

Biased landscapes for random Constraint Satisfaction Problems

Louise Budzynski,¹ Federico Ricci-Tersenghi,² and Guilhem Semerjian¹

¹*Laboratoire de physique théorique de l'École normale supérieure,
PSL Research University, CNRS, Sorbonne Universités,
24 rue Lhomond, 75231 Paris Cedex 05, France*

²*Dipartimento di Fisica, Sapienza Università di Roma, and Nanotec-CNR,
UOS di Roma, and INFN-Sezione di Roma 1, P.le Aldo Moro 5, 00185 Roma Italy*

The typical complexity of Constraint Satisfaction Problems (CSPs) can be investigated by means of random ensembles of instances. The latter exhibit many threshold phenomena besides their satisfiability phase transition, in particular a clustering or dynamic phase transition (related to the tree reconstruction problem) at which their typical solutions shatter into disconnected components. In this paper we study the evolution of this phenomenon under a bias that breaks the uniformity among solutions of one CSP instance, concentrating on the bicoloring of k -uniform random hypergraphs. We show that for small k the clustering transition can be delayed in this way to higher density of constraints, and that this strategy has a positive impact on the performances of Simulated Annealing algorithms. We characterize the modest gain that can be expected in the large k limit from the simple implementation of the biasing idea studied here. This paper contains also a contribution of a more methodological nature, made of a review and extension of the methods to determine numerically the discontinuous dynamic transition threshold.

PACS **11.10**

Contents

I. Introduction	2
II. Definition of the model and statistical mechanics formalism	4
A. Definition of the model	4
B. BP equations and Bethe free-energy	5
C. Replica symmetric cavity method	6
D. 1RSB formalism	7
1. 1RSB cavity equations	7
2. Simplifications for $m = 1$	9
3. The local instability of the RS solution (Kesten-Stigum bound)	11
4. The presence of hard-fields in the 1RSB solution with $m = 1$ (rigidity threshold)	12
III. On the numerical determination of the dynamic transition	13
A. Scalar bifurcations	13
B. Discontinuous functional bifurcations	14
C. The stability parameter λ in the functional case	16
IV. Results of the cavity method	18
A. The existence of a RS phase for $\alpha > \alpha_{d,u}$	18
B. More detailed zero temperature phase diagrams	20
V. Results of simulated annealing	22
A. Estimating the algorithmic threshold for Simulated Annealing	23
B. Performances of Simulated Annealing with optimal RS parameters	25
C. Performances of Simulated Annealing with the biased measure	25
VI. Large k asymptotics	27
VII. Discussion	29
Acknowledgments	30

I. INTRODUCTION

In Constraint Satisfaction Problems (CSPs) a set of N discrete-valued variables are subjected to M constraints; the decision version of the problem consists in answering yes or no to the question “is there an assignment of the variables that satisfies all the constraints simultaneously ?” Computational complexity theory [1, 2] classifies the difficulty of these problems according to the existence of efficient (polynomial time) algorithms able to solve all their possible instances.

Besides this worst-case point of view an important effort has been devoted to the characterization of the “typical” difficulty of CSPs, where typical is defined with respect to a random ensemble of instances, a property being considered typical if it occurs with a probability going to one in the thermodynamic (large size) limit. Many random CSPs have been studied, notably k -SAT and q -COL; in this paper we will focus on the bicoloring of random hypergraphs (related to k -NAESAT), in which the N variables can each take two values (colors), each of the M constraints is generated by choosing uniformly at random a k -uplet of distinct variables, and impose that both colors appear in the configuration of these k variables (i.e. they forbid monochromatic neighborhoods). The interactions induced by these constraints have thus the structure of an Erdős-Rényi random hypergraph, and the thermodynamic limit corresponds to $N, M \rightarrow \infty$ with fixed ratio $\alpha = M/N$ and arity parameter k . The random bicoloring problem exhibits the same rich phenomenology as k -SAT and q -COL, while being slightly simpler from a technical point of view.

Random CSPs exhibit threshold phenomena in the large size limit, the probability of some properties jumping abruptly from 1 to 0 as a function of the control parameter α . The most prominent of these phase transitions occur at the satisfiability threshold $\alpha_{\text{sat}}(k)$: for $\alpha < \alpha_{\text{sat}}(k)$ typical instances are satisfiable, i.e. admit configurations of variables that satisfy all constraints simultaneously, while for $\alpha > \alpha_{\text{sat}}(k)$ a random instance is typically unsatisfiable.

Random CSPs bear a formal similarity with mean-field spin-glasses, the interactions induced by the constraints being of a frustrating nature while lacking a finite-dimensional structure thanks to the randomness in their construction. This analogy has been exploited in depth through the application of methods first developed in the context of statistical mechanics of disordered systems, namely the replica and cavity method, to random CSPs [3–7]. This line of study has provided predictions of $\alpha_{\text{sat}}(k)$ for many models, but also unveiled many other phase transitions for the structure of the set of solutions in the satisfiable phase, and led to the proposal of new algorithms that exploit this detailed picture of the solution space. Many of these (heuristic) predictions have been confirmed rigorously later on [8–11].

In this paper we will pay a particular attention to the dynamic phase transition that occurs at some threshold $\alpha_{\text{d}}(k) < \alpha_{\text{sat}}(k)$, which is also known as the clustering and reconstruction transition. This transition can indeed be described from various perspectives; looking at the set of solutions of typical instances, $\alpha_{\text{d}}(k)$ separates a regime where this set is rather well-connected, any solution can be reached from any other one via nearby intermediate solutions, while for $\alpha > \alpha_{\text{d}}(k)$ the solutions break apart into distinct clusters (or pure states) which are internally well-connected, but separated one from the other by regions of the configuration space void from solutions. This transition marks also the birth of a specific type of long-range correlations between variables, known as point-to-set correlations, which implies the solvability of an information-theoretic problem called tree reconstruction [12]. These correlations forbid in turn the rapid equilibration of the stochastic processes that respect the detailed balance condition [13], hence the name dynamic given to $\alpha_{\text{d}}(k)$. As a matter of fact the static properties of the model are smooth at $\alpha_{\text{d}}(k)$ and are only sensitive to a further condensation transition $\alpha_{\text{c}}(k)$ that affects the number of dominant clusters [7].

Despite the rather detailed picture of the set of solutions of random CSPs sketched above, many questions remain open, in particular concerning the behavior of algorithms that attempts at finding such solutions. These algorithms can be of very different types, proceeding through a local search in the space of configurations [14–17], or assigning variables one by one, according to either simple heuristics [18–21] or detailed information provided by message passing algorithms (Belief or Survey Propagation) inspired by statistical mechanics considerations [5, 22–25]. These dynamics are most of the time “out-of-equilibrium” processes, either because their definition breaks explicitly the detailed balance (reversibility) conditions, or because they will not be able to remain equilibrated during their evolution on reasonably accessible time scales. The great diversity of these algorithms and their out-of-equilibrium character makes very challenging the attempts to characterize the putative algorithmic barrier $\alpha_{\text{alg}}(k)$ above which no algorithm is able to find a solution in polynomial time for a typical random instance (assuming of course $P \neq NP$), and to relate it to the structural phase transitions undergone by the set of solutions. For small values of k numerical experiments suggest that carefully designed local search algorithms [14, 16, 17] or Survey Propagation implementations [23] work (i.e. find solutions in polynomial time) up to densities very close to the satisfiability threshold, thus setting lowerbounds on $\alpha_{\text{alg}}(k)$ almost coinciding with the upperbound $\alpha_{\text{sat}}(k)$. The situation is quite different in the large k limit, that allows for some analytical simplifications. Let us recall that in this limit the satisfiability threshold occurs at $\alpha_{\text{sat}}(k) \sim 2^{k-1} \ln 2$, while the asymptotic expansion of the clustering threshold is $\alpha_{\text{d}}(k) \sim 2^{k-1} (\ln k)/k$ (the

quantitative statements are made here for the hypergraph bicoloring problem, but the qualitative picture is the same for many random CSPs). Numerical experiments cannot access directly the large k limit, but simple enough algorithms can be studied analytically for all k ; sequential assignment algorithms that use simple heuristics to guide their choices can be described in terms of differential equations [18–21] and shown to work in polynomial time up to densities of the order $\alpha_{\text{sat}}(k)/k$, with a constant prefactor depending on the heuristic chosen. This scaling was improved in [26], where an algorithm was shown to work up to densities of constraints coinciding at leading order with $\alpha_{\text{d}}(k)$. This leaves a multiplicative gap of order k (neglecting the sub-dominant logarithmic correction) from the satisfiability transition, hence a wide range of parameters α for which typical instances are known to have solutions, yet no provably efficient algorithm is known at present to find them. Some negative results have also been obtained, [27] proved that no “local” algorithm (in a precise sense) can find solutions in polynomial time for densities α larger than $\alpha_{\text{d}}(k) \ln k$ (asymptotically at large k), see also [24, 25, 28–31] for other analysis of some specific algorithms.

A further structural property of the set of solutions of a CSP has been studied under the name of “frozen variables”. Roughly speaking, a frozen variable is a variable that takes the same value in all the solutions of a cluster [32, 33], or equivalently a variable whose flip from a solution to another one requires rearranging an extensive number of other variables [34, 35]. Two distinct phase transitions can then be defined: the so-called rigidity transition α_{r} marks the appearance of a positive fraction of frozen variables in typical solutions of a random CSP, while above the freezing transition α_{f} all solutions have this property. In the intermediate regime $[\alpha_{\text{r}}, \alpha_{\text{f}}]$ unfrozen solutions still exist but are exponentially less numerous than the typical, frozen ones. The determination of α_{r} is relatively easy, as it concerns a property of the typical solutions, and in the large k limit the rigidity threshold α_{r} is very close to the dynamic one α_{d} [10, 36–39]. On the contrary the freezing transition at α_{f} is intrinsically a large deviation result, and its determination is quite challenging even for heuristic statistical mechanics methods. It has been predicted in [40] to occur close to the satisfiability threshold for large k (more precisely at $\alpha_{\text{f}} \sim \alpha_{\text{sat}}/2$), in line with the rigorous result $\alpha_{\text{f}} \leq (4/5)\alpha_{\text{sat}}$ from [8, 41].

Frozen variables induce a very strong form of correlations, and it seems impossible to construct in polynomial time a solution containing an extensive number of frozen variables, that must all be set in a consistent way (except in the very special case of XORSAT due to the linear structure of its constraints) [32, 42]. As a matter of fact the best performing solving algorithms in the range $\alpha \in [\alpha_{\text{r}}, \alpha_{\text{f}}]$ do not find typical solutions with frozen variables, but atypical solution without frozen variables [23, 43–45]. It is thus natural to conjecture that α_{f} is an upperbound on α_{alg} , but this still leaves a very wide gap between the threshold of the best known algorithms and α_{f} .

In this paper we study probability measures over the set of solutions of random CSPs, for which not all solutions are equally probable but have a weight depending on their tendency to form frozen variables. The same perspective has been used in a few recent works [40, 46–48], namely to consider a biased probability measure over the set of solutions of a random CSP (in [46, 47] the local entropy, or density of solutions in configuration space, is used to weight differently the solutions, in [40] this role is played by the number of frozen variables, while in [48] hard sphere particles are considered as a CSP, with a bias due to an additional pairwise interaction between particles). Indeed the structural phase transitions mentioned above concern the uniform measure over solutions, and it has been demonstrated that the threshold for properties that are typical in the uniform ensemble (in particular the existence of frozen variables) can be significantly moved by an appropriate bias [40] (biased measures with weights depending on the number of satisfying literals in a clause were also studied in [49] but with the different goal of allowing for a quiet planting of solutions in k -SAT instances). This opens some hope to diminish the algorithmic gap, by giving more weight to solutions that are “easier” to find, for instance because they contain less frozen variables, and to turn atypical properties of the uniform measure into typical ones of the biased measure.

We will concentrate in particular on the increase of the dynamic threshold α_{d} that results from a well-chosen bias between solutions. The algorithmic motivation for the study of this threshold comes from the Simulated Annealing [50] procedure: below α_{d} a Markov Chain reversible with respect to a finite-temperature probability distribution should be able to equilibrate in polynomial time, hence to find solutions (non-uniformly) once the temperature is lowered slowly enough (if there is no reentrance in temperature). We shall implement here a relatively simple version of this idea, introducing soft interactions between the variables inside a constraint of the original CSP. We will demonstrate that for small values of k this allows indeed to increase the dynamic threshold α_{d} , and check that this effect improves the performance of Simulated Annealing. On a more negative side we shall see that this simple implementation of the idea is not powerful enough to modify the leading order of the large k algorithmic gap, but motivates the study of more elaborate biasing strategies.

The rest of the paper is organized as follows. In Section II we define more precisely the model under study and present the equations that describe its behavior in the framework of the cavity method from statistical mechanics. Section III is of a more methodological nature, and contains a review and extension of the numerical procedures to determine accurately the dynamic threshold in models exhibiting a discontinuous 1RSB transition. This Section can be skipped by a reader mostly interested in the results we obtained, which are presented in Sec. IV for what concerns the phase diagrams predicted by the cavity method, in particular the increase of the dynamic threshold of the biased

measure with respect to the uniform one, and in Sec. V for numerical simulations on finite size samples via Simulated Annealing. We study in Section VI the limit of large k and derive an asymptotic upperbound on the possible increase of the dynamic threshold, before presenting our conclusions and perspectives in Sec. VII.

II. DEFINITION OF THE MODEL AND STATISTICAL MECHANICS FORMALISM

A. Definition of the model

We will consider in this paper the k -uniform hypergraph bicoloring problem (related to the k -NAESAT problem) [9, 40, 45, 51–56]. An instance of this constraint satisfaction problem (CSP) is defined by an hypergraph $G = (V, E)$ where V is a set of N vertices, and E a set of M hyperedges, each of them containing k vertices (see Fig. 1 for a representation of G as a factor graph). We shall denote ∂a the set of vertices contained in the a -th hyperedge, and similarly ∂i the set of hyperedges adjacent to the i -th vertex. The variables of this CSP are N Boolean variables, represented as Ising spins $\sigma_i \in \{-1, 1\}$, living on the vertices of G . We will denote $\underline{\sigma} = (\sigma_1, \dots, \sigma_N)$ the global configuration of the variables, and $\underline{\sigma}_S = \{\sigma_i\}_{i \in S}$ the configuration of the variables in a subset S of the vertices. A constraint (or clause) is associated to each hyperedge $a \in E$; the a -th constraint is satisfied by the configuration $\underline{\sigma}$ if and only if there is at least one $+1$ and one -1 among the k variables of $\underline{\sigma}_{\partial a}$, in such a way that the edge is not monochromatic (i.e. that not all variables adjacent to it are equal). A configuration $\underline{\sigma}$ is called a solution of the CSP if it satisfies the M constraints simultaneously.




FIG. 1: Example of a factor graph with $k = 3$, $N = 8$, $M = 4$: vertices are represented by circles, and hyperedges by squares. An edge is drawn between the a -th hyperedge and the i -th vertex if and only if $i \in \partial a$.

A convenient way to study the set of solutions $S(G)$ of a given instance (assuming it is non-empty) is to consider the uniform probability measure over the solutions,

$$\mu(\underline{\sigma}) = \frac{1}{Z(G)} \prod_{a=1}^M \omega(\underline{\sigma}_{\partial a}), \quad (1)$$

where the normalization factor $Z(G)$ (also called partition function) counts the number of solutions, and $\omega(\sigma_1, \dots, \sigma_k)$ is the indicator function of the event “the k variables $\sigma_1, \dots, \sigma_k$ are not all equal”; the a -th term in this product is thus equal to 1 if the a -th constraint is satisfied, and to 0 otherwise.

We shall actually study in this paper a measure of the form (1), but with a more generic form for the function $\omega(\sigma_1, \dots, \sigma_k)$. We will assume that ω is invariant under all permutation of its k arguments; as the latter are binary variables ω can only depend on the number p of -1 among its arguments, and we will denote $\omega_p \geq 0$ the value it then assumes. This translates into the formula

$$\omega(\sigma_1, \dots, \sigma_k) = \omega_p \quad \text{if} \quad \sum_{i=1}^k \sigma_i = k - 2p. \quad (2)$$

The uniform measure over the solutions of the bicoloring problem is recovered for the choice $\omega_0 = \omega_k = 0$, $\omega_1 = \dots = \omega_{k-1} = 1$. If one chooses instead ω_p to depend on $p \in \{1, \dots, k-1\}$, while keeping $\omega_0 = \omega_k = 0$, one obtains a probability measure μ that is still supported solely on the proper bicolorings of G , but is not uniform anymore. As explained in the introduction our goal in this paper is to explore the properties of μ that arises from this bias

between solutions of the CSP. We will sometimes relax the constraint $\omega_0 = \omega_k = 0$, to model the effect of a positive temperature that allows some constraints to be violated. We will in any case always assume that $\omega_p = \omega_{k-p}$: this ensures that the global spin-flip symmetry $\mu(-\underline{\sigma}) = \mu(\underline{\sigma})$ (which is indeed a property of the set of proper bicolorings) is preserved.

We shall characterize the properties of $\mu(\underline{\sigma})$ for “typical” hypergraphs, by studying random instances; even if most of the approach can be generalized to more generic random ensembles, for concreteness we will concentrate on Erdős-Rényi (ER) random hypergraphs. An instance of this ensemble is generated by drawing, independently for each of its hyperedges $a = 1, \dots, M$, the adjacent vertices ∂a as a uniformly random k -uplet among the $\binom{N}{k}$ possible ones. We will be interested in the large size (thermodynamic) limit, in which both N and M go to infinity at a fixed ratio $\alpha = M/N$. We recall that in this limit such random hypergraphs converge locally to hypertrees: the neighborhood within a fixed distance around a uniformly chosen vertex is, with a probability going to 1 in the large size limit, acyclic. More precisely, the local limit tree is a Galton-Watson branching process, in which the law p_d for the degree $|\partial i|$ of a uniformly chosen vertex i is the Poisson distribution of average αk . Thanks to the properties of the Poisson law this is also the probability that a uniformly chosen vertex among an uniformly chosen hyperedge has degree $d + 1$ (i.e. the offspring probability in the Galton-Watson tree is also given by p_d).

B. BP equations and Bethe free-energy

In order to determine the typical properties of the measure $\mu(\underline{\sigma})$, and of the free-entropy density $(\ln Z)/N$, we shall exploit the cavity method [57–59], a formalism particularly efficient for interacting particle models on sparse random structures. As recalled above these structures are locally tree-like; the first step of the cavity method amounts in consequence to study such models on finite trees. In that case one can exploit the recursive nature of trees to derive an exact description of $\mu(\underline{\sigma})$ in terms of its marginals, from which also $\ln Z$ can be expressed.

More precisely, let us introduce the messages $\eta_{i \rightarrow a}$ and $\hat{\eta}_{a \rightarrow i}$ on each edge (i, a) of the factor graph, that are the marginal probability laws of σ_i in amputated graphs where some interactions are discarded; $\eta_{i \rightarrow a}$ is the marginal of σ_i in the factor graph where one removes the hyperedge a , and $\hat{\eta}_{a \rightarrow i}$ is the marginal of σ_i in the factor graph where one removes all the hyperedges in $\partial i \setminus a$. Removing an interaction in a tree breaks it into independent subtrees, which allows to write recursive equations between these messages:

$$\eta_{i \rightarrow a}(\sigma_i) = \frac{1}{z_0^{ia}} \prod_{b \in \partial i \setminus a} \hat{\eta}_{b \rightarrow i}(\sigma_i), \quad (3)$$

$$\hat{\eta}_{a \rightarrow i}(\sigma_i) = \frac{1}{\hat{z}_0^{ai}} \sum_{\underline{\sigma}_{\partial a \setminus i}} \omega(\underline{\sigma}_{\partial a}) \prod_{j \in \partial a \setminus i} \eta_{j \rightarrow a}(\sigma_j), \quad (4)$$

where the constants z_0^{ia} and \hat{z}_0^{ai} are normalizing factors. These equations are valid for any (discrete) domain of the spins σ_i ; as we are studying the case where $\sigma_i = \pm 1$, we can parametrize the probability laws $\eta_{i \rightarrow a}$ and $\hat{\eta}_{a \rightarrow i}$ by their mean values, defining $\eta_{i \rightarrow a}(\sigma_i) = (1 + h_{i \rightarrow a} \sigma_i)/2$ and $\hat{\eta}_{a \rightarrow i}(\sigma_i) = (1 + u_{a \rightarrow i} \sigma_i)/2$, with $h_{i \rightarrow a}, u_{a \rightarrow i} \in [-1, 1]$. The recursive equations can be rewritten with this parametrization as

$$h_{i \rightarrow a} = f(\{u_{b \rightarrow i}\}_{b \in \partial i \setminus a}), \quad u_{a \rightarrow i} = g(\{h_{j \rightarrow a}\}_{j \in \partial a \setminus i}), \quad (5)$$

where the functions f and g read explicitly

$$f(u_1, \dots, u_d) = \frac{\prod_{i=1}^d (1 + u_i) - \prod_{i=1}^d (1 - u_i)}{z_0(u_1, \dots, u_d)}, \quad z_0(u_1, \dots, u_d) = \prod_{i=1}^d (1 + u_i) + \prod_{i=1}^d (1 - u_i), \quad (6)$$

$$g(h_1, \dots, h_{k-1}) = \frac{\sum_{\sigma_1, \dots, \sigma_k} \omega(\sigma_1, \dots, \sigma_k) \sigma_k \prod_{i=1}^{k-1} (1 + h_i \sigma_i)}{\hat{z}_0(h_1, \dots, h_{k-1})}, \quad \hat{z}_0(h_1, \dots, h_{k-1}) = \sum_{\sigma_1, \dots, \sigma_k} \omega(\sigma_1, \dots, \sigma_k) \prod_{i=1}^{k-1} (1 + h_i \sigma_i) \quad (7)$$

The function f has been written here for a vertex of degree $d + 1$.

On a tree the equations (5) admit a single solution, that can be found by iterating (5) from the leaves towards the interior of the graph. Once this solution is determined one can easily compute the marginal probability of σ_i under

μ (using the formula in (3) with all messages incoming onto i), as well as the partition function $Z(G)$:

$$\frac{1}{N} \ln Z(G) = \frac{1}{N} \sum_{i=1}^N \ln \mathcal{Z}_0^v(\{u_{a \rightarrow i}\}_{a \in \partial i}) + \frac{1}{N} \sum_{a=1}^M \ln \mathcal{Z}_0^c(\{h_{i \rightarrow a}\}_{i \in \partial a}) - \frac{1}{N} \sum_{(i,a)} \ln \mathcal{Z}_0^e(h_{i \rightarrow a}, u_{a \rightarrow i}), \quad (8)$$

where the last sum runs over the edges of the factor graph, and the local partition functions are defined as:

$$\mathcal{Z}_0^v(u_1, \dots, u_d) = \sum_{\sigma} \prod_{i=1}^d \left(\frac{1 + \sigma u_i}{2} \right), \quad (9)$$

$$\mathcal{Z}_0^c(h_1, \dots, h_k) = \sum_{\sigma_1, \dots, \sigma_k} \omega(\sigma_1, \dots, \sigma_k) \prod_{i=1}^k \left(\frac{1 + \sigma_i h_i}{2} \right), \quad (10)$$

$$\mathcal{Z}_0^e(h, u) = \sum_{\sigma} \left(\frac{1 + \sigma h}{2} \right) \left(\frac{1 + \sigma u}{2} \right). \quad (11)$$

The recursive equations (5) and the associated expression (8) of the log partition function (a.k.a. free-entropy density) are exact if the factor graph is a finite tree; they can however be used heuristically on any factor graph, even in the presence of cycles. The resulting message passing iterative algorithm that searches for fixed points of (5) is then known as Belief Propagation (BP) [60–62], and the expression (8) is called the Bethe-Peierls approximation for the free-entropy. As typical random graphs are locally tree-like one can reasonably hope that this approach is asymptotically exact in the thermodynamic limit for typical instances. This is indeed the basis of the cavity method, with nevertheless some subtleties in the treatment of the long loops that are present in random graphs.

C. Replica symmetric cavity method

The aim of the cavity method is to characterize the properties of the measure μ defined in (1), for typical random graphs in the thermodynamic limit, in particular the value of the quenched free-entropy density

$$\Phi(\alpha, \{\omega_p\}) = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\ln Z(G)], \quad (12)$$

around which $(\ln Z(G))/N$ concentrates thanks to the self-averaging phenomenon.

There are different versions of the cavity method, that rely on self-consistent hypotheses of various complexity on the effect of the long loops. In the simplest version, called replica symmetric (RS), one assumes a fast decay of the correlations between distant variables in the measure $\mu(\underline{\sigma})$, in such a way that the BP equations converge to a unique fixed point on a typical large instance, and that the measure is well described by the locally tree-like approximation. Consider then an uniformly chosen directed edge $i \rightarrow a$ in a random hypergraph, and call \mathcal{P}^{RS} the probability law of the fixed-point message $h_{i \rightarrow a}$ thus obtained. We shall denote similarly $\widehat{\mathcal{P}}^{RS}$ the probability of the messages $u_{a \rightarrow i}$; within the decorrelation hypothesis of the RS cavity method the incoming messages on a given vertex (resp. hyperedge) are i.i.d. with the law $\widehat{\mathcal{P}}^{RS}$ (resp. \mathcal{P}^{RS}). For this to be self-consistent the recursion equations (5) must become equalities in distribution, or in other words the laws \mathcal{P}^{RS} and $\widehat{\mathcal{P}}^{RS}$ must obey the following equations:

$$\mathcal{P}^{RS}(h) = \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d du_i \widehat{\mathcal{P}}^{RS}(u_i) \right) \delta(h - f(u_1, \dots, u_d)), \quad (13)$$

$$\widehat{\mathcal{P}}^{RS}(u) = \int \left(\prod_{i=1}^{k-1} dh_i \mathcal{P}^{RS}(h_i) \right) \delta(u - g(h_1, \dots, h_{k-1})). \quad (14)$$

The RS cavity prediction for the free-entropy (12) is then obtained by averaging the Bethe expression (8) with respect to the message distributions \mathcal{P}^{RS} and $\widehat{\mathcal{P}}^{RS}$, which yields:

$$\begin{aligned} \Phi^{RS} &= \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d du_i \widehat{\mathcal{P}}^{RS}(u_i) \right) \ln \mathcal{Z}_0^v(u_1, \dots, u_d) + \alpha \int \left(\prod_{i=1}^k dh_i \mathcal{P}^{RS}(h_i) \right) \ln \mathcal{Z}_0^c(h_1, \dots, h_k) \\ &\quad - \alpha k \int dh du \mathcal{P}^{RS}(h) \widehat{\mathcal{P}}^{RS}(u) \ln \mathcal{Z}_0^e(h, u). \end{aligned} \quad (15)$$

As we assume that $\omega_p = \omega_{k-p}$, i.e. that the model is invariant under the global spin-flip symmetry, the RS equations admit as a solution the uniform distributions $\mathcal{P}^{RS}(h) = \delta(h)$, $\widehat{\mathcal{P}}^{RS}(u) = \delta(u)$. For frustrated models with an antiferromagnetic character this is the relevant solution (a spontaneous breaking of the symmetry between positive and negative spins would induce an alternating order of the Néel type, that is admissible on a tree but incompatible with random graphs that contain cycles of odd lengths), hence one obtains explicitly the value of the free-entropy by inserting the trivial solution of the RS equations into (15):

$$\Phi^{RS}(\alpha, \{\omega_p\}) = \ln 2 + \alpha \ln \left(\frac{1}{2^k} \sum_{\sigma_1, \dots, \sigma_k} \omega(\sigma_1, \dots, \sigma_k) \right) = \ln 2 + \alpha \ln \left(\frac{1}{2^k} \sum_{p=0}^k \binom{k}{p} \omega_p \right). \quad (16)$$

Note that this expression actually coincides with the annealed (first moment) computation $\lim_{N \rightarrow \infty} (\ln \mathbb{E}[Z(G)])/N$.

In the special case $\omega_0 = \omega_k = 0$, $\omega_1 = \dots = \omega_{k-1} = 1$, for which $\mu(\underline{\sigma})$ corresponds to the uniform measure over proper bicolourings, the partition function $Z(G)$ counts the number of solutions, hence the free-entropy $\ln Z$ is equal to the entropy of the uniform measure. The prediction of the RS cavity method is thus (using a subscript u for uniform):

$$s_u^{RS}(\alpha) = \Phi_u^{RS}(\alpha) = \ln 2 + \alpha \ln \left(1 - \frac{1}{2^{k-1}} \right). \quad (17)$$

For a generic choice of parameters $\{\omega_p\}$ the free-entropy $\ln Z$ differs from the (Shannon) entropy of the measure $\mu(\underline{\sigma})$. The latter can be obtained by a Legendre transform with respect to the parameters $\{\omega_p\}$; one way to justify this statement is to remember that for a probability measure of the form (1) one has

$$S(\mu) = - \sum_{\underline{\sigma}} \mu(\underline{\sigma}) \ln \mu(\underline{\sigma}) = \ln Z - \sum_{a=1}^M \sum_{\underline{\sigma}} \mu(\underline{\sigma}) \ln \omega(\sigma_{\partial a}). \quad (18)$$

From the joint marginal of the variables around a constraint in the trivial RS solution, and from the RS free-entropy (16) one thus obtains the RS prediction for the entropy density

$$s^{RS}(\alpha, \{\omega_p\}) = \ln 2 + \alpha \ln \left(\frac{1}{2^k} \sum_{p=0}^k \binom{k}{p} \omega_p \right) - \alpha \frac{\sum_{p=0}^k \binom{k}{p} \omega_p \ln \omega_p}{\sum_{p=0}^k \binom{k}{p} \omega_p}. \quad (19)$$

This quantity is a decreasing function of α and becomes negative for $\alpha > \alpha^{s=0}$, with

$$\alpha^{s=0}(\{\omega_p\}) = \frac{\ln 2}{\frac{\sum_{p=0}^k \binom{k}{p} \omega_p \ln \omega_p}{\sum_{p=0}^k \binom{k}{p} \omega_p} - \ln \left(\frac{1}{2^k} \sum_{p=0}^k \binom{k}{p} \omega_p \right)}. \quad (20)$$

The negativity of the entropy for $\alpha > \alpha^{s=0}$ is a clear sign of the failure of the RS assumptions, the Shannon entropy of a discrete probability measure being non-negative.

D. 1RSB formalism

1. 1RSB cavity equations

The hypothesis underlying the RS cavity method must break down when the density of interactions per variable α becomes too large; a first hint of this phenomenon, called Replica Symmetry Breaking (RSB), is the negativity of the RS entropy at large enough α , which is impossible for a system with discrete degrees of freedom. As a matter of fact RSB can occur before $\alpha^{s=0}$; increasing α above a certain threshold causes the appearance of long-range correlations between distant variables under the measure μ , which contradicts the RS hypothesis. In such a case it becomes necessary to use more refined versions of the cavity method, that are able to deal with this RSB phenomenon [57]. At its first non-trivial level, called 1RSB for one step of RSB, the cavity method postulates the existence of a partition of the configuration space $\{-1, 1\}^N$ into “pure states”, or clusters, such that the restriction of the measure μ to a pure

state has good decorrelation properties. This restricted measure can then be treated as the full measure in the RS cavity method, i.e. with BP equations to describe its marginal probabilities, and the Bethe free-entropy to compute its partition function.

To be more quantitative let us index with γ the partition of the configuration space into clusters, and denote Z_γ the contribution to the partition function of the γ -th cluster, as well as $\{u_{a \rightarrow i}^\gamma, h_{i \rightarrow a}^\gamma\}$ the solution of the BP equations that describe it. The 1RSB cavity method aims at computing the potential

$$\Phi_1(m) = \lim_{N \rightarrow \infty} \frac{1}{N} \ln \left(\sum_\gamma (Z_\gamma)^m \right), \quad (21)$$

where the so-called Parisi parameter m allows to weight differently the various pure states, according to their relative weights. This quantity contains precious informations about the pure-state decomposition; suppose indeed that, at the leading exponential order, there are $e^{N\Sigma(\phi)}$ pure states γ with $Z_\gamma = e^{N\phi}$ (again neglecting sub-exponential corrections). The so-called complexity $\Sigma(\phi)$ plays thus the role of an entropy density, with pure states replacing usual configurations, and captures the RSB phenomenon quantitatively. The potential $\Phi_1(m)$ and the complexity $\Sigma(\phi)$ are Legendre transforms of each other [63]; evaluating (21) via the Laplace method yields indeed

$$\Phi_1(m) = \sup_\phi [\Sigma(\phi) + m\phi], \quad (22)$$

which can be inverted in terms of the conjugated parameter as

$$\Sigma(m) = \Phi_1(m) - m \frac{d}{dm} \Phi_1(m). \quad (23)$$

In order to compute Φ_1 one introduces, for a given sample and a given edge (i, a) of the factor graph, two distributions $P_{i \rightarrow a}$ and $\widehat{P}_{a \rightarrow i}$, that encode the laws of $h_{i \rightarrow a}^\gamma$ and $u_{a \rightarrow i}^\gamma$ when the pure state γ is chosen randomly with a probability proportional to Z_γ^m . These distributions are found to obey self-consistent equations of the form

$$P_{i \rightarrow a} = F(\{\widehat{P}_{b \rightarrow i}\}_{b \in \partial i \setminus a}), \quad \widehat{P}_{a \rightarrow i} = G(\{P_{j \rightarrow a}\}_{j \in \partial a \setminus i}), \quad (24)$$

where $P = F(\widehat{P}_1, \dots, \widehat{P}_d)$ is a shorthand for

$$P(h) = \frac{1}{z_1(\widehat{P}_1, \dots, \widehat{P}_d)} \int \left(\prod_{i=1}^d du_i \widehat{P}_i(u_i) \right) \delta(h - f(u_1, \dots, u_d)) z_0(u_1, \dots, u_d)^m, \quad (25)$$

and $\widehat{P} = G(P_1, \dots, P_{k-1})$ means

$$\widehat{P}(u) = \frac{1}{\widehat{z}_1(P_1, \dots, P_{k-1})} \int \left(\prod_{i=1}^{k-1} dh_i P_i(h_i) \right) \delta(u - g(h_1, \dots, h_{k-1})) \widehat{z}_0(h_1, \dots, h_{k-1})^m; \quad (26)$$

the functions f , z_0 , and g , \widehat{z}_0 were defined in Eqs. (6) and (7), respectively, and the factors z_1 and \widehat{z}_1 ensure the normalization of the distributions $P(h)$ and $\widehat{P}(u)$.

In order to deal with random hypergraphs one introduces the probability distributions over the 1RSB messages $\mathcal{P}^{1RSB}(P)$ and $\widehat{\mathcal{P}}^{1RSB}(\widehat{P})$ that obeys the consistency relations similar to (13),

$$\begin{aligned} \mathcal{P}^{1RSB}(P) &= \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d d\widehat{P}_i \widehat{\mathcal{P}}^{1RSB}(\widehat{P}_i) \right) \delta[P - F(\widehat{P}_1, \dots, \widehat{P}_d)], \\ \widehat{\mathcal{P}}^{1RSB}(\widehat{P}) &= \int \left(\prod_{i=1}^{k-1} dP_i \mathcal{P}^{1RSB}(P_i) \right) \delta[\widehat{P} - G(P_1, \dots, P_{k-1})]. \end{aligned} \quad (27)$$

The 1RSB potential for typical random hypergraphs can then be computed from the solution of these equations as

$$\begin{aligned} \Phi_1(m) &= \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d d\widehat{P}_i \widehat{\mathcal{P}}^{1RSB}(\widehat{P}_i) \right) \ln \mathcal{Z}_1^\gamma(\widehat{P}_1, \dots, \widehat{P}_d) + \alpha \int \left(\prod_{i=1}^k dP_i \mathcal{P}^{1RSB}(P_i) \right) \ln \mathcal{Z}_1^e(P_1, \dots, P_k) \\ &\quad - \alpha k \int dP d\widehat{P} \mathcal{P}^{1RSB}(P) \widehat{\mathcal{P}}^{1RSB}(\widehat{P}) \ln \mathcal{Z}_1^e(P, \widehat{P}), \end{aligned} \quad (28)$$

with:

$$\mathcal{Z}_1^v(\widehat{P}_1, \dots, \widehat{P}_d) = \int \left(\prod_{i=1}^d du_i \widehat{P}_i(u_i) \right) (\mathcal{Z}_0^v(u_1, \dots, u_d))^m, \quad (29)$$

$$\mathcal{Z}_1^c(P_1, \dots, P_k) = \int \left(\prod_{i=1}^k dh_i P_i(h_i) \right) (\mathcal{Z}_0^c(h_1, \dots, h_k))^m, \quad (30)$$

$$\mathcal{Z}_1^e(P, \widehat{P}) = \int dh du P(h) \widehat{P}(u) (\mathcal{Z}_0^e(h, u))^m. \quad (31)$$

Finally the 1RSB prediction for the free-entropy is

$$\Phi^{1RSB} = \inf_{m \in [0,1]} \frac{\Phi_1(m)}{m}. \quad (32)$$

Note that the 1RSB equations always admit the RS solution as a special case, when the distributions P in the support of \mathcal{P}^{1RSB} are Dirac measures. In most models this trivial solution of the 1RSB equations is the only one at small values of α ; then $\Phi_1(m) = m\Phi^{RS}$, and the thermodynamic prediction of the RS and 1RSB versions of the cavity method coincides. Increasing the number of constraints of the system non-trivial solutions of the 1RSB equations can appear; the dynamic threshold α_d is defined as the smallest value of α for which the 1RSB equations with $m = 1$ admit a solution distinct from the RS one. A further distinction has then to be made: if the associated complexity $\Sigma(m = 1)$ is positive the extremum in (32) is reached for $m = 1$ and $\Phi^{1RSB} = \Phi^{RS}$. In such a ‘‘dynamic 1RSB’’ situation the typical configurations of the Gibbs measure are supported on an exponentially large number of pure states, in such a way that the total free-entropy (or any correlation functions between a finite number of variables) is unable to detect the difference with a RS situation. On the contrary when $\Sigma(m = 1) < 0$ the extremum in (32) selects a non-trivial value $m_s < 1$ of the Parisi parameter, the Gibbs measure condenses on a sub-exponential number of clusters, and correlations between finite sets of variables unveil the RSB phenomenon. One calls condensation threshold α_c the smallest value of α for which a solution of the 1RSB equations with $\Sigma(m = 1) < 0$ exists, which corresponds to a point of non-analyticity of the free-entropy density.

2. Simplifications for $m = 1$

The complete 1RSB equations have a rather intricate structure, as they are self-consistent equations for probability distributions over probability distributions, which make in particular their numerical resolution rather cumbersome. Fortunately for special values of the parameter m (i.e. $m = 0$ and $m = 1$) they can be largely simplified. We shall sketch here this simplification procedure for $m = 1$, which as explained above is the important one for the determination of the dynamic and condensation phase transitions; for further details the reader is referred to [33] where the simplifications are explained in more details and in a general setting, and to [64] where their connections with the tree reconstruction problem are discussed.

The crucial technical property of the equations that opens the door to simplifications at $m = 1$ is the fact that, for this value, the normalization constant z_1 in (25) does not depend on the whole distributions $\widehat{P}_1, \dots, \widehat{P}_d$, but only on their average values $\int du_i \widehat{P}_i(u_i) u_i$ (a similar statement holds for \widehat{z}_1 in (26)). Conditional on this average values F is thus a multilinear function of its arguments, the equations (27) can then be averaged and closed on the mean distributions Q and \widehat{Q} defined as:

$$Q = \int dP \mathcal{P}^{1RSB}(P) P, \quad \widehat{Q} = \int d\widehat{P} \widehat{\mathcal{P}}^{1RSB}(\widehat{P}) \widehat{P}, \quad (33)$$

which must be symmetric probability laws (i.e. $Q(h) = Q(-h)$ and $\widehat{Q}(u) = \widehat{Q}(-u)$) for the global spin-flip symmetry to be preserved. These two quantities are solutions of

$$Q(h) = \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d du_i \widehat{Q}(u_i) \right) \delta(h - f(u_1, \dots, u_d)) \frac{z_0(u_1, \dots, u_d)}{z_0(0, \dots, 0)}, \quad (34)$$

$$\widehat{Q}(u) = \int \left(\prod_{i=1}^{k-1} dh_i Q(h_i) \right) \delta(u - g(h_1, \dots, h_{k-1})) \frac{\widehat{z}_0(h_1, \dots, h_{k-1})}{\widehat{z}_0(0, \dots, 0)}. \quad (35)$$

These equations are definitely simpler than the full 1RSB equations, as they bear on probability distributions instead of distributions of distributions; they have however one inconvenient feature, in particular for their numerical resolution, namely the reweighting terms z_0 and \hat{z}_0 which prevents their direct interpretation as recursive distributional equations. To get around this difficulty we shall define, for $\sigma = \pm 1$, the distributions $Q_\sigma(h) = (1 + h\sigma)Q(h)$ and $\hat{Q}_\sigma(u) = (1 + u\sigma)\hat{Q}(u)$. Thanks to the symmetry of Q and \hat{Q} these are well-normalized, and are related to the original distributions by $Q(h) = (Q_+(h) + Q_-(h))/2$. One can then show that they obey the following equations,

$$Q_\sigma(h) = \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d du_i \hat{Q}_\sigma(u_i) \right) \delta(h - f(u_1, \dots, u_d)), \quad (36)$$

$$\hat{Q}_\sigma(u) = \sum_{\sigma_1, \dots, \sigma_{k-1}} \tilde{p}(\sigma_1, \dots, \sigma_{k-1} | \sigma) \int \left(\prod_{i=1}^{k-1} dh_i Q_{\sigma_i}(h_i) \right) \delta(u - g(h_1, \dots, h_{k-1})),$$

with the conditional probability distribution:

$$\tilde{p}(\sigma_1, \dots, \sigma_{k-1} | \sigma) = \frac{\omega(\sigma_1, \dots, \sigma_{k-1}, \sigma)}{\sum_{\sigma'_1, \dots, \sigma'_{k-1}} \omega(\sigma'_1, \dots, \sigma'_{k-1}, \sigma)}. \quad (37)$$

Noting finally that the global flip-spin symmetry implies $Q_-(h) = Q_+(-h)$ and $\hat{Q}_-(u) = \hat{Q}_+(-u)$, one can write closed equations, without reweighting terms, on Q_+ and \hat{Q}_+ solely:

$$Q_+^{(t+1)}(h) = \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d du_i \hat{Q}_+^{(t)}(u_i) \right) \delta(h - f(u_1, \dots, u_d)), \quad (38)$$

$$\hat{Q}_+^{(t)}(u) = \sum_{\sigma_1, \dots, \sigma_{k-1}} \tilde{p}(\sigma_1, \dots, \sigma_{k-1} | +) \int \left(\prod_{i=1}^{k-1} dh_i Q_+^{(t)}(h_i) \right) \delta(u - g(\sigma_1 h_1, \dots, \sigma_{k-1} h_{k-1})),$$

where for future use we introduced discrete time indices t on these distributions, and where

$$\tilde{p}(\sigma_1, \dots, \sigma_{k-1} | +) = \frac{\omega(\sigma_1, \dots, \sigma_{k-1}, +)}{\sum_{\sigma'_1, \dots, \sigma'_{k-1}} \omega(\sigma'_1, \dots, \sigma'_{k-1}, +)} = \frac{\sum_{p=0}^{k-1} \omega_p \mathbb{I} \left[\sum_{i=1}^{k-1} \sigma_i = k - 1 - 2p \right]}{\sum_{p=0}^{k-1} \binom{k-1}{p} \omega_p}. \quad (39)$$

Let us now turn to the thermodynamic computations; at $m = 1$ one can check that $\Phi_1(m = 1) = \Phi^{RS}$ as given in (16). To compute the complexity at $m = 1$ from (23) we need to take the derivative with respect to m of Φ_1 from (28). Because of its variational character ((28) is stationary with respect to variations of \mathcal{P}^{1RSB} and $\hat{\mathcal{P}}^{1RSB}$ as long as the 1RSB cavity equations (27) are fulfilled) only the explicit dependency on m has to be differentiated. Doing the simplification at $m = 1$ yields then an expression in terms of Q_+ and \hat{Q}_+ :

$$\frac{d}{dm} \Phi_1(m) \Big|_{m=1} = \sum_{d=0}^{\infty} p_d \int \left(\prod_{i=1}^d du_i \hat{Q}_+(u_i) \right) \ln \mathcal{Z}_0^v(u_1, \dots, u_d) \quad (40)$$

$$+ \alpha \sum_{\sigma_1, \dots, \sigma_k} p(\sigma_1, \dots, \sigma_k) \int \left(\prod_{i=1}^k dh_i Q_+(h_i) \right) \ln \mathcal{Z}_0^c(\sigma_1 h_1, \dots, \sigma_k h_k) \quad (41)$$

$$- \alpha k \int dh du Q_+(h) \hat{Q}_+(u) \ln \mathcal{Z}_0^e(h, u),$$

with

$$p(\sigma_1, \dots, \sigma_k) = \frac{\omega(\sigma_1, \dots, \sigma_k)}{\sum_{\sigma'_1, \dots, \sigma'_k} \omega(\sigma'_1, \dots, \sigma'_k)}. \quad (42)$$

The form (38) of the 1RSB equations at $m = 1$ is particularly convenient for an approximate numerical resolution with a procedure known as population dynamics [57]. Suppose indeed that $\widehat{Q}_+^{(t)}(u)$ can be approximated by the empirical distribution over a large sample of representative elements:

$$\widehat{Q}_+^{(t)}(u) \approx \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \delta(u - u_i^{(t)}) . \quad (43)$$

Inserting this form in the r.h.s. of the first line of (38) yields an approximation for $Q_+^{(t)}(h)$ as

$$Q_+^{(t+1)}(h) \approx \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \delta(h - h_i^{(t+1)}) , \quad (44)$$

where each of the representants $h_i^{(t+1)}$ is constructed independently by drawing an integer d from the law p_d , then d indices i_1, \dots, i_d uniformly at random in $\{1, \dots, \mathcal{N}\}$ and setting $h_i^{(t+1)} = g(u_{i_1}^{(t)}, \dots, u_{i_d}^{(t)})$. The second line of (38) can similarly be translated into a rule for generating a population of fields $u^{(t)}$ from the population of fields $h^{(t)}$. The size \mathcal{N} of the population used controls the computational cost of the procedure, and the numerical accuracy (in the limit $\mathcal{N} \rightarrow \infty$ empirical distributions converge to the exact ones). Iterating these two steps many times one converges to a fixed point solution of (38), which can either be the trivial one $Q_+(h) = \delta(h)$, $\widehat{Q}_+(u) = \delta(u)$, or a non-trivial solution. The dynamic transition α_d is precisely the threshold that separates these two behaviors. It has been shown in [64] that the equations (38) can also be interpreted in terms of an information-theoretic problem called tree reconstruction [12, 65, 66]. In the latter one considers a rooted Galton-Watson random tree and use it as an information channel, with spin variables located on the vertices. The value σ of the spin at the root is broadcasted along the hyperedges, according to the free-boundary Gibbs measure with local interaction ω . The question in this context is whether the observation of the variables at distance t from the root contains a non-vanishing information on σ , in the limit $t \rightarrow \infty$; in which case one says that the problem is reconstructible. As $Q_+^{(t)}(h)$, after t iterations of (38), is the distribution of the posterior magnetization of the root conditional on the observation of the variables at distance t , in the broadcast process with $\sigma = +1$, the reconstructibility of the tree problem is equivalent to the existence of a non-trivial solution of the 1RSB equations with $m = 1$, and the dynamic threshold α_d coincides with the reconstruction transition of the associated tree problem. Moreover this connection unveils a natural initial condition for the iterative numerical resolution of (38),

$$Q_+^{(t=0)}(h) = \delta(h - 1) , \quad (45)$$

corresponding to the perfect observation of the variables at distance t from the root.

3. The local instability of the RS solution (Kesten-Stigum bound)

The properties of the measure $\mu(\underline{\sigma})$ change drastically when, upon increasing α , one moves from the RS phase to the 1RSB phase. The transition between the two situations, at the critical (dynamic) density α_d , shows up as the appearance of a non-trivial solution of the 1RSB equations at $m = 1$. Depending on the models this bifurcation, on which we shall come back in more details in Sec. III, can occur in a continuous or a discontinuous way. We shall discuss here a bound on α_d , known as the Kesten-Stigum [12, 67] transition in the context of the tree reconstruction problem, or as the de Almeida-Thouless [68] transition for mean-field spin-glasses, that is tight for continuous bifurcations and that in any case provide an easy to compute analytical upper bound on α_d (besides the bound $\alpha_d < \alpha^{s=0}$ we already discussed).

Let us recall that the 1RSB equations (27) always admit as a solution the RS distribution, with $\mathcal{P}^{1RSB}(P) = \delta[P - P_{\text{triv}}]$, $P_{\text{triv}}(h) = \delta(h)$ (and similarly for $\widehat{\mathcal{P}}^{1RSB}(\widehat{P})$). One way to test the existence of a non-trivial solution of the 1RSB equations is to investigate the local stability of the RS solution. Suppose indeed that the distributions in the support of \mathcal{P}^{1RSB} are close to P_{triv} , i.e. that they are supported on small values of h . One can then expand (25,26) and study the evolution of their average moments under the iterations of (27); the global spin-flip symmetry imposes that the average mean remains zero. The first non-trivial moment is thus the variance, which is found after a short computation (see for instance [69] or App. B in [70] for more details) to grow if and only if $\alpha k(k-1)\theta^2 > 1$, where θ is the derivative of $g(u_1, \dots, u_{k-1})$ with respect to one of its arguments, evaluated on the trivial fixed-point. From the expression (7) we thus obtain the Kesten-Stigum threshold α_{KS} above which the trivial solution of the 1RSB

equations is unstable (and there must then exist a non-trivial solution) as

$$\alpha_{\text{KS}} = \frac{1}{k(k-1)\theta^2}, \quad \theta = \frac{\sum_{\sigma_1, \dots, \sigma_k} \omega(\sigma_1, \dots, \sigma_k) \sigma_1 \sigma_2}{\sum_{\sigma_1, \dots, \sigma_k} \omega(\sigma_1, \dots, \sigma_k)} = \frac{\sum_{p=0}^{k-2} \binom{k-2}{p} (\omega_p - 2\omega_{p+1} + \omega_{p+2})}{\sum_{p=0}^k \binom{k}{p} \omega_p}. \quad (46)$$

4. The presence of hard-fields in the 1RSB solution with $m = 1$ (rigidity threshold)

For the special value $m = 1$ of the Parisi parameter we have obtained in (38) a simplified form of the 1RSB equations; even if much simpler than the general formalism, these equations bear on probability distributions ($Q_+(h)$ and $\widehat{Q}_+(u)$) and cannot be solved analytically in general. One can however make some more explicit statements when $\omega_0 = \omega_k = 0$ and $\omega_p > 0$ for $p \in [1, k-1]$, i.e. when the constraints forbid monochromatic hyperedges, but allow all bichromatic configurations (even if they can give different weights to these configurations). In this case the distribution $Q_+(h)$ can contain a Dirac peak in $h = 1$, corresponding to “hard fields” that impose strictly the value of some variables (that are called frozen, or rigid) inside one pure state. Let us denote q (resp. \widehat{q}) the weight of $h = 1$ (resp. $u = 1$) under Q_+ (resp. \widehat{Q}_+). One can then obtain from (38) closed equations on q and \widehat{q} :

$$q = 1 - \sum_{d=0}^{\infty} p_d (1 - \widehat{q})^d = 1 - e^{-\alpha k \widehat{q}}, \quad (47)$$

$$\widehat{q} = \widetilde{p}(-, \dots, -|+) q^{k-1}; \quad (48)$$

indeed the expression of f in (6) reveals that $h = 1$ as soon as one of the neighboring constraint sends the hard field $u_i = 1$, while (7) shows that a variable is forced to a certain value σ by a constraint only when the $k-1$ other variables are simultaneously forced to $-\sigma$. Eliminating \widehat{q} one sees that q is solution of

$$q = 1 - \exp(-\Gamma q^{k-1}), \quad \text{with } \Gamma = \alpha k \widetilde{p}(-, \dots, -|+). \quad (49)$$

For $k \geq 3$ a non-trivial solution to this equation appears discontinuously, when Γ exceeds a critical value Γ_r . The value of Γ_r , and the associated solution q_r , are the solutions of

$$\begin{cases} q_r = 1 - \exp(-\Gamma_r q_r^{k-1}) \\ 1 = (k-1) \Gamma_r q_r^{k-2} \exp(-\Gamma_r q_r^{k-1}) \end{cases}; \quad (50)$$

see Sec. III for more explanations on the origin of these equations. One can close the equation on q_r , that obeys $1 = (k-1) \ln(1 - q_r)(1 - 1/q_r)$, from which Γ_r is obtained as $\Gamma_r = -\frac{\ln(1 - q_r)}{q_r^{k-1}}$. Note that q_r and Γ_r depend solely on k .

Translating back to the parameters $\alpha, \{\omega_p\}$, one sees that for any choice of $\{\omega_p\}$ such that $\omega_0 = \omega_k = 0$ and $\omega_p > 0$ for $p \in [1, k-1]$ there exists a “rigidity threshold” $\alpha_r(k, \{\omega_p\})$ above which the equation on the probability of hard-fields admits a non-trivial solution, with (recalling the expression of \widetilde{p} from (39))

$$\alpha_r(k, \{\omega_p\}) = \frac{1}{k} \Gamma_r(k) \frac{\sum_{p=1}^{k-1} \binom{k-1}{p} \omega_p}{\omega_1} = \frac{1}{k} \Gamma_r(k) \frac{\sum_{p=1}^{k-1} \binom{k}{p} \omega_p}{2\omega_1}, \quad (51)$$

where in the last step we exploited the symmetry $\omega_p = \omega_{k-p}$. We shall denote $\alpha_{r,u}$ the value of this threshold for the uniform case $\omega_1 = \dots = \omega_{k-1} = 1$, i.e. when all proper bicolourings are weighted equally, in such a way that

$$\alpha_{r,u}(k) = \frac{1}{k} \Gamma_r(k) (2^{k-1} - 1). \quad (52)$$

This rigidity threshold is an upperbound on the dynamic transition: if there exists a solution to the 1RSB equations at $m = 1$ containing hard-fields, this is certainly a non-trivial solution of the 1RSB equations. The inequality $\alpha_d \leq \alpha_r$ is in general strict, i.e. there can be non-trivial solution of the 1RSB equations at $m = 1$ that do not contain any hard-field; this has been seen numerically in many problems, and proven rigorously for the graph q -coloring problem in the large q limit in [37, 39]. This rigidity threshold corresponds actually to a transition for a strong form of reconstructibility in the tree reconstruction interpretation: when $\alpha > \alpha_r$, with positive probability the observation of

far away variables allows to infer the value of the root without possibility of error. Let us also underline that among all the parameters $\{\omega_p\}$ that define the bias among proper bicolorings $\omega_1 = \omega_{k-1}$ plays a special role in the expression (51) of the rigidity threshold. Indeed hard fields are propagated along constraints that are “almost violated”, in the sense that they contain a single variable of a given color. Penalizing such “almost monochromatic” hyperedges tends thus to avoid the percolation of frozen variables.

III. ON THE NUMERICAL DETERMINATION OF THE DYNAMIC TRANSITION

The dynamic threshold α_d is the smallest value of α such that the 1RSB equations at $m = 1$ admit a non-trivial solution (besides the RS trivial one in which all fields are equal to zero). Depending on the models the appearance of a non-trivial solution can occur either in a continuous or a discontinuous way. In the former case one has $\alpha_d = \alpha_{KS}$, the bifurcation occurs via the local instability of the trivial fixed point studied in Sec. IID 3, and α_d is thus known analytically. In the latter case $\alpha_d < \alpha_{KS}$, the birth of the non-trivial solution occurs non-perturbatively and cannot be detected from the properties of the trivial fixed point. The accurate numerical determination of α_d when the transition is discontinuous is a rather difficult task. It corresponds to study a bifurcation for a fixed-point equation of the form $Q = F(Q, \alpha)$, where Q is a probability distribution and F a functional on this space, depending on the parameter α . We shall discuss later on the different numerical strategies that can be followed to determine α_d , in particular one that, to the best of our knowledge, is new in this context. To explain these different methods it is instructive to study first a much simpler case, in which the unknown Q is replaced by a real number, that we shall instead denote x for clarity.

A. Scalar bifurcations

Let us consider a function $f(x, \alpha)$, smooth in its two real arguments, and the associated discrete dynamical system $x^{(t+1)} = f(x^{(t)}, \alpha)$, parametrized by α . We recall some basic facts in this setting: the stationary configurations of the dynamical system are the solutions $x_*(\alpha)$ of the fixed point equation $x = f(x, \alpha)$. Their (linear) stability is determined by the coefficient $\lambda(\alpha) = (\partial_x f)(x_*(\alpha), \alpha)$ (here and below we denote $(\partial_x f)$, $(\partial_\alpha f)$ and so on the partial derivatives of the function f); a fixed point $x_*(\alpha)$ is indeed stable under iterations if $|\lambda(\alpha)| < 1$, and unstable if $|\lambda(\alpha)| > 1$. We also recall the implicit function theorem: if (x_0, α_0) is a solution of $f(x_0, \alpha_0) = x_0$, and if $(\partial_x f)(x_0, \alpha_0) \neq 1$, then there is a unique smooth function $x_*(\alpha)$ with $f(x_*(\alpha), \alpha) = x_*(\alpha)$ in a neighborhood of α_0 , with $x_*(\alpha_0) = x_0$. Hence the bifurcations of the fixed point equation, i.e. the modifications in the number of solutions, or the singularities of these solutions, are associated to points where $(\partial_x f)(x(\alpha_0), \alpha_0) = 1$, in order for the implicit function theorem to be unapplicable. At these points the stability parameter $\lambda(\alpha)$ reaches its critical value 1.

To be more concrete we shall make the additional hypotheses that the order parameter x is restricted to non-negative values ($x \geq 0$), and that $f(0, \alpha) = 0$ for all α . Let us assume that this trivial fixed point, that exists for all α , is the unique solution for small enough values of α , and becomes non-unique when α exceeds a threshold α_d . The two simplest ways to implement these hypotheses are sketched in Fig. 2, corresponding to a continuous bifurcation on the left panel, a discontinuous one on the right. Let us state a series of simple facts on these two types of phase transitions, that will be enlightening when we turn to the functional case.

Consider first the continuous case illustrated on the left panel of Fig. 2. The bifurcation occurs at the critical parameter value α_d defined by $(\partial_x f)(0, \alpha_d) = 1$, the trivial fixed point being stable (resp. unstable) for $\alpha < \alpha_d$ (resp. $\alpha > \alpha_d$). For $\alpha > \alpha_d$ there exists a non-trivial branch of stable fixed points $x_*(\alpha) > 0$; in the neighborhood of the bifurcation the latter behaves as

$$x_*(\alpha) = K(\alpha - \alpha_d) + o((\alpha - \alpha_d)) \quad \text{when } \alpha \rightarrow \alpha_d^+, \quad (53)$$

with $K = -2(\partial_{x\alpha} f)/(\partial_{xx} f)$, the derivatives being computed in $(0, \alpha_d)$ (here and in the following the expressions of the various constants K can be obtained by a Taylor expansion of the equation $x = f(x, \alpha)$ around the bifurcation point, at the lowest non-trivial order). The stability parameter of the non-trivial solution, $\lambda(\alpha) = (\partial_x f)(x_*(\alpha), \alpha)$, reaches its marginal value 1 at the bifurcation as

$$\lambda(\alpha) = 1 - K'(\alpha - \alpha_d) + o(\alpha - \alpha_d) \quad \text{when } \alpha \rightarrow \alpha_d^+, \quad (54)$$

with $K' = (\partial_{x\alpha} f)$.

Let us now turn to the discontinuous case (cf. the right panel of Fig. 2), and emphasize the main properties of the critical behavior of the bifurcation. The trivial fixed point is stable for all values of the parameter α ; the bifurcation

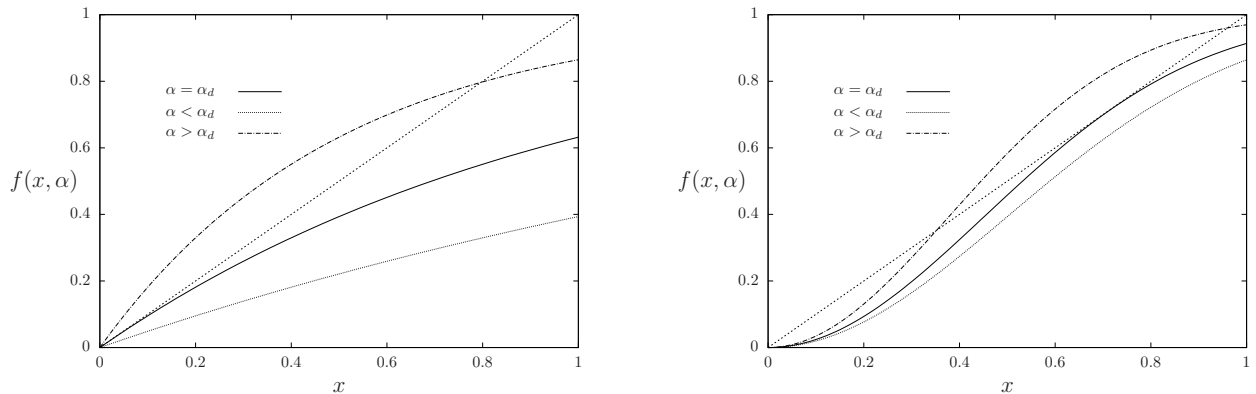


FIG. 2: Left panel: example of a continuous bifurcation with $f(x, \alpha) = 1 - e^{-\alpha x}$, for which $\alpha_d = 1$. Right panel: example of a discontinuous bifurcation with $f(x, \alpha) = 1 - e^{-\alpha x^2}$, for which $\alpha_d = 2.45541$.

occurs at α_d with the abrupt appearance of a solution $x_d > 0$. These two quantities can be determined by solving the system of equations

$$\begin{cases} x_d = f(x_d, \alpha_d) , \\ 1 = (\partial_x f)(x_d, \alpha_d) . \end{cases} \quad (55)$$

For $\alpha > \alpha_d$ there are two branches of non-trivial solutions $x_-(\alpha) < x_d < x_+(\alpha)$ that emerge from x_d (see the right panel of Fig. 3); in the neighborhood of α_d they behave as

$$x_{\pm}(\alpha) = x_d \pm K\sqrt{\alpha - \alpha_d} + o(\sqrt{\alpha - \alpha_d}) \quad \text{when } \alpha \rightarrow \alpha_d^+ , \quad (56)$$

where the coefficient K can be computed from the expansion of f around the bifurcation point (explicitly, $K = \sqrt{-2(\partial_\alpha f)/(\partial_{xx} f)}$, the derivatives being computed in (x_d, α_d)). For $\alpha > \alpha_d$ $x_+(\alpha)$ (resp. $x_-(\alpha)$) is linearly stable (resp. unstable); the stability parameter $\lambda(\alpha) = (\partial_x f)(x_+(\alpha), \alpha)$ of the stable non-trivial branch reaches its critical value 1 at the bifurcation, with a critical exponent 1/2:

$$\lambda(\alpha) = 1 - K'\sqrt{\alpha - \alpha_d} + o(\sqrt{\alpha - \alpha_d}) \quad \text{when } \alpha \rightarrow \alpha_d^+ , \quad (57)$$

with $K' = \sqrt{-2(\partial_\alpha f)(\partial_{xx} f)}$. We present in the left panel of Fig. 3 the iterates $x^{(t+1)} = f(x^{(t)}, \alpha)$, for a few values of $\alpha < \alpha_d$, starting from an initial condition $x^{(0)} > x_d$. Their long time limit is of course 0, the only fixed point in this phase, but when $\alpha \rightarrow \alpha_d^-$ the decay is slower and slower, with a large number of iterations spent around a plateau value at x_d . More quantitatively one can define $t_*(\alpha)$ as the minimal t such that $x^{(t)} \leq x_d - \epsilon$, and obtain that

$$t_*(\alpha) \sim K''(\alpha_d - \alpha)^{-1/2} \quad \text{when } \alpha \rightarrow \alpha_d^- , \quad (58)$$

with $K'' = 2\pi/K'$, independently of $x^{(0)}$ and ϵ (as long as $0 < \epsilon < x_d$). Actually a whole scaling function describing the evolution of $x^{(t)}$ around the plateau can be derived, see [34] for more details.

B. Discontinuous functional bifurcations

Let us now come back to our original goal, namely the determination of the dynamic threshold α_d above which appears a non-trivial solution of the 1RSB equations at $m = 1$. As in the scalar case this transition can occur either in a continuous or in a discontinuous way; the former case was analytically dealt with in Sec. IID 3, we shall hence concentrate now on the discontinuous transitions.

The 1RSB equations (38) can be written abstractly as a functional fixed point equation $Q = F(Q, \alpha)$; at variance with the scalar toy model discussed above they can only be solved approximately, for instance by the population dynamics numerical algorithm explained in Sec. IID 2. Some examples of typical numerical results that can be obtained in this way are presented in Fig. 4; we use as an observable to condense the distribution Q_+ into a single scalar the overlap $q_1 = \int Q_+(h)h dh$, which is equal to 0 for the trivial solution. On the left panel we plot the value




FIG. 3: Properties of the iterates and fixed points for the discontinuous bifurcation undergone by the function $f(x, \alpha) = 1 - e^{-\alpha x^2}$. Left panel: $x^{(t)}$ as a function of t for a few values of α close but strictly less than α_d . Right panel: the non-trivial solutions $x_{\pm}(\alpha)$ for $\alpha \geq \alpha_d$.

of q_1 as a function of the number of iterations, for a few values of α . One sees on this plot, reminiscent of the left panel of Fig. 3, the discontinuous birth of a non-trivial fixed point at α_d , with a longer and longer plateau in the low α phase as a precursor of the transition. On the right panel of Fig. 4 we present the asymptotic value of q_1 reached for large t , for different values of α around the dynamic transition (corresponding to the right panel of Fig. 3), that jumps discontinuously from 0 when α crosses α_d .




FIG. 4: Example of a discontinuous dynamic transition for $k = 6$, $\omega_0 = \omega_k = 0.005$, $\omega_1 = \omega_{k-1} = 0.92$, and $\omega_2 = \dots = \omega_{k-2} = 1$. Left: $q_1(\alpha, t)$ versus iteration time t for different values of α . Right: $q_1(\alpha)$ averaged over t after equilibration. The size of the population used is 10^6 .

It is not completely obvious how to extract a precise estimate of α_d from this kind of data. The simplest approach amounts to determine the curves $q_1^{(t)}$ for several closely spaced values of α , and assess that $\alpha_d \in [\alpha_<, \alpha_>]$, where $\alpha_<$ is the largest value for which $q_1^{(t)}$ drops to 0 at large t , $\alpha_>$ the smallest value for which a stable plateau is encountered. This determination suffers however from inaccuracies due to the finite number of α values one can investigate, the finite number of iterations one can perform (leading to an underestimation of $\alpha_>$) and to the finite size of the population that approximate the distribution Q_+ ($\alpha_<$ can thus be overestimated, finite size fluctuations having a destabilizing effect).

One can try to circumvent these difficulties by getting some inspiration from the much simpler scalar bifurcation studied above. We recall that the criticality at α_d showed up in three different ways: (i) $x_+(\alpha)$ exhibits a square root singularity when $\alpha \rightarrow \alpha_d^+$, see Eq. (56); (ii) the length of the plateau diverges when $\alpha \rightarrow \alpha_d^-$ with a critical exponent $-1/2$, cf. Eq. (58); (iii) the stability parameter $\lambda(\alpha)$ reaches 1 with a square root singularity when $\alpha \rightarrow \alpha_d^+$, as written in Eq. (57).

Assuming the same critical behavior to occur in the discontinuous functional bifurcation case (more complicated behaviors could occur in infinite dimensional spaces, but in absence of accidental degeneracies there should be a single critical direction at a bifurcation driven by a single parameter) one can try to exploit these scaling laws in order to

obtain more precise estimates of α_d . Point (i) translates into a square root singularity of the large t limit of q_1 in the limit $\alpha \rightarrow \alpha_d^+$; this does not seem very useful to us, as it would involve a fit of $q_1(\alpha)$ in which both α_d and $q_1(\alpha_d)$ are unknowns. On the contrary points (ii) and (iii) yield simpler fits for the determination of α_d . The aspect (ii) is very easy to exploit: from the curves of the left panel of Fig. 4 one can deduce immediately a value $t_*(\alpha)$ for the number of iterations necessary to fall below the plateau (as in the scalar case one can define $t_*(\alpha)$ with any threshold strictly between 0 and the plateau value). According to (58) $t_*(\alpha)^{-2}$ should vanish linearly at α_d ; this is indeed what we obtain with a rather good accuracy, see the left panel of Fig. 5. However one cannot reach in this way a very large number of iterations, the numerical rounding errors and finite population size fluctuations having the tendency to accumulate over time; this cutoff on t thus limits the accuracy of this determination of α_d . We have thus turned to the functional generalization of point (iii) above, namely the computation of a stability parameter $\lambda(\alpha)$ for the stable non-trivial branch $\alpha > \alpha_d$, and the determination of α_d as the parameter for which λ reaches 1. This extrapolation is done using the scaling anticipated in the scalar case in (57), and is illustrated in the right panel of Fig. 5. The functional nature of the unknown in the fixed point equation makes the definition of λ more complicated than in the scalar case, where it was simply $\partial_x f$; we give detailed explanations on the numerical computation of $\lambda(\alpha)$ in the functional case in the next section. Before that let us emphasize that the square root behavior of λ around α_d , guessed from the scalar bifurcation, is in very good agreement with the numerical results obtained in the functional case when the dynamic transition is discontinuous (see the right panel of Fig. 5). We believe the determination of α_d reached by the extrapolation of $\lambda(\alpha)$ is more reliable and accurate than the one based on $t_*(\alpha)$. Indeed the former quantity is defined from a stable fixed-point of the equations, averages can be performed in a steady-state (plateau) regime to reduce the statistical error on its computation, while $t_*(\alpha)$ is a measure of a transient regime more sensitive to numerical inaccuracies.




FIG. 5: Study of the discontinuous dynamic transition encountered as a function of α for the choice of parameters $k = 6$, $\omega_0 = \omega_k = 0.005$, $\omega_1 = \omega_{k-1} = 0.92$, $\omega_2 = \dots = \omega_{k-2} = 1$. Left panel: determination of α_d from the study of the decorrelation time $t_*(\alpha)$ for $\alpha < \alpha_d$. The plot displays $t_*(\alpha)^{-2}$ versus α , where one has defined $t_*(\alpha)$ as the first time for which the overlap drops below the value $q_1 = 0.4$. The line is a fit of the data of the form $t_*(\alpha)^{-2} = A(\alpha_d - \alpha)$, with fitting parameters A and α_d . The linear behavior confirms the divergence of t_* with a scaling exponent $-1/2$, as in the scalar case (58), the fit gives the estimation $\alpha_d = 19.47$. Right panel: determination of α_d from the study of the stability parameter $\lambda(\alpha)$ for $\alpha > \alpha_d$. The plot displays $(1 - \lambda(\alpha))^2$ versus α , the linear fit reproduces the scaling behavior (57) of the scalar case, and yields $\alpha_d = 19.467$. Only the first points are used for the fit, in order to avoid the higher order contributions in powers of $\alpha - \alpha_d$ that are clearly visible at the largest values of α .

As a consistency check we also present in Fig. 6 a similar study in the case of a continuous transition. We see that the stability parameter λ computed on the non-trivial solution, i.e. for $\alpha > \alpha_d$, reaches 1 with a linear behavior (as in the scalar case, see Eq. (54)), and that its extrapolation is in good agreement with the analytically computed value of α_{KS} from Eq. (46). Moreover the numerical computation of the stability parameter of the trivial fixed-point coincides for $\alpha < \alpha_d$ with the analytical one, $\lambda^{KS} = \alpha k(k-1)\theta^2$.

C. The stability parameter λ in the functional case

As an intermediate step in the generalization from the scalar to the functional case let us consider a fixed point equation of the form $\vec{x} = f(\vec{x}, \alpha)$, where the unknown \vec{x} is a finite-dimensional real vector. The stability of a branch of solutions $\vec{x}(\alpha)$ can be determined by considering the Jacobian matrix J of the first derivatives of f computed at




FIG. 6: The stability parameter $\lambda(\alpha)$ in the case of a continuous transition (here for $k = 4$, $\omega_0 = \omega_k = 0$, $\omega_1 = \dots = \omega_{k-1} = 1$). From the fit one obtains the result $\alpha = 4.088$, which is in good agreement with the analytical one $\alpha_{KS} = 4.083$.

the fixed point, that can be defined through the linearization

$$f(\vec{x}(\alpha) + \vec{\varepsilon}, \alpha) = \vec{x}(\alpha) + J\vec{\varepsilon} + o(\|\vec{\varepsilon}\|), \quad (59)$$

where $\vec{\varepsilon}$ is a small perturbation around the fixed point. The stability parameter $\lambda(\alpha)$ can then be defined as the spectral radius of J , i.e. the largest absolute value of the elements of its spectrum. This spectral radius can be expressed in terms of successive applications of J on a perturbation $\vec{\varepsilon}$ as

$$\lambda(\alpha) = \lim_{n \rightarrow \infty} \left(\frac{\|J^n \vec{\varepsilon}\|}{\|\vec{\varepsilon}\|} \right)^{\frac{1}{n}}, \quad (60)$$

where we assume that $\vec{\varepsilon}$ has a non-vanishing projection on the eigenspace associated to the relevant eigenvalue, and where $\|\bullet\|$ can be any norm. For future use let us define $\vec{\varepsilon}_n = J^n \vec{\varepsilon}$ and rewrite this expression as

$$\lambda(\alpha) = \lim_{n \rightarrow \infty} \left(\frac{\|\vec{\varepsilon}_n\|}{\|\vec{\varepsilon}_{n-1}\|} \frac{\|\vec{\varepsilon}_{n-1}\|}{\|\vec{\varepsilon}_{n-2}\|} \dots \frac{\|\vec{\varepsilon}_1\|}{\|\vec{\varepsilon}\|} \right)^{\frac{1}{n}} \quad (61)$$

We would like now to extend the computation of a stability parameter to the 1RSB equations (38) that can be rewritten as $Q_+ = F(Q_+, \alpha)$ by grouping the two lines together. Q_+ being a probability distribution the Jacobian of F is now an infinite-dimensional operator, which makes the study of its spectrum rather difficult. Even worse, we do not have at our disposal an exact description of the fixed point Q_+ around which we would like to expand F : we only have a sequence of approximations of $Q_+^{(t)}$ by the population representation written in Eq. (44). The individual elements of these representations still evolve at each iteration step, even when the observables computed as averages of $Q_+^{(t)}$ have reached convergence (within the numerical accuracy fixed by the population size \mathcal{N}). To circumvent these difficulties we have followed a strategy inspired by the expression (61): we consider $Q_+^{(t)}$ and a slight perturbation of it, $Q_+^{(t)} + \delta Q_+^{(t)}$, and assess the rate of growth of the perturbation along the iterations by the functional F . In order to implement this idea in practice one needs to choose a specific form for the perturbation; given that $Q_+^{(t)}$ is represented as a sum of Dirac deltas we perturb it by giving an infinitesimal width to each of the peaks, that we replace by Gaussian distributions with a small variance. We thus define

$$Q_+^{(t)}(h) \approx \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \delta(h - h_i^{(t)}), \quad (Q_+^{(t)} + \delta Q_+^{(t)})(h) \approx \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \mathcal{G}(h; h_i^{(t)}, \varepsilon_i^{(t)}), \quad (62)$$

$$\widehat{Q}_+^{(t)}(u) \approx \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \delta(u - u_i^{(t)}), \quad (\widehat{Q}_+^{(t)} + \delta \widehat{Q}_+^{(t)})(u) \approx \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \mathcal{G}(u; u_i^{(t)}, \widehat{\varepsilon}_i^{(t)}), \quad (63)$$

where $\mathcal{G}(\cdot; a, b)$ denotes the density of a Gaussian random variable of average a and variance b . Consider now the insertion of the form (63) in the right hand side of (38); the choice of the d peaks indexed by i_1, \dots, i_d produces a random variable h equal in distribution to $f(u_{i_1} + \sqrt{\widehat{\varepsilon}_{i_1}} z_1, \dots, u_{i_d} + \sqrt{\widehat{\varepsilon}_{i_d}} z_d)$, where z_1, \dots, z_d are independent standard

Gaussians (of zero mean and unit variance). As the $\widehat{\varepsilon}$ are infinitesimally small one can linearize f to compute the mean and variance of this random variable.

In summary, the determination of $\lambda(\alpha)$ is done by tracking the evolution of Q_+ , \widehat{Q}_+ and their perturbed versions with populations of couples of real numbers, (h_i, ε_i) and $(u_i, \widehat{\varepsilon}_i)$, that evolve in time according to the following generalization of the update rules given in Sec. IID 2. To obtain $(h_i^{(t+1)}, \varepsilon_i^{(t+1)})$ one repeats, independently for $i = 1, \dots, \mathcal{N}$, these steps:

- draw an integer d from the law p_d
- draw d indices i_1, \dots, i_d uniformly at random in $\{1, \dots, \mathcal{N}\}$
- set $h_i^{(t+1)} = f(u_{i_1}^{(t)}, \dots, u_{i_d}^{(t)})$ and $\varepsilon_i^{(t+1)} = \sum_{j=1}^d (\partial_j f)^2 \widehat{\varepsilon}_{i_j}^{(t)}$, where $\partial_j f$ denotes the derivative of f with respect to its j -th argument, computed in $(u_{i_1}^{(t)}, \dots, u_{i_d}^{(t)})$

Similarly the population $(u_i^{(t)}, \widehat{\varepsilon}_i^{(t)})$ is generated according to, again independently for $i = 1, \dots, \mathcal{N}$:

- draw $\sigma_1, \dots, \sigma_{k-1}$ from the probability law $\widetilde{p}(\sigma_1, \dots, \sigma_{k-1}|+)$
- draw $k-1$ indices i_1, \dots, i_{k-1} uniformly at random in $\{1, \dots, \mathcal{N}\}$
- set $u_i^{(t)} = g(\sigma_1 h_{i_1}^{(t)}, \dots, \sigma_{k-1} h_{i_{k-1}}^{(t)})$ and $\widehat{\varepsilon}_i^{(t)} = \sum_{j=1}^{k-1} (\partial_j g)^2 \varepsilon_{i_j}^{(t)}$

The rate of growth of the perturbation during the iteration $t \rightarrow t+1$ is estimated as the ratio of the L_1 norms of the perturbation parameters,

$$\lambda_t = \frac{\sum_{i=1}^{\mathcal{N}} \varepsilon_i^{(t+1)}}{\sum_{i=1}^{\mathcal{N}} \varepsilon_i^{(t)}} , \quad (64)$$

and the stability parameter is finally computed as

$$\lambda(\alpha) = (\lambda_{t_0} \lambda_{t_0+1} \dots \lambda_{t_0+n-1})^{\frac{1}{n}} . \quad (65)$$

Indeed the first t_0 iterations are done with the usual population dynamics algorithm, evolving only the h_i 's and u_i 's, in order to reach an approximate convergence in distribution of the populations to their fixed points, and the perturbation is then initialized with $\varepsilon_i^{(t_0)} = 1$. A large number n of additional iterations during which the growth rates are recorded are then performed, and averaged geometrically as in (61); in the large n limit the value of $\lambda(\alpha)$ should be independent of the norm used to define λ_t . In practice we divide the $\varepsilon_i^{(t+1)}$ by λ_t after each iteration in order to keep the norm constant and avoid numerical underflows.

This method is similar to the one presented in [71, 72] to determine the location of a continuous RSB transition from a non-trivial RS solution.

IV. RESULTS OF THE CAVITY METHOD

A. The existence of a RS phase for $\alpha > \alpha_{d,u}$

We shall address now the main question raised in the introduction, namely the evolution of the dynamic phase transition when the measure $\mu(\underline{\sigma})$ over the proper bicolourings of a typical Erdős-Rényi random hypergraph is not uniform anymore. In the setting considered in this paper this corresponds to take the parameters $\{\omega_p\}$ of the interaction function (2) different from the uniform choice $\omega_0 = \omega_k = 0$, $\omega_1 = \dots = \omega_{k-1} = 1$.

We will concentrate first on the “zero-temperature” case, i.e. on the measures that give a non-zero weight to proper bicolourings only, which implies $\omega_0 = \omega_k = 0$. The choice of the other parameters is constrained by the global spin-flip symmetry that we want to preserve, hence $\omega_p = \omega_{k-p}$; as it is obvious from (1), multiplying all the ω_p by a common constant does not change the properties of the model. One realizes that for $k=3$ there is no free parameter left, we will thus concentrate on the cases $k \geq 4$ from now on. For arbitrary large values of k there will be of the order of $k/2$ free parameters in the ω_p ; we will however make the following choice for the zero-temperature measure:

$$\omega_0 = \omega_k = 0 , \quad \omega_1 = \omega_{k-1} = 1 - \epsilon , \quad \omega_2 = \dots = \omega_{k-2} = 1 , \quad (66)$$

k	$\alpha_{d,u}$	α_{opt}	ϵ_{opt}
4	4.083	4.578	-0.10
5	9.465	9.636	0.06
6	18.088	18.879	0.12

TABLE I: Dynamic threshold for the uniform measure ($\alpha_{d,u} = \alpha_d(\epsilon = 0)$), and largest α reachable in the RS phase, this optimal point having coordinates $(\alpha_{opt}, \epsilon_{opt})$.

where ϵ is the sole parameter that quantifies the deviation from the uniform measure (that is recovered for $\epsilon = 0$). This slight loss in generality is made for the sake of simplicity, and motivated by considerations on the large k limit presented in Sec. VI. The parameter ϵ controls the relative weight given to the “almost monochromatic” constraints that contain a single vertex of one of the possible colors (positive values of ϵ disfavoring them); as discussed in Sec. IID 4 these are precisely those responsible for the existence of frozen variables, one of the mechanism of RSB.




FIG. 7: Phase diagram for $k = 4$, $k = 5$ and $k = 6$ (from left to right), in the plane (ϵ, α) , at zero temperature $\omega_0 = 0$. The RS phase, painted in gray, is on the left of $\alpha_d(\epsilon)$, the latter corresponds either to a continuous transition with $\alpha_d(\epsilon) = \alpha_{KS}(\epsilon)$ (solid line, see (68)) or to a discontinuous transition (black squares). The dashed horizontal line corresponds to $\epsilon = 0$, the uniform measure, which intersects α_d at $\alpha_{d,u}$. The arrow points to the optimal point that maximizes α_d . The dotted line is the rigidity threshold α_r from (69), the dot-dashed line marks the vanishing of the RS entropy (see Eq. (67)).

We present in Fig. 7 phase diagrams in the (α, ϵ) plane for $k = 4$, $k = 5$ and $k = 6$. The three lines in these plots correspond to the thresholds defined in (20) from the vanishing of the RS entropy, in (46) from the instability of the RS solution (Kesten-Stigum threshold), and in (51) from the appearance of hard fields in the solution of the 1RSB equations at $m = 1$ (rigidity threshold); specializing these three expression with the choice of parameters (66) yields

$$\alpha^{s=0}(k, \epsilon) = \frac{\ln 2}{\frac{k(1-\epsilon) \ln(1-\epsilon)}{2^{k-1}-1-k\epsilon} - \ln\left(1 - \frac{1+k\epsilon}{2^{k-1}}\right)}, \quad (67)$$

$$\alpha_{KS}(k, \epsilon) = \frac{1}{k(k-1)} \left(\frac{2^{k-1} - 1 - k\epsilon}{1 + (k-4)\epsilon} \right)^2, \quad (68)$$

$$\alpha_r(k, \epsilon) = \frac{1}{k} \Gamma_r(k) \frac{2^{k-1} - 1 - k\epsilon}{1 - \epsilon}. \quad (69)$$

In addition the black squares in Fig. 7 signal a discontinuous appearance of a non-trivial solution of the 1RSB equations at $m = 1$ upon increasing α , that we located by a numerical resolution of these equations following the methods explained in Sec. III. One can see on these plots that for all values of ϵ there is a critical density of constraints, $\alpha_d(\epsilon)$, such that a non-trivial solution of the 1RSB equations at $m = 1$ exist if and only if $\alpha > \alpha_d(\epsilon)$. To make this separation more visible the area on the left of $\alpha_d(\epsilon)$, i.e. the RS phase of the model, has been painted in gray in Fig. 7. Let us call $(\alpha_{opt}, \epsilon_{opt})$ the coordinates of the point on the line $\alpha_d(\epsilon)$ which maximizes the density α of constraints, $\alpha_{opt} = \max_{\epsilon} \alpha_d(\epsilon)$, that corresponds to an optimal choice of the bias parameter. The numerical values of these optimal parameters can be found in Table I for $k = 4, 5$ and 6 . By definition $\alpha_{opt} \geq \alpha_{d,u} = \alpha_d(\epsilon = 0)$, the dynamic transition of the usual model, with the uniform measure over the proper bicolorings; the non-trivial result here is that the inequality is strict, i.e. that a well chosen value of the biasing parameter ϵ is able to turn the clustered uniform measure into an unclustered biased one (for $\alpha \in [\alpha_{d,u}, \alpha_{opt}]$).

A further scrutiny of the phase diagrams reveals different scenarios depending on the value of k . For $k = 4$ the nature of the bifurcation on the line $\alpha_d(\epsilon)$ changes precisely at ϵ_{opt} : for $\epsilon > \epsilon_{opt}$ the transition is continuous and thus

α_d coincides with the Kesten-Stigum line α_{KS} , while it is discontinuous for $\epsilon < \epsilon_{opt}$ and there is a cusp at the optimal point (we shall come back on this point later on). It turns out that for $k = 4$, $\epsilon_{opt} < 0$: this is rather counterintuitive at first sight, as it means that favoring the almost violated configurations of variables actually makes the measure less frustrated. This peculiarity can be explained by noticing that for $k = 4$ the dynamic transition of the uniform measure ($\epsilon = 0$) is continuous and that α_{KS} decreases with ϵ . As the dynamic transition of the uniform measure is discontinuous for $k \geq 5$ [70] this peculiarity is restricted to $k = 4$, and one has $\epsilon_{opt}(k \geq 5) > 0$. Turning now to the phase diagram for $k = 5$ in Fig. 7 one observes similarly a cusp in $\alpha_d(\epsilon)$ at ϵ_{opt} , that separates a continuous and discontinuous branch of the dynamic transition line, but with now $\epsilon_{opt} > 0$. Finally for $k = 6$ the optimal point is on the discontinuous branch of $\alpha_d(\epsilon)$; increasing further ϵ one encounters a cusp at some value of $\epsilon > \epsilon_{opt}$ and then a continuous branch $\alpha_d(\epsilon) = \alpha_{KS}(\epsilon)$. The large k behavior of the model will be further discussed in Sec. VI; we can nevertheless anticipate that for large enough k the Kesten-Stigum threshold becomes irrelevant, as it happens in the negative RS entropy region (compare the leading orders of Eqs. (67,68)). In this case the whole line $\alpha_d(\epsilon)$ corresponds to a discontinuous bifurcation. As a last remark on the phase diagrams of Fig. 7 let us emphasize that for all ϵ one has necessarily $\alpha_d(\epsilon) \leq \min(\alpha_{KS}(\epsilon), \alpha_r(\epsilon), \alpha^{s=0}(\epsilon))$, these three thresholds implying a mechanism of failure for the hypotheses underlying a purely RS phase. This should easily convince the reader of the necessity of discontinuous branch of $\alpha_d(\epsilon)$ in some parts of the phase diagrams. For instance when $k = 4$ and $\epsilon \leq -0.3$ the rigidity and negative entropy bounds imply $\alpha_d(\epsilon) < \alpha_{KS}(\epsilon)$, in other words the dynamic transition must be discontinuous.

We have motivated earlier our study of the boundaries of the RS phase by algorithmic considerations, Monte Carlo Markov Chains being expected to equilibrate rapidly inside such a phase. However in a practical simulation one cannot assume that the initial configuration belongs to the support of a zero-temperature measure (otherwise the problem of finding a solution of the CSP would be already solved), it is thus necessary to make an annealing in temperature for a random initial condition to be allowed. For this reason we have also studied the evolution of the phase diagrams at positive temperature, modifying the parameters (66) with $\omega_0 = \omega_k > 0$, see the results in Fig. 8. These plots show the absence of “reentrance” in temperature, in the sense that the lines $\alpha_d(\epsilon)$ move towards higher density of constraints when ω_0 is increased. Hence in principle a simulated annealing procedure with parameters (α, ϵ) in the zero temperature RS domain, progressively decreasing ω_0 , should be able to remain equilibrated on polynomial time scales, hence finding solutions for $\alpha < \alpha_{opt}$ if the appropriate bias is used. A numerical test of this conjecture is presented in Sec. V.




FIG. 8: Phase diagrams for $k = 4$, $k = 5$ and $k = 6$ (from left to right) in the plane (α, ϵ) , giving the RS phase delimited by the KS bound and a dynamic line where the transition toward a non-trivial solution is discontinuous, for different temperatures. Left ($k = 4$): the dynamic line is given from left to right for $\omega = 0$ (filled square and solid line), $\omega = 0.1$ (empty triangle and dashed line), $\omega = 0.2$ (filled triangle and dotted line). Middle ($k = 5$): from left to right $\omega = 0$, $\omega = 0.002$, $\omega = 0.005$, $\omega = 0.01$. Right ($k = 6$): from left to right $\omega = 0$, $\omega = 0.005$, $\omega = 0.01$.

B. More detailed zero temperature phase diagrams

The extent of the RS domain in the (α, ϵ) phase diagram presented in Fig. 7 was the most interesting information to extract from the cavity formalism in the perspective of this paper. For the sake of completeness we shall nevertheless discuss with slightly more details some properties of the RSB phase, and present another version of the phase diagrams for $k = 4$ and $k = 5$ in Fig. 9.

The most important additional feature unveiled by these phase diagrams is that for some values of the parameters k , α , ϵ , there exists (at least) two different non-trivial solutions of the 1RSB equations at $m = 1$ (38). This type of behavior was described in [70] for a family of random CSPs generalizing the hypergraph bicoloring, and its consequences for inference problems (or planted CSPs) have been discussed in [69]. In order to reach numerically these different




FIG. 9: Phase diagram for $k = 4$ (left), and $k = 5$ (right), in the plane (α, ϵ) , at zero temperature ($\omega_0 = 0$): the solid line is the Kesten-Stigum (KS) transition where a non-trivial solution of the 1RSB equations emerge continuously from the trivial one upon increasing α , the filled (resp. empty) squares corresponds to the spinodal (sp) of the HO (resp. LO) branch that appears discontinuously when α is increased (resp. decreased). The filled circles are defined by the vanishing of the complexity of the HO branch.

solutions we used the population dynamics algorithm explained in Sec. IID 2 with an initial condition generalizing (45) into

$$Q_+^{(t=0)}(h) = (1 - \epsilon) \delta(h) + \epsilon \delta(h - 1) . \quad (70)$$

For each choice of the parameters we ran twice the population dynamics algorithm, once with $\epsilon = 1$ and once with a small value of $\epsilon > 0$ (in practice we used $\epsilon = 0.01$); in the tree reconstruction interpretation the latter correspond to a variant known as robust tree reconstruction [66], in which only a fraction ϵ of the variables at large distance from the root are revealed to the observer. We will call HO, for high overlap, the initialization with $\epsilon = 1$, and LO (low overlap) the small ϵ one. Depending on the parameters these two procedures can produce different solutions of the 1RSB equations, or not. More precisely, the different phases located in Fig. 9 are defined as follows:

- RS: both HO and LO initial conditions lead to the trivial solution.
- RS+HO: LO initial condition leads to the trivial solution, whereas HO initial condition leads to a non trivial solution.
- LO+HO: LO initial condition leads to a non trivial solution, HO initial condition leads to a non trivial solution with a higher overlap.
- HO/LO: both HO and LO initial conditions lead to the same non-trivial solution.

The frontiers between these different phases are:

- α_{KS} , the limit of stability of the trivial fixed point, that undergoes a bifurcation at the Kesten-Stigum transition.
- two spinodal lines (denoted sp HO and sp LO) that correspond to the limit of existence of the two non-trivial branches of solution of the 1RSB equations.

We invite the reader to consult also the top panels of Fig. 10 where the evolution of the overlap is plotted as a function of ϵ for different fixed α at $k = 5$, which should help to grasp the meaning and succession of the different phases. The frontiers between the various phases are indicated with the same names in Fig. 9 and Fig. 10.

This more complete study of the number and domain of existence of solutions of the 1RSB equations should clarify the cusp at ϵ_{opt} of the line $\alpha_d(\epsilon)$ found for $k = 4$ and $k = 5$ in Fig. 7: a first look at these figures could suggest that the two parts of the $\alpha_d(\epsilon)$ line join at a tri-critical point, in the sense that the discontinuous transition becomes less and less discontinuous before crossing over to a continuous transition. However from Fig. 9 one sees that this is not the case, the discontinuous branch of $\alpha_d(\epsilon)$ extends to the RSB phase as a spinodal unrelated to the continuous transition, which can only make sense in the context of coexistence of two non-trivial solutions. A tri-critical point does exist in these phase diagrams, but it is located strictly inside the RSB phase, not at the cusp, and corresponds to the merging of the two spinodals.

Finally we have also indicated on Fig. 9 the threshold for the cancellation of the complexity of the HO solution (see in addition the bottom panels in Fig. 10); this corresponds to the condensation transition of the model in the $\alpha < \alpha_{KS}$ part of the phase diagram (for $\alpha > \alpha_{KS}$ the LO solution has a negative complexity hence the problem is condensed, see [70] for a discussion of this point).




FIG. 10: Overlaps (top) and complexities (bottom) versus ϵ at $k = 5$, for $\alpha = 9.5$ (left), $\alpha = 9.6$ (center), and $\alpha = 9.7$ (right).

V. RESULTS OF SIMULATED ANNEALING

In this Section we present the results of extensive simulations, where we have used the Simulated Annealing (SA) algorithm [50] to find solutions of the hypergraph bicoloring problem. Our main aim is to show that SA finds solutions more easily if the biased measure is used: although the uniform measure ($\epsilon = 0$) has a larger entropy, the biased one is more concentrated on solutions that can be reached in an easier way by SA and thus the SA algorithmic threshold improves if $\epsilon \neq 0$ is used.

As in the rest of paper, we consider Erdős-Rényi random hypergraphs with $k = 4, 5, 6$ and sizes ranging from $N = 10^4$ to $N = 10^6$. The parameters α and ϵ are taken in the relevant region where we expect an algorithmic phase transition to take place, that is around α_d . Let us rewrite the biased measure that we are willing to sample via the SA at a generic finite temperature $T = 1/\beta$ as

$$\mu(\underline{\sigma}) = \frac{1}{Z(G)} e^{-\beta U(\underline{\sigma})} (1 - \epsilon)^{F(\underline{\sigma})} \quad (71)$$

where $U(\underline{\sigma})$ is the number of unsatisfied constraints (i.e. monochromatic hyperedges) and $F(\underline{\sigma})$ is the number of freezing clauses (i.e. hyperedges with exactly $k - 1$ variables of the same color). This corresponds to the measure (1) with parameters

$$\omega_0 = \omega_k = e^{-\beta}, \quad \omega_1 = \omega_{k-1} = 1 - \epsilon, \quad \omega_2 = \dots = \omega_{k-2} = 1. \quad (72)$$

The solutions of the CSP have $U = 0$, and non-uniform weights if $\epsilon \neq 0$. Our SA implementation uses the Metropolis algorithm with single-spin flip dynamics: at each time step one considers a configuration $\underline{\sigma}'$ that differs from the current configuration $\underline{\sigma}$ by the reversal of a uniformly chosen spin. The move $\underline{\sigma} \rightarrow \underline{\sigma}'$ is accepted with the probability

$$\min\left(\frac{\mu(\underline{\sigma}')}{\mu(\underline{\sigma})}, 1\right) = \min\left(e^{-\beta \Delta U} (1 - \epsilon)^{\Delta F}, 1\right), \quad (73)$$

where $\Delta U = U(\underline{\sigma}') - U(\underline{\sigma})$ and $\Delta F = F(\underline{\sigma}') - F(\underline{\sigma})$, in such a way that the detailed balance (reversibility) condition with respect to the measure (71) is ensured. We store the quantity $\sum_{i=1}^k \sigma_i$ for each clause, which allows a fast computation of the changes ΔU and ΔF when a spin is flipped.




FIG. 11: Lowest intensive energy reached by SA for $k = 4$ (left) and $k = 5$ (right) as a function of α for different cooling times τ . For each cooling time τ we show results for 3 problem sizes: $N = 10^4$ with filled circles, $N = 10^5$ with empty circles and $N = 10^6$ with triangles. The latter two values do always coincide (except for $k = 5$ and $\alpha = 9.8$, where the $N = 10^6$ datapoint is missing). Estimating the algorithmic threshold from these plots is very difficult due to the strong τ dependency.

We run SA with a very simple piecewise constant, uniformly spaced, temperature scheduling: the first Monte Carlo Sweep (MCS, i.e. N elementary steps described above) is performed with T fixed to T_{\max} (we used $T_{\max} = 0.5$ in all our simulations), then T is reduced by $\Delta T = T_{\max}/\tau$ and a new MCS is performed, T is again reduced by ΔT , and so on and so forth. We perform in this way $\tau + 1$ MCS, the last one being at zero temperature, the running time of the algorithm thus scales as $N\tau$ elementary steps.

The lowest value of $U(\underline{\sigma})$ is always reached at the end of each run, when the annealing has reached zero temperature. So we present results only for the quantity $U_0 = U(T = 0)$, that is the smallest number of violated clauses that the SA is able to reach in a running time of τ MCS. SA is successful as a solver if and only if $U_0 = 0$, but we will be interested in estimating the lowest energy reachable by SA even in the regime where it is not successful. In particular we are going to study the lowest intensive energy reached by SA, $u_0 = U_0/N$, in the large size limit where it becomes independent on the problem size N .

A. Estimating the algorithmic threshold for Simulated Annealing

We shall first discuss the problem of the estimation of the algorithmic threshold for a stochastic algorithm like SA, concentrating for simplicity on the unbiased ($\epsilon = 0$) case, the extension to $\epsilon \neq 0$ will be considered later on.

The behavior of the algorithm depends on the density of constraints α , the annealing time τ , and the size of the problem N ; it can be described in terms of the average energy density $u_0(\alpha, \tau, N)$ reached at the end of the run, or in terms of the probability (with respect to the random instance generation and the stochasticity of the algorithm) $p_{\text{succ}}(\alpha, \tau, N)$ that the algorithm discovers a solution of the instance. It is clear that the energy (resp. success probability) reached by SA is a decreasing (resp. increasing) function of the running time τ . We are interested in the limit of large times but sub-exponential with respect to the problem size N (on exponentially large timescales any Monte Carlo simulation of a finite size system is ergodic and $u_0 = 0$ as long as $\alpha < \alpha_{\text{sat}}$, but this is not the regime we are interested in). An idealized definition of the algorithmic threshold α_{algo} would be the smallest density of constraints such that

$$\lim_{N \rightarrow \infty} u_0(\alpha, \tau = N^c, N) > 0 \quad \text{or} \quad \lim_{N \rightarrow \infty} p_{\text{succ}}(\alpha, \tau = N^c, N) = 0, \quad (74)$$

for any fixed exponent c , corresponding to polynomial time algorithms. Of course time and space requirements impose strong constraints on the values of τ and N that can be used in practice. The limit above must thus be performed by an extrapolation from finite N results, and if c is free any running time could be considered as “polynomial” as long as N is finite. To resolve this ambiguity we shall restrict our study to linear times (this time scale is the only one practically accessible on very large problems), i.e. consider τ fixed (but arbitrary large) in the thermodynamic limit $N \rightarrow \infty$.

Even with this restriction the numerical extrapolation necessary to estimate α_{algo} is far from being an easy task. The definition given above relies on the behavior of the asymptotic intensive energy u_0 as a function of α . We plot the corresponding data in Figure 11 for $k = 4, 5$: these data have been obtained for the unbiased measure ($\epsilon = 0$)

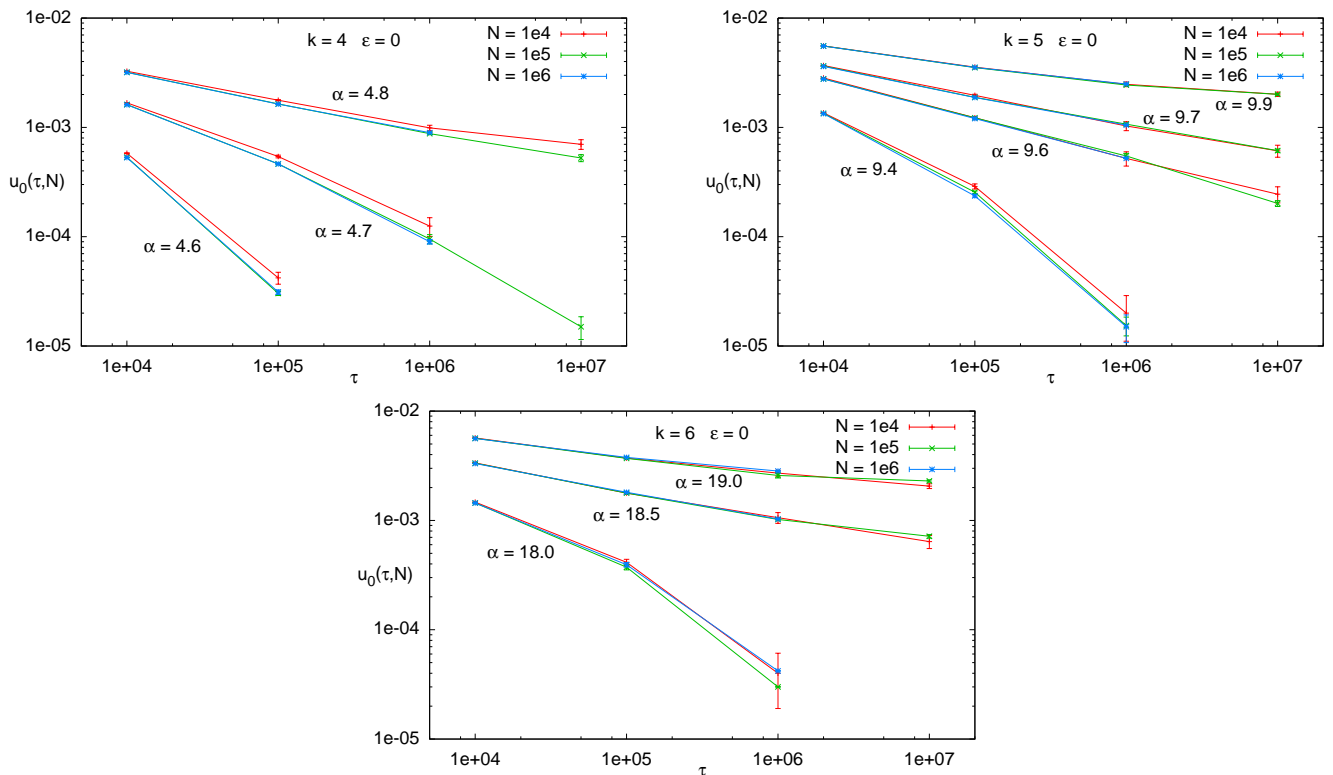


FIG. 12: Plotting the lowest intensive energy reached by SA with no bias ($\epsilon = 0$) as a function of τ provides a better way to estimate the algorithmic threshold α_{alg} : the decay is faster (resp. slower) than a power law if $\alpha < \alpha_{\text{alg}}$ (resp. $\alpha > \alpha_{\text{alg}}$).

and different problem sizes ($10^4 \leq N \leq 10^6$); note that data points with different N values are very close, i.e. the size dependence is very weak, and the values for $N = 10^5$ and $N = 10^6$ always coincide within errorbars. Unfortunately the asymptotic energy u_0 is strongly dependent on the running time τ and it is thus very difficult to extract from this figure the algorithmic threshold, i.e. the value of α where u_0 becomes positive in the large τ limit.

A much more convenient way of analyzing the same data is presented in Figure 12, where for each value of α we study the dependence of u_0 on τ . The α values shown are such that the relative difference between the smallest and the largest α values is around 5%. Again the size dependence is weak and we can mostly ignore it. Error bars have been computed only from sample to sample fluctuations. The main observation now — note the log-log scale in the plots — is that for the smallest α values shown in the plots the asymptotic energy is decreasing very fast with τ , faster than a power law (data not shown have $u_0 \simeq 0$); on the contrary, for the largest α values, u_0 decreases slower than a power law. In the latter case we even observe an upwards curvature, suggesting a non-zero value for u_0 in the $\tau \rightarrow \infty$ limit.

In practice, our best estimate for the SA algorithmic threshold is given by the α value such that u_0 decays as an inverse power law of τ , thus separating the regimes where u_0 decays faster and slower than a power law in τ . For $\epsilon = 0$, we find the following approximate values $\alpha_{\text{alg}}(k = 4) \approx 4.7$, $\alpha_{\text{alg}}(k = 5) \approx 9.6$ and $\alpha_{\text{alg}}(k = 6) \approx 18.5$. We notice that all these algorithmic thresholds are larger than the threshold $\alpha_{\text{d,u}}$ listed in Table I, where the “dynamic” phase transition, defined as the appearance of a solution of the 1RSB equations, takes place. This observation is consistent with the idea that sampling solutions uniformly is more difficult than just finding one or few solutions. Indeed, while a MCMC is expected to sample uniformly the solutions efficiently only for $\alpha < \alpha_{\text{d,u}}$, SA can find a solution in linear time until α_{alg} , which is greater than $\alpha_{\text{d,u}}$. In other words, in the range $[\alpha_{\text{d,u}}, \alpha_{\text{alg}}]$ the SA algorithm does not thermalize at the lowest temperatures explored during the annealing, but it can be seen as an efficient out of equilibrium process converging in linear time to a solution, as discussed for instance in [73].

In particular for $k = 4$ the model has a continuous phase transition and the SA algorithm seems to be very efficient in this case: the algorithmic threshold $\alpha_{\text{alg}}(k = 4) \approx 4.7$ is well beyond the dynamic threshold $\alpha_{\text{d,u}} = 4.083$ and not far from the 1RSB estimate of the satisfiability threshold $\alpha_{\text{sat}}(k = 4) \approx 4.9$ [70] (this is only expected to be an upperbound on the true satisfiability threshold due to an instability towards higher levels of RSB). On the contrary for $k \geq 5$ the phase transition taking place at $\alpha_{\text{d,u}}$ is of the random first order type (discontinuous) and this seems to




FIG. 13: The fast decays of u_0 as a function of τ for $\alpha = \alpha_{\text{opt}}$ and $\epsilon = \epsilon_{\text{opt}}$ confirms that for these “optimal” parameters SA is effective in finding the ground state in linear time.

have a dramatic effect on the performance of SA, which is able to find solutions only slightly beyond $\alpha_{\text{d,u}}$, stopping far from the α_{sat} threshold. For example, for $k = 5$ we have $\alpha_{\text{d,u}} = 9.465$, $\alpha_{\text{alg}} \approx 9.6$ and $\alpha_{\text{sat}} = 10.46$ [70].

B. Performances of Simulated Annealing with optimal RS parameters

As shown in Sec. IV A we can extend the RS phase in the region $\alpha > \alpha_{\text{d,u}}$ by tuning appropriately the bias ϵ . In Figure 13 we show the asymptotic energy as a function of the running time τ for the parameters $\alpha = \alpha_{\text{opt}}$ and $\epsilon = \epsilon_{\text{opt}}$ given in Table I, that are optimal from the point of view of extending the RS phase to the largest α possible. For each value of τ we report the results obtained with 3 sizes $N = 10^4, 10^5, 10^6$ although the different data points are hardly visible due to their strong overlap (for $\tau = 10^6$ fewer sizes are shown). Errors are computed from sample to sample fluctuations.

For all the values of $k = 4, 5, 6$ the behavior of the asymptotic energy is compatible with a power law decay or even faster than that (the straight line is just a guide to the eye with a slope -0.61). So, as expected, SA seems to be an efficient algorithm to find solutions in the RS phase, even when this phase extends beyond $\alpha_{\text{d,u}}$ via the optimization of the bias ϵ .

C. Performances of Simulated Annealing with the biased measure

In Section IV we have shown how the phase diagram and the corresponding thresholds change in presence of a non-zero bias ($\epsilon \neq 0$). The suggestion we get from this analytical study is that a non-zero bias should make easier for the SA algorithm to find solutions at large α values. However the connection between the phase diagram and the behavior of the SA is not obvious, as already shown in Section V A for the $\epsilon = 0$ case.

The aim of the present section is to show the results of extensive numerical simulations running SA with the biased measure in order to gather evidence that a non-zero bias is in general beneficial for the performances of SA in finding a solution to the random hypergraph bicoloring problem.

We have already shown that finite size effects are very small and slightly visible only for $N = 10^4$. So in the following we present uniquely data obtained with size $N = 10^5$. We have checked these are practically indistinguishable from the results with $N = 10^6$ on the time scales reachable in the latter case.

In Figure 14 we show the data collected at the three largest α values for $k = 4$ (upper row), $k = 5$ (middle row) and $k = 6$ (lower row). In each panel we plot u_0 as a function of ϵ for three different cooling rates $\tau = 10^4, 10^5, 10^6$ (from top to bottom in each panel). The plots provides a clear evidence that reweighting solutions with a bias $\epsilon > 0$ enhances the probability that SA reaches lower energies.

Already a simple qualitative analysis reveals the advantage of using $\epsilon > 0$. In every panel we see that u_0 reaches a minimum for a strictly positive value of ϵ . The value ϵ_{SA} that minimizes u_0 is only weakly dependent on the SA




FIG. 14: The lowest intensive energy reached by SA reaches its minimum for a positive bias parameter ϵ .

cooling time τ , so it is likely to assume that $\lim_{\tau \rightarrow \infty} \epsilon_{\text{SA}} > 0$ and the bias is effective even in the limit of large times.

The data in Figure 14 suggest that the SA algorithmic threshold may grow for moderately small values of ϵ with respect to its $\epsilon = 0$ value. For example for $k = 6$ the SA algorithmic threshold for $\epsilon = 0$ was estimated around $\alpha_{\text{alg}} \approx 18.5$, but looking at the plots in the lower row it is evident that at least for $\alpha = 19$ and $\epsilon \simeq 0.15$ SA reaches the ground state $u_0 = 0$ and for $\alpha = 19.5$ and $\epsilon \simeq 0.15$ the convergence to $u_0 = 0$ is very fast in τ .

We have done a more quantitative analysis for the largest α values, shown in the three panels on the right. We have interpolated the data of u_0 with a quadratic function of ϵ , the interpolating parabolas are shown in the right panels in Figure 14. Fitting the minimum of the parabola $u_0(\epsilon_{\text{SA}})$ as a power law in τ we find the results shown in Figure 15. For $(k = 5, \alpha = 9.9)$ the behavior is faster than a power law and thus we expect the SA algorithmic threshold with the biased measure to be slightly greater than $\alpha = 9.9$. On the contrary for both $(k = 4, \alpha = 4.8)$ and $(k = 6, \alpha = 20)$ the behavior looks slightly slower than a power law and thus we are tempted to believe $\lim_{\tau \rightarrow \infty} u_0(\epsilon_{\text{SA}}) > 0$ in these cases and the SA algorithmic threshold is slightly below.




FIG. 15: The lowest energy reached with the optimal bias ϵ_{SA} for the largest α value simulated as a function of τ .

Unfortunately the quantitative analysis cannot be made more robust, due to the strong τ dependence observed. Nevertheless we believe that the evidence that $\epsilon > 0$ makes ground states more accessible to Simulated Annealing is strong enough, both for finite τ values and in the large τ limit.

Let us finally compare the optimal value ϵ_{SA} of the bias that improves most the performances of SA with the optimal value ϵ_{opt} found in Section IV A that increases most the extent of the RS phase. We notice that ϵ_{SA} is always larger than ϵ_{opt} . Considering that, for the small values of k studied in the simulations, it is approximately true that for $\epsilon > \epsilon_{opt}$ the phase transition increasing α is continuous, while for $\epsilon < \epsilon_{opt}$ the model undergoes a random first order transition, we believe the most natural explanation for the observation $\epsilon_{SA} > \epsilon_{opt}$ is the following. The ergodicity breaking taking place at a discontinuous (i.e. random first order) transition is much more severe than the one taking place at a continuous phase transition. In the case of a discontinuous phase transition, the SA algorithm can find solutions only slightly above the dynamic threshold α_d , while in the continuous case SA remain efficient in finding solutions even well above α_{KS} . The analysis supporting this scenario has been presented in Section V A. Thus it is natural to expect that SA presents its best performances for $\epsilon > \epsilon_{opt}$ where the phase transition is continuous and the ergodicity breaking not too severe. The finding that $\epsilon_{SA}(k=4) > 0$ also resolves the rather counterintuitive result $\epsilon_{opt}(k=4) < 0$. So even for $k=4$ the SA algorithm finds the ground state more easily if frozen variables are avoided.

VI. LARGE k ASYMPTOTICS

The numerical resolution of the cavity equations presented in Sec. IV shows that for small values of k one has $\alpha_{opt}(k) > \alpha_{d,u}(k)$, in other words that distorting the measure over solutions can make it RS for larger densities of constraints than the uniform one. We want now to investigate the large k limit, for which the gap between the satisfiability threshold and the algorithmic ones is most clearly demonstrated. As $\alpha_{opt}(k)$ is defined from a discontinuous bifurcation of a functional equation (in other words a reconstruction problem for which the Kesten-Stigum bound is not tight) we do not have an analytical expression for it, we will in consequence aim at a more modest objective, namely deriving asymptotic bounds on its large k behavior.

As explained before for any ϵ one has $\alpha_d(\epsilon) \leq \min(\alpha_{KS}(\epsilon), \alpha_r(\epsilon), \alpha^{s=0}(\epsilon))$, these three upperbounds having simple expressions given in Eqs. (67-69). At large enough k it is easy to convince oneself that the Kesten-Stigum transition occurs after $\alpha^{s=0}$, hence is completely irrelevant (the dominant term in the asymptotic expansion of α_{KS} is of the order $2^{2(k-1)}$ instead of 2^{k-1} for $\alpha^{s=0}$ and α_r). We show in Fig. 16 the lines $\alpha_r(\epsilon), \alpha^{s=0}(\epsilon)$ for a large value of k , as well as a guess on the qualitative behavior of $\alpha_d(\epsilon)$. We define $(\alpha_*(k), \epsilon_*(k))$ as the coordinates of the intersection of the rigidity and the zero-entropy line, in such a way that $\alpha_{opt}(k) \leq \alpha_*(k)$: for larger densities either the RS entropy is negative, or there exists a 1RSB solution with hard-fields (or both), in any case no RS phase can exist for $\alpha \geq \alpha_*$. We will now derive an asymptotic expansion at large k of this upperbound $\alpha_*(k)$.

In the expression (69) of α_r the coefficient $\Gamma_r(k)$ is a series depending solely on k , that was defined in Eq. (50); in order to obtain more easily its asymptotic equation we define a series w_k by $q_r(k) = 1 - e^{-w_k}$, in such a way that w_k




FIG. 16: For $k = 20$, plot in the plane (α, ϵ) of the rigidity line and zero RS entropy line, that intersect at the point $(\alpha_*(20), \epsilon_*(20))$. The gray zone is a qualitative guess of the RS zone delimited by $\alpha_d(k, \epsilon)$.

is solution of the implicit equation $e^w = 1 + (k-1)w$. Taking the logarithm of this equation and iterating once yields

$$w_k = \ln k + \ln \ln k + O\left(\frac{\ln \ln k}{\ln k}\right). \quad (75)$$

One can then compute

$$\Gamma_r(k) = w_k \left(1 + \frac{1}{(k-1)w_k}\right)^{k-1} = \ln k + \ln \ln k + O(1). \quad (76)$$

This gives immediately the expansion of the rigidity threshold for the uniform measure ($\epsilon = 0$),

$$\alpha_{r,u}(k) = 2^{k-1} \frac{1}{k} (\ln k + \ln \ln k + O(1)). \quad (77)$$

We come back to the determination of $(\alpha_*(k), \epsilon_*(k))$; this intersection of the rigidity and zero-entropy line is solution of the two following equations, immediately obtained from (67,69):

$$\alpha = \frac{\ln 2}{\frac{k(1-\epsilon)\ln(1-\epsilon)}{2^{k-1}-1-k\epsilon} - \ln\left(1 - \frac{1+k\epsilon}{2^{k-1}}\right)} = \frac{1}{k} \Gamma_r(k) \frac{2^{k-1} - 1 - k\epsilon}{1 - \epsilon}. \quad (78)$$

At large k , as can be seen for instance on (77), the asymptotic expansions of constraint densities are organized in different scaling behaviors, namely exponential, polynomial and logarithmic in k . Neglecting only exponentially small corrections we can simplify (78) into

$$\frac{\alpha}{2^{k-1}} = \frac{\ln 2}{k(1-\epsilon)\ln(1-\epsilon) + 1 + k\epsilon} = \frac{1}{k} \Gamma_r(k) \frac{1}{1-\epsilon}. \quad (79)$$

Without making additional approximations we see that $\epsilon_*(k)$ is solution of

$$\epsilon + (1-\epsilon)\ln(1-\epsilon) = \frac{\ln 2}{\Gamma_r(k)} - \frac{1}{k} - \epsilon \frac{\ln 2}{\Gamma_r(k)}. \quad (80)$$

Given the asymptotic behavior of $\Gamma_r(k)$ stated in (76) it is easy to see that $\epsilon_*(k)$ must vanish in the limit; one can thus expand the l.h.s. of (80) and obtain

$$\epsilon_*(k) = \sqrt{\frac{2 \ln 2}{\Gamma_r(k)}} + O\left(\frac{1}{\ln k}\right). \quad (81)$$

Reinserting in (79) we obtain

$$\alpha_*(k) = 2^{k-1} \frac{1}{k} \Gamma_r(k) \left(1 + \sqrt{\frac{2 \ln 2}{\Gamma_r(k)}} + O\left(\frac{1}{\ln k}\right)\right). \quad (82)$$

Using the expansion (76) of $\Gamma_r(k)$ we have finally

$$\alpha_*(k) = 2^{k-1} \frac{1}{k} \left(\ln k + \sqrt{2 \ln 2} \sqrt{\ln k} + O(\ln \ln k) \right). \quad (83)$$

The comparison with the expansion of $\alpha_{r,u}(k)$ given in (77) shows that the leading order is not modified, the term $\ln \ln k$ in the correction being replaced by a (larger) term of order $\sqrt{\ln k}$.

The leading order expansion of $\alpha_{d,u}(k)$ has been rigorously derived in [38] for a family of model encompassing the hypergraph bicoloring one, yielding in this case $2^{k-1} \ln k/k$. By analogy with other rigorous results [37, 39] (obtained for the q -coloring problem) we shall assume that $\alpha_{d,u}(k)$ has the same asymptotic expansion (77) as $\alpha_{r,u}$ (with a strictly smaller constant hidden in the $O(1)$ term). We thus conclude that

$$\alpha_{d,u}(k) = 2^{k-1} \frac{1}{k} [\ln k + \ln \ln k + O(1)] \leq \alpha_{\text{opt}}(k) \leq \alpha_*(k) = 2^{k-1} \frac{1}{k} \left(\ln k + \sqrt{2 \ln 2} \sqrt{\ln k} + O(\ln \ln k) \right), \quad (84)$$

hence that the best improvement of $\alpha_{\text{opt}}(k)$ with respect to $\alpha_{d,u}(k)$ that can be hoped for with the bias considered in this paper is a replacement of $\ln \ln k$ by $\sqrt{\ln k}$ in the second order term of their asymptotic expansions.

We shall finally come back briefly on the choice of parameters we made in (66), where we used a single parameter ϵ for the bias instead of trying to exploit all the $k/2$ free values of ω_p . Let us define $\alpha'_{\text{opt}}(k) = \sup \alpha_d(k, \{\omega_p\})$, where the maximization is now over all possible values of the ω_p , under the conditions $\omega_0 = \omega_k = 0$ and $\omega_p = \omega_{k-p}$. We have certainly $\alpha'_{\text{opt}}(k) \geq \alpha_{\text{opt}}(k)$, and the inequality is probably strict at least for small enough values of k ; however we shall now show that $\alpha'_{\text{opt}}(k) \leq \alpha_*(k)$, hence that the larger freedom of choice of generic parameters does not allow to beat the upperbound derived and discussed in the special case (66). To see this more easily let us exploit the invariance of the measure (1) under a multiplication of all ω_p by a common constant, and fix their normalizations in such a way that

$$\sum_{p=1}^{k-1} \binom{k}{p} \omega_p = 1. \quad (85)$$

With this choice the expressions of the RS entropy (19) and rigidity threshold (51) become

$$s^{RS}(\alpha, \{\omega_p\}) = \ln 2 + \alpha \left(\sum_{p=1}^{k-1} \binom{k}{p} \omega_p \ln \omega_p - k \ln 2 \right), \quad \alpha_r(k, \{\omega_p\}) = \frac{1}{k} \Gamma_r(k) \frac{1}{2\omega_1}. \quad (86)$$

Consider now a given choice of the parameter ω_1 , and a value of $\alpha \leq \alpha_r$; for these to allow a RS phase the corresponding entropy should be positive. Maximizing the entropy in (86) with respect to $\{\omega_2, \dots, \omega_{k-2}\}$, under the normalization condition (85) and for a fixed value of ω_1 is easily seen to yield $\omega_2 = \dots = \omega_{k-2}$, i.e. precisely the choice of parameters (66). In other words this bias is the one that allows to tune the fraction of frozen variables while keeping the measure as uniform as possible, in order to minimize the entropy cost it induces.

VII. DISCUSSION

We have studied the problem of bicoloring random k -regular hypergraphs, where every hyper-edge joins exactly k variables, also known as NAE- k -SAT. Having in mind the observation that algorithms usually reach solutions with no frozen variables and the conjecture that the computational complexity of the problem is directly related to the presence of frozen variables, we have modified the uniform measure over solutions to a biased measure where solutions with frozen variables are disfavoured (the works [46, 47] are based on a similar idea, with a bias favouring the regions of configuration space with a high local entropy of solutions). We have studied this biased measure both analytically via the cavity method, and numerically, running extensive Simulated Annealing processes in the search for solutions.

As a byproduct of our study we have presented two technical tools of interest by themselves: (i) how to properly determine the dynamic transition threshold, that corresponds to a first order transition in a functional space, by studying the stability parameter close to a random first order transition; (ii) how to estimate the algorithmic threshold for Simulated Annealing in random CSP by studying the dependence of the asymptotic energy on the cooling rate.

The results we got are somehow different depending on whether k is small or very large. In the range of small values of k we have found that the use of the biased measure is rather effective in helping SA to find solutions, thus moving its algorithmic threshold to larger values of α . This numerical finding is supported by the analytical computation that predicts a shift of the dynamic threshold in presence of a bias. This is a clear result showing that,

both in the thermodynamic limit and on problems of finite size, the use of the bias favouring unfrozen configurations is effective. As the non-uniformity of the measure necessarily implies a reduction of its entropy, the present result supports the idea that many solutions are not necessarily useful from the point of view of searching algorithms and actually removing these solutions (or reducing their weight) helps in the search for the remaining ones. The analytic result on the shift of the dynamic threshold implies that for $\alpha \in [\alpha_{d,u}, \alpha_{opt}]$ there are long range correlations for $\epsilon = 0$ that actually disappear for a range of non-zero biases. This result suggests that long range correlations are due to a subset of solutions where variables are very strongly correlated, while in the remaining subset of solutions variables are much more weakly correlated. The lack of reentrance of critical lines in temperature is a good news for thermal algorithms, that is algorithms satisfying detailed balance with respect to a Gibbs-Boltzmann probability distribution, because it implies that this class of algorithms should not suffer a slowing down due to critical points at finite temperature below the threshold given by the zero-temperature critical point, and the latter should be the one eventually determining the asymptotic behavior of these algorithms.

In the large k limit the bounds of (84) show that in the most optimistic scenario the gain obtained by the present version of the biasing strategy is rather modest, corresponding to an increase from $\ln \ln k$ to $\sqrt{\ln k}$ in the second order of the expansion of the threshold density of constraints. We believe that a very interesting open problem would be to determine which of the two bounds on the expansion of α_{opt} in (84) is tight (or whether the scaling is actually intermediate between the two). If the lowerbound is not tight even this tiny increase would be an improvement of the algorithmic gap, and would beat the rigidity threshold of the uniform measure (that was sometimes thought to be the algorithmic barrier). Adapting the techniques of [37–39] to prove a lowerbound on the reconstruction threshold of the biased measure should clarify this issue.

Another possible direction we are currently investigating is the study of more generic bias on the set of solutions, induced by interactions between variables at larger distances; this was shown in [40] to have an even more dramatic effect on the location of the rigidity transition. It remains to see whether the dynamic transition can also be efficiently manipulated in this way.

An interesting question to investigate in the future is the effect of the decimation on the biased measure, which can be analyzed along the lines of [25]. This would probably prove analytically that the algorithmic threshold for the process that searches a solution fixing variables according to marginals provided by belief propagation is improved by a non zero bias.

In the present work we have focused our study of biased measures on the hypergraph bicoloring problem because it is technically simpler than k -SAT or q -COL, its degrees of freedom being binary and its replica symmetric solution being trivial, while its unbiased measure is known to exhibit all the transitions undergone by the more complicated random CSPs. We are convinced that the biasing idea can also be applied to the latter and we expect a similar increase of the clustering transition threshold; we leave an explicit verification of this expectation as an open problem for future work.

Acknowledgments

We thank Lenka Zdeborova for useful discussions. GS is part of the PAIL grant of the French Agence Nationale de la Recherche, ANR-17-CE23-0023-01.

-
- [1] M. Garey and D. Johnson. *Computers and intractability: A guide to the theory of NP-completeness*. Freeman, San Francisco, 1979.
 - [2] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
 - [3] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. 2+p-SAT: Relation of typical-case complexity to the nature of the phase transition. *Random Structures and Algorithms*, **15**, 414 (1999).
 - [4] G. Biroli, R. Monasson, and M. Weigt. A variational description of the ground state structure in random satisfiability problems. *Eur. Phys. J. B*, **14**, 551 (2000).
 - [5] M. Mézard, G. Parisi, and R. Zecchina. Analytic and Algorithmic Solution of Random Satisfiability Problems. *Science*, **297**, 812–815 (2002).
 - [6] S. Mertens, M. Mézard, and R. Zecchina. Threshold values of random K-SAT from the cavity method. *Random Struct. Algorithms*, **28**(3), 340–373 (2006).
 - [7] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborova. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, **104**(25), 10318–10323 (2007).
 - [8] D. Achlioptas and F. Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In *Proc. of 38th STOC*, pages 130–139, New York, NY, USA, 2006. ACM.

- [9] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. In *Proceedings of FOCS 2008*, page 793, 2008.
- [10] M. Molloy. The freezing threshold for k -colourings of a random graph. In *Proceedings of the 44th symposium on Theory of Computing*, page 921. ACM, 2012.
- [11] J. Ding, A. Sly, and N. Sun. Proof of the Satisfiability Conjecture for Large K . In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, STOC '15, pages 59–68, 2015.
- [12] E. Mossel and Y. Peres. Information flow on trees. *Ann. Appl. Probab.*, **13**(3), 817–844 (2003).
- [13] A. Montanari and G. Semerjian. Rigorous Inequalities between Length and Time Scales in Glassy Systems. *J. Stat. Phys.*, **125**, 23 (2006).
- [14] B. Selman, H. A. Kautz, and B. Cohen. Noise strategies for improving local search. In *Proc. 12th AAAI*, pages 337–343, Menlo Park, CA, USA, 1994. AAAI Press.
- [15] G. Semerjian and R. Monasson. Relaxation and metastability in a local search procedure for the random satisfiability problem. *Phys. Rev. E*, **67**(6), 066103 (2003).
- [16] J. Ardelius and E. Aurell. Behavior of heuristics on large and hard satisfiability problems. *Phys. Rev. E*, **74**, 037702 (2006).
- [17] M. Alava, J. Ardelius, E. Aurell, P. Kaski, S. Krishnamurthy, P. Orponen, and S. Seitz. Circumspect descent prevails in solving random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, **105**(40), 15253–15257 (2008).
- [18] J. Franco and M. Paull. Probabilistic analysis of the Davis-Putnam procedure for solving satisfiability. *Discrete Applied Mathematics*, **5**, 77–87 (1983).
- [19] J. Franco. Results related to threshold phenomena research in satisfiability: lower bounds. *Theor. Comput. Sci.*, **265**, 147 (2001).
- [20] D. Achlioptas. Lower bounds for random 3-SAT via differential equations. *Theor. Comput. Sci.*, **265**(1-2), 159–185 (2001).
- [21] S. Cocco and R. Monasson. Trajectories in Phase Diagrams, Growth Processes, and Computational Complexity: How Search Algorithms Solve the 3-Satisfiability Problem. *Phys. Rev. Lett.*, **86**(8), 1654–1657 (2001).
- [22] A. Braunstein, M. Mézard, and R. Zecchina. Survey propagation: An algorithm for satisfiability. *Random Struct. Algorithms*, **27**(2), 201–226 (2005).
- [23] R. Marino, G. Parisi, and F. Ricci-Tersenghi. The Backtracking Survey Propagation Algorithm for Solving Random K -SAT Problems. *Nature Communications*, **7**, 12996 (2016).
- [24] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Solving Constraint Satisfaction Problems through Belief Propagation-guided decimation. (2007). [arXiv:0709.1667](https://arxiv.org/abs/0709.1667), Proceedings of the 45th Allerton Conference.
- [25] F. Ricci-Tersenghi and G. Semerjian. On the cavity method for decimated random constraint satisfaction problems and the analysis of belief propagation guided decimation algorithms. *Journal of Statistical Mechanics: Theory and Experiment*, page P09001 (2009).
- [26] A. Coja-Oghlan. A Better Algorithm for Random k -SAT. *SIAM Journal on Computing*, **39**(7), 2823–2864 (2010).
- [27] D. Gamarnik and M. Sudan. Performance of Sequential Local Algorithms for the Random NAE- K -SAT Problem. *SIAM Journal on Computing*, **46**(2), 590–619 (2017).
- [28] A. Coja-Oghlan and A. Y. Pachon-Pinzon. The decimation process in random k -SAT. *SIAM Journal on Discrete Mathematics*, **26**, 1471–1509 (2012).
- [29] S. Hetterich. Analysing Survey Propagation Guided Decimation on Random Formulas. *arXiv preprint arXiv:1602.08519*, (2016).
- [30] A. Coja-Oghlan and A. Frieze. Analyzing Walksat on Random Formulas. *SIAM Journal on Computing*, **43**(4), 1456–1485 (2014).
- [31] A. Coja-Oghlan, A. Haqshenas, and S. Hetterich. Walksat Stalls Well Below Satisfiability. *SIAM Journal on Discrete Mathematics*, **31**(2), 1160–1173 (2017).
- [32] L. Zdeborová and F. Krzakala. Phase Transitions in the Coloring of Random Graphs. *Phys. Rev. E*, **76**, 031131 (2007).
- [33] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Clusters of solutions and replica symmetry breaking in random k -satisfiability. *J. Stat. Mech.*, **P04004** (2008).
- [34] A. Montanari and G. Semerjian. On the dynamics of the glass transition on Bethe lattices. *J. Stat. Phys.*, **124**, 103–189 (2006).
- [35] G. Semerjian. On the freezing of variables in random constraint satisfaction problems. *J. Stat. Phys.*, **130**, 251 (2008).
- [36] M. Molloy and R. Restrepo. Frozen variables in random boolean constraint satisfaction problems. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1306–1318, 2013.
- [37] A. Sly. Reconstruction of Random Colourings. *Communications in Mathematical Physics*, **288**(3), 943–961 (2009).
- [38] A. Montanari, R. Restrepo, and P. Tetali. Reconstruction and Clustering in Random Constraint Satisfaction Problems. *SIAM Journal on Discrete Mathematics*, **25**(2), 771–808 (2011).
- [39] A. Sly and Y. Zhang. Reconstruction of colourings without freezing. *arXiv preprint arXiv:1610.02770*, (2016).
- [40] A. Braunstein, L. Dall’Asta, G. Semerjian, and L. Zdeborova. The large deviations of the whitening process in random constraint satisfaction problems. *Journal of Statistical Mechanics: Theory and Experiment*, **2016**(5), 053401 (2016).
- [41] D. Achlioptas and F. Ricci-Tersenghi. Random formulas have frozen variables. *SIAM Journal on Computing*, **39**(1), 260–280 (2009).
- [42] F. Krzakala and L. Zdeborová. Phase Transitions and Computational Difficulty in Random Constraint Satisfaction Problems. *J. Phys.: Conf. Ser.*, **95**, 012012 (2008).
- [43] E. N. Maneva, E. Mossel, and M. J. Wainwright. A new look at survey propagation and its generalizations. *J. ACM*, **54**(4) (2007).

- [44] A. Braunstein and R. Zecchina. Survey propagation as local equilibrium equations. *J. Stat. Mech.*, **P06007** (2004).
- [45] L. Dall’Asta, A. Ramezanzpour, and R. Zecchina. Entropy landscape and non-Gibbs solutions in constraint satisfaction problems. *Phys. Rev. E*, **77**, 031118 (2008).
- [46] C. Baldassi, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina. Local entropy as a measure for sampling solutions in constraint satisfaction problems. *Journal of Statistical Mechanics: Theory and Experiment*, **2016**(2), 023301 (2016).
- [47] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina. Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes. *Proceedings of the National Academy of Sciences*, **113**(48), E7655–E7662 (2016).
- [48] T. Maimbourg, M. Sellitto, G. Semerjian, and F. Zamponi. Generating dense packings of hard spheres by soft interaction design. *SciPost Phys.*, **4**, 39 (2018).
- [49] F. Krzakala, M. Mézard, and L. Zdeborová. Reweighted belief propagation and quiet planting for random K-SAT. *Journal on Satisfiability, Boolean Modeling and Computation*, **8**, 149–171 (2014).
- [50] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, **220**, 671–680 (1983).
- [51] T. Castellani, V. Napolano, F. Ricci-Tersenghi, and R. Zecchina. Bicoloring Random Hypergraphs. *J. Phys. A*, **36**, 11037 (2003).
- [52] A. Coja-Oghlan and L. Zdeborová. The condensation transition in random hypergraph 2-coloring. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 241–250, 2012.
- [53] V. Bapst, A. Coja-Oghlan, and F. Rasmann. A positive temperature phase transition in random hypergraph 2-coloring. *Ann. Appl. Probab.*, **26**(3), 1362–1406 (2016).
- [54] D. Achlioptas and C. Moore. Two moments suffice to cross a sharp threshold. *SIAM Journal on Computing*, **36**, 740–762 (2006).
- [55] A. Coja-Oghlan and K. Panagiotou. Catching the k-NAESAT Threshold. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC ’12, pages 899–908. ACM, 2012.
- [56] J. Ding, A. Sly, and N. Sun. Satisfiability Threshold for Random Regular nae-sat. *Communications in Mathematical Physics*, **341**(2), 435–489 (2016).
- [57] M. Mézard and G. Parisi. The Bethe lattice spin glass revisited. *Eur. Phys. J. B*, **20**, 217 (2001).
- [58] M. Mézard and G. Parisi. The cavity method at zero temperature. *J. Stat. Phys.*, **111**, 1–34 (2003).
- [59] M. Mézard and A. Montanari. *Physics, Information, Computation*. Oxford Press, Oxford, 2009.
- [60] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.
- [61] F. R. Kschischang, B. Frey, and H.-A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory*, **47**(2), 498–519 (2001).
- [62] J. Yedidia, W. Freeman, and Y. Weiss. Understanding Belief Propagation and Its Generalizations. In *Exploring Artificial Intelligence in the New Millennium*, pages 239–236. Science & Technology Books, 2003.
- [63] R. Monasson. Structural glass transition and the Entropy of the Metastable States. *Phys. Rev. Lett.*, **75**, 2847 (1995).
- [64] M. Mézard and A. Montanari. Reconstruction on Trees and Spin Glass Transition. *J. Stat. Phys.*, **124**, 1317–1350 (2006).
- [65] E. Mossel. Reconstruction on Trees: Beating the Second Eigenvalue. *Ann. Appl. Probab.*, **11**(1), 285–300 (2001).
- [66] S. Janson and E. Mossel. Robust reconstruction on trees is determined by the second eigenvalue. *Ann. Probab.*, **32**, 2630–2649 (2004).
- [67] H. Kesten and B. P. Stigum. Additional Limit Theorems for Indecomposable Multidimensional Galton-Watson Processes. *The Annals of Mathematical Statistics*, **37**, 1463 (1966).
- [68] J. R. L. de Almeida and D. J. Thouless. Stability of the Sherrington-Kirkpatrick Solution of a Spin-Glass Model. *J. Phys. A*, **11**, 983–990 (1978).
- [69] F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová. Typology of phase transitions in Bayesian inference problems. *arXiv preprint arXiv:1806.11013*, (2018).
- [70] M. Gabrié, V. Dani, G. Semerjian, and L. Zdeborová. Phase transitions in the q-coloring of random hypergraphs. *Journal of Physics A: Mathematical and Theoretical*, **50**(50), 505002 (2017).
- [71] G. Parisi, F. Ricci-Tersenghi, and T. Rizzo. Diluted mean-field spin-glass models at criticality. *Journal of Statistical Mechanics: Theory and Experiment*, **2014**(4), P04013 (2014).
- [72] G. Parisi, F. Ricci-Tersenghi, and T. Rizzo. Erratum: Diluted mean-field spin-glass models at criticality. *Journal of Statistical Mechanics: Theory and Experiment*, **2015**(3), E03001 (2015).
- [73] L. Zdeborová and F. Krzakala. Generalization of the cavity method for adiabatic evolution of Gibbs states. *Phys. Rev. B*, **81**, 224205 (2010).