subvolumes. Although strictly speaking the measurement reported has a dynamical component, it may contain potentially interesting and currently unexplained information on the equilibrium spin glass. However, the limited nature of the measurements done to date seem to us insufficient grounds for ruling out the droplet-scaling picture, as asserted in [18,19].

[57] The possibility that finite size effects might be persistent in systems with quenched disorder was also noted in Ref. [19].

[58] We should point out the special properties, under these arguments, of *free* boundary conditions. Free BC's are not flip related to any others and our arguments in Secs. IV and V do not apply to them. We further note that in the SK model itself, free BC's are in some sense the only natural boundary condition available. So could it be the case that the nonstandard SK picture might appear under free BC's and no other? We do not find this to be a reasonable possibility because, unlike in the case of the infinite-ranged model, there is nothing particularly special about free BC's in finite-dimensional short-ranged models. Although for technical reasons our arguments apply to BC's such as periodic, antiperiodic, fixed, and so on, the crucial aspect of our arguments is more closely related to the property that these BC's are chosen *independently* of the couplings. In this respect free BC's for arbitrary volumes are no different from the others. In the highly disordered model, for example, we expect (but have not proved) that the periodic or antiperiodic BC metastate is identical to the free BC metastate (cf. Sec. IV).

[59] We discussed in the Appendix to Ref. [22] various subtleties associated with the precise method of construction of the overlap distribution. In this paper we have referred only to the case where the overlap is computed in finite volumes using the replica measure $\rho_{\mathcal{J}}^{n(L)}$ discussed in that paper. If replica nonindependence [21,22] were present, as would be the case if the chaotic pairs picture were to hold, then one could construct a different infinite-volume overlap distribution by breaking replica symmetry *after* the infinite-volume limit is taken (cf. construction 2 of Ref. [20]). This would be the replica overlap for the average $\rho_{\mathcal{J}}$ of the metastate, and it would be the same not only for almost all flip-related boundary conditions but also, at the same time, for almost every $\mathcal{J}$. Given that, the only reasonable possibilities for this overlap function within the chaotic pairs scenario would be either a single δ function at the origin, or (less likely, we believe) a continuous distribution between $\pm q_{\mathrm{EA}}$ with *no* δ-function spikes.