Nonequilibrium entropy and entropy distributions

Stephen A. Langer,* James P. Sethna, and Eric R. Grannan†
Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853
(Received 24 August 1989)

Glasses have nonzero zero-temperature entropies. Because they are out of equilibrium, the "thermodynamic" entropy, determined by heat flow, is not equal to the "statistical" entropy, which measures volumes in phase space. We discuss the relationship between the two kinds of entropy in nonequilibrium systems and show that the thermodynamic entropies measured by cooling and heating form lower and upper bounds to the statistical entropy. In a computer simulation of a glass, the distribution of thermodynamic entropies measured by repeated fast coolings provides information about the dynamics of the glass. Entropy distributions are presented for a spin glass and a simple two-level system, and the distributions are used as a tool to compare the dynamics of the two models.

I. INTRODUCTION

A. Equilibrium

In equilibrium, there is no history. In computing the properties of an equilibrium system, such as a gas, we can safely assume that states are populated according to their Boltzmann (or Fermi or Bose) weights and derive thermodynamic properties, even though we may know that half an hour ago all the gas molecules were in one half of the room. It does not matter that the molecules were unevenly distributed half an hour ago—they rearrange themselves quickly on the time scale of any measurement. On the other hand, nonequilibrium systems are sensitive to their history. The properties of a piece of glass at this moment *will* depend upon which half of the room it was in half an hour ago. Processes in nonequilibrium systems take place on time scales comparable to the time scale of any measurement, so measurements made at different times or under different conditions will yield different results.

Thermodynamics and statistical mechanics apply in situations where all processes are either so fast that the system is in equilibrium or so slow that it is static. This paper discusses the intermediate regime, where traditional thermodynamics and statistical mechanics do not apply. The primary motivation is the study of glasses, but the results could be applied to any nonequilibrium system. The examples we will study in detail are a spin glass and a simple two-level system.

B. Glasses

Glasses are ubiquitous nonequilibrium systems. Because atomic rearrangements in glasses take place on anything from phonon (10^{-12} sec.) to laboratory (min. or h.) to astronomical time scales, there will always be processes in a glass that are neither fast nor slow on the time scale of the measurement. Glasses, therefore, exhibit history dependence and hysteresis. History dependence means that the state of the glass, and its measurable properties, such as volume or heat capacity, depend on how it has been treated. Figure 1 shows the heat capacity of a sample of boron trioxide glass. The two solid curves were measured the same way, but at different times, and give different results. Hysteresis is a particular form of history dependence—properties measured while cooling a glass through temperature T will not be reproduced when the glass is subsequently heated back through T. The dotted curve in Fig. 1 is the heat capacity measured while cooling, and has a markedly different character from the two heating curves. The heating curves have a bump at 275 °C, while the cooling curve does not.

The bump in the specific heat in Fig. 1 is a signature of

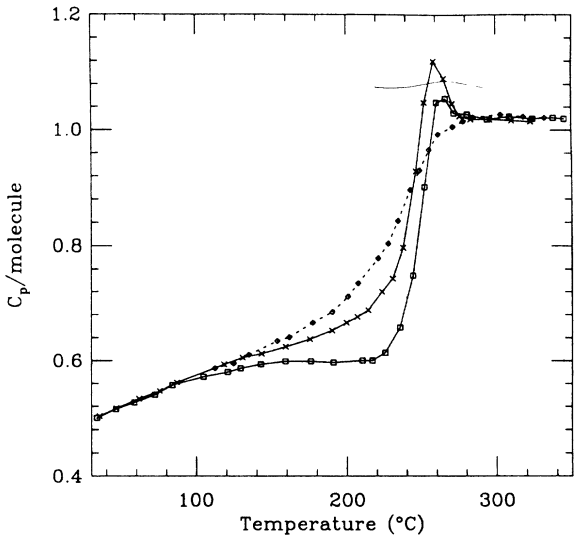

FIG. 1. Specific heat of B_2O_3 glass measured while heating and cooling. The glass was first rapidly cooled from the melt (500 °C→50 °C in a half hour), then heated from 33 °C to 345 °C in 14 h (solid curve with squares), cooled from 345 °C to room temperature in 18 h (dotted curve with diamonds), and finally heated from 35 °C to 325 °C (solid curve with crosses). Data are from Thomas and Parks.[27]

the glass transition, during which the sample falls out of equilibrium. The transition is marked by a diverging viscosity and rounded discontinuities in the thermodynamic derivative quantities, such as the specific heat and the thermal expansion coefficient. The glass transition looks like a second-order phase transition, except that it is not sharp, and the transition temperature T_g is history dependent—in particular, it depends on how fast the temperature is changing. Slower cooling rates lead to lower T_g's.

C. Entropy

It has been known at least since the early 1920's that the zero-temperature entropy (the "residual" entropy) of glasses is not zero. The residual entropies of a number of glasses are displayed in Table I. These measurements must be taken with a grain of salt, because the residual entropy, like all glassy properties, is history dependent, and we do not know the details of each measurement. However, it is easy to see that the entropy is of order one per molecular or atomic unit (here, and for the remainder of this paper, Boltzmann's constant $k_B \equiv 1$). Glasses apparently violate Nernst's theorem, which states that the entropy of a material at absolute zero is zero.[4] Nernst's theorem, however, applies only to equilibrium quantum systems in which the energy is a monotonically increasing function of the temperature. In glasses, the energy depends on history, as well as temperature, so the theorem does not apply.

The standard interpretation[3,5] for the residual entropy is that because entropy is the log of the number of states occupied by a system, the residual entropy is the log of the number of ground states of the glass. The measurements in Table I were all made by calorimetry. It is remarkable that measuring the heat flowing out of a sample as it is freezing into a *particular* glassy state gives any information about the *total* number of such states.

D. Summary

In Sec. II we discuss the relationship between the entropy that is measured in the laboratory (the "thermodynamic" entropy) and the entropy that counts the number of states (the "statistical" entropy). We present the correct definition of the statistical entropy, and discuss various approximations that have been used in the literature. Section III introduces the concept of an entropy distribution (a distribution of measured values of the thermodynamic entropy) in the context of a simple spin-glass simulation. The distribution is a measurable history-dependent quantity. Section IV introduces a toy model of a glass—the two-level system—and discusses the origins of its entropy distribution. In Sec. V, we return to the spin glass, attempting to use its entropy distribution as a probe of its dynamics. In Sec. VI, we discuss more realistic simulations and the applicability of these concepts to real materials. Finally, in Sec. VII, we mention some possible future applications.

TABLE I. Residual entropies for a variety of glasses, in units in which $k_B = 1$. Entropies are per molecule, except for *a*-Au:Ge:Si, which is per atom.

Material	$S(T=0)$
Cyclohexene[a]	1.4
Ethanol[a]	1.1
Glycerol[b]	2.8
Isopentane[a]	1.7
Isopropylbenzene[a]	1.4
Vitreous Ice[a]	<1.6
Vitreous Silica[c]	1.0
a-Au:Ge:Si[c]	0.7

[a]Reference 1.
[b]Reference 2.
[c]Reference 3.

II. THERMODYNAMIC AND STATISTICAL ENTROPY

A. Definitions and theorem

An experimentalist who wishes to know the entropy of a system uses the thermodynamic definition of temperature, $T^{-1} = dS/dE$, to measure what we call the "thermodynamic" entropy,

$$S_{\text{therm}}(T) = S(T_0) + \int_{T_0}^{T} \frac{dQ}{T} \ . \tag{1}$$

Q is the heat flow into the system in question, T is the equilibrium temperature, and T_0 is some initial reference temperature where the entropy is known. On the other hand, when a theorist talks about entropy, he or she is usually interested in the volume the system occupies in phase space, and computes what we call the "statistical" entropy,

$$S_{\text{stat}} = -\text{Tr}\,\hat{\rho}\ln\hat{\rho} \ , \tag{2}$$

where $\hat{\rho}$ is the density matrix. The two definitions of entropy are equivalent only in equilibrium.

The thermodynamic entropy (1) is easy to measure in experiments and computer simulations, whereas the statistical entropy (2) is impossible to measure in experiments and difficult to measure in simulations. Historically,[3,5] there has been some ambiguity in how it is even to be defined. How is the density matrix to be computed, and which states are to be included in the trace? We will discuss the correct definition below, but first we prove that the thermodynamic entropy measured on heating (S_{heat}) and cooling (S_{cool}) provides upper and lower bounds, respectively, on any reasonable definition of the statistical entropy.[6,7]

Theorem 1: The thermodynamic entropy (1) measured by heating and cooling provides upper and lower bounds, respectively, on the statistical entropy, provided that the statistical entropy meets the following three criteria: (1) it must be extensive; (2) it must equal the thermodynamic entropy in any equilibrium system or subsystem; and (3) in a closed system in the thermodynamic limit, it must increase with time.

Proof: Consider an experiment in which a glass forming material (by which we mean any material that falls out of equilibrium as the temperature is lowered) is held in con-

tact with an equilibrium heat bath, whose temperature $T(t)$ is controlled externally (see Fig. 2). The heat bath may be thought of as infinite succession of heat baths at infinitesimally different temperatures, so that the temperature of the bath in contact with the glass may be changed without doing any external work or adding any heat. Start at some high temperature T_M (the melting temperature, for example) where the glass former is an equilibrium liquid and its entropy (statistical and thermal) is $S_{\rm liq}(T_M)$. The liquid could have been obtained by slowly melting a crystal, and since the crystal had zero entropy at zero temperature, the liquid entropy can be found exactly from (1). Or, perhaps, the liquid could have been formed by cooling a gas from a high enough temperature that the ideal gas law was in effect. In any case, assume that $S_{\rm liq}(T_M)$ is known. Then cool the heat bath quickly through the glass transition temperature to some temperature T_1, possibly zero, and measure the heat flow $\dot{Q}<0$ from the bath to the glass. The measured, thermodynamic entropy of the glass at T_1 is

$$S_{\rm cool}(T_1)=S_{\rm liq}(T_M)+\int_{T=T_M}^{T=T_1}\frac{\dot{Q}}{T}dt \ . \tag{3}$$

The temperature T is always the temperature of the bath, since the glass is out of equilibrium and does not have a well-defined temperature. During the process, the entropy of the bath changes by $\int(\dot{Q}/T)dt$, so the statistical (and thermodynamic) entropy of the bath at T_1 is

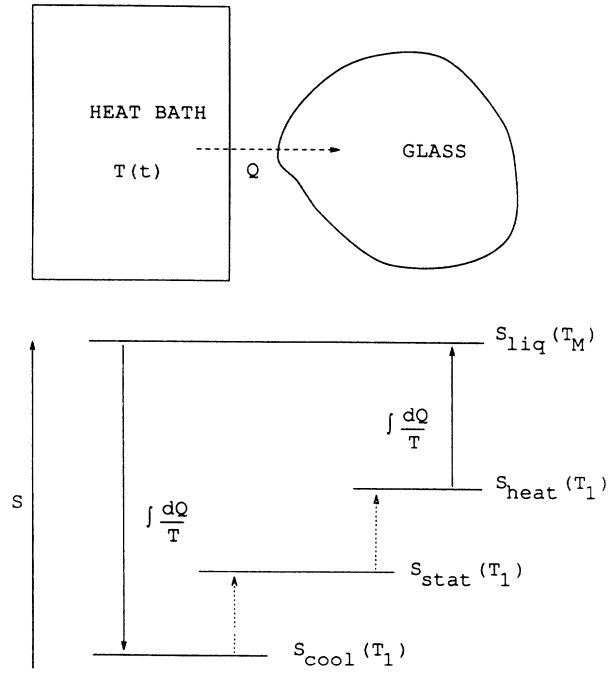

FIG. 2. Schematic of entropy flow for cooling and heating. A glass is held in contact with a heat bath at temperature $T(t)$ while the heat Q flowing between the bath and glass is measured. The vertical solid lines represent the measured thermodynamic entropy of the glass, and the dotted lines represent irreversible, unmeasurable processes within the glass contributing to the statistical entropy.

$$S_{\rm bath}(T_1)=S_{\rm bath}(T_M)-\int_{T=T_M}^{T=T_1}\frac{\dot{Q}}{T}dt \ . \tag{4}$$

Now, the total statistical entropy of the closed glass-bath system must increase with time, so

$$S_{\rm liq}(T_M)+S_{\rm bath}(T_M)\le S_{\rm stat}(T_1)+S_{\rm bath}(T_1) \ ,$$

or, using (3) and (4),

$$S_{\rm cool}(T_1)\le S_{\rm stat}(T_1) \ . \tag{5}$$

To measure the entropy on heating, $S_{\rm heat}$, one would take the glass at the low temperature T_1 and measure how much heat had to be added to reach the liquid state,

$$S_{\rm heat}(T_1)+\int_{T=T_1}^{T=T_M}\frac{\dot{Q}}{T}dt=S_{\rm liq}(T_M) \ . \tag{6}$$

Again, the measured quantity is the change in entropy of the heat bath, $-\int(\dot{Q}/T)dt$. Since the total statistical entropy of the whole system must increase,

$$S_{\rm stat}(T_1)+S_{\rm bath}(T_1)\le S_{\rm liq}(T_M)+\left[S_{\rm bath}(T_1)-\int_{T=T_1}^{T=T_M}\frac{\dot{Q}}{T}dt\right] \ ,$$

which, together with (6) and (5) gives

$$S_{\rm cool}(T_1)\le S_{\rm stat}(T_1)\le S_{\rm heat}(T_1) \ . \tag{7}$$

That is, the thermodynamic entropies measured on heating and cooling are bounds to the statistical entropy. Q.E.D.

A good question to ask is, "How far apart are the bounds?" If, in any real experiment or simulation, the upper and lower bounds always agree to within the available accuracy, then the issue is rather pedantic. Measuring the thermodynamic entropy will be as good as measuring the statistical entropy, and the distinction between the two is unimportant. In Sec. VI we will argue that the bounds should differ by on the order of 1% in a real experiment, and in Sec. III we will show that the bounds can differ enormously in computer simulations. There is then a real distinction between statistical and thermodynamic entropies.

We should point out that the statement "entropy always increases" is a statement about average systems in the thermodynamic limit. In a computer simulation, far from the thermodynamic limit, we may have to average many computations of $S_{\rm heat}$ or $S_{\rm cool}$ before establishing good bounds on $S_{\rm stat}$.

B. Statistical entropy

When computing the statistical entropy for a glass, it is clearly incorrect to use standard equilibrium thermodynamics and statistical mechanics. Equilibrium theories tell us to find the partition function

$$Z={\rm Tr}e^{-\beta H}=\sum_i e^{-\beta E_i} \ , \tag{8}$$

and the density matrix

$$\hat{\rho}_i = \frac{e^{-\beta E_i}}{Z}$$

in order to compute the free energy

$$F = -T \ln Z ,$$

the internal energy

$$U = \frac{1}{Z} \mathrm{Tr} H \hat{\rho} = \frac{1}{Z} \sum_i E_i e^{-\beta E_i} = -\frac{\partial}{\partial \beta} \ln Z ,$$

and the entropy

$$S = \frac{U - F}{T} .$$

This is incorrect because most glasses have a crystalline ground state—a state which has lower energy than all the others. Call the ground-state energy E_0. Then, as $T \to 0$,

$$Z = e^{-\beta E_0} \left[1 + \sum_{i \neq 0} e^{-\beta (E_i - E_0)} \right] \to e^{-\beta E_0} ,$$

since $E_i > E_0$. Hence, at zero temperature, $F = U = E_0$, $S = 0$, and the system is crystalline. A common solution to this difficulty is to exclude the crystalline state from the problem. In other words, just *declare* that the trace in (8) is to include glassy states only. This is an artificial approach, and is not the one we will use.

Palmer[5] and Jäckle[3] have discussed various approximate ways of computing thermodynamic quantities in nonergodic systems. A system is ergodic if (8) applies—that is, if its properties can be found by averaging over phase space. Below the glass transition, glasses are nonergodic. Palmer and Jäckle divide phase space up into "components" ("cells," in Jäckle's terminology). Within any one component, the glass is ergodic, but transitions between components never occur. The problem of defining the statistical entropy is reduced to the problem of finding the probability p_i of the system's being in the ith component, and deciding at what stage in the calculation it is appropriate to average over components. Some of the methods used in the literature follow:

(1) Each component is treated as an independent equilibrium system, and the free energy is the average free energy of the components. The residual entropy is zero.

(2) $\{p_i\}$ are frozen, at the "fictive" temperature T_f of the glass, near the glass transition. $\hat{\rho}_i(x)$ is the probability of a given state x within component i, $\hat{\rho}(x) = p_i \hat{\rho}_i(x)$ for x in component i, and S is found from (2). The residual entropy is

$$-\sum_i p_i(T_f) \ln p_i(T_f) > 0 .$$

(3) $\{p_i\}$ are determined by the equilibrium probabilities of the components. $p_i = e^{-\beta F_i}/Z$, with Z given by (8). F_i is the free energy of component i, treated as an independent equilibrium system.

(4) $\{p_i\}$ are computed from the entire history of the system.

Of course, it is possible to define components within components, and so on. This approach was used by Jäckle and Kinzel[8] in discussing the Edwards-Anderson Ising spin glass. They find an upper bound on the residual entropy, given by $\ln N_S(E_0)$, where $N_S(E)$ is the density of metastable states with energy E, and E_0 is the residual energy. Ettelaie and Moore[9] found agreement with this result for a one-dimensional Ising spin glass. Jäckle[10] actually derived $S_{\mathrm{cool}} < S_{\mathrm{stat}}$ for a version of item 2.

A different approach to finding the entropy of computer simulations is to assume that, although the simulated material is out of equilibrium, it can be described by an equation of state. Certain coefficients in the equation of state are determined by measuring P and V while holding T fixed, and then changing the temperature and repeating the measurement. Given the equation of state, the entropy is calculated by equilibrium thermodynamics. This technique was applied to the hard-sphere fluid by Gordon *et al.*,[11] and subsequently by Woodcock.[12] Both authors found residual entropies on the order of one per atom. Cape and Woodcock[13] applied the same method to a soft-sphere model, finding a residual entropy of 0.3 per atom.

Our point of view is that for (2) to hold, the density matrix $\hat{\rho}$ must give the probability of finding the system in a certain state, *given a specific thermal history.* That is, $\hat{\rho} = \hat{\rho}(t)$ is a function of the external variables $T(t)$, $P(t)$, $V(t)$, etc. To the extent that it is possible to define components at all, item (4) is correct. We will determine the statistical entropy by finding how $\hat{\rho}$ changes in time, starting in a state where $\hat{\rho}$ is unambiguously defined, and computing. The trace in (2) is taken over all states. We do not explicitly forbid the crystalline state, but if during its evolution the material rarely finds it, it will not contribute substantially to the statistics. We do not partition phase space, and we compute the entropy directly, without using equilibrium thermodynamics. Appendix A demonstrates that the statistical entropy defined in this way for a system governed by a master equation fulfills the requirements set out in Theorem 1.

The entropy given by (2), with time dependent $\hat{\rho}$, measures the volume of phase space occupied by an ensemble of glasses, each with the same *macroscopic* history (the same cooling schedule, pressure, etc.) but different *microscopic* histories (paths in phase space). If the system is large, it can effectively explore all of its histories at once: it can be divided into virtually identical subsystems, each of which will explore phase space independently. If the system is small, the thermodynamic entropy (defined for each member of the ensemble) fluctuates, forming a distribution. The statistical entropy is a single number, but is determined by the whole ensemble. For a small system, only the average of the distribution of thermodynamic entropies provides bounds on the statistical entropy.

III. SPIN GLASS I: ZERO-TEMPERATURE ENTROPY AND ENTROPY DISTRIBUTIONS

A. Simulation

To illustrate the entropy bounds introduced in the last section, we look at small Ising spin glass. The Ising spin

glass is often used as a model glass. Randomness in the model causes it to have some glasslike properties: a multiplicity of metastable zero-temperature states, nonzero zero-temperature entropy, hysteresis, and gradual freezing.

We are using the $\pm J$ Ising model because we need a pedagogical example, and this model is easy to simulate. The model consists of a collection of spins $\{S_i\}$ on a square lattice, where each spin can either point up ($S_i=1$) or down ($S_i=-1$). The randomness comes in via the coupling between spins. Nearest-neighbor spins S_i and S_j are connected by a bond J_{ij}, which with a specified probability is either $+J$ (ferromagnetic) or $-J$ (antiferromagnetic). Our simulations have equal ferromagnetic and antiferromagnetic probabilities. The Hamiltonian for a particular realization (set of J_{ij}'s) is

$$H_J=-\sum_{\langle ij\rangle} J_{ij}S_iS_j\ , \tag{9}$$

where the sum is over all pairs of nearest-neighbor sites.

The simulation uses Monte Carlo heat bath dynamics. The inputs are the annealing schedule $T(t)$, giving the temperature T at each time step t of the simulation, and the set of random bonds $\{J_{ij}\}$. The equilibrium energy $E_{\rm eq}$ and entropy $S_{\rm eq}$ at the initial temperature are computed from the equilibrium high-temperature expansion of the partition function.[14] When initializing a simulation, we first use the Hamiltonian (9) to find the energy E. This energy will not exactly equal the equilibrium energy $E_{\rm eq}$ computed from the high-temperature expansion, because the high-temperature expansion is for the average energy of a thermodynamic ensemble, and E is the energy of a particular spin configuration. Furthermore, the high-temperature expansion takes on continuous values, whereas in the simulation E is quantized in units of $2J$. Because the initial energy is not the equilibrium energy, it would be a mistake to initialize the entropy to its equilibrium value, $S_{\rm eq}$. Instead, we start the simulation with

$$S=S_{\rm eq}+(E-E_{\rm eq})/T\ .$$

As the simulation evolves, the energy E and the thermodynamic entropy $\int dQ/T$ are found by summing ΔE and $\Delta E/T$ for all those spins flipped at each time step. Averages are computed by averaging over runs of the simulation, i.e., repeating the whole process with the same $\{J_{ij}\}$ and $T(t)$, but choosing different random numbers during the time evolution. This is in contrast to the equilibrium method of computing thermodynamic quantities, which is to hold the temperature constant and average over time in a single run.

B. Average temperature dependence

Figure 3 shows the energy per spin of a sample spin glass, averaged over a number of runs of the simulation, for two different cooling rates. Glasslike properties are clearly exhibited. The faster cooling rate results in a higher residual energy, leaving the system farther from equilibrium. The glass behaves differently on heating and cooling—the heat capacity dE/dT measured during the fast run is always positive while cooling, but goes temporarily negative while heating. This is a history effect—the glass was above equilibrium when the temperature started to rise, so the energy was still relaxing downwards towards equilibrium.

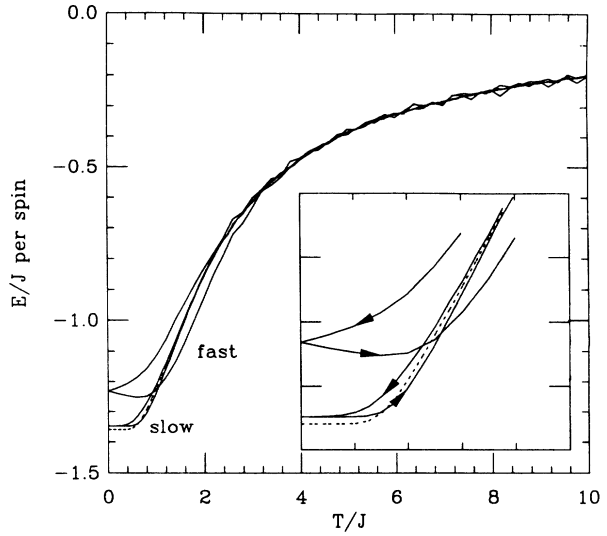

FIG. 3. Energy vs temperature for 5×5 spin-glass sample 789 cooled and heated between $T=10J$ to $T=0$ in 50 (fast) and 1000 (slow) Monte Carlo time steps. The curves represent an average over 1000 and 10 000 runs, respectively. The dotted line is the equilibrium energy, computed by the brute force computation of the spin-glass partition function.

Figure 4 shows the average entropy $\int dQ/T$ measured during the cooling and heating process. Two heating curves are plotted: what is actually measured,

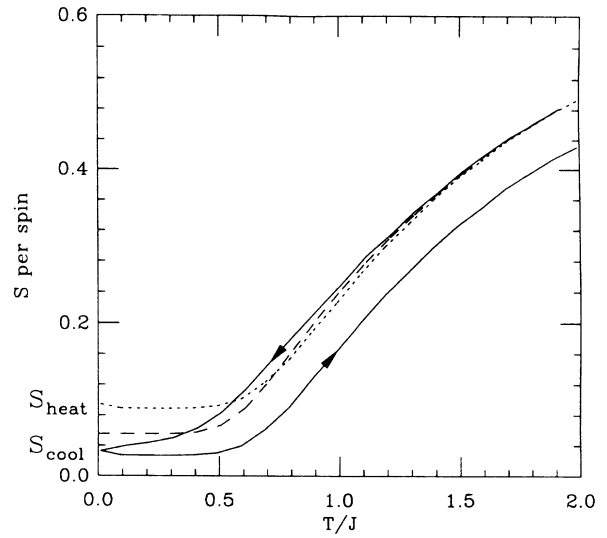

FIG. 4. Entropy vs temperature for 5×5 spin-glass sample 789 cooled and heated between $T=10J$ and $T=0$ in 1000 Monte Carlo time steps and averaged over 10 000 runs. The solid curves are the cumulative integral of dQ/T. To compute the upper bound on the zero-temperature statistical entropy, the heating curve must be shifted upwards (dotted line) to match the high-temperature cooling curve. The dashed line is the equilibrium entropy.

$$S'_{\text{heat}}(T)=S(0)+\int_{T=0}^{T} dQ/T$$

(solid line), and the same curve shifted upwards to match the cooling curve at high temperatures (dotted line). The entropy measured in heating, as defined in Sec. II, corresponds to the shifted curve, because changes in entropy must be referred to the equilibrium high-temperature value, not $S_{\text{cool}}(0)$. It is seen that at all temperatures the shifted entropy measured on heating is above the entropy measured on cooling, in agreement with Eq. (7).

C. Zero-temperature entropy

Table II shows the statistical entropy and the bounds derived from the thermodynamic entropy for a number of different realizations of a 5×5 spin glass at a variety of cooling rates. The thermodynamic entropies were determined by averaging over at least 1000 runs of the simulation, and the statistical entropies were found by counting states, as described below. Figure 5 summarizes the results for the first spin-glass sample. Notice that the lower bound S_{cool} is almost always positive, proving that the sample indeed has a positive statistical zero-temperature entropy, at least for some cooling rates. This by itself does not prove that the sample is glassy—a $\pm J$ spin glass typically has multiple ground states, so it will have a positive zero-temperature entropy even in equilibrium. The lower bound for the slowly cooled ferromagnet in Table II is positive because the ferromagnet has two ground states, and its total entropy per spin at zero temperature is ln2. On the other hand, if the lower bound to the statistical entropy is not only positive, but above the equilibrium entropy, then the system is definitely out of equilibrium. The slowest run of 5×5 sample 789 (10^4 MC steps) gives a lower bound greater than the equilibrium entropy.

Our spin glasses are small enough that we can actually compute their statistical entropies and compare them with the thermodynamic bounds, as in Table II and Fig. 5. As mentioned in Sec. II, to find the statistical entropy it is necessary to define the trace and density matrix in Eq. (2). We will trace over all states of the spin glass, and find the density matrix by computing its time evolution from a known state. To be precise, we assume that at the initial high temperature all states are equally likely, and the probability of a low-temperature state is proportional to the number of times it was found by the simulation. For the $\pm J$ spin glass there is one subtlety in this approach. Because all bonds in the model have the same strength, any spin with equal numbers of satisfied and unsatisfied bonds can flip at will, even at zero temperature. Spin configurations that differ only by such zero-

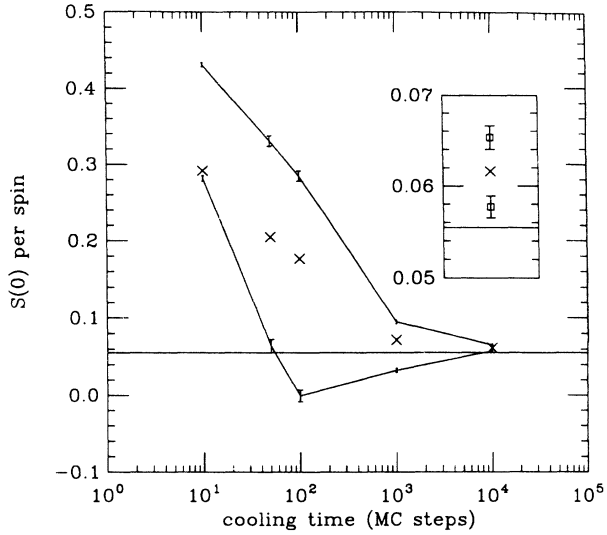

FIG. 5. Statistical entropy (crosses) and upper and lower bounds (solid lines) on the zero temperature entropy of 5×5 sample 789. The horizontal solid line is the equilibrium residual entropy. The insert is an expanded view of the points at 10^4 MC steps.

TABLE II. Zero-temperature entropies for a number of spin-glass configurations. The sample number is the random seed used to generate the bonds $\{J_{ij}\}$. The samples were cooled and heated between a $T=10J$ (well above freezing) to $T=0$ in the specified number of steps.

MC Steps	S_{cool}	S_{stat}	S_{heat}	S_{eq}
		Sample 789 (5×5 $\pm J$)		
10	0.282±0.004	0.292	0.431±0.003	(ln4)/25=0.055
50	0.064±0.008	0.205	0.331±0.007	
100	−0.001±0.008	0.176	0.285±0.007	
1000	0.033±0.003	0.072	0.095±0.003	
10000	0.057±0.001	0.062	0.065±0.001	
		Sample 1024 ($5\times5\pm J$)		
100	0.16±0.16	0.225	0.26±0.13	(ln 248)/25=0.2205
1000	0.21±0.02	0.220	0.22±0.02	
10000	0.220±0.006	0.2205	0.221±0.005	
		5×5 Ferromagnet		
10	−0.027±0.010	0.192	0.365±0.007	(ln2)/25=0.027 726
100	−0.036±0.003	0.028 285	0.070±0.002	
1000	0.023±0.001	0.027 725	0.032±0.002	

energy spin flips must be counted as the same state. The number of different configurations in a state σ is the *degeneracy* d_σ of the state. If each state was found n_σ times in a simulation of $N=\sum_\sigma n_\sigma$ coolings, then the statistical entropy is

$$S_{\text{stat}}=-\sum_\sigma \frac{n_\sigma}{N}\ln\frac{n_\sigma}{Nd_\sigma}\ .$$

Alternatively, if each spin configuration S was found m_S times in the $N=\sum_S m_S$ runs, then the entropy is simply

$$S_{\text{stat}}=-\sum_S \frac{m_S}{N}\ln\frac{m_S}{N}\ .$$

As $N\to\infty$, these two expressions agree, but the first is more accurate for finite N, because it does not rely on bad statistics to determine the degeneracy of each state. For a large sample, it may be difficult to determine d_σ, and the second form will be required.

For all glasses, the dependence of the upper and lower bounds on cooling rate should have the form shown in Fig. 5. It is the slow dynamics of the glass that determines the shapes of the curves. Transitions to higher-energy states while heating and to lower energy states while cooling take place later than they would if the system were in equilibrium. The entropy measured while cooling is always below the statistical entropy because these transitions release more entropy (E/T is larger because T is smaller) than they would if they were not delayed. Similarly, the entropy measured on heating is always too high.

If the cooling rate is slow and the system is nearly in equilibrium, the delays will be small and both the upper and lower bounds will approach the statistical entropy, which in turn will approach the equilibrium entropy. This effect is clearly seen in Fig. 5. If, as the cooling rate is decreased, the bounds S_{heat} and S_{cool} approach the statistical entropy S_{stat} faster than S_{stat} approaches the equilibrium entropy S_{eq}, there will be a cooling rate below which $S_{\text{cool}}>S_{\text{eq}}$.

If the cooling rate is somewhat faster, the glass will be farther from equilibrium, and the bounds will deviate more from the statistical entropy. The lower bound can be negative.

If the cooling rate is very fast, the system does not have time to react at all. Few transitions will take place, and the statistical entropy will tend towards the high-temperature equilibrium value, which is ln2 in these simulations. The effect is clearly exhibited in Fig. 5. Depending on whether or not the simulation updates the temperature before or after flipping the spins, as the rate diverges either $S_{\text{cool}}(0)\approx S_{\text{eq}}(\infty)$ and $S_{\text{heat}}(0)\to+\infty$, or $S_{\text{cool}}(0)\to-\infty$ and $S_{\text{heat}}\approx S_{\text{eq}}(\infty)$. This is not profound—it is an artifact of the discrete nature of the time in the Monte Carlo simulation.

D. Entropy distributions

There are error bars on the heating and cooling curves in Fig. 5 and Table II. These reflect the uncertainty in finding the mean of a set of values. Each time the simulation is cooled or heated, it evolves through a different set of states, and makes the transitions between these states at different times and temperatures. The measured thermodynamic entropy, $\int dQ/T$, must therefore differ from run to run of the simulation. (Figure 6 shows how the entropy fluctuates during three different runs of a spin-glass

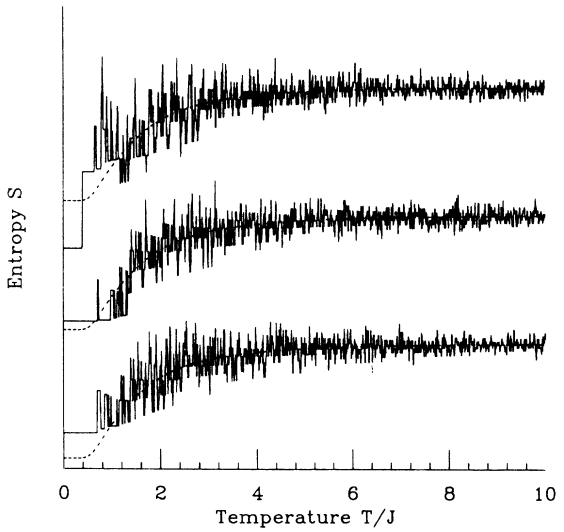

FIG. 6. Thermodynamic entropy vs temperature for three runs of a simulation of 5×5 sample 789, cooled from $T=10J$ to $T=0$ in 1000 MC steps. The dotted lines are the equilibrium entropy (the curves are displaced vertically for clarity). Although both of the top two traces are from runs that ended in the ground state of the system, the measured residual entropy is negative for one and positive for the other. The bottom trace is from a run that ended in an excited state.

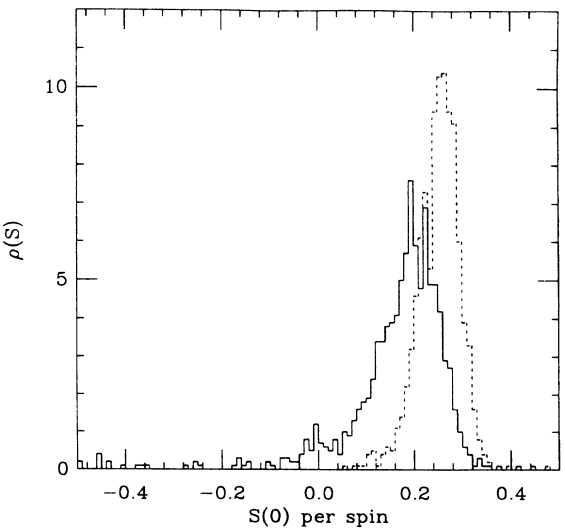

FIG. 7. Residual entropy distribution for 5×5 sample 1024 cooled and heated between $T=10J$ and $T=0$ in 100 Monte Carlo steps. The solid line is the distribution measured on cooling, so its mean is a lower bound on the statistical zero-temperature entropy. The dotted line is the distribution measured on heating, and its mean is the corresponding upper bound. Each distribution contains 1000 runs of the simulation.

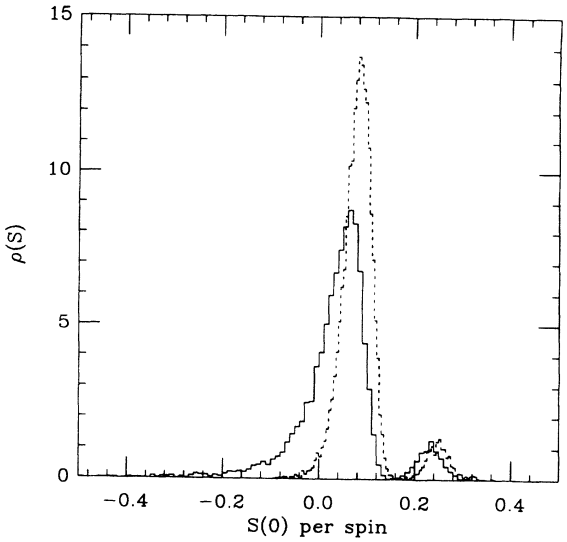

FIG. 8. Residual entropy distribution for 5×5 sample 789 under the same conditions as in Fig. 7. The clearly defined double peaks contrast to the single peaks in the distributions for sample 1024, and show that the sample is behaving like a two-level system (see Sec. V). Each distribution contains 10 000 runs of the simulation.

simulation.) To get the entropy defined in (1), we have to average many runs of the simulation. In other words, in order to measure $S_{\rm cool}$ or $S_{\rm heat}$, we have to take the mean of a *distribution* of values. The properties of glasses are history dependent, and this entropy distribution depends both on the dynamics of the glass and its thermal history. By the end of this paper, we will show that by varying the thermal history, we can extract information about the dynamics from the distribution.

We define the entropy distribution $\rho(S)$ so that $\rho(S)dS$ is the probability of measuring $\int dQ/T$ to be between S and $S+dS$. $\rho(S)$ is really a function of S and a functional of $T(t)$, but the $T(t)$ dependence will usually not be made explicit. On the computer, we measure $\rho(S)$ by running a simulation repeatedly and making a normalized histogram of the measured values of S.

Figures 7 and 8 are the zero temperature entropy distributions generated by cooling and heating two samples linearly [i.e., $T(t)$ is linear in t] between $T=10J$ and $T=0$. The detailed structure of the distributions will be discussed in later sections. For now it suffices to say that the distribution is clearly sensitive to the microscopic dynamics of the glass. The simulations differ only by the choice of bonds $\{J_{ij}\}$, yet the first set of distributions has one hump while the second set has two. The distributions are also sensitive to changes in the cooling schedule $T(t)$. We will return to this topic in Sec. V.

IV. TWO-LEVEL SYSTEMS I: A TOY MODEL

We are interested in the history-dependent dynamic properties of glasses and their relationship to the residual entropy. Glasses are not in equilibrium because they get hung up in metastable states—the transition rates are slow and ergodicity is broken. The simplest conceivable model glass is a two-level system (TLS), a system with a single metastable state.[10,15] In this section, we first demonstrate how and why a TLS exhibits history dependence. We then use the TLS model to estimate the difference in the entropy bounds for a real glass. Finally, we examine a TLS's entropy distribution, which will inspire experiments for spin-glass simulations.

A. The model

The TLS model is possibly the simplest model for a glass (see Fig. 9). It has only one degree of freedom, the population of the upper well, and, because the TLS is either in its excited state or not in its excited state, that degree of freedom is discrete. To jump from one well to the other, the system must be thermally activated over a barrier. (The actual transitions are assumed to have no duration.) The energy of the ground state is 0, and the energy of the excited state is the "asymmetry" ε. The energy barrier is of height V, measured from the bottom of the upper well. The transition attempt frequency, or small oscillation frequency at the bottom of the wells, is Γ_0, so the thermally activated transition rate for hopping from the excited state, over the barrier, and into the ground state is

$$\Gamma(T)=\Gamma_0\exp(-V/T)\ . \tag{10}$$

Consider an ensemble of TLS's, all with the same V, ε, Γ_0, and thermal history. If the average population of the upper wells in this ensemble is n, and the average population of the lower wells is $1-n$, then

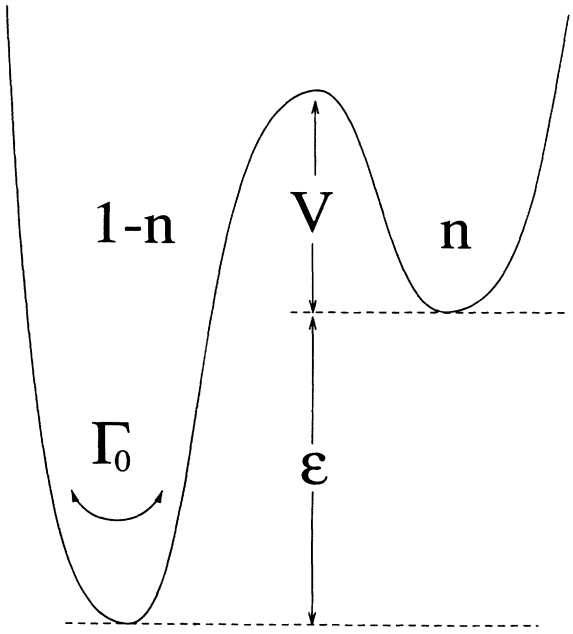

FIG. 9. A two-level system, showing the barrier height V and asymmetry ε. The small oscillation frequency in the bottom of the wells is Γ_0. An individual TLS is always either in its excited state ($\uparrow$) or in its ground state ($\downarrow$). By taking an ensemble of TLS's, we define the average population of the upper well, n.

$$\frac{dn}{dt}=\Gamma_0 e^{-\beta V}[e^{-\beta\varepsilon}(1-n)-n]\ , \tag{11}$$

where $\beta\equiv 1/T$. This is the master equation for the evolution of the population $n(t)$, or equivalently, the average energy $\varepsilon n(t)$. Transitions out of the lower well, at a rate $e^{-\beta\varepsilon}\Gamma(T)$, increase n, and transitions out of the upper well, at a rate $\Gamma(T)$, decrease n. The equilibrium population is

$$n_0(T)=\frac{e^{-\beta\varepsilon}}{1+e^{-\beta\varepsilon}}\ . \tag{12}$$

Figure 10 compares a numerical solution of (11) (dots) with the equilibrium population (dotted line) of a TLS.

As a TLS is cooled, while the temperature is high the transition rate $\Gamma(T)$ is much larger than the rate of change of n_0, so $n(T)\approx n_0(T)$. (see Figure 10). When the temperature becomes low, $\Gamma(T)$ gets very small, so $n(T)$ is nearly constant. The residual population $[n(T=0)]$ is roughly the population of the system at the "freezing" temperature, T^*, where the rate of transitions over the barrier is comparable to the cooling rate

$$\left|\frac{dn_0}{dt}\right|_{T=T^*}\approx\Gamma(T^*)n_0(T^*)\ . \tag{13}$$

The residual population will decrease as the cooling rate is decreased, which is characteristic of glasses.

Glasslike hysteresis effects show up when the TLS is now heated back up to high temperatures. At first, the population remains constant at its residual value, because the TLS is still frozen. Once T gets to be around T^*, the TLS unfreezes, and, because n is still greater than n_0, the population will decrease although the temperature is increasing. The measured specific heat, $\varepsilon\, dn/dT$, will be negative in this region. This is not a paradox, since the system is not in equilibrium. As the temperature rises further, n_0 increases, and when it becomes greater than n, n will begin to increase. At high temperatures $n\lesssim n_0$.

We define the statistical entropy of the TLS the same way we defined it for the spin glass—we use for the density matrix the probability of finding a state, given the thermal history. Therefore,

$$S_{\rm stat}(T)=-n\ln n-(1-n)\ln(1-n)\ . \tag{14}$$

$S_{\rm stat}$ defined this way is zero only if $n=0$ or $n=1$. For any nonzero coupling rate, there will be a nonzero residual population, which implies a positive zero-temperature entropy, a particularly glassy property.

The equilibrium entropy of the TLS may be found from the partition function, $Z=1+\exp(-\beta\varepsilon)$, using $E=(-\partial/\partial\beta)\ln Z$, $F=-T\ln Z$, and $F=E-TS$, resulting in

$$S_{\rm eq}=\ln(1+e^{-\beta\varepsilon})+\beta\varepsilon n_0\ . \tag{15}$$

The same result may be obtained from (14) with $n=n_0$.

B. Entropy distributions for a two-level system

Entropy distributions for a TLS are computed numerically via Monte Carlo simulation, the same way they are for the spin glass. The Monte Carlo simulations subject each TLS in the ensemble to the same thermal history, but allow them to make transitions independently. n and $\rho(S)$ are determined by averaging the final states and tabulating the final entropies, respectively. Figure 11 depicts the evolution of the distribution for a particular TLS, cooled at a rate of $0.01V$ per MC step.

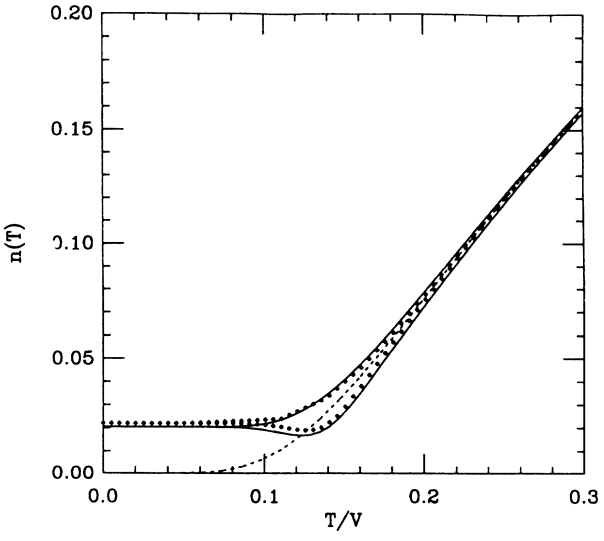

FIG. 10. Population vs temperature for a TLS with asymmetry $\varepsilon/V=0.5$ cooled linearly to zero temperature with dimensionless cooling rate $|\dot{T}|/\Gamma_0 V=10^{-5}$ and then heated at the same rate. The data points are the result of direct numerical integration of (11), and the solid line is the asymptotic result from Ref. 16. The dashed line is the equilibrium population, n_0.

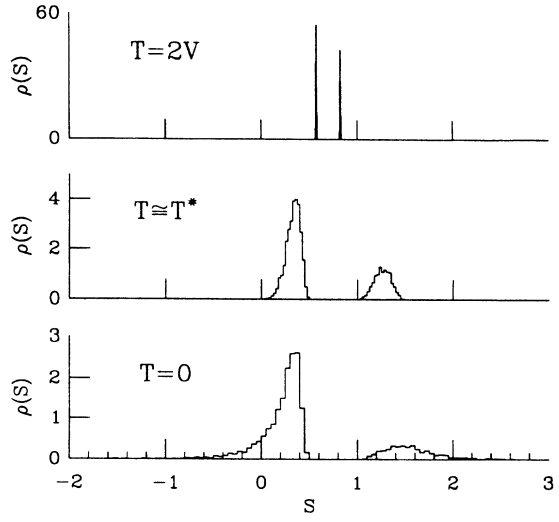

FIG. 11. Entropy distributions from a Monte Carlo simulation of a TLS with asymmetry $\varepsilon=0.5V$ at three different temperatures while cooling at a rate $\delta=0.01$ (100 MC steps per unit T/V).

What should we expect the entropy distribution to look like? A high temperatures, as in the top trace in Fig. 11, the TLS is in equilibrium. In equilibrium, history is unimportant, so the distribution should depend only on the temperature. At infinite temperature, $S_{\rm therm}=S_{\rm stat}$ and the population is evenly distributed between the two wells, so

$$\rho(S)=\delta(S-\ln 2)=\delta(S-S_{\rm eq}(T=\infty)) .$$

Cool the TLS very slowly to temperature $T>T^*$. Since the system is in equilibrium at all times, the transition rate is fast compared to the cooling rate, and all the TLS's in the ensemble make transitions at the same temperatures. Hence, the measured values of $\int dQ/T$ for each TLS in the ensemble will be the same, and $\rho(S)$ will be sharply peaked. Whenever a TLS jumps from the lower well to the upper well, however, it absorbs energy ε from the environment, and its entropy increases by $\beta\varepsilon$. The distribution must therefore have two delta function peaks, separated by $\beta\varepsilon$. The upper peak, at $S_\uparrow$, corresponds to members of the TLS ensemble that are in the upper well at the time of measurement, while the lower peak, at $S_\downarrow$, corresponds to TLS's in their ground state.

$$S_\uparrow - S_\downarrow = \beta\varepsilon$$

and

$$n_0 S_\uparrow + (1-n_0) S_\downarrow = S_{\rm eq} ,$$

implies

$$S_\uparrow = S_{\rm eq} + (1-n_0)\beta\varepsilon = \ln(1+e^{-\beta\varepsilon}) + \beta\varepsilon , \tag{16a}$$

$$S_\downarrow = S_{\rm eq} - n_0\beta\varepsilon = \ln(1+e^{-\beta\varepsilon}) , \tag{16b}$$

and

$$\rho(S)=n_0\delta(S-S_\uparrow)+(1-n_0)\delta(S-S_\downarrow) . \tag{17}$$

As the ensemble is cooled further, to $T\gtrsim T^*$, the TLS's will begin to fall out of equilibrium. The transition rate will slow, and not all members of the ensemble will make transitions at every temperature. The integrals $\int dQ/T$ will not be the same for the whole ensemble, so the peaks in the entropy distribution will broaden, as in the middle curve in Fig. 11. The amount of broadening, as well as the temperature at which it occurs, depends on the history $T(t)$.

At temperatures below T^*, the equilibrium positions of the peaks change rapidly ($S_\uparrow \to \infty$ as $T\to 0$). The peaks move by means of entropy transferring transitions between the wells of the TLS, which are infrequent when $T<T^*$. The peaks in the distribution will be frozen into place when the TLS freezes. At low temperatures they will be separated by $\beta^*\varepsilon$, rather than $\beta\varepsilon$. There will still be some transitions from the upper well to the lower well, but they will transfer an entropy $\beta\varepsilon>\beta^*\varepsilon$, so weight in the distribution will move from the upper peak to the low entropy side of the lower peak, creating a long tail. This effect is clearly seen in the lower curve in Fig. 11.

In the Monte Carlo simulations, the initial states of an ensemble of TLS's were chosen using Eq. (12). The initial entropies of the TLS's were assigned according to (16). After a run, comparing the measured n versus T to n_0 versus T verified that the initial temperature was high enough that the equilibrium approximation was valid. That is, if the ensemble fell out of equilibrium immediately, the initial temperature was too low for that cooling rate.

V. SPIN GLASS II AND TLS II: USING THE ENTROPY DISTRIBUTION

The entropy distribution of a glass depends on the dynamics of the glass. By simulating spin glasses and two-level systems, we have used the dynamics to predict the distribution. We would like to be able to go the other way—to simulate a glass and determine its dynamics from its entropy distribution. As a concrete example, we can ask how well the spin glasses of Sec. III are described as two-level systems. From their entropy distributions, can we extract their metastable states, energy barriers, and attempt frequencies?

We will first discuss how to extract information from a TLS entropy distribution. Then we will apply the same techniques to spin-glass distributions. Finally, we will explain the results by examining the microscopic states of the spin glass.

A. TLS entropy distributions

As shown in Sec. IV, the distance between the two peaks in a zero-temperature TLS entropy distribution, such as the upper curves in Fig. 12, is roughly $\beta^*\varepsilon$, where $T^*=1/\beta^*$ is the temperature at which the TLS falls out of equilibrium. Therefore, we can measure the ratio

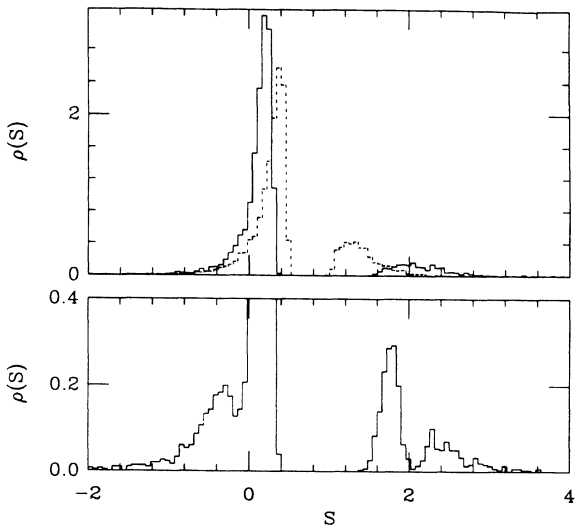

FIG. 12. Three entropy distributions for a TLS with $\epsilon=0.5$, $V=1$, and $\Gamma_0=1$ per MC step. The first (solid line in upper plot) was cooled to zero temperature continuously at a rate of $0.001V$ per MC step. The second (dotted line) was cooled twenty times faster. The shift in the peaks reflects the change in T^*. The third (lower plot) was cooled at the original rate, but was quenched from $T=0.30V$ to $T=0.25V$, and shows the resulting holes in the tails of the peaks. Each distribution was compiled from 5000 runs of the simulation. Note the different vertical scales.

ε/T^* directly from the distribution.

With a little more work, we can extract ε and T^* independently. The long tail on the low entropy side of the lower peak in the distribution comes from transitions from the upper well to the lower well at temperatures below T^*. Assume that the upper peak is narrow and is frozen at its equilibrium position $S_\uparrow^*$ when $T=T^*$. Then at a temperature $T<T^*$, transitions from up to down are building the tail at $S=S_\uparrow^*-\beta\varepsilon$. If we change the cooling schedule and quench instantaneously from some temperature $T_1<T^*$ to $T_2<T_1$, then there will be no transitions with changes in entropy between $\beta_1\varepsilon$ and $\beta_2\varepsilon$. The tail will therefore have no weight between $S_\uparrow-\beta_2\varepsilon$ and $S_\uparrow-\beta_1\varepsilon$ (see the lower curve in Fig. 12). The size of this "hole" in the tail of the distribution, $(\beta_2-\beta_1)\varepsilon$, determines ε independently of T^*. If the quench is done at a high enough temperature, a hole will appear in the upper peak as well, due to transitions (or the lack thereof) from the lower well to the upper well at temperatures between T_1 and T_2.

We now know the asymmetry ε and the freezing temperature T^* for the TLS. T^* is not an intrinsic quantity, but depends on the cooling rate $\dot{T}$. The unknown intrinsic quantities are the barrier height V and the attempt frequency Γ_0. Equation (13) for T^* can be solved for V in terms of Γ_0:

$$V\approx T^*\ln\left[\frac{\Gamma_0 T^{*2}}{\varepsilon|\dot{T}|n_0(T^*)}\right]-\varepsilon\ . \tag{18}$$

Once we know ε, then, we can get two equations for the two unknowns V and Γ_0 by using (18) with values of T^* obtained at two different cooling rates $\dot{T}$. To be precise, if the cooling rates are $\dot{T}_1$ and $\dot{T}_2$, and the associated freezing temperatures (measured from the distributions) are T_1^* and T_2^*, then

$$V=\frac{T_1^*T_2^*}{T_1^*-T_2^*}\ln\left[\left[\frac{T_2^*}{T_1^*}\right]^2\left|\frac{\dot{T}_1}{\dot{T}_2}\right|\frac{n_0(T_1^*)}{n_0(T_2^*)}\right]-\varepsilon \tag{19}$$

and

$$\Gamma_0=\frac{\varepsilon|\dot{T}_1|n_0(T_1^*)}{T_1^{*2}}\exp[(V+\varepsilon)/T_1^*]\ . \tag{20}$$

Three entropy distributions, from two linear coolings and one interrupted cooling, therefore suffice to determine ε, Γ_0, and V. These determinations will be imprecise because of the assumption that the peaks in the distribution are sharp at T^*, and because T^* itself is not a well defined quantity. A better theory would take into account the entire evolution of the distribution, rather than assuming it begins at T^*.

The preceding analysis can be applied to the distributions in Fig. 12. Actually, distributions were computed for this TLS for four different cooling rates. The values of ε/T^* for the four runs are summarized in the upper part of Table III. The slowest run (run 1 in the table) was repeated, but this time was quenched from $T=0.3V$ to $T=0.25V$, producing the distribution with two holes in Fig. 12. Measuring the width of the lower hole, and using the measured ε/T^* from run 1 gave $\varepsilon=0.5\pm0.1$, in good agreement with reality ($\varepsilon\equiv0.5$). The same calculation using the hole in the upper peak gave $\varepsilon=0.3\pm0.15$. Since this is a pedagogical exercise, the better the value of ε was used for the remaining calculations. The freezing temperatures T^* found from the peak separation for each cooling rate coincide with the actual freezing temperatures, as shown in Fig. 13. Equation (19) was used for each pair of runs to compute the barrier heights displayed in the middle part of Table III. The results are in agreement with the actual barrier ($V\equiv1$), but not spectacularly so. The attempt frequencies from (20) and shown in the lower part of Table III are not so good. Γ_0 in (20) depends exponentially on V, ε, and T^*, so we

TABLE III. Measurements taken from entropy distributions for a TLS with $\varepsilon=0.5$, $V=\Gamma_0=1$. The cooling rates are in units of V per MC step. Each distribution was compiled from 5000 runs of the simulation. The upper part gives the positions of the two peaks in the distribution and the value of ε/T^* derived from them. T^* in the last column was derived assuming $\varepsilon=0.5\pm0.1$ (see the text). The middle part shows the barrier heights computed from Eq. (19) for each pair of runs, and the lower part shows the attempt frequencies derived from (20) for each V in the middle part.

Run No.	Rate	Lower peak	Upper peak	ε/T^*	T^*
1	0.001	0.3±0.1	2.0±0.1	1.7±.2	0.3±0.1
2	0.004	0.3±0.1	1.7±0.1	1.3±.2	0.4±0.1
3	0.01	0.34±0.06	1.4±0.1	1.0±.1	0.5±0.1
4	0.02	0.38±0.06	1.2±0.1	0.8±.1	0.63±0.1

V	2	3	4
1	0.69±2.0	0.73±1.2	0.66±0.9
2		0.82±2.7	0.63±1.2
3			0.41±1.9

Γ_0	2	3	4
1	0.22±1.5	0.63±2.9	1.00±3.4
2		0.19±1.3	0.23±0.8
3			0.07±0.3

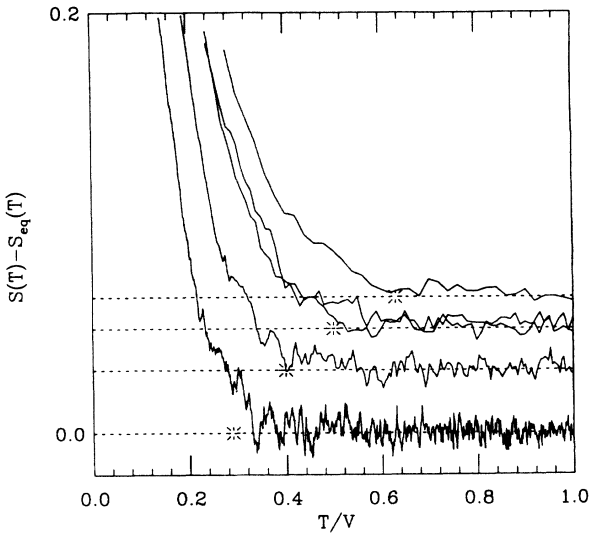

FIG. 13. Deviation of the average measured thermodynamic entropy from its equilibrium value for four runs of the TLS simulation used for the entropy distributions shown in Fig. 12. From top to bottom, the runs were cooled at rates of 0.02, 0.01, 0.004, and 0.001 V per MC step. The curves are displaced vertically for clarity by an amount proportional to the logarithm of the cooling rate. The stars (*) mark the freezing temperatures determined from the entropy distributions. Each curve is an average over 5000 runs—two such runs are superimposed for the second fastest simulation, showing that the apparent T^* can fluctuate greatly.

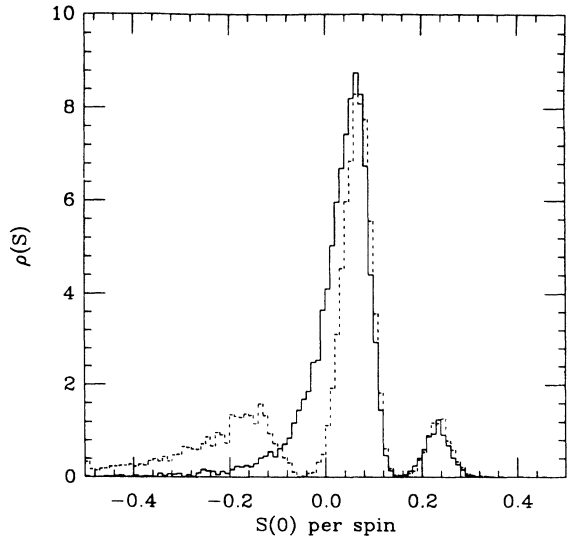

FIG. 14. Two entropy distributions for spin-glass sample 789. The first (solid line) was cooled continuously from $T=10J$ to $T=0$ in 1000 MC steps. The second (dotted line) was cooled at the same rate, but was quenched from $T=0.7J$ to $T=0.4J$. Each distribution was compiled from 10 000 runs of the simulation.

should not expect the results to be great, but the errors are too small. There also seems to be a systematic dependence on the cooling rate, which may arise from the approximate nature of (13).

B. Entropy distributions for 5×5 spin-glass sample 789

We now apply the same analysis to a small spin glass, 5×5 sample 789. We do not know that the spin glass is a two-level system; in fact, it almost certainly is not. On the other hand, it probably has metastable states, and at low temperatures it is possible that it behaves like a TLS. Examination of the entropy distributions should tell us the parameters of this effective TLS.

Two entropy distributions from sample 789 are shown in Fig. 14. For the first (solid line) the simulation was cooled to $T=0$ at a rate of $0.01J$ per MC step. The second (dotted line) was cooled at the same rate, but was quenched from $T=0.7$ to $0.4J$. Qualitatively, the distributions are similar to the TLS distributions in Fig. 12. The distance between the peaks in the first distribution shows that $\varepsilon/T^*=0.15\pm0.03$. The hole caused by the quench in the second distribution shows that $\epsilon=0.12\pm0.03$, and that for this cooling rate $T^*=0.8\pm0.3$. Peak positions and freezing temperatures for three different cooling rates are shown in Table IV. (The range of cooling rates is smaller than for the two-level systems—faster rates than the ones shown tend to blur the peaks together in the distributions, and slower rates leave no weight in the upper peak.) Using Eq. (19) to compute the barrier height for the effective TLS in the spin glass is amusing—the values obtained from the three pairs of runs are $V(1,2)=-0.14\pm0.03$, $V(1,3)=-0.11\pm0.03$, and $V(2,3)=8\pm49$. We can conclude, therefore, that Eq. (13) is not valid for the spin glass. Even if the spin glass does look like a two-level system, its equilibrium population does not behave like (12).

C. Spin-glass microscopy

To check the value of the asymmetry ε and to explain the inconsistent barrier height measurements obtained in the last section, we examine the configuration space of the spin glass. Given a list of the configurations found at the end of each cooling run of the simulation, and the entropies found on each run, we can discover which configurations contribute to which peaks of the entropy distribution. By making zero-energy spin flips, we can

TABLE IV. Same as the upper part of Table III, but for 5×5 spin-glass sample 789. To compute T^*, $\varepsilon=0.12\pm0.03$ was assumed.

Run No.	Rate	Lower peak	Upper peak	ε/T^*	T^*
1	0.001	0.042±0.006	0.27±0.01	0.23±0.01	0.028±0.007
2	0.01	0.08±0.02	0.23±0.02	0.15±0.03	0.8±0.3
3	0.04	0.08±0.04	0.22±0.3	0.14±0.05	0.9±0.4

group the configurations into states, and measure the degeneracy of each state (as defined in Sec. III). If all the spin configurations making up a state are stable with regard to single spin flips, then the state is stable. For example, analysis of the configurations found by the runs that produced the distribution shown in the solid curve in Fig. 14 produced the following list of states:

stable state 13522626 reached 4696 times	(degeneracy 2),
stable state 19999037 reached 4587 times	(degeneracy 2),
stable state 16963034 reached 355 times	(degeneracy 14),
stable state 12397093 reached 367 times	(degeneracy 14),
unstable state 12617282 reached 10 times	(degeneracy 157),
unstable state 19931453 reached 5 times	(degeneracy 157).

The identifying numbers for the states, when written in binary, specify the spins in one of the degenerate configurations (the configuration with the smallest binary encoding). There are two states with each degeneracy—one can be derived from the other by flipping all the spins. For purposes of comparing the spin glass to a TLS, we can ignore the duplicity of states. The energy of the first two stable states, those reached most often by the simulation, is $E=-1.36J$ per spin. The energy of the next two state is $E=-1.20J$ per spin. Since both of these states are stable, they form the TLS we were looking for in Sec. V B. The asymmetry ε is $0.16J$, not in bad agreement with the predicted value of $0.12\pm0.03J$.

The barrier between the two stable states can be mapped, but not in as much detail as the states themselves. There are two configurations of spins, one in the ground state and one in the metastable state, that can be transformed into one another by a sequence of eight spin flips. There are no two such states that can be joined by fewer spin flips. On going from the ground to the metastable state, the first spin flip breaks four bonds, increasing the energy per spin by $0.32J$. (Breaking a bond changes the total energy by $2J$.) The next six spin flips cost no energy. The eighth spin flip satisfies two bonds, decreasing the energy by $0.16J$. Therefore the barrier height V is $0.16J$, measured from the bottom of the upper well.

Two entropy distributions are plotted in Fig. 15. The first (solid line) is from spin-glass sample 789, and is identical to the solid curve in Fig. 14. The second is the best match with a TLS entropy distribution. The asymmetry and barrier height of the TLS were chosen to match the asymmetry and barrier height of the effective TLS found in the spin glass ($\varepsilon=0.16$ and $V=0.16$). The degeneracy of the spin-glass states was accounted for by adjusting the TLS transition rates—the rate for upward transitions was reduced by the degeneracy $d_{\downarrow}=2$ of the lower state, and the rate for downward transitions was reduced by the degeneracy $d_{\uparrow}=14$ of the upper state.[17] The cooling time was adjusted so that the residual population of the TLS was 0.070, which is the population derived from the state counts for the spin glass. The TLS entropies were divided by 25, the number of spins in the spin glass, before the histogram was compiled, and the whole histogram was shifted horizontally by 0.05 to match the positions of the upper peaks in the two distributions. (This shift is legitimate. There are contributions to the entropy of the spin glass from sources other than its effective TLS, so the absolute positions of the peaks are not important.) Despite all the manipulation, the two distributions are quantitatively dissimilar.

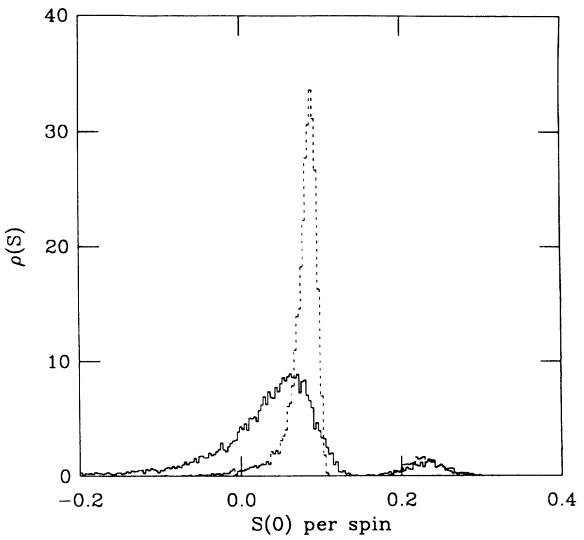

FIG. 15. A TLS residual entropy distribution (dotted line) chosen to match a distribution for spin-glass sample 789 as closely as possible. Both distributions were compiled from 10 000 runs of the corresponding simulations. See the text for the details of the matching procedure.

The differences in the dynamics of the spin glass and the TLS can be deduced (with a little hindsight) from the differences in their entropy distributions. While the upper peaks are almost identical, the lower peak in the spin-glass distribution is much broader than its counterpart in the TLS distribution. This would be the case if transitions *out* of the upper wells of the spin glass and the TLS happen in the same manner, but transitions *into* the lower well of the spin glass are delayed, relative to the corresponding transitions in the TLS. By spending a longer time in transit between the wells, the spin glass allows the temperature to drop farther during the transition, so the entropy released when it falls into the lower well is greater.

What is the spin glass doing when it is between states? The width of the barrier at its narrowest point is eight spin flips, so it is unlikely that the system will get across in a single MC time step (each spin is given one chance to flip in one time step). Furthermore, the top of the barrier actually consists of 5052 degenerate configurations, so it is possible for the system to wander around on top of the barrier for a while. Most of these configurations (3224 of them) are unstable, so it is unlikely that the spin glass will stay on top long. Although a few of these unstable barrier configurations are connected directly to the ground state, many of them are connected to a "plateau" of 157 degenerate configurations. Only 23 configurations on this plateau are unstable (these lead directly to the ground state), so it *is* likely that the spin glass will wander around on the plateau for a long time before finding its way into its ground state. This is the origin of the delayed transitions and the broadened peak in the entropy distribution.

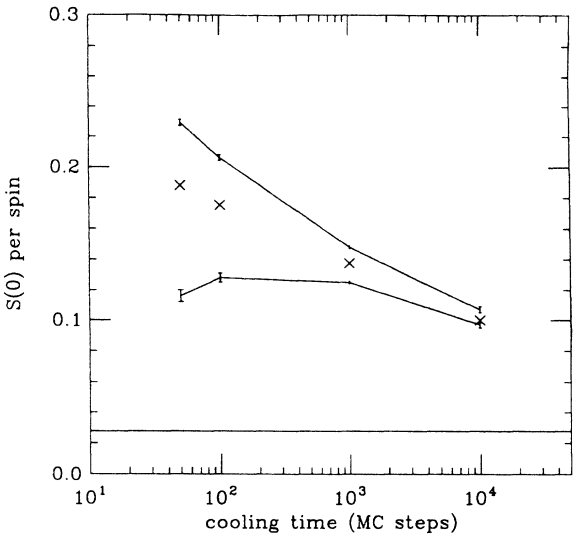

FIG. 16. Zero-temperature thermodynamic entropies (solid lines) and statistical entropies (crosses) for 5×5 Gaussian spin-glass sample 1012. The horizontal solid line is the equilibrium residual entropy, ln2/25.

D. Spin-glasses with Gaussian bonds

The preceding analysis applied to a particular realization of a $5\times5\pm J$ Ising spin glass. The same analysis could be applied to other realizations, and would do as well. It is important, though, to make sure that the structure in the distributions is not an artifact of the pathology of the $\pm J$ model—if there were no degenerate configurations, would the distributions be useful?

Rather than repeating the analysis for a new model, we simply present two plots indicating that a spin glass with Gaussian bonds ($\{J_{ij}\}$ chosen from a Gaussian distribution) has nontrivial thermodynamic entropy. Figure 16 shows the average thermodynamic entropy measured by heating and cooling 1000 copies of a particular realization of a 5×5 Gaussian spin glass. As in the $\pm J$ model in Fig. 5, the thermodynamic entropy bounds the statistical entropy, and is greater than the equilibrium entropy, so the system is out of equilibrium. The measured entropies are farther above the equilibrium entropy than they are for the 5×5 $\pm J$ spin glass in Fig. 5. The Gaussian spin glass has more metastable states (it has no plateaus), so it is harder to equilibrate. The multitude of metastable states leads to a multitude of peaks in the entropy distribution, which overlap and merge for the 5×5 sample, as in Fig. 17(a). The smaller samples shown in the figure have fewer states, and look more like two-level systems.

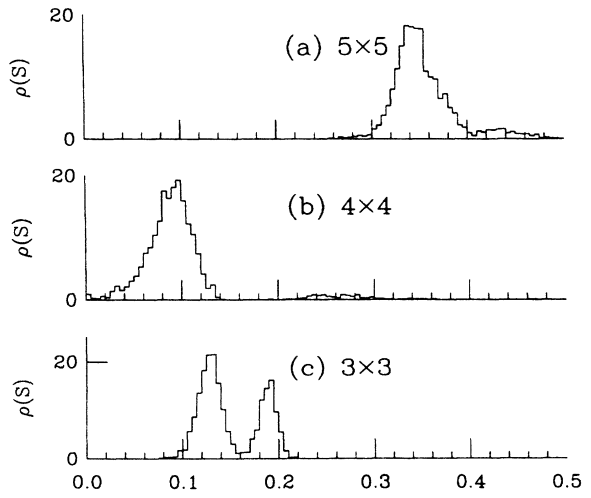

FIG. 17. Entropy distributions for Gaussian spin glasses of three sizes: (a) 5×5 sample 7890, (b) 4×4 sample 4884, and (c) 3×3 sample 5835. The cooling rate for all three was 0.001σ per MC step, where σ is the width of the bond distribution. The mean bond strength was zero.

VI. OTHER MODELS, AND REAL DATA

So far we have discussed only small spin glasses and the toy TLS model, both of which are rather far removed from real glasses. It is important to know whether the concepts introduced earlier will apply to real glasses and/or more realistic models.

A. Real glasses as ensembles of two-level systems

The low-temperature properties of glasses are understood in terms of two-level systems.[18,19] These models postulate that glasses at low temperatures contain a number of sites at which atoms can tunnel between two states, or wells (see Fig. 9). The two states differ in energy by an asymmetry ε. There is a barrier V to tunneling between the states. The thermodynamic properties of the glass are explained in terms of the distribution $f(V,\varepsilon)$ of barrier heights and asymmetries. At higher temperatures, but still below the glass transition, glassy properties ("β relaxations") can be understood in terms of thermal activation over the barriers of the two-level systems.[20–22] The TLS model introduced in Sec. IV describes a single one of these two-level systems.

In real glasses the distribution of barrier heights and asymmetries will have some temperature dependence.[22] We shall model a glass simply as a distribution of TLS's with *fixed* barriers and asymmetries, because we can then use directly the results of Sec. IV. In particular, all of the average properties of this extended TLS model can be determined by numerical integration of the master equation (11). This is a great advantage, because it allows us to do simulations on the time scales of real experiments. A Monte Carlo simulation could tell us how the barrier heights change, but Monte Carlo simulations require on the order of one computation per molecule per oscillation. If the oscillation frequency is phononlike, the simulation requires 10^{12} computations per molecule per simulated second, which obviously forbids simulations of real experiments. The master equation can be integrated on any time scale, simply by changing its parameters.

For the orientational glass $KBr_x:KCN_{1-x}$, the distribution $f(V,\varepsilon)$ of barrier heights and asymmetries is known.[20] In this system, football shaped cyanide molecules randomly replace bromine atoms in the crystal lattice. The cyanides introduce lattice strain[23] and interact

with each other as elastic dipoles. The elastic dipole interaction is symmetric with respect to 180° flips of the cyanides, so each cyanide has two stable orientations. Cyanide molecules are also electric dipoles, so the energies of the two orientations of each molecule are asymmetric. The data and simulations presented below are all for the concentration $x=0.5$. Experimentally, the barrier distribution is Gaussian, peaked at[24] $V=660$ K with a width of 210 K. The asymmetry distribution is assumed to be flat for $0<\varepsilon<350$ K and zero otherwise.[21] If the potential energy of the cyanide is $V\cos\theta$, where θ is its orientation, then the small oscillation frequency $\Gamma_0 \propto \sqrt{V}$. Γ_0 may be taken to be 35 K for $V=660$ K.

Figure 18 shows the entropy bounds for $KBr_{0.5}$:$KCN_{0.5}$ as a function of cooling time, found by numerically integrating the TLS master equation and averaging over the distribution f. The statistical entropy, computed from (14), falls between the upper and lower bounds, but is not shown on the plot. The upper and lower bounds differ by about 1% for experimentally accessible cooling rates. Because this model ignores the structural entropy of the processes that create the two-level systems (see Sec. VI C), the curve should not be interpreted as a prediction of the zero-temperature entropy of $KBr_{0.5}$:$KCN_{0.5}$. It simply suggests that the difference between the upper and lower bounds may be experimentally detectable, and therefore that the distinction between the statistical and thermodynamic entropies is of more than pedagogical interest.

B. Memory effects

The extended TLS model is also useful in illustrating the long-term history dependence of glasses. Figure 19 depicts the energy and heat capacity of the $KBr_{0.5}$:$KCN_{0.5}$ simulation during cooling and heating on a reasonable experimental time scale. Upon cooling, the glass was allowed to anneal at a constant temperature for three (simulated) months. Some of the TLS's in the ensemble were in equilibrium during these months, and did nothing interesting. Some were completely frozen, and also did nothing interesting. Some, on the other hand, continued to relax, producing a spike in the specific heat ($dE\neq 0$ although $dT=0$). More interesting, though, is the behavior of the specific heat when the sample was subsequently heated continuously through the annealing temperature. The TLS's which contributed to the spike on cooling had a lower population than they would have led there been no annealing, so they absorbed more energy when they thawed. The specific heat consequently has a pronounced bump. The glass has memory—its behavior on heating depends upon how it was treated while cooling. This effect is seen in real glasses too. See, for example, Brawer's discussion[25] of sub-T_g relaxation in B_2O_3.

C. High-temperature relaxations

Bounds on the statistical entropy can be found for processes that do not derive from two-level systems. Figure 20 shows the upper and lower bounds on the statistical entropy derived from a Monte Carlo simulation of $KBr_{0.5}$:$KCN_{0.5}$ by Eric Grannan.[23] The entropy plotted is the configurational entropy from the formation of the TLS's, not the entropy of the TLS's themselves. The simulation allows the cyanides to interact through their elastic dipole strain fields, which create two stable orientations for each cyanide, but ignores electric dipole forces, which create asymmetries in the orientations. Therefore, this simulation describes the falling out of

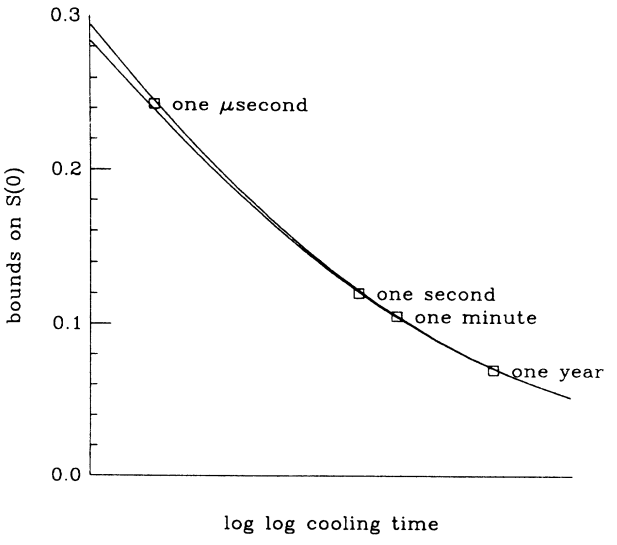

FIG. 18. Upper and lower bounds on the zero-temperature entropy per TLS as a function of cooling time for an ensemble of TLS's chosen to mimic $KBr_{0.5}$:$KCN_{0.5}$. The ensemble was cooled from $T=100$ K to $T=0$ K in the time shown. Note that the plot is log log, in a way.

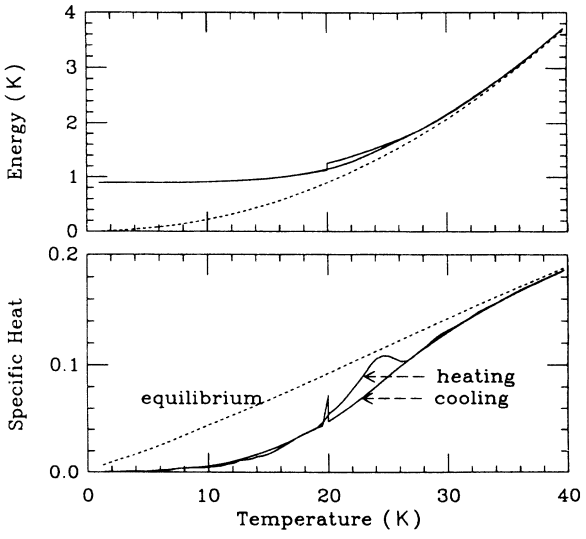

FIG. 19. Energy and specific heat vs temperature for the TLS ensemble mimicking $KBr_{0.5}$:$KCN_{0.5}$. The ensemble was cooled from high temperature to 0 K at a rate of 45 K per hour but along the way was annealed at 20 K for 3 months. Upon reheating at the original rate, a large bump appeared in the specific heat.

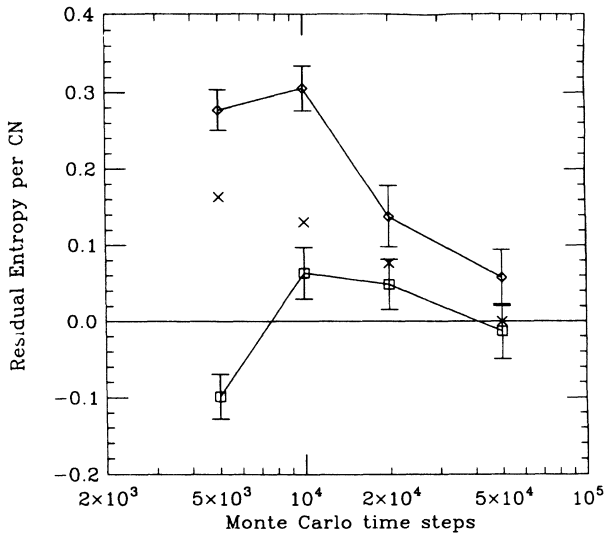

FIG. 20. Bounds on the statistical entropy (crosses) given by thermodynamic entropy on cooling (squares) and heating (diamonds) for a Monte Carlo simulation of $KBr_{0.5}$:$KCN_{0.5}$. The system contained 16 cyanides, and was cooled and heated between 1000 K and 0 K in the specified number of steps. The error bars represent the uncertainty in the mean after averaging over a number of runs. If there were enough data, we could plot entropy distributions as well.

equilibrium that takes place at the glass transition (via "α relaxations"), not the behavior of the TLS's below the transition. The differences in the bounds are much larger than in Fig. 18 because the cooling times are much faster. Appendix B provides details on how to find the zero-temperature entropy of a system with continuous degrees of freedom.

D. Entropy in real glasses

Are the bounds on the entropy really distinguishable in real glasses? The extended TLS model implies that they are. Most authors that publish entropy measurements neglect to provide data on both heating and cooling.[26] An exception is a paper by Thomas and Parks,[27] which has the added benefit of being from 1931, when people still published numbers, rather than just graphs. Thomas and Parks measured the specific heat of B_2O_3 by heating and cooling it between room temperature and 345 °C. The glass transition temperature is about 275 °C. Of the three curves shown in Fig. 1, the last two are the sort of experiment described in Sec. II—cooling through the glass transition to a low temperature, followed by heating back to the original state. Integrating Thomas' and Park's specific-heat data from $T=112$ to 345 °C (the range covered by the two runs) gives a change in entropy per molecule of 0.869 on cooling, and of 0.858 on heating. Assuming that the material is in equilibrium at the higher temperature, we see that the entropy per molecule at the lower temperature measured on cooling is 0.012 below the entropy measured on heating, in accordance with (7). Unfortunately we do not know the entropy at 345 °C, so we cannot determine the actual bonds on the statistical entropy.[28] Assuming that the residual entropy is approximately one per molecule, though, we see that the bounds on the statistical entropy differ by about 1%, as predicted.

What about entropy distributions? Can entropy distributions be measured in the laboratory? Probably not. The entropy distribution gets its structure from the fluctuations that occur during the cooling and heating processes. A macroscopic system is made up of many microscopic systems. Although each microsystem may have an interesting entropy distribution, only the macroscopic heat flow can be measured. Therefore, only the collective mean of the microscopic entropy distributions can be measured. We are forced to conclude that although the entropy distribution is a useful probe of the dynamics of glassy computer simulations, it will be of little use in real glasses.

VII. FUTURE WORK

A. Susceptibility to the cooling schedule

The methods discussed in Sec. V for extracting information from the entropy distribution rely upon making large changes in the cooling schedule—either changing the rate or quenching. The response of the distribution to infinitesimal changes in the cooling schedule might provide a more systematic method of extracting the same information. For slow cooling rates, we can think of the imposed temperature change as a sort of field that is pushing the TLS out of equilibrium, and the tails of the entropy distributions as the linear response to this field.

B. Asymptotics for more complicated models

The TLS model describes the behavior of glasses at temperatures below the glass transition, but does not describe the transition itself. Ultimately, we would like to find a model with a real glass transition (a 3D spin glass?) and solve it asymptotically. A more modest calculation would consider generalizations of TLS model which may reproduce glassy properties ignored in the current theory, such as the plateau in the thermal conductivity, nonexponential relaxation, and the Kauzmann paradox. Possible generalizations are an n-level system, an interacting network of two-level systems, a hierarchical model,[30] a mean-field spin glass (the infinite-range SK model or the Bethe lattice spin glass[30]) or simply a TLS with a plateau.

ACKNOWLEDGMENTS

We would like to thank A. Dorsey for his assistance and advice. This work was supported by National Science Foundation (NSF) Grant Nos. DMR-8451921 and DMR-8217227A-01 through the Materials Science Center at Cornell University. The dipole glass simulations were done using the Cornell National Supercomputer Facility, a resource of the Center for Theory and Simulation in Science and Engineering at Cornell University, which is funded in part by the NSF, the State of New York, and the IBM Corporation.

APPENDIX A: MASTER EQUATIONS AND STATISTICAL ENTROPY

Many systems in statistical mechanics are described by master equations. We wish to show that for these systems the statistical entropy defined in Sec. II

$$S_{\text{stat}} = -\text{Tr}\hat{\rho}\ln\hat{\rho} , \tag{2}$$

with $\hat{\rho}_i$ defined to be the occupation probability of state i, given a certain thermal history—satisfies the criteria set out in Theorem 1: (1) S_{stat} is extensive, (2) $S_{\text{stat}} = S_{\text{therm}}$ in equilibrium, and (3) S_{stat} increases with time in a closed system.

Consider a system of N states with energies E_i and populations $p_i (i = 1, \ldots N)$. If the transition rate from state j to state i is γ_{ij}, then the master equations for the populations are

$$\frac{dp_i}{dt} = \sum_j \gamma_{ij} p_j . \tag{21}$$

Detailed balance implies

$$\gamma_{ij} e^{-\beta E_j} = \gamma_{ji} e^{-\beta E_i} , \tag{22}$$

while conservation of probability implies

$$\gamma_{ii} = -\sum_{j \neq i} \gamma_{ji} , \tag{23}$$

and normalization implies

$$\sum_i p_i = 1 . \tag{24}$$

The statistical entropy is

$$S_{\text{stat}} = -\sum_i p_i \ln p_i . \tag{25}$$

S_{stat} is obviously extensive. When p_i is the equilibrium population of state i, (25) is exactly the equilibrium statistical entropy, so Criterion 2 is satisfied.

To prove Criterion 3, we need to close the system by introducing a heat bath. The bath is in equilibrium, so its thermodynamic entropy equals its statistical entropy. The entropy of the bath changes only through heat exchanged with the N state system:

$$S_{\text{bath}} = -\sum_i \int \frac{E_i dp_i}{T} .$$

We have

$$\frac{dS_{\text{stat}}}{dt} = -\sum_i \left[\frac{dp_i}{dt} + \ln p_i \frac{dp_i}{dt} \right] = -\sum_i \ln p_i \frac{dp_i}{dt}$$

[using (24)] and

$$\frac{dS_{\text{bath}}}{dt} = -\sum_i \beta E_i \frac{dp_i}{dt} .$$

Letting $S_{\text{tot}} = S_{\text{stat}} + S_{\text{bath}}$ and using (21) and (23) gives

$$\begin{aligned}\frac{dS_{\text{tot}}}{dt} &= -\sum_i (\ln p_i + \beta E_i) \frac{dp_i}{dt} \\ &= -\sum_i (\ln p_i + \beta E_i) \sum_j \gamma_{ij} p_j \\ &= -\sum_i (\ln p_i + \beta E_i) \left[\sum_{j \neq i} \gamma_{ij} p_j - \left[\sum_{j \neq i} \gamma_{ji} \right] p_i \right] .\end{aligned}$$

Defining $w_i \equiv e^{\beta E_i} p_i$, and using (22) to convert γ_{ji} to γ_{ij}, we have

$$\frac{dS_{\text{tot}}}{dt} = -\sum_i \sum_{j \neq i} \gamma_{ij} e^{-\beta E_j} \ln w_i (w_j - w_i) ,$$

which, when written as a sum over all pairs of states $\langle ij \rangle$, is clearly positive:

$$\frac{dS_{\text{tot}}}{dt} = 2 \sum_{\langle ij \rangle} \gamma_{ij} e^{-\beta E_j} (\ln w_i - \ln w_j)(w_i - w_j) .$$

Q.E.D. A similar proof, but for a closed system without a heat bath, is given by van Kampen.[31]

APPENDIX B: RESIDUAL ENTROPY OF A CLASSICAL SYSTEM

The thermodynamic entropy (1) of a classical system with continuous degrees of freedom diverges to minus infinity as $T \to 0$, because $dQ = C\,dT$ and the heat capacity C for such a system is constant. The statistical entropy (2) also diverges as $T \to 0$ because the system sits at the very bottoms of the wells in configuration space, and these minima have zero volume. In order to compare the statistical and thermodynamic entropies we have to subtract off the divergences. This method was used to produce Fig. 20.

We are interested in part of the entropy that arises from the multiplicity of metastable zero-temperature states. Once the system is trapped in a single metastable state, at a low enough temperature it behaves more or less like a set of harmonic oscillators. The divergence of the thermodynamic entropy for the full system is the same as the divergence of entropy for the harmonic oscillators; the discrepancy between the full thermodynamic entropy and the oscillator thermodynamic entropy, as $T \to 0$, is the sought after residual entropy. (The residual statistical entropy is found simply by counting the number of wells in configuration space weighted by the probability of finding them, as in Sec. V C.) Because the kinetic energy factors out of the thermodynamics, for our purposes the Hamiltonian can be written $H = \sum_i \frac{1}{2} k_i x_i^2$, where the sum is over the normal modes of the system. The equilibrium entropy is then

$$S_{\text{cl}} = \frac{1}{2} \sum_i \left[1 + \ln \left[\frac{2\pi T}{k_i} \right] \right] . \tag{26}$$

The Monte Carlo simulation of $KBr_{0.5}$:$KCN_{0.5}$ described in Sec. VI C is a classical simulation. The orientations of the dipoles are continuously variable. The measured entropy S_{therm} and energy E are found by integrating dQ/T and dQ while cooling from the initial

high temperature to some small but nonzero temperature T'. The initial value of $S_{\rm therm}$ is determined from a high-temperature expansion,[32] and T' must be low enough that the system is trapped in a single well in configuration space. The normal modes of this well are found, and the classical diverging equilibrium entropy (26) is found. The system is forced to make a last transition so that its energy equals the equilibrium energy $E_{\rm cl}$ of the normal modes at T' (in the same spirit as the initial jump away from equilibrium at high temperature—see Sec. III A). The residual entropy, plotted in Fig. 20 is then

$$S_{\rm therm}(0)=S_{\rm therm}(T')+\frac{E_{\rm cl}(T')-E(T')}{T'}-S_{\rm cl}(T') .$$

---

*Present address: James Franck Institute, University of Chicago, 5640 S. Ellis Ave., Chicago, IL 60637.

†Present address: AT&T Bell Laboratories, Murray Hill, NJ 07974.

[1]G. P. Johari, Philos. Mag. B **41**, 41 (1980).

[2]G. E. Gibson and W. F. Giauque, J. Am. Chem. Soc. **45**, 93 (1923).

[3]J. Jäckle, Philos. Mag. B **44**, 533 (1981).

[4]L. D. Landau and E. M. Lifschitz, *Statistical Physics* (Addison Wesley, New York, 1969).

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

[6]P. Sulewski (private communication).

[7]S. A. Langer and J. P. Sethna, Phys. Rev. Lett. **61**, 570 (1988).

[8]J. Jäckle and W. Kinsel, J. Phys. A **16**, L163 (1983).

[9]R. Ettelaie and M. A. Moore, J. Phys. Lett. **46**, L-893 (1985).

[10]J. Jäckle, Physica **127B**, 79 (1984).

[11]J. M. Gordon, J. H. Gibbs, and P. D. Fleming, J. Chem. Phys. **65**, 2771 (1976).

[12]L. V. Woodcock, Ann. (N.Y.) Acad. Sci. **371**, 274 (1981).

[13]J. N. Cape and L. V. Woodcock, J. Chem. Phys. **72**, 976 (1980).

[14]C. Domb, in *Phase Transitions and Critical Phenomena,* edited by C. Domb and M. S. Green (Academic, New York, 1974), Vol. 3, p. 375.

[15]D. A. Huse and D. S. Fisher, Phys. Rev. Lett. **57**, 2203 (1986).

[16]S. A. Langer, A. T. Dorsey, and J. P. Sethna, Phys. Rev. B **40**, 345 (1989).

[17]If the degeneracies of the upper and lower states of the TLS are $d_\uparrow$ and $d_\downarrow$, respectively, then the equilibrium population is

$$n_0(T)=\frac{d_\uparrow e^{-\beta\epsilon}}{d_\downarrow+d_\uparrow e^{-\beta\epsilon}}$$

and the equilibrium entropies are [compare to (16)]

$$S_\uparrow=\ln(d_\downarrow+d_\uparrow e^{-\beta\epsilon})+\beta\epsilon ,$$

$$S_\downarrow=\ln(d_\downarrow+d_\uparrow e^{-\beta\epsilon}) .$$

Multiplying the nondegenerate TLS transition rates by factors of $1/d_\uparrow$ and $1/d_\downarrow$ guarantees the correct equilibrium behavior.

[18]P. Anderson, B. Halperin, and C. Varma, Philos. Mag. **25**, 1 (1972).

[19]W. Phillips, J. Low Temp. Phys. **7**, 351 (1972).

[20]J. P. Sethna and K. S. Chow, *Phase Transitions* **5**, 317 (1985).

[21]M. Meissner, W. Knaak, J. P. Sethna, K. S. Chow, J. J. De Yoreo, and R. O. Pohl, Phys. Rev. B **32**, 6091 (1985).

[22]N. O. Birge, Y. H. Jeong, S. R. Nagel, S. Bhattacharya, and S. Susman, Phys. Rev. B **30**, 2306 (1984).

[23]For details of the model, see E. R. Grannan, M. Randeria, and J. P. Sethna, Phys. Rev. Lett. **60**, 1402 (1988), and (unpublished).

[24]$35\text{K}=4.4\times10^{12}$ Hz. The numbers reported here are those used in the simulation. We were not attempting to duplicate detailed properties of $KBr_{0.5}$:$KCN_{0.5}$, and the parameters are all rounded to two significant figures.

[25]S. Brewer, *Relaxation in Viscous Liquids and Glasses* (The American Ceramic Society, Columbus, 1985), p. 93.

[26]Ref. 1, for example.

[27]S. B. Thomas and G. S Parks, J. Phys. Chem. **35**, 2091 (1931).

[28]Thomas and Parks say that their specific-heat data have a probable absolute error of less than 4%, and an reproducibility of 0.5%. This gives us some faith in the 1% difference in the bounds, but little faith in any absolute statement about the residual entropy.

[29]R. G. Palmer, D. L. Stein, E. Abrahams, and P. W. Anderson, Phys. Rev. Lett. **53**, 958 (1984).

[30]J. M. Carlson, Ph.D. thesis, Cornell Univeristy (1988).

[31]N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, New York, 1981), p. 118.

[32]At infinite temperature $S=\ln2\pi$ per dipole. The heads and tails of elastic dipoles are identical, so the phase space of each dipole is a hemisphere.