Going after the k-SAT Threshold

[Extended Abstract]

Amin Coja-Oghlan*
Goethe University
Mathematics Institute
Frankfurt 60054, Germany
acoghlan@math.uni-frankfurt.de

Konstantinos Panagiotou†
University of Munich
Mathematics Institute
Theresienstr. 39, 80333 München, Germany
kpanagio@math.lmu.de

ABSTRACT

Random k-SAT is the single most intensely studied example of a random constraint satisfaction problem. But despite substantial progress over the past decade, the threshold for the existence of satisfying assignments is not known precisely for any $k \geq 3$. The best current results, based on the second moment method, yield upper and lower bounds that differ by an additive $k \cdot \frac{\ln 2}{2}$, a term that is unbounded in k (Achlioptas, Peres: STOC 2003). The basic reason for this gap is the inherent asymmetry of the Boolean value 'true' and 'false' in contrast to the perfect symmetry, e.g., among the various colors in a graph coloring problem. Here we develop a new *asymmetric second moment method* that allows us to tackle this issue head on for the first time in the theory of random CSPs. This technique enables us to compute the k-SAT threshold up to an additive $\ln 2 - \frac{1}{2} + O(1/k) \approx 0.16$. Independently of the rigorous work, physicists have developed a sophisticated but non-rigorous technique called the "cavity method" for the study of random CSPs (Mézard, Parisi, Zecchina: Science 2002). Our result matches the best bound that can be obtained from the so-called "replica symmetric" version of the cavity method, and indeed our proof directly harnesses parts of the physics calculations.

Categories and Subject Descriptors

G.2 [**Discrete Mathematics**]: General; F.2 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms

Theory

Keywords

Random Structures, Phase Transitions, k-SAT, Second Moment Method, Belief Propagation

*Supported by ERC Starting Grant 278857–PTCC (FP7).

†Supported by DFG grant PA 2080/2-1.

Submitted to STOC 2013.

1. INTRODUCTION

Since the early 2000s physicists have developed a sophisticated but highly non-rigorous technique called the "cavity method" for the study of random constraint satisfaction problems. This method allowed them to put forward a very detailed *conjectured* picture according to which various phase transitions affect both computational and structural properties of random CSPs. In addition, the cavity method has inspired new message passing algorithms called *Belief/Survey Propagation guided decimation*. Over the past few years there has been significant progress in turning bits and pieces of the physics picture into rigorous theorems. Examples include results on the interpolation method [2, 7] or the geometry of the solution space [1, 26, 27] and their algorithmic implications [3, 9].

In spite of this progress, substantial gaps remain. Perhaps most importantly, in most random CSPs the threshold for the existence of solutions is not known precisely. In the relatively simple case of the random k-NAESAT ("Not-All-Equal-Satisfiability") problem the difference between the best current lower and upper bounds is as tiny as $2^{-\Omega(k)}$ [11]. By contrast, in random graph k-coloring, a problem already studied by Erdős and Rényi in the 1960s, the best current bounds differ by $\Theta(\ln k)$ [5]. Hence, the difference is *unbounded* in terms of the number of colors. Even worse, in random k-SAT the gap is as big as $\Theta(k)$ [6]. Yet random k-SAT is probably the single most important example of a random CSP, not least due to the great amount of experimental and algorithmic work conducted on it (e.g., [21, 23]).

The reason for the large gap in random k-SAT is that the satisfiability problem lacks a certain *symmetry property*. This property is vital to the current rigorous proof methods, particularly the *second moment method*, on which most of the previous work is based (e.g., [4, 5, 6]). More precisely, in random graph coloring the different colors all play the exact same role: for any proper coloring of a graph, another proper coloring can be obtained by simply permuting the color classes (e.g., color all red vertices blue and vice versa). Similarly, in k-NAESAT, where the requirement is that in each clause at least one literal must be true and at least one false, the binary inverse of any NAE-solution is a NAE-solution as well. By contrast, in k-SAT there is an inherent asymmetry between the Boolean values 'true' and 'false'.

As has been noticed in prior work [4, 6], the second moment

method is fundamentally ill-posed to deal with such asymmetries. Roughly speaking, the second moment method is based on the assumption that in a random CSP instance, two randomly chosen solutions are perfectly uncorrelated. But in random k-SAT, this is simply not the case. Indeed, suppose that a variable x appears much more often positively than negatively throughout the formula. Then it seems reasonable to expect that most satisfying assignments set x to 'true', thereby satisfying all clauses where x appears positively. More generally, define the *majority vote* σ_{maj} to be the assignment that sets variable x to true if it appears more often positively than negatively, and to false otherwise. Then we expect that the satisfying assignments of a random formula "gravitate toward" σ_{maj}. Unfortunately, the correlations among satisfying assignments induced by this drift toward σ_{maj} doom the second moment method. Previously this issue was sidestepped by symmetrizing the problem artificially [4, 6]. But this inevitably leaves a $\Theta(k)$ gap.

The main contribution of the present work is a new *asymmetric second moment method* that enables us to tackle this problem head on. A key feature of this method is that we harness the Belief Propagation calculation from physics, called the "replica symmetric case" of the cavity method in physics jargon. We are going to employ Belief Propagation directly as an "educated guess" in the design the random variable upon which our proof is based in order to quantify how much a typical satisfying assignment leans toward σ_{maj}.

This is in contrast to most prior work on the subject, where individual statements hypothesized on the basis of physics arguments were proved via completely different methods (with the notable exception of the interpolation technique [2, 7, 16]). Hence, we view the present work as a pivotal step in the long-term effort of providing a rigorous foundation for the physicists' cavity method. In fact, the general approach developed here does not hinge on particular properties of the k-SAT problem, and thus we expect the technique will extend to other asymmetric problems as well. Examples include not only other random CSPs that are asymmetric per se, but also instances of random problems that arise at intermediate steps of message passing algorithms such as *Belief/Survey Propagation guided decimation*, even if the initial problem is symmetric. In particular, we believe that getting a handle on asymmetric problems is a necessary step to analyze such message passing algorithms accurately.

To state our results precisely, we let $k \geq 3, n > 0$ be integers and we let $V = \{x_1, \ldots, x_n\}$ be a set of n Boolean variables. Further, let $\mathbf{\Phi} = \mathbf{\Phi}_k(n, m)$ denote a Boolean formula with m clauses of length k over the variables V chosen uniformly at random among all $(2n)^{km}$ such formulas. Let $r = m/n$ denote the *density*. We say that an event occurs *with high probability* ('w.h.p.') if its probability tends to 1 as $n \to \infty$.

Friedgut [17] showed that for any $k \geq 3$ there exists a *threshold sequence*[1] $r_{k-\mathrm{SAT}}(n)$ such that for any (fixed) $\varepsilon > 0$ w.h.p. $\mathbf{\Phi}$ is satisfiable if $m/n < (1-\varepsilon) r_{k-\mathrm{SAT}}(n)$, while for $m/n > (1+\varepsilon) r_{k-\mathrm{SAT}}(n)$ $\mathbf{\Phi}$ is unsatisfiable w.h.p.

Upper bounds on $r_{k-\mathrm{SAT}}$ can be obtained via the *first moment method*. The best current ones [16, 22] are

$$r_{k-\mathrm{SAT}} \leq r_{\mathrm{upper}} = 2^k \ln 2 - (1+\ln 2)/2 + o_k(1), \qquad (1)$$

where $o_k(1)$ hides a term that tends to 0 for large k. The best prior lower bound is due to Achlioptas and Peres [6], who used a "symmetric" second moment argument to show

$$r_{k-\mathrm{SAT}} \geq r_{\mathrm{bal}} = 2^k \ln 2 - k \cdot \frac{\ln 2}{2} - \left(1 + \frac{\ln 2}{2}\right) + o_k(1). \quad (2)$$

The bounds (1) and (2) leave an additive gap of $k \cdot \frac{\ln 2}{2} + \frac{1}{2} + o_k(1)$, i.e., the gap is unbounded in terms of k.

THEOREM 1. *There is $\varepsilon_k = o_k(1)$ such that*

$$r_{k-\mathrm{SAT}} \geq r_{\mathrm{BP}} = 2^k \ln 2 - \frac{3 \ln 2}{2} - \varepsilon_k. \qquad (3)$$

Achlioptas and Peres asked whether the gap $r_{\mathrm{upper}} - r_{k-\mathrm{SAT}}$ is bounded by an absolute constant (independent of k). Theorem 1 answers this question, reducing the gap to $\ln 2 - \frac{1}{2} \approx 0.19$. No attempt at optimizing the error term ε_k has been made, but our proofs yield rather directly that $\varepsilon_k = O(1/k)$.

Apart from the quantitative improvement, the main point of this paper is that we manage to solve the problem of asymmetry in random CSPs for the first time. To explain this point, we start by discussing what we mean by asymmetry and how it derails the second moment method. That this is so was already intuited in [4, 6]. In the next section, we are going to verify and elaborate on those discussions.

2. ASYMMETRY AND THE SECOND MOMENT METHOD

The second moment method. In general, the second moment method works as follows. Suppose that $Z = Z(\mathbf{\Phi})$ is a non-negative random variable such that $Z > 0$ only if $\mathbf{\Phi}$ is satisfiable. Moreover, suppose that for some density $r > 0$ there is a number $C = C(k) > 0$ that may depend on k but not on n such that

$$0 < \mathrm{E}\left[Z^2\right] \leq C \cdot \mathrm{E}\left[Z\right]^2. \qquad (4)$$

We claim that then $r_{k-\mathrm{SAT}} \geq r$. Indeed, the *Paley-Zygmund inequality*

$$\mathrm{P}\left[Z > 0\right] \geq \frac{\mathrm{E}\left[Z\right]^2}{\mathrm{E}\left[Z^2\right]} \qquad (5)$$

implies that $\mathrm{P}\left[\mathbf{\Phi} \text{ is satisfiable}\right] \geq \mathrm{P}\left[Z > 0\right] \geq 1/C$. Because the right hand side remains bounded away from 0 as $n \to \infty$, the following simple consequence of Friedgut's sharp threshold result implies $r_{k-\mathrm{SAT}} \geq r$.

LEMMA 1 ([17]). *Let $k \geq 3$. If for some r we have*

$$\liminf_{n\to\infty} \mathrm{P}\left[\mathbf{\Phi} \textit{ is satisfiable}\right] > 0,$$

then $r_{k-\mathrm{SAT}} \geq r - o(1)$.

Hence, we "just" need to find a random variable that satisfies (5). Let $\mathcal{S}(\mathbf{\Phi})$ denote the set of satisfying assignments;

[1] It is widely conjecture but as yet unproved that $r_{k-\mathrm{SAT}}(n)$ converges for any $k \geq 3$.

then certainly $Z = |\mathcal{S}(\mathbf{\Phi})|$ is the most obvious choice. However, this "vanilla" second moment argument turns out to fail spectacularly. We need to understand why.

Asymmetry and the majority vote. The origin of the problem is that k-SAT is asymmetric in the following sense. Suppose that all we know about the random formula $\mathbf{\Phi}$ is for each variable x the number d_x of times that x appears as a positive literal in the formula, and the number $d_{\neg x}$ of negative occurrences. Then our best stab at constructing a satisfying assignment seems to be the "majority vote" assignment σ_{maj} where we set x to true if $d_x > d_{\neg x}$ and to false otherwise. Indeed, by maximizing the total number of true literal occurrences, of which a satisfying assignment must put one in every clause, σ_{maj} also maximizes the probability of being satisfiable.

Our proof of Theorem 1 allows us to formalize this observation, thereby verifying a conjecture from [6]. Let $\mathrm{dist}(\cdot,\cdot)$ denote the Hamming distance.

COROLLARY 1. *There is a number $\delta = \delta(k) > 0$ such that for $2^k/k < r < r_{\mathrm{BP}}$ w.h.p. we have*

$$\sum_{\sigma \in \mathcal{S}(\mathbf{\Phi})} \frac{dist(\sigma, \sigma_{\mathrm{maj}})}{|\mathcal{S}(\mathbf{\Phi})|} \leq \left(\frac{1}{2} - \delta\right) \cdot n. \tag{6}$$

Hence, the average Hamming distance of $\sigma \in \mathcal{S}(\mathbf{\Phi})$ from σ_{maj} is strictly smaller than $n/2$, i.e., the set $\mathcal{S}(\mathbf{\Phi})$ is "skewed toward" σ_{maj} w.h.p.

This asymmetry dooms the second moment method. To see why, let

$$w_{\mathrm{maj}} = w_{\mathrm{maj}}(\mathbf{\Phi}) = \sum_{x \in V} \frac{\max\{d_x, d_{\neg x}\}}{km}$$

denote the ***majority weight*** of $\mathbf{\Phi}$. Then the larger w_{maj}, the more likely σ_{maj} and assignments close to it are to be satisfying. In effect, the bigger w_{maj}, the more satisfying assignments we expect to have. The consequence of this is that the number $|\mathcal{S}(\mathbf{\Phi})|$ of satisfying assignments behaves like a "lottery": its expectation is driven up by a tiny fraction of "lucky" formulas with w_{maj} much bigger than expected.

Let us highlight this tradeoff, as it is characteristic of the kind of trouble that asymmetry causes. For $\xi > 0$ independent of n but sufficiently small it turns out that for a certain constant $c > 0$,

$$\mathrm{P}\left[w_{\mathrm{maj}} \sim \mathrm{E}\left[w_{\mathrm{maj}}\right] + \xi\right] = \exp\left[-(c\xi^2 + O(\xi^3))n\right]. \tag{7}$$

That is, the probability is exponentially small but, like in the Chernoff bound, the exponent is a *quadratic* function of ξ. By comparison, increasing the majority weight by ξ boosts the expected number of satisfying assignments by a *linear* exponential factor: there is $c' > 0$ such that

$$\begin{aligned}\mathrm{E}\left[\,|\mathcal{S}(\mathbf{\Phi})| \;\middle|\; w_{\mathrm{maj}} \sim \mathrm{E}\left[w_{\mathrm{maj}}\right] + \xi\right] \;&= \\ \exp\left[(c'\xi + O(\xi^2))n\right] \cdot \mathrm{E}\left[|\mathcal{S}(\mathbf{\Phi})| \,\middle|\, w_{\mathrm{maj}} \sim \mathrm{E}\left[w_{\mathrm{maj}}\right]\right].\end{aligned} \tag{8}$$

The exponent in (8) is linear because for a typical assignment τ at distance $(\frac{1}{2} - \delta)n$ from σ_{maj} increasing w_{maj} by ξ boosts the number of literals that are true under τ by $2\delta\xi \cdot km$, a term that is linear in ξ.

Since the exponent is linear in (8) but quadratic in (7), there is a (small but) strictly positive $\xi > 0$ such that the "gain" $\exp\left[(c'\xi + O(\xi^2))n\right]$ in the expected number of satisfying assignments exceeds the "penalty" $\exp\left[-(c\xi^2 + O(\xi^3))n\right]$ for deviating from $\mathrm{E}\left[w_{\mathrm{maj}}\right]$. With little extra work, this observation leads to

LEMMA 2. *For any $k \geq 3$ and $r > 2^k/k$ we have*

$$|\mathcal{S}(\mathbf{\Phi})| \leq \exp\left(-\Omega(4^{-k}) \cdot n\right) \cdot \mathrm{E}\left[|\mathcal{S}(\mathbf{\Phi})|\right] \quad \textit{w.h.p.}$$

Lemma 2 entails rather easily that the "vanilla" second moment argument fails dramatically. Indeed, as already noticed in [4, 6], we have $\mathrm{E}\left[|\mathcal{S}(\mathbf{\Phi})|^2\right] \geq \exp(\Omega(n)) \cdot \mathrm{E}\left[|\mathcal{S}(\mathbf{\Phi})|\right]^2$. Hence, we miss our mark (4) by an exponential factor. But Lemma 2 is witness to an even worse failure: not only does (4) fail to hold, but even the normally much more dependable *first* moment overshoots the "actual" number of satisfying assignments by an exponential factor! (Lemma 2 is an improvement of an observation from [1], showing that $|\mathcal{S}(\mathbf{\Phi})| \leq \exp(-\xi n) \cdot \mathrm{E}\left[|\mathcal{S}(\mathbf{\Phi})|\right]$ w.h.p. for some tiny $\xi = \xi(k) > 0$; we conjecture that the 4^{-k} term in Lemma 2 is tight.)

In summary, the drift toward σ_{maj} and the resulting fluctuations of the majority weight induce a tremendous source of variance, derailing the "vanilla" second moment argument.

Balanced assignments. A natural way to sidestep this issue is to work with a 'symmetric' subset of $\mathcal{S}(\mathbf{\Phi})$. Perhaps the most obvious choice is the set $\mathcal{S}_{\mathrm{NAE}}(\mathbf{\Phi})$ of NAE-solutions. In a landmark paper, Achlioptas and Moore [4] proved that indeed there is $C = C(k) > 0$ such that for $Z_{\mathrm{NAE}} = |\mathcal{S}_{\mathrm{NAE}}(\mathbf{\Phi})|$ we have

$$\mathrm{E}\left[Z_{\mathrm{NAE}}^2\right] \leq C \cdot \mathrm{E}\left[Z_{\mathrm{NAE}}\right]^2 \quad \text{for } r \leq 2^{k-1}\ln 2 - O_k(1). \tag{9}$$

As we saw above (cf. Lemma 1), this implies that $r_{k-\mathrm{SAT}} \geq 2^{k-1}\ln 2 - O(1)$. However, a simple (first moment) calculation shows that for $r > 2^{k-1}\ln 2$, the set $\mathcal{S}_{\mathrm{NAE}}(\mathbf{\Phi})$ is empty w.h.p. Thus, the idea of working with NAE-solutions stops working at $r \sim 2^{k-1}\ln 2$, about a factor of two below the satisfiability threshold.

Achlioptas and Peres [6] obtained a better bound by precipitating symmetry in a more subtle manner. Let us call $\sigma \in \{0,1\}^n$ ***balanced*** if under σ out of the km literal occurrences in $\mathbf{\Phi}$ *exactly half* are true (i.e., $\frac{km}{2} \pm 1$). Thus, balanced assignments are expressly forbidden from pandering to σ_{maj}. Now, let $\mathcal{S}_{\mathrm{bal}}(\mathbf{\Phi})$ be the set of all balanced satisfying assignments, and set $Z_{\mathrm{bal}} = |\mathcal{S}_{\mathrm{bal}}(\mathbf{\Phi})|$. Achlioptas and Peres used a clever weighting scheme to prove that

$$\mathrm{E}\left[Z_{\mathrm{bal}}^2\right] \leq C \cdot \mathrm{E}\left[Z_{\mathrm{bal}}\right]^2 \qquad \text{for } r \leq r_{\mathrm{bal}} \quad \text{(cf. (2))}. \tag{10}$$

As before, this implies that $r_{k-\mathrm{SAT}} \geq r_{\mathrm{bal}}$ (Lemma 1).

Yet as in the case of NAE-solutions, balanced satisfying assignments cease to exist way before the satisfiability threshold. Indeed, Achlioptas and Peres observed that $\mathcal{S}_{\mathrm{bal}}(\mathbf{\Phi}) = \emptyset$

for $r > 2^k \ln 2 - k\frac{\ln 2}{2}$ w.h.p. In effect, to close in further on $r_{k-\mathrm{SAT}}$ we will have to accommodate assignments that lean toward σ_{maj}. How can this be accomplished?

A quick fix? We saw that to make an asymmetric second moment argument work, we need to rule out fluctuations of the majority weight. A sensible way of implementing this is by actually fixing the entire vector $\boldsymbol{d} = (d_x, d_{\neg x})_{x \in V}$ that counts the positively/negatively occurrences of each variable. More precisely, given a non-negative integer vector $\boldsymbol{d} = (d_x, d_{\neg x})_{x \in V}$ with $\sum_{x \in V} d_x + d_{\neg x} = km$ let $\boldsymbol{\Phi_d}$ denote a uniformly random k-CNF in which each variable x appears d_x times positively and $d_{\neg x}$ times negatively. Then we can split the generation of a random formula $\boldsymbol{\Phi}$ into two steps:

> First, choose the occurrence vector $\boldsymbol{d}$ randomly from the "correct" distribution $\boldsymbol{D}$.
>
> Then, choose a random formula $\boldsymbol{\Phi_d}$.

The "correct" $\boldsymbol{D}$ is as follows. Let $\boldsymbol{e} = (e_x, e_{\neg x})_{x \in V}$ be a family of independent Poisson variables with mean $kr/2$ each. Moreover, let $\mathcal{E}$ be the event that $\sum_{x \in V} e_x + e_{\neg x} = km$. Let $\boldsymbol{D}$ be the conditional distribution of $\boldsymbol{e}$ given $\mathcal{E}$. Then standard arguments show that the outcome of first choosing $\boldsymbol{d}$ and then $\boldsymbol{\Phi_d}$ is exactly the uniformly random $\boldsymbol{\Phi}$.

The point of generating $\boldsymbol{\Phi}$ in two steps as above is that given the outcome $\boldsymbol{d}$ of the first step, the majority weight is *fixed*. Hence, if we could show that *given* a "typical" $\boldsymbol{d}$, the second moment succeeds for $|\mathcal{S}(\boldsymbol{\Phi_d})|$ we would obtain a lower bound on r_{k-SAT}. Unfortunately, matters are not so simple.

LEMMA 3. *W.h.p. for a vector $\boldsymbol{d}$ chosen from $\boldsymbol{D}$ we have* $\mathrm{E}[|\mathcal{S}(\boldsymbol{\Phi_d})|^2] \geq \exp(\Omega(n)) \cdot \mathrm{E}[|\mathcal{S}(\boldsymbol{\Phi_d})|]^2$.

Let us stress the two levels of randomness in Lemma 3. First, there is the choice of $\boldsymbol{d}$. Then, for a *given* $\boldsymbol{d}$, we compare $\mathrm{E}[|\mathcal{S}(\boldsymbol{\Phi_d})|^2]$ and $\mathrm{E}[|\mathcal{S}(\boldsymbol{\Phi_d})|]^2$. Of course, both of these quantities depend on $\boldsymbol{d}$, and we find that w.h.p. $\boldsymbol{d}$ is such that the first exceeds the second by an exponential factor.

The explanation for this is that even if we fix $\boldsymbol{d}$, various other types of fluctuations remain, turning $|\mathcal{S}(\boldsymbol{\Phi_d})|$ into a "lottery". For instance, even given $\boldsymbol{d}$ the number of clauses that are unsatisfied under σ_{maj} fluctuates. Hence, the inherent asymmetry of k-SAT puts not only the majority weight but also various other parameters on a slippery slope. What we need is a way of controlling all these fluctuations simultaneously. We will present our solution in Section 4.

Catching the k-SAT threshold? Before we come to that, let us discuss what it would take to eliminate the (small but non-zero) gap left by Theorem 1, i.e., how far we are from "catching" the k-SAT threshold. The physicists' cavity method comes in two installments. The (relatively speaking) simpler "replica symmetric" version is based on Belief Propagation. Theorem 1 provides a rigorous proof of the best possible bound on the k-SAT threshold that can be obtained from this version of the cavity method (up to possibly the precise error term ε_k) [24].

Unfortunately, for $r > r_{\mathrm{BP}}$ the replica symmetric version (and in particular the Belief Propagation predictions that we depend upon) are conjectured to break down. According to the more sophisticated "1-step replica symmetry breaking" (1RSB) version of the cavity method, the reason for this is that at $r \sim r_{\mathrm{BP}}$ a new type of correlation amongst satisfying assignments arises. To deal with these correlations, the physics methods replace Belief Propagation by the *much* more intricate Survey Propagation technique.

In [11] we managed to prove rigorously that the 1RSB prediction for the random k-NAESAT threshold is correct (up to an additive $2^{-\Omega(k)}$). However, [11] depends *heavily* on the fact that k-NAESAT is symmetric. While it would be very interesting to combine the merits of the present paper with those of [11], this appears to be quite challenging. Thus, putting the 1RSB calculation for random k-SAT on a rigorous foundation remains an important open problem. That said, we believe that any such attempt would need to build upon the techniques developed in this paper.

3. RELATED WORK

The interest in random k-SAT originated largely from the *experimental* observation that there seems to be a sharp threshold for satisfiability and, moreover, that for certain densities $r < r_{k-SAT}$ no polynomial time algorithm is known to find a satisfying assignment w.h.p. [21, 23]. Currently, the precise k-SAT threshold is known (rigorously) only in two cases. Chvatal and Reed [8] and Goerdt [20] proved independently that $r_{2-\mathrm{SAT}} = 1$. Of course, 2-SAT is special because there is a simple criterion for (un)satisfiability, which enables the proofs of [8, 20]. Unsurprisingly, these methods do not extend to $k > 2$. Additionally, the threshold is known precisely when $k > \log_2 n$, i.e., the clause length *diverges* as a function of n [19]. In this case, the problem of asymmetry evaporates because the majority weight is sufficiently concentrated for the "vanilla" second moment method to succeed. (Note that Proposition 2 holds for any *fixed* k, but not for $k = k(n) \to \infty$.) The issue of asymmetry also disappears in the case of *strongly regular* formulas [29] where for some fixed d we have $d_x = d_{\neg x} = d$ for all $x \in V$.

Also in random k-XORSAT (random linear equations mod 2) the threshold for the existence of solutions is known precisely [14]. The proof relies on computing the second moment of the number of solutions (after the instance has been stripped down to a suitable core). In contrast to random k-SAT, the random k-XORSAT problem is symmetric (cf. Remark 3 below), albeit in a more subtle way than k-NAESAT.

Other problems where the second moment method succeeds are symmetric as well. Pioneering the use of the second moment method in random CSPs, Achlioptas and Moore [4] computed the random k-NAESAT threshold within an additive $1/2$. By enhancing this argument with insights from physics this gap can be narrowed to a mere $2^{-\Omega(k)}$ [11, 12]. Moreover, the best current bounds on the random (hyper)graph k-colorability thresholds are based on "vanilla" second moment arguments as well [5, 15]. In summary, in all the previous second moment arguments, the issue of asymmetry either did not appear at all by the nature of the problem [4, 5, 11, 12, 14, 15, 19], or it was sidestepped [6].

The best current algorithms for random k-SAT find satisfying assignments w.h.p. for densities up to $1.817 \cdot 2^k/k$ (better for small k) resp. $2^k \ln(k)/k$ (better for large k) [9, 18], a factor of $\Theta(k/\ln k)$ below the satisfiability threshold. By comparison, the Lovász Local Lemma and its algorithmic version succeed up to $r = \Theta(2^k/k^2)$ [28].

Apart from experimental work [23], very little is known about the physics-inspired message passing algorithms ("Belief/Survey Propagation guided decimation") [25]. The most basic variant of Belief Propagation guided decimation is known to fail w.h.p. on random formulas if $r > c \cdot 2^k/k$ for some constant $c > 0$ [10]. However, it is conceivable that Survey Propagation and/or other variants of Belief Propagation perform better.

4. THE RANDOM VARIABLE

Our goal is to make the second moment method work for a random variable that counts "asymmetric" satisfying assignments. In this section, we develop this random variable. The starting point, and the key ingredient, is simply a map $p : \mathbf{Z} \to [0, 1]$. For the sake of clarity, we start by setting up the framework for generic maps p; below we will use the Belief Propagation formalism to pick the "optimal" p.

The idea is that p prescribes how strongly the assignments that we work with lean toward the majority vote. Informally speaking, we are going to work with assignments such that a variable x that occurs d_x times positively and $d_{\neg x}$ times negatively has a $p(d_x - d_{\neg x})$ chance of being set to 'true'. Before we give a formal definition, we need to fix the number of times that each variable appears positively or negatively.

Fixing the majority weight. As we saw in Section 2, in order to make the second moment argument work, we need to rule out fluctuations of the majority weight. To achieve this, we follow the strategy outlined in Section 2. That is, we are going to work with formulas $\boldsymbol{\Phi_d}$ with a given vector $\boldsymbol{d} = (d_x, d_{\neg x})_{x \in V}$ of occurrence counts, where each variable x appears precisely d_x times positively and $d_{\neg x}$ times negatively. As in Section 2, we let $\boldsymbol{D}$ denote the (conditional Poisson) distribution over sequences $\boldsymbol{d}$ such that first choosing $\boldsymbol{d}$ from $\boldsymbol{D}$ and then generating $\boldsymbol{\Phi_d}$ is equivalent to choosing a k-CNF $\boldsymbol{\Phi}$ uniformly at random.

Fixing the marginals. Now, fix one such vector $\boldsymbol{d}$. Then the map $p : \mathbf{Z} \to [0, 1]$ induces a map $p_{\boldsymbol{d}}$ from the set $L = \{x, \neg x : x \in V\}$ of literals to $[0, 1]$ in the natural way: we let

$$p_{\boldsymbol{d}}(x) = p(d_x - d_{\neg x}) \text{ and } p_{\boldsymbol{d}}(\neg x) = 1 - p(x). \tag{11}$$

The idea is that, given $\boldsymbol{d}$, we should set variable x to 'true' with probability $p_{\boldsymbol{d}}(x)$.

To formalize this, we call $p_{\boldsymbol{d}}(l)$ the $p_{\boldsymbol{d}}$***-type*** of the literal l. Let $\mathcal{T}_{\boldsymbol{d}} = \{p_{\boldsymbol{d}}(l) : l \in L\}$ be the set of all possible $p_{\boldsymbol{d}}$-types. We say that $\sigma : V \to \{0, 1\}$ has $p_{\boldsymbol{d}}$***-marginals*** if for any type $t \in \mathcal{T}_{\boldsymbol{d}}$ we have[2]

$$\sum_{l \in L : p_{\boldsymbol{d}}(l) = t} \sigma(l) \cdot d_l = t \cdot \sum_{l \in L : p_{\boldsymbol{d}}(l) = t} d_l.$$

[2]In this extended abstract we disregard rounding issues, i.e., we systematically omit all floor and ceiling signs.

i.e., among all occurrences of literals of type t, a t fraction is true under σ. This definition captures the above idea that variable x has a $p_{\boldsymbol{d}}(x)$ chance of being 'true'.

Fixing the clause types. We define the $p_{\boldsymbol{d}}$***-type*** of a clause $l_1 \vee \cdots \vee l_k$ as the k-tuple $(p_{\boldsymbol{d}}(l_1), \ldots, p_{\boldsymbol{d}}(l_k)) \in [0, 1]^k$ comprising of the individual literal types. Let $\mathcal{L}_{\boldsymbol{d}} = \mathcal{T}_{\boldsymbol{d}}^k$ be the set of all possible clause types. For each $\ell \in \mathcal{L}_{\boldsymbol{d}}$ let $M_{\boldsymbol{\Phi_d}}(\ell)$ be the set of indices $i \in [m]$ such that the ith clause of $\boldsymbol{\Phi_d}$ has type ℓ, and let $m_{\boldsymbol{\Phi_d}}(\ell) = |M_{\boldsymbol{\Phi_d}}(\ell)|$.

In addition to fluctuations of the majority weight, we also need to suppress fluctuations of the numbers $m_{\boldsymbol{\Phi_d}}(\ell)$. We are going to use the same trick as in the case of the majority weight. Namely, we split the generation of a random formula $\boldsymbol{\Phi_d}$ into two steps:

> First, choose a vector $\boldsymbol{m} = (m(\ell))_{\ell \in \mathcal{L}}$ from the "correct" distribution $\boldsymbol{M_d}$.
>
> Then, generate a formula $\boldsymbol{\Phi_{d,m}}$ uniformly at random in which each variable x appears exactly d_x times positively and exactly $d_{\neg x}$ times negatively and that has exactly $m(\ell)$ clauses of type ℓ for all $\ell \in \mathcal{L}$.

Formally, the "correct" $\boldsymbol{M_d}$ is just the distribution of the random vector $\boldsymbol{m_{\Phi_d}} = (m_{\boldsymbol{\Phi_d}}(\ell))_{\ell \in \mathcal{L}}$ that counts the clauses by types in the "unrestricted" formula $\boldsymbol{\Phi_d}$. It is easily verified that the overall outcome of the above experiment is identical to $\boldsymbol{\Phi_d}$. From now on, we fix both $\boldsymbol{d}$ and $\boldsymbol{m}$.

Given $\boldsymbol{d}, \boldsymbol{m}$ there is a simple way of generating the random formula $\boldsymbol{\Phi_{d,m}}$. Namely, create d_l clones of each literal l, and put all the clones of a given $p_{\boldsymbol{d}}$-type on a pile. Then the formula $\boldsymbol{\Phi_{d,m}}$ is simply the result of matching the clones on the type t pile randomly to all the clauses where a literal of type t is required.

An assignment σ with p-marginals splits each pile into two subsets, namely the clones that are true under σ and those that are false. For each type among the clones on the type t pile a t-fraction are true (because σ has p-marginals). Therefore, we expect that under the random matching for each clause type ℓ and each index j in an ℓ_j-fraction of clauses the jth literal is matched to a 'true' clone.

Judicious assignment. This observation motivates the following definition. We say that an assignment σ is $p_{\boldsymbol{d}}$***-judicious*** in $\boldsymbol{\Phi_{d,m}}$ if for all clause types $\ell = (\ell_1, \ldots, \ell_k) \in \mathcal{L}$ and all $j \in [k]$ we have

$$\sum_{i \in M_{\boldsymbol{\Phi_{d,m}}}(\ell)} \sigma(\boldsymbol{\Phi_{d,m}}_{,i,j}) = m(\ell) \cdot \ell_j, \tag{12}$$

where $\boldsymbol{\Phi_{d,m}}_{,i,j}$ denotes the jth literal of the ith clause of $\boldsymbol{\Phi_{d,m}}$, and the sum is over all i such that the ith clause has type ℓ. Let $\mathcal{S}_p(\boldsymbol{\Phi_{d,m}})$ be the set of p-judicious satisfying assignments, and set $Z_p(\boldsymbol{\Phi_{d,m}}) = |\mathcal{S}_p(\boldsymbol{\Phi_{d,m}})|$.

Given that σ is p-judicious, in order for σ to be satisfying we just need that for each type ℓ the 'true' clones are distributed so that each clause receives at least one. Thus, the event of being satisfying is merely a matter of how exactly the 'true' clones are "shuffled" amongst the clauses of type ℓ, while for each j the total number of 'true' clones of type ℓ_j is

fixed. In particular, this shuffling occurs independently for each clause type. Such random shuffling problems tend to be amenable to the second moment method. Therefore, it seems reasonable to expect that a second moment argument succeeds for $Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$. This is indeed the case for $r < r_{\mathrm{BP}} - 1 + \ln 2 \approx r_{\mathrm{BP}} - 0.3$. However, to actually reach r_{BP} we need to control one further parameter.

Fixing the cluster size. According to the physics predictions [24, 25], for $r_{\mathrm{bal}} < r < r_{\mathrm{BP}}$ the set of satisfying assignments decomposes into an exponential number of well-separated 'clusters'. More precisely, we expect that w.h.p. for any two satisfying σ, τ either $\mathrm{dist}(\sigma, \tau) < 0.01n$ (if σ, τ belong to the same cluster), or $\mathrm{dist}(\sigma, \tau) > 0.49n$ (different clusters). Formally, we simply define the ***cluster of*** σ as

$$\mathcal{C}_\sigma(\Phi) = \{\tau \in \mathcal{S}(\Phi) : \mathrm{dist}(\sigma, \tau) \leq 0.01n\}.$$

The intuitive reason why the second moment argument for $Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$ breaks down for r close to r_{BP} is that the cluster sizes $|\mathcal{C}_\sigma(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})|$ fluctuate. A similar problem occurred in prior work on random k-NAESAT [11, 12].

As in those papers, the problem admits a remarkably simple solution: let us call an assignment σ ***good*** in $\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}$ if

$$|\mathcal{C}_\sigma(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})| \leq \mathrm{E}\left[Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]. \tag{13}$$

Let $\mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$ be the set of all good $\sigma \in \mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$. To avoid fluctuations of the cluster size, we are just going to work with $Z_{p,\mathrm{good}} = |\mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})|$.

The second moment bound. We now face the task of estimating the first and the second moment of $Z_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$. The result can be summarized as follows.

THEOREM 2. *Suppose $r_{\mathrm{bal}} < r < r_{\mathrm{BP}}$. There exists $C = C(k)$ and a map $p = p_{\mathrm{BP}} : \mathbf{Z} \to [0,1]$ such that for $\boldsymbol{d}$ chosen from $\boldsymbol{D}$ and for $\boldsymbol{m}$ chosen from $\boldsymbol{M}_{\boldsymbol{d}}$ w.h.p.*

$$0 < \mathrm{E}\left[Z_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})^2\right] \leq C \cdot \mathrm{E}\left[Z_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]^2.$$

Together with Paley-Zygmund (5), Theorem 2 shows that with $\boldsymbol{d}$ chosen from $\boldsymbol{D}$ and $\boldsymbol{m}$ chosen from $\boldsymbol{M}_{\boldsymbol{d}}$ w.h.p.

$$\begin{aligned}&\mathrm{P}\left[\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}} \text{ is satisfiable}\right] \\ &\geq \mathrm{P}\left[Z_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}) > 0\right] \geq \frac{\mathrm{E}\left[Z_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]^2}{\mathrm{E}\left[Z_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})^2\right]} \geq \frac{1}{C}.\end{aligned} \tag{14}$$

The construction of $\boldsymbol{D}$, $\boldsymbol{M}_{\boldsymbol{d}}$ ensures that choosing $\boldsymbol{\Phi}$ at random is the same as first picking $\boldsymbol{d}$ from $\boldsymbol{D}$ and $\boldsymbol{m}$ from $\boldsymbol{M}_{\boldsymbol{d}}$ and then generating $\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}$. Therefore, (14) implies $\liminf_{n\to\infty} \mathrm{P}\left[\boldsymbol{\Phi} \text{ is satisfiable}\right] > 0$, so that Lemma 1 yields $r_{k-\mathrm{SAT}} \geq r_{\mathrm{BP}}$. Hence, we are left to prove Theorem 2. We begin by constructing the map p_{BP}.

Guessing the marginals. For a set $\emptyset \neq S \subset \{0,1\}^V$ and a variable x we define the S***-marginal of*** x as

$$\mu_S(x) = \sum_{\sigma \in S} \frac{\sigma(x)}{|S|}. \tag{15}$$

The definition of '$p_{\boldsymbol{d}}$-judicious' is guided by the idea that $p_{\boldsymbol{d}}(x)$ should prescribe the marginal of x in the set of all $p_{\boldsymbol{d}}$-judicious satisfying assignments. Hence, in order to make the set of $p_{\boldsymbol{d}}$-judicious assignments as good an approximation of the *entire* set of satisfying assignments as possible, we better pick p so that $p_{\boldsymbol{d}}(x)$ is a good approximation to the actual marginal $\mu_{\mathcal{S}(\boldsymbol{\Phi}_{\boldsymbol{d}})}(x)$ of x in the set of *all* satisfying assignments. The problem is that, because of the asymmetry of the k-SAT problem, these marginals are highly non-trivial quantities. Indeed, on general formulas Φ the marginals $\mu_{\mathcal{S}(\Phi)}(x)$ are $\#P$-hard to compute.

However, according to the physicists' cavity method, on random formulas with density $r < r_{\mathrm{BP}}$ the marginals can be computed by means of an efficient message passing algorithm called *Belief Propagation* [24]. While the mechanics of this are not important in our context, the result is.

CONJECTURE 1. *Suppose that $r_{\mathrm{bal}} < r < r_{\mathrm{BP}}$. Let $\boldsymbol{d}$ be chosen from $\boldsymbol{D}$ and let x be a variable. Then w.h.p.*

$$\mu_{\mathcal{S}(\boldsymbol{\Phi}_{\boldsymbol{d}})}(x) = \frac{1}{2} + \frac{d_x - d_{\neg x}}{2^{k+1}} + O\left(\frac{d_x - d_{\neg x}}{2^k}\right)^2. \tag{16}$$

We observe that (16) is in line with the notion that $\mathcal{S}(\boldsymbol{\Phi}_{\boldsymbol{d}})$ is "skewed toward" σ_{maj}. Indeed, the conjecture quantifies how much so. Motivated by Conjecture 1, we define

$$p_{\mathrm{BP}}(z) = \begin{cases} \frac{1}{2} + \frac{z}{2^{k+1}} & \text{if } |z| \leq 10\sqrt{k2^k \ln k}, \\ \frac{1}{2} & \text{otherwise.} \end{cases} \tag{17}$$

Under the distribution $\boldsymbol{D}$, the random variables $d_x, d_{\neg x}$ are asymptotically independent Poisson with mean $kr/2$ (cf. Section 2). Therefore,

$$\mathrm{E}_{\boldsymbol{d}}\left[(d_x - d_{\neg x})^2\right] = kr \leq k2^k \ln 2,$$

and standard concentration inequalities show that w.h.p. there are no more than n/k^{30} variables x with $(d_x - d_{\neg x})^2 > 100k2^k \ln k$. Hence, $p_{\boldsymbol{d}} = p_{\mathrm{BP},\boldsymbol{d}}$ is (asymptotically) equal to the conjectured value on the bulk of variables w.h.p.

In summary, the problem with the "vanilla" second moment argument is that the drift toward the σ_{maj} induces correlations amongst the satisfying assignments. Indeed, they are correlated with the majority assignment and thus with each other. We circumvent this problem by explicitly prescribing the marginal probability that each variable is set to 'true'. One could think of this as working with the intersection of $\mathcal{S}(\boldsymbol{\Phi})$ with a particular "surface" within the Hamming cube $\{0,1\}^n$, namely the assignments with $p_{\boldsymbol{d}}$-marginals. Within this surface, all assignments are slanted equally toward σ_{maj}. The Belief Propagation-informed definition of p_{BP} is meant to ensure that the surface that we consider with is (about) the most populous one, i.e., the one with the largest number of satisfying assignments in it. The core of our argument will be to show that *with respect to the marginal distribution* p_{BP}, i.e., within the surface that p_{BP} defines, two random elements of $\mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$ are typically uncorrelated (Proposition 2 below). But before we come to that, we need to compute the "first moment", i.e., the expected number of good p_{BP}-judicious satisfying assignments.

REMARK 1. *Belief Propagation actually leads to a stronger prediction than Conjecture 1. Namely, it yields a conjecture for $\mu_{\mathcal{S}(\boldsymbol{\Phi}_{\boldsymbol{d}})}(x)$ up to an additive error then tends to 0 as*

$n \to \infty$. *However, (a) this stronger conjecture is not in explicit form, and (b) it does not only depend on $d_x, d_{\neg x}$, but also on various other parameters. In any case, even a more accurate prediction would not yield a better constant than $\frac{3}{2}\ln 2$ in Theorem 1.*

REMARK 2. *In the present framework, the notion of balanced satisfying assignments from [6] simply corresponds to working with the constant map $p_{bal} : \mathbf{Z} \to [0,1],\ z \mapsto \frac{1}{2}$. This hightlights that the improvement that we obtain here stems from choosing the non-constant map p_{BP} inspired by Belief Propagation.*

REMARK 3. *The definition (15) of the marginal of a set gives rise to a formal notion of 'symmetric problem'. Namely, we could call a (binary) random CSP* ***symmetric*** *if its set $\mathcal{S}_{\mathrm{CSP}}(\mathbf{\Phi})$ of solutions is such that for each variable x w.h.p. we have $\mu_x(\mathcal{S}_{\mathrm{CSP}}(\mathbf{\Phi})) = \frac{1}{2} + o(1)$. Clearly, k-NAESAT passes this test as $\mu_x(\mathcal{S}_{\mathrm{NAE}}(\mathbf{\Phi})) = \frac{1}{2}$ for all x with certainty. Similarly, the problem of having a balanced satisfying assignment is symmetric [6], as is random k-XORSAT.*

From here on out we keep the assumptions of Theorem 2. In particular, we assume $r_{\mathrm{bal}} < r < r_{\mathrm{BP}}$. Let $\boldsymbol{d}$ be chosen from $\boldsymbol{D}$, and let $\boldsymbol{m}$ be chosen from $\boldsymbol{M_d}$. Let $p = p_{\mathrm{BP}}$ be as in (17) and $p_{\boldsymbol{d}}$ as in (11).

5. THE FIRST MOMENT

Let $\rho > \frac{3}{2}\ln 2$ be such that $r = 2^k \ln 2 - \rho$.

PROPOSITION 1. *W.h.p. $\boldsymbol{d}, \boldsymbol{m}$ are such that*

$$\mathrm{E}\left[Z_{p,\mathrm{good}}(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] = \exp\left[\frac{n}{2^k}\left(\rho - \frac{\ln 2}{2} + o_k(1)\right)\right].$$

We begin by computing $\mathrm{E}\left[Z_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]$. By definition, any assignment that is $p_{\boldsymbol{d}}$-judicious has $p_{\boldsymbol{d}}$-marginals. Thus, let $\mathcal{H}_p(\boldsymbol{d}) \subset \{0,1\}^V$ denote the set of all assignments that have $p_{\boldsymbol{d}}$-marginals. Then by the linearity of expectation,

$$\mathrm{E}\left[Z_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] = \sum_{\sigma \in \mathcal{H}_p(\boldsymbol{d})} \mathrm{P}\left[\sigma \in \mathcal{S}_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]. \quad (18)$$

Hence, we need to compute $|\mathcal{H}_p(\boldsymbol{d})|$ and the probability $\mathrm{P}\left[\sigma \in \mathcal{S}_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]$ for any $\sigma \in \mathcal{H}_p(\boldsymbol{d})$. Using basic properties of the entropy, we obtain

LEMMA 4. *Let $\chi(z) = -z\ln z - (1-z)\ln(1-z)$ denote the entropy function. Then w.h.p. $\boldsymbol{d}$ is such that*

$$\ln|\mathcal{H}_p(\boldsymbol{d})| \sim n \cdot \sum_{x \in V} \chi(p(x)).$$

Taylor expanding $\chi(z)$ around $z = 1/2$ and plugging in the definition (17) of p, we obtain that w.h.p. $\boldsymbol{d}$ is such that

$$\frac{1}{n}\ln|\mathcal{H}_p(\boldsymbol{d})| = \ln 2 - \frac{k\ln 2}{2^{k+1}} + o_k(2^{-k}). \quad (19)$$

As a next step, we compute the probability of $\sigma \in \mathcal{S}_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$ for $\sigma \in \mathcal{H}_p(\boldsymbol{d})$.

LEMMA 5. *W.h.p. $\boldsymbol{d}$, $\boldsymbol{m}$ are such that for any $\sigma \in \mathcal{H}_p(\boldsymbol{d})$,*

$$\frac{1}{n}\ln \mathrm{P}\left[\sigma \in \mathcal{S}_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] = -\ln 2 + \frac{k\ln 2}{2^{k+1}} + 2^{-k}\left[\rho - \frac{\ln 2}{2} + o_k(1)\right]. \quad (20)$$

Let us defer the proof of Lemma 5, which is the core of the first moment computation, for a little while. Combining (18)–(20), we see that w.h.p. over the choice of $\boldsymbol{d}, \boldsymbol{m}$ we have

$$\begin{aligned}\ln \mathrm{E}\left[Z_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] &= \ln|\mathcal{H}_p(\boldsymbol{d})| + \ln \mathrm{P}\left[\sigma \in \mathcal{S}_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] \\ &\sim 2^{-k}\left[\rho - \frac{\ln 2}{2} + o_k(1)\right] \cdot n \quad (21)\end{aligned}$$

To obtain the expectation of $Z_{p,\mathrm{good}}$, we show the following.

LEMMA 6. *W.h.p. over the choice of $\boldsymbol{d}, \boldsymbol{m}$ we have*

$$\mathrm{E}\left[Z_{p,\mathrm{good}}(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] \sim \mathrm{E}\left[Z_p(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right].$$

The proof of Lemma 6 is based on arguments developed in [1] for analyzing the geometry of the set of satisfying assignments. Combining (21) and Lemma 6 yields Proposition 1.

Proof of Lemma 5. In the random formula $\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}}$ there are dependencies amongst the clauses, arising, e.g., from prescribing the precise number of occurrences for each literal. The key idea of the proof is to work with a different probability space in which the clauses and even the individual literals behave independently.

The elements of this new space $\hat{\Omega}$ are all 0/1 vectors

$$(\hat{\sigma}_{ij}(\ell))_{\ell \in \mathcal{L}, i \in [m(\ell)], j \in [k]}.$$

The idea is that $\hat{\sigma}_{ij}(\ell)$ is going to represent the truth value of the jth literal of the ith clause of type ℓ. The probability distribution $\hat{\mathrm{P}}$ on $\hat{\Omega}$ is defined via a parameter vector

$$\boldsymbol{q} = (q_{\ell,j})_{\ell \in \mathcal{L}, j \in [k]}$$

with entries in $[0,1]$: each $\hat{\sigma}_{ij}(\ell)$ is the result of a Bernoulli experiment with success probability $q_{\ell,j}$, and these experiments are independent for all ℓ, i, j. We will choose $\boldsymbol{q}$ below so that we can compute the first moment (relatively) easily.

For each $\ell = (\ell_1, \ldots, \ell_k) \in \mathcal{L}$ and each $i \in [m(\ell)]$ let $S_i(\ell)$ be the event that $\max_{j \in [k]} \hat{\sigma}_{ij}(\ell) = 1$. This is going to mirror the event that the ith clause of type ℓ is satisfied. Let

$$S = \bigcap_{\ell \in \mathcal{L}} \bigcap_{i \in [m(\ell)]} S_i(\ell),$$

which is going to reflect the event that *all* clauses are satisfied. Further, for any $j \in [k]$ let $B_j(\ell)$ be the event that

$$\sum_{i \in m(\ell)} \hat{\sigma}_{ij}(\ell) = m(\ell) \cdot \ell_j.$$

Let $B = \bigcap_{\ell \in \mathcal{L}, j \in [k]} B_j(\ell)$ to capture the event of being judicious.

We observed in Section 4 that given $\boldsymbol{d}, \boldsymbol{m}$, choosing $\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}}$ amounts to generating a random matching between literal occurrences of each type t and clauses where literals of type t are required. This observation leads to

LEMMA 7. *Assume that* $\boldsymbol{q}$ *satisfies* $\hat{\mathrm{P}}[B] > 0$. *Then for any* $\sigma \in \mathcal{H}_p(\boldsymbol{d})$ *we have*

$$\mathrm{P}[\sigma \in \mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}) | \sigma \text{ is } p_{\boldsymbol{d}}\text{-judicious}] = \hat{\mathrm{P}}[S|B].$$

Suppose that σ has p-marginals. Viewing $\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}$ as a random matching, we see that for each $\ell \in \mathcal{L}$ and for each $j \in [k]$ the *expected* fraction of clauses of type ℓ whose jth literal is true under σ equals ℓ_j. Hence, a local limit theorem yields

$$\mathrm{P}[\sigma \text{ is } p_{\boldsymbol{d}}\text{-judicious}] = n^{-O(1)} \quad \text{for all } \sigma \in \mathcal{H}_p(\boldsymbol{d}). \tag{22}$$

We are left to compute $\hat{\mathrm{P}}[S|B]$. By the definition of $\hat{\mathrm{P}}$, it is quite easy to compute $\hat{\mathrm{P}}[S]$, $\hat{\mathrm{P}}[B]$ individually. Indeed, because the $\hat{\sigma}_{i,j}(\ell)$ are mutually independent, for any $\ell \in \mathcal{L}$ and any $i \in [m(\ell)]$ we have

$$\hat{\mathrm{P}}[S_i(\ell)] = 1 - \prod_{j=1}^{k}(1 - q_{\ell,j}), \text{ whence} \tag{23}$$

$$\hat{\mathrm{P}}[S] = \prod_{\ell \in \mathcal{L}} \left[1 - \prod_{j=1}^{k}(1 - q_{\ell,j})\right]^{m(\ell)}. \tag{24}$$

Once more by independence, for any $\ell \in \mathcal{L}$, $j \in [k]$ the sum $\sum_{i \in m(\ell)} \hat{\sigma}_{ij}(\ell)$ has a binomial distribution $\mathrm{Bin}(m(\ell), q_{\ell,j})$. As a consequence,

$$\begin{aligned}\hat{\mathrm{P}}[B] &= \prod_{\ell,j} \mathrm{P}[\mathrm{Bin}(m(\ell), q_{\ell,j}) = m(\ell)\cdot\ell_j + O(1)] \\ &= \exp\left[o(n) + \sum_{\ell,j} m(\ell)\,\psi(q_{\ell,j}, \ell_j)\right], \text{ with} \end{aligned} \tag{25}$$

$$\psi(x,y) = -y\ln\left(\frac{y}{x}\right) - (1-y)\ln\left(\frac{1-y}{1-x}\right) \tag{26}$$

the Kullback-Leibler divergence. We stress that (24) and (25) hold for *any* $\boldsymbol{q}$.

Hence, if we could choose $\boldsymbol{q}$ so that

$$\hat{\mathrm{P}}[B|S] = \exp(-o(n)), \tag{27}$$

then

$$\begin{aligned}\hat{\mathrm{P}}[S|B] &= \hat{\mathrm{P}}[B|S]\cdot\hat{\mathrm{P}}[S]/\hat{\mathrm{P}}[B] \\ &= \exp(-o(n))\cdot\hat{\mathrm{P}}[S]/\hat{\mathrm{P}}[B].\end{aligned} \tag{28}$$

LEMMA 8. *There is* $\boldsymbol{q}$ *such that (27) holds and*

$$q_{\ell,j} = \ell_j - 2^{-k-1} + \tilde{O}(2^{-3k/2}) \quad \textit{for all } \ell, j. \tag{29}$$

PROOF. By (23), for any ℓ, j we have

$$\mathrm{E}\Big[\sum_{i \in m(\ell)} \hat{\sigma}_{ij}(\ell) \big| S\Big] = \frac{m(\ell)q_{\ell,j}}{1 - \prod_{j=1}^{k}(1-q_{\ell,j})}.$$

Hence, if we could choose $\boldsymbol{q}$ so that

$$\frac{q_{\ell,j}}{1 - \prod_{j=1}^{k}(1-q_{\ell,j})} = \ell_j \quad \text{for all } \ell, j, \tag{30}$$

then the local limit theorem for sums of independent random variables would imply $\hat{\mathrm{P}}[B|S] = \exp(-o(n))$. Using the inverse function theorem from analysis, one can show that there is indeed a $\boldsymbol{q}$ that satisfies both (29) and (30). □

Plugging the definition (17) of $p = p_{\mathrm{BP}}$ into (24)–(29), we can now compute $\ln\hat{\mathrm{P}}[S|B] \sim \ln\hat{\mathrm{P}}[S] - \ln\hat{\mathrm{P}}[B]$ asymptotically. The result of this is the expression given in (20). Thus, Lemma 5 follows from Lemma 7 and (22).

6. THE SECOND MOMENT

We aim to prove the second part of Theorem 2. Thus, we need to estimate the expected number $\mathcal{Z} = |\mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})|^2$ of *pairs* of good p-judicious satisfying assignments. The very definition of 'good' entails an easy bound if $\mathrm{dist}(\sigma,\tau)$ is small. Indeed, let $\mathcal{Z}_{\mathrm{near}}$ be the number of pairs $(\sigma,\tau) \in \mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$ with $\mathrm{dist}(\sigma,\tau) \le 0.01n$.

LEMMA 9. *We have* $\mathrm{E}[\mathcal{Z}_{\mathrm{near}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})] \le \mathrm{E}[Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})]^2$.

PROOF. The bound (13) ensures that with certainty for any $\sigma \in \mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$ we have

$$\begin{aligned}&|\{\tau \in \mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}) : \mathrm{dist}(\sigma,\tau) \le 0.01n\}| \\ &\qquad \le |\mathcal{C}_\sigma(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})| \le \mathrm{E}[Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})].\end{aligned}$$

Hence, by the linearity of expectation,

$$\begin{aligned}&\mathrm{E}[\mathcal{Z}_{\mathrm{near}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})] \\ &\quad = \mathrm{E}\sum_{\sigma \in \mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})} |\{\tau \in \mathcal{S}_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}) : \mathrm{dist}(\sigma,\tau) \le 0.01n\}| \\ &\quad \le \mathrm{E}[Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})]\cdot\mathrm{E}[Z_{p,\mathrm{good}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})] \le \mathrm{E}[Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})]^2,\end{aligned}$$

as claimed. □

There is another range of distances for which we can get a quick bound. Let $\xi = k2^{-k/8}$ and let $\mathcal{Z}_{\mathrm{med}}$ be the number of pairs $(\sigma,\tau) \in \mathcal{S}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$ such that $\mathrm{dist}(\sigma,\tau) > 0.01n$ and $|\mathrm{dist}(\sigma,\tau) - \frac{n}{2}| > \xi n$. A rather crude calculation yields the following.

LEMMA 10. *W.h.p. over the choice of* $\boldsymbol{d}, \boldsymbol{m}$ *we have*

$$\mathrm{E}[\mathcal{Z}_{\mathrm{med}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})] = o(1).$$

Thus, we are left to analyze the number $\mathcal{Z}_{\mathrm{centre}}$ of pairs $(\sigma,\tau) \in \mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})^2$ such that $|\mathrm{dist}(\sigma,\tau) - \frac{n}{2}| \le \xi n$. We zoom in on this range by quantifying the similarity between two assignments more accurately. Define the ***overlap*** of $\sigma,\tau \in \{0,1\}^V$ in $\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}$ to be the vector

$$\begin{aligned}\omega(\sigma,\tau) &= \omega_{\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}}(\sigma,\tau) = (\omega_{\ell,j}(\sigma,\tau))_{\ell\in\mathcal{L}, j\in[k]}, \text{ with} \\ \omega_{\ell,j}(\sigma,\tau) &= \sum_{i \in M_{\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}}(\ell)} \frac{\sigma(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m},i,j})\tau(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m},i,j})}{m(\ell)}.\end{aligned}$$

Thus, $\omega_{\ell,j}(\sigma,\tau)$ is the fraction of clauses of type ℓ in $\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}$ whose jth literal is true under both σ,τ. Let

$$\omega^* = (\omega^*_{\ell,j}) \quad \text{with} \quad \omega^*_{\ell,j} = \ell_j^2 \quad \text{for all } \ell, j. \tag{31}$$

LEMMA 11. *W.h.p. over the choice of* $\boldsymbol{d}$, $\boldsymbol{m}$ *the following is true. Choose two assignments* $\boldsymbol{\sigma}, \boldsymbol{\tau}$ *with p-marginals uniformly and independently. Then* $\mathrm{E}[\omega_{\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}}(\boldsymbol{\sigma},\boldsymbol{\tau})] = \omega^*$.

For a vector $\omega = (\omega_{\ell,j})$ we let $\mathcal{Z}_\omega$ denote the number of all $(\sigma,\tau) \in \mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})^2$ such that $\|\omega(\sigma,\tau) - \omega\|_\infty \leq O(1/n)$. The key step of the second moment analysis can be summarized as follows.

PROPOSITION 2. *There is $\zeta = \zeta(k) > 0$ such that w.h.p. over the choice of $\boldsymbol{d}, \boldsymbol{m}$ for all ω such that $\|\omega - \omega^*\|_\infty \leq 2\xi$ we have*

$$\mathrm{E}\left[\mathcal{Z}_\omega(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] \leq \exp\left[-\zeta n \cdot \|\omega - \omega^*\|_2^2\right] \cdot \mathrm{E}\left[\mathcal{Z}_{\omega^*}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right].$$

Before we come to the proof of Proposition 2, let us indicate how it implies the second moment bound. For $\delta > 0$ let $\mathcal{Z}_{\omega,\delta}$ be the number of $(\sigma,\tau) \in \mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})^2$ such that $\|\omega(\sigma,\tau) - \omega\|_\infty \leq \delta$. Proposition 2 shows that the contribution of $\mathcal{Z}_\omega$ decays exponentially in $\|\omega - \omega^*\|_2^2$. Therefore, with a bit of calculus we obtain

$$\mathrm{E}\left[\mathcal{Z}_{\omega^*,2\xi}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] \leq C \cdot \mathrm{E}\left[\mathcal{Z}_{\omega^*,1/\sqrt{n}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]. \tag{32}$$

Furthermore, for σ,τ with $\mathrm{dist}(\sigma,\tau) \leq \xi$ we have

$$\left\|\omega^* - \mathrm{E}\left[\omega_{\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}}(\sigma,\tau)\right]\right\| \leq \xi.$$

Indeed, $\omega_{\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}}(\sigma,\tau)$ is sufficiently concentrated about its expectation to obtain

$$\mathrm{E}\left[\mathcal{Z}_{\mathrm{centre}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] \sim \mathrm{E}\left[\mathcal{Z}_{\omega^*,2\xi}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]. \tag{33}$$

As ω^* is the expected overlap of two perfectly uncorrelated assignments σ,τ with $p_{\boldsymbol{d}}$-marginals (cf. Lemma 11), it relatively straightforward to verify that

$$\mathrm{E}\left[\mathcal{Z}_{\omega^*,1/\sqrt{n}}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] \leq C' \cdot \mathrm{E}\left[Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]^2 \tag{34}$$

for some $C' = C'(k) > 0$. Combining (32)–(34), we find

$$\mathrm{E}\left[\mathcal{Z}(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right] \quad \leq \quad C'' \cdot \mathrm{E}\left[Z_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]^2 \tag{35}$$

for some $C'' = C''(k) > 0$. Finally, Theorem 2 follows by combining (35) with Lemmas 6, 9 and 10.

Establishing Proposition 2. Assume that $\|\omega - \omega^*\|_\infty \leq \xi$. To compute $\mathrm{E}\left[\mathcal{Z}_\omega(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]$, we need to estimate for each pair σ,τ with $\omega_{\boldsymbol{\Phi}_{\boldsymbol{d}}}(\sigma,\tau) = \omega$ the probability that $\sigma,\tau \in \mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}})$. Similarly as in the computation of the first moment, we are going to work with a different probability space $(\hat{\Omega},\hat{\mathrm{P}})$. This time, $\hat{\Omega}$ consists of all vectors of 0/1 *pairs*

$$(\hat{\sigma}_{ij}(\ell),\hat{\tau}_{ij}(\ell))_{\ell\in\mathcal{L},i\in[m(\ell)],j\in[k]}$$

The distribution $\hat{\mathrm{P}}$ is defined by means of a vector

$$\boldsymbol{q} = (q_{\ell,j}^{ab})_{\ell\in\mathcal{L},j\in[k],a,b\in\{0,1\}}$$

that satisfies $\sum_{a,b\in\{0,1\}} q_{\ell,j}^{ab} = 1$ and $q_{\ell,j}^{01} = q_{\ell,j}^{10}$ for all ℓ, j. Given $\boldsymbol{q}$, we let

$$\hat{\mathrm{P}}\left[\hat{\sigma}_{ij}(\ell) = a,\, \hat{\tau}_{ij}(\ell) = b\right] = q_{\ell,j}^{ab} \qquad (a,b \in \{0,1\}),$$

independently for all ℓ, i, j. The intended semantics is that $\hat{\sigma}_{ij}(\ell),\hat{\tau}_{ij}(\ell)$ represent the truth values of the jth literal of the ith clause of type ℓ under a pair of assignments. Let

$$q_{\ell,j} = q_{\ell_j}^{10} + q_{\ell_j}^{11}$$

denote the marginal probability that $\hat{\sigma}_{ij}(\ell) = 1$.

Let $S_i(\ell)$ denote the event that

$$\max_{j\in[k]} \hat{\sigma}_{ij}(\ell) = \max_{j\in[k]} \hat{\tau}_{ij}(\ell) = 1,$$

and let $S = \bigcap_{\ell,i} S_i(\ell)$. Further, let $B_j(\ell)$ denote the event

$$\sum_{i\in m(\ell)} \hat{\sigma}_{ij}(\ell) = \sum_{i\in m(\ell)} \hat{\tau}_{ij}(\ell) = m(\ell)\cdot\ell_j,$$

and let $B = \bigcap_{\ell,j} B_j(\ell)$. Finally, let $\hat{\omega} = (\hat{\omega}_{\ell,j})$ denote the vector with entries

$$\hat{\omega}_{\ell,j} = \sum_{i\in[m(\ell)]} \hat{\sigma}_{ij}(\ell)\hat{\tau}_{ij}(\ell)/m(\ell).$$

In analogy to Lemma 7, we obtain

LEMMA 12. *Let $\sigma,\tau \in \{0,1\}^V$ have $p_{\boldsymbol{d}}$-marginals. Then*

$$\mathrm{P}\left[\sigma,\tau \in \mathcal{S}_p(\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}) | \omega_{\boldsymbol{\Phi}_{\boldsymbol{d},\boldsymbol{m}}}(\sigma,\tau) = \omega\right] = \hat{\mathrm{P}}\left[S|B,\,\hat{\omega} = \omega\right].$$

As in the proof of Lemma 7, to compute $\hat{\mathrm{P}}\left[S|B,\,\hat{\omega} = \omega\right]$ we are going to choose $\boldsymbol{q}$ so as to maximize $\hat{\mathrm{P}}\left[B,\,\hat{\omega} = \omega\right]$. This allows us to express the $\hat{\mathrm{P}}\left[S|B,\,\hat{\omega} = \omega\right]$ as a quotient of a term that factorises over clauses and a bunch of binomial large deviations.

LEMMA 13. *There exists $\boldsymbol{q} = \boldsymbol{q}(\omega)$ such that*

$$\begin{aligned}
\ell_j &= \frac{q_{\ell,j} - (q_{\ell,j} - q_{\ell,j}^{11})\prod_{h\neq j}(1-q_{\ell,h})}{1 - 2\prod_{h=1}^k(1-q_{\ell,h}) + \prod_{h=1}^k(1-2q_{\ell,h}+q_{\ell,h}^{11})},\\
\omega_{\ell,j} &= \frac{q_{\ell,j}^{11}}{1 - 2\prod_{h=1}^k(1-q_{\ell,h}) + \prod_{h=1}^k(1-2q_{\ell,h}+q_{\ell,h}^{11})}
\end{aligned}$$

for all $\ell \in \mathcal{L}, j \in [k]$. For this $\boldsymbol{q}$, and with ψ as in (26) let

$$\begin{aligned}
\mathcal{P}_\ell(\omega) \quad = \quad & \ln\left[1 - 2\prod_{j=1}^k(1-q_{\ell,j}) + \prod_{j=1}^k(1-2q_{\ell,j}+q_{\ell,j}^{11})\right]\\
& - \sum_{j\in[k]}\Big[\psi(q_{\ell,j}^{11},\omega_{\ell,j}) +\\
& \qquad (1-\omega_{\ell,j})\psi\left(\frac{1-2q_{\ell,j}+q_{\ell,j}^{11}}{1-q_{\ell,j}^{11}}, \frac{1-2\ell_j+\omega_{\ell,j}}{1-\omega_{\ell,j}}\right)\Big]
\end{aligned}$$

for any $\ell \in \mathcal{L}$, and set $\mathcal{P}(\omega) = \sum_{\ell\in\mathcal{L}} \frac{m(\ell)}{n}\mathcal{P}_\ell(\omega)$. Then

$$\hat{\mathrm{P}}\left[S|B,\,\hat{\omega} = \omega\right] = \Theta(1)\cdot\exp\left[n\cdot\mathcal{P}(\omega)\right].$$

To complete the proof, we need to analyze the function $\mathcal{P}(\omega)$.

LEMMA 14. *The function $\mathcal{P}$ has two continuous derivatives. These satisfy*

$$\frac{\partial}{\partial\omega_{\ell,j}}\mathcal{P}(\omega^*) = 0, \qquad \frac{\partial^2}{\partial\omega_{\ell,h}\omega_{\ell,j}}\mathcal{P}(\omega_\ell) \leq 4^{-k+o(k)}$$

for all $\ell \in \mathcal{L}, j, h \in [k]$ and all ω such that $\|\omega - \omega^*\|_\infty \leq 2\xi$.

The fact that the first derivative of $\mathcal{P}$ vanishes at ω^* is crucial. Indeed, because of this we can use Taylor's theorem to bound $\mathcal{P}$ around ω^* by a *quadratic* function, obtaining

$$\mathcal{P}(\omega) \leq \mathcal{P}(\omega^*) + 4^{-k+o(k)} \sum_{\ell \in \mathcal{L}} \frac{m(\ell)}{n} \left\|\omega_\ell - \omega_\ell^*\right\|_2^2, \qquad (36)$$

where $\omega_\ell = (\omega_{\ell,j})_{j\in[k]}$. To obtain the desired bound on $\mathrm{E}[Z_\omega]$, we need to estimate the number of pairs (σ, τ) with overlap ω. A somewhat delicate analysis shows that this number can be bounded by a function $\mathcal{H}(\omega)$ such that

$$\frac{\mathcal{H}(\omega)}{\mathcal{H}(\omega^*)} \leq \exp\left[-2^{-k+o(k)} \sum_{\ell \in \mathcal{L}} \frac{m(\ell)}{n} \left\|\omega_\ell - \omega_\ell^*\right\|_2^2\right]. \qquad (37)$$

Combining (36) and (37) yields

$$\begin{aligned} \frac{\mathrm{E}\left[\mathcal{Z}_\omega(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]}{\mathrm{E}\left[\mathcal{Z}_{\omega^*}(\mathbf{\Phi}_{\boldsymbol{d},\boldsymbol{m}})\right]} &\leq C' \cdot \frac{\mathcal{H}(\omega)\mathcal{P}(\omega)}{\mathcal{H}\left(\omega^*\right)\mathcal{P}(\omega^*)} \\ &\leq C' \cdot \exp\left[-\zeta \left\|\omega - \omega^*\right\|_2^2\right] \end{aligned}$$

for certain $C' = C'(k), \zeta = \zeta(k) > 0$, as desired.

Acknowledgment. The first author thanks Dimitris Achlioptas for helpful discussions on the second moment method.

7. REFERENCES

[1] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.

[2] D. Achlioptas, R. Menchaca-Mendez: Unsatisfiability Bounds for Random CSPs from an Energetic Interpolation Method. Proc. 39th ICALP (2012) 1–12.

[3] D. Achlioptas, R. Menchaca-Mendez: Exponential lower bounds for DPLL algorithms on satisfiable random 3-CNF formulas. Proc. 15th SAT (2012) 327–340.

[4] D. Achlioptas, C. Moore: Random k-SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.

[5] D. Achlioptas, A. Naor: The two possible values of the chromatic number of a random graph. Annals of Mathematics **162** (2005) 1333–1349.

[6] D. Achlioptas, Y. Peres: The threshold for random k-SAT is $2^k \ln 2 - O(k)$. Journal of the AMS **17** (2004) 947–973.

[7] M. Bayati, D. Gamarnik, P. Tetali: Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. Proc. 42nd STOC (2010) 105–114.

[8] V. Chvátal, B. Reed: Mick gets some (the odds are on his side). Proc. 33th FOCS (1992) 620–627.

[9] A. Coja-Oghlan: A better algorithm for random k-SAT. SIAM J. Computing **39** (2010) 2823–2864.

[10] A. Coja-Oghlan: On belief propagation guided decmation for random k-SAT. Proc. 22nd SODA (2011) 957–966.

[11] A. Coja-Oghlan, K. Panagiotou: Catching the k-NAESAT threshold. Proc. 43rd STOC (2012) 899–908.

[12] A. Coja-Oghlan, L. Zdeborová: The condensation transition in random hypergraph 2-coloring. Proc. 23rd SODA (2012) 241–250.

[13] O. Dubois, Y. Boufkhad: A general upper bound for the satisfiability threshold of random r-SAT formulae. J. Algorithms **24** (1997) 395–420.

[14] O. Dubois, J. Mandler: The 3-XORSAT threshold. Proc. 43rd FOCS (2002) 769–778.

[15] M. Dyer, A. Frieze, C. Greenhill: On the chromatic number of a random hypergraph. Preprint (2012).

[16] S. Franz, M. Leone: Replica bounds for optimization problems and diluted spin systems. J. Statist. Phys. **111** (2003) 535–564.

[17] E. Friedgut: Sharp Thresholds of Graph Proprties, and the k-SAT Problem. J. AMS **12** (1999) 1017–1054.

[18] A. Frieze, S. Suen: Analysis of two simple heuristics on a random instance of k-SAT. Journal of Algorithms **20** (1996) 312–355.

[19] A. Frieze, N. Wormald: Random k-Sat: a tight threshold for moderately growing k. Combinatorica **25** (2005) 297–305.

[20] A. Goerdt: A threshold for unsatisfiability. Proc. 17th MFCS (1992) 264–274.

[21] S. Kirkpatrick, B. Selman: Critical behavior in the satisfiability of random boolean expressions. Science **264** (1994) 1297–1301.

[22] L. Kirousis, E. Kranakis, D. Krizanc, Y. Stamatiou: Approximating the unsatisfiability threshold of random formulas. Random Structures Algorithms **12** (1998) 253–269.

[23] L. Kroc, A. Sabharwal, B. Selman: Message-passing and local heuristics as decimation strategies for satisfiability. Proc 24th SAC (2009) 1408–1414.

[24] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Gibbs states and the set of solutions of random constraint satisfaction problems. Proc. National Academy of Sciences **104** (2007) 10318–10323.

[25] M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. Science **297** (2002) 812–815.

[26] M. Molloy: The freezing threshold for k-colourings of a random graph. Proc. 43rd STOC (2012) 921–930.

[27] A. Montanari, R. Restrepo, P. Tetali: Reconstruction and clustering in random constraint satisfaction problems. SIAM J. Discrete Math. **25** (2011) 771–808.

[28] R. Moser, G. Tardos: A constructive proof of the general Lovász local lemma. J. ACM **57** (2010).

[29] V. Rathi, E. Aurell, L. K. Rasmussen, M. Skoglund: Bounds on threshold of regular random k-SAT. Proc. 12th SAT (2010) 264–277.