

Ordering and broken symmetry in short-ranged spin glasses

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2003 J. Phys.: Condens. Matter 15 R1319

(<http://iopscience.iop.org/0953-8984/15/32/202>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 171.66.16.125

The article was downloaded on 19/05/2010 at 14:59

Please note that [terms and conditions apply](#).

TOPICAL REVIEW

Ordering and broken symmetry in short-ranged spin glasses

C M Newman¹ and D L Stein²¹ Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA² Departments of Physics and Mathematics, University of Arizona, Tucson, AZ 85721, USAE-mail: newman@cims.nyu.edu and dls@physics.arizona.edu

Received 31 January 2003, in final form 29 May 2003

Published 1 August 2003

Online at stacks.iop.org/JPhysCM/15/R1319**Abstract**

In this topical review we discuss the nature of the low-temperature phase in both infinite-ranged and short-ranged spin glasses. We analyse the meaning of pure states in spin glasses, and distinguish between physical, or ‘observable’, states and (probably) unphysical, ‘invisible’ states. We review replica symmetry breaking (RSB), and describe what it would mean in short-ranged spin glasses. We introduce the notion of thermodynamic chaos, which leads to the metastate construct. We apply these tools to short-ranged spin glasses, and conclude that RSB, in any form, cannot describe the low-temperature spin glass phase in any finite dimension. We then discuss the remaining possible scenarios that *could* describe the low-temperature phase, and the differences they exhibit in some of their physical properties—in particular, the interfaces that separate them. We also present rigorous results on metastable states and discuss their connection to thermodynamic states. Finally, we discuss some of the differences between the statistical mechanics of homogeneous systems and those with quenched disorder and frustration.

Contents

1. Introduction	1320
2. A brief history of early theoretical developments	1323
2.1. The EA Hamiltonian	1323
2.2. Mean-field theory, the Sherrington–Kirkpatrick model and the Parisi solution	1324
3. Open problems	1326
4. Nature of ordering in the infinite-ranged spin glass	1327
4.1. Thermodynamic pure states	1328
4.2. Overlap functions and distributions	1329
4.3. Non-self-averaging	1330
4.4. Ultrametricity	1332

5. Detection of many states in spin glasses	1332
5.1. Chaotic size dependence in the SK model	1333
5.2. Chaotic size dependence in the EA model	1335
6. Metastates	1338
6.1. Motivation and mathematical construction	1338
6.2. Physical meaning and significance	1339
7. Can a mean-field scenario hold in short-ranged models?	1340
7.1. Translation ergodicity	1340
7.2. The standard SK picture	1342
7.3. The non-standard SK picture	1343
7.4. What non-self-averaging really means	1345
7.5. Differences between the standard and non-standard pictures	1345
7.6. Invariance of the metastate	1345
8. Structure of the low- T spin glass phase	1347
8.1. Remaining possibilities	1347
8.2. The problem with $P(q)$	1348
9. Interfaces	1349
9.1. Space-filling interfaces and observable states	1350
9.2. Invisible states	1350
9.3. Relation between interfaces and pure states	1351
9.4. Low-lying excited states	1351
10. Summary and discussion	1352
10.1. Summary	1352
10.2. Comparison with other work	1354
10.3. Effects of a magnetic field	1356
11. Other topics	1356
11.1. Metastable states	1356
11.2. The statistical mechanics of homogeneous versus disordered systems	1359
Acknowledgments	1361
References	1361

1. Introduction

Despite decades of intensive investigation, the statistical mechanics of systems with both quenched disorder and frustration remains an open problem. Among such systems, the spin glass is arguably the prototype, and inarguably the most studied.

Spin glasses are systems in which competition between ferromagnetic and antiferromagnetic interactions among localized magnetic moments (or more colloquially, ‘spins’) leads to a magnetically disordered state (figure 1). The prime example of a metallic spin glass is a dilute magnetic alloy, in which a magnetic impurity (typically Fe or Mn) is randomly diluted within a non-magnetic metallic host, typically a noble metal. The competition between ferromagnetic and antiferromagnetic interactions arises in these systems from the RKKY interactions [1–3] between the localized spins, mediated by the conduction electrons.

But many other types of spin glass, with different microscopic mechanisms for their ‘spin glass-like’ behaviour, exist. Certain insulators, in which low concentrations of magnetic impurities are randomly substituted for non-magnetic atoms, also display spin glass behaviour. A well-known example [4] is $\text{Eu}_x\text{Sr}_{1-x}\text{S}$, with x roughly between 0.1 and 0.5, where the competition arises largely from nearest-neighbour ferromagnetic and next-nearest-neighbour antiferromagnetic interactions. There are many other materials that exhibit spin glass




Figure 1. A rough sketch of the classical ground states of: (a) a crystal; (b) a glass; (c) a ferromagnet; (d) an Ising spin glass. In (a) and (b) the dots represent atoms; in (c) and (d) the arrows represent localized magnetic moments. (In the case of (b), it is more accurate to describe the configuration as a frozen metastable state.)

behaviour, both metallic and insulating, crystalline and amorphous. They can display Ising, planar or Heisenberg behaviour, and come in both classical and quantum varieties. In this review we consider only classical spin glasses [5].

What are the main experimental features of spin glasses? One is the presence of a cusp in the low-field ac susceptibility (figure 2), as first observed in AuFe alloys by Cannella and Mydosh [6]. This cusp becomes progressively rounded as the external magnetic field increases [7]. The specific heat, however, rather than showing a similar singularity, typically displays a broad maximum (figure 3) at temperatures somewhat higher than the ‘freezing temperature’ T_f defined via the susceptibility peak (see e.g. [8]).

Probes of the low-temperature magnetic structure using neutron scattering, Mössbauer studies, NMR and other techniques confirm the absence of long-range spatial order coupled with the presence of temporal order in so far as the spin orientations appear to be frozen on the timescale of the experiment. An extensive description of these and related experiments is presented in the review article by Binder and Young [9].

Spin glasses are also characterized by a host of irreversible and non-equilibrium behaviours, including remanence, hysteresis, anomalously slow relaxation, ageing and related phenomena. Because this review will focus on static equilibrium behaviour, these topics will not be treated here, but it is important to note that explaining these phenomena is essential to any deeper understanding of spin glasses. For reviews, see [9–12].

The quest to attain a theoretical understanding of spin glasses has followed a tortuous path, and to this day many of the most basic and fundamental issues remain unresolved. An extensive discussion of theoretical ideas can be found in a number of reviews [9–11, 13–16]. The good news is that research into spin glasses has uncovered a variety of novel and sometimes stunning concepts; the bad news is that it is not clear how many of these apply to real spin glasses themselves. In this review we will explore some of these issues, in particular the nature of ordering and broken symmetry in the putative spin glass phase.




Figure 2. Low-field magnetic susceptibility $\chi(T)$ in AuFe alloys at varying concentrations of iron impurity (from [6]).




Figure 3. Magnetic specific heat C_m of CuMn at 1.2% manganese impurity. The arrow indicates the freezing temperature T_f as discussed in the text (from [8]).

As a first step, one needs to capture mathematically the absence of orientational spin ordering in space with the presence of frozenness, or order in time. This was achieved early on by Edwards and Anderson (EA) [17], who noted that a low-temperature pure phase of spin glasses should be characterized by a vanishing magnetization per spin

$$M = \lim_{L \rightarrow \infty} \frac{1}{|\Lambda_L|} \sum_{x \in \Lambda_L} \langle \sigma_x \rangle \quad (1)$$

accompanied by a non-vanishing ‘EA order parameter’ (as it is now called)

$$q_{\text{EA}} = \lim_{L \rightarrow \infty} \frac{1}{|\Lambda_L|} \sum_{x \in \Lambda_L} \langle \sigma_x \rangle^2 \quad (2)$$

where σ_x is the spin at site x , Λ_L is a cube of side L centred at the origin, and $\langle \cdot \rangle$ denotes a thermal average.

However, it was later discovered that, while the EA order parameter q_{EA} plays a central role in describing the spin glass phase, it is insufficient to describe the low-temperature ordering—at least in a mean-field version of the problem. In the following sections we will explain this statement, explore the relationship between the mean-field spin glass problem and its short-ranged version, and discuss some new and general insights and tools that may turn out to be useful in unravelling the complexities of the statistical mechanics of inhomogeneous systems.

2. A brief history of early theoretical developments

Questions regarding spin glass behaviour fall naturally into two classes: the first pertains to properties of a system in thermal equilibrium, and the second to those related to non-equilibrium dynamics. It is still not certain whether spin glasses possess non-trivial equilibrium properties, but they surely do possess non-equilibrium ones. As already noted, throughout this review we will focus primarily on spin glasses in thermal equilibrium.

Perhaps the most fundamental question that can be asked is whether there exists such a thing as a true, thermodynamic spin glass phase; that is, is there a sharp phase transition from the high-temperature paramagnetic state to a low-temperature spin glass state in zero external magnetic field? Of course, this is presumably the simplest transition that could in principle occur; one could also ask about field-induced transitions, ferromagnetic to spin glass transitions, and others. But given that all of these questions remain open, we will confine our attention here to the simplest of these. And even here, although experimental and numerical studies have tended to favour an affirmative answer, the issue is by no means resolved.

2.1. The EA Hamiltonian

In order to proceed, we need a specific model to study. The majority of theoretical investigations begin with a Hamiltonian proposed by EA [17]:

$$\mathcal{H} = - \sum_{\langle x,y \rangle} J_{xy} \sigma_x \sigma_y - h \sum_x \sigma_x, \quad (3)$$

where (to keep things as general as possible) x is a site in a d -dimensional cubic lattice, σ_x is the spin at site x , the spin couplings J_{xy} are independent, identically distributed random variables, h is an external magnetic field, and the first sum is over nearest-neighbour sites only. We will usually take the spins to be Ising variables, i.e. $\sigma_x = \pm 1$. Throughout most of the paper we will choose $h = 0$ and the spin couplings J_{xy} to be symmetrically distributed about zero; as a result, the EA Hamiltonian in equation (3) has global spin inversion symmetry.

Popular choices for the distribution of the couplings J_{xy} are bimodal and Gaussian. Most of what we discuss below will be independent of which of these is chosen, but for specificity (and to avoid accidental degeneracies when discussing ground states) we will choose the couplings from a Gaussian distribution with mean zero and variance one. It is important to note that the disorder is *quenched*: once chosen, the couplings are fixed for all time. We denote by \mathcal{J} a particular realization of the couplings, corresponding physically to a specific spin glass sample. Proper averaging over quenched disorder is done on extensive quantities only [18], i.e. at the level of log Z rather than Z , where Z is the partition function.

Of course, the EA Hamiltonian looks nothing like a faithful microscopic description of spin interactions in a dilute magnetic alloy, or an insulator like $\text{Eu}_x\text{Sr}_{1-x}\text{S}$ —and because the statistical mechanics of the EA Hamiltonian remains to be worked out, a direct comparison with experiment remains elusive. However, a central assertion of [17] is that the essential physics of spin glasses is the competition between quenched ferromagnetic and antiferromagnetic interactions, regardless of microscopic details; and so the EA Hamiltonian remains the usual launching point for theoretical analyses of real spin glasses.

2.1.1. Frustration. A striking feature of random-bond models like the EA spin glass is the presence of *frustration*, in this case meaning the inability of any spin configuration to satisfy all couplings simultaneously. It is easily verified that, in any dimension larger than one, all of the spins along any closed circuit \mathcal{C} in the edge lattice cannot be simultaneously satisfied if

$$\prod_{(x,y) \in \mathcal{C}} J_{xy} < 0. \quad (4)$$

The definition of frustration given above was first suggested by Toulouse [19]. A different formulation due to Anderson [20] received less notice when it was first proposed, but its underlying ideas may prove useful in more recent spin glass research. The basic notion is that frustration manifests itself as free energy fluctuations (e.g. with a change in boundary conditions from periodic to antiperiodic) that scale as the *square root* of the surface area of a typical sample.

Hence the spin glass is characterized by both *quenched disorder* and *frustration*. The presence of frustration, leading to a complicated geometry of entangled frustration contours, suggests the possibility that spin glasses, in at least some dimensions, may possess multiple infinite-volume ground or pure states unrelated by any simple symmetry transformation. We will return to this question later. We note here, though, that there exists at least one (unrealistic) spin glass model where the number of ground states can be computed in all finite dimensions. This is the *highly disordered model* of the authors [21, 22] (see also [23]) in which the coupling magnitudes scale non-linearly with the volume (and so are no longer distributed independently of the volume, although they remain independent and identically distributed for each volume). It can be shown [21, 22] that there exists a *transition* in ground state multiplicity in this model: below eight dimensions, it has only a single (globally spin reversed) pair of ground states, while above eight it has uncountably many ground state pairs. Interestingly, the high-dimensional ground state multiplicity can be shown to be *unaffected* by the presence of frustration, although frustration still plays an interesting role: it leads to the appearance of *chaotic size dependence*, to be discussed in section 5.

2.2. Mean-field theory, the Sherrington–Kirkpatrick model and the Parisi solution

Mean-field models often provide a useful first step towards understanding the low-temperature phase of a condensed matter system; in the case of spin glasses, the usual procedure seems to have taken a particularly interesting twist. The mean-field theory of spin glasses turns out to be far more intricate than those of most homogeneous systems, and as a result several different approaches have been tried. Also noteworthy are even simpler, soluble spin-glass-like models, in particular the random energy model of Derrida [24]. However, here we will confine ourselves to a discussion of the Sherrington–Kirkpatrick (SK) model [25], an infinite-ranged version of the EA model in which mean-field theory is presumably exact.

The SK Hamiltonian for a system of N spins is (as usual, we take external field to be zero)

$$\mathcal{H}_N = -(1/\sqrt{N}) \sum_{1 \leq i < j \leq N} J_{ij} \sigma_i \sigma_j \quad (5)$$

where we again take the spins to be Ising and choose the (independent, identically distributed) couplings J_{ij} from a Gaussian distribution with zero mean and variance one; the latter rescaling ensures a sensible thermodynamic limit for free energy per spin and other thermodynamic quantities.

It was shown in [25] that this model has a sharp phase transition at $T_c = 1$. At this temperature the static susceptibility has a cusp—but so does the specific heat.

SK solved for the low-temperature spin glass phase, using the EA order parameter to describe the broken symmetry. However, their solution was unstable [25]; in particular, the entropy was found to become negative at sufficiently low temperature.

The following four years saw intensive efforts to solve for the low-temperature phase of the SK model. Of particular note is the direct mean-field approach of Thouless *et al* [26], who pointed out the necessity of including the Onsager reaction field term, and the paper by de Almeida and Thouless [27], who studied the stability of the SK solution in the $T-h$ plane and calculated the boundary between the regions where a single (i.e. paramagnetic) phase is stable and the region where the low-temperature phase resides. One important question that remains open to this day is whether such an ‘AT line’ exists for more realistic models (see section 10.3 for further discussion).

We will mainly focus, however, on what is today believed to be the correct solution for the low-temperature phase of the SK model. In a series of papers, Parisi and collaborators [28–31] proposed, and worked out the consequences of, an extraordinary *ansatz* for the nature of this phase. Following the mathematical procedures underlying the solution, it came to be known as *replica symmetry breaking* (RSB).

We will not review those mathematical procedures here; they are worked out in detail in several review articles and books (see e.g. [9–11, 13, 14, 16]). We will also omit discussion of important related developments, such as the dynamical interpretation of Sompolinsky and Zippelius [32–34]. We will concern ourselves instead with both the physical and mathematical interpretations of the Parisi solution, and the type of ordering that it implies.

These interpretations took several years to work out, culminating in the work of Mézard *et al* [30, 31] that introduced the ideas of overlaps, non-self-averaging and ultrametricity as a means of understanding the type of order implied by the Parisi solution. These terms, and their relevance for the Parisi solution, will be described in section 4. For now, we simply note that the solution of the infinite-ranged SK model generated tremendous excitement; as described by Rammal *et al* [35], it displayed a new type of broken symmetry ‘radically different from all previously known’. This is not an overstatement.

The starting point is the observation that the low-temperature phase consists not of a single spin-reversed pair of states, but rather of ‘infinitely many pure thermodynamic states’ [29], not related by any simple symmetry transformations. This possibility had already been foreshadowed by the Thouless–Anderson–Palmer approach [26], whose mean-field equations were known to have many solutions (not necessarily all free energy minima, except at zero temperature [36]). The existence of many states meant that the correct order parameter needed to reflect their presence, and to describe the relations among them. The single EA order parameter was therefore insufficient to describe the low-temperature phase (although it retained an important role, as we will see); instead one needed an order parameter *function*.

Before we describe these ideas in more detail, we will first step back and consider the basic outline of the problem that interests us. In particular, we need to ask: what is it that we want to know? What are the fundamental open questions? And how do they tie in with the broader areas of condensed matter physics and statistical mechanics? These questions will be considered in the following section.

3. Open problems

In this review we concern ourselves with perhaps the most basic questions that can be asked: is there a true spin glass phase, and if so, what is its nature?

For the infinite-ranged SK model, these questions appear largely resolved, though some open questions remain. For the EA model (as a representative of more general ‘realistic’ models, i.e. finite-dimensional and non-infinite-ranged), the primary question of whether a thermodynamic phase transition exists remains open. There is suggestive analytical [37, 38] and numerical [9, 39–41] evidence that a phase transition to a broken spin-flip symmetric phase is present in three-dimensional and, even more likely, in four-dimensional Ising spin glasses. However, no one has yet been able to prove or disprove the existence of a phase transition, and the issue remains unsettled [42].

Of course, existence of a phase transition does not necessarily imply more than a single low-temperature phase; one could, for example, have a transition above which correlations decay exponentially and below which they decay as a power law, with $q_{\text{EA}} = 0$ at all non-zero temperatures. However, most numerical simulations and theoretical pictures that point to a low-temperature spin glass phase suggest broken spin-flip symmetry. We are therefore led to:

Open question 1. Does the EA Ising model have an equilibrium phase transition above some lower critical dimension d_c ; and if so, does the low-temperature phase have broken spin-flip symmetry?

If the answer to this question turns out to be no, then subsequent research will need to focus on dynamical behaviour, which—as in ordinary glasses—presents a range of difficult and important problems. However, given the reasonable possibility that there is indeed a sharp phase transition, it is worthwhile to ask:

If there is a phase transition for the EA Ising model at some $d > d_c$, what is the nature of the ordering of the low-temperature phase?

Because of the open-ended nature of this question, it will not be assigned a number; instead, we will break it down into several parts. The remaining questions assume that there is an equilibrium phase transition critical temperature $T_c > 0$, below which there is broken spin-flip symmetry (equivalently, a phase with $q_{\text{EA}} > 0$), but we make no assumptions as to whether d_c is less than or equal to three.

Open question 2. What is the number of equilibrium pure state pairs (at non-zero temperature) and ground state pairs (at zero temperature) in the spin glass phase?

We have seen that the mean-field RSB picture assumes infinitely many such pairs. A competing picture, known as the droplet/scaling picture, due to Macmillan, Bray and Moore, Fisher and Huse, and others [43–49], asserts that there is only a single pair of pure/ground states in the spin glass phase in *any* finite dimension.

Because of the importance of this picture, we discuss it briefly here. ‘Domain wall renormalization group’ studies [43, 44] led to a scaling *ansatz* [43, 44, 46] that in turn led to the development of a corresponding physical droplet picture [46–49] for spin glasses. In this picture, thermodynamic and dynamic properties 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 thermally activated clusters form a dilute gas and can be considered as non-interacting two-level systems. The resulting two-state picture (in which there is no non-trivial RSB) is therefore significantly different from the mean-field picture arising from the SK model.

So do spin glasses in finite dimensions have many equilibrium states or a single pair? Except in the highly disordered model [21, 22], the answer is not known. In one dimension

(where there is no internal frustration), there is only a single pair of ground states, and a single paramagnetic phase at all non-zero temperatures. In an infinite number of dimensions, there presumably would be an infinite number of pure state pairs for $T < T_c$. Recent numerical experiments [50–52] *seem* to indicate a single pair of ground states in two dimensions (where it is believed that $T_c = 0$), but given that lattice sizes studied are still not very large, the question is not completely settled. Recent rigorous work by the authors [53, 54] has led to a partial result that supports the notion that only a single pair of ground states occurs in two dimensions. In three dimensions numerical simulations give conflicting results [55, 56].

While the mean-field-like RSB many-state picture and the two-state droplet/scaling picture have historically been the main competing pictures, there are others as well. At least one of these will be discussed later. One often sees in the literature an unspoken assumption that the presence of many states is synonymous with RSB, and similarly that the presence of only a single pair is equivalent to droplet/scaling. We emphasize, however, that while these are *necessary* requirements for each picture, respectively, they are not sufficient: each has considerable additional structure (which, in the mean-field RSB case, will be discussed in the next section). This then leads to our next question:

Question 3. If there do exist infinitely many equilibrium states in some dimensions and at some temperatures, are they organized according to the mean-field RSB picture?

Treatment of this question is the main theme of the remainder of this review. A series of both rigorous and heuristic results, due to the authors, has largely answered this question in the negative, and it is therefore not listed as open. (A complete discussion is given in section 7.) However, there are remaining questions, such as:

(Semi-)open question 4. What are the remaining possibilities for the number and organization of equilibrium states in the low-temperature spin glass phase? This question is examined in section 8.1.

In discussing this, we will not consider every logical alternative to the mean-field picture, but rather what we consider to be the most likely remaining scenarios for the low-temperature phase of finite-dimensional spin glasses.

The discussion so far has considered only equilibrium pure or ground states, with a view towards determining the nature of broken symmetry in realistic spin glasses. However, a more general discussion of thermodynamics and dynamics, particularly with a view towards explaining experimental observations, needs to include questions about other types of states, such as:

Question 5. How are energetically low-lying excitations above the ground state(s) characterized? (Section 9.4.)

Question 6. What can be proven about numbers and overlaps of metastable states? Do they have any connection(s) to thermodynamic pure states? (Section 11.1.)

Although many other important questions remain open, we close here with a question of more general interest than for spin glass physics alone:

Question 7. In what ways do we now understand how the statistical mechanical treatment of frustrated, disordered systems differs in fundamental ways from that of homogeneous systems? (Section 11.2.)

4. Nature of ordering in the infinite-ranged spin glass

We now return to a more detailed discussion of the nature of ordering implied by Parisi's solution of the SK model. As noted in section 2.2, the RSB scheme introduced by Parisi assumes the existence of many equilibrium pure states. Because the notion of pure states has generated some confusion in the literature, we detour to clarify exactly what is meant by this and related terms. The discussion here closely follows appendix A of [57].

4.1. Thermodynamic pure states

The notion of pure states is well defined for short-ranged, finite-dimensional systems, but is less clear for infinite-ranged ones like the SK model. We therefore begin with a discussion of the EA model (in arbitrary $d < \infty$), and then briefly discuss application of these ideas to the SK model.

Consider first $\mathcal{H}_{J,L}$, the EA Hamiltonian (3) restricted to a finite volume of linear extent L . We will always take such a volume, hereafter denoted as Λ_L , to be an L^d cube centred at the origin. In addition, we need to impose boundary conditions, which we will often take to be periodic; other possibilities include antiperiodic, free, fixed (e.g. all spins on the boundary set equal to +1) and so on. Given a specified boundary condition, the finite-volume Gibbs state $\rho_{\mathcal{J},T}^{(L)}$ on Λ_L at temperature T is defined by

$$\rho_{\mathcal{J},T}^{(L)}(\sigma) = Z_{L,T}^{-1} \exp\{-\mathcal{H}_{\mathcal{J},L}(\sigma)/k_B T\}, \quad (6)$$

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

All equilibrium quantities of interest can be computed from $\rho_{\mathcal{J},T}^{(L)}(\sigma)$, which is simply a probability measure: it describes at fixed T the probability of a given spin configuration obeying the specified boundary condition appearing within Λ_L . Such a (well-behaved) probability distribution is completely specified by its moments, which in this case is the set of all correlation functions within Λ_L : $\langle \sigma_{x_1} \cdots \sigma_{x_m} \rangle$ for arbitrary m and arbitrary $x_1, \dots, x_m \in \Lambda_L$.

Consider next the $L \rightarrow \infty$ limit of a sequence of such finite-volume Gibbs states $\rho_{\mathcal{J},T}^{(L)}(\sigma)$, each with a specified boundary condition (which may remain the same or may change with L). Of course, such a sequence may or may not have a limit; existence of a limit would require that every m -spin correlation function, for $m = 1, 2, \dots$, itself possesses a limit [58]. A *thermodynamic* state $\rho_{\mathcal{J},T}$ is therefore an *infinite-volume* Gibbs measure, providing information such as the probability of any finite subset of spins taking on specified values; and of course it determines global properties such as magnetization per spin, energy per spin and so on.

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

$$\rho_{\mathcal{J},T} = \lambda \rho_{\mathcal{J},T}^1 + (1 - \lambda) \rho_{\mathcal{J},T}^2, \quad (7)$$

where $0 < \lambda < 1$ and ρ^1 and ρ^2 are also thermodynamic states (distinct from ρ), then $\rho_{\mathcal{J},T}$ is a *mixed* thermodynamic state or simply, a mixed state. The meaning of the decomposition in equation (7) is easily understood: it simply means that any correlation function computed using $\rho_{\mathcal{J},T}$ can be decomposed as

$$\langle \sigma_{x_1} \cdots \sigma_{x_m} \rangle_{\rho_{\mathcal{J},T}} = \lambda \langle \sigma_{x_1} \cdots \sigma_{x_m} \rangle_{\rho_{\mathcal{J},T}^1} + (1 - \lambda) \langle \sigma_{x_1} \cdots \sigma_{x_m} \rangle_{\rho_{\mathcal{J},T}^2}. \quad (8)$$

A mixed state may or may not be further decomposed into as many as an uncountable infinity of disjoint other states.

We are now ready to define the idea of a thermodynamic pure state. If a distinct thermodynamic state cannot be written as a convex combination of any other thermodynamic states, it is then a thermodynamic *pure* state, or simply a pure state. So the paramagnetic state is a pure state, as are each of the positive and negative magnetization states in the Ising ferromagnet. In that system at $T < T_c$, the Gibbs state produced by a sequence of increasing volumes, with either periodic or free boundary conditions, is a mixed state, decomposable into the positive and negative magnetization states each with $\lambda = 1/2$. On the other hand, a sequence of increasing volumes with all boundary spins fixed at +1 leads to the positive magnetization pure state.

A thermodynamic, or infinite-volume, *ground* state is a pure state at zero temperature consisting of a single spin configuration, with the property that its energy cannot be lowered by flipping any finite subset of spins³.

A pure state ρ_P can be intrinsically characterized by a *clustering property* (see e.g. [58, 59]), which implies that for any fixed x ,

$$\langle \sigma_x \sigma_y \rangle_{\rho_P} - \langle \sigma_x \rangle_{\rho_P} \langle \sigma_y \rangle_{\rho_P} \rightarrow 0, \quad |y| \rightarrow \infty, \quad (9)$$

and similar clustering for higher-order correlations.

Without getting into technical details, it is already clear that a problem exists with extending the concept of a pure (or for that matter thermodynamic) state to the SK model. Such states are defined for a fixed realization \mathcal{J} of *all* the couplings, but in the SK model the couplings all scale to zero as $N \rightarrow \infty$. However, one can still talk in some rough sense about an SK state \mathcal{P} with a ‘modified’ clustering property

$$\langle \sigma_i \sigma_j \rangle_{\rho_P} - \langle \sigma_i \rangle_{\rho_P} \langle \sigma_j \rangle_{\rho_P} \rightarrow 0, \quad N \rightarrow \infty, \quad (10)$$

for any pair i and j [9]. The problem with this is that there may be no measurable (i.e. effective) way to construct such a state (see section 5 below).

For now, we will ignore such difficulties and use the terms ‘pure state’, ‘thermodynamic state’ and so on, for the SK model also, as has been done extensively in the physics literature. It should always be kept in mind, though, that such usage is rough, and any attempt to make the notion more precise runs into serious difficulties (although we will propose a heuristic method for detecting the existence of many states in the SK model in section 5). Note that some of the difficulties discussed here are not present in the Curie–Weiss model of the uniform ferromagnet. Although couplings scale to zero there also, they ‘reinforce’ each other, being non-random, so that $N \rightarrow \infty$ positive and negative magnetization states still make sense. It turns out that this fundamental difference between finite and infinite systems in the *disordered* case will have profound consequences.

4.2. Overlap functions and distributions

As noted in section 2.2, Parisi’s mathematical RSB scheme led to the physical interpretation of many pure states below T_c . All such states at fixed T have vanishing magnetization per spin as N gets large, and all have the same (non-zero) $q_{EA}(T)$ [9, 13], so these alone are insufficient to describe the ordering. What is needed is a way to describe the relations among the different states, and this can be accomplished by means of *overlaps*.

The usual interpretation of the Parisi solution is as follows [9, 13, 30, 31]. For large N , the Gibbs distribution is a mixture of many pure states:

$$\rho_{\mathcal{J},N}(T) \approx \sum_{\alpha} W_{\mathcal{J}}^{\alpha}(T) \rho_{\mathcal{J}}^{\alpha}(T), \quad (11)$$

where $\rho_{\mathcal{J}}^{\alpha}$ is pure state α , $W_{\mathcal{J}}^{\alpha}$ its weight in the decomposition of $\rho_{\mathcal{J},N}(T)$, and the approximate equality sign indicates that the notion of pure state is not precise in this model. Although equation (11) involves in principle an infinite sum, most weights are vanishingly small; only $O(1)$ states have appreciable weights as $N \rightarrow \infty$.

The overlap between state α and β is defined as

$$q_{\alpha\beta} \approx \frac{1}{N} \sum_{i=1}^N \langle \sigma_i \rangle_{\alpha} \langle \sigma_i \rangle_{\beta}, \quad (12)$$

³ The term ‘ground state’ is also sometimes used more broadly to denote measures on multiple ground state configurations. In this review, we avoid this terminology, and use the term ‘ground state’ to refer only to a single, infinite-volume, spin configuration with the properties described above.




Figure 4. Overlap distribution function $P(q)$ for the Curie–Weiss Ising ferromagnet below T_c , or for short-ranged ferromagnets with periodic boundary conditions. In this figure, $M(T)$ is the magnetization per spin and the spikes at $\pm M^2(T)$ are δ -functions.

where $\langle \cdot \rangle_\alpha$ is a thermal average in pure state α , and dependence on \mathcal{J} and T has been suppressed. So $q_{\alpha\beta}$ is a measure of the similarity between states α and β . It is clear that

$$-q_{EA} \leq q_{\alpha\beta} \leq q_{EA} \quad (13)$$

because $q_{EA} = q_{\alpha\alpha}$ and $-q_{EA} = q_{\alpha, -\alpha}$, where $-\alpha$ is the global flip of α (i.e. all odd-spin correlation functions in α and $-\alpha$ have the same magnitude and opposite sign, and all even-spin correlations in the two are equal).

Because there is no spatial structure in the infinite-ranged model, the overlap function does seem to capture the essential relations among the different states. However, it might already be noticed that such a global measure may well miss important information in short-ranged models—assuming that such models also have many pure states.

Quantities referring to individual pure states are problematic, since there is no known procedure for constructing such things in the SK model. However, what is really of interest is the *distribution* of overlaps; that is, if one randomly and independently chooses two pure states from $\rho_{\mathcal{J}, N}(T)$ in equation (11), the probability that their overlap lies between q and $q + dq$ is $P_{\mathcal{J}}(q) dq$, where

$$P_{\mathcal{J}}(q) = \sum_{\alpha} \sum_{\beta} W_{\mathcal{J}}^{\alpha} W_{\mathcal{J}}^{\beta} \delta(q - q_{\alpha\beta}). \quad (14)$$

As before, we suppress the dependence on T and N for ease of notation. $P_{\mathcal{J}}(q)$ is commonly referred to as the *Parisi overlap distribution*.

What does $P_{\mathcal{J}}(q)$ look like? When there is a single pure state with zero magnetization per spin, such as the paramagnet, it is simply a δ -function at $q = 0$. For the Curie–Weiss Ising ferromagnet below T_c (as well as in short-ranged ferromagnets with, say, periodic boundary conditions), there are two pure states: the uniform positive and negative magnetization states, each appearing with weight 1/2 in a pure state decomposition of the type described in equation (11). The resulting overlap distribution is shown in figure 4.

For the SK model, the overlap distribution for fixed \mathcal{J} is non-trivial, due to the presence of many non-symmetry-related pairs of states in the decomposition of the Gibbs distribution $\rho_{\mathcal{J}}$ (cf equation (11)). Now there are several δ -functions of non-trivial weight, distributed symmetrically about zero in the interval $[-q_{EA}, q_{EA}]$, as sketched in figure 5.

4.3. Non-self-averaging

One of the most interesting and peculiar features of the Parisi solution is the *non-self-averaging* of the overlap distribution function. For large N and fixed \mathcal{J} , $P_{\mathcal{J}}(q)$ has the form shown




Figure 5. Sketch of the overlap distribution function $P_{\mathcal{J}}(q)$ for the SK model below T_c .




Figure 6. Sketch of the averaged overlap distribution function $P(q)$ for the SK model below T_c . The spikes at $\pm q_{EA}$ are δ -functions.

in figure 5. What happens if one looks at a \mathcal{J}' different from \mathcal{J} ? Surprisingly, the overlaps (except for the two at $\pm q_{EA}$, which are present for almost every \mathcal{J}) will generally appear at different values of q , and the set of corresponding weights will also differ. This is true no matter how large N becomes. If one then averages $P_{\mathcal{J}}(q)$ over all \mathcal{J} as $N \rightarrow \infty$, the resulting distribution will be supported on all values of q in the interval $[-q_{EA}, q_{EA}]$.

Let $P_N(q) = \overline{P_{\mathcal{J},N}(q)}$, where the overbar indicates a quenched average over coupling realizations \mathcal{J} , and let $P(q) = \lim_{N \rightarrow \infty} P_N(q)$. A sketch of the averaged overlap distribution $P(q)$ is shown in figure 6.

Together $P_{\mathcal{J}}(q)$ and $P(q)$ can be thought of as describing the nature of ordering in the low-temperature phase of the SK model. Instead of $P(q)$, one can study the function

$$x(q) = \int^{-q} P(q') dq' \tag{15}$$

where $x(q)$ is essentially the fraction of states with overlap smaller than q . This function, or more commonly its (monotonic) inverse $q(x)$, is also commonly referred to as the Parisi order parameter. We will focus on $P(q)$ here.

The sample-to-sample variation, even in the thermodynamic limit, implied by the non-self-averaging of the overlap distribution function, may seem somewhat disturbing at first, since it violates our thermodynamic intuition of a ‘typical’ sample. However, it should be remembered that $P_{\mathcal{J}}(q)$ is not directly measurable in the laboratory (though possibly information about it could be obtained indirectly). Measurable thermodynamic quantities, such as free energy, magnetization and so on, remain self-averaging.

Of course, $P(q)$ is measurable through numerical simulations (see e.g. [60]). Problems in the numerical determination of $P(q)$ for short-ranged models will be discussed in section 8.2.




Figure 7. Rough sketch of the hierarchical structure of distances between states. A faithful representation of the actual tree corresponding to the Parisi solution would contain a continuous sequence of branchings.

4.4. Ultrametricity

The discussion in the preceding subsection concerned random choices of pairs of states taken from the Gibbs distribution $\rho_{\mathcal{J},N}(T)$. A striking prediction of the Parisi solution concerns the expected outcome when one independently chooses *triples* of states from $\rho_{\mathcal{J},N}(T)$. The disorder-averaged overlap probability distribution for triples of states, $P(q_1, q_2, q_3) dq_1 dq_2 dq_3$, gives the probability that the spin overlap between one of the three pairs of states lies in $[q_1, q_1 + dq_1]$, the second pair overlap in $[q_2, q_2 + dq_2]$, and the third in $[q_3, q_3 + dq_3]$. An explicit calculation of $P(q_1, q_2, q_3)$ using the Parisi RSB *ansatz* yields [30, 31]

$$\begin{aligned}
 P(q_1, q_2, q_3) = & \frac{1}{2} P(q_1) x(q_1) \delta(q_1 - q_2) \delta(q_1 - q_3) + \frac{1}{2} P(q_1) P(q_2) \theta(q_1 - q_2) \delta(q_2 - q_3) \\
 & + \frac{1}{2} P(q_2) P(q_3) \theta(q_2 - q_3) \delta(q_3 - q_1) + \frac{1}{2} P(q_3) P(q_1) \theta(q_3 - q_1) \delta(q_1 - q_2).
 \end{aligned}
 \tag{16}$$

That is, if one chooses three states independently from $\rho_{\mathcal{J},N}(T)$, then as $N \rightarrow \infty$ there is a probability 1/4 that all overlaps are equal, probability 1/4 that $q_1 = q_2 < q_3$, probability 1/4 that $q_2 = q_3 < q_1$, and probability 1/4 that $q_3 = q_1 < q_2$. That is, if $d_{\alpha\beta} = q_{EA} - q_{\alpha\beta}$ is defined as the ‘distance’ between states α and β , then the distances among any three states chosen randomly from ρ form the sides of an equilateral (probability 1/4) or acute isosceles (probability 3/4) triangle. A distance metric with these properties is called an ‘ultrametric’; a detailed discussion is given in [35]. So there are strong correlations among the states in the SK model, corresponding to a tree-like or hierarchical structure among their overlaps [61]. This is illustrated in figure 7.

We have seen that the novel nature of the broken symmetry and ordering implied by the Parisi solution of the SK model is significantly different from anything observed in the low-temperature phases of more familiar, homogeneous condensed matter systems. It is therefore not surprising that the Parisi solution, once it was understood, generated considerable excitement. Because mean-field theory has almost always provided the correct physics of the broken symmetry, ordering and low-temperature properties of more realistic models (typically providing poor quantitative results only close to T_c), it was natural to suppose that the RSB mean-field picture should similarly provide a correct description of the EA and other short-ranged spin glass models. In the next few sections we will consider this issue.

5. Detection of many states in spin glasses

In describing the order parameters and broken symmetries of the SK spin glass phase, our starting point was the presumed existence of ‘infinitely many pure thermodynamic states’ [29]. However, in section 4.1 we cautioned about potential difficulties in applying these notions to

the SK model. Are the objections raised there merely pedantic, or do they carry substantial physical significance? If the latter, how do we understand the physical meaning of the Parisi solution, and how can we apply it to the EA and other short-ranged models, where these concepts are better defined?

A second difficulty, also alluded to in section 4.1, is the question of how one could actually construct a pure state in either the SK *or* the EA model. In homogeneous systems like the uniform ferromagnet, this presents no difficulties: to select the positively magnetized pure state, for example, one could introduce either a bulk term (a positive, homogeneous external field) whose magnitude goes to zero, or a surface term (fixed boundary condition consisting of all plus spins), with volume going to infinity. However, even in the EA model, where the existence of well-defined pure states can be proved, there is no known \mathcal{J} -independent general procedure for constructing a single pure state if many are present. This difficulty has also been posed in the context of broken ergodicity [62, 63].

In approaching these questions, it is worthwhile to ask: how can one even know whether there *are* multiple pure state pairs? Just as, in section 4.2, one forgoes examination of individual pure states in the SK model for statistical information regarding their overlaps, we will similarly abandon the idea of attempting to construct individual pure states for a broader study of how the properties of the system are affected if many pure states exist. To accomplish this we need to ask whether one can construct a simple and unambiguous procedure both for detecting the existence of multiple states and for studying their properties. In doing so, we will arrive at a deeper understanding of the meaning and consequences of the existence of many states in disordered materials, and we will simultaneously achieve a clearer, and deeper, understanding of the meaning and significance of the Parisi solution for the SK model. The first step is showing the connection between the existence of multiple states and the presence of what we have called *chaotic size dependence* [64].

5.1. Chaotic size dependence in the SK model

5.1.1. States. We noted in section 4.1 that a thermodynamic state, pure or mixed, can be constructed as the infinite-volume limit of a sequence of finite-volume Gibbs distributions. In the Curie–Weiss model, and again ignoring various technical complications in definitions of pure states, one can study the effects of adding a spin to an N -spin system, which simultaneously requires the introduction of N new couplings. (In spin glasses, this approach has been termed the ‘cavity method’ [13].) It is clear that addition of a single spin will not substantially alter a fixed correlation function, whether single- or multi-spin. One can therefore sensibly conceive of a limiting thermodynamic state (in this case a mixture of positive and negative uniformly magnetized states, each with probability 1/2).

In the SK model, though, the situation is radically different. As $N \rightarrow \infty$, any specified correlation function should not depend on any finite set of couplings, because the overall magnitude of any such finite set scales to zero and each of the new couplings accompanying each additional spin is chosen independently of the previous ones. What this suggests is that (below T_c) any specified correlation function will not settle down to a limit as $N \rightarrow \infty$; equivalently, the cavity method does not result in a limiting thermodynamic state. This heuristic argument introduces us to the pervasive presence of *chaotic size dependence* in spin glasses. For recent results on this issue, see [65].

If such a straightforward procedure does not result in a thermodynamic state, then does the notion of many pure states in the SK model make sense? In one way, at least, the answer is yes, even though the meaning of an individual pure state is still imprecise. In an operational sense, one could in principle keep a record of the values of a finite set of even correlation functions

(with, say, free boundary conditions) as N grows (at fixed T). If these values approach a limit, one can with some justification describe the low-temperature phase as consisting of a single pair of states (because each finite-spin state would simply be an equal mixture of the two). If, however, the correlation functions persist in changing their values as N grows, then one can infer that there must exist many such pairs—presumably infinitely many, if there is no repetitive pattern.

Now it is true, by compactness arguments, that in the latter case there must exist (infinitely many) subsequences that result in distinct limiting states. But it is also almost certainly true that there is no *measurable* way of doing this, i.e. no procedure for selecting subsequences that is independent of a specific coupling realization \mathcal{J} . The consequence is that there is *no finite procedure* for selecting convergent subsequences, and therefore generating thermodynamic states.

However, while generating individual states seems to be problematic, it is nonetheless possible to devise simple procedures for measuring *statistical* properties of such states. So the focus should not be on individual states, as in the conventional statistical mechanics of ordered systems, but instead on some larger construct that measures such statistical properties. This will be further discussed in section 6.

5.1.2. Overlaps. States represent local quantities (i.e. correlation functions), and so depend on details of the coupling realization \mathcal{J} . *Global* quantities, on the other hand, such as free energy per spin, should depend only on general parameters of the coupling *distribution* (such as its mean and variance), and therefore should have (the same) thermodynamic limit for almost every \mathcal{J} [66–68]. (From here on, we use the term ‘almost every’, abbreviated ‘a.e.’, in the strict probabilistic sense: that is, a result holding ‘for a.e. \mathcal{J} ’ means it holds for every coupling realization except for a set of measure zero in the space of all coupling realizations.)

What about a quantity like $P_{\mathcal{J},N}(q)$? The fact that it is non-self-averaging already indicates that, although also a global measure, it may not have an $N \rightarrow \infty$ limit for a.e. \mathcal{J} (using any coupling-independent procedure, i.e. one that is both measurable and finite). This intuition was made rigorous in [64] as the following result:

Theorem 1 ([64]). *If $P_{\mathcal{J},N}(q)$ has a limit $P_{\mathcal{J}}(q)$ for a.e. \mathcal{J} , then it is self-averaged; that is, $P_{\mathcal{J}}(q)$ is independent of \mathcal{J} for a.e. \mathcal{J} .*

Proof. To prove this, we study $Y_{\mathcal{J},N}(t)$, the Laplace transform of $P_{\mathcal{J},N}(q)$:

$$Y_{\mathcal{J},N}(t) \equiv \int dq P_{\mathcal{J},N}(q) e^{tq}. \quad (17)$$

Consider two coupling realizations, \mathcal{J} and \mathcal{J}' , that differ in only finitely many couplings. Then because the coupling magnitudes in the SK model scale as $N^{-1/2}$, it follows that

$$Y_{\mathcal{J}',N}(t)/Y_{\mathcal{J},N}(t) = e^{O(N^{-1/2})}, \quad (18)$$

so if $Y_{\mathcal{J}}(t) = \lim_{N \rightarrow \infty} Y_{\mathcal{J},N}(t)$ exists for a.e. \mathcal{J} , $Y_{\mathcal{J}'}(t) = Y_{\mathcal{J}}(t)$ for any \mathcal{J} and \mathcal{J}' differing by a finite number of couplings. But by the Kolmogorov zero–one law [69], it follows that for each fixed t , $Y_{\mathcal{J}}(t)$ is constant for a.e. \mathcal{J} , which in turns implies that, for any q , $P_{\mathcal{J}}(q)$ is independent of \mathcal{J} for a.e. \mathcal{J} . \square

Here chaotic size dependence seems to follow from the scaling of the coupling magnitudes to zero as $N \rightarrow \infty$, with the result that no local state properties depend on a finite set of couplings. While this is part of the story (and, in fact, an important part), it is not the entire story. In short-ranged models like the EA spin glass, coupling magnitudes are fixed independently of

volume, but as we shall see in the next section, chaotic size dependence is present there also, although certain differences now appear between the infinite- and finite-ranged models.

5.2. Chaotic size dependence in the EA model

5.2.1. ‘Observability’ of states. We already noted in the previous section that one important difference between the EA and the SK models is that coupling magnitudes scale to zero in the latter but not the former. While this is also true in infinite-ranged models in homogeneous systems, the *randomness* in the couplings in the disordered case results in significant differences between short-ranged and infinite-ranged models.

As before, consider a volume Λ_L in the EA model. Boundary conditions, such as free, periodic, antiperiodic or fixed, now need to be specified. (In some rough sense, one can think of the SK model as having free ‘boundary’ conditions.) Recall that in the homogeneous Ising ferromagnet below T_c , different pure states can be generated by using sequences of plus versus minus boundary conditions. So one test of multiple states is the *sensitivity* of correlation functions, deep in the interior, upon change of boundary conditions far away.

If there is only a single pure state (e.g. a paramagnet), then any sequence of boundary conditions results in that single limiting state, and there is no chaotic size dependence. If broken spin-flip symmetry exists, i.e. there is a phase with $q_{EA} > 0$, then there must be different sequences of boundary conditions leading to different limiting states. But how can we use this fact to determine whether there is only a single pair of pure states, as in the droplet/scaling picture, or infinitely many pairs, as would be required, for example, by an RSB-like picture?

We start by defining an *observable* state. We do not expect observable physical properties to depend on the microscopic details of the couplings, but instead on macroscopic properties. For example, in a dilute magnetic alloy, one needs to study the system without knowing the microscopic locations of all the magnetic impurity atoms. As a consequence, measurements on a spin glass are necessarily made in a manner independent of the microscopic coupling realization corresponding to that particular sample.

This suggests that physical properties should be associated with states that are mathematically constructed in a *coupling-independent* manner; in particular, by using boundary conditions that do not depend on the couplings [70, 71]. We are therefore led to the following:

Working definition: An *observable state* is a thermodynamic state, pure or mixed, whose existence is manifested through some sequence of *coupling-independent* boundary conditions.

This working definition is not mathematically precise, but will be made so in section 6.2.1. The definition suggests the possible existence of other, ‘invisible’, states; these will be briefly discussed in section 9.2.

Clearly, the many states appearing in the Parisi solution of the SK model are observable. But this then leads to another question: instead of bothering with boundary conditions, why not use the overlap measure $P_{\mathcal{J}}(q)$ to detect many states in the EA model, just as it was successfully used in the SK model? As we will see in section 8.2, $P_{\mathcal{J}}(q)$ turns out in general to be an unreliable indicator of the existence of many states in short-ranged models (whether homogeneous or disordered); it can give the appearance of many states where there is only a single pair, and the appearance of only a single pair—indeed, sometimes even only a single state—when there are infinitely many.

So how can one distinguish between the presence of infinitely many pure states versus only a single pair for the EA model? In the next two subsections, we will show that the answer lies in studying *spin-flip symmetric* boundary conditions.

5.2.2. *Sensitivity to boundary conditions and ‘windows’.* The definition of observable states given in the previous section still seems somewhat impractical, or at least not ‘finite’, since it involves the existence of infinite-volume limits. However, its consequences are indeed practical *and* finite, in particular for numerical experiments. For example, one implication is that the presence of multiple observable states manifests itself as a sensitivity of correlation functions deep in the interior of a finite volume Λ_L to changes in the boundary conditions on $\partial\Lambda_L$.

In order to make this idea precise, consider a fixed volume Λ_{L_0} within which correlation functions are studied, centred inside a much larger volume Λ_L . We will call Λ_{L_0} a ‘window’ [72].

Suppose now that there exists only a single pure state at some temperature T ; for sufficiently high T in either the EA or SK models, this would be the paramagnet. As noted in the previous section, any infinite sequence of boundary conditions will result in this single limiting pure state. (In the above sentence, and throughout the paper, ‘infinite sequence of boundary conditions’ is shorthand for an infinite sequence of finite-volume Gibbs states, as in equation (6), generated by a sequence of volumes with the specified boundary conditions.) As a consequence, if one looks at a fixed window Λ_{L_0} inside a much larger volume Λ_L , any change in the boundary conditions on $\partial\Lambda_L$ will leave the correlation functions inside the window largely unaffected (with any small effects vanishing as the boundaries recede to infinity).

Suppose now that there are two pure states, say the positive and negative magnetization states in the Ising ferromagnet. If one switches from the fixed boundary condition on $\partial\Lambda_L$ with all boundary spins $+1$, to the fixed boundary condition with all boundary spins -1 , correlation functions *everywhere* inside the volume will change (corresponding to a change from the positively magnetized to the negatively magnetized state), no matter how large L becomes.

So the presence of multiple Gibbs states results in sensitivity to boundary conditions. Demonstration of such sensitivity in spin glasses at sufficiently low but non-zero temperature in a given dimension would already be sufficient to answer open question 1. More is needed, however, if one wants to address the issue of whether there is a single pair of pure states or infinitely many.

5.2.3. *Domain walls and free energies.* As noted in section 4.2, an infinite sequence of periodic boundary conditions in the (uniform or random) Ising ferromagnet, above the lower critical dimension and below T_c , does lead to a limiting mixed state $\rho_{\text{per}} = \frac{1}{2}\rho_+ + \frac{1}{2}\rho_-$, where ρ_+ is the positively magnetized and ρ_- the negatively magnetized state. If these are the *only* pure states, then any sequence of antiperiodic boundary conditions (b.c.s) will also have a limit and yield the *same* mixed state. Of course, there will be a relative domain wall between periodic and antiperiodic b.c. states in a fixed volume, but as the volume size increases, the domain wall eventually moves outside of any fixed window. In such a scenario we say the domain wall has ‘deflected to infinity’.

One can prove a similar result in the EA Ising spin glass with periodic, antiperiodic, free or other ‘symmetric’ b.c.s. By symmetric b.c. we mean one for which the resulting Gibbs state ρ_L in any Λ_L is spin-flip invariant; that is, all odd correlations vanish. In a spin glass, if there is only a single pair of pure states ρ and $\bar{\rho}$ that transform into each other under a global spin flip (as in the droplet/scaling picture), then any sequence of symmetric boundary conditions has a limiting thermodynamic state [64], which is simply the mixed state $\rho_{\text{sym}} = \frac{1}{2}\rho + \frac{1}{2}\bar{\rho}$. As a consequence, switching from periodic to antiperiodic b.c.s in a large volume leaves the Gibbs state unaffected inside a window deep in the interior.

What happens if there are many pure state pairs, as would be the case in a mean-field-like picture? In this case, two arbitrarily chosen pure states (not from the same pair) would have relative domain walls that do *not* deflect to infinity. If the free energy cost of such domain walls

does not exceed the free energy difference in a typical Λ_L incurred by switching from, say, periodic to antiperiodic b.c.s, then these different pairs of states should be observable under such a switch. In other words, if many pure state pairs exist, and their relative free energy differences are not too large, then switching from periodic to antiperiodic b.c.s in a typical Λ_L should change the Gibbs state deep in the interior.

How large is ‘not too large’? A trivial upper bound is a domain wall whose free energy scales as L^{d-1} or smaller, since a switch in b.c.s obviously cannot result in a greater free energy change. But in the EA model, a much better bound can be found rigorously. At fixed temperature, if one switches from periodic to antiperiodic b.c.s then in Λ_L , the root mean square free energy difference is bounded from above by $O(L^{(d-1)/2})$, i.e. the square root of the surface area [64, 73]. This rigorous result, which we present here without proof, and which applies to other ‘flip-related’ [64] coupling-independent b.c.s, supports Anderson’s [20] intuitive notion of frustration described in section 2.1.1.

This suggests the following picture: in dimensions with broken spin-flip symmetry, there exists an exponent $\theta(d)$ with $0 < \theta(d) \leq (d-1)/2$ [46, 49] that governs free energy fluctuations upon a switch from periodic to antiperiodic b.c.s. If many pure state pairs exist, and their lowest-lying relative domain wall energies scale as L^θ or less, then the Gibbs state ρ_L in a typical Λ_L will be sensitive throughout the entire volume to changes in coupling-independent b.c.s, irrespective of the size of L . In other words, these many pure state pairs are observable in the sense described in section 5.2.1.

What is expected in a mean-field-like scenario for the EA model? If the many pure state pairs there are not observable, a case could be made that they are unphysical and not of great interest. However, because they *are* seen in the SK model under what can roughly be thought of as (coupling-independent) free b.c.s, one would expect the analogous states—if they exist—in the EA model to be observable also. In fact, RSB theory predicts that domain walls between such states are both ‘space filling’ (to be discussed in section 9) and with energies of $O(1)$ independent of lengthscale [74–77]. They therefore must manifest themselves under a change from, say, periodic to antiperiodic b.c.s [74] in a typical large Λ_L .

5.2.4. Many states and chaotic size dependence. The discussion so far indicates that the existence of many pairs of ‘SK-like’ pure states in the EA model leads to chaotic size dependence: if one takes a sequence of volumes all with, for example, periodic b.c.s (with the volumes chosen deterministically, i.e. independently of the couplings) then there will almost surely be no limiting thermodynamic state. That is, a typical n -spin correlation function, computed in each Λ_L using the corresponding finite-volume Gibbs state, will change continually as $L \rightarrow \infty$, never settling down to a limit.

This leads to an interesting alternative formulation of the problem of existence of many pure state pairs in the EA model. Consider any arbitrary deterministic sequence of volumes with symmetric (periodic, antiperiodic, free or other) b.c.s. *If there exists only a single pair $\alpha, \bar{\alpha}$ of (observable) states, as in the droplet/scaling model, then for fixed \mathcal{J} such a sequence will (with probability 1 over the coupling realizations) possess a limit: the thermodynamic state $\rho = \frac{1}{2}\alpha + \frac{1}{2}\bar{\alpha}$. If instead there are two or more (observable) pure state pairs, then there will be chaotic size dependence: with probability 1, such a sequence will not converge to a limiting thermodynamic state.*

The precise statement underlying these conclusions is the following:

Theorem 2 ([64]). *Given a symmetric coupling distribution, suppose that a sequence of coupling-independent boundary conditions results in a limiting thermodynamic state ρ . If in every volume in the sequence (or an arbitrary subset chosen independently of the couplings),*

the original boundary condition is replaced by one that is flip related, then (1) the new sequence will also have a limit, which (2) will be the same thermodynamic state ρ .

Remark. By ‘flip-related’ b.c.s we mean that for each L , there is some $B_L \subset \partial\Lambda_L$ whose flip transforms one b.c. into the other. So, for example, periodic and antiperiodic b.c.s are flip related; so are two different fixed b.c.s. However, periodic and fixed b.c.s are not flip related.

The proof of the second conclusion is fairly straightforward, and will be useful later, so we present it here. For a proof of the first conclusion, see [64].

Proof of part 2. Consider the correlation function of the spins at the m sites i_1, \dots, i_m . For each L , call this m -point correlation function X_L^1 for b.c. 1 and X_L^2 for b.c. 2, which is flip related to b.c. 1. X_L^1 and X_L^2 are bounded between -1 and $+1$ for all L and all \mathcal{J} . The limits $X^1 = \lim_{L \rightarrow \infty} X_L^1$ and $X^2 = \lim_{L \rightarrow \infty} X_L^2$ exist by assumption, and are functions of \mathcal{J} . Let $X(\mathcal{J}) \equiv X^1(\mathcal{J}) - X^2(\mathcal{J})$.

Consider now the conditional expectation $E_r[\cdot]$, defined as the expectation of a quantity after averaging over all couplings *outside* Λ_r , so $E_r[\cdot]$ depends on the couplings only inside Λ_r . Then

$$E_r[X] = \lim_{L \rightarrow \infty} E_r[X_L] = \lim_{L \rightarrow \infty} (E_r[X_L^1] - E_r[X_L^2]) \quad (19)$$

exists. Because r and m are fixed, and $L \rightarrow \infty$, eventually $L > r$ and $\{i_1, \dots, i_m \subset \Lambda_L\}$ for any r and m . But because $E_r[\cdot]$ averages over the boundary bonds of the cube, and the b.c.s 1 and 2 are flip related, it is easy to see that $E_r[X_L^1] = E_r[X_L^2]$ for large L and hence that $E_r[X] = 0$ for every r . If a bounded random variable X has zero conditional expectation for every r , then $X = 0$ (for a.e. \mathcal{J}). The claimed result follows, because the same argument holds for all correlation functions. \square

Chaotic size dependence in the presence of many states is intuitively plausible. Suppose that there exists infinitely many pure state pairs. For a given Λ_L , some subset of those states will have larger weights than others; in a rough sense, they will be those which best ‘match’ the b.c. for that L . As L varies, the selected states with large weights should vary in some unpredictable fashion. The theorem simply proves this plausible scenario.

There is a numerical consequence of this observation. Just as the presence of many states is in principle detectable numerically by looking at the sensitivity of the state in the deep interior to changes in the b.c. in a *fixed* volume (section 5.2.3), one can also look numerically for chaotic size dependence; that is, study a given set of correlation functions as the volume size changes for a *fixed* b.c. (such as periodic).

It may now seem that chaotic size dependence adds a layer of complexity to the study of thermodynamic states in spin glasses. How can one even talk about pure states when there now seems to be no measurable way to construct them (if there are infinitely many)? A new thermodynamic tool seems to be needed; such a tool will be presented in the next section.

6. Metastates

6.1. Motivation and mathematical construction

If there exist many observable pure states, a sequence of coupling-independent b.c.s will generally *not* converge to a limiting thermodynamic state: there is chaotic size dependence (hereafter denoted CSD). That is, a typical correlation $\langle \sigma_{i_1} \cdots \sigma_{i_n} \rangle$, computed in Λ_L from the finite-volume Gibbs state, will not have a single limit as $L \rightarrow \infty$ but rather many different limits along different subsequences of the L (chosen in a coupling-dependent manner).

Such behaviour in L is analogous to chaotic behaviour in time t along the orbit of a dynamical system. Of course, in each case the behaviour is deterministic but effectively unpredictable, and appears to be a random sampling from some distribution κ on the space of states. In the case of dynamical systems, one can in principle reconstruct κ by keeping a record of the proportion of time the particle spends in each coarse-grained region of state space. Similarly, one can prove [15, 78] that for inhomogeneous systems like spin glasses, a similar distribution exists: roughly speaking, the fraction of the Λ_L in which a given thermodynamic state Γ appears converges, even in the presence of CSD. By saying that a thermodynamic state Γ (which is an infinite-volume quantity) ‘appears’ within a finite volume Λ_L , we mean the following: within a window deep inside the volume, all correlation functions computed using the finite-volume Gibbs state ρ_L are the same as those computed using Γ (with negligibly small deviations). The state Γ can be either pure or mixed, depending on the b.c.s.

Mathematically, a metastate κ is a probability measure on the space of all (fixed \mathcal{J}) thermodynamic states. Of course, the metastate depends on the b.c.s used: we will refer to the metastate constructed from a deterministic sequence of volumes, all with periodic b.c.s, as the ‘periodic b.c. metastate’, and similarly for the antiperiodic b.c. metastate, the free b.c. metastate and so on. One can also construct metastates in which the b.c.s vary with L .

A simple empirical construction of κ would be as follows: consider a ‘microcanonical’ ensemble κ_L , in which each of the finite-volume Gibbs states $\rho^{(1)}, \dots, \rho^{(L)}$ in volumes $\Lambda_1, \dots, \Lambda_L$ has weight L^{-1} . Then $\kappa = \lim_{L \rightarrow \infty} \kappa_L$. The meaning of the limit is that for every well-behaved function $g(\cdot)$ on states,

$$\lim_{L \rightarrow \infty} L^{-1} \sum_{\ell=1}^L g(\rho^{(\ell)}) = \{g(\Gamma)\}_\kappa, \quad (20)$$

where the bracket $\{\cdot\}_\kappa$ denotes the average over κ .

There is an alternative (and earlier) construction due to Aizenman and Wehr [79]. In this construction, one replaces the microcanonical ensemble κ_L by the ensemble of states obtained by varying the couplings *outside* Λ_L . The limit here means that for every well-behaved function F of finitely many couplings and finitely many correlations,

$$\lim_{L' \rightarrow \infty} [F(\mathcal{J}, \rho^{(L')})]_{\text{av}} = \{[F(\mathcal{J}, \Gamma)]_{\kappa(\mathcal{J})}\}_{\text{av}}. \quad (21)$$

Here, $[\cdot]_{\text{av}}$ denotes the average over the quenched coupling distribution. In fact it has not yet been proved that these relatively simple limits, using all the ℓ and the L , are valid. However, it can be proved [15] that there exist deterministic (i.e. \mathcal{J} -independent) subsequences of the ℓ and the L for which limits such as in both equations (20) and (21) exist and yield the *same* $\kappa(\mathcal{J})$.

6.2. Physical meaning and significance

6.2.1. Observable states and thermodynamic chaos. Like a thermodynamic state Γ , the metastate κ is an infinite-volume probability measure. But while Γ is a measure on spin configurations, κ is a measure on the *thermodynamic states* themselves. That is, Γ provides the probability that a given spin configuration appears inside a finite region, while the metastate κ provides the probability that a given pure or mixed state appears inside a typical large volume Λ_L with specified b.c.s. As such, the metastate contains far more information than any single thermodynamic state.

So, instead of treating CSD as a problem and trying to do an ‘end run’ around it, introducing metastates allows us to exploit the vast amount of information contained within CSD; for fixed \mathcal{J} , a metastate allows us to analyse how the finite-volume Gibbs states ρ_L with given b.c.s sample from the available set of thermodynamic states. Although the metastate concept is

equally applicable to situations where CSD does *not* occur, it is most useful as a tool for analysing ‘thermodynamic chaos’.

As always, we are interested in metastates constructed using coupling-independent b.c.s. This allows us to redefine somewhat more precisely the notion of an observable state roughly defined in section 5.2.1:

Definition. *An observable state is a thermodynamic state, pure or mixed, that lies in the support of some coupling-independent metastate.*

6.2.2. Finite versus infinite volumes. Another useful consequence of using metastates is that they enable us to relate the observed behaviour of a system in large but finite volumes to the system’s thermodynamic properties. This relation is relatively straightforward for systems with a 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. Here the metastate approach may be not only useful but necessary.

Occasionally a distinction is drawn between finite- and infinite-volume states (see e.g. [80]), where it is argued that the first is more physical and the second merely mathematical in nature. While we will see below (see also [15, 64, 72, 78, 81]) that the relation between the two may be more subtle than previously realized, we also argue that such a distinction 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.

7. Can a mean-field scenario hold in short-ranged models?

We have now developed the tools we need to analyse whether the type of ordering present in the RSB solution of the SK model can hold in more realistic short-ranged models. The two most-discussed scenarios have been the many-state mean-field picture described above, and the two-state droplet/scaling picture introduced in section 3. We will see below that application of the RSB picture to short-ranged models is not at all straightforward, resulting in considerable confusion in the literature. Before we turn to that subject, however, we need to examine a preliminary question: at fixed dimension and temperature, why cannot it be that the mean-field scenario holds for, say, half of all coupling realizations and droplet/scaling for the other half?

7.1. Translation ergodicity

In fact, such a possibility *cannot* occur: any type of ordering, based on multiplicity of states, whether it is one of the above or something else entirely, *must occur either for every \mathcal{J}* (except for a set of measure zero) *or for none*. The proof of such a statement lies in a straightforward use of the ergodic theorem [82]. Because the ideas used here will be useful later on, we make a small detour to explain them more fully.

As always, we assume a fixed coupling distribution $\nu(\mathcal{J})$, in which the couplings are independent, identically distributed random variables. At some fixed dimension d and temperature T , let $\mathcal{N}(d, T, \mathcal{J})$ denote the number of pure states (one or two or $\dots \infty$). Then it can be proved that $\mathcal{N}(d, T, \mathcal{J})$ is constant almost surely; that is, it is the same for a.e. \mathcal{J} , at fixed T and d . (Of course, $\mathcal{N}(d, T, \mathcal{J})$ can and—if there is a low-temperature spin glass phase—will have some dependence on both d and T . It also clearly can depend on ν , although

that dependence is not explicitly indicated.) In the language of spin glass theory, \mathcal{N} is a self-averaged quantity.

The proof depends on three ingredients: measurability of \mathcal{N} as a function of \mathcal{J} , translation invariance of \mathcal{N} with respect to a uniform shift of the couplings, and translation ergodicity of the underlying disorder distribution ν . A discussion of measurability (in the mathematical sense) would be somewhat technical and will be avoided here, and a precise definition will not be given; but, roughly speaking, it implies, for a function on random variables, that there is an explicit realization-independent procedure for constructing it. A proof that \mathcal{N} is measurable is given in [15].

The concepts of translation invariance and translation ergodicity are relatively straightforward, and we discuss them informally here. Let a be any lattice translation; then \mathcal{J}^a indicates the coupling realization with the locations of all couplings in \mathcal{J} uniformly shifted by a . A translation-invariant function f on \mathcal{J} is one where $f(\mathcal{J}^a) = f(\mathcal{J})$ and a translation-invariant distribution of the \mathcal{J} is defined similarly. Clearly, both \mathcal{N} and the disorder distribution ν are translation invariant.

Translation *ergodicity* of a probability measure, such as ν , is analogous to the more familiar notion of time ergodicity. Consider again a function g of \mathcal{J} , with \mathcal{J} chosen from some distribution μ ; g may or may not be translation invariant. The *distribution* average $E_\mu[g]$ of g is what we normally refer to simply as the average; that is,

$$E_\mu[g] = \int d\mathcal{J} \mu(\mathcal{J}) g(\mathcal{J}). \quad (22)$$

But a *translation* average $E_t[g]$ can also be defined for any realization \mathcal{J} , as simply the spatial average of $g(\mathcal{J})$ under all lattice translations \mathcal{J}^a . Then the distribution μ is translation ergodic if

$$E_\mu[g] = E_t[g], \quad (23)$$

for a.e. \mathcal{J} .

By the ergodic theorem, any (measurable) translation-invariant function of \mathcal{J} chosen from a translation-ergodic distribution is constant almost surely (i.e. it is the same for a.e. realization of \mathcal{J}).

Informally, this is easy to see: by definition, a translation-invariant function is constant with respect to any lattice translation of the realization on which it depends. Suppose that $g(\mathcal{J})$ is translation invariant for a.e. \mathcal{J} , and suppose (for example) that it equals the constant g_1 for half of all realizations and g_2 for the other half, with $g_1 \neq g_2$. Then the distribution average $E_\mu[g] = (1/2)(g_1 + g_2)$, which is not equal to $g(\mathcal{J})$ for *any* \mathcal{J} (outside of a possible set of measure zero). But this violates equation (23), contradicting the supposition that μ is a translation-ergodic distribution.

Returning to the question of the possible variability of the number of pure states with \mathcal{J} , we recall that the couplings are independent, identically distributed random variables. The distribution for such random variables is translation ergodic [83], and so $\nu(\mathcal{J})$ is translation ergodic. So, because \mathcal{N} is a translation-invariant function of \mathcal{J} , which is drawn from the translation-ergodic distribution $\nu(\mathcal{J})$, it follows that \mathcal{N} is the same for a.e. \mathcal{J} (at fixed T and d). Of course, this argument does not tell us the *value* of \mathcal{N} at a given T and d , only that it is constant with respect to \mathcal{J} .

This argument has been presented in some detail not because there is any controversy on this particular question—there is not—but because similar arguments can be used to resolve issues that heretofore *had* been controversial. We will turn to these issues in the following sections. Before doing so, we note that these and similar results allow us to make statements like ‘The number of pure states in a spin glass at low temperatures in three dimensions is x .’

What this statement really means is that, with probability 1, for any particular realization of the couplings on a three-dimensional lattice, there are x pure states at that temperature. (Even a statement such as this, however, should specify whether one is talking about observable states only, but that will be assumed in what follows. We conjecture, but have not proved, that the same number of pure states will be in the support of the a.e. coupling-independent metastate.)

7.2. The standard SK picture

We now turn to an examination of whether a mean-field-like picture can hold in short-ranged spin glasses like the EA model. It may seem initially that the outlines of such a picture should be clear. A typical description is given in [84]:

Hence the Gibbs equilibrium measure decomposes into a mixture of many pure states. This phenomenon was first studied in detail in the mean field theory of spin-glasses, where it received the name of RSB . . . But it can be defined in a straightforward way and easily extended to other systems, by considering an order parameter function, the overlap distribution function. This function measures the probability that two configurations of the system, picked up independently with the Gibbs measure, lie at a given distance from each other . . . RSB is made manifest when this function is nontrivial.

But if there are many pure states, then there must be CSD; and if there is CSD, what does one mean by the equilibrium Gibbs measure? This question was first addressed in [85], where it was shown that the most natural and straightforward interpretation of statements like the one above—what we have called the standard SK picture—*cannot* hold in short-ranged models in any dimension and at any temperature.

The difficult part in studying this problem is that of constructing limiting Gibbs states, given the presence of CSD when there are many states—but the notion of the metastate now makes it easy (or at least easier). In [85], two constructions of overlap distributions were given, but we use metastates here to simplify the discussion. Before proceeding, we will use the description above to construct the standard SK picture.

The extract from [84] quoted above, as applied to the EA model at fixed temperature T , requires a Gibbs equilibrium measure $\rho_{\mathcal{J}}(\sigma)$ which is decomposable into many pure states $\rho_{\mathcal{J}}^{\alpha}(\sigma)$:

$$\rho_{\mathcal{J}}(\sigma) = \sum_{\alpha} W_{\mathcal{J}}^{\alpha} \rho_{\mathcal{J}}^{\alpha}(\sigma). \quad (24)$$

One then considers the overlap distribution function $P_{\mathcal{J}}(q)$, constructed as described above. That is, one chooses σ and σ' from the product distribution $\rho_{\mathcal{J}}(\sigma)\rho_{\mathcal{J}}(\sigma')$, and then the overlap

$$Q = \lim_{L \rightarrow \infty} |\Lambda_L|^{-1} \sum_{x \in \Lambda_L} \sigma_x \sigma'_x \quad (25)$$

has $P_{\mathcal{J}}$ as its probability distribution. Here $|\Lambda_L|$ is the volume of the cube Λ_L .

Given equation (24), there is a non-zero probability that σ and σ' will be chosen from different pure states. If σ is drawn from $\rho_{\mathcal{J}}^{\alpha}$ and σ' from $\rho_{\mathcal{J}}^{\beta}$, then the expression in equation (25) equals its thermal mean,

$$q_{\mathcal{J}}^{\alpha\beta} = \lim_{L \rightarrow \infty} |\Lambda_L|^{-1} \sum_{x \in \Lambda_L} \langle \sigma_x \rangle_{\alpha} \langle \sigma_x \rangle_{\beta}. \quad (26)$$

Thus $P_{\mathcal{J}}$ is given by

$$P_{\mathcal{J}}(q) = \sum_{\alpha, \beta} W_{\mathcal{J}}^{\alpha} W_{\mathcal{J}}^{\beta} \delta(q - q_{\mathcal{J}}^{\alpha\beta}). \quad (27)$$

In the mean-field picture, the $W_{\mathcal{J}}^{\alpha}$ and the $q_{\mathcal{J}}^{\alpha\beta}$ are non-self-averaging quantities, except for $\alpha = \beta$ or its global flip, where $q_{\mathcal{J}}^{\alpha\beta} = \pm q_{\text{EA}}$. As in section 4.3, the average $P(q)$ of $P_{\mathcal{J}}(q)$ over the disorder distribution ν of the couplings is a mixture of two δ -function components at $\pm q_{\text{EA}}$ and a continuous part between them.

The problem, as already noted, is constructing $\rho_{\mathcal{J}}(\sigma)$ given the presence of CSD: simply taking a sequence of cubes with periodic b.c.s, for example, will not work. However, consider the periodic b.c. *metastate* $\kappa_{\mathcal{J}}^{\text{PBC}}$ (in fact, any coupling-independent metastate would do). One can construct a state $\rho_{\mathcal{J}}(\sigma)$ which is the *average* over the metastate:

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

One can also think of this $\rho_{\mathcal{J}}$ as the average thermodynamic state, $N^{-1}(\rho_{\mathcal{J}}^{(L_1)} + \rho_{\mathcal{J}}^{(L_2)} + \dots + \rho_{\mathcal{J}}^{(L_N)})$, in the limit $N \rightarrow \infty$. It can be proved [15, 79] that $\rho_{\mathcal{J}}(\sigma)$ is indeed a Gibbs state.

One can also construct overlaps without constructing Gibbs states at all, as is done numerically. Such a construction (similar to that above) is described in [85], and leads ultimately to the same conclusion.

It is easy to show, given the torus-translation symmetry inherent in periodic b.c.s, that the Gibbs state $\rho_{\mathcal{J}}(\sigma)$ is translation *covariant*; that is, $\rho_{\mathcal{J}^a}(\sigma) = \rho_{\mathcal{J}}(\sigma^{-a})$, or in terms of correlations, $\langle \sigma_x \rangle_{\mathcal{J}^a} = \langle \sigma_{x-a} \rangle_{\mathcal{J}}$. Translation covariance of $\rho_{\mathcal{J}}$ immediately implies, via equations (25)–(27), translation invariance of $P_{\mathcal{J}}$. But, given the translation ergodicity of the underlying disorder distribution ν , *it immediately follows that $P_{\mathcal{J}}(q)$ is self-averaging, and equals its distribution average $P(q)$ for a.e. \mathcal{J}* . The same result can be shown for other coupling-independent b.c.s, where torus-translation symmetry is absent, using methods described in [86]. A simple argument, given in [85], shows further that non-trivial ultrametricity, as in the Parisi solution of the SK model, cannot hold among three arbitrarily chosen states.

So the most natural interpretation of the RSB picture cannot be applied to short-ranged spin glasses. The question then becomes: are there alternative, less straightforward interpretations?

7.3. The non-standard SK picture

The standard SK picture follows a traditional approach in its focus on thermodynamic states. We argued in section 6 that for inhomogeneous systems like spin glasses, such a focus is too restrictive if many pure states are present. Instead, the metastate approach is far better suited as a guide for analysing these kinds of systems.

We have mentioned above (section 5.2.1), and will describe in detail below, the issue of the overlap function $P(q)$ being a poor tool for determining numbers of pure states in short-ranged models. For the moment, however, let us assume that this is not the case, and that numerical simulations on the EA model in three and higher dimensions detect a Parisi-like overlap structure (see e.g. [87]). Does the interpretation given in the cited extract in section 7.2 necessarily follow? The answer is *no*; it was shown in [78] that any evidence for RSB arising from numerical studies of $P(q)$ can correspond to more than one thermodynamic picture. This leads to a reinterpretation not only of what broken replica symmetry might mean in short-ranged systems, but also of what it *does* mean in the SK model.

In numerical computations, overlaps are by necessity computed in finite volumes. Because of CSD, it cannot be assumed that a simple $L \rightarrow \infty$ extrapolation leads to a single thermodynamic mixed state whose decomposition includes all of the observable pure states of the system. But if one is indeed observing a non-trivial (i.e. decomposable into many pure states) mixed state Γ in *one* volume, why should one expect that a similar observation in a




Figure 8. The overlap distribution, at fixed \mathcal{J} , in two different volumes Λ_1 and Λ_2 in the non-standard SK picture.

different volume corresponds to the same Γ ? In fact, the presence of CSD indicates that this *cannot* be the case for all large volumes.

As a mathematical aside, the standard SK picture effectively corresponds to breaking the replica symmetry *after* the thermodynamic limit has been taken. Numerical studies, however, essentially break the replica symmetry *before* taking the thermodynamic limit. Guerra [88] has noted that the order of these limits can be significant. That an interchange of such limits can lead to a new thermodynamic picture of the spin glass phase does not seem to have been appreciated prior to [78].

Based on these considerations, a new, non-standard interpretation of the mean-field RSB picture was introduced in [78] and described in detail in section 7 of [81]. It is a maximal mean-field picture, preserving mean-field theory's main features, although in an unusual way. The most natural description of this non-standard interpretation is in terms of the metastate.

We now summarize, informally, the non-standard SK picture. Formal treatments can be found in [15, 78]; other detailed descriptions appear in [57, 72, 81].

We again consider the PBC metastate, although, as always, almost any other coupling-independent metastate will suffice. In order to resemble the structure of ordering in the SK model as closely as possible, the non-standard SK picture assumes that in each Λ_{L_i} , the finite-volume Gibbs state $\rho_{\mathcal{J}, L_i}$ is well approximated deep in the interior by a mixed thermodynamic state $\Gamma^{(L_i)}$, decomposable into many pure states $\rho_{\alpha_{L_i}}$:

$$\Gamma^{(L_i)} = \sum_{\alpha_{L_i}} W_{\Gamma^{(L_i)}}^{\alpha_{L_i}} \rho_{\alpha_{L_i}}. \quad (29)$$

In this equation, explicit dependence on \mathcal{J} is suppressed. $\Gamma^{(L_i)}$ is a thermodynamic mixed state decomposable into pure states $\rho_{\alpha_{L_i}}$; the index L_i is meant only to indicate the volumes in which $\Gamma^{(L_i)}$ appears.

Each mixed state $\Gamma^{(L_i)}$ is presumed to have a non-trivial overlap distribution

$$P_{\Gamma^{(L_i)}} = \sum_{\alpha_{L_i}, \beta_{L_i}} W_{\Gamma^{(L_i)}}^{\alpha_{L_i}} W_{\Gamma^{(L_i)}}^{\beta_{L_i}} \delta(q - q_{\alpha_{L_i} \beta_{L_i}}) \quad (30)$$

of the form shown in figure 5. Moreover, the distances among any three pure states *within a particular* Γ are assumed to be ultrametric.

As already noted, the presence of CSD requires that $\Gamma^{(L_i)}$ change in some ‘chaotic’ fashion with L_i . Hence, if one computes the overlap distribution in a *particular* Λ_{L_i} , one would see something like figure 5. However, if one looks at two typical volumes of very different sizes, one would see something like figure 8.

7.4. What non-self-averaging really means

In this section we discuss why the non-standard SK picture forces a redefinition of the meaning of non-self-averaging. An important consequence of this redefinition is that most quantities of interest can now be defined for a *single* realization of the disorder: this includes the overlap distribution function. So it is possible after all to focus on a particular sample rather than an ensemble of samples.

The same argument, given in section 7.2, that showed the translation covariance of the thermodynamic state $\rho_{\mathcal{J}}(\sigma)$ applies equally to the metastate $\kappa_{\mathcal{J}}$. As a consequence, the resulting *ensemble* of overlap distributions $P_{\Gamma^{(L_i)}}$ is independent of \mathcal{J} . The dependence of the overlap distribution on Λ_{L_i} (as $\Gamma^{(L_i)}$ varies within the metastate ensemble), no matter how large L_i becomes, is the redefinition of non-self-averaging. It replaces dependence of overlap distributions on \mathcal{J} with dependence on L_i for *fixed* \mathcal{J} .

So, instead of averaging the overlap distribution over \mathcal{J} , the averaging *must now be done over the states Γ within the metastate $\kappa_{\mathcal{J}}$* , all at fixed \mathcal{J} :

$$P(q) = [P_{\Gamma}(q)]_{\kappa_{\mathcal{J}}} = \int P_{\Gamma}(q) \kappa_{\mathcal{J}}(\Gamma) d\Gamma. \quad (31)$$

The $P(q)$ thus obtained for a single \mathcal{J} has the form shown in figure 6, and is the same for a.e. \mathcal{J} .

7.5. Differences between the standard and non-standard pictures

The non-standard SK picture differs from the usual one in several important respects. One is the lack of dependence of overlap distributions on \mathcal{J} , and the replacement of the usual sort of non-self-averaging with that of dependence on states within the metastate. Another important difference is that, in the non-standard SK picture, a continuum (and therefore an uncountable infinity) of pure states *and* their overlaps must be present; therefore, *ultrametricity would not hold in general among any three pure states chosen at fixed \mathcal{J}* , for the same reason ultrametricity breaks down in the standard SK picture [85]. Instead, the pure states at fixed \mathcal{J} are split up into a continuum of families, where each family consists of those pure states occurring in the decomposition of a particular Γ , and only *within* each such family would ultrametricity hold.

This has an important zero-temperature implication, because *there is no difference between the standard and non-standard SK pictures at $T = 0$* . The reason is that, for any finite L and, say, periodic b.c.s, there will be only a single ground state pair $\pm\sigma_0^L$ in Λ_L . (We assume as usual that the coupling distribution is continuous, such as Gaussian, to avoid accidental degeneracies.) *It follows that overlaps of ground states cannot display non-trivial ultrametricity, or any other non-trivial structure.*

We have presented the non-standard SK picture as a replacement for the more standard mean-field picture; if realistic spin glasses display any mean-field features, something like it must occur. However, this leaves open the question of what *does* happen in realistic spin glasses. In particular, can the non-standard SK picture actually occur in short-ranged spin glasses? In the next section we show that the answer, again, is no.

7.6. Invariance of the metastate

In this section we present a theorem about the metastate whose proof is relatively simple but whose implications are powerful and far reaching, not only for spin glasses but for disordered and inhomogeneous systems in general. However, we restrict the discussion here to the EA model in any finite dimension and at any temperature, in zero field and with a symmetric

coupling distribution. We consider again flip-related b.c.s, such as periodic and antiperiodic, or any two fixed b.c.s.

Theorem 3 ([72]). *Consider two metastates constructed along a deterministic sequence of the Λ_L , using two different sequences of flip-related, coupling-independent b.c.s. Then with probability 1, these two metastates are the same.*

The proof, given in [72], is relatively straightforward, and uses two ingredients. The first is that, as proved in [15], along some deterministic subsequence of volumes both the histogram construction of metastates and the Aizenman–Wehr construction [79] result in the same metastate. But because the latter method averages over couplings outside of each volume, it rigorously follows (exactly as in the proof of theorem 2 in section 5.2.4) that two metastates constructed with flip-related b.c.s must be the same.

Despite the straightforward nature of the theorem and proof, it is a striking result, and its consequences for the nature of the spin glass state are immediate. Not only are the periodic and antiperiodic metastates the same; if one were to choose, in a \mathcal{J} -independent manner, two *arbitrary* sequences of periodic and antiperiodic b.c.s, the metastates (with probability 1) would *still* be identical. That is, the metastate, and corresponding overlap distributions constructed from it, at any fixed temperature and in any dimension are highly *insensitive* to b.c.s.

This invariance with respect to different sequences of periodic and antiperiodic b.c.s means that the frequency of appearance of various thermodynamic states $\Gamma^{(L)}$ in finite volumes Λ_L is (with probability 1) *independent* of the choice of b.c.s. Moreover, this same invariance property holds (with probability 1) among any two sequences of *fixed* b.c.s (and the fixed b.c. of choice may even be allowed to vary arbitrarily along any single sequence of volumes)! It follows that, with respect to changes of b.c.s, the metastate is extraordinarily robust.

This insensitivity would be unsurprising if there were only a single thermodynamic state, such as paramagnetic, or a single pair of flip-related states, as in droplet/scaling. 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.

However, while it does not rule out the possibility of many states, theorem 3 does put severe constraints on the form of the metastate and its overlap distribution function. In light of this strong invariance property, any metastate constructed via coupling-independent b.c.s should be able to support only a very simple overlap structure, effectively ruling out the non-standard SK picture.

The non-standard SK picture requires (cf equation (29)) that the Γ appearing in the metastate be of the form $\sum_{\alpha} W_{\Gamma}^{\alpha} \rho_{\Gamma}^{\alpha}$, with the weights W_{Γ}^{α} in each Γ non-zero and unequal. That implies that with periodic b.c.s, say, the fraction of the L_i for which the finite volume Gibbs state in Λ_{L_i} puts (for example) at least 84% of its weight in one pair of pure states is, say, 0.39. But then it must also be the case that with antiperiodic b.c.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; that is, one constructs the periodic b.c. 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 b.c.s, and similarly (but possibly separately) among all pairs of fixed b.c. states.

The only sensible way in which this could happen would be for the selection of states to be insensitive to the choice of b.c.s, i.e. a particular sequence of b.c.s should not prefer any states over any others, so that $\rho_{\mathcal{J}}$, the average over the metastate, would be some sort of uniform mixture of the pure states. However, this *cannot* happen when the Γ are non-trivial mixed states, as in equation (29).

The essential reason for this is that the weights in equation (29) must change with \mathcal{J} , as seen heuristically by the following argument. Choose a particular coupling J_{xy} and consider the transformation $J_{xy} \rightarrow J'_{xy} = J_{xy} + \Delta J$. Then the weight W^α of the pure state α within Γ will transform as

$$W^\alpha \rightarrow W'^\alpha = r_\alpha W^\alpha / \sum_\gamma r_\gamma W^\gamma \quad (32)$$

where

$$r_\alpha = \langle \exp(\beta \Delta J \sigma_x \sigma_y) \rangle_\alpha = \cosh(\beta \Delta J) + \langle \sigma_x \sigma_y \rangle_\alpha \sinh(\beta \Delta J), \quad (33)$$

for each pure state α .

Now if each Γ has a non-trivial decomposition over many pure state pairs, as in equation (29), then the different pure state pairs differ in at least some even correlation functions; that is, they have relative domain walls. Equations (32) and (33) would then lead to a change of the relative weights of the different pure state pairs. As a consequence, even a finite change in a *single* coupling J_{xy} would change a uniform mixture (e.g. in the metastate average) to a non-uniform one. But, of course, the mixture must be the same for a.e. \mathcal{J} .

To summarize, the invariance of the metastate with respect to b.c.s appears to be inconsistent with the transformation properties of the Γ of the form equation (29) with respect to finite changes in \mathcal{J} . This leads to a contradiction, ruling out not only non-standard SK, but *any picture* in which the Γ are a non-trivial mixture of pure states. When combined with our previous elimination of more standard versions of the mean-field picture, it removes the possibility of *any* version of mean-field ordering in short-ranged spin glasses.

The invariance property of the metastate requires that both the pure state structure and the overlap structure of realistic spin glasses should be relatively simple. So what are the remaining, plausible possibilities for the structure of the spin glass phase? In the next section, we will introduce a new picture, guided once again by metastate concepts.

8. Structure of the low- T spin glass phase

8.1. Remaining possibilities

Assuming that spin-flip symmetry is broken below some $T_c(d) > 0$, with d greater than some lower critical dimension d_c , what then are the possible structures that the spin glass phase can assume?

A two-state picture, in which the only observable pure states are global flips of each other, remains completely consistent with all rigorous results described in the preceding section. As noted in section 3, the droplet/scaling scenario is such a two-state picture. At the same time, it should be emphasized that droplet/scaling makes additional assumptions that our results do not address. If it were to be proved that there exists only a single pair of pure states in the spin glass phase in any finite dimension above d_c , this would lend strong support to droplet/scaling, but such a proof alone would not be sufficient to demonstrate its correctness. For fuller discussions of droplet/scaling, see [43–49].

Are there any many-state pictures that *are* consistent with the strong constraints imposed by the invariance of the metastate? There is such a picture, introduced by the authors in [64], and discussed in detail in [15, 57, 72, 78, 81]. It is called the *chaotic pairs* picture, for reasons that will become apparent below.

In a simple two-state picture like droplet/scaling, the overlap distribution function $P_{\mathcal{J}}^L(q)$ in a volume Λ_L will simply approximate a sum of two δ -functions, as in figure 4, but with $\pm M^2(T)$ replaced by $\pm q_{EA}$. In the chaotic pairs picture, each finite-volume Gibbs state

$\rho_{\mathcal{J}}^L$ will still be approximately a mixture of a *single* pair of spin-flip-related pure states, *but now the pure state pair will vary chaotically with L* . Then for each Λ_L , $P_{\mathcal{J}}^L(q)$ will again approximate a sum of two δ -functions at $\pm q_{\text{EA}}$. This picture is fully consistent with metastate invariance.

So the chaotic pairs picture resembles the droplet/scaling picture in finite volumes, but has a very different thermodynamic structure. It is a many-state picture, but unlike the mean-field picture, only a *single* pair of spin-reversed pure states $\rho_{\mathcal{J}}^{\alpha L}, \rho_{\mathcal{J}}^{\bar{\alpha} L}$, appears in a large volume Λ_L with symmetric b.c.s, such as periodic. In other words, for large L , one finds that

$$\rho_{\mathcal{J}}^{(L)} \approx \frac{1}{2}\rho_{\mathcal{J}}^{\alpha L} + \frac{1}{2}\rho_{\mathcal{J}}^{\bar{\alpha} L}. \quad (34)$$

Here, the pure state pair (of the infinitely many present) appearing in finite volume Λ_L depends chaotically on L . Unlike the droplet/scaling picture, this new possibility exhibits CSD with periodic b.c.s. So, like non-standard SK, the periodic b.c. metastate is dispersed over (infinitely) many Γ , but unlike non-standard SK each Γ is a *trivial* mixture of the form $\Gamma = \Gamma^{\alpha} = \frac{1}{2}\rho_{\mathcal{J}}^{\alpha} + \frac{1}{2}\rho_{\mathcal{J}}^{\bar{\alpha}}$. The overlap distribution for each Γ is the same: $P_{\Gamma} = \frac{1}{2}\delta(q - q_{\text{EA}}) + \frac{1}{2}\delta(q + q_{\text{EA}})$. It is interesting to note that the highly disordered spin glass model [21–23], mentioned in section 2.1.1, appears to display just this behaviour in its ground state structure in sufficiently high dimension.

Why does an argument similar to that used to rule out non-standard SK not also rule out chaotic pairs? Because in the chaotic pairs picture, as in droplet/scaling, there are in each Γ only two pure states (depending on Γ in chaotic pairs), each with weight $1/2$. All even correlations are the same in any pair of flip-related pure states, so, by equations (32) and (33), any change in couplings leaves the weights unchanged.

There is an interesting additional piece of information that metastate invariance supplies for many-state pictures like chaotic pairs: the number of pure state pairs (if infinite) must be an *uncountable* infinity. If there were a *countable* infinity, one could not have a uniform distribution consisting of all equal, positive weights within the metastate.

The term ‘chaotic pairs’ was chosen in reference to spin-symmetric b.c.s, such as periodic. If one considers a fixed b.c. metastate, then it would be more appropriate to refer to this picture as ‘chaotic pure states’, because the Gibbs state in a typical large volume Λ_L with fixed b.c.s will be (approximately) a single pure state that varies chaotically with L . But the thermodynamic picture is the same, and just manifests itself slightly differently in different metastates. In particular, the *average* over the metastate $\rho_{\mathcal{J}}$ (cf equation (28)) should be the same for periodic and fixed b.c.s.

We conclude with a brief note about zero temperature, where we observed in section 7.5 that in each Λ_L there is only a single ground state pair. If droplet/scaling holds, then this pair is the same for all large L ; if infinitely many ground state pairs exist, 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 the Γ in each volume exhibiting a non-trivial mixture of pure state pairs as in equation (29), while in chaotic pairs the Γ appearing in any Λ_L consists of a single pure state pair, as in equation (34). So the periodic b.c. metastate in the chaotic pairs picture looks similar at zero and non-zero temperatures, like droplet/scaling but unlike non-standard SK.

8.2. The problem with $P(q)$

In section 7.3, we discussed the usual numerical procedure for constructing overlaps. In both chaotic pairs and droplet/scaling, the overlap distribution function in almost every large volume

(computed in a window far from the boundaries, to avoid non-universal boundary effects) is simply a single pair of δ -functions at $\pm q_{\text{EA}}$, as in figure 4. *So using $P(q)$ in the usual way cannot differentiate between a many-state picture like chaotic pairs and a two-state picture like droplet/scaling.*

One could try an alternative approach, such as looking at the average state over many volumes, as an approximation to studying the average over the metastate $\rho_{\mathcal{J}}(\sigma)$. This average looks very different in the two pictures: in droplet/scaling, it is still a trivial mixture of two states, of the form equation (34), while in chaotic pairs, it is presumably a uniform mixture over uncountably many states. This approach, in effect, takes the thermodynamic limit *before* breaking the replica symmetry, as discussed in section 7.3.

Now there *will* be a difference between overlap functions in droplet/scaling and chaotic pairs. In droplet/scaling, $P_{\mathcal{J}}(q)$ is again a pair of δ -functions at $\pm q_{\text{EA}}$. In chaotic pairs, $P_{\mathcal{J}}(q)$ would now most likely equal $\delta(q)$: it was proven in [89], and will be discussed in section 11.1, that $P_{\mathcal{J}}(q) = \delta(q)$ for the spin overlaps of M -spin-flip-stable metastable states for any finite M . If there are infinitely many ground state pairs, we expect the same to be true for ground states, i.e. for $M = \infty$. But, although there is now a difference between overlap functions in droplet/scaling and chaotic pairs, there is now also *no* difference between overlap functions in chaotic pairs and in the simple paramagnet!

So, although the form of the overlap function can depend on how its computation is done, *the overlap structure in spin glasses must be simple*, regardless of whether there are infinitely many pure states or only a single pair. Moreover, the overlap function cannot distinguish between many states, and one or two states, in an unambiguous manner.

Similar problems with overlap functions, in particular their sensitivity to b.c.s for short-ranged systems, are discussed in [47]. There it is noted that in the three-dimensional random field Ising model at low temperature and weak field magnitude, $P_L(q)$ computed in a given Λ_L will miss one of the relevant pure states. Conversely, in the two-dimensional Ising ferromagnet on a square lattice, at low temperature, and with antiperiodic b.c.s in both directions, $P_L(q)$ gives the appearance of many states where there are only two.

A final, interesting example was suggested to us by van Enter and appears in the appendix of [81]. Consider $P_L(q)$ for an Ising antiferromagnet in two dimensions with periodic b.c.s. For *odd*-sized squares $P_L(q)$ is the same as that of the ferromagnet with *periodic* b.c.s, and for *even*-sized squares it is equivalent to that of the ferromagnet with *antiperiodic* b.c.s. The overlap distribution computed in the full volume therefore oscillates between two different answers, an example of CSD for *overlap* distributions. This illustrates the importance of computing quantities in windows that are much smaller than the system size. In this case, restricting $P(q)$ to such a window gives rise to a well-defined answer: a pair of δ -functions.

9. Interfaces

Our discussion up to now has focused largely on the numbers and organization of pure states in the spin glass phase, if it exists. Even though these are fundamental constructs from the point of view of thermodynamics, and their multiplicity and organization directly affect observable equilibrium [9] and dynamical [90] spin glass properties, they may still seem somewhat abstract. In this section we will draw on recent results that provide a concrete link between the structure of interfaces on the one hand, and numbers of pure states, the relationships among them and their low-lying excited states, on the other. At the same time, we will extend and clarify our earlier distinction (section 5.2.1) between observable and invisible states.

9.1. Space-filling interfaces and observable states

In section 5.2.3, we introduced the notion of ‘space-filling’ interfaces; these play a central role in what follows. The *interface* between two spin configurations σ^L and σ'^L in Λ_L is simply the set of couplings satisfied in one and not the other spin configuration; a domain wall is a connected component of an interface. So the interface between two configurations is the union of all domain walls, and may consist of one or many. We usually envision domain walls as lines or surfaces in the *dual* lattice, cutting those bonds satisfied in one but not the other spin configuration. Domain walls therefore separate regions (in the real lattice) in which spins in σ^L and σ'^L agree from regions where they disagree.

We will confine our remarks here to zero temperature, although it is possible to extend the discussion to non-zero temperatures. We are interested only in interfaces whose linear spatial extent l is $O(L)$ in Λ_L . At zero temperature, these are the *only* interfaces between ground state configurations, if the coupling distribution is continuous; that this is so follows from the same argument proving that domain walls separating infinite spin configurations cannot end in any finite region and cannot have closed loops [53]. However, the restriction against closed loops need not apply to interfaces separating ground and excited states.

If the number of couplings in an interface of linear extent l scales as l^{d_s} , we call d_s the dimension of the interface. A *space-filling* interface is one with $d_s = d$. One of the interesting features of spin glasses is the possibility of ground or pure states separated by space-filling interfaces, unlike in ferromagnets where $d_s < d$ always. For example, the interface ground states generated from Dobrushin b.c.s (all boundary spins on one side of a plane bisecting Λ_L fixed at +1 and the remaining boundary spins -1) have a single domain wall relative to the uniform ground states, with $d_s = d - 1$.

The first result of this section is a theorem first proved in [53] indicating that interfaces separating *observable* ground states are space filling.

Theorem 4 ([53]). *Suppose that in the EA model in some finite dimension d , there exists more than a single pair of thermodynamic ground states in a coupling-independent metastate. Then the interface between two ground states chosen from different pairs must have $d_s = d$.*

Proof. We sketch the proof here for periodic b.c. metastates; the extension to other coupling-independent metastates can be obtained using the procedure presented in [54]. Consider the periodic b.c. metastate $\kappa_{\mathcal{J}}$. By taking two ground state pairs chosen randomly from $\kappa_{\mathcal{J}}$, one obtains a configuration of interfaces. This procedure yields a measure $\mathcal{D}_{\mathcal{J}}$ on domain walls. By integrating out the couplings, one is left with a translation-invariant measure \mathcal{D} on the domain wall configurations themselves. By the translation invariance of \mathcal{D} , any ‘geometrically defined event’, e.g. that a bond belongs to a domain wall, either occurs nowhere or else occurs with strictly positive density. This immediately yields the result. \square

Remark. Theorem 4 extends to pure states at non-zero temperatures.

9.2. Invisible states

Our emphasis in this review has been on observable pure states, which we have identified as the ‘physical’ states; that is, we expect only those states to influence outcomes of laboratory measurements or numerical simulations. Still, it may be interesting to consider also the characteristics of ‘invisible’ states, should they be present.

Recall that an invisible pure state is one that does not lie in the support of any coupling-independent metastate. Invisible pure states can only be generated by sequences of coupling-dependent b.c.s. There is presently no known method for constructing such states, and it is not

clear whether such constructions, if found, would be measurable (section 5). Consequently, it is not presently known whether such states actually exist in spin glasses, or if they do, whether they would be of any interest other than mathematical. Nevertheless, if they should be found to exist, one can still make some predictions about their properties.

From the discussion in section 5.2.3, the interfaces between invisible states would necessarily have energies scaling as l^θ , with $\theta > (d - 1)/2$; any interface with lower energy would almost certainly show up in a coupling-independent metastate. So a ‘high-energy’ interface, satisfying the above condition, would separate invisible pure states from observable ones (or each other), regardless of the interface dimensionality [91].

It also follows, by theorem 4, that any interface with dimension $d_s < d$ separating pure states would necessarily appear only from coupling-dependent b.c.s. Such an interface would necessarily have high energy; if its energy exponent $\theta \leq (d - 1)/2$, it would not correspond to an interface between invisible and observable pure states, but rather to one between excited and ground states. This is discussed further in section 9.4.

9.3. Relation between interfaces and pure states

We conclude from the previous discussion that interfaces separating observable ground (or pure) states are both space filling and have energies scaling as l^θ , where $0 \leq \theta \leq (d - 1)/2$. In [57] we proved the converse, namely that the existence of such interfaces is a sufficient condition for the existence of more than one thermodynamic pure state pair. Specifically, the presence of space-filling interfaces is already sufficient for the existence of multiple pure state pairs; the energy condition is necessary for the pure states to be observable.

As noted in section 5.2.3, RSB theory predicts that interfaces between ground states in Λ_L (say, under a switch from periodic to antiperiodic b.c.s) are both space filling and have energies of $O(1)$ independent of lengthscale [74–77]; that is, they have exponent $\theta = 0$. By the above theorem, the space-filling property requires the existence of multiple thermodynamic pure states, and the $O(1)$ energy property implies that typical large Λ_L with, say, periodic b.c.s, would exhibit thermodynamic states that are non-trivial mixtures of different pure state pairs (see also [72]). But in section 7.6 it was shown that such states cannot appear in the EA model. It follows that space-filling interfaces with $O(1)$ energy are ruled out.

9.4. Low-lying excited states

The discussion of the preceding two sections was restricted to pure or ground states. One can construct other types of metastates, such as ‘excitation metastates’ [54], whose support includes both ground and low-lying excited states. One example of an excitation metastate is the ‘uniform perturbation metastate’ [92], which we now describe.

Once again, consider a deterministic sequence Λ_L of cubes with periodic b.c.s; at zero temperature, each such cube has a single pair $\pm\sigma_0^L$ of ground states. For each L , consider a second spin configuration σ'^L , generated by some prespecified procedure (examples will be given below). Then do for the pair (σ_0^L, σ'^L) what was done for σ_0^L in the original metastate; that is, measure the relative frequency of occurrence of each pair (inside a fixed window, as always). The resulting uniform perturbation metastate gives for both infinite-volume ground and excited states their relative frequency of appearance inside large volumes.

The uniform perturbation metastate is useful when considering recent numerical studies [76, 77, 93, 94] on EA Ising spin glasses in three and four dimensions that may have uncovered new types of states. In each volume, their interface with the ground state has $d_s < d$ with an energy of $O(1)$, independently of lengthscale. (It should be noted that questions have

been raised over the correct interpretation of these numerical results [74, 95], so we consider these states only as an interesting possibility.) Because the two new constructions that lead to these states are done in a translation-invariant manner, a simple extension of theorem 4 to uniform perturbation metastates [92] rules out the possibility that these new states can be ground or pure states; if they exist at all, they must be excited states.

Although the two procedures are different, they seem to lead to similar outcomes. The Krzakala–Martin [77] procedure forces a random pair of spins $(\sigma_z, \sigma_{z'})$ to assume a relative orientation opposite to that in the ground state pair $\pm\sigma_0^L$; the rest of the spins are then allowed to relax to the available lowest-energy configuration. This ensures that at least some bonds in the excited spin configuration, σ'^L , must be changed (i.e. satisfied \leftrightarrow unsatisfied) from σ_0^L . It also ensures that the energy of σ'^L is no more than $O(1)$ above that of σ_0^L , regardless of L . In the Palassini–Young [76] procedure, a novel coupling-dependent bulk perturbation \mathcal{H}_{PY} is added to the Hamiltonian (3) in Λ_L , where

$$\mathcal{H}_{\text{PY}} = (\epsilon/N_b) \sum_{\langle x,y \rangle \in \Lambda_L} (\sigma_0^L)_x (\sigma_0^L)_y \sigma_x \sigma_y, \quad (35)$$

ϵ is independent of L , and N_b is the number of bonds $\langle x, y \rangle$ in Λ_L . So here too the energy of σ'^L is no more than $O(1)$ above that of σ_0^L .

Krzakala–Martin and Palassini–Young excitations have $d_s < d$ and $\theta = 0$. In the two-state droplet picture of Fisher and Huse, however, excitations have $d_s < d$ and $0 < \theta \leq (d-1)/2$. One focus of current research is to determine which, if either, of these pictures holds in realistic spin glasses.

To summarize, spin configurations with space-filling interfaces correspond to new ground or pure states, observable if their energy exponent $\theta \leq (d-1)/2$ and invisible if $\theta > (d-1)/2$. If the ‘space-filling’ exponent $d_s < d$, then the corresponding configurations *could* correspond to new ground or pure states only if $\theta > (d-1)/2$. If $d_s < d$ and $\theta \leq (d-1)/2$, such states can only be excitations [92], and do not signify the presence of additional pure states (for further discussion, see section 10.1.1 below).

10. Summary and discussion

In this section we provide a brief summary of our conclusions and discuss how they tie in to various other approaches.

10.1. Summary

The central theme of this review is that a new set of concepts and methods are needed in the treatment of statistical models with both disorder and frustration. While we have mostly focused on the strong disconnect between infinite- and short-ranged spin glass models, we emphasize that this conclusion essentially results from an application—although a very important one—of our ideas and techniques. These results should be viewed as part of a more extensive and general framework for approaching the study of disordered systems. The unifying concept is that of the *metastate* (section 6), which broadens the focus of study from the conventional one of thermodynamic states, to *distributions* (or probability measures) on thermodynamic states, i.e. *ensembles* of thermodynamic states. Within this framework, universal themes that clarify the study of disordered systems can be considered: the detectability of many states, chaotic size dependence, the relations between interface types and pure states, ‘windows’ in numerical experiments, and others.

As noted, a prominent application of these methods leads to the conclusion that the spin glass differs from most other statistical mechanical systems in that its infinite-ranged version

displays unique low-temperature properties, not shared by corresponding short-ranged models in any dimension and at any temperature (for other differences between homogeneous and disordered systems in general, see section 11.2). A likely underlying cause of this difference is the combination of the following features in the infinite-ranged spin glass: the presence of both ferromagnetic and antiferromagnetic couplings, the statistical independence of all of the couplings, and the scaling of their magnitudes to zero as $N \rightarrow \infty$ (cf [65] and section 5.1.1). The simultaneous presence of all three of these properties seems essential. In other systems (e.g. a number of combinatorial optimization problems) with analogous features, one should expect similar behaviour.

In the following subsections, we expand upon these remarks. We begin by clarifying some essential issues. Although the contents of the following two subsections can be placed naturally into sections 4.1 and 7.2, respectively, the issues they treat have generated sufficient confusion that it seems worthwhile to address them separately.

10.1.1. Are the pure states we discuss the ‘usual’ ones? The spin glass literature is replete with well-defined terms like ‘pure state’ used in imprecise ways, and often interchangeably with other terms such as ‘valley’. A careful discussion has been given in [59], to which we refer the reader. We note that the definition of pure states given in section 4.1 does in fact correspond, at least on a heuristic level, to the intuitive notion of ‘valley’ as a collection of configurations separated from all other configurations by barriers that diverge in the thermodynamic limit.

This last statement can be made more precise, as in [59], by considering a specific dynamics; in our case, the dynamics could be the lattice animal M -spin-flip dynamics discussed in section 11.1.1, for any finite M . Then a pure state may be thought of as a collection of configurations that can be reached from each other via the dynamics in finite time. At a given dimension and temperature, there may be only a single pure state or multiple, disjoint pure states.

For an N -spin system, with N finite, it is often useful to think of spin configurations as belonging to different ‘pure states’ if the dynamical pathway (here it is necessary that $M \ll N$) connecting them requires a time that scales as the exponent of N to some power (see e.g. [96]). This notion has been highly successful in constructing dynamics-based solutions to the SK model [32–34].

10.1.2. Is the $P_{\mathcal{J}}(q)$ used to rule out the standard SK picture the ‘correct’ one? There are two natural objections that are sometimes raised to the conclusion of [85], in which the standard SK picture is ruled out (see e.g. [97]). The specific overlap distributions constructed in [85] and described above involve certain types of averaging, for example, over lengthscale for fixed \mathcal{J} . Are these in fact different from the $P_{\mathcal{J}}(q)$ described by RSB theory, or studied in numerical experiments? Moreover, might it not be that the averaging process used in [85] is itself the cause of the self-averaged nature of the resultant $P_{\mathcal{J}}(q)$?

The answer to these objections appears in [98]. The first point is that, because of chaotic size dependence, the presence of many states probably forbids the construction of any (infinite-volume) $P_{\mathcal{J}}(q)$ with the following two properties:

- (i) the construction can be defined for a.e. \mathcal{J} (necessary if *any* sort of averaging over coupling realizations is to make sense), and
- (ii) the construction entails no averaging of any kind.

So, for example, the construction proposed in [97], where $P_{\mathcal{J}}(q)$ is first defined for finite L , followed by taking a straightforward $L \rightarrow \infty$ limit, cannot work—there is no limiting $P_{\mathcal{J}}(q)$ for this construction (if there are many states).

However, even if there were a construction with the above two properties, the conclusion would be unchanged, because *how $P_{\mathcal{J}}(q)$ is constructed is irrelevant*—there cannot exist any $P_{\mathcal{J}}(q)$ which is both a physical infinite-volume object *and* non-self-averaging. In fact, the only way out would be to construct a clearly unphysical $P_{\mathcal{J}}(q)$ that depends on the choice of origin of the coordinate system. This is an immediate (and rigorous) consequence of the spatial ergodicity of the underlying disorder distribution, as explained in [85].

10.2. Comparison with other work

Our conclusions regarding the disconnect between the SK and EA models are supported by rigorous results. Nevertheless, there have been numerous studies that claim to support the basic features of the RSB, mean-field pure state structure in EA models in finite dimensions. In this section we discuss some of these studies and examine how they can be reconciled with the theorems reviewed above. Given the large number of such studies, we confine ourselves to a small but representative sample. For the sake of conciseness, we will not review studies that claim to support pictures that are consistent with our results [99], such as droplet/scaling. (We emphasize once again, however, that our results make no claim about the correctness of droplet/scaling, and are equally compatible with some other competing pictures.)

10.2.1. Numerical studies. There have been numerous simulations of the EA model, mostly in three and four dimensions, that claim to find results consistent with the RSB picture. We will restrict ourselves here to those studies that examine equilibrium properties. The number of such studies is still quite large, and even with this restriction, we will not attempt to cover them all here; an extensive review is given in [87]. A reasonably representative sample is given by [42, 80, 100–103].

Most of these studies directly measure the spin overlap $P(q)$ (and in some cases, also edge overlaps) for relatively small samples, and generally at temperatures that are not too low. Some problems associated with using $P(q)$ have already been discussed in section 8.2. At least as important, most of these studies measure the overlap in the *full* volume under consideration, rather than in a much smaller ‘window’. But as noted in section 5.2.2, boundary effects—especially on these relatively small samples—will almost certainly dominate, so that no reliable conclusions can be drawn about whether one is really ‘seeing’ the pure state structure. An illuminating example of how false conclusions can be drawn in such a case is given in section 6 of [72].

A further problem concerns the temperature at which many of these studies is carried out. As discussed in section 5 of [57], thermal effects can contribute heavily to non-trivial structure in either spin or edge overlaps, and can therefore also lead to misleading conclusions. Some suggestions for overcoming these problems were presented in [57].

Given all of these issues, it is unclear exactly what is actually being measured in these simulations, and extreme caution needs to be exercised before any conclusions addressing the relevance of RSB to short-ranged systems can be drawn from these studies.

10.2.2. Analytical studies. There have been comparatively few analytical studies that try to extend directly the RSB picture to short-ranged spin glass models. An early attempt examined the averaged free energy in a $1/d$ expansion about $d = \infty$ [104], and found thermodynamics consistent with that of the SK model. Near the critical temperature, an enhancement of RSB effects was found.

A related but more detailed study was done by DeDominicis *et al* [105] who used a field-theoretic approach to extend the RSB theory to short-ranged models via a one-loop expansion.

Here, the expansion was done (at fixed dimension d) in $1/z$ about $z = \infty$, where z is the coordination number. Above six dimensions, the leading mean-field term does not appear to be significantly modified by the loop corrections.

While an interpretation of the $1/d$ or $1/z$ expansion results that supports the standard SK picture is ruled out by our own results, these calculations may nevertheless indicate interesting behaviour. It should first be noted, though, that the calculations, though extensive, involve a number of uncontrolled approximations that need to be understood before any conclusions can be reached. One issue to resolve is the validity of an expansion done about a singular solution ($d = \infty$ for simple nearest-neighbour hypercubic lattices in [104], and $z = \infty$ at fixed dimension d in [105]). It is also unclear what would happen if one went to higher orders in perturbation theory, and even whether these series converge.

Another interesting possibility was raised by Moore [106], who noted that the standard RSB approach expands about only one of the saddles of the mean-field solution. However, experience with other systems (the disordered ferromagnet treated using replica methods [107], and the randomly diluted ferromagnet treated using non-replica methods [108]) implies that one may need to consider all saddle points, including nominally subdominant ones, to get the correct result.

10.2.3. Renormalization group studies and types of chaoticity. Our emphasis on the centrality of the connection between chaotic size dependence and the presence of many states leads naturally to questions regarding whether our result is connected in some way to the presence of other types of chaos that have been encountered in earlier spin glass studies. One well-known example arose from studies [109, 110] of Migdal–Kadanoff renormalization-group transformations on frustrated Ising systems on hierarchical lattices. In these studies, chaotic renormalization-group trajectories were observed, possibly suggesting a type of spin-glass-like behaviour. The specific behaviour uncovered was a chaotic sequence of alternating strong and weak spin–spin correlations as distance increased.

We suspect that this behaviour, if it carries over to spin glasses on Euclidean lattices, arises from different physical origins than those giving rise to chaotic size dependence. Roughly speaking, the renormalization-group studies of [109, 110] uncover how correlation strengths change as one looks at spins increasingly farther apart. In contrast, chaotic size dependence focuses on a *fixed* correlation function, e.g. $\langle \sigma_0 \sigma_1 \rangle$, calculated for each Λ_L using an ordinary Gibbs measure with, say, periodic b.c.s. If $\langle \sigma_0 \sigma_1 \rangle$ changes chaotically as L grows, a clear signal for many pure states is provided. If it does not, then the system likely has no more than a single pair of pure states.

We are not aware of any studies attempting to correlate the presence of many states with changes in $\langle \sigma_x \sigma_y \rangle$ as $|x - y|$ increases, as in [109, 110], but we suspect that such chaoticity may be present *regardless* of the multiplicity of pure states.

Another type of chaoticity in spin glasses is *chaotic temperature dependence* (CTD). This has been the subject of numerous studies, and again can be detected using Migdal–Kadanoff renormalization-group techniques [111]. Roughly speaking, temperature chaos refers to the erratic behaviour of correlations, upon changing temperature, on lengthscales that diverge as the temperature increment goes to zero. It was predicted [45, 49, 111] for the EA spin glass as a consequence of the scaling/droplet *ansatz*, but seemed to be implied as well by the RSB theory [112–114]. More recent numerical and analytical work (see [115] and references therein) have led to claims that CTD is not present in either the SK or the EA model (see also [116]), although [117] allows the possibility of a weak effect at large lattice sizes. Other work [118] claims to see a very small effect at ninth order in perturbation theory. At this time the issue remains unresolved.

CTD and chaotic size dependence (CSD) are clearly different, given that the latter is seen at *fixed* temperature. Nevertheless, the intriguing question arises: are the two somehow related? We do not know the answer at this time. It seems at least plausible that a chaotic change in a fixed correlation function as the volume increases (CSD) would imply a similar chaoticity as temperature changes (CTD), given that the former corresponds to surface changes and the latter to bulk changes. On the other hand, CSD only occurs (for, say, periodic b.c.s) when there exist many competing pure states, and is absent for a single pair, as in droplet/scaling. At the same time, the droplet/scaling picture seems to require CTD, which arises there from a lack of ‘rigidity’ (compared with a ferromagnet) in the spin glass phase. So at the very least, CTD does not seem to imply CSD. Whether the converse is true remains an open question.

10.3. Effects of a magnetic field

In section 2.2 the stability analysis of de Almeida and Thouless [27] was discussed. Although the original intent was to study the stability of the replica-symmetric SK solution [25] in the $T-h$ plane, the consequences of the analysis of [27] remain important for short-ranged models as well. In particular, the phase separation boundary between the paramagnetic and spin glass phases in the $T-h$ plane is now generically referred to as an ‘AT’ line, and debate over its existence for realistic spin glasses remains strong. In the RSB picture, the AT line begins at $(T_c, 0)$ and extends upwards through non-zero values of h (see e.g. figure 48 of [9]). In contrast, the droplet/scaling model predicts that the spin glass phase is unstable to *any* external magnetic field, no matter how small, resulting in no AT line.

However, in the droplet/scaling picture even a small magnetic field will have dramatic dynamical consequences. The droplet theory predicts a ‘magnetic correlation length’ ξ_h that diverges as an inverse power of the field as $h \rightarrow 0$ (the power itself being a function of the domain wall energy exponent θ defined in section 5.2.3). The magnetic correlation length describes roughly the lengthscale over which a field h will destroy local correlations in a spin glass phase. On this lengthscale, the characteristic relaxation time τ_{ξ_h} grows exponentially with ξ_h [49]. Consequently, spin glass correlations can persist for unmeasurably long times in small fields. This makes it difficult to establish for a real spin glass sample whether a true thermodynamic AT line exists or whether one is instead observing a non-equilibrium dynamical effect. In the droplet/scaling picture, one therefore replaces the equilibrium AT line with a dynamical one, separating regions in the $T-h$ plane where the system falls out of equilibrium on accessible laboratory timescales (cf figure 3 in [49]). This feature, however, can create difficulties in the interpretation not only of laboratory experiments but also of numerical simulations (see e.g. [119, 120]).

We turn now to our own results. We point out first that all of our theorems regarding self-averaging of overlaps, and therefore lack of viability of the RSB picture in short-ranged models, are not affected by addition of a magnetic field. However, our results do *not* rule out the presence of a true, thermodynamic AT line. In particular, the many-state chaotic pairs picture remains perfectly consistent with the presence of such a line. Whether an AT line ultimately exists will almost certainly depend on the resolution of the problem of the number of states in *zero* external field: in all likelihood, there is no line if only a single pair of pure states exists, but it should be present if many states exist.

11. Other topics

11.1. Metastable states

Our analysis has focused largely on pure state structure and ordering in short-ranged spin glass models. While our emphasis has centred on equilibrium thermodynamics, we have argued

elsewhere [90] that pure states also deeply influence non-equilibrium dynamical phenomena, such as coarsening, ageing, and others related to dynamical evolution following a deep quench.

Another prominent, and much studied, feature of spin glasses is the presence of many *metastable* states, i.e. states that are energetically stable up to M -spin flips for some finite M , but unstable to flips of clusters of more than M spins. Colloquially speaking, these are spin configurations that are ‘local’ rather than ‘global’ energy (or free energy, at positive temperature) minima; their confining barriers are of height $O(1)$ rather than $O(N^\delta)$ for some $\delta > 0$. These states are believed to be responsible for much of the anomalously slow relaxation features of spin glasses [121, 122], and their presence has led to new numerical techniques such as simulated annealing [123, 124].

In [89], we studied the properties of metastable states not only in EA spin glasses with continuous coupling distributions such as Gaussian, but also in disordered ferromagnets. We will confine ourselves here to spin glasses, and will only present our own results; the reader is referred to [89] for a more detailed discussion, as well as references to other work on metastability.

To clarify the discussion, we define a 1-spin-flip stable state as an infinite-volume spin configuration whose energy, as given by equation (3), cannot be lowered by the flip of any single spin. Similarly, an M -spin-flip stable state ($M < \infty$) is an infinite-volume spin configuration whose energy cannot be lowered by the flip of any cluster of $1, 2, \dots, M$ spins. Recall (section 4.1) that a *ground* state is an infinite-volume spin configuration whose energy cannot be lowered by the flip of *any* finite cluster of spins (i.e. $M \rightarrow \infty$).

11.1.1. A new dynamical method. Our approach to studying metastable states is rather unusual, and based on a new dynamical technique: we construct a natural ensemble (the ‘ M -stable ensemble’) on states that evolve from an initial spin configuration generated through a deep quench via a zero-temperature ‘lattice animal’ dynamics.

We briefly describe this method, starting with the single-spin-flip case. Let σ^0 denote the initial (time zero) infinite-volume spin configuration on Z^d . It is chosen from the infinite-temperature ensemble in which each spin is equally likely to be $+1$ or -1 , independently of the others. The continuous-time dynamics is given by independent, rate-1 Poisson processes at each site x corresponding to those times t at which the spin at x looks at its neighbours and determines whether to flip. It does so only if a flip lowers the system energy (i.e. we consider only zero-temperature dynamics). We denote by ω_1 a given realization of this zero-temperature single-spin-flip dynamics; so a given realization ω_1 would then consist of a collection of random times $t_{x,i}$ ($x \in Z^d, i = 1, 2, \dots$) at every x when spin flips for the spin σ_x are considered.

Given the Hamiltonian (3) and a specific \mathcal{J} , σ^0 and ω_1 , a system will evolve towards a single well-defined spin configuration σ^t at time t (this uses the fact that the individual couplings come from a continuous distribution such as a Gaussian one). It is important to note that these three realizations (coupling, initial spin configuration and dynamics) are chosen independently of one another. The continuous coupling distribution and zero-temperature dynamics together guarantee that the energy per spin $E(t)$ is always a monotonically decreasing function of time.

We now consider *multiple*-spin flips, in which we allow rigid flips of all lattice animals (i.e. finite connected subsets of Z^d , not necessarily containing the origin) up to size M . The case $M = 1$ is the single-spin-flip case just described; $M = 2$ corresponds to the case where both single-spin flips and rigid flips of all nearest-neighbour pairs of spins are allowed; and the case of general M corresponds to flips of 1-, 2-, 3-spin, \dots , M -spin connected clusters. A specific realization of this M -spin-flip dynamics will be denoted ω_M .

There is a technical issue that needs to be addressed, because we wish the dynamics to remain sensible even in the limit $M \rightarrow \infty$. We require that the probability that any fixed spin considers a flip in a unit time interval remains of order 1, uniformly in M . Such a choice would guarantee, for example, that the probability that a spin considers a flip in a time interval Δt vanishes as $\Delta t \rightarrow 0$, again uniformly in M . A further requirement for the dynamics to be well defined is that information should not propagate arbitrarily fast throughout the lattice as M becomes arbitrarily large. It is not difficult to construct such a dynamics, but we omit the technical details here; they can be found in [89].

11.1.2. Results. In this section we present some of the results found in [89]; we omit all proofs and detailed discussions. Some of these results are expected, while others are surprising and at variance with the ‘folk wisdom’ that has evolved over the years. Apart from the intrinsic interest in the structure of metastability in spin glasses, we believe that the information obtained also sheds light on some aspects of pure state structure, if only by way of contrast. All results given below are rigorous, and hold for a.e. \mathcal{J} , σ^0 and ω_M .

Existence and number of metastable states. In an infinite system, the Hamiltonian (3) displays *uncountably* many M -spin-flip stable states, for all finite $M \geq 1$ and for all finite $d \geq 1$.

Convergence of dynamics following a deep quench. It is not obvious a priori whether the system evolves towards a single, final metastable state σ^∞ . It can be proved, however, that such a final state exists almost surely. Equivalently, every spin flips only finitely many times. (This is in contrast to, say, the two-dimensional Ising ferromagnet, where every spin flips *infinitely* many times [90].)

How much information is contained in the initial state? For $M = 1$ in one dimension, precisely half the spins in σ^∞ are completely determined by σ^0 , with the other half completely undetermined by σ^0 . For higher d and the same dynamics, it can be shown that a dynamical order parameter q_D , measuring the percentage dependence of σ^∞ on σ^0 , is strictly between 0 and 1.

Size of basins of attraction. The basins of attraction of the individual metastable states are of negligible size, in the following sense: almost every σ^0 is on a boundary between two or more metastable states. Equivalently, the union of the domains of attraction of *all* of the metastable states forms a set of measure zero, in the space of all the σ^0 . (A similar result for pure states was proved in [90].)

Distribution of energy densities. For any M , almost every M -spin-flip stable state has the same energy density, E_M . Moreover, the dynamics can be chosen so that $E_1 > E_2 > E_3 > \dots$, and furthermore E_M for any finite M is larger than the ground state energy density, which is the limit of E_M as $M \rightarrow \infty$.

Overlaps. Almost every pair of metastable states, whether two M -spin-flip stable states or one M - and one M' -spin-flip stable state, has zero spin overlap. So the overlap distribution function, for either *all* of the metastable states or only all those with any fixed M , is simply a δ -function at zero.

Scaling of number of metastable states with volume. For sufficiently large volumes, the number of metastable states in finite samples scales exponentially with the volume in general d for states of any stability.

Remanent magnetization. If the initial state is uniformly +1 in one dimension, the remanent magnetization is known to be 1/3 (for $M = 1$) [125]. In higher dimensions, a heuristic calculation for the Gaussian spin glass gives a rigorous lower bound on the remanent

magnetization that for large d behaves like $\exp(-2d \log(d))$. Exact results can be obtained in all d for some other models [89].

Correspondence between pure and metastable states. More precisely, if there are multiple pure or ground state pairs, does the spin configuration corresponding to a typical metastable state ‘live in’ the domain of attraction of a single pure or ground state, as is commonly thought? The answer is *no* in both cases: almost every metastable state will be on a boundary in configuration space between multiple pure or ground states.

Finally, we ask: does knowledge of metastable states provide information on the number or structure of thermodynamic ground or pure states? So far, the answer seems to be largely no. For example, in the one-dimensional spin glass there is only a single pair of thermodynamic ground states, but an uncountable number of infinite-volume M -spin-flip stable states for any finite M . This example illustrates a potential difficulty with numerical studies in higher dimensions, aimed at determining the number of ground or pure states: the presence of many metastable states could complicate interpretations of these studies.

11.2. The statistical mechanics of homogeneous versus disordered systems

Our hope is that this topical review has conveyed some of the depth and richness of the physics and mathematics of the equilibrium statistical mechanics of spin glasses. We conclude by returning briefly to question 7, raised in section 3: in what ways do we now understand how the statistical mechanical treatment of systems with quenched disorder differs in fundamental ways from that of homogeneous systems? In some sense, much of this review focused on one or another aspect of this question, but it may be helpful for us to tie together some of the common threads that have run through much of this review. We emphasize that we will not attempt to provide a comprehensive or even an extensive answer to this question, which remains largely open, but will only focus on some of the insights that our analysis of the spin glass problem may have uncovered. (A similar discussion appears also in [126].)

Most homogeneous systems that can be treated by classical equilibrium statistical mechanics share several salient features. Whether one is studying crystals, ferromagnets, ferroelectrics, liquid crystals, or even some quantum systems such as superconductors and superfluids, the analysis of the low-temperature phase is simplified by various *spatial* symmetries, such as translational, orientational, gauge and others. Of course, no such symmetries are manifestly apparent in many disordered systems.

It has also been known for many years [18] that averaging over the quenched disorder presents additional complications. But this issue has been extensively discussed in many other reviews [9–14, 16] and will not be further addressed here. Another complicating feature present in some (but not all) disordered systems is frustration. While frustration may also occur in homogeneous systems (e.g. triangular antiferromagnets), its joint presence with quenched disorder may result in more profound effects. We hereafter confine ourselves to systems with both disorder *and* frustration, taking the spin glass as their prototype.

The nature of the low-temperature phase is a clear point of departure between homogeneous systems and spin glasses (assuming they have a low-temperature phase). Homogeneous systems typically display a relatively simple order parameter, representing the nature of ordering in a pure or ground state that is unique up to an overall simple symmetry transformation that leaves the Hamiltonian invariant. It remains unknown whether the spin glass similarly possesses a simple order parameter in the EA model in finite dimensions, but it certainly does not in the infinite-ranged SK model.

The striking contrast between the nature of broken symmetry in the RSB theory, as opposed to that present in most homogeneous systems, is clearly an important difference

between them and infinite-ranged spin glasses. However, as extensively discussed in this review, it now appears that this type of broken symmetry is absent in short-ranged spin glass models.

However, rather than seeing this as a disappointing piece of news, our viewpoint is that it is perhaps indicative of a far deeper contrast between spin glasses and simple homogeneous systems: in all such latter systems of which we are aware, mean-field theory has been invaluable in providing the basic information concerning the nature of broken symmetry, order parameter(s) and low-temperature behaviour in general. Mean-field theory's main shortcoming lies in its behaviour very close to the transition temperature, but even this failure usually occurs only in sufficiently low dimensions (which often includes the physically interesting case of three dimensions): there usually exists an upper critical dimension above which mean-field theory also provides the correct critical exponents.

If the infinite-ranged SK model becomes exact for the EA model in infinite dimensions, as is commonly believed, then we have a new feature whose contrast with homogeneous systems is perhaps even more striking than the presence of RSB: that is, the failure of mean-field theory to provide a correct description of the low-temperature phase, even far from the critical point, in *any* finite dimension. Equivalently, the $d \rightarrow \infty$ limit of the EA model is singular. This possibility was broached by Fisher and Huse [48], and our work confirms their conjecture. Some of the reasons why mean-field theory fails in any finite dimension are presented in section 5.1 and reviewed in section 10.1.

It is important to emphasize that the disconnect between the SK and EA models is profound [65]: when one tries to transfer concepts in *either* direction, contradictory and even absurd results can ensue. One interesting question discussed in [65], worth repeating here, is whether a new type of thermodynamic object in place of states may be appropriate for understanding the meaning of RSB in the SK model. The usual notions of states are local ones: that is, they describe the behaviour of correlation functions (at fixed locations) and related quantities. These notions do not seem to work for the SK model. This suggests the intriguing possibility that a more global object might be constructible that would provide a more natural 'fit' for the SK model. We do not know whether this will turn out to be the case, but it is clear that if it does, such an object would be substantially different from the thermodynamic states that have been used up until now.

A different issue concerns the non-existence of a thermodynamic limit for states (or equivalently, correlations). This is manifested as CSD, and occurs when coupling-independent b.c.s are used and there exist many observable pure states. This is really a reflection of the lack of any spatial symmetries that allow one to choose b.c.s, or an external symmetry-breaking field, that can lead to the existence of such a limit.

One could, of course, artificially obtain CSD even in a simple homogeneous ferromagnet below T_c by, say, choosing random b.c.s independently in each Λ_L . But there one also knows how to choose b.c.s to obtain a limiting pure state, including the interface states, whose analogues would be invisible pure states in spin glasses (section 9). The option of choosing b.c.s that lead to a limiting pure state does not now exist for spin glasses if they possess many pure state pairs, and there may or may not be fundamental reasons preventing that option from ever becoming available. Regardless of whether it does, it remains interesting and useful that the presence of CSD provides a clear signal of the existence of many states (section 5).

Besides CSD, it has been speculated that spin glasses, both short ranged and infinite ranged, display CTD, as discussed in section 10.2.3. As indicated there, whether CTD actually exists in spin glasses remains an open question. However, its potential presence in spin glasses represents a qualitatively new thermodynamic feature of at least some types of disordered systems.

A further contrast between homogeneous and disordered systems is provided by the presence of observability versus invisibility of different types of pure states. This is intimately interwoven with CSD, and also with the nature of the interfaces separating these states; these interfaces can, in principle, be different from those seen in homogeneous systems (section 9). These issues all arise from the use of coupling-dependent versus independent b.c.s. No analogue for this distinction exists in homogeneous systems, but it is very possibly of basic importance in systems with quenched disorder.

Finally, we would like to emphasize that if this entire discussion has uncovered any fundamental unifying principle, it is the set of ideas and techniques encapsulated in the construct of the *metastate*. It is our strong belief that any final understanding of the spin glass phase, and possibly that of other inhomogeneous systems, will make extensive use of this concept.

Acknowledgments

We thank T Rosenbaum for suggesting we write this review, and A van Enter and P Contucci for useful remarks on the manuscript. We are also indebted to M A Moore for a number of illuminating comments and for pointing us towards several useful references. CMN and DLS were partially supported by the National Science Foundation under grants DMS-01-02587 and DMS-01-02541 respectively.

References

- [1] Ruderman M A and Kittel C 1954 *Phys. Rev.* **96** 99
- [2] Kasuya T 1956 *Prog. Theor. Phys.* **16** 45
- [3] Yosida K 1957 *Phys. Rev.* **106** 893
- [4] Maletta H and Felsch W 1979 *Phys. Rev. B* **20** 1245
- [5] For a recent review of quantum spin glasses, see Bhatt R 1997 *Spin Glasses and Random Fields* ed A P Young (Singapore: World Scientific) pp 225–49, and references therein
- [6] Cannella V and Mydosh J A 1972 *Phys. Rev. B* **6** 4220
- [7] Mydosh J A 1987 *Heidelberg Colloquium on Glassy Dynamics* ed J L van Hemmen and I Morgenstern (Berlin: Springer) pp 24–39
- [8] Wenger L E and Keesom P H 1976 *Phys. Rev. B* **13** 4053
- [9] Binder K and Young A P 1986 *Rev. Mod. Phys.* **58** 801
- [10] Chowdhury D 1986 *Spin Glasses and Other Frustrated Systems* (New York: Wiley)
- [11] Fischer K H and Hertz J A 1991 *Spin Glasses* (Cambridge: Cambridge University Press)
- [12] Bouchaud J P, Cugliandolo L, Kurchan J and Mézard M 1998 *Spin Glasses and Random Fields* ed A P Young (Singapore: World Scientific) pp 161–223
- [13] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
- [14] Stein D L 1989 *Lectures in the Sciences of Complexity* ed D L Stein (New York: Addison-Wesley) pp 301–55
- [15] Newman C M and Stein D L 1997 *Mathematics of Spin Glasses and Neural Networks* ed A Bovier and P Picco (Boston, MA: Birkhäuser) pp 243–87
- [16] Dotsenko V 2001 *Introduction to the Replica Theory of Disordered Statistical Systems* (Cambridge: Cambridge University Press)
- [17] Edwards S and Anderson P W 1975 *J. Phys. F: Met. Phys.* **5** 965
- [18] Brout R 1959 *Phys. Rev.* **115** 824
- [19] Toulouse G 1977 *Commun. Phys.* **2** 115
- [20] Anderson P W 1978 *J. Less-Common Met.* **62** 291
- [21] Newman C M and Stein D L 1994 *Phys. Rev. Lett.* **72** 2286
- [22] Newman C M and Stein D L 1996 *J. Stat. Phys.* **82** 1113
- [23] Banavar J R, Cieplak M and Maritan A 1994 *Phys. Rev. Lett.* **72** 2320
- [24] Derrida B 1980 *Phys. Rep.* **67** 29
- [25] Sherrington D and Kirkpatrick S 1975 *Phys. Rev. Lett.* **35** 1792
- [26] Thouless D J, Anderson P W and Palmer R J 1977 *Phil. Mag.* **35** 593
- [27] de Almeida J R L and Thouless D J 1978 *J. Phys. A: Math. Gen.* **11** 983

- [28] Parisi G 1979 *Phys. Rev. Lett.* **43** 1754
- [29] Parisi G 1983 *Phys. Rev. Lett.* **50** 1946
- [30] Mézard M, Parisi G, Sourlas N, Toulouse G and Virasoro M 1984 *Phys. Rev. Lett.* **52** 1156
- [31] Mézard M, Parisi G, Sourlas N, Toulouse G and Virasoro M 1984 *J. Physique* **45** 843
- [32] Sompolinsky H and Zippelius A 1981 *Phys. Rev. Lett.* **47** 359
- [33] Sompolinsky H 1981 *Phys. Rev. Lett.* **47** 935
- [34] Sompolinsky H and Zippelius A 1982 *Phys. Rev. B* **25** 6860
- [35] Rammal R, Toulouse G and Virasoro M A 1986 *Rev. Mod. Phys.* **58** 765
- [36] Bray A J and Moore M A 1980 *J. Phys. C: Solid State Phys.* **13** L469
- [37] Fisher M E and Singh R R P 1990 *Disorder in Physical Systems* ed G Grimmett and D J A Welsh (Oxford: Clarendon) pp 87–111
- [38] Thill M J and Hilhorst H J 1996 *J. Physique I* **6** 67
- [39] Ogielski A T 1985 *Phys. Rev. B* **32** 7384
- [40] Ogielski A T and Morgenstern I 1985 *Phys. Rev. Lett.* **54** 928
- [41] Kawashima N and Young A P 1996 *Phys. Rev. B* **53** R484
- [42] Marinari E, Parisi G and Ritort F 1994 *J. Phys. A: Math. Gen.* **27** 2687
- [43] McMillan W L 1984 *J. Phys. C: Solid State Phys.* **17** 3179
- [44] Bray A J and Moore M A 1985 *Phys. Rev. B* **31** 631
- [45] Bray A J and Moore M A 1987 *Phys. Rev. Lett.* **58** 57
- [46] Fisher D S and Huse D A 1986 *Phys. Rev. Lett.* **56** 1601
- [47] Huse D A and Fisher D S 1987 *J. Phys. A: Math. Gen.* **20** L997
- [48] Fisher D S and Huse D A 1987 *J. Phys. A: Math. Gen.* **20** L1005
- [49] Fisher D S and Huse D A 1988 *Phys. Rev. B* **38** 386
- [50] Middleton A A 1999 *Phys. Rev. Lett.* **83** 1672
- [51] Palassini M and Young A P 1999 *Phys. Rev. B* **60** R9919
- [52] Hartmann A K 2000 *Eur. Phys. J. B* **13** 539
- [53] Newman C M and Stein D L 2000 *Phys. Rev. Lett.* **84** 3966
- [54] Newman C M and Stein D L 2001 *Commun. Math. Phys.* **224** 205
- [55] Palassini M and Young A P 1999 *Phys. Rev. Lett.* **83** 5126
- [56] Marinari E and Parisi G 2000 *Phys. Rev. B* **62** 11677
- [57] Newman C M and Stein D L 2002 *J. Stat. Phys.* **106** 213
- [58] Infinite-volume Gibbs measures $\rho_{J,T}$ 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 Georgii H O 1988 *Gibbs Measures and Phase Transitions* (Berlin: de Gruyter Studies in Mathematics)
- [59] van Enter A C D and van Hemmen J L 1984 *Phys. Rev. A* **29** 355
- [60] Young A P 1983 *Phys. Rev. Lett.* **51** 1206
- [61] An interesting example of a *non-random* system with an ultrametric ground state structure is discussed in van Enter A C D, Hof A and Miękisz J 1992 *J. Phys. A: Math. Gen.* **25** L1133
- [62] Palmer R G 1982 *Adv. Phys.* **31** 669
- [63] Palmer R G and Stein D L 1985 *Relaxations in Complex Systems* ed K L Ngai and G B Wright (Washington, DC: US Government Printing Office) pp 253–9 (1985–461–700/20001)
- [64] Newman C M and Stein D L 1992 *Phys. Rev. B* **46** 973
- [65] Newman C M and Stein D L 2003 *Nonrealistic Behavior of Mean Field Spin Glasses* (available as Preprint cond-mat/0301202)
- [66] Aizenman M, Lebowitz J and Ruelle D 1987 *Commun. Math. Phys.* **112** 3
- [67] Pastur L and Shcherbina M 1991 *J. Stat. Phys.* **62** 1
- [68] Guerra F and Toninelli F L 2002 *Commun. Math. Phys.* **230** 71
- [69] Feller W 1971 *An Introduction to Probability Theory and Its Applications* vol 2 (New York: Wiley) p 124
- [70] See also the discussion in section 3 of van Enter A C D and Fröhlich J 1985 *Commun. Math. Phys.* **98** 425
- [71] An example of non-physical behaviour arising from the use of coupling-*dependent* boundary conditions can be found in Gandolfi A, Newman C M and Stein D L 1993 *Commun. Math. Phys.* **157** 371
- [72] Newman C M and Stein D L 1998 *Phys. Rev. E* **57** 1356
- [73] Aizenman M and Fisher D S 1991 unpublished
- [74] Marinari E and Parisi G 2000 *Phys. Rev. B* **62** 11677
- [75] Marinari E and Parisi G 2001 *Phys. Rev. Lett.* **86** 3887
- [76] Palassini M and Young A P 2000 *Phys. Rev. Lett.* **85** 3017
- [77] Krzakala F and Martin O C 2000 *Phys. Rev. Lett.* **85** 3013

- [78] Newman C M and Stein D L 1996 *Phys. Rev. Lett.* **76** 4821
- [79] Aizenman M and Wehr J 1990 *Commun. Math. Phys.* **130** 489
- [80] Marinari E, Parisi G and Ruiz-Lorenzo J J 1997 *Spin Glasses and Random Fields* ed A P Young (Singapore: World Scientific) pp 59–98
- [81] Newman C M and Stein D L 1997 *Phys. Rev. E* **55** 5194
- [82] See, for example, Nadkarni M G 1998 *Basic Ergodic Theory* (Basle: Birkhäuser)
- [83] Wiener N 1939 *Duke Math. J.* **5** 1
- [84] Franz S, Mézard M, Parisi G and Peliti L 1998 *Phys. Rev. Lett.* **81** 1758
- [85] Newman C M and Stein D L 1996 *Phys. Rev. Lett.* **76** 515
- [86] Newman C M and Stein D L 2000 *Phys. Rev. Lett.* **84** 3966
- [87] Marinari E, Parisi G, Ricci-Tersenghi F, Ruiz-Lorenzo J J and Zuliani F 2000 *J. Stat. Phys.* **98** 973
- [88] Guerra F 1995 private communication
see also Guerra F 1996 *Int. J. Mod. Phys. B* **19** 1675
- [89] Newman C M and Stein D L 1999 *Phys. Rev. E* **60** 5244
- [90] Newman C M and Stein D L 1999 *J. Stat. Phys.* **94** 709
- [91] A very simple example, suggested to us by A van Enter, of invisible states can be found in the Mattis site-disordered Ising model
Mattis D C 1976 *Phys. Lett. A* **56** 421
After gauging away the randomness, one ends up with a ferromagnetic Ising model with random boundary conditions, whose interior state chaotically varies with volume between the positively and negatively magnetized states
van Enter A C D 1990 *J. Stat. Phys.* **60** 275
While interface states exist in this model, just as in the ferromagnet, they are not seen when using boundary conditions not conditioned on the disorder realization. The free energy of an interface scales as L^{d-1} , while the free energy difference between, say, the uniformly plus and minus boundary conditions on the site-disordered system scales only as $L^{(d-1)/2}$, in accord with our discussion in the text.
- [92] Newman C M and Stein D L 2001 *Phys. Rev. Lett.* **87** 077201
- [93] Houdayer J, Krzakala F and Martin O C 2000 *Eur. Phys. J. B* **18** 467
- [94] Katzgraber H G, Palassini M and Young A P 2001 *Phys. Rev. B* **63** 184422
- [95] Middleton A A 2001 *Phys. Rev. B* **63** 060202(R)
- [96] Mackenzie N D and Young A P 1982 *Phys. Rev. Lett.* **49** 301
- [97] Parisi G 1996 unpublished (available as *Preprint cond-mat/9603101*)
- [98] Newman C M and Stein D L 1996 unpublished (available as *Preprint adap-org/9603001*)
- [99] See, for example, Moore M A, Bokil H and Drossel B 1998 *Phys. Rev. Lett.* **81** 4252
- [100] Caracciolo S, Parisi G, Patarnello S and Sourlas N 1990 *J. Physique* **51** 1877
- [101] Reger J D, Bhatt R N and Young A P 1990 *Phys. Rev. Lett.* **64** 1859
- [102] Marinari E, Parisi G, Ruiz-Lorenzo J J and Ritort F 1996 *Phys. Rev. Lett.* **76** 843
- [103] Marinari E, Parisi G, Ricci-Tersenghi F and Ruiz-Lorenzo J J 1998 *J. Phys. A: Math. Gen.* **31** L481
- [104] Georges A, Mézard M and Yedidia J 1990 *Phys. Rev. Lett.* **64** 2937
- [105] See DeDominicis C, Kondor I and Temesvári T 1997 *Spin Glasses and Random Fields* ed A P Young (Singapore: World Scientific) pp 119–60 and references therein
- [106] Moore M A 2003 private communication
- [107] Bray A J, McCarthy T J, Moore M A, Reger J D and Young A P 1987 *Phys. Rev. B* **36** 2212
- [108] McKane A J 1994 *Phys. Rev. B* **49** 12003
- [109] McKay S R, Berker A N and Kirkpatrick S 1982 *Phys. Rev. Lett.* **48** 767
- [110] McKay S R and Berker A N 1984 *J. Appl. Phys.* **55** 1646
- [111] Banavar J R and Bray A J 1987 *Phys. Rev. B* **35** 8888
- [112] Kondor I 1989 *J. Phys. A: Math. Gen.* **22** L163
- [113] Ney-Nifle M and Hilhorst H J 1993 *Physica A* **193** 48
- [114] Ritort F 1994 *Phys. Rev. B* **50** 6844
- [115] Billoire A and Marinari E 2000 *J. Phys. A: Math. Gen.* **33** L265
- [116] Bouchaud J-P, Dupuis V, Hammann J and Vincent E 2002 *Phys. Rev. B* **65** 024439
- [117] Billoire A and Marinari E 2002 *Preprint cond-mat/0202473*
- [118] Rizzo T and Crisanti A 2002 *Preprint cond-mat/0209333*
- [119] Singh R R P and Huse D A 1991 *J. Appl. Phys.* **69** 5225
- [120] Grannan E R and Hetzel R E 1991 *Phys. Rev. Lett.* **67** 907
- [121] Dasgupta C, Ma S-K and Hu C-K 1979 *Phys. Rev. B* **20** 3837

- [122] See, for example,
Binder K and Kinzel W 1983 *Heidelberg Colloquium on Spin Glasses* ed J L van Hemmen and I Morgenstern (Berlin: Springer) pp 279–304
Morgensetern I 1983 *Heidelberg Colloquium on Spin Glasses* ed J L van Hemmen and I Morgenstern (Berlin: Springer) pp 305–27
Young A P 1983 *Heidelberg Colloquium on Spin Glasses* ed J L van Hemmen and I Morgenstern (Berlin: Springer) pp 328–45
- [123] Kirkpatrick S, Gelatt C D and Vecchi M P 1983 *Science* **220** 671
- [124] Černý V 1985 *J. Optim. Theory Appl.* **45** 41
- [125] Fernandez J F and Medina R 1979 *Phys. Rev. B* **19** 3561
- [126] Stein D L 2003 Spin glasses: still complex after all these years? *Quantum Decoherence and Entropy in Complex Systems* ed H-T Elze (Berlin: Springer) at press