Journal of Statistical Physics, Vol. 106, Nos. 1/2, January 2002 (© 2002)

The State(s) of Replica Symmetry Breaking: Mean Field Theories vs. Short-Ranged Spin Glasses[1]

C. M. Newman[2] and D. L. Stein[3]

Received May 17, 2001; accepted August 29, 2001

We prove the impossibility of recent attempts to decouple the Replica Symmetry Breaking (RSB) picture for finite-dimensional spin glasses from the existence of many *thermodynamic* (i.e., infinite-volume) pure states while preserving another signature RSB feature—space filling relative domain walls between different finite-volume states. Thus revisions of the notion of pure states cannot shield the RSB picture from the internal contradictions that rule out its physical correctness in finite dimensions at low temperature in large finite volume.

KEY WORDS: Spin glass; Edwards–Anderson model; replica symmetry breaking; mean-field theory; pure states; ground states; metastates; domain walls; interfaces; incongruence.

1. INTRODUCTION

In this paper we will describe what the mean-field picture and its central component, replica symmetry breaking (RSB), *must* mean for short-ranged spin glasses in all finite dimensions—and why it *cannot* hold for these systems.

In a recent paper[(1)] (hereafter [MPRRZ]), Marinari, Parisi, Ricci-Tersenghi, Ruiz-Lorenzo, and Zuliani have provided the most extensive description of the mean-field RSB picture offered to date. In response to earlier demonstrations[(2–6)] (hereafter [NS96a], [NS96b], [NS97a], [NS97b], and [NS98], respectively) by the authors that the mean-field RSB picture

[1] An earlier version of this paper was posted as cond-mat/0105282 under the title, "Replica Symmetry Breaking's New Clothes."

[2] Courant Institute of Mathematical Sciences, New York University, New York, New York 10012; e-mail: newman@cims.nyu.edu

[3] Departments of Physics and Mathematics, University of Arizona, Tucson, Arizona 85721; e-mail: dls@physics.arizona.edu

0022-4715/02/0100-0213/0

cannot describe the structure of thermodynamic, i.e., infinite-volume, pure states at temperatures $T > 0$ or ground states at $T = 0$ of spin glasses in any finite dimension d, [MPRRZ] proposed (see also Appendix 1 of the paper by Marinari *et al.*[7]) that RSB is not meant to provide such a description, but instead applies only to the structure of "finite-volume pure states," which are the relevant physical objects. An unambiguous definition of finite-volume pure states was not provided, but it was emphasized that they were different from the "pure states in an infinite system," i.e., the usual thermodynamic pure states.

However, in Section 6, we present a new proof (whose applications go well beyond spin glasses alone), which, when applied in the current context, shows rigorously that the primary claims of [MPRRZ][1] are *incompatible* with each other. That is, if the claim of nontrivial link overlap ($P_e^L(q)$) for large L is valid, *it must give rise to multiple ground and pure states in the usual thermodynamic sense* (cf. Appendix A), as traditionally envisioned (see, for example, the review article[8] by Binder and Young, hereafter [BY], or the book[9] by Mezard, Parisi, and Virasoro, hereafter [MPV]). Therefore, whether or not a new interpretation of the mean-field RSB theory in terms of "finite-volume pure states" can ever be precisely formulated, the more usual infinite-volume interpretations cannot be avoided. But a mean-field structure for multiple *thermodynamic* states has been ruled out by the authors' previous arguments.[2–6] Mean-field RSB theory therefore cannot apply to realistic (i.e., finite-dimensional short-ranged) spin glasses.

The paper is organized as follows. In Section 2 we introduce many of the terms and concepts needed for later sections, and provide an abbreviated review of the issues concerning pure states within the framework of mean-field RSB theory. Section 3 discusses the behavior of interfaces, or domain walls, in finite volumes and how they can (or cannot) give rise to multiple thermodynamic ground or pure states. Section 4 discusses the notion of "finite-volume pure states" introduced in [MPRRZ],[1] and provides an initial critique of this concept, and Section 5 reviews the predictions of mean-field RSB theory for interface properties. In Section 6 we formally state our theorem that the mean-field RSB theory in fact *must* predict multiple thermodynamic pure state pairs with properties that have been previously ruled out; and in Section 7 we discuss the implications of our theorem and present our conclusion.

We also include two appendices. Appendix A is a brief summary of the definitions and properties of finite-volume Gibbs states, infinite-volume Gibbs states, and pure states. These play a major role in the text. Appendix B is a glossary providing brief definitions of other terms frequently used in the text; some are in common usage in the literature, but most are less so.

2. BACKGROUND AND REVIEW

In this section we provide a review of the concepts and terms that will be used throughout, and provide an abbreviated overview of recent developments in the equilibrium theory of finite-dimensional spin glasses, as they pertain to the current discussion. Detailed presentations can be found in the references cited. This section is included so that this paper is reasonably self-contained; readers familiar with these topics may want to skip ahead to Section 3.

For specificity, we focus on the Edwards–Anderson[10] (EA) Ising spin glass, whose Hamiltonian is given by

$$\mathscr{H} = -\sum_{\langle x,y\rangle} J_{xy}\sigma_x\sigma_y, \tag{1}$$

where the couplings J_{xy} are independently chosen from a Gaussian distribution with mean zero and variance one, the sum is over only nearest neighbors on the d-dimensional cubic lattice $\mathbf{Z}^d$, and the spins $\sigma_z = \pm 1$. We note for later, however, that our results apply to a wide range of models, including systems other than spin glasses.

2.1. Essential Features of the Mean-Field RSB Picture

Nontrivial replica symmetry breaking within the "mean-field picture" is associated with a number of remarkable properties, including the existence of many equilibrium states, non-self-averaging of overlap functions, ultrametricity of pure state overlaps, and several others less relevant to the current discussion. This picture is believed to describe the low-temperature phase of the infinite-ranged Sherrington–Kirkpatrick[11] model, where the sum in the Hamiltonian Eq. (1) now runs over *all* pairs of spins, and the variance of the coupling distribution is rescaled to provide a sensible thermodynamic limit. We assume that the reader is largely familiar with this picture, and refer her/him to [BY][8] or [MPV][9] for an extensive and detailed description. Throughout this paper we will refer to this picture and its variations as the mean-field picture, to adhere to common usage in the literature; but it should be kept in mind that it is based on the Parisi solution[12–15] of the Sherrington–Kirkpatrick infinite-ranged model.

Numerous authors have asserted that mean-field-like RSB should describe the broken symmetry of the low-temperature phase of more realistic short-ranged, finite-dimensional spin glass models as well. Its basic features, for a fixed $T < T_c$, have been described in many places (see, e.g., refs. 1, 7–9, 16–23)) and can be summarized as follows: (1) the existence of many equilibrium states not related by any simple symmetry transformation, and whose number grows without bound as system size $L \to \infty$; (2) for fixed coupling realization $\mathcal{J}$, a nontrivial probability distribution $P_{\mathcal{J}}(q)$, supported on countably many values $q_{\alpha\beta}$, for the spin overlap between two different replicas; (3) non-self-averaging, i.e., $\mathcal{J}$-dependence, of $P_{\mathcal{J}}(q)$, so that averaging $P_{\mathcal{J}}(q)$ over all $\mathcal{J}$ yields a $P(q)$ supported on a continuum of values between $\pm q_{EA}$, with nonzero weight at $q = 0$ and δ-function spikes at $\pm q_{EA}$; (4) ultrametricity of the spin overlaps $q_{\alpha\beta}$ among the equilibrium states; (5) nontrivial edge overlap $P_e(q_e)$; for example, if one chooses at fixed $\mathcal{J}$ a ground state from a cube with periodic boundary conditions, and a second ground state from the same cube with antiperiodic boundary conditions, then there would be a nonvanishing density (as $L \to \infty$) of bonds satisfied in one but not the other ground state.

To arrive at these features, the mean-field RSB picture postulates, as in [MPRRZ],[1] that at fixed T the finite-volume Gibbs state $\rho^L_{\mathcal{J}}$ in Λ_L, the cube of side-length L centered at the origin (we henceforth assume periodic boundary conditions for specificity, but in fact our arguments will apply to any boundary conditions chosen independently of $\mathcal{J}$), is approximately a mixture of many pure states (we defer until later the question of what this actually means):

$$\rho^{(L)}_{\mathcal{J}} \approx \sum_{\alpha} W^{\alpha}_{\mathcal{J},L} \rho^{\alpha}_{\mathcal{J}} \tag{2}$$

where $W^{\alpha}_{\mathcal{J},L}$ represents the Boltzmann weight in $\rho^{(L)}_{\mathcal{J}}$ of pure state $\rho^{\alpha}_{\mathcal{J}}$. The finite-volume overlap distribution $P^L_{\mathcal{J}}(q)$ is approximately the corresponding mixture of many δ-functions:

$$P^L_{\mathcal{J}}(q) \approx \sum_{\alpha,\gamma} W^{\alpha}_{\mathcal{J},L} W^{\gamma}_{\mathcal{J},L} \delta(q - q^{\alpha\gamma}_{\mathcal{J}}), \tag{3}$$

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

$$q^{\alpha\gamma}_{\mathcal{J}} \approx |\Lambda_L|^{-1} \sum_{x \in \Lambda_L} \langle \sigma_x \rangle^{\alpha} \langle \sigma_x \rangle^{\gamma}; \tag{4}$$

here $|\Lambda_L|$ is the number of sites in Λ_L.

An edge, or link, overlap distribution function can be similarly constructed.[1, 16–18] For simplicity, consider a ground state pair $\pm\sigma^L$ in Λ_L with periodic boundary conditions, and a second ground state pair $\pm\sigma'^L$ obtained in Λ_L, e.g., with antiperiodic boundary conditions. There will be a relative domain wall (or walls) between the pairs $\pm\sigma^L$ and $\pm\sigma'^L$, consisting of the set of bonds $\langle xy \rangle$ in the dual lattice satisfied in one and not the other ground state pair; that is, they obey

$$\sigma_x^L \sigma_y^L = -\sigma'^L_x \sigma'^L_y. \tag{5}$$

The link overlap between $\pm\sigma^L$ and $\pm\sigma'^L$ is

$$q_{\mathscr{J},e}^{\sigma^L \sigma'^L} = |E_L|^{-1} \sum_{\langle x,y \rangle \in E_L} \sigma_x^L \sigma_y^L \sigma'^L_x \sigma'^L_y \tag{6}$$

which is equal to one when $\sigma^L = \pm\sigma'^L$ and smaller than one otherwise. Here E_L is the edge set of, and $|E_L|$ is the number of edges in, Λ_L. The edge overlap distribution function is then given by

$$P^L_{\mathscr{J},e}(q_e) \approx \sum_{\alpha,\gamma} W^{\alpha}_{\mathscr{J},L} W^{\gamma}_{\mathscr{J},L} \delta(q - q^{\alpha\gamma}_{\mathscr{J},e}) \quad . \tag{7}$$

While this interesting picture at first seems reasonably clear, on closer inspection there are numerous problems in interpretation when applied to realistic models. Much of [MPRRZ][1] is devoted to arriving at a definition of an "equilibrium" or "pure" state within the mean-field RSB picture; but leaving that issue aside for now, there are numerous other questions that could affect interpretation of numerical measurements. For example, by what procedure are states, or replicas, chosen, and from what distribution? In computing $P(q)$, what does one mean by the "infinite-volume limit"? What is meant by non-self-averaging when its presence or absence may depend on the sequence of steps used to compute overlaps? Until these questions are clarified, we are forced to leave most of the above equations as approximate relations.

To illustrate, consider the situation at $T = 0$. For fixed $\mathscr{J}$ and Λ_L with periodic boundary conditions, there will be a single pair of ground states $\pm\sigma^L$. The overlap function will therefore be a pair of δ-functions at ± 1 for all L, and so the limiting $P(q)$ is that same pair of δ-functions, independently of $\mathscr{J}$. Does this imply a single pair of ground states, as predicted by the droplet/scaling picture of Macmillan,[24] Bray and Moore,[25, 26] Fisher and Huse,[27–30] and others? Not necessarily, because if there were many ground state pairs then $\pm\sigma^L$ would change chaotically with L, though for any single L one would see only a single pair. The presence or absence of this *chaotic size dependence*[31] (hereafter [NS92]) is a reliable test[5, 31] of whether there are, respectively, many ground state pairs or only a single pair. But if there are many ground state pairs, can one construct $P(q)$ in order to see them?

2.2. Standard Interpretation of the RSB Mean-Field Picture

The most straightforward, and natural, interpretation of the features of the RSB mean-field picture described above is that the "pure" states in Eq. (2) are the usual thermodynamic pure states, which are easily and unambiguously defined for the EA model (see, e.g., Appendix A of this paper and [NS97a, NS97b, NS98][4–6]). In this interpretation, for almost every fixed $\mathcal{J}$ there would be an infinite number of these states. The spin overlap distribution function would be nontrivial in the sense described in the preceding section, and would satisfy the properties of non-self-averaging and ultrametricity (including at $T=0$). The edge overlap distribution function would similarly be nontrivial and non-self-averaging. Procedures for constructing overlap distributions are provided in [NS96a].[2]

This interpretation has generally been the standard view (see, e.g., refs. 8, 9, 20, 21, 32, and 33), and is one way to answer the questions posed in the previous section. It allows us to replace the approximate relation Eq. (2) with an equality

$$\rho_{\mathcal{J}}(\sigma)=\sum_{\alpha} W^{\alpha}_{\mathcal{J}}\rho^{\alpha}_{\mathcal{J}}(\sigma), \tag{8}$$

where $\rho_{\mathcal{J}}(\sigma)$ is an infinite volume mixed Gibbs state for a particular coupling realization $\mathcal{J}$, the $\rho^{\alpha}_{\mathcal{J}}$ are infinite-volume pure states for that $\mathcal{J}$, and the $W^{\alpha}_{\mathcal{J}}$ their corresponding weights in $\rho_{\mathcal{J}}$.

The other equations in Section 2.1 are similarly replaced with exact relations. The overlap random variable becomes

$$Q=\lim_{L\to\infty}|\Lambda_L|^{-1}\sum_{x\in\Lambda_L}\sigma_x\sigma'_x \tag{9}$$

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

$$q^{\alpha\gamma}_{\mathcal{J}}=\lim_{L\to\infty}|\Lambda_L|^{-1}\sum_{x\in\Lambda_L}\langle\sigma_x\rangle^{\alpha}\langle\sigma_x\rangle^{\gamma}. \tag{10}$$

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

$$P_{\mathcal{J}}(q)=\sum_{\alpha,\gamma} W^{\alpha}_{\mathcal{J}}W^{\gamma}_{\mathcal{J}}\,\delta(q-q^{\alpha\gamma}_{\mathcal{J}}). \tag{11}$$

Edge overlap distribution functions are similarly defined.

However, it was rigorously shown in [NS96a][(2)] that *this standard interpretation of the mean-field picture cannot hold at any temperature in any finite dimension*, because the $P_{\mathscr{J}}(q)$ of Eq. (11) must be self-averaging, i.e., the same for almost every $\mathscr{J}$. This also rules out the possibility of nontrivial ultrametricity among the thermodynamic pure states. We note that the (nonrigorous) arguments of Parisi and Ricci-Tersenghi,[(33)] claiming to support ultrametricity among *all* the equilibrium states of finite-dimensional spin glasses, are in fact also consistent with *trivial* ultrametricity, such as that displayed in the droplet/scaling two-state picture or the many-state chaotic pairs picture (cf. Section 2.4 later). One must therefore adopt an unconventional interpretation of the mean-field RSB picture if there is to be any hope of its application to realistic spin glasses.

2.3. The Nonstandard Interpretation of the Mean-Field RSB Picture

The standard interpretation just described is a natural extrapolation to large lengthscales of numerical simulations necessarily done on cubes Λ_L with relatively small L. What is typically done numerically, of course, is to generate (usually with periodic boundary conditions, assumed here for specificity) finite-volume equilibrium Gibbs states (see Appendix A) in Λ_L and then to measure the overlap distribution for fixed $\mathscr{J}$; then repeat the procedure for different $\mathscr{J}$'s and compute the disorder-averaged $P_L(q)$. Doing this for several different L's allows one to examine finite-size scaling and other properties of overlap functions.

Numerically, one has no choice but to follow this or some similar procedure; but in [NS96b][(3)] it was shown that evidence for RSB arising from this approach can correspond to more than one thermodynamic picture. The above procedure, if extrapolated to arbitrarily large L, gives rise to a $P(q) = \lim_{L \to \infty} P_L(q)$ without any explicit or prior construction of thermodynamic states. (Another procedure that *does* first construct states and then computes overlaps is given in [NS96a][(2)]).

In these numerical computations, replica symmetry is of necessity broken *before* the $L \to \infty$ limit is taken. Guerra[(34)] has pointed out that changing the order of these limits can be quite significant. That this interchange of limits[(3)] can lead to a new thermodynamic picture of the spin glass phase does not seem to have been appreciated prior to [NS96a, NS96b].

Based on these considerations, a new, nonstandard interpretation of the mean-field RSB picture was described in detail in Section 7 of [NS97b];[(5)] we provide only a brief summary here. It is a maximal mean-field picture, preserving mean-field theory's main features, as discussed in Section 2.1, although in an unusual way. The most natural description of this non-standard interpretation is in terms of the *metastate*, described in the next section, but in order to simplify the discussion we forego use of the metastate here.

As a starting point, then, this interpretation would mean that in any large Λ_L, the Gibbs state is an approximate decomposition over many *thermodynamic* pure states:

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

where a few states dominate the sum in any fixed L. The overlap distribution in Λ_L for fixed $\mathcal{J}$, given by Eq. (3), is nontrivial: for any fixed, large L it would be a sum of several δ-functions, and the locations of these δ-functions would satisfy ultrametricity increasingly accurately as $L \to \infty$. When averaged over $\mathcal{J}$ at fixed L, this distribution would broaden into a continuum between two δ-functions at $\pm q_{EA}$.

Equation (12) and the properties listed after it all hold equally well for the standard interpretation of mean-field RSB described in Section 2.2. The difference between the standard and nonstandard interpretations arises from their thermodynamics; the straightforward extrapolation to infinite volumes characteristic of the standard interpretation is absent in the nonstandard picture. In the latter case, the infinitely many thermodynamic pure states are grouped into "families" of mixed states (the Γ's of Section 2.4), each of which *individually* has the properties listed in the preceding paragraph. The union of all of these families, which loses these properties, comprises the thermodynamic structure of the nonstandard interpretation.

The crucial conceptual point is that the resulting *ensemble* of overlap distributions remains independent of $\mathcal{J}$. So while overlap distributions still do *not* depend on $\mathcal{J}$, one now replaces the usual notion of non-self-averaging over $\mathcal{J}$'s with a nonstandard one: that is, averaging over L's for fixed $\mathcal{J}$. It can be shown that this picture must have uncountably many pure states and overlaps, so that *ultrametricity would not hold in general*[2] *among any three pure states chosen at fixed* $\mathcal{J}$, unlike in the standard interpretation (see, for example, the papers of Vincent *et al.*[35] and of Badoni *et al.*[36]). Instead, each large Λ_L would pick out a subset of these (one of the families discussed above) that *do* satisfy ultrametricity.

Unfortunately for the mean-field approach, it can also be shown that this picture cannot hold at any temperature in any dimension, as discussed in the next section.

2.4. Metastates, Chaotic Pairs, and the Simplicity of *P(q)*

To explain why even the nonstandard interpretation cannot be valid, we need to introduce the concept of *metastate*, discussed in detail in [NS96b, NS97a, NS97b, NS98].[3–6] (For some uses of this concept in mean field models, see the papers of Külske,[37, 38] Bovier and Gayrard,[39] and Bovier *et al.*[40]) 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, as typically occurs in homogeneous systems. Experience with these has mostly guided intuition in the case of disordered systems. However, one of our early results is that, in the presence of many pure states not related by any clear-cut symmetry transformations, the relation between the system's thermodynamic properties and its behavior in large but finite volumes may be non-obvious.

This is primarily due to the following result of [NS92]:[31] if a system has many, non-symmetry-related, pure states, the sequence of finite-volume Gibbs measures generated using coupling-independent boundary conditions will generally *not* converge to a single limiting thermodynamic state as $L \to \infty$. This is the phenomenon of *chaotic size dependence*, mentioned in Section 2.1. In the metastate approach, rather than avoid this problem, we exploit it by focusing on an *ensemble* of (pure or mixed) thermodynamic states. This approach, based on an analogy to chaotic dynamical systems, allows the construction of a limiting measure. Hence the term *metastate*—while a thermodynamic state is a probability measure on infinite-volume spin configurations (see Appendix A), this new limiting measure is one *on the thermodynamic states themselves.*

This infinite-volume measure has a particular usefulness in the context of finite volumes because it 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 n-point correlation functions contained in a box (or "window"), centered at the origin, whose sides are sufficiently far from any of the boundaries so that finite size or boundary effects do not appreciably affect the result. (We discuss this in more detail in Section 2.6.)

There are several ways of constructing metastates. In [NS96b, NS97a, NS97b][3–5] we introduced the empirical distribution approach. This considers, at fixed $\mathcal{J}$, a sequence of volumes with coupling-independent boundary conditions. Each finite-volume Gibbs state $\rho_{\mathcal{J}}^{(L_1)}$, $\rho_{\mathcal{J}}^{(L_2)}$,..., $\rho_{\mathcal{J}}^{(L_N)}$ in the sequence is given weight N^{-1}. This allows us to construct a histogram of finite-volume Gibbs states; it was shown in [NS96b, NS97a] that this histogram converges to a probability measure $\kappa_{\mathcal{J}}$ on the thermodynamic states as $N \to \infty$. A finite-volume Gibbs state in a particular (large) volume approximates (deep in its interior—cf. the remarks in the preceding paragraph) some infinite-volume thermodynamic state Γ restricted to that volume. The resulting metastate $\kappa_{\mathcal{J}}$ therefore specifies the fraction of cube sizes that the system spends in each different thermodynamic state Γ. An individual Γ may be either pure or mixed, depending on the system and the boundary conditions used.

The empirical distribution approach presented above was shown in [NS96b, NS97a][3,4] to be equivalent to an earlier construction of Aizenman and Wehr.[41] In this alternative approach, the randomness of the couplings is used directly to generate an ensemble of states. It can be proved that the two approaches are basically equivalent, in that there exists at least a $\mathcal{J}$-independent subsequence of volumes along which both methods yield the *same* limiting metastate.[4,5]

The metastate approach is specifically designed to consider both finite and infinite volumes together and to unify the two cases. In essence, the metastate provides the probability, for a randomly chosen large L, of various thermodynamic pure (or ground, at $T=0$) states appearing inside any fixed Λ_{L_0}.

We return now to the nonviability of the nonstandard interpretation of the mean-field RSB theory in realistic spin glass models. Our claim is based on a simple theorem, presented in [NS98],[6] with a powerful implication—that (at fixed d and T) the metastate $\kappa_{\mathcal{J}}$ is *invariant* with respect to flip-related boundary conditions, chosen independently of the couplings. That is, the metastate constructed using periodic boundary conditions on the Λ_L's is the same as that constructed using antiperiodic boundary conditions. Even if one were to choose two *arbitrary* sequences of periodic and antiperiodic boundary conditions, the metastates would *still* be identical. The metastate, and its corresponding overlap distributions, is therefore highly insensitive to boundary conditions.

This metastate invariance has profound consequences. It means that the frequency of appearance of various thermodynamic states in finite volumes is *independent* of the choice of periodic or antiperiodic boundary conditions. Moreover, this same invariance property holds among any two sequences of *fixed* boundary conditions; the fixed boundary condition 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.

If there were only a single thermodynamic state, such as paramagnetic, or a single pair of states as in droplet/scaling, this would be expected. But can this result can be reconciled with the presence of *many* thermodynamic states?

The answer is yes, but it puts severe constraints on the form of the metastate and overlap distribution functions. In light of this strong invariance property, any metastate constructed via coupling-independent boundary conditions should be able to support only a very simple structure, effectively ruling out the nonstandard interpretation of the mean-field RSB picture.

How can this invariance property be reconciled with the presence of many non-symmetry-related pure or ground states? The only plausible possibility is that in any metastate constructed from coupling-independent boundary conditions (periodic, antiperiodic, free, fixed, etc.), all pure thermodynamic states are equally likely. That is, each of these metastates should be supported *uniformly* in some appropriate sense (which can be made precise only with detailed knowledge about the pure states—see, e.g., the discussion in Section 4 of [NS98][(6)]) on the pure state pairs in that metastate. This is the only plausible way in which all sorts of different boundary conditions could give rise to the same pure state distribution.

Such a uniform distribution, though, is inconsistent with the features of the nonstandard mean-field picture. That picture requires a *non*uniform distribution over the pure states (for further discussion, see [NS98][(6)]), as does *any* picture in which a nonzero fraction of Γ's consists of a nontrivial mixture of pure state pairs. There is only one many-state picture of which we are aware that is consistent with this theorem. This is the *chaotic pairs* picture, introduced in [NS92][(31)] and [NS96b][(3)] and further developed in [NS97a, NS97b, NS98].[(4–6)]

The chaotic pairs picture resembles the scaling/droplet picture in finite volumes, but has a very different thermodynamic structure. It has *infinitely* many thermodynamic pure states, but, unlike any mean-field picture, in each large volume with periodic boundary conditions one "sees" only one pair of pure states at a time. That is, for large L, one finds that

$$\rho_{\mathscr{J}}^{(L)} \approx \tfrac{1}{2}\,\rho_{\mathscr{J}}^{\alpha_L} + \tfrac{1}{2}\,\rho_{\mathscr{J}}^{-\alpha_L} \tag{13}$$

where $-\alpha$ refers to the global spin-flip of pure state α. So each L picks out a single pure state pair from the infinitely many present. If all Λ_L have periodic boundary conditions, then the chaotic pairs picture would exhibit chaotic size dependence, unlike the droplet/scaling picture. In other words, in the scaling/droplet picture, the low-temperature, periodic boundary condition metastate is supported on one thermodynamic mixed state Γ consisting of a single pure state pair, and this Γ is seen in a fraction one of the Λ_L's. In chaotic pairs, the metastate is dispersed over infinitely many Γ's, of the form $\Gamma = \Gamma^{\alpha} = \frac{1}{2}\,\rho_{\mathscr{J}}^{\alpha} + \frac{1}{2}\,\rho_{\mathscr{J}}^{-\alpha}$; here, two different Λ_L's will typically see *different but single* pure state pairs.

The overlap distribution for each Γ, hence each Λ_L (for L large) is the same:

$$P_\Gamma = \tfrac{1}{2}\,\delta(q - q_{EA}) + \tfrac{1}{2}\,\delta(q + q_{EA}). \tag{14}$$

So the disorder-averaged spin overlap function $P(q)$ *and* link overlap function $P_e(q_e)$, when constructed by breaking the replica symmetry *before* taking the thermodynamic limit, as in Section 2.3, must have the same, simple structure whether there exists a single pair of thermodynamic pure states or infinitely many: the spin overlap function $P(q)$ would be a pair of δ-functions at $\pm q_{EA}$, and the link overlap function $P_e(q_e)$ a single δ-function at some $q_e(T)$ (see later), in *either* case.

There *is* a difference in overlap functions in the two pictures, however, if the thermodynamic limit is taken *before* replica symmetry is broken, as in Section 2.2. Here, as already noted, $P_{\mathcal{J}}(q)$ and $P_{\mathcal{J},e}(q_e)$ are, at fixed T and d, the same for almost every $\mathcal{J}$, regardless of which of the two pictures actually occurs. In droplet/scaling, $P_{\mathcal{J}}(q)$ is again a pair of δ-functions at $\pm q_{EA}$, whereas the link overlap function (computed in a box small compared to Λ_L and far from the boundaries—cf. Section 2.6) $P_{\mathcal{J},e}(q_e) = \delta(q_e - 1)$ at $T = 0$ and presumably remains a single δ-function at all temperatures, though the q_e value where the spike occurs decreases due to thermal fluctuations as T increases. In chaotic pairs, $P_{\mathcal{J}}(q)$ would now most likely equal $\delta(q)$: it was proven by the authors[42] that $P_{\mathcal{J}}(q) = \delta(q)$ for the spin overlaps of M-spin-flip-stable metastable states for any finite M, and, if there are infinitely many ground state pairs, we expect the same to be true for ground states, i.e., for $M = \infty$. The form of the edge overlap function in the chaotic pairs picture, when replica symmetry breaking occurs after taking $L \to \infty$, is less clear; the contribution coming from relative interfaces between the many pairs of pure states may well be a δ-function, but unlike any two-state picture, would be supported on a link overlap $q_e < 1$ even at $T = 0$. (A lengthier discussion of link overlap functions in given in Section 5.)

Our conclusions are therefore that *the thermodynamic overlap structure in spin glasses must be simple*, regardless of whether there are infinitely many pure states or only a single pair. The form of the overlap function, however, can depend on how the computation is done. Our results for the spin overlap function $P(q) = P_{\mathcal{J}}(q)$ are summarized in Fig. 1.

In Fig. 1, the overlap function $P(q)$ is shown for two very different physical pictures—one a single pure state pair picture, as in droplet/scaling, and the other the chaotic pairs picture, which presupposes an uncountable infinity of pure states. When comparing the overlap function for different scenarios in general, it is important that computations be done in the same way. Figures 1a and 1b represent overlap computations done on cubes Λ_L

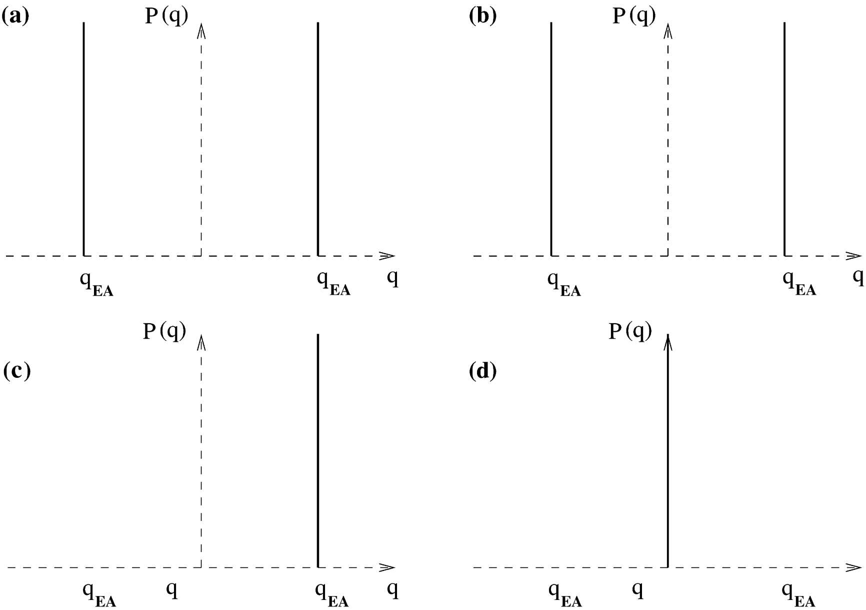

Fig. 1. The spin overlap function $P(q)$ at $T < T_c$ for: (a) a two-state picture when replica symmetry is broken *before* taking the thermodynamic limit; (b) the many-state chaotic pairs picture when replica symmetry is broken before taking the thermodynamic limit; (c) a two-state picture when replica symmetry is broken *after* taking the thermodynamic limit; (d) the many-state chaotic pairs picture when replica symmetry is broken after taking the thermodynamic limit (conjectured).

with periodic boundary conditions, while Figs. 1c and 1d represent overlap computations done in infinite volume on states randomly chosen from the respective periodic boundary condition metastates.

The insensitivity of $P(q)$ (with all else remaining equal) to these very different physical pictures indicates one potential problem with using $P(q)$ for determining ground or pure state structure. Figures 1a and 1b are identical because in either case a typical finite volume Λ_L "contains" only a single pure state pair. If one instead looks at the overlap of *all* of the infinite-volume pure states chosen (in this example) from the periodic boundary condition metastate, as in Figs. 1c and 1d, the difference is evident. In a two-state picture, one again sees a single pair of delta-functions (Fig. 1c); the thermodynamic limit here is straightforward because chaotic size dependence (periodic boundary conditions again are assumed) is absent. A many-state picture cannot have the same $P(q)$ as the two-state picture when the replica symmetry is broken *after* the thermodynamic limit is taken.

However, rather than the nontrivial $P(q)$ one might expect, the invariance of the metastate requires a very simple structure, as in Fig. 1d.

The forms of $P(q)$ sketched in Figs. 1c and 1d, however, are computed using a procedure different from that used in numerical measurements, which always use procedures corresponding to Figs. 1a and 1b. Therefore the usual measurements of $P(q)$ seem unable to provide unambiguous information on pure state multiplicity or structure in realistic spin glasses. (Other numerical methods for distinguishing between two-state and many-state pictures are described in [NS92, NS98][31, 6] and in a more recent paper[43] by us.) If a measurement of $P(q)$ in a simple geometry (e.g., a cube) and with simple boundary conditions (e.g., free or periodic) results in a complicated structure, it is likely that one is not restricting the computation to a sufficiently small box far from the boundaries (cf. Section 2.6).

2.5. Behavior at $T = 0$

If the coupling distribution is continuous, such as Gaussian, then for any finite L and, say, periodic boundary conditions, there will be only a single ground state pair $\pm\sigma^L$ in Λ_L. If droplet/scaling holds, then this pair will be the same (when restricted to a region far from the boundaries—see later) for all large L; if there are infinitely many ground state pairs, then the pair changes chaotically with L. This will be true at $T = 0$ for *any* many-state picture, whether chaotic pairs, mean-field RSB, or some other such picture. The metastates, hence overlap functions, of these many-state pictures differ only at positive temperature: the mean-field RSB picture at $T > 0$ consists of a nontrivial mixture of pure state pairs as in Eq. (12), while chaotic pairs looks similar at nonzero T to its $T = 0$ behavior. That is, in chaotic pairs at $T > 0$, the Γ appearing in any Λ_L consists of a single pure state pair, as in Eq. (13).

The overlap distributions in Fig. 1 should therefore apply (either (a) and (c) or else (b) and (d), depending on whether droplet/scaling or the chaotic pairs picture is correct) to both zero *and* nonzero temperatures less than T_c. The only temperature dependence is in the magnitude of q_{EA}.

It is particularly important to note that *there is no difference between the standard and nonstandard interpretations of the mean-field RSB pictures at* $T = 0$. It follows that overlaps of ground states *cannot* display nontrivial ultrametricity, or any other nontrivial structure.

Recent numerical results of Hed *et al.*[44] have claimed to see a nontrivial, hierarchical (though not ultrametric) ground state structure for the $\pm J$ model in $3D$. It is important to note that the theorems described in previous sections apply to discrete coupling models such as $\pm J$ as well as to continuous ones. For all of these, both the absence of non-self-averaging of overlap functions and the invariance of the metastate are rigorous conclusions. Overlap distributions must therefore have a simple, or even trivial, form regardless of the number of ground or pure states. It seems likely, therefore, that the results of Hed *et al.* are attributable to local degeneracies that appear in the $\pm J$ model, rather than to any nontrivial large-scale structures.

2.6. Windows

Before we finish this review, we briefly mention that there remain subtleties, alluded to in previous sections, in interpreting the results of overlap measurements. We refer the reader to the Appendix of [NS97b][5] for a detailed discussion of the effects of boundary conditions and of different methods of constructing $P(q)$. We also wish to emphasize a point discussed in detail in Section 6 of [NS98],[6] where we discuss why, 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. A window (always defined in reference to the volume Λ_L under investigation) is a fixed cube Λ_{L_0} in d dimensions, centered at the origin, and with $1 \ll L_0 \ll L$. The window lengthscale L_0 may be arbitrarily large, but must always be small compared to the lengthscale L of the entire volume Λ_L under consideration. In particular, when examining the pure state structure of a metastate, a window is simply any large cube centered at the origin with fixed side $L_0 \gg 1$. This is because the metastate examines pure state structure in a sequence of finite cubes Λ_L with $L \to \infty$.

When calculating $P^L(q)$ and $P_e^L(q_e)$ in Λ_L, therefore, one needs to do the overlap computation in a cube Λ_{L_0} with $L_0 \ll L$, rather than in the entire volume as is usually done. This condition is difficult to achieve numerically, but cannot be avoided if one wants to draw inferences about ordering of the low-temperature phase using overlap functions.

This is not to say that computations done in the entire volume carry no relevant or interesting information, only that their interpretation may be unclear. Such an example occurs in the numerical studies of Krzakala and Martin[45] (hereafter [KM]) and Palassini and Young[46] (hereafter [PY]). These studies claim to have uncovered a new type of excitation, which we have called *KMPY excitations*[47] (hereafter [NS01]; however, the numerical procedures used have been questioned—see the paper by Middleton[48]). It was rigorously shown in [NS01] that KMPY excitations do not yield new ground or pure states, but, if they persist on large lengthscales, could be relevant to the *excitation* spectrum in finite volumes.

3. PINNING VS. DEFLECTION OF INTERFACES, AND THERMODYNAMIC STATES

To pursue further the above idea, and also in preparation for the next section, we briefly review a basic physical feature that distinguishes thermodynamic pure states from (putative) non-thermodynamic ones. The discussion here will closely follow that of [NS01];[47] see also Section 6 of [NS98].[6]

To simplify the discussion, we focus on ground states; these are the thermodynamic pure states at $T=0$. The discussion can be extended to $T>0$ pure states by considering interfaces, equivalently domain walls, between two spin configurations chosen from different pure states.

Suppose one considers the finite-volume GSP $\pm\sigma_P^L$ corresponding to a cube Λ_L with periodic boundary conditions and L large. If one then switches to antiperiodic boundary conditions, a new GSP $\pm\sigma_{AP}^L$ is generated. The two ground state pairs will have one or more relative interfaces, or domain walls, consisting of the set of bonds $\langle x, y\rangle$ (in the dual lattice) satisfying Eq. (5). This finite-size domain wall consists of bonds that are satisfied in one but not the other GSP; it is the boundary of the set of spins that are flipped in going from $\pm\sigma_P^L$ to $\pm\sigma_{AP}^L$.

The question then arises: how could one know in principle whether there exists more than one *thermodynamic* GSP? These are infinite-volume spin configurations whose energy cannot be lowered by the flip of any *finite* subset of spins, and are generated by any convergent sequence of finite-volume ground state pairs, such as $\pm\sigma^L$ with $L\to\infty$ (see Appendix A).

The answer is that if the domain wall between $\pm\sigma^L$ and $\pm\sigma'^L$ is *pinned*, then there are multiple ground state pairs. By pinning we mean the following. Consider a fixed window of size L_0, which though finite can be arbitrarily large. Apply the procedure of generating ground state pairs $\pm\sigma_P^L$ by using periodic boundary conditions on cubes Λ_L, with $L \gg L_0$, and ground state pairs $\pm\sigma_{AP}^L$ generated with antiperiodic boundary conditions on the same cubes. Observe $\pm\sigma_P^{(L,L_0)}$ and $\pm\sigma_{AP}^{(L,L_0)}$, which are the two ground state pairs restricted to Λ_{L_0}. If their relative interface remains inside Λ_{L_0} as $L\to\infty$, then the interface is *pinned*. If there are many ground state pairs, then the interface would converge, along different subsequences of L's, to different well-defined limits inside Λ_{L_0}.

These pinned domain walls are interfaces between true thermodynamic ground state pairs. This follows because the corresponding spin configurations are limits of finite-volume ground state pairs (see Appendix A). However, another method of constructing interfaces uses a single boundary condition (typically periodic) and adds a perturbation, either by forcing a pair of spins to take an opposite relative orientation from that in the ground state, as in [KM],[45] or by adding a bulk perturbation to the Hamiltonian, as in [PY];[46] the two methods are believed to give equivalent results. Consider, e.g., the method of Krzakala-Martin. If the interface is *pinned*, then one can prove again that it separates true thermodynamic ground state pairs, as follows. As in [NS01],[47] let the two spins be chosen randomly for each Λ_L from the uniform distribution on its sites. Because in the $L \to \infty$ limit the two sites will, with probability one, be outside any fixed window, the spin configuration inside the window will have minimum energy (given its configuration on the boundary of the window). This proves the desired result, because the resulting infinite-volume spin configuration cannot have its energy lowered by flipping any finite subset of spins (which would necessarily be inside some fixed window).

Pinning of interfaces by quenched disorder occurs in disordered ferromagnets[49–51] for sufficiently large d; but these interfaces have lower dimension than the embedding space. One interesting feature of RSB is the prediction of interfaces with dimension d_s *equal* to that of the embedding space; this will be discussed in more detail in the following sections.

On the other hand, if the interface is *not* pinned, we say it "deflects to infinity." Here, for any fixed L_0, the interface, for all L above some L', will be *outside* Λ_{L_0}. This is what occurs with interface ground states in disordered ferromagnets[49–51] for small d; see Fig. 2 for a schematic illustration. If an interface deflects to infinity, then it does not give rise to new thermodynamic pure or ground states.

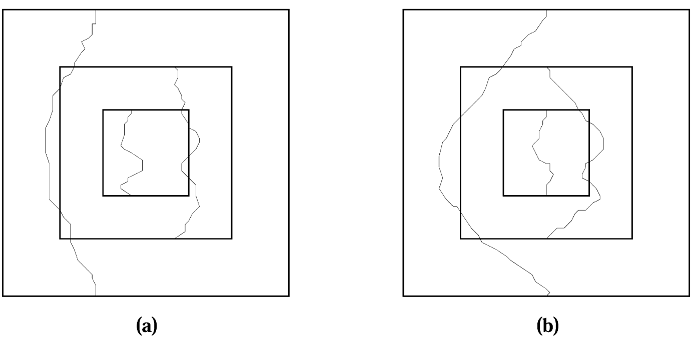

Fig. 2. A sketch of interface deflection to infinity for a $2D$ disordered ferromagnet under (a) a change from periodic to antiperiodic boundary conditions, and (b) a change from uniform plus boundary conditions to Dobrushin boundary conditions (i.e., plus spins on the left half-boundary of each square and minus on the right). As L increases, the interface recedes from the origin in each case. The interfaces eventually are completely outside any fixed square. (The deflection can scale more slowly with L than in the figure.)

4. FINITE-VOLUME PURE STATES—REPLICA SYMMETRY BREAKING'S NEW CLOTHES

The various interpretations in Section 2 of what the mean-field RSB theory could mean for realistic spin glasses all used the usual concept of "pure state" in its well-defined, traditional thermodynamic sense (see Appendix A). However, in [MPRRZ][1] it was asserted that the "pure states" that have played a central role over the past 20 years in the physical interpretation of the Parisi replica symmetry breaking scheme[9, 12–15] should *not* be thought of in this way; instead, the relevant physical objects are "finite-volume pure states" (not to be confused with the usual finite-volume *Gibbs* states, as discussed in Appendix A.) Because of the potential importance of this re-interpretation of the meaning of RSB in terms of finite-volume pure states, we now review (and critique) this claim in this section. The theorem and proof that in fact RSB *must* involve the more traditional *thermodynamic* pure states, which is the central result of this paper, will be presented in subsequent sections.

The main new theoretical idea of [MPRRZ], introduced and discussed in its Section 3, is the attempt to formalize the relation between RSB and state structure for short-ranged models via the notion of *finite volume pure states*. This interpretation is contrasted with the "not appropriate use of Eq. (35) to describe an infinite system." Equation (35) of [MPRRZ] is simply:

$$\langle \cdot \rangle = \sum_{\alpha} w_{\alpha} \langle \cdot \rangle_{\alpha}, \tag{15}$$

where α is a "pure state" index and w_α its Boltzmann weight.

Such a decomposition of course *can* be done and is well-defined (see, e.g., the book of Georgii[52]) for the usual thermodynamic pure states α in infinite volume. It can also be done in a well-defined sense for *finite* volumes and is closely related to the idea of "window overlaps;" these ideas are introduced and discussed in [NS98].[6] In both cases, our theorems[2–6] apply and rule out any of the interpretations in Section 2 of the RSB mean-field picture in finite-dimensional systems. But a central point of [MPRRZ] is that in RSB theory the decomposition Eq. (15) of a finite-volume Gibbs state does *not* involve these traditional pure states, but rather a decomposition into something else that does not have an infinite-volume definition or meaning.

This is problematic in that it not only contradicts earlier statements of the same authors, but also conflicts with other sections within the same paper. As to the former, references to pure states within the mean-field RSB scheme that imply the usual thermodynamic definition are too numerous to list. As just one example, the RSB literature, whether discussing infinite-ranged or finite-ranged spin glasses, repeatedly refers to pure states as those having the clustering property given in Eq. (A4) (see, for example, [P83],[13] Section 3.6 of [BY],[8] Section 3.1 in [MPV],[9] [P92][21]); but clustering appears to be an exclusive property of thermodynamic pure states. Much of the other terminology frequently used in the literature, such as "valleys," is vaguer. Nevertheless, it is hard to interpret the often-made claim that RSB pure states are separated by infinitely high barriers (see, e.g., Section 4.5 of [BY] and Section 7.1 (and 3.1) of [MPV]), or the dynamical assertion that a spin glass in one of these states would thereafter spend an infinite amount of time in that state (see, e.g., Section 4.5 of [BY] and Section 7.1 (and 3.1) of [MPV]), as referring to anything other than thermodynamic pure states.

It is implied in Section 3.1 of [MPRRZ] that finite volume pure states first appeared in a 1987 paper of Parisi[53] (hereafter [P87]), more than a decade earlier than [MPRRZ]. We believe that this is unjustified. There does not appear to be any discussion about finite volume pure states in [P87], but rather discussion about "pure clustering states." As already noted, clustering is a property that belongs to standard thermodynamic pure states. Moreover, in [P87] the pure state decomposition Eq. (15) is justified several times on the basis of a theorem in Ruelle's book[54] that is explicitly about thermodynamic pure states.

Even more seriously, there is a direct contradiction between the claim that RSB refers only to "nonthermodynamic" pure states and the discussions in Sections 2.2, 2.3, 8.4, and 8.5 of [MPRRZ], where window overlaps (cf. Section 2.6) are discussed (see also a paper of Marinari *et al.*[55]). Equations (24)–(27) of [MPRRZ] concern the predictions of RSB for the spin overlap distribution confined to a small region in the center of the cube; i.e., a window. But these predictions are precisely those that would be made by either the standard or the nonstandard interpretation of the mean-field RSB picture described in Sections 2.2 and 2.3; window overlaps were especially constructed [NS98][6] so as to separate properties arising from *thermodynamic* pure state structure from those due to boundary or other effects.

The final sentence of [MPRRZ] asserts that "the recent rigorous results by Newman and Stein strongly support RSB." Given the discussion in this section and the earlier demonstration that our results imply that the *only* sensible many-state picture is chaotic pairs, it should be clear that we strongly dispute such a claim. Its basis is discussed in Sections 7.4 and 7.5 of [MPRRZ] (some of which appeared earlier as an unpublished posting of Parisi[56]). Our discussion about why the arguments used in ref. 56 (and [MPRRZ]) do not support this assertion is provided in our own unpublished posting,[57] to which we refer the reader.

Indeed, we will present later a rigorous result that, regardless of any new intepretation of RSB based on finite volume pure states, excludes the possibility that the mean-field RSB theory describes the low-temperature spin glass phase in any finite dimension.

5. RSB PREDICTIONS FOR LINK OVERLAPS AND DOMAIN WALLS

We turn now to the central argument of our paper. An unambiguous prediction of the mean-field RSB approach, applied to finite-dimensional spin glasses, is that the edge overlap distribution function $P_e(q_e)$ (Eq. (7)) is nontrivial on all length scales. Extensive discussion of the predictions of the RSB theory for realistic spin glasses is given in several places; see in particular [MPRRZ],[1] [PY],[46] [MP00a][16] and [MP00b],[17] to which we refer the reader for details.

It was noted in [MP00a, MP00b] that nontriviality in $P_e(q_e)$ at $T=0$ can be ascribed to the presence of space-filling domain walls between ground states generated from different boundary conditions, or between ground and excited states with the latter generated through a perturbation (cf. Section 3). This important feature of RSB theory can be described in the following way.

Consider a d-dimensional cube Λ_L, centered at the origin and with periodic boundary conditions, so that for a given coupling realization $\mathcal{J}^L$ inside Λ_L there exists a ground state pair $\pm\sigma^L$. Consider as before the spin configuration generated by forcing a random pair of spins to take on an opposite orientation from that in $\pm\sigma^L$, and then letting the resulting configuration relax to a new state σ'^L with minimum energy subject to this constraint (alternatively, one could add a bulk perturbation as in [PY, MP00b]. Then a central *physical* feature of the mean-field RSB picture ([MPRRZ, MP00a, MP00b]) is that $\pm\sigma^L$ and $\pm\sigma'^L$ differ in the following ways: (1) their difference is global, i.e., there are $O(L^d)$ spins flipped in going from $\pm\sigma^L$ to $\pm\sigma'^L$; (2) the lengthscale l of their relative interface is $O(L)$, and the number of bonds in the interface scales as L^{d_s} with $d_s=d$, i.e., the interface is space-filling; and (3) the energy of the relative interface remains of order one independently of $l=O(L)$ so that the domain wall energy scales as $l^{\theta'}$ with $\theta'=0$. Domain walls with these properties will henceforth be called *RSB interfaces*. It is easy to see that, at $T=0$, properties (1) and (2) already give rise to nontrivial $P_e(q_e)$ (and conversely that nontrivial $P_e(q_e)$ implies the existence of interfaces with those properties).

What about $T>0$? Now, because RSB asserts that each individual low-temperature Gibbs state is a mixture of several states, it predicts that even without any perturbations, there will be nontrivial $P(q)$ and $P_e(q_e)$ inside Λ_L as described in Section 2.1. This prediction relies heavily on property (3) of the RSB interfaces, namely that their energies are $O(1)$; if properties (1) and (2) were valid, but not (3), this would lead to the chaotic pairs picture (cf. Section 2.4) with many states but trivial overlap distribution.

We note, however, that there is now a problem in interpretation, especially for $P_e(q_e)$, because one needs to disentangle effects due to potential multiple states from those due to normal thermal fluctuations. One way of doing this was discussed in Section 5 of [NS92];[31] here one looks at two different cube sizes and uses the presence or absence of chaotic size dependence to differentiate between the two effects. We propose another way here. It is known that, if the probability density function of the couplings is bounded by a constant C, as in the usual Gaussian coupling case, then[31]

$$1-\overline{\langle \sigma_x \sigma_y \rangle^2} \leqslant 2Ck_BT \tag{16}$$

in a cube Λ_L with coupling-independent boundary conditions, such as periodic. Here an overbar represents an average over coupling realizations. This bound is rigorously obeyed by a Gibbs state generated from a *single* boundary condition, regardless of how many pure states it contains.

But now suppose that one generates *two* Gibbs states at $T>0$ in Λ_L, e.g., one with and one without a Palassini–Young bulk perturbation ([PY]), as in [MP00b]. Then it should still be true that the contribution to $P_e(q_e)$ from trivial thermal excitations would remain bounded by $O(T)$, but the contribution from multiple RSB-like states, if present, would not obey any such bound. Therefore, at sufficiently low T, thermal contributions to $P_e(q_e)$ should be negligible compared to putative RSB contributions. (As a consequence, we suggest that results obtained at higher temperatures, like $0.7T_c$ as in some of the simulations in [MPRRZ], are not useful in verifying the applicability of RSB theory to realistic spin glasses.)

We now address the central question of this paper: are RSB interfaces, which comprise a central feature of mean-field RSB theory, compatible with the claims of [MPRRZ] that a thermodynamic interpretation of RSB pure states can be avoided? In other words, can RSB domain walls *avoid* giving rise to many traditional thermodynamic pure states?

We will provide a proof in the next section that the answer is *no*; these central predictions of mean-field RSB theory are *rigorously incompatible* with each other. The prediction of RSB interfaces means that multiple *thermodynamic* pure states, with properties that have been ruled out in our previous papers,[2–6] must appear. The mean-field RSB theory is therefore *inconsistent* in any finite dimension.

Before we turn to our rigorous result, we provide a heuristic argument that illustrates the central idea of our theorem and makes clear why RSB interfaces must give rise to thermodynamic pure states. Recall from Section 3 that if a domain wall, generated by switching from periodic to antiperiodic boundary conditions, is pinned, then it *must* give rise to thermodynamic pure states whose relative interface is that same domain wall. So in order for RSB interfaces to be both space-filling and *not* give rise to thermodynamic pure states, *they must deflect out of any fixed region* as

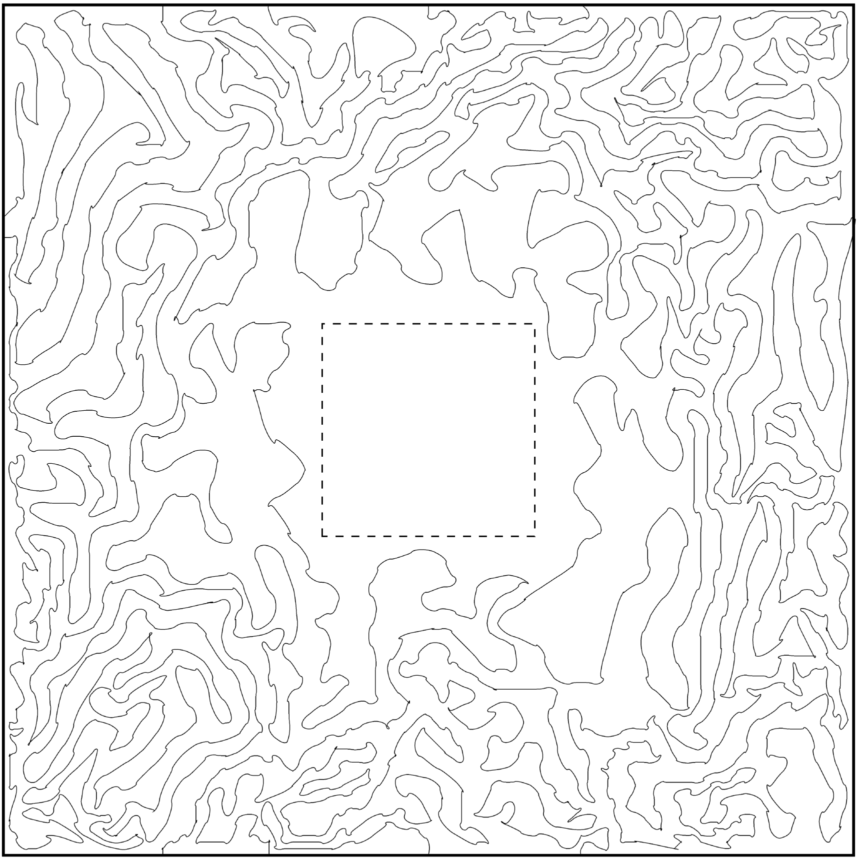

Fig. 3. A sketch of RSB interface deflection to infinity, in the situation where a bulk perturbation is applied to a volume with periodic boundary conditions at zero temperature, as described in the text. (Two dimensions is shown only for illustrative simplicity, and is not meant to imply that mean-field RSB theory is expected to apply there.) In this figure a single, positive-density interface is depicted. Whether the interface consists of a single or many domain walls is irrelevant, so long as their union has density that scales as L^d.

$L \to \infty$. The resulting situation is shown in Fig. 3, which would have to occur on all large lengthscales.

It seems already clear that such a situation is highly improbable; in the next section we prove that, indeed, it cannot occur.

6. THEOREM AND PROOF

Consider again the pair of spin configurations σ^L and σ'^L discussed in Section 5. We will use the *uniform perturbation metastate* introduced in [NS01].[47] Here one does for the *pair* $(\pm\sigma^L, \pm\sigma'^L)$ what was done for $\pm\sigma^L$ in the original metastate. The resulting metastate gives, among other things, a translation-invariant $(\mathcal{J}, \widetilde{\mathcal{D}}_{\mathcal{J}})$, where $\widetilde{\mathcal{D}}_{\mathcal{J}}$ is a domain wall measure that provides the $(L \to \infty)$ probability that a given bond belongs to the relative domain wall between $\pm\sigma^L$ and $\pm\sigma'^L$ inside an arbitrary large cube Λ_L.

Now let $(\pm\sigma, \pm\sigma')$ be chosen randomly from the $T = 0$ metastate described in the last paragraph. The argument in [NS01] shows that, for almost every $(\mathcal{J}, \pm\sigma, \pm\sigma')$, either $\pm\sigma' = \pm\sigma$ or else the two infinite-volume ground state pairs have a relative interface of strictly positive density (i.e., σ and σ' are *incongruent*, in the terminology of Huse and Fisher[28]). We therefore already know that if there is a pinned interface at all, it must have strictly positive density, i.e., $d_s = d$. We now prove the converse as well.

That is, we prove that if such a space-filling interface exists, then the situation depicted in Fig. 3 *cannot* happen. Such interfaces must be *pinned*, i.e., they have *strictly positive* probabilities of remaining inside any large fixed window Λ_{L_0} as the outer cube size $L \to \infty$. Moreover, the fraction of bonds in the domain wall that remain *inside* the window scales as $(L_0/L)^d$.

Theorem. On each cube Λ_L, consider torus-translation-invariant $(\mathcal{J}^L, \widetilde{\mathcal{D}}_{\mathcal{J}^L})$, a sequence of random couplings and domain wall measures (from, e.g., the triple $(\mathcal{J}^L, \pm\sigma^L, \pm\sigma'^L)$). Let $(\mathcal{J}, \widetilde{\mathcal{D}}_{\mathcal{J}})$ be any limit in distribution as $L \to \infty$ of $(\mathcal{J}^L, \widetilde{\mathcal{D}}_{\mathcal{J}^L})$. Then if the probability that a particular edge belongs to a domain wall is bounded away from zero as $L \to \infty$, there must be at least a positive fraction of the ergodic components of $(\mathcal{J}, \widetilde{\mathcal{D}}_{\mathcal{J}})$ that have a positive density of domain walls.

Proof. Because the joint distribution of $(\mathcal{J}^L, \widetilde{\mathcal{D}}_{\mathcal{J}^L})$ is, for every L, invariant under torus translations, any limiting distribution $(\mathcal{J}, \widetilde{\mathcal{D}}_{\mathcal{J}})$ is invariant under all translations of the infinite-volume cubic lattice $\mathbf{Z}^d$. The translation-invariance of $(\mathcal{J}, \widetilde{\mathcal{D}}_{\mathcal{J}})$ allows its decomposition into components in which translation-ergodicity holds (see, e.g., [NS96a, NS97b][2,5]). For each bond $\langle x, y\rangle$ consider the event $A_{\langle x, y\rangle}$ that $\langle x, y\rangle$ is in the domain wall. If the probability, with respect to $(\mathcal{J}^L, \tilde{\mathcal{D}}_{\mathcal{J}^L})$, of the event $A_{\langle x, y\rangle}$ occurring in Λ_L is larger than some fixed $\rho > 0$ independent of L, then any limiting measure must also have the probability that $A_{\langle x, y\rangle}$ occurs being strictly positive (and greater than ρ). Because the translation-invariant measure $(\mathcal{J}, \tilde{\mathcal{D}}_{\mathcal{J}})$ therefore has $A_{\langle x, y\rangle}$ occurring with a probability $P(A_{\langle x, y\rangle}) > \rho > 0$, it follows that in a positive fraction of its ergodic components, the probability of $A_{\langle x, y\rangle}$ occurring is also strictly greater than zero. In each of these ergodic components, by the spatial ergodic theorem (see, e.g., [NS96a, NS97b]) the spatial density of $\langle x, y\rangle$'s such that $A_{\langle x, y\rangle}$ occurs must equal a strictly positive number, i.e., the interface has a nonzero density.

Remark 1. Although the theorem as formulated here addresses domain walls between ground states, it should be extendable to domain walls between spin configurations chosen from different pure states at low temperature, by "pruning" small thermally induced droplets.[47]

Remark 2. Note that the third property of RSB interfaces, namely that their energy remains of $O(1)$ independently of L, was not needed; the theorem applies to *any* space-filling domain wall constructed as discussed. The theorem therefore applies to RSB excitations as a special case, but to other kinds as well. These will be discussed further in Section 7.

We now apply the theorem to RSB interfaces. These satisfy the condition that the probability that an arbitrarily chosen bond belongs to a $(\pm\sigma^L, \pm\sigma'^L)$ domain wall is bounded away from zero as $L \to \infty$. It therefore follows that, if RSB interfaces exist, and if one chooses a random *infinite-volume* GSP $(\pm\sigma, \pm\sigma')$ from the metastate, then there must be a positive probability that any given bond belongs to that interface. This is equivalent to the statement that, for every L as $L \to \infty$, there is a positive probability of finding an RSB interface inside any fixed window Λ_{L_0} of arbitrary (but finite) size L_0.

Moreover, if pinned RSB interfaces are present at $T = 0$, then they would presumably give rise to multiple pure state pairs at low but nonzero temperature. Here it would be the case that, in addition to the expected thermal fluctuations, two spin configurations, each randomly chosen from different pure states (not globally flip-related), would have a relative RSB interface. Equivalently, one could determine the existence of these positive-temperature interfaces by examining the thermal expectations of, e.g., two-point correlation functions.

It follows that the mean-field RSB theory *must give rise to multiple thermodynamic ground state pairs at* $T = 0$, *and by extension, pure state pairs as conventionally defined (cf. Appendix A) also at low* T.

7. DISCUSSION AND CONCLUSION

We have shown that the claim in [MPRRZ][1] (elaborated on in Sections 3 and 7 of that paper) that mean-field RSB theory does *not* give rise to the usual thermodynamic pure states is in rigorous contradiction with the simultaneous claim that the theory also predicts interfaces (equivalently, nontrivial link overlap distribution functions) with the properties delineated in Section 5. Then, in short-ranged spin glasses in finite dimensions, either there are *no* interfaces that are both space-filling, with $d_s = d$, and have energy of $O(1)$, or else there are and they comprise domain walls between distinct thermodynamic pure state pairs. We investigate each of these possibilities in turn.

The first possibility is that there are no domain walls that are both space-filling and have $O(1)$ energy. Suppose we relax the second requirement, so that the energy of the interface increases with L. The simplest possibility is that this energy scales as $L^{\theta'}$ with $\theta' > 0$, although more slowly increasing functions are also possible (e.g., $\log(L)$); but because the argument is the same for both, we simply examine the case $\theta' > 0$. But when $\theta' > 0$, one recovers[6] the chaotic pairs picture, as noted in Section 5; multiple pure state pairs cannot now coexist at any T within a single Λ_L with large L.

Relaxing the first requirement implies that there are either no domain walls at all, in which case one recovers a two-state picture, or else there are interfaces with $d_s < d$. The latter would not be seen in any $T = 0$ metastate constructed using coupling-independent boundary conditions,[43, 58] and may or may not give rise to new pure states at $T > 0$ depending on how they are constructed; an extensive discussion is given in [NS01].[47] In either case any metastate constructed using coupling-independent boundary conditions would see at most only a single pair of pure states (cf. the discussion of "invisible" states within the metastate given in Section 5 of [NS98][6]).

We turn now to the second possibility in which space-filling domain walls of $O(1)$ energy *are* present on all lengthscales. Now the theorem in Section 6 necessitates the existence of multiple pure state pairs so that the thermodynamic states Γ (cf. Section 2.4) would be mixed states at $T > 0$.

The conclusion is that regardless of any re-interpretations of the "meaning" of the RSB ansatz for finite-dimensional spin glasses, it has an unambiguous prediction for the structural difference in ground states generated in a large cube when a KMPY-type perturbation ([KM, PY][45, 46]), as examined in [MP00b],[17] is applied. This prediction is the presence of space-filling interfaces. But we have shown here that this feature gives rise to multiple *infinite-volume* pure states. So the RSB ansatz predicts these multiple thermodynamic states whether it was originally intended to or not.

Moreover, given that these RSB interfaces have energies of $O(1)$, the thermodynamic states they give rise to would necessarily appear as mixed states in large finite volumes, as in Eq. (12). But this possibility was ruled out in our earlier papers[2–6] (and was disavowed in [MPRRZ]).

Mean-field RSB theory therefore can not describe the low-temperature structure of the spin glass phase in any *finite* dimension, although of course RSB theory presumably remains valid for infinite-ranged models. If the low-temperature spin glass phase displays multiple pure states in any finite dimension, their structure would have to be given by the chaotic pairs picture of [NS96b, NS97, NS98][3, 5, 6] and spin overlap structures inside any window would be trivial regardless of how the overlaps are constructed.

APPENDIX A. GIBBS STATES IN FINITE AND INFINITE VOLUME

In this appendix, we present some background information about Gibbs states, closely following the discussion in [NS97].[5] Given the EA Hamiltonian (1) on Λ_L with a specified boundary condition (e.g., free, fixed, periodic, etc.), the finite-volume Gibbs state $\rho^{(L)}_{\mathcal{J},T}$ on Λ_L at temperature T is defined by:

$$\rho^{(L)}_{\mathcal{J},T}(\sigma) = Z^{-1}_{L,T} \exp\{-\mathcal{H}_{\mathcal{J},L}(\sigma)/k_B T\}, \tag{A1}$$

where the partition function $Z_{L,T}$ is such that the sum of $\rho^{(L)}_{\mathcal{J},T}$ over all spin configurations in Λ_L yields one.

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

A *thermodynamic* state $\rho_{\mathcal{J},T}$ is defined as an *infinite*-volume Gibbs measure, containing information such as the probability of any finite subset of spins taking on specified values. Thermodynamic states can be constructed by taking the $L \to \infty$ limit of a sequence of finite-volume Gibbs states $\rho^{(L)}_{\mathcal{J},T}(\sigma)$, each with a specified boundary condition (which may remain the same or may change with L). 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. Infinite-volume Gibbs measures $\rho_{\mathcal{J},T}$ can also be characterized independently of any limiting process, as probability measures on infinite-volume spin configurations that satisfy the Dobrushin–Lanford–Ruelle (DLR) equations (for a mathematically detailed presentation, see the book of Georgii[52]).

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

$$\rho_{\mathcal{J},T} = \lambda \rho^1_{\mathcal{J},T} + (1-\lambda)\, \rho^2_{\mathcal{J},T}, \tag{A2}$$

where $0 < \lambda < 1$ and ρ^1 and ρ^2 are also infinite-volume Gibbs states (distinct from ρ), then $\rho_{\mathcal{J},T}$ is a *mixed* thermodynamic state or simply, mixed state. A mixed state may have as few as two or as many as an uncountable infinity of states in its decomposition. The meaning of Eq. (A2) can be understood as follows: any correlation function computed using the thermodynamic state $\rho_{\mathcal{J},T}$ can be decomposed in the following way:

$$\langle \sigma_{x_1} \cdots \sigma_{x_m} \rangle_{\rho_{\mathcal{J},T}} = \lambda \langle \sigma_{x_1} \cdots \sigma_{x_m} \rangle_{\rho^1_{\mathcal{J},T}} + (1-\lambda) \langle \sigma_{x_1} \cdots \sigma_{x_m} \rangle_{\rho^2_{\mathcal{J},T}}. \tag{A3}$$

If a state cannot be written as a convex combination of any other infinite-volume Gibbs states, it is then a thermodynamic *pure* 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 $\lambda = 1/2$. A thermodynamic pure state ρ_P can be intrinsically characterized by a *clustering property* (see, e.g., refs. 52 and 59), which implies that for any fixed x,

$$\langle \sigma_x \sigma_y \rangle_{\rho_P} - \langle \sigma_x \rangle_{\rho_P} \langle \sigma_y \rangle_{\rho_P} \to 0, \qquad |y| \to \infty, \tag{A4}$$

and similar clustering for higher order correlations.

Finite-volume Gibbs states, which are well-defined probability measures, should not be confused with the putative "finite-volume pure states" of [MPRRZ],[1] which have not been clearly defined. A finite-volume Gibbs state can have an *approximate* decomposition into *thermodynamic* pure states restricted to a "window"[6] deep inside Λ_L, as in Eq. (2). Whether a similar decomposition of finite-volume Gibbs states into "finite-volume pure states" can be made is unclear; it would at the least require making the notion of finite-volume pure state more precise.

APPENDIX B. GLOSSARY

We include this glossary for the reader's convenience. All definitions here are informal. Terms that have appeared only recently in the literature, or that may be less familiar, are also defined within the text; in such cases, the section where they are first defined is also noted.

Chaotic Pairs Picture. A scenario for the low-temperature spin glass phase in which there exist infinitely many (incongruent) pure state pairs (for a.e. $\mathcal{J}$) at all temperatures below T_c, but with probability one only a single one of these pairs would be seen in any large volume with periodic boundary conditions. The overlap function computed in any large volume is therefore indistinguishable from a two-state picture like droplet/scaling (cf. Fig. 1). However, the pair chosen varies chaotically with volume. (Section 2.4.)

Chaotic Size Dependence. Inside any large volume Λ_L with specified boundary conditions, the Gibbs state is approximately either a single pure state (e.g., in a homogeneous Ising ferromagnet, one has a paramagnet above T_c for any boundary condition, the magnetized plus state below T_c for all plus spins at the boundary, etc.), or else an approximate decomposition over pure states as in Eq. 12 (e.g., in the same system below T_c the Gibbs state is an equal mixture of the magnetized plus and minus states). Chaotic size dependence occurs when the pure states and/or weights vary persistently as L is increased, so that there is no limiting infinite-volume Gibbs state. (Sections 2.1 and 2.4.)

Deflection to Infinity. Consider an interface between two ground or pure states in Λ_L generated either by a change in boundary condition (e.g., periodic to antiperiodic), or by addition of a perturbation with a single boundary condition. Consider a volume Λ_{L_0} of arbitrary but fixed side L_0. If, for any L_0, the relative interface eventually moves (and stays) outside of Λ_{L_0} as $L \to \infty$, the interface has "deflected to infinity." See Fig. 2 for an illustration. (Section 3.)

Droplet/Scaling. A two-state picture (see later) whose properties follow from a scaling *ansatz* developed by Macmillan,[24] Bray and Moore,[25,26] and Fisher and Huse;[27–30] the last of these fully developed the physical droplet picture corresponding to the scaling ansatz, which followed from "domain wall" renormalization-group studies of the first two groups. In this picture, the thermodynamic and dynamic properties of spin glasses at low temperature are dominated by low-lying excitations corresponding to clusters of coherently flipped spins. The density of states of these clusters at zero energy falls off as a power law in lengthscale L, with exponent bounded from above by $(d-1)/2$. At low temperatures and on large lengthscales the density of thermally activated clusters is dilute and they can be considered as non-interacting two-level systems.

Ground State. In a finite volume Λ_L, the lowest-energy state(s) consistent with the boundary conditions. A convergent sequence of finite-volume ground states yields an infinite-volume ground state, which is simply a pure state (as in Appendix A) at $T = 0$. An infinite-volume ground state can alternatively (and often more usefully) be defined as an infinite-volume spin configuration whose energy cannot be lowered by the flip of any finite subset of spins. (Section 3.)

Ground State Pair (GSP). In the absence of an external field or spin-flip symmetry breaking boundary conditions, ground states occur in pairs related by a global spin flip.

Incongruence. Two spin configurations are incongruent (a notion introduced by Huse and Fisher[28]) if they differ by a relative flip along a space-filling interface; that is, a nonzero density of bonds is satisfied in one but not the other spin configuration. If the relative interface has zero density, the spin configurations are said to be *regionally congruent*. (Section 6.)

Pinning. Given the same scenario as in the definition of *deflection to infinity* above, the interface remains inside a sufficiently large volume of fixed size L_0 as $L \to \infty$. (Section 3.)

Metastate. A probability measure on infinite-volume thermodynamic states that carries all relevant thermodynamic information about a system. In the current context, the metastate provides, among other things, the probability of appearance of various pure or ground states appearing within a large finite volume Λ_L with specified boundary conditions. (Section 2.4.)

RSB Interface. An interface between two globally different spin configurations (that are not global flips of each other) that has the properties of being both space-filling and also of having approximately order one energy independently of lengthscale. (Section 5.)

Two-State Picture. A scenario for the low-temperature spin glass phase in which there exists only a single pair of (spin-flip-related) pure states at all temperatures below T_c. Although not used in the text, we note that these can be divided into at least two kinds. A *strong* two-state picture is one where there are no more than two pure states at any temperature, as in the droplet/scaling picture. A *weak* two-state picture is one where there exists a "special" pair of pure states that supports any metastate generated by coupling-independent boundary conditions, but in which there also exist other pure states that can be generated only by coupling-dependent boundary conditions. These latter states are "invisible" in any coupling-independent b.c. metastate. (This possibility for spin glasses is very briefly discussed in Section 7.) An example of a weak two-state picture could be the homogeneous Ising ferromagnet, at $T = 0$ in two dimensions and below the roughening temperature in three and higher dimensions. Here the special pair is the uniformly magnetized plus and minus states, while the others are the interface states.

Ultrametricity. In the spin glass context, the property that the joint overlap statistics of any three pure states with overlaps q_1, q_2, and q_3 satisfy the condition $q_1 = q_2 \leqslant q_3$, consistent with a hierarchical pure state structure.

Window. Given a large volume Λ_L with specified boundary conditions, a window is an interior volume Λ_{L_0} with $1 \ll L_0 \ll L$. Both Λ_L and Λ_{L_0} are centered at the origin. We argue in [NS98][(6)] that an overlap computation must be done inside a window if it is to reveal any unambiguous information about pure state structure. (Section 2.6.)

ACKNOWLEDGMENTS

C.M.N. was partially supported by the National Science Foundation under Grant DMS-01-02587. D.L.S. was partially supported by the National Science Foundation under Grant DMS-01-02541.

We are grateful to Peter Young for useful correspondence, and to Michael E. Fisher for numerous helpful comments and suggestions on the manuscript. We also thank David Huse and Joel Lebowitz for their comments.

REFERENCES

1. E. Marinari, G. Parisi, F. Ricci-Tersenghi, J. J. Ruiz-Lorenzo, and F. Zuliani, *J. Statist. Phys.* **98**:973 (2000). (Referred to as [MPRRZ] in the text.)
2. C. M. Newman and D. L. Stein, *Phys. Rev. Lett.* **76**:515 (1996). (Referred to as [NS96a] in the text.)
3. C. M. Newman and D. L. Stein, *Phys. Rev. Lett.* **76**:4821 (1996). (Referred to as [NS96b] in the text.)
4. C. M. Newman and D. L. Stein, in *Mathematics of Spin Glasses and Neural Networks*, A. Bovier and P. Picco, eds. (Birkhäuser, Boston, 1997), pp. 243–287. (Referred to as [NS97a] in the text.)
5. C. M. Newman and D. L. Stein, *Phys. Rev. E* **55**:5194 (1997). (Referred to as [NS97b] in the text.)
6. C. M. Newman and D. L. Stein, *Phys. Rev. E* **57**:1356 (1998). (Referred to as [NS98] in the text.)
7. E. Marinari, G. Parisi, and J. J. Ruiz-Lorenzo, in *Spin Glasses and Random Fields*, A. P. Young, ed. (World Scientific, Singapore, 1997), pp. 59–98.
8. K. Binder and A. P. Young, *Rev. Mod. Phys.* **58**:801 (1986). (Referred to as [BY] in the text.)

9. M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987). (Referred to as [MPV] in the text.)
10. S. Edwards and P. W. Anderson, *J. Phys. F* **5**:965 (1975).
11. D. Sherrington and S. Kirkpatrick, *Phys. Rev. Lett.* **35**:1792 (1975).
12. G. Parisi, *Phys. Rev. Lett.* **43**:1754 (1979).
13. G. Parisi, *Phys. Rev. Lett.* **50**:1946 (1983). (Referred to as [P83] in the text.)
14. M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, *Phys. Rev. Lett.* **52**:1156 (1984).
15. M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, *J. Phys. (Paris)* **45**:843 (1984).
16. E. Marinari and G. Parisi, *Phys. Rev. B* **62**:11677 (2000). (Referred to as [MP00a] in the text.)
17. E. Marinari and G. Parisi, *Phys. Rev. Lett.* **86**:3887 (2001). (Referred to as [MP00b] in the text.)
18. S. Caracciolo, G. Parisi, S. Patarnello, and N. Sourlas, *J. Phys. France* **51**:1877 (1990).
19. J. D. Reger, R. N. Bhatt, and A. P. Young, *Phys. Rev. Lett.* **64**:1859 (1990).
20. S. Caracciolo, *Physica A* **185**:261 (1992).
21. G. Parisi, *Physica A* **185**:316 (1992). (Referred to as [P92] in the text.)
22. G. Parisi, *Physica A* **194**:28 (1993).
23. E. Marinari, G. Parisi, J. J. Ruiz-Lorenzo, and F. Ritort, *Phys. Rev. Lett.* **76**:843 (1996).
24. W. L. McMillan, *J. Phys. C* **17**:3179 (1984).
25. A. J. Bray and M. A. Moore, *Phys. Rev. B* **31**:631 (1985).
26. A. J. Bray and M. A. Moore, *Phys. Rev. Lett.* **58**:57 (1987).
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. C. M. Newman and D. L. Stein, *Phys. Rev. B* **46**:973 (1992). (Referred to as [NS92] in the text.)
32. R. Rammal, G. Toulouse, and M. A. Virasoro, *Rev. Mod. Phys.* **58**:765 (1986).
33. G. Parisi and F. Ricci-Tersenghi, *J. Phys. A* **33**:113 (2000).
34. F. Guerra, private communication; see also *Int. J. Mod. Phys. B* **19**:1675 (1996).
35. E. Vincent, J. Hammann, and M. Ocio, in *Recent Progress in Random Magnets*, D. H. Ryan, ed. (World Scientific, Singapore, 1992), p. 207.
36. D. Badoni, J. C. Ciria, G. Parisi, F. Ritort, J. Pech, and J. J. Ruiz-Lorenzo, *Europhys. Lett.* **21**:495 (1993).
37. C. Külske, *J. Statist. Phys.* **88**:1257 (1997).
38. C. Külske, *J. Statist. Phys.* **91**:155 (1998).
39. A. Bovier and V. Gayrard, *Math. Phys. Anal. Geom.* **1**:107 (1998).
40. A. Bovier, A. C. D. van Enter, and B. Niederhauser, *J. Statist. Phys.* **95**:181 (1999).
41. M. Aizenman and J. Wehr, *Commun. Math. Phys.* **130**:489 (1990).
42. C. M. Newman and D. L. Stein, *Phys. Rev. E* **60**:5244 (1999).
43. C. M. Newman and D. L. Stein, *Phys. Rev. Lett.* **84**:3966 (2000).
44. G. Hed, A. K. Hartmann, D. Stauffer, and E. Domany, *Phys. Rev. Lett.* **86**:3148 (2001).
45. F. Krzakala and O. C. Martin, *Phys. Rev. Lett.* **85**:3013 (2000). (Referred to as [KM] in the text.)
46. M. Palassini and A. P. Young, *Phys. Rev. Lett.* **85**:3017 (2000). (Referred to as [PY] in the text.)
47. C. M. Newman and D. L. Stein, *Phys. Rev. Lett.* **87**, 077201 (2001). (Referred to as [NS01] in the text.)

48. A. A. Middleton, *Phys. Rev. B* **63**, 060202(R) (2001).
49. D. A. Huse and C. L. Henley, *Phys. Rev. Lett.* **54**, 2708 (1985).
50. M. Kardar, *Phys. Rev. Lett.* **55**, 2923 (1985).
51. A. Bovier and C. Külske, *Rev. Math. Phys.* **6**, 413 (1994).
52. H.O. Georgii, *Gibbs Measures and Phase Transitions* (de Gruyter Studies in Mathematics, Berlin, 1988).
53. G. Parisi, *Physica Scripta* **35**:123 (1987). (Referred to as [P87] in the text.)
54. D. Ruelle, *Statistical Mechanics: Rigorous Results* (W. A. Benjamin, Reading, MA, 1969).
55. E. Marinari, G. Parisi, F. Ricci-Tersenghi, and J. J. Ruiz-Lorenzo, *J. Phys. A* **31**:L481 (1998).
56. G. Parisi, cond-mat/9603101.
57. C. M. Newman and D. L. Stein, adap-org/9603001.
58. C. M. Newman and D. L. Stein, to appear in *Commun. Math. Phys.*
59. A. C. D. van Enter and J. L. van Hemmen, *Phys. Rev. A* **29**:355 (1984).